

Artificial Intelligence and Legal Reasoning

A Discussion of the Field & Gardner's Book

Edwina L. Rissland

In this article, I discuss the emerging field of artificial intelligence and legal reasoning and review the new book by Anne v.d.L. Gardner, An Artificial Intelligence Approach to Legal Reasoning, published by Bradford/MIT Press (1987, 225 pp., \$22.50) as the first book in its new series on the subject.

Legal reasoning is an intriguing field for the researcher in artificial intelligence because it demands that many deep and nettlesome problems, for instance, those of knowledge representation and analogy, be addressed head on. It presents a range of interesting reasoning skills, some of which seem tantalizingly tractable with currently well-understood methodologies such as those of expert systems, and others of which seem to subsume exceedingly hard problems such as natural language understanding.

The law is an attractive domain for AI research for several reasons. First, it has a tradition of examining its own reasoning process. Second, its reasoning is stylized; in Anglo-American common law, one reasons according to *stare decisis*, or the doctrine of precedent, in which similar cases are to be decided similarly. Central aspects of such reasoning involve analogy and reasoning with cases. Third, much of the knowledge is readily accessible and some of it is very well structured and codified.

Legal scholars and philosophers have a long tradition of grappling with the sort of issues—like the status of rules and exceptions—that are of interest to AI. Legal philosophy, or jurisprudence, seeks to examine the basis and workings of the legal system and to elucidate, among other things, the nature of legal concepts, rules and principles, the process of *stare decisis*, and the role of society and morality in law. While several legal scholars, (for example, Levi 1949; Llewellyn 1981;

Dworkin 1977, 1985) offer insights valuable to AI, their jurisprudential analyses often raise more questions than they answer and their insights, couched in philosophical discourse, are difficult to harness computationally.

Much of the knowledge used in legal reasoning is published, codified, and highly indexed. The legal system maintains extremely detailed records of its cases and commentary on them and except for the lowest-level courts, all cases are published and indexed commercially. For example, *Shepard's Citations* records and updates all forward and backward pointers for cases; that is, for a given case, all the succeeding cases citing it as well as all the cases it cited. Some areas of the law are structured in ways that greatly facilitate knowledge acquisition and representation. For instance, some areas such as tax law have statutes so carefully drafted that they read almost like computer pseudocode or sets of logical propositions. In some areas, such as contract law, there are well-organized efforts of legal scholars and practitioners to tease out, debate, and codify the principles and rules of the legal domain. For example, the *Restatement of the Law, Second: Contracts* (1981) is a compendium of principles, illustrative cases and hypotheticals, discussion, and case citations. Its 385 principles are abstractions of the holdings of thousands of contract cases. These were worked out in a series of meetings, held between 1962 and 1979 for the second *Restatement*. Certain sections of the first *Restate-*

ment, published in 1932, were hotly debated. Such areas look like a knowledge engineer's dream, since much of the epistemological homework has already been done.

Despite all these attractive resources, much work remains to represent any area of the law in a way usable by an AI program. No matter how well drafted, statutes use terms whose meaning is not clear without the sort of interpretive reasoning at the heart of reasoning with cases. For example, there are still cases on the meaning of "income" as it pertains to certain sections of the code governing the IRS. Statutes often contain conflicting and ambiguous provisions. The great codification efforts, even though they involved an enormous amount of effort by the best legal thinkers, are not perfect; for example, there is circularity in some of the definitions. Thus, the central question of "What is the knowledge?" is only partially answered by such material. Furthermore, one should not fall prey to the seductive suggestion implicit in such seemingly well-worked out materials: the idea that the law works in a purely mechanically logical way is worse than superficial; it is false.

Thus, because of its stylized reasoning, its penchant for philosophical self-examination, and the efforts of its own domain experts to inventory and organize its knowledge, the law is an attractive domain for AI research. The challenge to AI is to produce models of legal reasoning that are faithful to the law's characteristic features, are not jurisprudentially flawed, and are computationally tractable.

