Within the Letter of the Law: open-textured planning

Kathryn E. Sanders *
Department of Computer Science
University of Maryland
College Park, MD 20742
sanders@cs.umd.edu

Abstract

Most case-based reasoning systems have used a single "best" or "most similar" case as the basis for a solution. For many problems, however, there is no single exact solution. Rather, there is a range of acceptable answers. We use cases not only as a basis for a solution, but also to indicate the boundaries within which a solution can be found. We solve problems by choosing some point within those boundaries.
In this paper, I discuss this use of cases with illustrations from CHIRON, a system T have implemented in the domain of personal income tax planning.

1 Introduction

Most case-based reasoning (CBR) systems to date have used a single "best" or "most similar" case as the basis for a solution. For many problems, however, there is no single exact solution. Rather, there is a range of acceptable answers. We use cases not only as a basis for a solution, but also to indicate the boundaries within which a solution can be found. We solve problems by choosing some point within those boundaries.

For example, suppose you want to sell your house. Two similar houses have been sold recently in the same neighborhood, one for $150,000 and one for $200,000. You can estimate that your house will sell for between $150,000 and $200,000. Or suppose you want to make chili for dinner. You have made the recipe with one tablespoon of chili powder and with three, and while you prefer the latter, both were acceptable. You can add the chili powder without measuring it precisely, as long as you stay within these limits.

'This work has benefited from the comments of Eugene Charniak, Tom Dean, Jim Hendler, Leslie Kaelbling, Robert McCartney, Leora Morgenstern, Edwina Rissland, David Skalak, and the members of the PLUS group at the University of Maryland. It was supported in part by grants from ONR (N00014-J-9I-1451), NSF PYI Award (IRI-8957601) to Thomas Dean, AFOSR and ARPA (F30602-91-C-0041), ONR (N00014-91-J-4052), ARPA Order 8225, NSF and ARPA (IRI-8905436), IBM (17290066, 17291066, 17292066, 17293066), and by NSF (IRI-8801253).

In this paper, I discuss this use of cases in the legal domain. In law, the reported cases are nearly always boundary cases in some respect; central cases are easy, so they are usually not taken to court.

Specifically, I focus on the way in which cases indicate the boundaries of the possible interpretation of statutes. Statutes are rules that have been created formally, by legislation. They are published by the government and often by private companies as well. For example, consider §1034(a) of the Internal Revenue Code, governing the tax treatment of income from the sale of a personal residence, given in Figure 1. The Internal Revenue Code is the most important statute for United States tax planning. It contains approximately 7000 sections.

Detailed as the Internal Revenue Code is, it still contains phrases that are not defined within the statute, for example, the phrase "principal residence" in § 1034. To qualify for the benefit of §1034, a taxpayer must show, among other things, that he bought and sold properties which belong to this category. These phrases are partially, but not completely, defined by commonsense knowledge about the meaning of the words used and by examples.

What makes reasoning about these phrases difficult — and interesting — is that defining them is not just a matter of inferring defining characteristics from a set of examples. Generally speaking, there is no set of essential characteristics shared by all positive instances of the statutory predicate. Some examples are typical, and others are more or less similar to them along various dimensions. As a result, classifying a particular object as an instance or noninstance of one of these categories is not always a simple task.

Planners using such open-textured rules need some way of determining whether their plans satisfy the rules. This issue arises throughout legal reasoning, not just in tax. Indeed, it is part of a general natural language problem. Many ordinary categories, such as "tiger" or "cup," are surprisingly difficult to define. This indeterminacy has been studied in linguistics and philosophy, where it is labelled *open texture* [Waismann, 1965; Hart, 1961]. Any planning rule expressed in natural language, such as "be careful," "never get involved in a land war in Asia," or "buy low, sell high," suffers from the same problem.

In any domain, open-textured rules can be partially

§1034. Rollover of gain on sale of principal residence. (a) Nonrecognition of gain.-If property (in this section called "old residence") used by the taxpayer as his principal residence is sold by him and, within a period beginning 2 years before the date of such sale and ending 2 years after such date, property (in this section called "new residence") is purchased and used by the taxpayer as his principal residence, gain (if any) from such sale shall be recognized only to the extent that the taxpayer's adjusted sales price (as defined in subsection (b)) of the old residence exceeds the taxpayer's cost of purchasing the new residence.

