

Some Arguments About Legal Arguments

L. Thorne McCarty
Rutgers University
New Brunswick, New Jersey 08903

1 Introduction

Everyone knows that lawyers are trained in the art of argument. But what exactly does this mean? Is a legal argument a chain of valid inferences, grounded in authoritative rules? Or is it merely a chain of plausible inferences? Does it require the citation of cases, pro and con? Are there canons of correct, or acceptable, argumentation? Can lawyers agree about what counts as a persuasive argument in a particular case, even if they disagree about the correct outcome?

These are important questions, and the literature in AI and Law has exploded recently with books and articles that purport to answer them. The titles alone are revealing: *Modeling Legal Argument: Reasoning with Cases and Hypotheticals* [Ashley, 1990]; *Arguments and Cases: An Inevitable Intertwining* [Skalak and Rissland, 1992]; *Logical Tools for Modeling Legal Argument* [Prakken, 1993]; *A Formal Model of Legal Argumentation* [Sartor, 1994]; *The Pleadings Game: An Exercise in Computational Dialectics* [Gordon, 1994]; *A Dialectical Model of Assessing Conflicting Arguments in Legal Reasoning* [Prakken and Sartor, 1996a]. It would appear from these examples that “argument” is the centerpiece of our subject.

In this paper, I will argue (and the self reference is intentional) that most of this work is misleading, and misrepresents the true nature of legal argument. I will set the stage in Section 2 by reviewing the rudiments of civil procedure in the United States, and describing the contexts in which legal arguments occur. Sections 3 and 4, which follow, are the main critical sections of the paper. Section 3 discusses “rule-based” theories of legal argument, as represented by Gordon, Prakken and Sartor. Section 4 discusses “case-based” theories of legal argument, as represented by Rissland, Ashley and Skalak. Section 5 then discusses my own theory briefly. In the polemical spirit in which this paper is offered, I call this simply “The Correct Theory”.

Throughout the discussion, I will assume that our main goal is theoretical, not practical. That is, I will assume that we are trying to acquire an understanding

of legal argument by building and studying computational models. If our main goal is practical — that is, if we are primarily interested in building intelligent legal information systems — then different considerations might apply. I will address this point briefly in the concluding section.

2 What is a Legal Argument?

I will begin with an elementary account of the progress of a civil lawsuit through the federal court system in the United States,¹ so that we can understand clearly what it is that lawyers are arguing about.

A lawsuit begins when one party files a *complaint*. According to the Federal Rules of Civil Procedure, a complaint must contain:

... (1) a short and plain statement of the grounds upon which the court’s jurisdiction depends, ... (2) a short and plain statement of the claim showing that the pleader is entitled to relief, and (3) a demand for judgment for the relief the pleader seeks. *Fed. R. Civ. P., Rule 8(a)*.

The most important part of this requirement is in 8(a)(2): the *claim*. Although no technical forms of pleading are required, there is a basic structure to every claim. The plaintiff must allege certain *facts*, and then state a conclusion of *law* that justifies the relief being sought. The distinction between fact and law is essential to the operation of the entire system. For example, if the defendant files an *answer*, there are usually two things that he or she can do: (i) deny the allegations of fact in the plaintiff’s complaint; or (ii) assert additional facts which, together with the necessary conclusions of law, would constitute an affirmative defense to the plaintiff’s original claim.

At this point in our scenario, though, nothing has yet happened that a lawyer would call an “argument”. There are only allegations and denials. In the legal vernacular, the argument has not yet been “joined”. The argument would have been joined, however, if the

Reprinted with permission from *Proceedings, Sixth International Conference on Artificial Intelligence and Law*, pages XXX–XXX, ACM Press, 1997. Copyright 1997 by the Association for Computing Machinery, Inc.

¹Of course, the legal systems of continental Europe are quite different in many respects, and I would welcome comments from my European colleagues about these differences. Since the arguments in this paper depend only on the most abstract structural features of a legal system, however, I doubt that they would be significantly altered by a move to Europe.

defendant had filed a motion to dismiss the plaintiff’s complaint for “failure to state a claim upon which relief can be granted” under Rule 12(b)(6). The point of a motion under Rule 12(b)(6) is to test the sufficiency of the complaint as a matter of law. The defendant must *assume* that the factual allegations in the plaintiff’s complaint are true, and yet still argue that the plaintiff’s legal conclusions do not follow. Thus, this is an argument about what the law *is*. A similar argument about matters of law would be triggered by a motion for summary judgment under Rule 56. On this motion, additional facts may be taken into account, by affidavits, etc., but only if they are uncontested. Here is the standard:

The judgment sought shall be rendered forthwith if the pleadings, depositions, answers to interrogatories, and admissions on file, together with the affidavits, if any, show that there is no genuine issue as to any material fact and that the moving party is entitled to a judgment as a matter of law. *Fed. R. Civ. P., Rule 56(c)*.

Again, under this standard, the plaintiff and the defendant would argue about questions of law, such as: What are the substantive rules? What are the valid defenses? Who has the burden of proof?

If the lawsuit proceeds to trial, however, the argument shifts to questions of fact. It is the job of the “trier of fact” (either a jury or a judge) to determine what actually happened in the case, and it is the job of the trial lawyer to persuade the trier of fact that his or her version of the truth is the correct one. This is a very different kind of argument than the argument over questions of law, and it requires very different skills. The facts at issue might be objective facts that cannot now be directly observed (e.g., what did the defendant say?), or subjective facts that are inherently unobservable (e.g., what was the defendant’s state of mind?). Or they might be “mixed questions of law and fact” that combine subjective standards with one or more objective facts (e.g., did the defendant use reasonable care?). In each case, the lawyer’s job is to marshal the raw evidence and explain it in such a way that the trier of fact takes that last inferential step to the desired conclusion.

