Artificial Intelligence and Law **1**, 3–44, 1992

Arguments and Cases: An Inevitable Intertwining[1]

DAVID B. SKALAK AND EDWINA L. RISSLAND

Dept. of Computer Science, University of Massachusetts, Amherst, Massachusetts 01003, U.S.A.
Skalak@cs.umass.edu,Rissland@cs.umass.edu

(Received 16 November 1991, accepted 14 February 1992)

Abstract. We discuss several aspects of legal arguments, primarily arguments about the meaning of statutes. First, we discuss how the requirements of argument guide the specification and selection of supporting cases and how an existing case base influences argument formation. Second, we present our evolving taxonomy of patterns of actual legal argument. This taxonomy builds upon our much earlier work on 'argument moves' and also on our more recent analysis of how cases are used to support arguments for the interpretation of legal statutes. Third, we show how the theory of argument used by CABARET, a hybrid case-based/rule-based reasoner, can support many of the argument patterns in our taxonomy.

Key words: argument, case-based reasoning, mixed paradigm, statutory interpretation

1. Introduction

Good supporting cases make good arguments. Selecting the best cases possible is crucially important to advancing one's interests, especially in an adversarial domain such as law that requires advocates to support their positions with previous cases. Deciding precisely what constitutes a good, better, or best case requires addressing a variety of theoretical and pragmatic questions.

In this article, we raise and address issues of case selection in the context of statutory interpretation, which we view as argument about the application of a stated legal rule to a new problem situation. Some of the central concerns of statutory interpretation are addressed by our attempt to provide a computational model for selecting the cases that can be used to argue for a particular interpretation of a legal rule. For instance, we distinguish reasoning about the necessity and sufficiency of the antecedents in a rule from reasoning about the meaning of the words used therein. As an example of the former, argument over the status of the antecedents in a rule, consider an advocate whose client has satisfied all but one of a rule's prerequisites, but who nonetheless wants to receive a benefit provided by the rule. The advocate must argue, ideally with the aid of supporting

[1] This work was supported in part by the National Science Foundation, contract IRI-890841, the Air Force Office of Sponsored Research under contract 90–0359, the Office of Naval Research under a University Research Initiative Grant, contract N00014-87-K-0238, and a grant from GTE Laboratories, Inc., Waltham, Mass.

cases, that the unsatisfied prerequisite is not strictly necessary to satisfy the rule's requirements. On the other hand, statutory interpretation often takes the form of argument about the meaning of a rule's individual terms. For example, 'Is taxpayer's home office his "principal place of business," as that term is used in Section 280A(c)(1)(A) of the Internal Revenue Code?' Arguments about the status of rules, the necessity and sufficiency of their conditions, the presence or absence of tacit exceptions and requirements, and the scope and meaning of ingredient terms is the heart of statutory interpretation. These and related issues have been considered in depth by legal philosophers from a spectrum of jurisprudential positions, such as Llewellyn, Hart, Fuller, Twining and Miers [Llewellyn, 1989], [Hart, 1958], [Fuller, 1958], [Twining & Miers, 1982]. Their treatments have provided us with such classic examples of statutory ambiguity as Hart's 'No vehicles in the park' and 'Take off your hat when entering a church.'

To put the matter simply, the best case for your argument about the meaning of a legal rule is a case that is available to you, that is similar to yours, and that went your way. Providing computational definitions of 'similar' and 'went your way' is one theme of this article. In trying to locate such cases, one must temper a specification of the ideal supporting case with the cases that are actually present in a case library. The characterization of the best case that can be found is a mixture of 'case-base-dependent' and 'case-base-independent' influences. Varying the proportions of the mixture yields a continuum of possible approaches to how to find cases that support a legal rule interpretation.

Although neither extreme approach is pursued in its pure form in legal practice, it may be instructive to observe that each approach is manifested in the way attorneys sometimes work. Consider the task of writing a legal brief. At one extreme, one could specify the ideal case independently of what cases are in the case base and then try to find that ideal case. (In the absence of an actual precedent meeting one's specifications, one could create a hypothetical case that makes the point.) One would specify 'top down' the features the ideal case must have and perhaps, must not have,[2] and which features may be ignored.[3] This case-base-independent approach is at work when a senior litigator leaves space in a brief for supporting citations for well-established propositions of law, and then directs her associates to go to the library to fill in the placeholders with citations to actual cases.[4]

At the other extreme approach, case-base dependence, one can determine what the best cases are through exhaustive examination of the case base, choosing cases according to some construction of 'best.' This is a case-driven, 'bottom-up', approach to best case

---

2. In effect, this approach attempts to reduce the problems of partial matching of cases to the simpler one of a complete match. The problems of partial matching intrude on this attempt to require 'complete matches,' however. If one specifies only those features that the ideal case is required to have, one must state also whether additional features must be absent or can simply be ignored.

3. Treatment of any unspecified or unanticipated features implicitly entails a decision whether — and to what extent – to adopt the closed world assumption. See also [Gibbons, 1991].

4. This specification-driven approach may be used to plan a sequence of 'deformations' of a current fact situation to yield a series of hypotheticals. See [McCarty & Sridharan, 1982]. The current case may be linked with an ideal case in such a 'connect-the-dots' manner. Of course, planning the intermediate hypothetical case points involves subtleties from each of the aspects of legal argument we will enumerate. For example, 'How much will the court tolerate in the way of deformations?' would be an important jurisprudential question.

selection. A best case is best only relative to the cases actually present in the case base. The case-base-dependent approach is evident in brief writing when the legal propositions advanced are actually derived from generalizations of the common law contained in the cases. In particular, the dependence of argument on the constituents of the case base is clear where argument is restricted to cases from a given jurisdiction, time period or court.

Of course, brief-writing, like many other argument tasks, is usually a mixture of both approaches. Empty placeholders for string cites to cases stating well-accepted propositions of law may be left for routine completion, but more tenuous and arguable propositions often must be derived bottom-up from relevant cases.

Selecting the best cases to use in an argument actually requires confronting a variety of aspects of legal argument: precedential aspects, rhetorical aspects, doctrinal aspects and jurisprudential aspects.

- Precedential aspects of argument include an analysis of the current problem (the 'current fact situation'), cases from the current case base and their inter-connections, the range of possible hypothetical variations on cases, and approaches to determining case similarity.
- Rhetorical aspects include persuasive strategies, such as the short-term strategy of how to argue the current fact situation, the repertoire of available argument techniques and forms, and any long range strategy of how to argue future cases addressing the same issue.
- Doctrinal considerations include the legal claim at issue and its elements, the recognized factors in the domain, and any evidence of doctrinal drift.
- Lastly, jurisprudential aspects include a model of what constitutes a *bona fide* argument, a model of the adherence required to decided cases ('*stare decisis*'), a theory of the relative pedigree or importance of cases, and a recognition of the varying procedural postures of precedents.

In this article, we will limit our attention to precedential aspects of legal argument, particularly on collecting, organizing, and implementing detailed argument patterns for using cases to support rule interpretation.

1.1 ORGANIZATION OF THE ARTICLE

We explore in this article how case-base-independent and case-base-dependent approaches influence each other. We provide a theoretical framework for combining them and present an empirical collection of argument patterns, observed in practice, and show how the two approaches manifest themselves in the patterns. We have applied this framework in CABARET, a domain-independent architecture that heuristically combines rule-based and case-based reasoning [Rissland & Skalak, 1991].

The desire to understand legal argument, particularly precedential aspects of statutory law, has motivated much of our work on CABARET. While we usually have viewed CABARET as a primary example of a 'hybrid' or 'mixed paradigm' reasoning shell – integrating the rule-based and case-based reasoning 'paradigms' – the system is also a

study of how to use available cases and rules to formulate arguments. Ultimately, both views of CABARET fundamentally concern the issue of program control: how the program decides what to do when.

We make several related points in this article.

1. A computational model of case-based legal argument must account for the complex control problems involved in simultaneously selecting cases and making arguments based on them.
2. Legal argument in practice often takes stereotypical forms, and we present a partial list of argument patterns found in legal practice.
3. These stereotypical forms may provide a partial solution to the problems of point (1) by providing top-down direction as to what types of cases to look for next and how to argue with them once found.
4. Primitive argument techniques already incorporated in CABARET are useful for implementing such stereotypical argument forms, even though CABARET's current implementation does not necessarily produce output in the form of the patterns we inventory.

We anticipate that such stereotypical patterns of argument could be used to extend CABARET and support the control core of a new hybrid architecture.

We interleave these points by organizing the paper in the following way. We complete Section 1 with a a brief survey of previous work on computational legal argument. Section 2 is a general introductory discussion of how 'specification-driven' (top-down, case-base-independent) and 'case-driven' (bottom-up, case-base-dependent) considerations influence argument. Our view is that these influences are intertwined, with each providing feedback to the other.

In Section 3 we provide some background on CABARET and its current domain of the so-called 'home office deduction' [Rissland & Skalak, 1991]. In Section 4, we present our partial inventory of argument patterns found in actual legal practice; we provide several examples of each. Section 5 describes CABARET's implemented approach to argument-generation in some detail, but the general point is that CABARET uses a control strategy incorporating both top-down and bottom-up processing to generate skeletal legal arguments. We conclude Section 5 with an example of a straightforward argument generated by CABARET.

In Section 6 we discuss work slated for the immediate future. There we demonstrate how some of the less straightforward argument forms discussed in Section 4 could be implemented using CABARET's framework. Section 7 retraces our steps in an attempt to remind ourselves where we have been and where we would like to go in future work.

1.2 PREVIOUS WORK ON COMPUTATIONAL LEGAL ARGUMENT

Past work in artificial intelligence (AI) and in law has addressed the precedential, rhetorical, doctrinal and jurisprudential aspects of legal arguments.[5] Precedential aspects have been addressed by a growing community of AI-and-law researchers, including [Ashley, 1990, 1991; Ashley & Rissland, 1988a], [Branting, 1989, 1991], [Bellairs, 1989],

[Berman & Hafner, 1991], [Gardner, 1987], [Goldman, Dyer & Flowers, 1987], [McCarty & Sridharan, 1982], [Stucky & Gidley, 1990] and ourselves [Rissland et al., 1984, 1985; Rissland & Skalak, 1991; Skalak & Rissland, 1991].

Much of the present work builds upon the model of case-based reasoning of the HYPO program by Ashley and developed in a series of papers by Rissland and Ashley. Cases in HYPO are represented at two levels: *factual features*, which represent the input description of a case; and *factual predicates*, which are derived features computed from the factual features. The second level is often called the 'interpretation frame' level [Ashley, 1990]. Cases are contained in a case knowledge base (*CKB*), which is an unstructured memory. Indices for retrieving cases from the CKB are called *dimensions* [Rissland, Valcarce and Ashley, 1984], [Ashley, 1990]. Dimensions capture the observation that cases are usually analyzed and argued about using a set of important domain-specific factors, and that a case's strength or weakness with respect to each factor can be assessed. In effect, dimensions are the thread of argument running through a 'line of cases' that addresses a particular issue in a similar way. Dimension-based analysis draws on both general knowledge of the domain and past, specific cases to which the factor applied. Each dimension has a set of prerequisites that determine its applicability – these encode the facts necessary for mounting an argument with respect to this dimension – and a means for comparing cases along it. In HYPO, the metric for assessing case relevancy ('on-pointness') is based on the intersection between the sets of dimensions applicable to the problem case and those applicable to a case from the CKB. Intuitively, the larger this overlap, the more on-point a precedent is. However, HYPO does not rely merely on the number of dimensions in this intersection. HYPO partially orders cases according to the precise subset of dimensions in common with the problem case into a 'claim lattice' [Ashley & Rissland, 1988a], [Ashley, 1990]. Construction of this partial order gives HYPO the ability to recognize the lines of cases bearing on a specific subset of relevant dimensions: they are linear suborderings within the claim lattice partial ordering. For a detailed exposition of HYPO's model of case-based reasoning and its implementation, see [Ashley, 1990].

The CABARET system incorporated into a hybrid architecture a case-based reasoner that uses many of the same data structures developed in HYPO (e.g., case, interpretation-frame, dimension and claim lattice) and also takes some of its procedural mechanisms from HYPO (e.g., initial analysis of the input case resulting in an interpretation frame, indexing of cases by dimensions, and sorting of cases by similarity into a lattice.)

In our own past work, the second author has compiled a collection of strategic argument moves, such as 'bolstering,' 'obfuscating,' 'red herring,' and 'slippery slope'

---

5. Some of these aspects really deserve more attention than research has given them, however. For instance, how the argument made in the current fact situation affects and is affected by long range strategic goals has not been addressed by computational approaches to argument, as far as we know. Strategic considerations would be especially important for a specialist litigating several cases in one area of law. Existing work might be extended to this end, however. For instance, McCarty and Sridharan's work on prototypes and deformations [McCarty & Sridharan, 1982] may be used to plot a sequence of cases. HYPO [Ashley, 1990] could be used to answer long range planning questions such as, "If this case were decided for the claimant, how will that affect the complexion of our backlogged cases in light of the resulting case base?".

[Rissland, 1985]. The current article builds on that foundation by identifying additional argument forms actually used by practitioners and providing explicit computational mechanisms for effecting some of them. In recent work, [Rissland & Skalak, 1991; Skalak & Rissland, 1991], we have attempted to tease statutory argument into a tiered set of 'strategies' and 'moves,' which are ultimately implemented through combinations of generic argument 'primitives,' such as analogizing and distinguishing.