Some Interesting Characteristics of the Legal Domain

There are some characteristics of the legal domain that make it unique. Many of these concern legal cases and rules. While other domains, like medicine, may share some of these characteristics, no others do so in quite the mix or degree that the law does. Before discussing them, I need to say a word about cases and rules.

For our purposes, a case is a legal

dispute with supporting collections of facts and arguments, which has been decided by a court. A hypothetical case or "hypo" is the same as a real case except that it hasn't been decided by a court. The word "case" is often used as a synonym for a court's opinion, a written essay that presents the court's "holding" and its supporting reasoning. Even though the opinion, if there is one, is a readily available representation of a case, a case, of course, is more than the opinion. There are also the legally established facts, not always recited in the opinion, various legal motions, trial records, procedural histories, briefs, oral arguments, and so on. Thus what counts as a case is a central question that must be resolved before one can construct a program that reasons with cases.

The word "rule" is used in manifold ways in the law. There is the "rule of the case," that is, a short statement or proposition of what the case stands for. There are legal rules, as in the principles abstracted from several cases or a whole body of law; sometimes, these are catalogued in sources like the *Restatement* and other times, they are well-known but not necessarily written down. Some rules are maxims, such as "one should not profit from one's misdeeds." Some rules are statutory, that is, adopted by a legislative body (for example, the Internal Revenue Code of 1954). Some rules are regulations adopted by an administrative agency (for example, the Treasury regulations governing administration of the IRS code). In this discussion, I will not be especially careful about keeping these usages distinct; however, for the most part, anything I say applies to all these senses of the word "rule." Of course, there are also various usages of "rule" in AI, such as rule, as an object of knowledge, and rule as in the if-then implementation in expert systems.

One way to characterize the law—for the most part, I am speaking of Anglo-American law—is by how it tries to find answers to a legal problem such as a dispute.

Answers are Not Derived by Logical Deduction

Legal rules do not have the force that rules or their equivalents have in some

other domains, mathematics being an extreme example; one does not logically chain them together to reach an answer. Legal rules have a status more like that of heuristics than of theorems, in the sense that the joining of antecedents and conclusions is not ironclad and all rules have exceptions. Further, deciding whether antecedent conditions are met often involves much interpretation and analogical reasoning. Such qualities are shared by rules in other domains familiar to AI, such as medicine. One quality not typical in other domains is that one can always argue about the correctness or goodness of a rule (for example, with moral or public policy arguments) and whether it should be followed, even if it applies.

In some areas, there are rules that lead to conflicting answers. Even very stable areas of law, like contract law, have rules that for certain facts, will lead to two logically opposite conclusions. One example, concerning an "acceptance with proposal to modify," is discussed by Gardner and touched on later in this article.

Yet in the law, there are situations in which one can profitably act as if the reasoning is actually rule-driven rather than "rule-guided" (to use Gardner's term) and as if rules can be manipulated in a more traditional logical manner. However, one typically needs to supplement deductive reasoning with other modes of reasoning involving cases and analogies. How and when to combine rule-based and case-based reasoning is a challenging problem.

The Terms Used are Open-Textured

Legal reasoning involves a great deal of interpretation of legal terms and predicates, many of which have meanings or definitions that are inherently indeterminate. Such "open-textured" predicates admit no necessary and sufficient conditions to classify examples into instances and noninstances. This problem of open-textured predicates is related to the idea of natural-kind classes in philosophy. For example, while it is true that most chairs and dogs have four legs, this is neither a necessary nor a sufficient condition for being a positive instance of either concept. Legal concepts are like open

Legal concepts are like open sets in mathematics: there is always room to sneak in a little epsilon neighborhood near the boundary.

sets in mathematics: there is always room to sneak in a little epsilon neighborhood near the boundary. Furthermore, they will never have hard boundaries no matter how much one tries to force them to be otherwise (even if this were a wise idea). Besides not having sharp boundaries, all concepts have exceptions, and thus to continue with the point set metaphor, the set of positive instances of a legal concept has holes in it. Also, the meanings of concepts change over time.