There are institutional differences as well. Findings of fact at trial have no precedential value, and there is no requirement of consistency from case to case. Also, within broad limits, a finding of fact cannot be overturned on appeal. By contrast, the express purpose of an appeal to a higher court is to challenge the decisions of the lower court on matters of law. Thus, if a lawyer has lost a motion for summary judgment, or if a significant piece of evidence has been excluded, or if a proposed jury instruction has been rejected, there

might be grounds for an appeal. If so, the argument now shifts back to questions of law.

It may be helpful to characterize the various components of this system in a way that would be familiar to most researchers in AI. The analysis of claims is mostly *deductive*. The alleged facts must lead by the application of acceptable rules to the desired legal conclusions. The arguments at trial are mostly *abductive*. The problem, in many cases, is to reason backwards from observed evidence to unobserved causes. Finally, appellate legal argument is mostly *inductive*, or more generally *constructive*. An appellate judge sits at the pinnacle of the legal system, after all, where the substance of the law is debated most intensively. Lawyers at this level are arguing about what the law ought to be, and this is necessarily a process of theory construction.

With this background, we can now turn to the two main theories of legal argument in the recent literature on AI and Law.

3 Rule-Based Theories of Legal Argument

The current “rule-based” theories of legal argument originated, not in the study of law, but in the study of nonmonotonic reasoning. The volume of literature on nonmonotonic reasoning is now almost unmanageable, and I will simply refer the reader to a survey article [Reiter, 1987] and a collection of readings [Ginsberg, 1987] for a snapshot of the field as it existed *circa* 1987. It is interesting to note that the word “argument” does not appear as a technical term in either of these references. Since that time, however, there have been two main theoretical developments in which the abstract concept of an argument plays a central role.² I will describe these briefly, and then discuss their relevance to AI and Law.

One development was in the field of logic programming [Lloyd, 1987], where a long-standing goal has been to find a suitable semantics for “negation-as-failure”. If a logic program is *stratified* (e.g., if it consists of the clauses $p :- \text{not } q$ and $q :- \text{not } r$), then there is universal agreement about its meaning.³ But if a logic program is *unstratified* (e.g., if it includes the clauses $p :- \text{not } q$ and $q :- \text{not } p$), then there are many possible interpretations. One important contribution to this subject was the work of Phan Minh Dung, who first described the meaning of an unstratified logic program in terms of an abstract set of *arguments* and an abstract *attacks*

²Actually, both developments can be traced back to the seminal papers of David Poole [Poole, 1985] [Poole, 1988], although Poole does not use the concept of an argument explicitly in his work.

³Technically, a stratified logic program has a *unique stable model*. In the example in the text, this model would consist of the single positive literal: q .

relation between arguments [Dung, 1993] [Dung, 1995]. Robert Kowalski and his colleagues have also expanded upon this point of view, and a very readable exposition of their ideas can be found in [Kowalski and Toni, 1996].

The main contribution of Dung's work, however, was not to provide an understanding of human argument, but rather to provide a uniform mathematical framework for a wide variety of combinatorial problems. Dung's system is justified by showing that it can represent both the well-founded semantics [Van Gelder *et al.*, 1991] and the stable model semantics [Gelfond and Lifschitz, 1988] of logic programs; by showing that it can represent default logic [Reiter, 1980] and autoepistemic logic [Moore, 1985]; by showing that it can encode the stable marriage problem [Hall, 1935]; and so on. For these purposes, it is irrelevant that the primitive objects in the system are called "arguments". Indeed, there is a closely related system studied by Marek, Nerode and Remmel [Marek *et al.*, 1990], with an equally broad range of applications, in which the primitive objects are called simply "nonmonotonic rules". To call them "arguments" is just to use a metaphor, in the same way that mathematics has used metaphors from ordinary language (e.g., "group", "manifold", "sheaf", "category") throughout its history.

The second development is closely related, but it comes from a different community of researchers. Early in the study of nonmonotonic reasoning, it became apparent that certain patterns of inference were strongly preferred by human subjects despite the fact that they were not sanctioned by traditional logics. To take the (notoriously overworked) canonical example, if we are told that "birds normally fly" and "penguins normally do not fly", we will tend to conclude that a particular penguin does not fly, even though we know that all penguins are birds, by definition, and thus subject to the broader rule as well. The explanation often advanced is *specificity*. We assume that the negative conclusion for the specific subclass, "penguins", is intended to override the positive conclusion for the general class, "birds". Numerous formalisms were developed to account for this phenomenon in the late 1980's [Delgrande, 1988] [Horty *et al.*, 1990], and among these were several that depended on an explicit notion of argument [Pollock, 1987] [Loui, 1987] [Simari and Loui, 1992] [Geffner and Pearl, 1992].

For example, in [Simari and Loui, 1992], an *argument* is (roughly) a subset of a designated set of "defeasible rules" that can consistently derive a particular *conclusion*, and the set of all possible *argument/conclusion* pairs (in a particular language) is ordered by a syntactic specificity condition. This order is then used, in a complex series of definitions, to pick out the "better" arguments in a particular case. Simari and Loui show that their system produces a reasonable answer in several familiar examples — including, of course, the

case of the "nonflying birds" and the case of the "republican quakers". But there is a problem. There are other systems, with different definitions [Poole, 1985] [Pollock, 1987], and it always seems possible to find an example on which these systems disagree. How, then, do we decide which one is correct? Unfortunately, this is precisely the "clash of intuitions" that has plagued the field of nonmonotonic reasoning since its inception [Touretzky *et al.*, 1987].