Branting has provided an alternative framework for integrating cases and rules, particularly to improve case matching through rule-based reasoning. Branting's program embodying this framework ('GREBE') uses a detailed semantic network representation of each of the case arguments in his case base. This article takes a wider, more taxonomic perspective for argument that does not assume the presence of detailed relational representations for case opinions. Similarly, Bellairs's extensive work on analogical reasoning in the law assumes a relatively strong domain theory, where 'significant relationships between the facts [of problem situations and of case situations] are known.' [Bellairs, 1989, p.50]. Berman and Hafner have argued that the procedural posture of a case is a crucial aspect of reasoning with precedents and must be incorporated into a model of legal argument. Goldman, Dyer and Flowers have emphasized memory organization and case representation in their research modeling aspects of contract law. Stucky and Gidley point out that pragmatic considerations for legal argument have been overlooked in current models. See generally [Rissland, 1990] and [Stucky, 1986] for overviews of computational legal argument. Stucky also argues for a top-down and bottom-up control strategy for argument creation.

Other AI models of legal argument have been put forth in seminal work by McCarty and by Gardner. McCarty's longstanding work on legal knowledge representation models legal concepts as 'prototypes plus deformations,' which are logical templates incorporating necessary conditions for a concept, a set of examples inside or outside the concept class, and a collection of transformations that map one example to another. Gardner's approach to legal reasoning and her program have influenced the design of our CABARET program in several ways. First, Gardner used several simple heuristics to determine when a case was 'hard,' such as a failure to establish a result using rules only, or a conflict between rules or between cases. Second, her program made use of positive and negative[6] examples of predicates used in the rules, (although the examples were abstract fact patterns, rather than particular cases). Third, Gardner's work incorporated both reasoning with rules and reasoning with these abstract examples. On the other hand, the implementation of Gardner's model did not create arguments, and all of the cases in CABARET's case base would be 'hard' according to the characterization used by her model and program.

---

[6] The terms 'positive' and 'negative' are used in the sense given by machine learning. A positive example of a concept is a member of that concept class; a negative example is not a member of the class. In the legal examples we will give, the relevant concept is often implicit: the category defined by the result desired by the advocate. For example, if the advocate argues to receive a deduction for an office in a home, then the category is 'deductible home office expenses.' See [Skalak, 1989] for a discussion of models of classification applied to the law.

The literature on jurisprudential aspects of argument is vast, and since we concentrate on other aspects of argument, we point arbitrarily to a few of our favorite sources: [Levi, 1949], [Llewellyn, 1989] and [Twining & Miers, 1982]. [Kennedy, 1989] presents a taxonomy of argument types at a level of abstraction higher than that of our classification. [Perelman & Olbrechts-Tyteca, 1969] and [Toulmin, 1958] are classic and encompassing works on argument generally. Several researchers have used Toulmin's analysis as a basis for their own work [Marshall, 1989; Storrs, 1991].

1.3 PREVIOUS WORK ON COMPUTATIONAL ARGUMENT IN GENERAL

While our focus in this article is the task of creating legal arguments, many AI researchers have made contributions to the computational theory of argument applied to a variety of task domains, including [Alvarado, 1990, 1991] (editorial comprehension), [August & McNamee, 1991] (editorial comprehension), [Birnbaum, Flowers & McGuire, 1980; McGuire, Birnbaum & Flowers, 1981; Flowers, McGuire & Birnbaum, 1982; Birnbaum, 1982, 1985] (argument understanding and response), and [Wu & Lytinen, 1991] (advertisement comprehension). While Clark [1988] and Sycara [1985, 1987, 1989, 1991] have not aimed their models at the legal domain, each has made contributions to the theory of case-based argument. Clark has presented a formal, logic model of argument that represents background domain knowledge as sets of arguments for or against a problem hypothesis. Models for argument in a distributed, multi-agent setting have been presented by Sycara, in which arguments are constructed with both case-based and decision-theoretic, utility-based support.

Of this work on argument outside the legal domain, that of Birnbaum, Flowers and McGuire is closest to our own approach. Their research recognized that argument creation requires a mixture of bottom-up and top-down control and tried to identify higher-level argument structures ('molecules' in their terminology) and tactics. However, this line of research apparently recognized only two higher-level argument forms. The tactics they suggest are more abstract than our argument moves, and appear tantamount to such strategies as attacking the 'warrant' and attacking the 'backing' of a claim, to use Toulmin's classic terminology [Toulmin, 1958].

2. Intertwining case specification and argument

While one can hypothesize about the ideal case for an argument independently of an existing case corpus, the cases that are actually available temper this vision.[7] To a lesser extent, what hypothetical cases are reasonable also constrains argument.[8] We maintain

---

7. We assume that the 'available' cases are contained in a 'case base.'

8. Much of the past work of our research group has focussed on hypotheticals and the important role accorded them in case-based reasoning [Rissland & Ashley, 1986], [Ashley, 1990]. Nonetheless, in legal argument one cannot rely on hypotheticals for authoritative statements of existing law. One can use them, however, to illustrate a point, poke holes in an argument, or show the ramifications of a line of reasoning. We shall treat 'cases' as including hypotheticals.

that computational argument with cases requires at least two types of processes: 'specification-driven' and 'case-driven'.

2.1 TWO ROUTES TO THE THE IDEAL CASE

By 'specification-driven' we mean that the arguer specifies what the ideal case for his situation might be, independently of what cases exist in the case base. By 'case-driven' we mean that the arguer creates and defines arguments according to what cases actually exist and are available to him.

In the law, it is certainly the case that both specification-driven and case-driven approaches are at work. The doctrine of *stare decisis* trivially requires that an advocate must at a minimum be case-driven. Cases must be cited and used according to the modes permitted by the prevailing institutional interpretation of the doctrine of precedent. On the other hand, an arguer is also driven by the need to establish certain goals, for instance, that some requirement of a relevant legal rule has been satisfied. The needs of argument define the desiderata for the perfect case. Thus, an arguer – computer system or human litigator – must interleave at least these specification-driven and case-driven processes. One can start from either vantage point but feedback will be given and direction exerted by the other.[9]

A system that uses a bottom-up approach is HYPO [Ashley, 1990], where case selection is driven by cases in the case base. In HYPO, a 'best case' for a user is a 'most on-point' case that is decided for his point of view (his side), and shares a factor with the current fact situation that is favorable to his point of view [Ashley, 1989]. A 'most on-point' case is maximal in the ordering of cases by similarity to the problem situation. HYPO does have a specification-driven aspect in that it limits the factors on which the similarity ordering is based to those that are present in the current situation. To that extent, a best case is implicitly specified (top-down) as one possessing the same applicable dimensions as the problem. But primarily HYPO works bottom up from the available case base by percolating best cases from the case base towards the upper levels of the similarity ordering, which contain the cases maximally similar to the current fact situation.

In a new tutoring system based on HYPO, Ashley and Aleven are re-implementing case-based argument terms such as 'best case' and 'most on-point case' as defined and manipulable object relations, rather than as the output of a particular procedure [Ashley & Aleven, 1991]. This change in representation from HYPO is designed to enable their system to function top-down: to retrieve argument fragments from a specification of the relations that must be present between the cases involved.

A mixture of case-driven and specification-driven approaches is evident in the way CABARET works. CABARET uses a tiered framework of argument strategies, moves and primitives, which sets up retrieval specifications for precedents to be used in

---

9. For a description of how lawyers alternate between constructing a theory of a case and assessing its emerging facts, see [Gibbons, 1991] and [Morris, 1937].

argument. CABARET's framework lays out specifications for potentially useful classes of cases and ways to use them in argument. After retrieving cases with the required characteristics, the system's case-based reasoning module, which is modelled closely after HYPO, sorts those cases, compares them, and selects the cases to use. Thus, the argument moves and strategies in this framework specify what the ideal case would be for a particular argument strategy-move combination. The cases actually found then determine which of the specific moves can be used to implement the argument strategies given the limitations of the current case base. Details of CABARET's tiered argument framework are given in Section 5.

The need to mix specification-driven and case-driven approaches is present in domains outside the law as well. For instance, in developing computer programs, specifying the program and developing concrete code are intertwined [Swartout & Balzer, 1982]. Mathematics interleaves deductive reasoning and reasoning with examples [Polya, 1973; Lakatos, 1976; Rissland et al., 1984]. We believe the interplay of reasoning with cases with other modes of reasoning is ubiquitous.

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.

3. Background on CABARET and the home office deduction

CABARET is a mixed paradigm reasoner that combines case-based and rule-based reasoning. The model of case-based reasoning used is based upon HYPO; the rule-based reasoning used involves the standard forward and backward inference methods. CABARET couples these two styles of reasoning with control heuristics that post to an agenda tasks for each of the reasoners to perform. Some of these heuristics encode various techniques used in statutory argument. For instance, if after attempting to satisfy the antecedents of a rule, one finds that only one of the rule's antecedents is unsatisfied, then one can try to argue that the missing condition is not necessary. This argument would be supported by the case-based task of finding and reasoning with cases where there was a similar 'near miss' but where the result of the rule was obtained nevertheless.

CABARET has several components, including a case-based reasoning module and a rule-based reasoning module, and for each, a dedicated monitoring process that makes observations on the progress and achievements of the individual reasoners. These observations are recorded in a 'control description' language that is also used to encode the control heuristics. The control module applies these heuristics, which are encoded as rules, to the observations to determine the next tasks to be posted to the agenda. A full description of CABARET's architecture and a detailed discussion of the control heuristics is provided in [Rissland & Skalak, 1991].

One of the domains in which the CABARET shell is currently instantiated is an area of U.S. Federal income tax law referred to as the 'home office deduction.' The home

office deduction domain is governed primarily by Section 280A of the Internal Revenue Code, and we have focused on Section 280A(c)(1), which contains the heart of the statute:

> [A deduction may be taken for] any item to the extent such item is allocable to a portion of the dwelling unit which is EXCLUSIVELY USED on a REGULAR basis –
>
> (A) [as] the PRINCIPAL PLACE OF BUSINESS for any trade or business of the taxpayer,
>
> (B) as a place of business which is used by patients, clients, or customers in MEETING OR DEALING with the taxpayer in the normal course of his trade or business, or
>
> (C) in the case of a SEPARATE STRUCTURE which is not attached to the dwelling unit, in connection with the taxpayer's trade or business.
>
> In the case of an employee, the preceding sentence shall apply only if the exclusive use referred to in the preceding sentence is for the CONVENIENCE OF HIS EMPLOYER.
>
> *[capitalization supplied]*

The home office deduction deals with the circumstances under which taxpayers may legitimately deduct on a U.S. Federal income tax return expenses relating to an office maintained at the taxpayer's residence. We have supplied capitalization to certain terms in the statute to emphasize their role as *statutory predicates*, important words or phrases on which the meaning of the statute turns.

CABARET's case base in the home office deduction domain currently contains representations of 23 actually litigated tax cases. In addition, there are 6 hypothetical cases in the case base. To give the flavor of this area, several examples of litigated cases follow.

- The case of David Weissman, *Weissman v. Commissioner*, 751 F.2d 512 (2d Cir. 1984) who was a professor of philosophy at City College in New York City. Although he was provided with a shared office at City College, it was not 'a safe place to leave teaching, writing, or research materials and equipment,' [Court of Appeals Opinion, p.513, quoting the lower Tax Court opinion.] In his 10-room apartment, Professor Weissman maintained a home office, consisting of two rooms and adjoining bathroom. He estimated working between 64 and 75 hours each week, and spent 80% of that time writing and researching in his home office. The Internal Revenue Service challenged Professor Weissman's deduction of $1540 of rent and other expenses relating to his home office. The IRS claimed that Weissman's home office did not satisfy the requirements of the statute, in particular that it was not his principal place of business and that it was not for the convenience of his employer, City College. On appeal of a decision against Weissman, the Court of Appeals for the Second Circuit determined that in these circumstances, the home office could be considered Weissman's principal place of business, especially where most of his work was done at home and a home office was necessitated by the lack of suitable workspace on campus. Under *Drucker*, below,

Weissman also satisfied the convenience of employer requirement for the deduction.

- The case of Earnest Drucker, *Drucker v. Commissioner*, 715 F.2d 67 (2d Cir. 1983), who was a violinist with the Metropolitan Opera Orchestra, concerns the 'convenience of employer' requirement for employees (among other requirements of Section 280A). Drucker was not provided with a practice room at Lincoln Center (or elsewhere, for that matter). Drucker and others in the orchestra therefore maintained home practice studios. Practice was absolutely necessary to the fulfillment of the musicians' responsibilities. The Second Circuit held that the home practice areas were for the convenience of the Met, since they relieved the Met from providing the necessary practice space for the musicians to do their jobs.
- The case of Yolanda Baie, *Baie v. Commissioner*, 74 T.C. 105 (1980), a woman who operated a foodstand (the 'Gay Dog') near her residence in Los Angeles, also addresses the 'principal place of business' requirement. Appearing for herself in Tax Court (and apparently doing an excellent job), Ms. Baie argued that – since she used her kitchen for preparing food, particularly hot dogs for her stand – her kitchen (and not the stand) was her principal place of business. In denying the deduction, the court used the 'focal point test', which looked to the 'focal point' of the taxpayer's activities, which was determined to be the Gay Dog itself.

We return in Sections 5 and 6 to discuss CABARET's approach to modeling argument. As a prelude to that discussion, we first consider some of the typical forms that legal argument actually takes, with a view to determining the types of cases that are required to drive them.

4. Recognized forms of legal argument: a partial inventory

Legal argument often takes stereotypical forms. The law student or legal scholar who has read the usual large volume of legal cases can often recognize immediately the overall form of an argument. Some easily-recognized forms include 'slippery slope arguments', 'balancing arguments' and 'make weight' arguments. Occasionally these arguments involve the creation of hypothetical cases, if real cases are not available to meet the requirements of the argument form. Slippery slope arguments, for example, require a sequence of hypothetical cases. Some of the argument forms we inventory apparently have not been named (or previously characterized). We ask the reader's indulgence where we have supplied our own neologisms.

We begin to make our inventory a taxonomy by considering the following dimensions on which arguments may be placed, which are based in part on the cases used in support:

- whether the cited case is a favorable (positive) example or an unfavorable (negative) example for the advocate making the argument;
- whether the cited case is a strong one or a weak one for the arguer;
- whether the argument is used offensively, to establish a result, or defensively, to respond to an opponent's case citation;
- whether a single case or a set of cases is used to complete the argument; and
- if more than one case is used, whether the set of cases is ordered in some way.