These characteristics of legal concepts mean that there is frequently reason to argue whether a given instance should be included in a class. While other domains like medicine also have open-textured concepts (for example, hypertension), in the law, the whole case might depend on the interpretation of a term. Thus, not to confront the centrality of reasoning with open-textured and exception-riddled terms is to invite disaster. In the legal domain, no amount of definitional craftsmanship will fix such problems, since they often spring from societal sources. There is no hope of ever pinning down an underlying causal model or set of axioms, as there is in other domains like biology or mathematics.

These caveats notwithstanding, there is recognition that many concepts do have stereotypical or clear instances and noninstances and that it is silly or wasteful of resources to argue about them. The focus of legal reasoning in such cases is not on the resolution of the terms but on other matters.

Open-textured predicates, exceptions, and conflicting rules are sources of what are called hard questions in the law. A hard question is one on which the experts—the judges, lawyers, scholars, courts—disagree and where there is room for judicial interpretation. An easy question or clear question is one on which there is consensus.

There Is Often More than One Answer

Conflict, disagreement, and argument are part and parcel of the law. The adversarial nature of the law derives from at least two sources: (1) the ill-defined nature of rules and concepts

implies that answers are more often “maybe” than “yes or no” and (2) our Anglo-American legal institutions work out problems through an adversarial process. As Gardner points out, not only are actors in our legal system free to argue, they are expected to do so. Experts are expected to be able to deal with hard questions and to see and uncover the hard aspects of even those that appear easy. Law school training emphasizes this almost to excess. Most interesting cases (for example, those that have been appealed, especially, as far as the United States Supreme Court) present at least two alternative answers, or they wouldn’t have ended up as legal disputes or under appeal in the first place. Furthermore, legal arguments occur on many planes ranging from what the facts are, to what rules and cases are relevant, to whether or not such rules or cases should in fact be followed, to what the societal implications are and how we as a society ought to resolve them.

But the Law Must Provide an Answer

The legal system must give a timely answer for one side or the other; it can’t say “maybe” or “come back in 20 years” (although it may seem like this to some involved parties). In reaching its decision, a court will countenance only so much argument; this amount might be large, but it is definitely limited, and there are real costs associated with litigation. Thus, legal reasoning is resource-limited and reasoners must pick their points of contention carefully. One cannot afford to turn every question into a federal case; some are easy and should be treated as such and others are hard and should be argued. The skillful lawyer knows how to distinguish the hard from the easy, when to concede a hard question as an easy one, and when to turn a seemingly easy one into a hard one. The trick is how to pick and choose. As in many classical problem-solving situations having a large search space coupled with significant resource constraints, a heuristic approach can provide both answers and a focus for attention.

Answers Change

Through the accretion of cases and statutes and the changing needs and values of society, legal concepts and rules evolve. Sometimes the change is incremental, as in the mode of Kuhn's (1970) normal science. At other times it is abrupt and discontinuous, as in a Kuhnian paradigm shift (examples are rulings on desegregation and abortion). The implications of this for an AI approach are that there are hard learning issues lurking not too far beneath the surface and that to accommodate even gradual change one needs to allow for reasoning with a changing base of knowledge: at the very minimum, a growing base of cases and changing indices, rules, and norms that manipulate the cases. Eventually one must confront change in the predicates and the representation itself, for instance, through emergence of new legal concepts and substantial modification of old ones. Legal reasoning sometimes involves the hardest problems in learning, like bias and the new-term problem.

Thus, legal reasoning requires certain capabilities: (1) the ability to reason with cases and examples, particularly through analogy; (2) the ability to handle ill-defined, open-textured predicates; (3) the ability to handle exceptions; (4) the ability to handle fundamental conflicts between rules; (5) the ability to argue and justify in a case-based manner; and (6) the ability to handle change and nonmonotonicity.

Some AI Specialties Relevant to Legal Reasoning

Several specialties within AI immediately come to mind as relevant: (1) case-based reasoning, (2) expert (rule-based) systems, (3) logic, (4) natural language processing, (5) nonmonotonic reasoning, and (6) learning from examples. Each specialty has something to offer to AI and legal reasoning and also to gain from such research. The law is a domain par excellence for providing such specialties with examples of interesting reasoning.