Despite these problems, it is easy to see why the nonmonotonic reasoners were attracted to the law, and *vice versa*. For the formalists, such as Loui, legal argument seemed to provide a concrete instantiation of their increasingly abstract ideas. For the lawyers, such as Gordon, Prakken and Sartor, nonmonotonic reasoning systems seemed to provide a tool that was needed for the representation of complex legal rules. Statutes are replete with "unless" clauses, which sometimes signal a shift in the burden of proof, and it is natural to encode these clauses using "negation-as-failure". *Lex Specialis*, that hoary maxim of statutory interpretation, suggests that there may be a special role for specificity in the law. Thus, in a period of great ferment around 1991, many of the extant nonmonotonic reasoning systems were applied to the law by a growing contingent of AI and Law researchers. For examples of this early work, see [Gordon, 1991] [Prakken, 1991] [Sartor, 1991].

In the intervening years, the emphasis has shifted slightly. Specificity has declined in importance, and this seems to me to be a positive development.⁴ Tom Gordon combed through Article 9 of the Uniform Commercial Code looking for examples of implicit exceptions ordered by specificity, and found almost none [Gordon, 1994]. Henry Prakken used a specificity ordering in his dissertation [Prakken, 1993], but in his more recent system he works with a general priority mechanism that can be used for *Lex Specialis*, *Lex Posterior*, *Lex Superior*, or anything else [Prakken and Sartor, 1996a]. Once we take this step, however, the priority mechanism itself ends up on the endangered species list. Kowalski and Toni have shown that priorities can be defined away if we have a general mechanism for metalevel reference plus an appropriate encoding of "negation-as-failure" [Kowalski and Toni, 1996]. And if this leads to a sim-

⁴In his review of Prakken's dissertation, Ron Loui takes issue with my assertion that a well-drafted statute should mark exceptions explicitly, instead of relying on an implicit ordering according to specificity [Loui, 1995]. According to Loui [p. 147]: "Prakken maintains that in Dutch law he found many examples of implicit specificity, and that some are in the thesis." Loui then cites an example from page 139 of Prakken's dissertation. This is a bizarre reference! Example 6.43 on page 139 of Prakken's dissertation is an artificial example, intended to illustrate the concept of "soft rebutting defeaters". It does not claim to be a statement of Dutch law, nor does it claim to be a representation of a Dutch statute. It is a modification of Example 6.2 on page 122, which is simpler, and which also neither claims to be a statement of Dutch law nor a representation of a Dutch statute.

pler system, then it is a better system, in my opinion.

I am less happy about another trend in the current literature. Occasionally, one sees a tendency to push rule-based theories of legal argument far beyond their reasonable limits. Prakken and Sartor described the field this way in their introduction to a special issue of *Artificial Intelligence and Law*:

For legal philosophy a new form of legal logic is emerging, whose features (its ability of dealing with conflicting sets of premises, the possibility of deriving defeasible conclusions, the adherence to the intuitive structure of legal reasoning) make it attractive for the theoretical analysis of legal reasoning. [Prakken and Sartor, 1996b], p. 158.

This statement seems to me to be grossly exaggerated, and seriously mistaken about the nature of legal reasoning.

Here is why: Take a look at some of the examples in these papers, such as the Italian priority rule on building regulations, Example 6.4 in [Prakken and Sartor, 1996a], or the European Community rule on pasta, Example 6.5 in [Prakken and Sartor, 1996a]. Each example consists of a handful of rules that include some nonmonotonic negation operators and/or some multiple priorities. Now, it is conceivable that rules of this sort could serve as the basis of a *claim*, as described in Section 2 above,⁵ in which case we could reason deductively from some alleged facts to some desired legal conclusion. But this will only be the case if both the negations and the priorities in these rules are *stratified*, for only then will a judge be able make sense of the complaint. In other words, statutes have to be clear and deductions have to be simple if they are to be used at this stage of the legal process.

What if this is not the case? Prakken and Sartor want us to believe that a judge would apply their formal theory to evaluate the conflicting claims, and thereby arrive at the correct result. But if so, we are in deep trouble, because the “clash of intuitions” is still with us. Here are some recent instances: (1) Example 6.5 in [Prakken, 1993]. Here Prakken resurrects an earlier disagreement between Loui [1987] and Poole [1985], and sides with Poole. (2) Example 5.8 in [Prakken and Sartor, 1996a]. Here Prakken and Sartor derive a conclusion at variance with [Horty *et al.*, 1990] and [Nute, 1994]. They write: “We think that in a dispute a judge would feel compelled to determine whether $r_1 < r_3$ before deciding in favour of P .” (3) Example 7.2 in [Kowalski and Toni, 1996]. Here Kowalski and Toni disagree with a derivation in [Prakken and Sartor, 1995]. They write: “Note that our approach takes into account the given

⁵This is essentially Tom Gordon’s approach in his formulation of the “Pleadings Game” [Gordon, 1994] [Gordon, 1995]. I will discuss this work briefly in Section 5 below.

priority, while their approach does not.” Can anyone imagine a lawyer making an argument like this before a real judge?

But suppose the statute really is a mess. What would we do then? The answer is easy. We would construct a new version of the legal rules, which would be simple and clear — and which would, incidentally, favor our client — and we would argue that this was what the legislators had really intended.