It will be useful to have certain additional terminology at our disposal. In this paper we use the term *disposition* to refer to the outcome of a case on the ultimate result argued for, regardless of whether that result was supported by a rule in issue. The disposition of a personal income tax case might be whether a desired deduction was allowed or denied, for example. When we want to refer explicitly to the conclusion warranted by a specific rule, we will refer to the rule's *consequent*. Finally, we will refer to the *status* of a rule with respect to a precedent and mean whether the rule's preconditions were or could be satisfied under a strict application of the rule to the precedent's facts.

We use the word 'rule' to encompass a legal statute, a regulation or other rule promulgated pursuant to a statute, a common law rule derived from a group of cases, or the rule of the case stating a rule of law implicit in a single case. In general, we also intend that 'rule' include rules of thumb that are heuristically gleaned and applied by experts in order to characterize the law in a particular domain. 'Teachers often have difficulty satisfying the principal place of business requirement of the home office deduction,' is an example.

The distinction between 'disposition' and 'consequent' can be subtle and an anticipatory example may be helpful. Consider a rule in the Internal Revenue Code (the 'Code') that entitles one to a deduction of a particular kind, for example, a deduction for certain travel expenses under Section 162. The consequent of the rule may be expressed as 'taxpayer is entitled to a travel expense deduction.' The ultimate disposition desired by the taxpayer is that he is entitled to a travel deduction by law, regardless of whether it is sanctioned by this particular Code rule or by some other justification (e.g., another Code rule or even a Supreme Court case holding that a social policy is more important than strict compliance with Section 162). The taxpayer aims to argue his way into the concept class 'taxpayer entitled to a travel expense deduction' named in the Code rule, by any warrant available. Due to the presence of alternative ways of arguing for a desired result, establishing a particular rule consequent may not be a *necessary* condition for obtaining the desired case disposition.

This distinction between rule consequent and disposition can cut against taxpayer's case for a deduction, of course. There may be impediments to his legitimately receiving the deduction, notwithstanding the satisfaction of the Code section (e.g., a more powerful controlling statute with an opposite result, or a case invalidating the rule on constitutional grounds). Thus, establishing the rule consequent is not always a *sufficient* condition for obtaining the desired case disposition. This conflict between a rule result and other norms dealing with the rule consequent can be the source of 'hard' questions as that term is used in jurisprudence (see, e.g., [Gardner, 1987]).

We will press on with examples of the forms of argument with reference to the following hypothetical situation:

> SITUATION: Alice, a single parent of two young children, is a junior faculty member at an urban law school. To obtain affordable housing and decent schooling for her children she lives in a suburban community located 25 miles from her office. Crowded highways extend the one-way trip commuting time to approximately one hour during rush hours from the half hour trip at other times. Given the lengthy commute and her parenting responsibilities, Alice asked her dean whether she could spend two days a

week working at home since she is able to do much of her research by accessing online legal databases from home. She ensured she would be available for faculty, committee, and student meetings. Her dean agreed and asked whether, given enormous space constraints, she would mind sharing her office with another junior faculty member. Alice readily accepted that arrangement. The evidence shows that it became difficult for Alice to conduct her research while at the school. Sharing an office interferes with her concentration and the library does not provide adequate computing facilities. The evidence also shows that several other junior faculty members have been required to share offices.

Alice converted a bedroom into a study. She uses the study mostly for her work as law professor except on rare occasions when she has house guests who sleep on a couch in the study. She also uses the study for some personal correspondence and paying her bills. Alice claimed a deduction for the costs attributable to her new study. The Commissioner of the Internal Revenue Service ('Commissioner') assessed a tax deficiency, which the Tax Court sustained. Alice has appealed.

In the following sections, we describe each argument form, give an example of its use in Alice's case and finally provide an instance of its use taken from an actual judicial opinion. A suggestive illustration is also provided for most forms.

4.1 ARGUMENTS INVOLVING THE CURRENT FACT SITUATION AND A SINGLE CASE

The following arguments typically involve reasoning with the current fact situation and one other case. If one thinks of these argument forms as procedures, their inputs are the current fact situation and another case, and the output is an argument or some data structure from which an argument in natural language can be prepared.

4.1.1 Straightforward argument. We suppose that a lawyer's first impulse is to analogize cases that have the desired disposition on the particular point in issue and were decided in what would be his favor. Also, he would be prepared to distinguish cases that were decided oppositely both as to the point he is trying to establish and as to ultimate disposition of the case. *Analogizing* and *distinguishing* are the fundamental processes of straightforward argument. Each process typically entails taking the current fact situation and analogizing it to or distinguishing it from a single case at a time.

EXAMPLE: In the hypothetical situation given, Alice may analogize the *Weissman* case described in Section 3: 'This case is governed by the *Weissman* case. In *Weissman*, a faculty member was prevented from working on campus by practical considerations, such as an inadequate on-campus office and inappropriateness of the library facilities, and the principal place of business of the taxpayer was found to be the home office. The deduction was allowed in *Weissman* and so should it be in my case.'

The Commissioner of the IRS might distinguish *Weissman*: '*Weissman* is distinguishable from Alice's case in that (1) the home office is not for the convenience of Alice's

employer, whereas Weissman satisfied the statute's convenience-of-employer requirement; (2) she only uses it 40% of the time, whereas Weissman used his office 80% of his working time; and (3) Alice's use of the home office is not exclusively related to her faculty position, whereas Weissman performed no other business in the home office.'

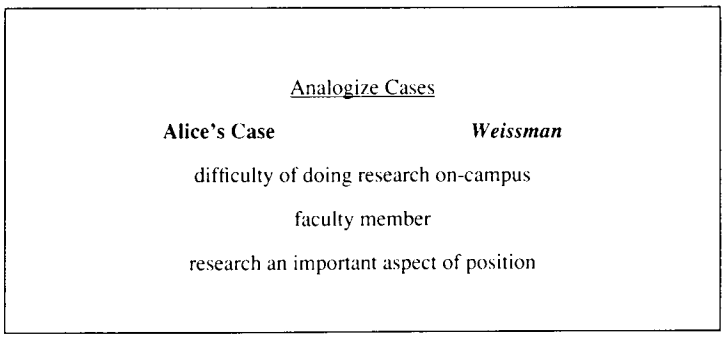

Fig. 1: Common features that may be used to form an analogy between Alice's case and *Weissman*.

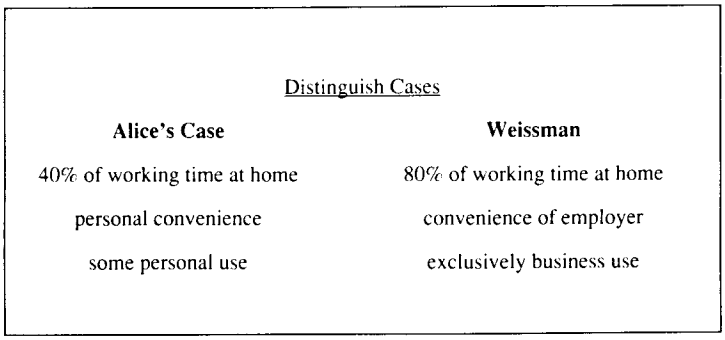

Fig. 2: Contrasting features that may be used to distinguish Alice's case and *Weissman*.

There are many ways to make an analogy between Alice's case and *Weissman*. For example, one can draw analogies via shared ground-level features (e.g., both Alice and Weissman were provided with shared campus offices), important derived, domain-level factors (e.g., both spent a substantial amount of time in their respective home office), the taxonomic class of the parties (e.g., both Alice and Weissman are faculty members), or shared functionality (e.g., the home office supports research).

The implicit model of analogy used by HYPO relies on intersections of domain factors present in the two cases being compared. HYPO's model makes a distinction between the ground-level features at which problem cases are input and derived features, 'factual predicates,' that are used as prerequisites for the presence of important domain factors.

CABARET extends HYPO's model of analogy in several small ways. Term hierarchies are used in CABARET to store domain knowledge that is not conveniently represented as rules or as cases. Under certain conditions CABARET's control heuristics

direct its case-based reasoning component to search these taxonomic hierarchies so that cases involving related terms in the hierarchy can be retrieved. CABARET also expands HYPO's claim lattice datatype; it was re-implemented in CABARET's case-based reasoning component to admit other similarity metrics, to use an improved case sorting algorithm, and to associate a lattice for each predicate in a rule.

Also, to find analogous or dissimilar cases, CABARET can apply case similarity metrics other than HYPO's 'maximize overlap' metric. For instance, CABARET can use the 'minimize differences' metric that minimizes the set of factors that are present in either case but not in the other.

In addition, CABARET also applies control rules that heuristically change from one similarity metric to another. These control heuristics examine the similarity ordering of cases created using one metric and, if various anomalies are found in the lattice reflecting that ordering, the system will apply another similarity metric in an attempt to ameliorate the anomaly. Example anomalies include (1) a trivial similarity ordering consisting of one node, (2) a surfeit of most-on-point cases, and (3) an approximately equal distribution of outcomes of the most-on-point cases between the two sides in a litigation.

EXAMPLE: An actual example of straightforward legal argument is one that HYPO has successfully re-created. This example is taken from [Ashley, 1988] and is discussed in detail there and in [Ashley, 1990]. The example is useful in that it applies both analogizing and distinguishing as straightforward tactics. In this case, which deals with misappropriation of trade secrets regarding a rivet-making machine, *USM Corp. v. Marson Fastener Corp.*, 379 Mass. 90 (1979), ('USM'), the primary issue was whether the plaintiff had taken adequate measures to protect its trade secrets. The Massachusetts court set out the relevant factors explicitly, quoting Kubik, Inc. v. Hull, 56 Mich. App. 335, 356 (1974):

> Relevant factors to be considered include (1) the existence or absence of an express *agreement restricting disclosure*, (2) the nature and extent of *security precautions* taken by the possessor to prevent acquisition of the information by unauthorized third parties, (3) the circumstances under which the information was disclosed... to [any] employee to the extent that they give rise to a reasonable inference that further disclosure, without the consent of the possessor, is prohibited, and (4) the degree to which the information has been placed in the public domain or rendered 'readily ascertainable' by the third parties through patent applications or unrestricted product marketing.
>
> *USM*, at 98. [*Underlining* supplied].

The court then treats each factor in turn, citing analogous cases and sometimes distinguishing contrary cases. With respect to the first factor, for example, the presence of agreements restricting disclosure, *Eastern Marble* is analogous and the court sets up the analogy by using the 'see' citation signal:[10]

> USM required... personnel...to sign nondisclosure agreements... [S]pecificity [of the agreement] is not required to put employees on notice that their work involves access to trade secrets and confidential information. *See Eastern Marble Prods. Corp. v. Roman Marble, Inc.*, 372 Mass. 835, 840 (1977) [and other cases].
>
> *USM*, at 99.

With respect to the second factor, security precautions, the court in USM simultane-

10. See [Ashley & Rissland, 1987].

ously analogizes a favorable case and distinguishes an unfavorable case by using the introductory case citation form 'compare <case1> with <case2>.' The major issue for the court with respect to security precautions was that USM provided tours of its operations:

> The fact that USM conducted escorted tours . . . does not militate against a finding that USM denied public access to the USM machine. *Compare Plant Indus., Inc v. Coleman*, 287 F.Supp 636, 643 (C.D. Cal. 1968) (tours by women's clubs and customer's representatives do not constitute failure to maintain secrecy) . . . *with Motorola, Inc. v. Fairchild Camera & Instrument Corp*, 366 F.Supp. 1173, 1186 (D. Ariz. 1973) (security inadequate where competitors toured plant, operated 'secret' machine and 'observed its 'secret' process in a separate microscope placed there for this purpose') . . .
>
> USM, at 100–101.

Straightforward arguments are sufficiently common that a system of stylized citation signals has evolved to flag the types of comparisons (analogize, distinguish) that are made between cases [BlueBook, 1986]. See [Ashley & Rissland, 1987] for a discussion of the use of citation signals in case-based argument.

Although analogy is an area of intense AI research, we treat 'analogizing' and 'distinguishing' as primitive procedures in this article. We recognize that analogizing and distinguishing may nonetheless be implemented in a variety of ways that stem from different theories of analogy, such as structure-mapping [Gentner, 1983; Falkenhainer, Forbus & Gentner, 1989], constraint-mapping [Holyoak & Thagard, 1989], transformational and derivational analogy [Carbonell, 1983, 1986] and mapping of corresponding features [Winston, 1980, 1982].

4.1.2 Make-weight argument. A make-weight argument cites cases, facts or factors that are not necessary to reach the conclusion desired, but merely provide additional support, often of a less probative sort. See Figure 3.This style is a weak version of what has been called 'bolstering' [Rissland, 1985] and 'strengthening' [Rissland & Ashley, 1986]. Legitimate strengthening or bolstering involves adding or deleting facts or improving the value of a 'focal slot' that determines a dimension's strength [Ashley, 1990].

EXAMPLE: To use a make-weight argument, the Commissioner, after distinguishing *Weissman*, might cite *Henry C. Smith*, 40 B.T.A. 1038 (1939) aff'd per curiam, 113 F.2d 114 (2d Cir. 1940) (denying child care expenses as a ordinary and necessary

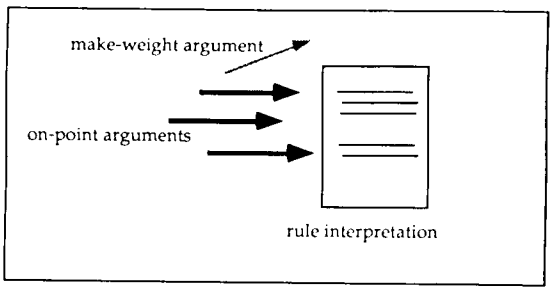

Fig. 3: A make-weight argument is a weak argument that does not squarely address the issue at hand.

business expenses) and argue: 'Furthermore, taxpayer implicitly attempts to convert the personal expenses of child care into ordinary business expenses.' This line of attack does not deal directly with the satisfaction of the home office provision, but merely attempts to darken Alice's case with the shadow of a tax shibboleth that personal expenses are not deductible.