Figure 1: §1034(a) of the Internal Revenue Code.

defined by examples. The legal domain has the advantage that examples are recorded and published. Each court case is an example — an application of the law to a particular set of facts. Facts and results are recorded by the courts in "opinions" and published, both by the government and by private companies. Thousands of examples are readily available in any law library.

The open-textured nature of legal rules enables them to cover a range of possibilities, without specifying them in complete detail. For example, a house can be your principal residence whether you've lived there for one year or fifty, whether it has one bedroom or a hundred, and so forth.

The fact that terms like "principal residence" are underspecified also means that the courts can respond to changing circumstances. For example, they can interpret "principal residence" to cover cooperatives and condominiums, even if those forms of ownership did not exist at the time the statute was passed. Similarly, the First Amendment protection of freedom of speech can be extended to cover television and radio, as well as newspapers.

Because the system is flexible, it is also uncertain. In law, unlike domains such as chess, it is impossible to prove a plan correct. This uncertainty is not due to lack of factual information (we can assume complete knowledge of the facts); but to the underspecified nature of the rules.

Planners need some ability to predict how the law will be applied. This is supplied by the doctrine of precedent, which provides that, unless there is some strong reason to do otherwise, courts must follow their own previous decisions. Similar cases must be decided similarly.

Suppose you want to take advantage of one of these open-textured rules. Say for example that you want to sell your house and obtain the benefits of § 1034(a). You will consider your situation in relation to past cases under this provision and try to construct a plan for selling your house that is supported by previous successful cases. Because the courts are bound by precedent, if your case is similar to or stronger than past successful cases, it will be decided the same way. If it is weaker than past successful cases, you might still win, but this is an adversarial domain. To avoid challenge by the government, you will try to stay within the boundaries indicated by those cases.

At least, that is what you will do if you are a typical conservative tax planner. If you are a little more aggressive, and the successful cases indicate a trend, you may go beyond them in the same direction. For example, if you have a successful case where someone sold a house they only lived in for a year, you may try selling your house after six months. In other words, you will extrapolate from the cases in your casebase, rather than interpolating between them.1

1 have implemented this approach in CHIRON, a system that constructs plans in the domain of personal income tax planning [Sanders, 1994]. In Section 2 of this paper, I give an overview of the system; in Section 3, 1 describe the representation it uses for rules, cases, and other knowledge; in Section 4, I give a brief example; in Section 5 I discuss related work; and in Section 6, I summarize the results of this paper.

2 Overview of CHIRON

In order to reason about open-textured rules and cases, CHIRON combines hierarchical and case-based planners in a hybrid system. The hierarchical planner (adapted from a version of Nonlin implemented at the University of Maryland [Ghosh *et ai,* 1991]) calls the case-based module for guidance in choosing rules to use in refining its plans and refining the open-textured predicates that remain when the rules run out. The case-based planner uses the hierarchical planner to help with indexing, controlling adaptations, and combining plans.

Because law is an adversarial domain, it is useful to have a safe interpretation of the rules, what tax planners refer to as a "safe harbor plan." If this plan is executed, it will probably be safe from challenge by the government; the more a plan differs from it, the more likely it is to be challenged (and to lose). Accordingly, CHIRON associates a prototype "safe-harbor plan" with each of the hierarchical planner's strategies.

CHIRON uses cases to indicate how much a plan can differ from the prototype. It starts by instantiating the prototype and adapting it only as much as necessary to fit the facts of the current situation. Some adaptations consist simply of adding a fact to the prototype; others involve varying the parameters of a given fact. Any fact in one of the previous cases can potentially be added to the prototype to make a new plan. Similarly, facts can be subtracted from the prototype, and parameters can be varied, as suggested by the cases in the casebase. If the resulting plan is weaker than any of the previous successful cases along some dimension, the system rejects it, unless there is a trend towards weakening cases along that dimension.2

The prototype, cases, and adaptations define a space of possible plans for satisfying a rule. In general, the system generates plans that fall between the prototype and previous successful cases of a given strategy. Within

1 Like the other examples in this paper, this one is strictly hypothetical, and not to be relied upon as tax advice.

2 For this purpose, a "trend" is defined simply as two cases, one decided earlier than the other, where both decided in favor of the taxpayer, both were weaker than the prototype along a given dimension, and the second was weaker than the first.