I will discuss this point further in Section 5 below. But first, I will consider the second main theory of legal argument in the recent literature: the “case-based” approach.

4 Case-Based Theories of Legal Argument

Most of the work on “case-based” theories of legal argument today can be traced back to the pioneering research of Edwina Rissland and her students, Kevin Ashley and David Skalak. The earliest work in this tradition is purely case-based, and is exemplified by Ashley’s system HYPO [Ashley, 1990]. More recent work, which combines rules and cases, is exemplified by Rissland and Skalak’s system CABARET [Rissland and Skalak, 1991].

The great virtue of the articles on HYPO and CABARET is their extensive use of real examples. The focus is on appellate legal argument, as I have defined it in Section 2 above, and the reader is given many illustrations in which real lawyers and real judges argue real cases. Not surprisingly, when you look closely at real data, interesting patterns emerge. In one of their articles, for example, Skalak and Rissland construct an extensive typology of legal arguments [Skalak and Rissland, 1992], sometimes using standard terminology and sometimes inventing their own. They discuss the familiar argument forms, such as “drawing distinctions” and “making analogies” [pp. 15-18], and they discuss the important class of “slippery slope” arguments [pp. 22-23]. But they also discuss some novelties: “reductio loops” [pp. 23-25] and “turkey/chicken/fish” arguments [pp. 19-20]. And they include a carefully balanced appraisal of the class of “weighing” and “balancing” arguments [pp. 25-26]. This is useful and insightful work, and it makes a significant contribution to our understanding of legal reasoning.

The problem, though, is that the computational model underlying HYPO and CABARET is not sufficient to support the theory. HYPO is based on a representation of cases in terms of predefined “factors” that are “oriented” towards either the plaintiff or the defendant in a particular area of the law. For example, in a trade secret misappropriation case, *Agreed-Not-To-Disclose* and *Security-Measures* are factors that

favor the plaintiff, while *Agreement-Not-Specific* and *Employee-Sole-Developer* are factors that favor the defendant. Most of these factors are boolean, i.e., they are either present or absent in a particular case, although some are scalar, i.e., they have a numerical range. In technical terms, whether the factors are boolean or scalar or mixed, the operational representation of a case in HYPO is just a *feature vector*. The same is true of CABARET. In CABARET, although the representation also includes logical rules, these are stored in a separate module. The case-based module of CABARET is based on the feature vector model of HYPO, with the trade secret factors replaced by a set of “home office deduction” factors that are deemed relevant under Section 280A(c)(1) of the Internal Revenue Code.

Now there is nothing wrong, *per se*, with feature vector models. Indeed, most contemporary machine learning programs are based on feature vector models [Quinlan, 1993]. Also, there is an ancient tradition in the field of legal studies in which cases are encoded as feature vectors. The earliest work in this tradition was by Fred Kort and Reed Lawlor in the 1960’s [Kort, 1963] [Lawlor, 1963], and their approach was further developed by the political scientist Glendon Schubert in the 1970’s [Schubert, 1975]. There are major differences between this early work and the work being done today, of course. Schubert’s research was based on linear statistical models, in which the features associated with a case are weighted numerically, and he was primarily interested in predicting the outcome of a case by looking at political factors, rather than legal factors. Kevin Ashley’s work [Ashley, 1990] is much more sophisticated on both counts. Ashley rejects simple numerical weightings in favor of a partial order on the feature space, called a *claim lattice*, and he uses the claim lattice, not to generate outcomes, but to generate *3-ply arguments* for both sides. Nevertheless, the underlying representation of cases used by Kort, Lawlor and Schubert thirty years ago is essentially the same as the representation used by Rissland, Ashley and Skalak today.⁶

So what? The problem is that a restricted input representation ultimately limits the possibilities of modeling legal argument. HYPO can represent similarities and dissimilarities between cases, and use these in arguments, if and only if the relevant similarities and dissimilarities have been encoded in its set of features. The system is thus relatively successful with “weighing”

⁶This is not true of the most recent work by Kevin Ashley. In a series of papers on moral reasoning written jointly with Bruce McLaren [McLaren and Ashley, 1995][Ashley and McLaren, 1995], Ashley uses a full semantic network that includes representations of (i) the actors (i.e., the truth teller, truth receiver, and others affected by the decision), (ii) the relationships between the actors (e.g., familial, professional, etc.), (iii) the possible actions and (iv) the reasons that support the possible actions. There is obviously much more structure encoded here than in the list of factors that encodes a case in HYPO and CABARET.

and “balancing” arguments, as described in [Skalak and Rissland, 1992] on pp. 25–26, since courts typically write down all the factors to be considered in such arguments, while saying very little about how these factors are to be taken into account. We would not expect HYPO to be as successful with the other argument forms.

Trade secret law in the United States, which is state law, not federal law, is a good example. The original *Restatement of Torts* defined a trade secret to be: “Any formula, pattern, device or compilation of information which is used in one’s business, and which gives him [*sic*] an opportunity to obtain an advantage over competitors who do not know or use it.” But instead of attempting to refine this definition further, the *Restatement* simply added the following explanation:

An exact definition of a trade secret is not possible. Some factors to be considered in determining whether given information is one’s trade secret are: (1) the extent to which the information is known outside of his business; (2) the extent to which it is known by employees and others involved in his business; (3) the extent of measures taken by him to guard the secrecy of the information; (4) the value of the information to him and to his competitors; (5) the amount of efforts or money expended by him in developing the information; (6) the ease or difficulty with which the information could be properly acquired or duplicated by others. *Restatement of Torts*, §757, comment (b) (1939).