For instance, since so much of legal reasoning is case-based, it is natural to look to the emerging specialty of case-based reasoning (CBR) for useful insights and techniques. Research in

case-based reasoning concentrates on several issues that are central to legal reasoning: in particular, case memory and indexing, assessment of similarity and relevancy, analogy of various sorts, context-dependent comparison of cases, generation and evaluation of arguments and plans, and generation and reasoning with hypothetical situations.

Some of the specialties will require some modification to accommodate the law's special characteristics. For instance, the open-textured nature of legal concepts and the arguable applicability or status of rules has implications for the use of both logic and expert systems. Incomplete or conflicting rules and the lack of black-and-white definitions means, for example, that the use of backchaining to resolve certain legal predicates will often fail because the "rules will run out," as Gardner puts it, before predicates have been resolved, or because there will be two conflicting ways of resolving them. Thus, if one uses logic or expert systems techniques to deal with the rulelike aspects of legal reasoning, one will need to supplement them with other techniques. One also must not forget that the style of justification, even in legal reasoning situations where one manipulates rules, is different in the legal domain, where the gold standard is the doctrine of precedent. Compare this with mathematics, where justification is based on logical deduction.

Much of the law seems intimately bound up with language, and some researchers feel that one cannot tackle legal reasoning without addressing language. Regardless of one's stance on that issue, natural language is just as difficult in the legal domain as it is anywhere else and for the usual reasons. The law is a microcosm of all human experience. In legal cases, people promise, buy and sell, steal, kill, appeal to higher moral authority, and so on, presenting all the hardest aspects of language and other well-recognized problems in AI, such as common-sense reasoning. Language understanding in a statutory domain might appear easier because a statute has its own defined terms and structure, but this is not as much of a blessing as one would hope because terms used in

statutes raise all the interpretative problems of case law. Further, the use of certain words that in other contexts are well defined can be problematic. For example, even the logical connective "and" has been used in a nonstandard way and given a disjunctive sense. Layman Allen has pointed out that there can be an awesome multiplicity of alternative interpretations of the logical structure of even the most carefully drafted regulation. For one statutory example, he finds 48 different structural interpretations stemming from interpretation of words such as "except," "if," "unless," and even "and." This also has serious implications for knowledge engineering. Thus, restricting oneself to natural language as used in the law is no simplifying assumption; using such language might even be more difficult because understanding the technical legal terms presents an additional processing burden.

While natural language is never easy, some of the techniques for story understanding might be applicable to certain areas where there are stereotypical fact patterns and actions by the involved parties. (Gardner used some of these techniques, but worked from individual sentences and phrases, not a continuous story.) With suitable restrictions, for instance, script-based understanding techniques might prove successful for understanding short case summaries such as those found in the headnotes of certain "reporters" that publish and index cases. Other areas sharing overlap of interest with natural language processing researchers are discourse, argumentation, and explanation.

Two other areas of AI that should eventually offer significant insights or techniques are nonmonotonic reasoning and example-based learning. The law exemplifies a nonmonotonic system since it often limits or overturns its past results. Because it is constantly refining old concepts and rules as well as carving out new ones, the law engages in activities that look like learning. Example-based learning, in particular, seems a natural area for mutual cross-fertilization, especially research on prototypes, explanation-based generalization, and similarity-based methods.

Research on legal reasoning has much to gain from its sister AI specialties. However, despite the obvious attraction of certain methods, *caveat emptor*. In particular, the researcher needs to make sure to use them only for problems for which they are well-suited. Perhaps more importantly, the researcher must be careful that using such methods does not lead to de-emphasizing certain hallmark characteristics of legal reasoning or shifting the focus of research away from critical reasoning tasks such as analogy and argument, where I feel it largely belongs.

Applying AI in the legal realm clearly is not easy but there is a growing community of researchers who believe it is also not impossible. They believe that what can be accomplished will be of interest from both the AI and legal points of view. They are willing to use what they can from the existing AI arsenal, extend it where necessary, and in the process contribute to AI, in general. Future contributions to AI should include findings on analogy, argumentation, and the melding of rule-based and case-based reasoning, to name a few. Because of the potential for shedding light on core questions in both AI and jurisprudence, the emergence of this new community of researchers—many of whom have both a J.D. and Ph.D. or significant expertise in both disciplines—and in particular, publication of Gardner's book, is a welcome event.