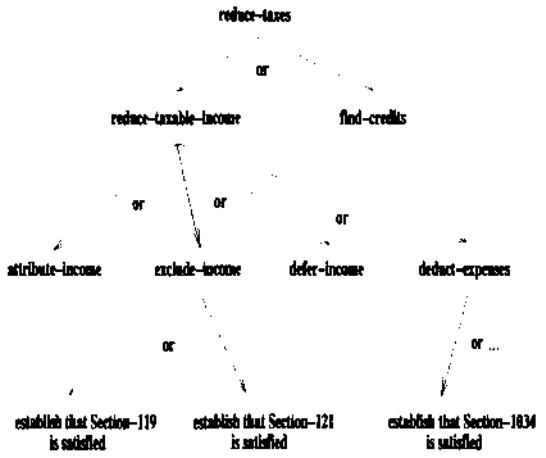

Figure 2: Part of the plan hierarchy.

that space, it generates plans that are as similar to the prototype as possible.

3 CHIRON's Knowledge Representation

CHIRON's knowledge base includes representations of the Internal Revenue Code and approximately twenty-four cases under various provisions of that statute. It also includes safe harbor plans, or prototypes, that satisfy the rules; a representation of the relationship between the rules, prototypes, and cases; and finally, a representation of the input description of the taxpayer's goals and current situation.

CHIRON has a number of tax-reduction strategies, including, for example, selling a house and buying another to get the §1034 rollover, making a charitable donation, and selling a house and taking the §121 exclusion. These strategies can generally be used either individually or in combination. All of these strategies and all of their possible combinations are candidate plans. Each strategy corresponds to a leaf node in the hierarchical planner's plan hierarchy, a portion of which is given in Figure 2.

The rules in the hierarchy are based on statutory provisions: to reduce your taxes, you can reduce your taxable income or find credits; to reduce your taxable income, you can attribute your income to someone else, exclude it, defer it to a later year, or deduct expenses; and so forth. Each leaf node, or strategy, corresponds to some open-textured portion of the rules. Associated with each leaf node, there is a prototype, a set of deformations, or ways in which the prototype can be adapted, and a set of past cases in which the taxpayer attempted (successfully or unsuccessfully) to satisfy the given rule.

Cases are represented using a structure that has fields for the name (the official legal citation), short-name (an abbreviated form of the citation for use in text), court, date, facts, strategies involved, and holdings of the case, as well as various indices. Information about which features strengthen or weaken a case with regard to a particular strategy is stored separately.

The facts of both previous cases and the current situation are represented as lists of propositions, including for each case a precise representation of the facts as given in the official reports for that case and in addition, abstractions from those facts that seem useful, based on the reasoning in each case, other cases, or general domain knowledge. Representing this knowledge required features such as space, time, action, permission, obligation, knowledge, belief, and intention, as other researchers in this domain have noted [Cuthill and McCartney, 1993; McCarty, 1989b]. Particularly characteristic of United States income tax law are transactions: sales, gifts, loans, and other transfers of money or property. In addition, unpredictable idiosyncratic details occur in each new case, such as the fact that the kitchen in a house has a tile floor, or the taxpayer is a war veteran, or the taxpayer has two children, one of whom is ten years old.

To represent these facts as accurately as possible, it is necessary to have a representation language that is both flexible and extensible. I chose an extension of the temporal logic developed by Shoham in [Shoham, 1988], modified to incorporate the modal operators "know," "believe," "want" (or "goal"), and "obligated." This language is influenced by McDermott's temporal logic and also by the work of McCarty in representing legal concepts, especially [McCarty, 1989b].

Part of the representation of one of CHIRON's cases is given in Figure 3; for the full representation of this case and the original text of the case, see [Sanders, 1994].