This definition, and the six factors, were picked up by the courts, and promptly became the standard in the majority of the states. See M. Jager, *Trade Secrets Law* §5.05 [1985]. Notice that the *Restatement* does not attempt to explain the significance of these factors, or how they fit together, or even whether they favor the plaintiff or the defendant.

Contrast this with the definition of a “trade secret” in the Uniform Trade Secrets Act. (This is a model statute that began to spread through the state legislatures in the early 1980’s, and is now the law in 39 states.) The Uniform Trade Secrets Act takes a more categorical approach:

§1(4). “Trade secret” means information, including a formula, pattern, compilation, program, device, method, technique, or process, that:

(i) derives independent economic value, actual or potential, from not being generally known to, and not being readily ascertainable by proper means by, other persons who can

obtain economic value from its disclosure and use, and

(ii) is the subject of efforts that are reasonable under the circumstances to maintain its secrecy.

This definition has a structure. It has two subparts, each of which contributes a distinct element: “derives independent economic value . . . from . . .” and “is the subject of efforts . . . to . . .”. The first subpart uses two linguistic expressions — “not . . . generally known” and “not . . . readily ascertainable by proper means” — that are intentionally vague, but have a core meaning in ordinary language. The second subpart uses the extremely vague qualifier “reasonable under the circumstances”. Although the components of this definition are certainly *open-textured* [Hart, 1961] [Gardner, 1987], they can be interpreted by reference to the whole, which gives clues about the principles and policies underlying trade secret law.

It is perhaps significant that all the cases in the HYPO data base were decided prior to the enactment of the Uniform Trade Secrets Act. Since the courts in those cases were “weighing” and “balancing” factors in the style of the *Restatement*, the HYPO feature vector representation does a reasonably good job. The initial cases under the Uniform Trade Secrets Act look quite different. See, e.g., *American Credit Indemnity Co. v. Sacks*, 213 Cal. App. 3d 622, 262 Cal. Rptr. 92 (Cal. Ct. App. 1989); *S.O.S., Inc., v. Payday, Inc.*, 886 F.2d 1081 (9th Cir. 1989); *Sheets v. Yamaha Motors Corp., U.S.A.*, 849 F.2d 179 (5th Cir. 1988). The *Restatement* factors have now disappeared, and the courts are struggling to understand the open-textured terms in the statute as part of a general theory of trade secret law. We would not expect the HYPO model to work as well with the argument forms in such cases.

For example, it does not seem possible to use HYPO to represent “slippery slope” arguments at all. In this argument form, the lawyer starts with the facts and the legal conclusion advocated by her opponent, and constructs a sequence of hypothetical cases that leads inexorably to a set of facts that seem to require precisely the opposite result. See [Skalak and Rissland, 1992], pp. 22-23. Although HYPO is able to generate hypothetical cases, as its name suggests, it does so in a very restrictive way. The user can ask the system to vary one of the features in the current fact situation so that the prospects for the plaintiff (or the defendant) are strengthened (or weakened) relative to the current constellation of cases. See [Ashley, 1990], pp. 78-86. This can be useful if the lawyer has control over the facts of her case, as she might at the time she is drafting the complaint, and simply wants to know which facts she should plead. See the discussion in Section 2 above. But at the appellate level, the facts are (es-

entially) fixed. The hypothetical cases in a slippery slope argument have to be generated and tested against some independent legal theory that has been advocated by the lawyer’s opponent, in order for the argument to work at all. Since HYPO does not maintain a separate representation of facts and legal theories, however, it cannot represent this argument form.⁷

CABARET does not seem to be able to represent slippery slope arguments either, but Skalak and Rissland do provide a much broader characterization of legal argument in their general commentary [Skalak and Rissland, 1992]. Discussing the interplay between rules and cases, in which specifications driven top-down by rules are combined with restrictions imposed bottom-up by cases, they write:

This interplay can be viewed as a special kind of theory formation. In mounting an argument, writing a proof, designing a computer program, one is putting forth a theory of how things are, or ought to be; this theory is tempered by actual examples of what is. In the case where the examples function as counterexamples to the posited theory, the tempering can be quite severe. [Skalak and Rissland, 1992], p. 11.

Is it possible to construct a computational theory of legal argument with these properties?

5 The Correct Theory

My own theory of legal argument dates back quite a number of years. It was first published in the collection of papers from the Swansea Conference in 1979 [McCarty, 1980], and followed up shortly thereafter by several papers written jointly with N.S. Sridharan [McCarty and Sridharan, 1980] [McCarty and Sridharan, 1981] [McCarty and Sridharan, 1982]. There was a long hiatus through the late 1980’s and early 1990’s, during which time I worked primarily on knowledge representation problems, but I returned to the subject again two years ago [McCarty, 1995]. The problem all along has been to build a computational model of the arguments that occur in the majority and dissenting opinions in a series of corporate tax cases in the United States dating from the 1920’s and 1930’s, and beginning with *Eisner v. Macomber*, 252 U.S. 189 (1920). The reader should consult the original papers, especially [McCarty, 1995], for the full story. These papers are dense with technical detail, however, and it may be difficult to see in them my overall thesis about legal argument. I will try to explain it simply here.

⁷It remains to be seen whether Kevin Ashley’s more recent work on moral reasoning [McLaren and Ashley, 1995][Ashley and McLaren, 1995] can represent these alternative argument forms.