An Artificial Intelligence Approach to Legal Reasoning

Anne Gardner's book grew out of her effort to extend AI to deal with certain key legal issues.

The Task

In her Stanford computer science Ph.D. thesis, which is the basis of her book, Gardner chose to address the questions of how to recognize what the legal questions are, how to distinguish the easy from the hard questions and answer only the easy ones, and how to deal with the problem of what to do "when the rules run out" in attempts to resolve open-textured

On July 1 Buyer sent the following telegram to Seller: "Have customers for salt and need carload immediately. Will you supply carload at \$2.40 per cwt?" Seller received the telegram the same day.

On July 12 Seller sent Buyer the following telegram, which Buyer received the same day: "Accept your offer carload of salt, immediate shipment, terms cash on delivery."

On July 13 Buyer sent by Air Mail its standard form "Purchase Order" to Seller. On the face of the form Buyer had written that it accepted "Seller's offer of July 12" and had written "One carload" and "\$2.40 per cwt." in the appropriate spaces for quantity and price. Among numerous printed provisions on the reverse of the form was the following: "Unless otherwise stated on the face hereof, payment on all purchase orders shall not be due until 30 days following delivery." There was no statement on the face of the form regarding time of payment.

Later on July 13 another party offered to sell Buyer a carload of salt for \$2.30 per cwt. Buyer immediately wired Seller: "Ignore purchase order mailed earlier today; your offer of July 12 rejected." This telegram was received by Seller on the same day (July 13). Seller received Buyer's purchase order in the mail the following day (July 14).

Briefly analyze each of the items of correspondence in terms of its legal effect, and indicate what the result will be in Seller's action against Buyer for breach of contract.

—Gardner, p. 5

Figure 1. Example of an Issue Spotter Question.

predicates. Besides developing a computational embodiment of the hard/easy distinction, Gardner's work advances our understanding on such points as how examples can be used to augment and inform rule-based reasoning. In this work, Gardner, who also holds a J.D. from Stanford Law School and has had many years of legal experience, brings to bear her insights and expertise as both lawyer and computer scientist.

Gardner created a program "to recognize issues a problem raises and to distinguish between those it has enough information to resolve and those on which competent human judgements might differ" (p. 4). Her program is designed to sift the comparatively clear-cut issues, where the experts will not disagree, from the legally debatable ones, where there is room for judicial interpretation. She required her program to be somewhat self-aware in that it should not resolve certain issues as easy when it has insufficient expertise to recognize them as hard and that it should determine whether an issue is easy in a computationally inexpensive manner.

Gardner's program works in the subfield of offer and acceptance from contract law. An offer (for example,

"I'll sell you my car for \$1200") followed by an acceptance (for example, "It's a deal; I'll buy it") is a standard way to form a contract. It is a classic subdomain that is covered in any basic course on contracts and is a standard topic on bar exams.

Contract law is a well-worked out, stable body of law, what the legal scholar Edward Levi (1949) would call a "stage 2" domain, which corresponds roughly to Kuhn's (1970) normal science stage, with a full body of cases and nicely worked out rules and principles. These have been compiled and synthesized by contract experts into the *Restatement* treatise, which as mentioned is a compendium of principles, illustrative cases and hypotheticals, and case citations. Contracts is a common law—that is, case-based—and not a statutory domain, although certain contractual situations involving the sale of goods are covered by the Uniform Commercial Code.

The specific task Gardner chose to address is that of analyzing an "issue-spotter" type of question found on law school and bar exams (see figure 1). The task in such questions is to analyze a description of the facts of the case, pinpoint events that raise inter-

esting or important legal issues, and explore how resolution of them affects resolution of other issues, in particular whether or not there is a contract. Law school questions of this form are usually rather long hypotheticals specially crafted to raise difficult legal points. Thus, for a law student or bar examinee, the exercise entails understanding a fairly long legal story, analyzing it, and then writing an essay explaining and justifying the legal analysis. Gardner's program performs the analysis.