One does not usually set out to create a make-weight argument: one relies on an opponent to characterize an argument disparagingly as merely 'make-weight'.

4.2 ARGUMENTS INVOLVING THE CURRENT FACT SITUATION AND A SET OF TWO OR MORE CASES

The following arguments involve reasoning with the current fact situation and a set of cases or a representative of a set of cases.

4.2.1 Turkey, chicken and fish or 'double negative' argument. This convoluted argument takes the following form: the case at bar is so *unlike* the cases where a rule's conditions were held *not* to have been established that the rule *should* apply to the current case. The argument may be used more generally to claim that an instance is within some category, whether or not that category is implicitly defined by a rule consequent. This slightly expanded form of the argument maintains that something is so unlike negative examples of a category – unlike things outside the category – that it should be considered as a positive example – as within the category. See Figure 4.

EXAMPLE: The name 'turkey, chicken and fish' stems from a hypothetical example in which a turkey farmer in Delaware tries to receive an entitlement reserved by a Federal Department of Agriculture regulation for chicken farmers. If the only cases denying the entitlement dealt with fish farmers, then the turkey farmer could argue that his turkey farms are so unlike fish farms that the FDA regulation should apply to his turkey business as well. Cf. *Frigaliment Importing Co. v. B.N.S. Int'l Sales Corp.*, 190 F.Supp. 116 (1980) (Friendly) (dealing with the definition of a 'chicken').

This fairly weak form of argument may be effective where the set of on-point cases is small and contains instances only of egregious attempts to crawl within the ambit of a statute that is inapplicable on its face (the fish cases in the example). In this setting, the advocate then claims that – unlike those flawed cases – his (turkey) claim is reasonable and comes within the statute. The advocate might support this line by appealing to a legislative history that revealed, for example, that when the regulation was promulgated, 'chicken farm' was merely a compromise term to clearly omit the (weak lobby) fish farmers, and should be construed to include all (strong lobby) fowl farms.

EXAMPLE: Alice may use this argument form to claim: 'The precedents where the home office deduction was denied involved situations where the home office was used almost 100% of the time for *personal* purposes. Business use was *de minimis* there. My own situation is not one of these abusive situations where the claim of a home office was merely a facade, which led Congress to tighten the requirements. In my

own case, the personal use was infrequent and minor. A deduction should not be denied in my case, which is clearly distinguishable from those egregious precedents.'

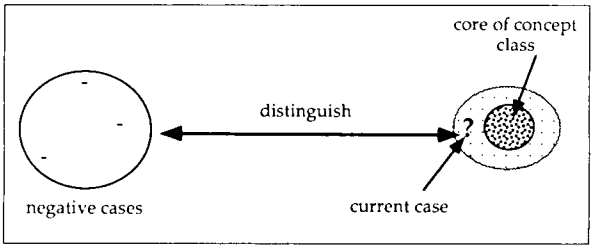

Fig. 4: The turkey, chicken and fish strategy distinguishes negative cases to argue that the current case is within the penumbra of the concept class of cases implicitly defined by the rule interpretation.

EXAMPLE: *Landreth Timber Co. v. Landreth*, 471 U.S. 681, 105 S.Ct. 2297, 85 L.Ed. 2d 692 (1985), uses a variation on this argument form. The issue in *Landreth* was whether the sale of all the stock of a company was held to be a 'security' subject to the antifraud provisions of the Federal securities laws. The definition of a security is given in the Securities Act of 1933 (15 U.S.C. §77b(1)): '*T*he term "security" means any note, stock, treasury stock, bond, debenture ... investment contract ... or, in general, any interest or instrument commonly known as a "security" ....' In *Landreth*, the Court stated that different tests might be used for (a) the prototypical case of an instrument that was called 'stock' on its face and possessed the usual characteristics of stock and (b) other enumerated items. Previous Supreme Court cases had applied tests looking to the 'economic reality' of various financial instruments in the 'security' determination. *Landreth*, however, looked to the face of the security, which was clearly that of stock. In arguing in dicta that a formal test may be used for things that are labeled 'stock' and an economic reality test may be used for other enumerated items, the Court applied a chicken, turkey and fish argument variant. The case of an instrument entitled 'Stock' (the 'turkey') is so unlike the other enumerated categories of security where the economic substance of the transaction has to be evaluated (the 'fish'), that the instrument should be considered stock (a 'chicken').

Statutes that contain enumerated lists by way of definition often present opportunities for arguing with turkey, chicken and fish.

4.2.2 'Throw the dog a bone' or 'gratuitous' argument. A gratuitous argument is similar to the foregoing double-negative argument in that it applies primarily in a setting where one is trying to broaden a rule by distinguishing a negative case. Here, however, the negative case is a hypothetical one that clearly does not satisfy the rule. A weak example (the 'bone'), whose negative classification is conceded, is retrieved or concocted and the arguer proceeds to show how the current situation is a stronger example than the artificial weak one. In practice, use of this argument may be advisable in a 'defensive' position where one perceives a case is weak.

EXAMPLE: In Alice's case, the Commissioner could argue that to allow Alice to claim a home office deduction would permit almost any single parent who does 'office' work to make a deal with the employer which would permit the employee to deduct expenses which for most other taxpayers are non-deductible personal expenses under Section 262, thereby eviscerating the intent of Congress when it enacted Section 280A. The taxpayer would then reply, 'Your honor, I agree with the Treasury that one cannot permit any taxpayer who has an office and does some work at home to claim a home office deduction. Such a result would clearly undermine Congress's intent in passing Section 280A. But we are not asking for anything resembling the judicial invalidation of Section 280A. We are not advocating making the deduction available to a worker with ideal working conditions at their office and who spends only 10% of their time in their home office to deduct personal living expenses. We are only arguing that the deduction be available where the taxpayer spends a significant amount of her time working at home because of overcrowding at her employer's office.'

In this example, Alice throws a gratuitous bone to the court in admitting that she would not be entitled to a deduction had she spent only 10% of her time working at home when an ideal office was available to her elsewhere.

4.2.3 Straw man argument. Straw man arguments are familiar to students of law. The term is used disparagingly by an opposing counsel, a judge, or a commentator to describe an argument with an allegedly obvious flaw. One does not deliberately set out to create a straw man argument. Various sorts of disingenuous or exaggerating arguments may be characterized as straw men arguments, but the term at least encompasses the creation of a hypothetical case that is analogized by one side to the current situation, but which is easily distinguished from the current problem. See Figure 5.

EXAMPLE: In the home office area, if the Commissioner were to seek disallowance on the ground that the study was not used 'exclusively' for business purposes the taxpayer might create the following straw man. 'Your honor, the statute cannot mean what the Commissioner says it means. Certainly, Congress did not intend that someone otherwise entitled to a deduction for a home office should lose the deduction merely because they write a single personal check at a desk that is otherwise used entirely for business.' To which the Commissioner might appropriately respond, 'Your honor, the taxpayer has raised a total straw man. We're not talking about the writing of a single check. We're talking about the use for house guests, for personal correspondence and for handling personal financial transactions. Taxpayer does not satisfy the statutory requirement of exclusive use.'

EXAMPLE: In *Schneckloth v. Bustamonte*, 412 U.S. 218, 93 S.Ct. 2041 (1973), a leading constitutional case on the definition of voluntary 'consent' to search, Justice Marshall, in dissent, said on the issue of the allocation of the burden of proof (at 285–286),

[Under one view, t]he question then is a simple one: must the Government show that the subject knew of his rights, or must the subject show that he lacked such knowledge?

I think that any fair allocation of the burden would require that it be placed on the prosecution. On this ques-

tion, the Court indulges in what might be called the 'straw man' method of adjudication. The Court responds to this suggestion by overinflating the burden. [The Court wrote in part 'For it would be thoroughly impractical to impose on the normal consent search the detailed requirements of an effective warning.'] And, when it is suggested that the *prosecution's* burden could be easily satisfied if the police informed the subject of his rights, the Court responds by refusing to require the *police* to make a 'detailed' inquiry. Ante, at 245. If the Court candidly faced the real question of allocating the burden of proof, neither of these maneuvers would be available to it.

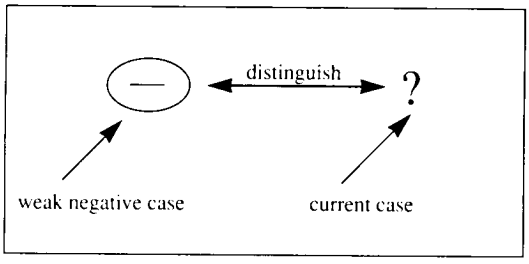

Fig. 5: A straw man argument often sets up a weak analogy that is easily distinguished by an opponent.

4.3 ARGUMENTS INVOLVING THE CURRENT FACT SITUATION AND A SEQUENCE OF CASES

The following argument forms involve not simply a set of cases, but an ordered sequence of them.

4.3.1 Slippery Slope. This form of argument may be familiar to students of the law almost to the point of banality.[11] An advocate suggests a sequence of cases, real or hypothetical, whose starting point is the current situation and whose endpoint is an untenable application of the rule or predicate in question. The difficulty of making a principled distinction among the cases in the sequence supports the argument that the interpretation of the rule or predicate in issue must be limited to the precedent and not expanded to the current fact situation. See Figure 6.

EXAMPLE: The Commissioner may pose: 'Consider your honor, a sequence of cases before you – all similar to Alice's case, but where the percentage of total work time the home office decreases: not the 80% found in *Weissman*, but the 40% in the current case, then 25%, 10%... How little use will be conceded? This court must stem the flow of home office litigation and require at least 80% of work be done in the home.'

[11] One of the authors attended a first-year law school class in which this form of argument was invoked on a daily basis.

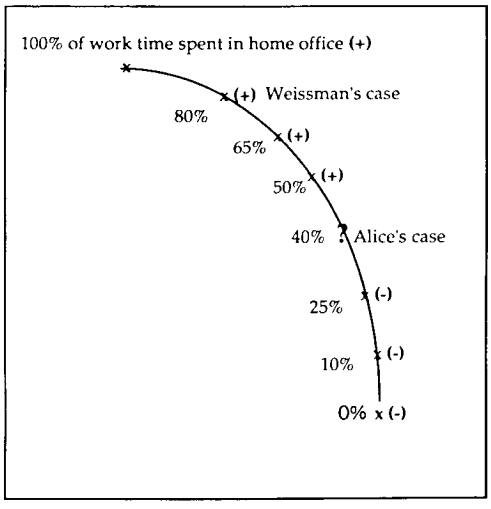

Fig. 6: A slippery slope argument suggests that the lessening of a threshold will lead to further decreases, until clearly negative cases are illegitimately included as positive examples.

EXAMPLE: In [Ely, 1975, p.1501], John Ely mimics a jurist questioning an attorney about a series of hypotheticals: 'I understand that you would protect sound-trucks. But what about a hospital zone? What about the middle of the night? Surely you wouldn't let a mayoral candidate aim a bullhorn at your window at three in the morning. Surely you have to balance, or employ a clear and present danger test, at *some* point.'

Finally, a slippery slope argument provided a light moment during oral argument before the Supreme Court in the *State of Texas v. Gregory Lee Johnson*, 491 U.S. 397 (1989), in response to the contention by counsel for Texas that a Texas statute proscribed the burning of an American flag. A member of the Court asked if burning a flag with 48 stars would be reachable by the Texas law. Texas counsel thought that it would. The Justice then asked if burning a flag with 47 stars was proscribed. When Texas counsel admitted that it would not be, the Justice observed, 'So, all you have to do is take one star out of a – out of the flag, and it's okay.'

4.3.2 Reductio[12] *Loop.* This argument form is the same as the slippery slope except that it leads to the conclusion that a case residing on what was assumed to be unimpeachable ground ('further up the slope') is questioned as well. One way to perform a *reductio* loop is to show how the supposedly unimpeachable case can be linked to the current fact situation, which in turn slides farther down the slope to the *reductio* case. The conclusion is that the allegedly unimpeachable case is as suspect as the *reductio* case and the current fact situation, since the difference among all of these cases is just a matter of degree. The

12. A 'reductio ad absurdum' argument is 'the method of disproving an argument by showing that it leads to an absurd consequence.' Black's Law Dictionary (5th Edition).

reductio argument might be used to subvert the effect of an established ruling in a precedent relied upon by opposing counsel. See Figure 7.

EXAMPLE: As an example of a *reductio* argument, the Commissioner might argue that the taxpayer's claim in this case shows the problems that inhere in decisions like *Weissman*. 'Your honor, once you allow an employee who has an office provided by her employer to claim that her principal place of business is the home then one can not draw a defensible line – a line that is both fair to all taxpayers and capable of efficient administration. Therefore, *Weissman* should not be followed in this court [or it should be reversed].'

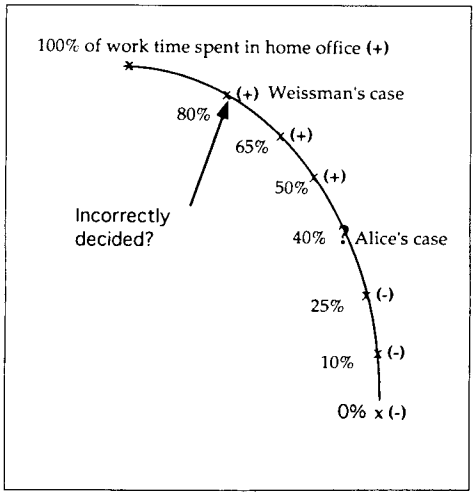

Fig. 7: The *reductio* argument is similar to a slippery slope argument, but questions a previous decision near the top of the slippery slope. *Weissman* is questioned as incorrectly decided in view of its slippery consequences for Alice's case and others.