Legal reasoning is a form of *theory construction*, as Skalak and Rissland say in the passage quoted above. A judge rendering a decision in a case is constructing a *theory* of that case. It follows, then, that a lawyer's job (in an appellate argument) is to construct a theory of the case, too, and one that just happens to coincide with his client's interest. Since the opposing lawyer is also constructing a theory of the case, which coincides with her client's interest, the argument boils down to this: Which theory should the judge accept? There are several constraints here, but they are weak ones. Legal theories should be consistent with past cases, and they should give acceptable results on future cases. They should not be too hard to understand, or too hard to administer, which means that they should be free of major anomalies. One term that is sometimes used to describe these requirements is "coherence". Although difficult to specify precisely, the link between coherent legal theories and persuasive legal arguments seems to be quite strong.⁸

If we are looking for a computational analogue of this phenomenon, the first field that comes to mind is machine learning. Machine learning, broadly conceived, is also a study of the construction of "theories" that classify "cases", so it ought to be relevant. However, the current state of the art in machine learning is not sufficiently advanced to be applied directly to the understanding of legal argument, for several reasons. First, machine learning algorithms work on large data sets, with hundreds or thousands of cases, but legal argument requires only a handful of cases. How is it possible to construct stable theories with such sparse data? Second, machine learning algorithms are based on very simple data representations: e.g., feature vectors, for most algorithms [Quinlan, 1993], or Horn clauses, for the most complex algorithms yet discovered [Muggleton, 1992]. But to adequately represent the cases in a legal knowledge base, we need an extremely rich knowledge representation language [McCarty, 1989]. How is it possible to learn anything in such a complex environment?

Finally, there is a difference that I have tended to

⁸My views about legal reasoning may have shifted slightly since 1979, but not by much. Here is a list of the "Major Research Issues" in the TAXMAN II project from the first slide of my talk at IJCAI'81:

- The most important legal concepts are *open-textured*: they are never logical and precise, but amorphous and poorly defined. How should we represent this structure in an AI system?
- Legal concepts are not static, but dynamic: they are typically *constructed* and *modified* as they are applied to a particular set of facts. How should we model this process?
- What role does the structure and the dynamics of legal concepts play in the process of *legal argument*?

For better or worse, these issues have remained remarkably constant over the years.

emphasize disproportionately in my previous papers, sometimes to the exclusion of everything else. Most machine learning algorithms assume that concepts have "classical" definitions, with necessary and sufficient conditions, but legal concepts tend to be defined by *prototypes*. When you first look at prototype models [Smith and Medin, 1981], they seem to make the learning problem harder, rather than easier, since the space of possible concepts seems to be exponentially larger in these models than it is in the classical model. But empirically, this is not the case. Somehow, the requirement that the exemplar of a concept must be "similar" to a prototype (a kind of "horizontal" constraint) seems to reinforce the requirement that the exemplar must be placed at some determinate level of the concept hierarchy (a kind of "vertical" constraint). How is this possible? This is one of the great mysteries of cognitive science.

It is also one of the great mysteries of legal theory. Most researchers in AI and Law will cite H.L.A. Hart for a discussion of the problems of *open texture* [Hart, 1961], or perhaps cite Ronald Dworkin for the distinction between *rules* and *principles* [Dworkin, 1967]. But these are just minor skirmishes with a very difficult problem. Dworkin's subsequent formulation of the problem, in *Hard Cases* [Dworkin, 1975], imagined a judge named Hercules who could decide cases by constructing a theory of the entire legal system, if necessary, in one supreme act of intellectual will. Is there only one best theory? Dworkin wants to answer: "Yes". But to defend his position, he needs a standard of "coherence" or "integrity" to constrain the possible theories. He tries to articulate such a standard in *Law's Empire* [Dworkin, 1986], by resorting to a metaphor. The integrity of the law is like the integrity of a novel composed in sequence by a chain of novelists, he explains [pp. 225-275]. But this answer just replaces one mystery by an even greater one.

I would be the first to admit that I do not yet have a solution to the problem of coherence, but I think that we should be looking for a computational explanation here rather than a literary metaphor. I suggested a possible approach in [McCarty, 1991], based on the research programme of Marcus [Marcus, 1980] and Berwick [Berwick, 1985] for natural language grammars. Natural language grammars should be (i) easy to parse and (ii) easy to learn, according to this research programme. Likewise, coherent concepts should be (i) *easy to compute with* and (ii) *easy to learn*. The trick is to formalize these ideas within (a) a knowledge representation language that is expressive enough for the legal domain [McCarty, 1989], in which (b) the problematic concepts are defined by prototypes [McCarty, 1995]. In my 1995 paper, I extended my formal model of legal argument up to the edge of the theory of coherence, but no further. The situation is essentially the

same today.

There are other researchers in the field of AI and Law who seem to share my view of legal argument as a form of theory construction. I have already discussed the insightful article by Skalak and Rissland [Skalak and Rissland, 1992], and I have mentioned the recent work of Kevin Ashley [McLaren and Ashley, 1995] [Ashley and McLaren, 1995], which appears to go beyond the limits of HYPO and CABARET (see notes 6 and 7, *supra*). Tom Gordon's work [Gordon, 1995] is usually grouped together with the rule-based theories of legal argument, but his own view is that the "Pleadings Game" is a form of theory construction [pp. 5–6]. Gordon's model of civil pleading is drawn, roughly, from the English common law practice rather than the modern American practice described in Section 2 above. Without further constraints, it would allow the parties to interject arbitrary rules and premises into the proceedings — in a sense, constructing legal theories "on the fly". The game is constrained, however, by the "norms of procedural justice" [p. 6], and the main body of Gordon's work is devoted to formalizing these norms. Karl Branting's recent work on the *ratio decidendi* of a case [Branting, 1994] can be viewed in a similar way. Each "reduction-graph" representation of a legal decision in Branting's analysis is a distinct "theory of the case" which might very well lead to a different result in an argument about a subsequent decision. Although Branting does not offer any criteria for choosing among different reduction graphs (e.g., for choosing between Figure 4 and Figure 5 on pp. 14–15), he does show us how to focus our attention on the facts that are relevant to such a choice, and this is itself an important contribution.