Associated with each strategy is a prototype. A prototype is represented using the same structure as a case, but with a conservative, safe set of facts, very likely to satisfy the open-textured provisions setting forth that strategy. The prototypes' facts are derived from general domain knowledge, including the commonsense meaning of statutory phrases, treatises, and the cases themselves. In any law case, there are some easy questions that are not at issue, and these questions provide some information about the prototypical case. The prototypes' facts are templates, or lists of generalized fact patterns, instantiated for each situation in which CHIRON attempts to use them. Similarly, in a treatise or regulation, a lawyer would find a generalized plan, with the details of particular cases to be filled in by the lawyer.

The relationship between the prototypes and cases is captured in the system's adaptation knowledge. Associated with each predicate in CHIRON's vocabulary is a com pare-function, which determines whether two facts with that predicate are comparable (matching, or one stronger than the other) with respect to a particular strategy. Some facts, where present, are known to strengthen or weaken the taxpayer's case. In some cases, there is domain knowledge concerning how a fact can be made stronger or weaker. Where known, this information is also associated with the predicates. A plan can be adapted by changing any of its facts into a stronger or weaker fact or by adding or subtracting facts.

4 Example

CHIRON solves a cluster of problems having to do with

```
(make-case
 :name "Hughston v. Commissioner, T.C.M. (P-H) 50,188 (1950)"
 :court "Tax Court"
 :date 1950
 :facts '((occurs (individual-return return2) (1947) (1947))
  (occurs (taxpayer return2 Hughston) (1947) (1947))
  (occurs (house house1) (October 1945) (December 1947))
  (occurs (real-property house1) (October 1945) (December 1947))
  (occurs (spatial-part room1 house1) (1947) (1947))
  (occurs (bathroom room1) (1947) (1947))
  (occurs (room room1) (1947) (1947))
  (occurs (real-property room1) (1947) (1947))
  (occurs (floortype room1 tile) (1947) (1947))
  (occurs (house house2) (October 1945) (December 1947))
  (occurs (real-property house2) (October 1945) (December 1947))
  (occurs (physically-occupy Hughston house1) (October 1945)
          (February 27 1947))
  (occurs (physically-occupy Hughston house2)
          (February 28 1947) (1951))
  (occurs (employment employment3) (1946) (1950))
  (occurs (employer employment3 Shell-Oil-Company) (1946) (1950))
  (occurs (corporation Shell-Oil-Company) (1946) (1951))
  (occurs (employee employment3 Hughston) (1946) (1951))
  (occurs (location employment3 (Houston Texas USA)) (1946)
          (February 1947))
  (occurs (location employment3 (Midland Texas USA))
          (March 1947) (1950))
  (occurs (selling selling1) (February 27 1947) (February 27 1947))
  (occurs (seller selling1 Hughston)
          (February 27 1947) (February 27 1947))
  (occurs (object selling1 house1)
          (February 27 1947) (February 27 1947))
  (occurs (selling selling2) (February 28 1947) (February 28 1947))
  (occurs (buyer selling2 Hughston)
          (February 28 1947) (February 28 1947))
                          ... )
 :transfer-types
       '((:sale :real-property :cash)(:sale :cash :real-property))
 :strategies '(:like-kind-exchange)
 :holdings '((:like-kind-exchange :government)))
```

Figure 3: Part of CHIRON's representation of *Hughston v. Commissioner.*

buying, selling, renting, and owning residential housing. Many of the provisions of the United States Internal Revenue Code affect these transactions, directly or indirectly. In addition to the general provisions governing sale of capital assets, payment of interest on loans, and so forth, there are special benefits for residential housing. For example, under certain circumstances, the gain from the sale of an individual's principal residence can be deferred or completely excluded from taxation. §§121, 1034.

CHIRON's examples have been chosen in part because they are simple; most tax planners would agree on at least the obvious solutions. In addition, they illustrate problems such as timing and satisfaction of open-textured rules that are typical of the domain. Finally, since they all involve transfers of residential housing, the commonsense knowledge that must be formalized to handle them involves the same cluster of concepts.

Consider a straightforward example: the taxpayer, whose name is Greenlee, wants to sell a house. She has owned the house, identified by the token 32-Eleventh-Street, since October 30, 1972, and has occupied the house during that entire period. She is now fifty-two years old. CHIRON takes as input the internal representation of these facts.