EXAMPLE: In *Commissioner v. Flowers*, 326 U.S. 465 (1945), a taxpayer who lived in Jackson, Mississippi, but who worked for days at a time at the headquarters of his railroad employer in Mobile, Alabama, attempted to deduct meals and hotel expenses while in Mobile. In dissenting from a decision in which the deduction was denied, Justice Rutledge wrote (at 478–479),

Congress gave the deduction for traveling away from home on business. The commuter's case, rightly confined, does not fall in this class. One who lives in an adjacent suburb or city and by usual modes of commutation can work within a distance permitting the daily journey and return, with time for the day's work and a period at home, clearly can be excluded from the deduction on the basis of the section's terms equally with its obvious purpose. But that is not true if 'commuter' is to swallow up the deduction by the same sort of construction which makes 'home' mean 'business headquarters' of one's employer. If the line may be extended somewhat to cover doubtful cases, it need not be lengthened to infinity or to cover cases as far removed from the prevailing connotation of commuter as this one. Including it pushes 'commuting' too far, even for these times of rapid transit.

The thrust of J. Rutledge's reasoning is that the majority's position will support an improper *reductio* argument. Rutledge argues that the 'top of the slope' prototype of a

'commuter' is weakened when decisions are made to deny the travel deduction in cases that are far removed – far down the slope – from that prototype are justified as 'commuting expenses'.

4.4 MISCELLANEOUS ARGUMENT FORMS

4.4.1 Hedging argument. Suppose that the advocate wants to argue in his brief that an apparently unfavorable rule does not apply or does not control the decision in the current case, but if the rule *is* found to be dispositive, the desired result follows nevertheless. This form of argument can be seen in legal briefs, where an attorney will hedge his rhetorical bet. See Figure 8.

EXAMPLE: The Commissioner may argue 'Your honor, we have established in the immediately preceding argument (the *reductio* argument) that the rule of *Weissman* should not apply because it is bad law. But if *Weissman* does apply, then Alice's case is distinguishable because of her smaller percentage use and her non-exclusive use. In addition, Weissman created a home office partly due to concerns for his safety.'

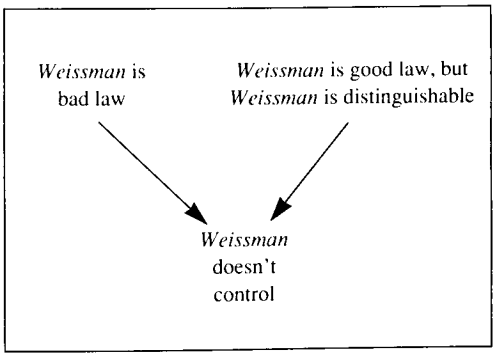

Fig. 8: A hedging argument can use a forked argument that the result in *Weissman* should not control the disposition in the current case.

4.4.2 Weighing or balancing argument. Perhaps the form of argument seen most by law students is the 'weighing' argument. Weighing arguments pre-suppose the identification of domain-important factors that are somehow 'weighed' against each other. Compare [Rissland, Valcarce & Ashley, 1984; Ashley, 1990] (arguing for a least-commitment, implicit weighting scheme). A closely related and possibly co-extensive form of argument is the 'balancing' argument, where one attempts to place policies supporting two legal ideas or ideals 'in balance'. In the typical application of this form of argument, the competing policies never balance (equally) and the balance arm dips in favor of the weightier policies.

EXAMPLE: To invoke a weighing or balancing argument, the Commissioner might argue, 'Your honor, this case requires you to balance the policy of preventing taxpayers from deducting purely personal expenses against the policy of only taxing net

income, income earned after the expenses incurred to produce that income. Where the taxpayer who is an employee has NO office provided by the employer then the balance tips in the taxpayer's favor. But where the employer does provide an office, then the balance shifts the other way. To balance these policies in this case in favor of the taxpayer would erode the tax base and create inequities between those taxpayers who can do a considerable amount of their employer's work at home and those who just work there on evenings and weekends.'

EXAMPLE: In *Bob Jones University v. U.S.*, 461 U.S. 574 (1983), dealing with whether schools that use racially discriminatory admissions standards can qualify as tax-exempt organizations, Justice Powell, concurring in part and concurring in the judgment, argued (at 610–611),

Congress, of course, may find that some organizations do not warrant tax-exempt status. In this case I agree with the Court that Congress has determined that the policy against racial discrimination in education should override the countervailing interest in permitting unorthodox private behavior.

I would like to emphasize, however, that the balancing of these substantial interests is for *Congress* to perform.

Professor Henkin described the balancing process in First Amendment cases:

The Court will hold the scales, attribute different weights to different communications (taking account of content, form, and context) and to values affected by the communication, will determine the level of tolerance imposed on government by the first amendment, and read the balance. [Henkin, 1968, p.81]

Henkin goes on to say, parenthetically, however,

(I use the metaphor of balancing but the point is the same however one describes the scope and limits of the amendment, or the judicial process in determining them.) [Henkin, 1968, p.81]

As Henkin observes, in practice this form of analysis provides a metaphor more than a rhetorical method. Legal opinions, which frequently invoke metaphorical language appealing to the 'physical mass' of competing factors, rarely – if ever – provide a scale that can gauge the 'weight' of these factors. And so the weighing argument often smacks of contrivance. Its role in actual use may be to dress up in objective clothing a decision made on other, unstated grounds. Often, the arguer presses his thumb on the scale: the words chosen to state the opposing claims are an immediate key as to which side possesses greater credence in the author's eyes. In a comment dealing with the means of deciding flag desecration cases under the first amendment, John Ely shores up this view:

The categorizers were right: where messages are proscribed because they are dangerous, balancing tests inevitably become intertwined with the ideological predispositions of those doing the balancing – or if not that, at least with the relative confidence or paranoia of the age in which they are doing it – and we must build barriers as secure as words are able to make them. [Ely, 1975, p. 1501]

Balancing or weighing arguments thus provide a familiar framework for stating legal problems, rather than suggesting an adversarial approach for their resolutions.

In this section we have given a partial taxonomy of legal argument forms as they are applied in legal practice. The next section changes pace to discuss a distinct computational implementation of some of the argument terms and methods in our CABARET

system. The reader may see connections between CABARET and the argument forms just presented in this Section. In Section 6 we make explicit these connections between the argument forms and the computational model we are about to present in Section 5.

5. Argument strategies and moves in CABARET

In this section, we describe a model of legal argument implemented in CABARET that involves three components: **argument strategies, argument moves,** and **argument primitives**. The facet of legal argument on which we concentrate is how to apply cases to argue for an interpretation of a term used in a legal rule. We define the term 'rule' quite expansively to include not only a statute and a supporting regulation, but also a blackletter rule of law induced from a group of cases, such as a rule from the Restatement of Contracts, or a 'rule of a case,' a rule-like statement capturing the holding of a single case.

CABARET's argument strategies, moves and primitives may be seen as inhabiting a hierarchy of argument techniques (Figure 9):

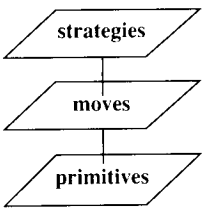

Fig. 9: The hierarchy of argument strategies, moves and primitives.

In the following sections we discuss each of the three levels of argument techniques, starting with argument strategies.

5.1 ARGUMENT STRATEGIES

How one argues about a rule depends on one's point of view. For instance, with respect to a deduction against gross income, a taxpayer would likely take a *pro* (in the sense of 'in favor of') position whereas the IRS may protect the tax revenues of the nation by adopting a *con* strategy. In our model, the pro or con point of view is the initial determinant of how one argues.

Suppose that the IRS has challenged a home office deduction on the grounds that the taxpayer did not meet the 'use on a regular basis' requirement, but the IRS admits that the other requirements of Section 280A(c)(1) have been met. The task confronting the taxpayer is to extend the rule so that it covers a situation – the taxpayer's – which is outside the strict scope of the rule as construed by the IRS. Taxpayer's task can be captured in the following rule of thumb:

If a rule's conditions are not met and the arguer wants the rule to succeed, then broaden the rule.

On the other hand, the IRS has adopted the contrary point of view ('con'): against allowing the home office deduction. The IRS would argue that the rule is to be strictly applied. To bolster that contention, it would attempt to find cases interpreting that section that confirm that the rule's conclusion – the deduction – should not follow in this situation. The parallel heuristic from the IRS's point of view is:

If a rule's conditions are not met and the arguer wants the rule to fail, then confirm the miss.

In section 5.2 we discuss how certain precedents can be used to support rule broadening and rule miss confirmation. The two points of view and the two possibilities that a rule's conditions are met or are not met by a fact situation suggest a 2 × 2 matrix of top-level **argument strategies** for the interpretation of a rule (Table I).

Table I. A matrix of argument strategies.

Point of View/ Rule Conditions	**Point of View = Pro**	**Point of View = Con**
Rule Conditions Met	Confirm the Hit	Discredit the Rule
Rule Conditions Not Met	Broaden the Rule	Confirm the Miss

Each of the individual cells represents a strategy useful in circumstances that depend on the arguer's goal and a strict interpretation of the rule:

- **broadening** is used to argue that a rule applies to a situation where strict application indicates it does not, or to argue for membership in the concept class implied by the rule consequent, irrespective of the rule;
- **discrediting** is used to argue that a rule does not apply to a situation where strict application indicates it does;
- **confirming a hit** is used to argue that a rule does apply to a situation just as strict application indicates;
- **confirming a miss** is used to argue that a rule does not apply to a situation just as strict application of the rule indicates it does not.

This presentation of argument strategies is directed at appellate argument. The model assumes that a determination has been made whether a rule's conditions have strictly been met. This determination may have been given by administrative interpretation (as by the IRS) or by a lower court's interpretation (as by the Tax Court). Alternatively, on questions of first impression, the arguer may make a hypothetical assumption regarding whether a rule's conditions have been met and 'play out' the arguments using the resulting strategies.

In this approach to statutory argument, each of the argument strategies is effected through a set of *argument moves*,[13] which are described in the next subsection.

5.2 ARGUMENT MOVES

An arguer can carry out a chosen argument strategy in several ways. While the choice of strategy looks in part to the current fact situation, the choice of move looks also to the available precedents. For each of the four argument strategies, there are four possible moves in turn. Each move depends on two factors: (1) the disposition of a precedent – whether or not the precedent held for the ultimate desired result; and (2) the status of the rule with respect to the precedent – whether or not the rule's conditions were satisfied by the facts of that precedent. We give some informal examples of possible argument moves and then provide a systematic treatment.

Consider again the example of a home office deduction claim that is challenged on the basis of taxpayer's failure to meet the 'regular use'[14] requirement of the statute. On the basis of available precedents, a taxpayer could attempt to carry out an argument for broadening the home office rule in several ways. She can argue, for example:

1. that the unmet condition actually has been satisfied: that the allegedly missing condition has been met in fact and thus the rule should apply. A useful precedent would be a case in which (a) the regular-use condition and all other rule conditions have been satisfied, (b) the home office deduction conclusion was therefore established, and (c) the taxpayer's situation is similar – especially with respect to the regularity of home office use. Or
2. that the unmet condition is not necessary: that the result of the rule follows without meeting the condition with which she is having difficulty. A good precedent would be a case in which the regular-use condition was not met but where the home office deduction was granted nonetheless. Or
3. that her case is so *unlike* the cases where the rule's conditions have *not* been established that the rule's conditions should be interpreted as being met in this situation. Useful precedents would be a group of cases that did not satisfy the regular-use condition and were not granted the deduction, but that were markedly different from cases that did satisfy the regular-use requirement, and distinguishable from the current situation.

Choices (1) and (2) are variations of arguments based on precedents with the desired disposition: the taxpayer in the retrieved cases was allowed the deduction. The arguer would draw analogies with these favorably disposed cases to justify why the desired dis-

---

13. In a recent article, [Gordon, 1991] refers briefly to 'argument moves' as analogs of three roles for cases identified by Ashley: 'cited cases', 'distinguished cases' and 'counter-examples'. We reserve the term 'argument moves' for specific, active tactics that an advocate can use to support his stance. The term was also previously used by Rissland [1985] and Ashley and Aleven [1991].

14. No well-formulated test for 'use on a regular basis' has been applied by the courts. The 'frequency and repetitive nature' of the business use appear to be the most important considerations. See [Knobbe, 1986] for a description for practitioners of the home office deduction.

position should be reached in his situation as well. In (1), the arguer works with a precedent where the rule applied and in (2), with one where it did not. Choice (3) provides one way to use a case offensively that has been decided in favor of the opposing point of view.

In (1), the arguer uses the precedent to establish an unmet condition in her current case in order to argue for the consequent of the rule. The more her facts support the unmet condition relative to the precedent's, the better her argument.[15] ('Since the other taxpayer met the requirement and I am at least as well situated as he, I also meet the requirement.')

In (2), one argues directly for the disposition by soft pedaling or denying outright the necessity of meeting the unmet rule precondition. Thus, in (2) the arguer will analogize the current case and the precedent to justify the deduction. She might go further to argue that the rule is invalid on its face, as opposed to incorrectly applied to this particular situation. The better the match between the unmet condition in the current fact situation and in the precedent, the easier the argument. ('Since the other taxpayer was allowed the deduction even though he failed the same alleged requirement, so it should be allowed in my case.')

Choice (3) might be used in light of precedents whose outcome is the opposite to that desired: the deduction was not allowed and just as in the current case, a condition was deemed unmet. The arguer would distinguish those cases: that they should not govern her case because her case is much stronger than those cases. She would need to argue by 'double negative' that the desired disposition should be reached in her case. ('Since I am so unlike those cases where the deduction was disallowed, I should be allowed the deduction.')