Finally, the central role of coherence in a theory of legal argument is emphasized by Aleksander Peczenik in his commentary in the special issue of *Artificial Intelligence and Law* on "Logical Models of Legal Argumentation" [Peczenik, 1996]. Peczenik strikes a cautionary note here, since he believes that the concept of coherence is inherently "circular" [pp. 323–325]. In the abstract, this is not an insurmountable problem for a computational model, since an apparent circularity can often be replaced by a well-founded recursive definition. So far, however, we do not know how to do this.

6 Conclusion

Although I have criticized much of the recent work in AI and Law for its deficiencies as a *theory* of legal argument, it does not follow that this work is devoid of important practical applications. Ideally, the two goals should coincide: We would like to build intelligent legal information systems on a solid theoretical foundation. But, in practice, simplifications of legal theory

that would be unacceptable to a legal philosopher might still be extremely useful for a system builder.

The work discussed in this paper may be a case in point. Specifically:

- The rule-based theories of legal argument discussed in Section 3 have taught us a great deal about the structure of legal rules. If we used only stratified negation-as-failure with metalevel references in our representation language, we would have a powerful normal form for statutes and regulations.
- The case-based theories of legal argument discussed in Section 4 have also contributed a number of useful representational devices: factors, claim lattices, 3-ply arguments, etc. To the extent that we can describe cases in this way, we have a powerful tool for organizing legal databases.
- The theories discussed in Section 5 are still incomplete, since the coherence criteria are so poorly understood. Nevertheless, research on this problem has forced us to develop extremely detailed knowledge representation languages, tuned to the nuances of legal discourse. These representation languages alone should have important practical applications.

The hope is that all of this work will continue, at a theoretical level, and that additional applications will become apparent as time goes on.

References

- [Ashley and McLaren, 1995] K.D. Ashley and B.M. McLaren. Reasoning with reasons in case-based comparisons. In M. Veloso and A. Aamodt, editors, *Case-Based Reasoning Research and Development (ICCB-95)*, number 1010 in Lecture Notes in Artificial Intelligence, pages 133–144. Springer Verlag, 1995.
- [Ashley, 1990] K.D. Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. MIT Press, 1990.
- [Berwick, 1985] R.C. Berwick. *The Acquisition of Syntactic Knowledge*. MIT Press, 1985.
- [Branting, 1994] L.K. Branting. A computational model of ratio decidendi. *Artificial Intelligence and Law*, 2:1–31, 1994.
- [Delgrande, 1988] J. Delgrande. An approach to default reasoning based on a first-order conditional logic: Revised report. *Artificial Intelligence*, 36:63–90, 1988.

- [Dung, 1993] P.M. Dung. An argumentation semantics for logic programming with explicit negation. In *Logic Programming: Proceedings of the Tenth International Conference*, pages 616–630. MIT Press, 1993.
- [Dung, 1995] P.M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- [Dworkin, 1967] R. Dworkin. The model of rules. *University of Chicago Law Review*, 35:14–46, 1967.
- [Dworkin, 1975] R. Dworkin. Hard cases. *Harvard Law Review*, 88:1057–1109, 1975.
- [Dworkin, 1986] R. Dworkin. *Law's Empire*. Harvard University Press, 1986.
- [Gardner, 1987] A.v.d.L. Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. MIT Press, 1987.
- [Geffner and Pearl, 1992] H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53:209–244, 1992.
- [Gelfond and Lifschitz, 1988] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proceedings, Fifth International Conference and Symposium on Logic Programming*, pages 1070–1080, 1988.
- [Ginsberg, 1987] M.D. Ginsberg, editor. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, 1987.
- [Gordon, 1991] T.F. Gordon. An abductive theory of legal issues. *International Journal of Man-Machine Studies*, 35:95–118, 1991.
- [Gordon, 1994] T.F. Gordon. The pleadings game: An exercise in computational dialectics. *Artificial Intelligence and Law*, 2:239–292, 1994.
- [Gordon, 1995] T.F. Gordon. *The Pleadings Game*. Kluwer, 1995.
- [Hall, 1935] P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 10:26–30, 1935.
- [Hart, 1961] H.L.A. Hart. *The Concept of Law*. Oxford: Clarendon Press, 1961.
- [Horty et al., 1990] J.F. Horty, R.H. Thomason, and D.S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence*, 42:311–348, 1990.
- [Jager, 1985] M. Jager. *Trade Secrets Law*. 1985.
- [Kort, 1963] F. Kort. Content analysis of judicial opinions and rules of law. In G. Schubert, editor, *Judicial Decision Making*, pages 133–197. 1963.
- [Kowalski and Toni, 1996] R.A. Kowalski and F. Toni. Abstract argumentation. *Artificial Intelligence and Law*, 4:275–296, 1996.
- [Lawlor, 1963] R.C. Lawlor. Foundations of logical legal decision making. *Modern Uses of Logic in Law*, pages 98–114, 1963.
- [Lloyd, 1987] J.W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, second edition, 1987.
- [Loui, 1987] R.P. Loui. Defeat among arguments: A system of defeasible inference. *Computational Intelligence*, 3:100–106, 1987.
- [Loui, 1995] R.P. Loui. Book review: Foucault, Derida, Women's speaking justified, and Modelling legal argument. *Artificial Intelligence and Law*, 3:143–150, 1995.
- [Marcus, 1980] M.P. Marcus. *A Theory of Syntactic Recognition for Natural Language*. MIT Press, 1980.
- [Marek et al., 1990] W. Marek, A. Nerode, and J.B. Remmel. Nonmonotonic rule systems. I. *Annals of Mathematics and Artificial Intelligence*, 1:241–273, 1990.
- [McCarty and Sridharan, 1980] L.T. McCarty and N.S. Sridharan. The representation of an evolving system of legal concepts: I. Logical templates. In *Proceedings of the Third Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, pages 304–311, Victoria, British Columbia, 1980.
- [McCarty and Sridharan, 1981] L.T. McCarty and N.S. 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 246–53, 1981.
- [McCarty and Sridharan, 1982] L.T. McCarty and N.S. Sridharan. A computational theory of legal argument. Technical Report LRP-TR-13, Computer Science Department, Rutgers University, 1982.
- [McCarty, 1980] L.T. McCarty. The TAXMAN project: Towards a cognitive theory of legal argument. In B. Niblett, editor, *Computer Science and Law: An Advanced Course*, pages 23–43. Cambridge University Press, 1980.
- [McCarty, 1989] L.T. McCarty. A language for legal discourse. I. Basic features. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 180–189. ACM Press, 1989.