Next, CHIRON determines that the current situation involves one type of transaction: a sale of real property in return for some other unspecified type of property. CHIRON then retrieves the cases in its casebase that also involved a sale of real property. There are four, *Welch v. Commissioner,* T.C.M. (P-H) 79,010 (1979), *Trisko v. Commissioner,* 29 T.C. 515 (1957), *Sayre v. United States,* 163 F.Supp. 495 (S.D.W.Va. 1972), and *Hughston v. Commissioner,* T.C.M. (P-H) 50,188 (1950).

Part of the representation of *Hughston* is shown above in Figure 3.

These cases suggest three strategies: a §121 exclusion, which permits a taxpayer who is over 55 years old to sell his or her principal residence and exclude part of the gain from tax, if he or she meets certain tests; a rollover, where the taxpayer sells one principal residence and buys another; and a like-kind exchange, where the taxpayer exchanges one piece of property for another similar one.

CHIRON now attempts to construct a plan for each of these strategies in the current situation, starting with the like-kind exchange. It rejects this strategy on the grounds that there is no support for the plan along the dimension "duration-of-occupancy." The taxpayer has been living in her house for many years. The properties involved in a like-kind exchange must be "held for productive use in a trade or business or for investment...." §1031. Personal use of of the property weakens the taxpayer's case, and there are no cases in CHIRON'S casebase with this weakness, so the system rejects the plan. Similarly, it rejects the §121 exclusion, because the statute sets a minimum age of 55 for taking advantage of the provision, the taxpayer is only 52, and there is no case support for a lower age.

With regard to the rollover, however, the taxpayer is in a strong position. Her case is no weaker, and in some respects stronger, than the prototype for this strategy. As a result, the system suggests the following plan. The taxpayer is advised to live in the house until the date of sale, sell it and buy another on the same date, and occupy the new house immediately. In addition, the taxpayer should have only one residence at the time of sale. This is a very conservative plan, but it can be executed by the taxpayer, given the simple facts of our example (from a tax point of view, at least; the realities of the housing market may make selling more difficult).

In support of the plan, the system notes its similarities with the prototype. As in the prototype, Greenlee has bought and occupied a house, and she will (if she executes this plan) sell a house that she has been occupying until the date of sale. In both the prototype and the current situation, the taxpayer has one residence at the time of sale. And she has been living in her house longer than the minimal amount of time required by the prototype. Finally, the system notes, the time elapsed between the sale of one house and purchase of another, if Greenlee follows the plan, will be shorter than that required by the prototype.

Next, the system prints out a response for the government. It cannot find any weaknesses in the plan (which is indeed quite standard), but it does find a case, *Welch v. Commissioner,* in which the taxpayer lost even though in some respects, the taxpayer's plan was even stronger than the prototype.

In *Welch,* the system points out, the government won even though the taxpayer had occupied his first residence for seven years, longer than the prototype requires, bought a new piece of property only ten months and 22 days after the sale of his old residence, and occupied the new property a year, eight months, and two days after the sale, both less than the two-year interval allowed by the prototype.

Since the government has cited a new case in response, the taxpayer offers a rebuttal. The best rebuttal here would be that *Welch* was decided under an old version of the statute. In the time period covered by this case, the statutory time limit was eighteen months. Because the taxpayer's house was not completed within the time limit, he was unable to occupy it until two months afterwards. The prototype case, on the other hand, is within the current statutory time limit.

CHIRON has no means of reasoning about the dynamic nature of statutes, however, so it must look for other distinctions. In addition to some minor differences between the transactions involved in the two cases, the system notes that the current plan is stronger than *Welch* along exactly the dimensions that the government had cited in comparing *Welch* to the prototype: Greenlee has lived in her house longer than the taxpayer in *Welch* lived in his, and, if she executes the plan, she will buy and occupy a new house sooner than the taxpayer in *Welch* did. Thus, even though the system assumes all its cases are decided under the same version of the statute, it provides strong support for the right conclusion in this case.

The system has now considered all the strategies recommended by the case-guided search mechanism and recommended one of them. It reminds the user which strategy was successful, and offers the user a chance to consider other possibilities. From here on, the user can request any strategy or combination of strategies he or she would like the system to consider; requests will be processed in a similar manner.