This task can be viewed as a search through a space of possible alternative interpretations of the facts. Since not all the choice points are worth arguing about, and in the context of an exam there are time and space constraints, the nature of the search must of necessity be heuristic. The task is how to select which interpretations can be defensibly argued, or in the language of search how to discern the plausible from the possible. In the words of Karl Llewellyn, a legal scholar of renown: "For while it is possible to build a number of divergent logical ladders up out of the same cases and down again to the same dispute, *there are not so many that can be built defensibly*. And of these few there are some, or there is one, toward which the prior cases pretty definitely press" (Llewellyn 1981, p. 81; the emphasis is Llewellyn's).

Thus Gardner's task is a central one in legal reasoning. Her program takes a moderate jurisprudential stand between those who believe there are no easy questions and those who question whether there are any hard ones. Although she does not use the language of "false positives" and "false negatives," her performance criterion—and that of Llewellyn—is that her program should minimize both. It should not raise a false alarm by calling an easy question hard and it should not miss any hard question by dismissing it as easy.

Gardner wisely circumvents the very difficult task of understanding a raw natural language version of such an exam question or the related task of establishing the facts. This is quite reasonable since both on issue-spotter questions and in appellate-level analysis one concentrates on questions of

interpretation. (For better or worse, the exam writer or the trial court established the facts and one must now work with them.)

Some of this interpretation is very much like full-blown natural language understanding in that it must understand what certain words mean, such as that "mailing" is a kind of "sending;" that such an act implies that there are several participants, the sender and the receiver; and that what was sent has some "content." However, it does not reach the level of difficulty reached in traditional story understanding or discourse research. The understanding is more on the level of making restricted inferences about certain types of speech acts, such as assertions or yes-no questions, and common-sense facts.

The Program

Gardner's program—which is never given a persona with a name—starts with a set of facts that have been entered by hand into a form acceptable to her program (for example, slot fillers). From these, the program generates a legal analysis which is then reported back to the user in the form of graphs summarizing the branchpoints and their underlying legal analyses. For instance, there is a question of whether the second telegram of figure 1 is to be interpreted legally as an acceptance or a counteroffer and if it is the second, whether the purchase order can be interpreted as a further counteroffer or an acceptance with proposal to modify. For each of these competing interpretations, there are many prerequisite legal antecedents to grind through. For instance, antecedents for the existence of an offer (based on Sections 24 and 33 of *Restatement [Second]*) include: There must be an act, with some symbolic content, done by some agent (the offeror), about some exchange, with terms, specified with reasonable certainty, such that the offeror manifests that he or she may-be-willing-to-enter the described exchange. Some of these ingredient prerequisites are well defined (for example, offeror, terms) and others are squarely of the open-textured variety (for example, reasonably certain, may-be-willing).

The relevant legal rules are also of varying inferential conclusiveness. Some, like the offer rule above, can resolve a legal point; others can only suggest conflicting legal theories. For instance, some contract scholars insist that either the offeree agrees exactly to what the offeror proposed or there is no acceptance at all; others tolerate some degree of mismatch and admit the concept of an acceptance with proposal to modify. These two legal theories do not agree as to whether such a thing as acceptance with proposal to modify even exists. Because it can accommodate and reason with conflicting rules, Gardner's program is able to deal with such unsettled issues in the law. Note that there is no question about the facts: a telegram was sent, a purchase order was sent. Rather, the question is what such facts represent from a legal perspective.

The output of her program is a two-level analysis represented in two graphs (see figure 2). The upper level is a summary of the interpretations of the events in the fact situation; branchpoints represent hard questions, that is, points where there are alternative, competing interpretations. In effect, the upper level is a decision tree, whose branching nodes represent hard questions and whose leaf nodes correspond to separate sequences of interpretations of the events. From such a representation of issues, it is a short inferential hop to answering the big questions of whether a contract can be said to exist and what the arguable issues are. For each upper-level branchpoint, there is a lower-level detailed analysis supporting the diverging interpretations.