- [McCarty, 1991] L.T. McCarty. On the role of prototypes in appellate legal argument (abstract). In *Proceedings of the Third International Conference on Artificial Intelligence and Law*, pages 185–190, 1991.
- [McCarty, 1995] L.T. McCarty. An implementation of Eisner v. Macomber. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 276–286, 1995.
- [McLaren and Ashley, 1995] B.M. McLaren and K.D. Ashley. Context sensitive case comparisons in practical ethics: Reasoning about reasons. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 316–325, 1995.
- [Moore, 1985] R.C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.
- [Muggleton, 1992] S. Muggleton. *Inductive Logic Programming*. Academic Press, 1992.
- [Nute, 1994] D.N. Nute. A decidable quantified defeasible logic. In D. Prawitz, B. Skyrms, and D. Westerstahl, editors, *Logic, Methodology and Philosophy of Science*, volume IX, pages 263–284. 1994.
- [Peczenik, 1996] A. Peczenik. Jumps and logic in the law. *Artificial Intelligence and Law*, 4:297–329, 1996.
- [Pollock, 1987] J.L. Pollock. Defeasible reasoning. *Cognitive Science*, 11:481–518, 1987.
- [Poole, 1985] D. Poole. On the comparison of theories: Preferring the most specific explanation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 144–147. Morgan Kaufmann, 1985.
- [Poole, 1988] D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.
- [Prakken and Sartor, 1995] H. Prakken and G. Sartor. On the relation between legal language and legal argument: Assumptions, applicability and dynamic priorities. In *Proceedings of the Fifth International Conference on Artificial Intelligence and Law*, pages 1–10, 1995.
- [Prakken and Sartor, 1996a] H. Prakken and G. Sartor. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence and Law*, 4:331–368, 1996.
- [Prakken and Sartor, 1996b] H. Prakken and G. Sartor. Editor’s introduction: Logical models of legal argumentation. *Artificial Intelligence and Law*, 4:157–161, 1996.
- [Prakken, 1991] H. Prakken. A tool in modelling disagreement in law: Preferring the most specific argument. In *Proceedings of the Third International Conference on Artificial Intelligence and Law*, pages 165–174, 1991.
- [Prakken, 1993] H. Prakken. *Logical Tools for Modeling Legal Argument*. PhD thesis, Vrije Universiteit, Amsterdam, 1993.
- [Quinlan, 1993] J.R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993.
- [Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [Reiter, 1987] R. Reiter. Nonmonotonic reasoning. In *Annual Reviews of Computer Science*, volume 2, pages 147–187. 1987.
- [Rissland and Skalak, 1991] E.L. Rissland and D.B. Skalak. CABARET: Rule interpretation in a hybrid architecture. *International Journal of Man-Machine Studies*, 34:839–887, 1991.
- [Sartor, 1991] G. Sartor. The structure of norm conditions and nonmonotonic reasoning in law. In *Proceedings of the Third International Conference on Artificial Intelligence and Law*, pages 155–164, 1991.
- [Sartor, 1994] G. Sartor. A formal model of legal argumentation. *Ratio Juris*, 7:177–211, 1994.
- [Schubert, 1975] G. Schubert. *Human Jurisprudence: Public Law as Political Science*. University Press of Hawaii, 1975.
- [Simari and Loui, 1992] G.R. Simari and R.P. Loui. A mathematical treatment of defeasible argumentation and its implementation. *Artificial Intelligence*, 53:125–157, 1992.
- [Skalak and Rissland, 1992] D.B. Skalak and E.L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial Intelligence and Law*, 1:3–44, 1992.
- [Smith and Medin, 1981] E. Smith and D.L. Medin, editors. *Categories and Concepts*. Harvard University Press, 1981.
- [Touretzky et al., 1987] D.S. Touretzky, J.F. Horty, and R.H. Thomason. A clash of intuitions: The current state of nonmonotonic multiple inheritance systems. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 476–482, 1987.
- [Van Gelder et al., 1991] A. Van Gelder, K.A. Ross, and J.S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38:620–650, 1991.