5 Related Work

Most CBR systems have used a single "best" or "most similar" case as the basis for a solution (see, e.g., [Hammond, 1986; Koton, 1988; Sycara, 1989; Kolodner, 1989]). Some systems can adapt and combine pieces of several old cases to solve a problem (Sec, e.g., [Alterrnan, 1986; Branting, 1990; Redmond, 1990; Zito-Wolf and Alterrnan, 1992; McCartney, 1993; Kettler, 1995]), but few have attempted to place the current situation in the context of a group of past cases.

The CBR systems that place the current situation in the context of a group of past cases have done so in order to reason about open-textured legal concepts. Gardner's program retrieves multiple cases and uses them to determine whether in a given situation a contract has been formed [Gardner, 1987]. Both HYPO and CABARET also retrieve multiple cases. Their case representation is more detailed than Gardner's and incorporate dimensions, so that cases can be compared with each other and with the current situation. These comparisons are used as a basis for generating arguments for and against the satisfaction of open-textured concepts, again in a given situation [Ashley, 1991; Rissland and Skalak, 1991].

Besides using multiple cases, other solutions to the problem of reasoning about open-textured rules have included prototypes [Bareiss, 1988; McCarty, 1980; 1989a]; annotating rules with arguments for and against their application in borderline cases [Bench-Capon and Sergot, 1988; Sergot *ei a*/., 1986]; and the single "best

match" from a set of cases with a rich, detailed semantic-net representation [Branting, 1990].

CHIRON is the first case-based planner to reason about its plans in the context of a group of previous cases. Its solution builds on previous approaches to reasoning about open-textured rules, combining prototypes, cases, dimensions, and arguments for and against the application of the rules. Prototypes are particularly useful in a planning system where, unlike legal analysis systems, the facts are not yet known. Accordingly, CHIRON uses prototypes and deformations, or adaptations, corresponding to its open-textured rules. In order to limit the extent to which prototypes can be adapted, it uses cases. Each rule corresponds to a prototype and a set of cases in which that rule is interpreted. The cases are related to the prototype, and to each other, by the possible adaptations. In addition, CHIRON compares the current plan to cases interpreting the same rule in order to generate arguments for and against the plan's success. The system's detailed representation language for the facts of cases supports detailed comparisons and adaptations. None of the previous systems combines all of these features.

6 Conclusions

Most case-based reasoning systems have used a single "best" or "most similar" case as the basis for a solution. For many problems, however, there is no single exact solution. Rather, there is a range of acceptable answers. We use cases not only as a basis for a solution, but also to indicate the boundaries within which a solution can be found. We solve problems by choosing some point within those boundaries. Some case-based legal analysis systems have used the technique of placing the current situation within the context of a set of past cases, but it has not been used in a problem-solving system.

In this paper, I discuss this use of cases with illustrations from CHIRON, a system 1 have implemented in the domain of personal income tax planning. CHIRON uses a prototype, or safe harbor plan, representations of actual legal cases, and HYPO-like dimensions to define a space of possible plans and constructs plans to fit within that space. This solution is especially well-suited to adversarial domains, where varying too far from the prototype may cause a plan to be challenged. The basic idea of planning within boundaries could also be translated into a domain without explicit adversaries, such as cooking.

References

[Alterman, 1986] Richard Alterman. An adaptive planner. In *Proceedings of the Fifth National Conference on Artificial Intelligence,* Philadelphia, Pennsylvania, 1986. (reprinted in *Proceedings of a Workshop on Case-Based Reasoning,* 1988, and *Readings in Planning,* edited by James Allan, James Hendler, and Austin Tate, Morgan Kaufman, 1990).

[Ashley, 1991] Kevin D. Ashley. *Modelling legal argument: reasoning with cases and hypotheticals.* MIT Press, 1991.

[Bareiss, 1988] Ray Bareiss. Protos: a unified approach to concept representation, classification, and learning. Technical Report AI88-83, University of Texas at Austin, 1988. (PhD Thesis).

[Bench-Capon and Sergot, 1988] Trevor Bench-Capon and Marek J. Sergot. Towards a rule-based representation of open texture in law. In *Computer power and legal language,* pages 39-60. Quorum Books, New York, 1988.

[Branting, 1990] L. Karl Branting. Integrating rules and precedents for classification and explanation: automating legal analysis. Technical Report AI90-146, Artificial Intelligence Laboratory, Department of Computer Sciences, University of Texas at Austin, 1990. (PhD Thesis).