More difficult to make even than (3) would be an argument using a precedent with an undesired disposition despite the fact that the rule conditions were met. To use such a case, one would have to both distinguish the negative disposition and make use of its positive aspect concerning rule satisfaction. This strategy is tricky: it tries to make a silk purse out of a sow's ear. Besides, an obvious counterattack exists. The opponent could use such a precedent to argue that the standard is even stricter than the rule specifies. In fact, the opposition may well rely on this precedent in the first place and the advocate would have to respond to it. It is in this defensive role that an advocate may be forced to reply with such delicate arguments.

Typically, one's first line of argument is to analogize cases with the disposition that one desires for the current fact situation. However, since the only relevant precedents might have been decided oppositely, argument moves are needed to use them to build offensive arguments. And as suggested in the last paragraph, since one's opponent is sure to cite cases with unfavorable outcomes, it is also necessary to respond defensively to cases with unfavorable dispositions by distinguishing them.

Argument choices (1), (2) and (3) are particular instances of a scheme of argument moves developed for the CABARET system. We now place these three choices in a

---

15. The stereotypical form of this argument is supported by cases that hold on the predicate in question and whose supporting criterial facts can be matched, analogized or favorably compared. See [Branting, 1991].

larger plan of argument moves and tie all such moves into the broad argument strategies presented in Section 5.1.

First consider the broadening argument strategy. In general there are two possibilities for the disposition of a relevant precedent and two possibilities for the success of strict application of the rule to it, yielding four possible types of precedents to be dealt with. These four possibilities provide the setting for the argument moves for broadening, which are listed in Table II.

Table II: Argument Moves for Broadening. (*Cell numbers in braces are included for reference only.*)

Disposition of Precedent/ Rule Consequent is established in Precedent	**Precedent has Desired Disposition**	**Precedent doesn't have Desired Disposition**
Yes	Analogize Case {1}	Distinguish Case {3} Disposition and Analogize Rule Consequent
No	Analogize Case {2} Disposition (and Distinguish Rule Consequent)	Distinguish Case {4}

A word as to terminology in Table II: 'Analogize Case Disposition' means to draw analogies between the precedent and the current situation of any sort to argue that the outcome of the two cases should be the same. 'Analogize Rule Consequent' means to draw analogies between two cases using features related to the rule's preconditions, in order to argue that the rule should hold – or fail to hold – in both cases. 'Analogize Case' means to do both: 'Analogize Case Disposition' and 'Rule Consequent' Analogize. 'Distinguish Case', 'Distinguish Case Disposition', and 'Distinguish Rule Consequent' are parallel to their three 'analogize' counterparts, except that distinctions between cases are exploited, not similarities.

Distinguishing moves are applied in the two situations for broadening where the retrieved cases have the wrong disposition. These moves are contained in cells {3} and {4} in Table II. There, the precedent may have a different factual complexion or a different status with respect to the strict application of the rule. Hence, an advocate may need to distinguish a case decided 'for the other side' with respect to the precedent's disposition, the result of strict rule interpretation in that situation, or both.

The broadening table also summarizes how to carry out the argument strategy to confirm a hit, that is, to confirm that a rule's preconditions are satisfied. That these strategies share moves makes sense in that the argument moves do not depend on the application of the rule to the current problem situation, but depend only on the retrieved precedents. The difference between being in a broadening mode and in a confirming mode depends only on the application of the rule to the current problem.

In confirming a hit, one would prefer to analogize a similar case with the right disposition where the rule's conditions were satisfied. However, depending on the cases that are available to an arguer, he may have to rely on cases that have unfavorable dispositions or reflect unfavorable strict interpretations under the rule. Various other argument moves to confirm a hit are available as with broadening. With a merely confirmatory strategy, however, it is less incumbent upon the arguer to use precedents that are in any way unfavorable. There is no need to stretch arguments very far when the received rule interpretation favors one's own side.

The remaining two argument strategies, discredit and confirm a miss, also can be effected through similar argument moves. These two strategies also share a table of argument moves, which is presented in Table III.

Table III: Argument Moves for Discredit a Rule and Confirm a Miss. (*Cell numbers in braces are included for reference only.*)

Disposition of Precedent/ Rule Consequent established in Precedent	**Precedent has Desired Disposition**	**Precedent doesn't have Desired Disposition**
Yes	Analogize Case {1} Disposition and Distinguish Rule Consequent	Distinguish Case {3}
No	Analogize Case {2}	Distinguish Case {4} Disposition and Analogize Rule Consequent

The tables for all the argument moves use the same generic argument tasks of analogizing and distinguishing, which for convenience we therefore call 'primitive' argument tasks. These primitives are considered briefly in the next section.

5.3 ARGUMENT PRIMITIVES

So far in our discussion of argument moves we have made a simplifying assumption: that there is no difference between similarly disposed cases aside from their status with respect to a rule. However, in order to carry out the details of analogizing and distinguishing cases, an arguer must give close consideration to the facts of the available precedents, the current fact situation, and the degree of match between an available case and the current fact situation. These are central concerns of case-based argument. The selection of a strategy, and then of a move, focuses attention on how the rules constrain the use of case-based reasoning to argue case similarity and difference. We now address some of the case-based concerns of how distinguishing and analogizing are implemented.

In performing moves that require analogies, an arguer would usually rely on 'best' cases in the sense defined by Ashley [1990]. Best cases are most on-point cases for which

there is no 'trumping counterexample,' which in turn is defined as a 'case with the opposite outcome that contains all of the cited case's similarities and then some'[16] [Ashley, 1989]. Relying on cases other than best cases can be rhetorically dangerous. Features present in a precedent but not in the current fact situation permit a counterargument: that the result was obtained in the precedent due to the presence of those other mitigating or contributing factors [Ashley, 1987]. The use of best cases to implement a strategy, such as broadening, forestalls a contrary argument about credit assignment. For instance in our taxpayer example, careful selection of the precedents on which the taxpayer relies can hinder a counterargument in which the taxpayer's adversary could 'distinguish away' taxpayer's case citations. Given a choice of ways to pursue an argument strategy, an

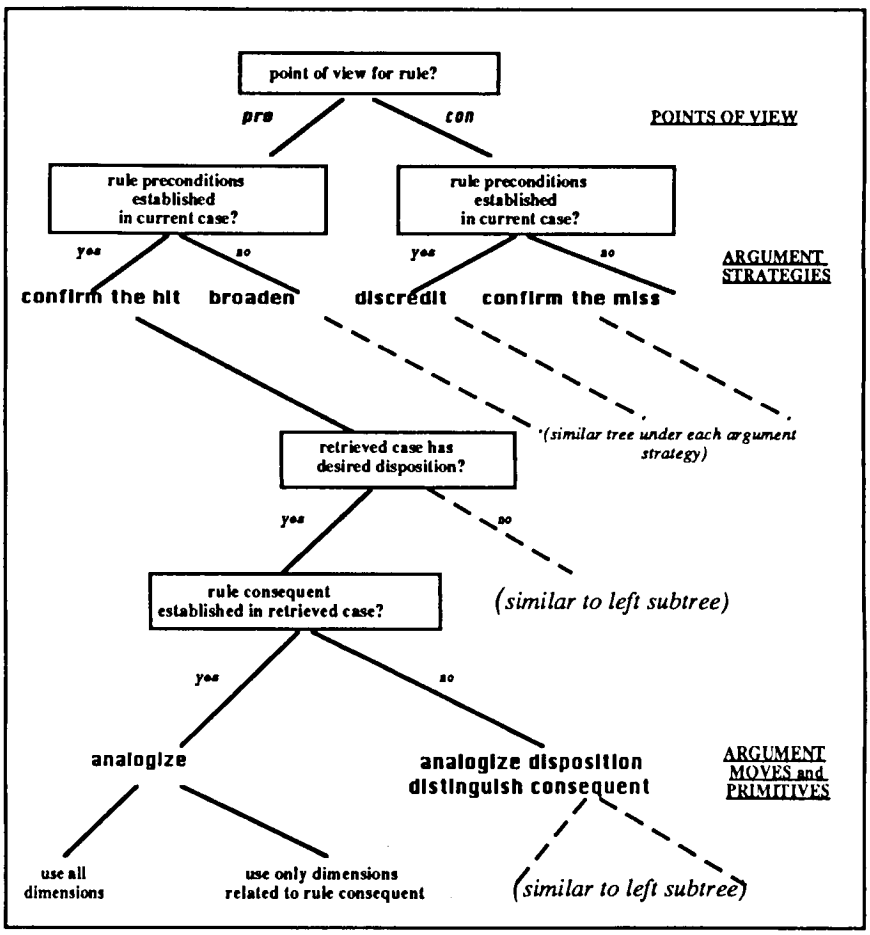

Fig. 10. Decision tree fragment showing points of view, argument strategies, argument moves and primitives.

16. A best case must also share some factor with the problem that affimatively favors one's own side. [Ashley, 1989].

arguer chooses a specific tactical move based on the cases actually available. But given a variance in the closeness of match of these cases with the current case, some moves might be better than others, but no one case or move may be clearly best.

When no best cases or even favorably disposed cases are available, the arguer might usefully employ hypothetical cases [Rissland & Ashley, 1986; Ashley, 1990], but in strict precedent-based domains hypotheticals lack the pedigree of genuine precedents.

One way to implement this theory of strategies, moves and primitives is to form arguments in a top-down manner. In this way, the entire sequence of choices of strategies, moves, and generic argument tasks can be shown in a decision tree whose branch points involve point of view, status of the current fact situation with respect to the rule, disposition of available precedents, status of precedents with respect to the rule, and degree of match of precedents with the current fact situation. Figure 10 shows a portion of the decision tree, in particular that fragment relevant to the confirm-a-hit strategy. Although the figure includes only two leaves below the choice of argument move, there may be many ways to effect the moves. The means chosen will involve consideration of the details of the cases. In this top-down approach, the degree of case match is only considered after the choice of move.

In the next subsection we give an example of an argument generated by CABARET by applying its implemented strategies, moves and primitives.

5.4 ARGUMENT GENERATION IN CABARET

The current implementation of CABARET generates argument skeletons that report the processing of the system to suggest lines of possible attack. Figure 11 gives an example. A complete argument could be based on such a skeleton. A correspondence can be seen between portions of the argument on one hand and the argument strategies, moves and primitives and the control heuristics that invoke them, on the other. For example, the first sub-argument for the *Weissman* case in Figure 11 corresponds to the rule-based-near-miss control rule, the broadening argument strategy, and a straightforward 'analogize case strategy' (as in cell {1} of Table I). The other pieces of the argument in Figure 11 are direct reflections of other tasks performed by CABARET.

The argument skeleton can be seen as a direct reporting of its controlled processing. This transparency (to use a hackneyed phrase) may be a feature or a bug. That arguments may be outlined simply by reporting the processing of the system emphasizes the primary role of argument techniques in controlling what tasks CABARET creates and performs. On the other hand, one of the points of this article is the simple one that creating arguments in a task domain that is guided by both rules and precedents presents a complex control problem. As Golding and Rosenbloom [1991] point out, CABARET presents all the evidence that it acquires to support a user's viewpoint, but it is important to note that part of that evidence is gathered in response to the failure of one reasoner or the other to achieve its goals alone. CABARET would profit from an ability to plan arguments and to prune arguments that are weak. The first of these potential extensions to CABARET is the subject of the next section.

;;; ARGUMENT for the WEISSMAN case with respect to the predicate
;;; PRINCIPAL-PLACE-OF-BUSINESS:
;;; ==========

While the rule PRINCIPAL-PLACE-OF-BUSINESS-RULE did not fire and the consequent of the rule, PRINCIPAL-PLACE-OF-BUSINESS, WAS NOT established, we may appeal to the following arguments to support a claim for the predicate PRINCIPAL-PLACE-OF-BUSINESS:

(1.) Note that only one conjunct of that rule,
((WEISSMAN PRIMARY-RESPONSIBILITY-IN-HOME-OFFICE T)), was missing.
For cases where that domain rule did fire and the result of the case
was our own, consider the following cases as analogies:
ADAMS, DRUCKER, FRANKEL, JUNIORXCHAMBER, MEIERS, SCOTT, ...

. . .

To analogize DRUCKER and WEISSMAN, consider the following factors possessed by them in common:
there was evidence as to the frequency of usage of the home office by the taxpayer,
the home office was necessary to perform the taxpayer's duties. . . .

(2.) Looking at case-based analysis, . . . dimensional analysis on the WEISSMAN case yields for the predicate PRINCIPAL-PLACE-OF-BUSINESS:
The APPLICABLE factors are: income was derived from activities in the home office; there was evidence as to the relative use of the home office and other work places: . . .
The NEAR MISS factors are: the home office was the location where the primary responsibilities were discharged.

The UNSATISFIED factors are: NONE.

For a pure COMMON LAW argument, the best cases to cite with respect to the predicate PRINCIPAL-PLACE-OF-BUSINESS are: BELLS MEIERS WEISSMAN_EMR
To analogize BELLS and WEISSMAN, consider the following factors in common: . . . .

(3.) The best cases for the OPPOSING side with respect to the predicate PRINCIPAL-PLACE-OF-BUSINESS are: BAIE CRISTO HONAN LOPKOFF POMARANTZ.

To distinguish BAIE from WEISSMAN, consider the following factors that were present in WEISSMAN but not in BAIE:
there was evidence as to the frequency of usage of the home office by the taxpayer;
there was evidence as to the relative use of the home office and other work places;
the home office was physically separated from the living area . . .

On the other hand, also consider the following factors that were present in BAIE but not in WEISSMAN: . . .

Fig. 11: Excerpts from CABARET's argument for the *Principal Place of Business predicate*[17].

[17] This Figure also appeared in [Rissland & Skalak, 1991], which discusses CABARET's implementation at a more detailed level than this article.

6. Future work: connect CABARET with the argument forms

6.1 USING ARGUMENT FORMS TO CONTROL PROCESSING

Argument moves and argument strategies are closely tied to the stereotypical argument forms. We revisit some of the argument forms described in Section 4 and briefly explain how the model of argument strategies and moves presented above can account for many of them. However, the stereotypical argument forms presented in Section 4 have not been implemented in CABARET, with the exception of straightforward arguments. We do envision several possible ways they may be used to direct the computation of arguments in the CABARET framework.

First, they may be used as control plans provided by the user in a top-down fashion. For instance, a user could direct the system to provide a slippery slope argument with respect to the exclusive use predicate of Section 280A(c)(1). The program would then create tasks to fill in a slippery slope argument frame. Such a template (Figure 12) would include slots for the predicate at issue, the dimension along which the predicate will be monotonically varied, the features that must be kept approximately constant from case to case, the case at the top of the slope (whose disposition is acknowledged), a list of cases with the appropriate disposition that vary along the dimension specified, and a foot-of-the-slope case whose decision is arguably untenable:

```
(define-argument slippery-slope-argument
    rule:
    rule-predicate:
    varied-dimension:
    direction-to-vary-dimension:
    unvaried-dimensions:
    case-at-top-of-slope:
    case-at-foot-of-slope:
    case-list:
    ...)
```

Fig. 12: A hypothetical definition of a template for the slippery-slope argument.

We envision one such template for each argument form. But each template may have different instantiations, of course. Using a slippery-slope argument for one dimension of the exclusive-use predicate does not preclude using slippery-slope on another predicate. (The repetition may be tedious as a matter of rhetorical presentation, however.)

Argument forms can be used as control plans to coordinate processing. Using control plans may serve to coordinate knowledge sources explicitly, flexibly and transparently (see [Cohen, 1987]). In particular, using argument forms as top-down plans would ameliorate at least two deficiencies in CABARET's current implementation. (1) There is minimal coordination of knowledge sources in CABARET. The system proposes tasks

that are useful in a given context, but has no knowledge of how to coordinate a sequence of tasks. (2) The selection of which task to execute next is done through an *ad hoc* numerical assessment of priority, rather than in a way that does not attempt to capture a complex control context in a single number.

The second application for the argument forms is to use them in a bottom-up fashion. During CABARET's normal processing, cases and relationships are discovered that can fill slots of the various argument templates. In this way, the system could opportunistically pursue arguments not envisioned at the outset of the run.

For example, suppose CABARET has already created a satisfactory argument (by some standard) with respect to the home office statutory predicate 'principal place of business' based on cases A through F. (See Figure 13.) If CABARET unearths a favorable case G that is the only case that shares a particular dimension with the current problem case (and that dimension is a factor with respect to principal place of business), then G may provide the foundation for a make-weight argument. Since G shares only a single dimension with the current case, it does not provide a very strong analogy. However, if it is the only case that shares that one dimension with the problem case, it may nonetheless be cited as a favorable case without fear of an opponent's responding with a 'trumping' case that is more on-point but has the opposite outcome [Ashley, 1990].

Such a case would be a leaf node that appears in level 1 of a 'claim lattice', the partial ordering of cases by similarity originally developed in HYPO [Ashley, 1990]. See Fig. 13.

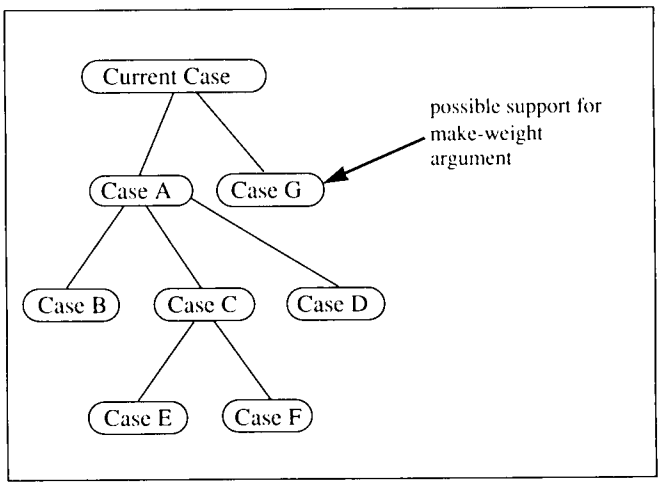

Fig. 13: The skeleton of a claim lattice for a hypothetical Current Case. If Case G shares a single dimension with Current Case, it may be a candidate for the basis for an additional, make-weight argument.

Consider another example of how argument forms can be used to control a system's processing from the bottom up. Suppose a number of cases (A-F) are most on-point, but they share few dimensions in common with the current fact situation (and each shares a different small subset of dimensions with it). Suppose further that cases A to F have been

decided in favor of the opposing viewpoint. A claim lattice based on this configuration of cases would be shallow, with a large branching factor at the first level (Figure 14). In this situation the control template representing the argument form 'turkey, chicken and fish' may be opportunistically invoked. For at this point the system would have the means to argue that the Current Case (the 'turkey') is sufficiently dissimilar from the oppositely decided cases A to F (the 'fish') – because the Current Case shares only a couple of dimensions in common with each of them – that the Current Case should be decided consistently with the desired viewpoint (i.e., should be considered a 'chicken').

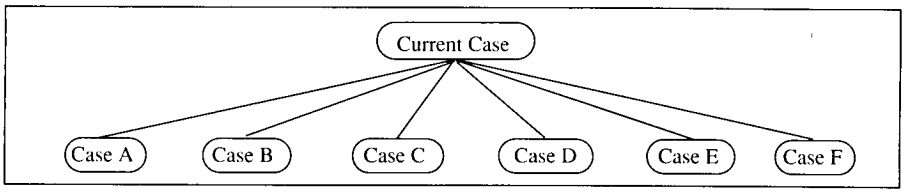

Fig. 14: Another possible claim lattice for a hypothetical Current Case. Case A to F are most on-point cases but may fail nonetheless to have many dimensions in common with Current Case.

We have argued in this article that the demands of argument require the use of (something such as) argument forms to control top-down and bottom-up processing. The argument forms could coordinate both specification-driven aspects of argument in which one attempts to find cases fitting specified criteria, with the precedent-driven aspects of argument, in which cases that have been unearthed by whatever means can be arranged to form a coherent argument.

6.2 HOW THE CURRENT FRAMEWORK OF STRATEGIES AND MOVES SUPPORTS ARGUMENT FORMS

Regardless of how argument forms might be implemented to direct processing, the thrust of this section is that the argument strategies and moves as described in Section 5 (a) provide examples of certain forms of arguments (e.g., straightforward, double-negative, hedging, make-weight), and (b) provide the means to create more complicated forms of argument (e.g., slippery slope, *reductio*). An attempt to place an argument forms plane in Figure 9 would require links both to the strategies and moves planes in that figure. We now discuss some of these links between the stereotypical forms of legal argument found in practice and the three-tiered model of argument previously developed for CABARET.

Straightforward argument. Analogizing favorable cases and distinguishing unfavorable ones are the basis of many straightforward arguments. These tactics are captured in the analogize case and distinguish case moves on the main diagonal of the Broaden/Confirm a Hit Table (Table II). Similarly, analogize case and distinguish case – counterdiagonal argument moves of the Discredit/Confirm a Miss Table (Table III) – also represent straightforward methods to argue that rule should not have fired. In these moves, cases

that have the appropriate outcome for the kind of move that's being made (favorable for analogy, opposite for distinguishing) are compared or contrasted. No additional qualification is required to account for unfavorable aspects of such cases, as in some of the hedging arguments below.

Of course, the term 'straightforward' argument is somewhat misleading. Analogizing and distinguishing are straightforward only within the context of this argument taxonomy, where they are distinguished from more complex forms of argument involving several cases or favorable pieces of unfavorable cases. Analogy and distinguishing are complex moves in themselves. They depend on a complex assessment of the features that are relevant to the comparison and of how relative similarity of cases is determined once the appropriate features have been identified [Ashley, 1987].

Turkey, chicken and fish argument. Recall that this form of argument uses negative cases to establish a positive result. For instance, one might use a case where a statute was held not satisfied to establish a positive result – that the statute's requirements are satisfied in the current case. One way to accomplish this rhetorical feat is to establish striking differences between the negative cases and the current fact situation, and argue that the point of the statute is to exclude the 'obviously egregious' negative cases. Closely similar cases – such as one's own – are within the penumbra of cases to which the statute applies.

The move in cell {4} of the Broaden/Confirm a Hit Table (Table II), is designed to use this strategy to expand the scope of a rule. One argues that the cases where the rule did not fire are so easily distinguished from the present case that the rule should apply to the present case.

Note that this use of the distinguishing move in cell {4} is an offensive[18] use of the move – a move to establish a particular interpretation. This argument form thus provides an unusual role for distinguishing. The usual role for distinguishing is a defensive response to the precedents cited by an opposing side. HYPO, for example, applies distinguishing in this role only, as a response to a case cited in a previous ply of a 3-ply point-counterpoint-rebuttal argument [Ashley, 1990].

Make-weight argument. A simple make-weight argument may be required if for example (a) the application of a controlling rule is clearly in one's favor or (b) if one has advanced other strong arguments. In the first situation, an advocate will have adopted either of the main diagonal strategies – confirm the hit or confirm the miss – of the Argument Strategies Table (Table I). In that setting, the weight of interpretation is presumably on the arguer's side since the initial interpretation of the statute has favored his case. As for (b), the example given in Figure 13 of using a case alone on a branch provides an example of a situation that CABARET can identify to yield a make-weight argument.

---

18. 'Offensive' is used as an antonym to 'defensive,' but we grant that other constructions of 'offensive' may be appropriate in this context.

By contrast to the confirming strategies, broadening and discrediting require arguing against an initial interpretation. There, the onus is upon the advocate who attempts to change the previously accorded interpretation. Lightweight make-weight arguments are usually inappropriate in these situations, where one is trying to overturn a previously established interpretation. However, an opponent may disparagingly refer to an opponent's argument as make-weight.

Straw man argument. One form of straw man argument uses a weak, irrelevant, or inappropriate example allegedly to bolster an opponent's position. Since we assume that an advocate does not deliberately set out to create a straw man argument (except possibly on rare occasions where a flamboyant straw man is created for its theatrical effect on a jury), we shall not examine how straw man arguments are reflected in the argument taxonomy. It should be clear, however, that straw man arguments may be the target of a distinguishing move in cell {3} of the Discredit Table, Table III. To discredit or limit the application of a rule, one may distinguish such a weak, hypothetical case to which the rule allegedly would apply. As we have noted, distinguishing may be applied to weak or to strong examples and so is not limited to defeating straw man arguments.

Hedging argument. As outlined in Section 4, a hedging argument argues that rule or case A controls a situation and yields a favorable outcome, but even if rule or case B controls, the outcome is still favorable. Presumably one's opponent would argue that rule or case B controls and the result is favorable to himself. A hedging argument offers two paths, both of which lead to a favorable result.

Two such paths are reflected in some of the more complex moves in the Broadening Table. For example, when an advocate is presented with a precedent whose overall result is favorable, but whose holding on a rule in issue is unfavorable, Cell {2} suggests using a combination of analogizing and distinguishing. The precedent is effectively split into two aspects. As to the features of the case that lead to a failure of the rule to fire, the precedent should be distinguished, if one is trying to argue for the rule's firing. However, as to the ultimate disposition, the move requires analogizing the aspects that lead to that favorable outcome. These two moves together reflect an attempt to hedge one's argument. The distinguishing part of the move argues that the precedent does not control the interpretation of the predicate in issue. But even if that precedent is deemed by the court to be controlling, the second half of the move argues by analogy for the favorable overall disposition of the current case.

Weighing argument. Rather than residing in any of the argument tables presented above, weighing arguments inhere in the model of case-based reasoning assumed by CABARET, which was originally implemented in HYPO [Ashley, 1990]. HYPO used important domain factors called 'dimensions' to index, retrieve and compare cases. Numerical weights were not used by HYPO, and the importance of dimensions depended

on the argument context, in particular the cases that are retrieved from the case base for support. The same is true of the weighing arguments that appear in many legal opinions. A court may list the various competing factors that are taken into account in deciding some issue, but numbers never enter the picture. The HYPO model does specify that dimensions (and values along dimensions) are considered to be favorable to one party or the other. So when cases are retrieved and sorted according to the dimensions present that are relevant to the current fact situation, the 'weight' of a set of dimensions depends on the relative similarity of the cases retrieved to the problem situation and the outcomes of those cases. Favorable and unfavorable dimensions are implicitly weighed, just as legal opinions may purport to weigh competing factors without the intervention of any quantitative aspects. See [Ashley & Rissland, 1988b] for an in-depth discussion of the implications of weighting.

7. Summary

Creating arguments is hard. It is hard for humans and it is hard for computer programs (Section 1). It is difficult for both of them partly due to this task's control problems: how does one decide what to do and when to do it? In particular, control decisions must resolve competing 'specification-driven' and 'case-driven' influences. To argue with cases effectively, one has to arrive at an idea of the overall form of the argument one would like to make, the legal categories that must be addressed and the kinds of precedents that would be useful. In those determinations one is constrained all the while by the precedents (and the rules) that can be unearthed (Section 2).

We have tried to elucidate some of the difficulties of creating arguments by presenting two sets of related ideas. The first set of ideas stems from the observation that legal argument often comes in stereotypical forms and that these forms can be placed in a taxonomy (Section 4). The other set of ideas arises out of the computational model implemented in CABARET of how to use cases to interpret terms in legal rules (Sections 3 and 5). The model distills aspects of statutory argument into three parts: assuming argument strategies, making argument moves, and implementing primitive argument techniques. The connection between the two sets of ideas is that this model appears to contain the vocabulary and methods necessary to implement actual legal argument forms (Section 6). Finally, we envision that these implemented forms could be applied as control devices to combine the top-down and bottom-up facets of case-based argument formation.