[Cuthill and McCartney, 1993] Barbara Cuthill and Robert McCartney. Issue spotting in legal cases. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law,* Amsterdam, pages 245-253, 1993.

[Gardner, 1987] A nne v.d.L. Gardner. *An artificial intelligence approach to legal reasoning.* MIT Press, Cambridge, Massachusetts, 1987.

[Ghosh *et al.,* 1991] Subrata Ghosh, James Hendler, Subbarao Kambhampati, and Brian Kettler. *Common Lisp Implementation of NONLIN USER MANUAL.* Computer Science Department, University of Maryland, College Park, Maryland 20742, February 1991.

[Hammond, 1986] Kristian J. Hammond. Case-based planning: an integrated theory of planning, learning, and memory. Technical Report YALEU/CSD/RR 488, Yale University Department of Computer Science, 1986. (PhD Thesis).

[Hart, 1961] H. L. A. Hart. *The concept of law.* Clarendon Press, Oxford, 1961.

[Kettler, 1995] Brian P. Kettler. Case-based planning with a massively parallel memory, ph.d. dissertation. Technical report, University of Maryland Department of Computer Science, 1995. (In preparation).

[Kolodner, 1989] Janet L. Kolodner. Judging which is the 'best' case for a case-based reasoner. In *Proceedings of a Workshop on Case-based Reasoning,* pages 77-86,1989.

[Koton, 1988] Phyllis Koton. Reasoning about evidence in causal explanations. In *Proceedings of a Workshop on Case-based Reasoning,* pages 260 270, 1988.

[McCartney, 1993] Robert McCartney. Planning from partial and multiple episodes in a case-based planner. In *Proceedings of the 1993 AAAI Workshop on Case-Based Reasoning,* Washington, D.C., pages 94-100, 1993.

[McCarty, 1980] L. Thorne McCarty. The TAXMAN project: towards a cognitive theory of legal argument. In Bryan Niblett, editor, *Computer Science and Law,* pages 23 43. Cambridge University Press, Cambridge, England, 1980.

[McCarty, 1989a] L. Thorne McCarty. Computing with prototypes (preliminary report). In Proceedings of the Bar-Han Symposium on the Foundations of Artificial Intelligence, 1989.

[McCarty, 1989b] L. Thorne McCarty. A language for legal discourse: I. basic features. In Proceedings of the Second International Conference on Artificial Intelligence and Law, Vancouver, British Columbia, pages 180-189, 1989.

[Redmond, 1990] Michael Redmond. Distributed cases for case-based reasoning; facilitating use of multiple cases. In Proceedings of the Eighth National Conference on Artificial Intelligence, Boston, Massachusetts, pages 304-309, 1990.

[Rissland and Skalak, 1991] Edwina L. Rissland and David B. Skalak. CABARET: rule interpretation in a hybrid architecture. International Journal of Man-Machine Studies, pages 839-887, 1991.

[Sanders, 1994] Kathryn E. Sanders. Chiron: planning in an open-textured domain. Technical Report 94-38, Computer Science Department, Brown University, 1994. (PhD Thesis).

[Sergot et ai, 1986] M. J. Sergot, F. Sadri, R. A. Kowalski, F. Kriwaczek, P. Hammond, and H. T. Cory. The British Nationality Act as a logic program. Communications of the Association for Computing Machinery, 29(5):370-386, 1986.

[Shoham, 1988] Yoav Shoham. Reasoning about change: time and causation from the standpoint of artificial intelligence. MIT Press, Cambridge, Massachusetts, 1988.

[Sycara, 1989] Katia Sycara. Using case-based reasoning for plan adaptation and repair. In Proceedings of a Workshop on Case-based Reasoning, pages 425-434, 1989.

[Waismann, 1965] Friedrich Waismann. Verifiability. In Antony Flew, editor, Logic and language: first and second series, pages 122-151. Anchor Books, Garden City, New Jersey, 1965. (first published in Proceedings of the Aristotelian Society, 119-150 (1945).).

[Zito-Wolf and Alterman, 1992] Roland Zito-Wolf and Richard Alterman. Multicases: a case-based representation for procedural knowledge. In Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society, 1992.