Acknowledgements

In a review of a previous draft of this article, Donald Berman provided many of the Alice examples of argument styles in the home office deduction area. We are grateful to him for expanding beyond recognition his role as reviewer. The rule-based reasoner and many of the other software tools used by CABARET were written by Dan Suthers. Elizabeth Gene and Tony Reish also programmed elements of CABARET. We thank Scott Anderson, Timur Friedman and an anonymous reviewer for their comments on a draft of this article.

References

Alvarado, S.J. 1990. *Understanding Editorial Text: A Computer Model of Argument Comprehension*. Boston, MA: Kluwer Academic Publishers.

Alvarado, S.J. 1991. Interrelationships Between Reasoning and Planning in One-Sided Arguments. *Working Notes, AAAI Spring Symposium Series: Argument and Belief*. Palo Alto, CA.

Ashley, K.D. 1987. Distinguishing – A Reasoner's Wedge. *Proceedings, Ninth Annual Cognitive Science Society Conference*. Seattle, WA.

Ashley, K.D. 1988. *Modelling Legal Argument: Reasoning with Case and Hypotheticals*. Ph.D. Thesis, Department of Computer and Information Science, University of Massachusetts, Amherst, MA.

Ashley, K.D. 1989. Toward a Computational Theory of Arguing with Precedents: Accommodating Multiple Interpretations of Cases. *The Second International Conference on Artificial Intelligence and Law*, 93–102. Vancouver, BC. Association for Computing Machinery.

Ashley, K.D. 1990 . *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. Cambridge, MA: M.I.T. Press.

Ashley, K.D. 1991 . Toward an Intelligent Case-Based Tutorial Program for Teaching Students to Argue with Cases. *Working Notes, AAAI Spring Symposium Series: Argument and Belief*. Palo Alto, CA.

Ashley, K.D. & Aleven, V. 1991. A Computational Approach to Explaining Case-Based Concepts of Relevance in a Tutorial Context. *Proceedings, Case-Based Reasoning Workshop 1991*, 257–268. Washington, DC. Morgan Kaufmann.

Ashley, K.D. & Rissland, E.L. 1986 . Towards Modelling Legal Argument. In A.A. Martino & F.S. Natali (Eds.), *Automated Analysis of Legal Texts, Logic, Informatics, Law* (pp. 19–30). Elsevier (North-Holland).

Ashley, K.D. & Rissland, E.L. 1987 . But, See, Accord: Generating Blue Book Citations in HYPO. *Proceedings, First International Conference on Artificial Intelligence and Law*, 67–74. Northeastern University, Boston, MA. Association for Computing Machinery.

Ashley, K.D. & Rissland, E.L. 1988a . A Case-Based Approach to Modelling legal Expertise. *IEEE Expert*, 3(3), 70–77.

Ashley, K.D. & Rissland, E.L. 1988 . Waiting on Weighting: A Symbolic Least Commitment Approach. *Proceedings, Seventh National Conference on Artificial Intelligence*, 239–244. Minneapolis, MN. American Association for Artificial Intelligence.

August, S.E. & McNamee, L.P. 1991 . ARIEL: A Model of Analogy Understanding in Arguments. *Working Notes, AAAI Spring Symposium Series: Argument and Belief*. Palo Alto, CA.

Bellairs, K. 1989 . *Contextual Relevance in Analogical Reasoning: A Model of Legal Argument*. Ph.D. Thesis, University of Minnesota. St. Paul, MN.

Berman, D. & Hafner, C. 1991 . Incorporating Procedural Context into a Model of Case-Based Legal Reasoning. *Proceedings of the Third International Conference on Artificial Intelligence and Law*, 12–20. Oxford, England. Association for Computing Machinery.

Birnbaum, L. 1982 . Argument Molecules: A Functional Representation of Argument Structure. *AAAI-82, Proceedings of the National Conference on Artificial Intelligence*. Pittsburgh, PA. American Association for Artificial Intelligence.

Birnbaum, L. 1985 . A Short Note on Opportunistic Planning and Memory in Arguments. *Proceedings, Ninth International Joint Conference on Artificial Intelligence*, 281–283. Los Angeles, CA. American Association for Artificial Intelligence.

Birnbaum, L., Flowers, M. & McGuire, R. 1980 . Towards an AI Model of Argumentation. *AAAI-80, Proceedings of the National Conference on Artificial Intelligence*. Palo Alto, CA. American Association for Artificial Intelligence.

BlueBook. 1986 . *A Uniform System of Citation* (14th ed.). Cambridge, MA: Harvard Law Review Association.

Branting, L.K. 1989 . Integrating Generalizations with Exemplar-Based Reasoning. *The Second DARPA Case-Based Reasoning Workshop*. Pensacola Beach, FL. Morgan Kaufmann.

Branting, L.K. 1990 . *Integrating Rules and Precedents for Classification and Explanation: Automating Legal Analysis*. Ph.D. Thesis, University of Texas, Austin, TX. Available as Artificial Intelligence Laboratory Technical Report AI90-46.

Branting, L.K. 1991 . Building Explanations from Rules and Structured Cases. *International Journal of Man-Machine Studies*, vol. 34(6), 797-838.

Carbonell, J.G. 1983 . Learning by Analogy: Formulating and Generalizing Plans from Past Experience. In

R.S. Michalski, J.G. Carbonell & T.M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach* (pp. 137–162). Los Altos, CA: Morgan Kaufmann.

Carbonell, J.G. 1986. Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition. In R.S. Michalski, J. G. Carbonell & T. M. Mitchell (Eds.), *Machine Learning: An Artificial Intelligence Approach* (pp. 371–392). Los Altos, CA: Morgan Kaufmann.

Clark, P. 1988. Representing Arguments as Background Knowledge for the Justification of Case-Based Inferences. *Proceedings, Case-Based Reasoning Workshop, AAAI-88*. Minneapolis-St. Paul, MN.

Cohen, P.R. 1987. The Control of Reasoning under Uncertainty. *The Knowledge Engineering Review*, 2, 5–25.

Ely, J.H. 1975 . Comment, Flag Desecration: A Case Study in the Roles of Categorizational Balancing in First Amendment Analysis. *Harvard Law Review*, 88, 7(May), 1482.

Falkenhainer, B., Forbus, K. & Gentner, D. 1989. The Structure-Mapping Engine: Algorithm and Examples. *Artificial Intelligence*, 41, 1 (November), 1–63.

Flowers, M., McGuire, R. & Birnbaum, L. 1982 . Adversary Arguments and the Logic of Personal Attacks. In W. Lehnert & M. Ringle (Eds.), *Strategies for Natural Language Processing* (pp. 275-294). Hillsdale, NJ: Lawrence Erlbaum.

Fuller, L.L. 1958 . Positivism and Fidelity to Law: A Reply to Professor Hart. *Harvard Law Review*, 71, 630–672.

Gardner, A. vdl. 1987 . *An Artificial Intelligence Approach to Legal Reasoning*. Cambridge, MA: M.I.T. Press.

Gentner, D. 1983. Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7, 155–170.

Gibbons, H. 1991 . *The Death of Jeffrey Stapleton: An Exploration of Legal Reasoning*. (In Press).

Golding, A. R. & Rosenbloom, P. S. 1991 . Improving Rule-Based Systems through Case-Based Reasoning. *Proceedings, Ninth National Conference on Artificial Intelligence*, Anaheim, CA. American Association for Artificial Intelligence.

Goldman, S.R., Dyer, M.G. & Flowers, M. 1987. Precedent-based Legal Reasoning and Knowledge Acquisition in Contract Law. *Proceedings of the First International Conference on Artificial Intelligence and Law*. Boston, MA. Association for Computing Machinery.

Gordon, T.F. 1991 . An Abductive Theory of Legal Issues. *International Journal of Man-Machine Studies*, 35, 95–118.

Hart, H.L.A. 1958 . Positivism and the Separation of Law and Morals. *Harvard Law Review*, 71, 593-629.

Henkin, L. 1968. The Supreme Court, 1967 Term – Forward: On Drawing Lines. *Harvard Law Review*, 82, 63.

Holyoak, K.J. & Thagard, P. 1989. Analogical Mapping by Constraint Satisfaction. *Cognitive Science*, 13 (July–September, 1989), 295–355.

Kennedy, D. 1982. The Stages of the Decline of the Public/Private Distinction. Draft, Prepared as a comment for a seminar on 'The Public/Private Distinction' sponsored by the University of Pennsylvania Law Review on January 23, 1982.

Kennedy, D. 1989 . *A Semiotics of Legal Argument*. Draft manuscript.

Knobbe, K.G. 1986 . *Hobby and Home Office Deductions – Sections 183 and 280A*. Edited by L.L. Silverstein, Tax Management Portfolio, vol. 241–4th. Washington, DC: Tax Management, Inc.

Lakatos, I. 1976 . *Proofs and Refutations*. London: Cambridge University Press.

Levi, E.H. 1949. *An Introduction to Legal Reasoning*. Chicago, IL: University of Chicago Press.

Llewellyn, K.N. 1989 . *The Case Law System in America*. Chicago, IL: University of Chicago Press.

Marshall, C.C. 1989. Representing the Structure of Legal Argument. *Proceedings, Second International Conference on Artificial Intelligence and Law*. Vancouver, BC. Association for Computing Machinery.

McCarty, L.T. & Sridharan, N.S. 1982. *A Computational Theory of Legal Argument* (LRP-TR-13). Laboratory for Computer Science Research, Rutgers University.

McGuire, R., Birnbaum, L. & Flowers, M. 1981. Opportunistic Processing in Arguments. *Proceedings, Seventh International Joint Conference on Artificial Intelligence*, 58–60, Vancouver, BC. International Joint Conferences on Artificial Intelligence.

Morris, C. 1937 . *How Lawyers Think* (Swallow Paperbooks ed.). Cambridge, MA: Harvard.

Perelman, C. & Olbrechts-Tyteca, L. 1969. *The New Rhetoric: A Treatise on Argumentation*. Notre Dame, Indiana: University of Notre Dame Press.

Polya, G. 1973. *How to Solve It* (Second Ed.). Princeton, NJ: Princeton University Press.

Rissland, E.L. 1990. Artificial Intelligence and Law: Stepping Stones to a Model of Legal Reasoning. *Yale Law Journal*, 99(8), 1957–1982.

Rissland, E.L. 1985. Argument Moves and Hypotheticals. In C. Walter (Ed.), *Computing Power and Legal Reasoning*, St. Paul, MN: West Publishing Co.

Rissland, E.L. & Ashley, K.D. 1986. Hypotheticals as Henristic Device. Proceedings, Fifth National Conference on Artificial Intelligence, 289–297. Philadelphia, PA. American Association for Artificial Intelligence.

Rissland, E.L & Skalak, D.B. 1989. Combining Case-Based and Rule-Based Reasoning: A Heuristic Approach. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 524–530. Detroit, MI. International Joint Conferences on Artificial Intelligence.

Rissland, E.L. & Skalak, D. B. 1991. CABARET: Rule Interpretation in a Hybrid Architecture. *International Journal of Man-Machine Studies* 34(6):839–887.

Rissland, E.L., Valcarce, E.M. & Ashley, K.D. 1984. Explaining and Arguing with Examples. *AAAI-84, Proceedings of The National Conference on Artificial Intelligence*. Austin, TX. American Association for Artificial Intelligence.

Skalak, D.B. 1989 . Taking Advantage of Models for Legal Classification. *Proceedings, Second International Conference on Artificial Intelligence and Law*, 234–241. Vancouver, BC. Association for Computing Machinery.

Skalak, D.B. & Rissland, E.L. 1991. Argument Moves in a Rule-Guided Domain. *Proceedings, Third International Conference on Artificial Intelligence and Law*, 1–11. Oxford, England. Association for Computing Machinery.

Storrs, G. 1991. Extensions to Toulmin Form for Capturing Real Arguments. *Working Notes, AAAI Spring Symposium Series: Argument and Belief*. Palo Alto, CA.

Stucky, B.K. 1986. *Understanding Legal Argument*. Counselor Project Technical Report 13. Department of Computer Science, University of Massachusetts, Amherst, MA.

Stucky, B.K. & Gidley, J.M. 1990. Legal Argument with Counterplans. *Workshop Notes, Artificial Intelligence and Legal Reasoning Workshop, AAAI-90*. Boston, MA.

Swartout, W. & Balzer, R. 1982. On the Inevitable Intertwining of Specification and Programming. *Communications of the ACM*, 25(7), July.

Sycara, K.P. 1985. Arguments of Persuasion in Labour Mediation. *Proceedings, Ninth International Joint Conference on Artificial Intelligence*, 294–296. Los Angeles, CA. International Joint Conferences on Artificial Intelligence.

Sycara, K.P. 1987. *Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods*. Ph.D. Thesis, School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA.

Sycara, K.P. 1989. Argumentation: Planning Other Agents' Plans. *Proceedings, Eleventh International Joint Conference on Artificial Intelligence*, 517–523. Detroit, MI. International Joint Conferences on Artificial Intelligence.

Sycara, K.P. 1991. Pursuing Persuasive Argumentation. *Working Notes, AAAI Spring Symposium Series: Argument and Belief*. Palo Alto, CA.

Toulmin, S. 1958 . *The Uses of Argument*. Cambridge, England: Cambridge University Press.

Twining, W. & Miers, D. 1982 . *How To Do Things With Rules* (Second Ed.). London: Weidenfeld and Nicolson.

Winston, P.H. 1980. Learning and Reasoning by Analogy. *Communications of the ACM*, 23(12).

Winston, P.H. 1982. Learning New Principles from Precedents and Exercises. *Artificial Intelligence*, 19, 321–350.

Wu, H.J.P. & Lytinen, S.L. 1991. Attitude and Coherence Reasoning in Persuasive Discourse, *Working Notes, AAAI Spring Symposium Series: Argument and Belief*. Palo Alto, CA.