To produce the analysis of the legal choices, Gardner's program employs several sources of legal, linguistic (for example, speech acts), and common-sense knowledge. Commonsense knowledge is represented in two ways: (1) by a slot-filler language used to describe fact situations; and (2) by a hierarchy of such things as events, states, and objects, implemented through a mechanism of common sense knowledge (CSK) rules. All are encoded in a standard fashion using Genesereth's MRS language. For example, "one carload of salt" is "(salt S1) (quantity S1 C1) (carloads C1) (number

and lost cases.

These ideas are embodied in an algorithm, which forms the backbone of Gardner's program. The mechanisms of applying CSK and legal rules to work one's way around the ATN are quite involved and occupy a large part of the discussion in chapters 5 and 6. In chapter 3, a preview of the algorithm is given on pages 54-55. In chapter 6, the algorithm is revisited and described in detail at the MRS level on pages 160-162.

The algorithm may be summarized as follows. First, determine if there is a tentative answer from the CSK rules. Then:

1. If there is a tentative answer, try to find "opposite" case examples that point to the opposite answer or resolution of a predicate.

a. If there are no opposites, then the question is easy and the tentative answer holds.

b. If there are opposites, check whether there are both positive and negative example cases, that is, both opposite and similarly aligned cases.

i. If there are both, then there are cases that point in both directions, thus the law is unsettled and subject to judicial interpretation, and the question is hard.

ii. Otherwise, there is a CSK rule that points one way and a technical legal answer that points the opposite. Let the technical legal meaning overrule and prevail. Therefore, there is no hard question; that is, the question is actually easy and its answer is the technical legal one.

2. If there is no tentative answer from the CSK rules, try to match for both positive and negative examples.

a. If no match can be found or there are both positive and negative examples, then the question is hard because either there are no cases suggesting how to resolve the issue or there are cases pointing to opposite resolutions.

b. Otherwise, the question is easy and the answer is that of the example matched.

Gardner's model can analyze issue-spotter questions such as that shown in figure 1, which has nine events requiring interpretation. For this problem, her program produces nine analyses, arising from the eight major two-way branchpoints shown in figure 2.

As she discusses, this is perhaps more than a human lawyer is likely to consider, but it compares well to her estimate of a possible search space of 5^9 . The program performed creditably on an assortment of problems from *Gilbert's Law Summaries* (Eisenberg 1982). It probably did well enough to pass these as bar exam questions, but not well enough to get an A or B at a place like Harvard or Stanford, or so my students tell me.

The Book

Gardner's book presents an explication of her program, its underlying architecture and knowledge representation; fundamental background on offer and acceptance law; and a discussion of the jurisprudential underpinnings, particularly on the nature of legal rules, mechanical jurisprudence and the paradigm of hard and easy cases. The following is a brief indication of what the chapters cover.

Chapter 1, "Introduction," describes what makes the legal domain special. Why it is an interesting domain for AI research. Some general remarks about the law, about contracts, about Gardner's approach.

Chapter 2, "Design I: The Place of Rules," gives a good survey discussion of the jurisprudential underpinnings and a quick review of some relevant episodes in legal history including the failure of "mechanical jurisprudence." It discusses the lessons of legal realism and the *Restatement, (Second), of Contracts* as a resource for knowledge engineering.

Chapter 3, "Design II: When the Rules Run Out," presents a thoughtful discussion of what happens when rule-based reasoning doesn't do the trick with particular attention to the problem of applying rules and predicates to a set of facts. It gives a good introduction to the exchange, known as the "Hart-Fuller debate," between Lon Fuller and H. L. A. Hart, which spawned such famous hypos as "the vehicle in the park" and "sleeping in the railroad station." The chapter covers how to handle legal words used to signal that there is a variable standard and room for judicial interpretation (for example, "reasonable") and ordinary English.

Chapter 4, "Related Work," provides the mandatory survey of who's doing what. It gives a complete picture of the field up to the summer of 1986.

Chapter 5, "Representing Problems" is the first of two chapters detailing the implementation. This first installment concentrates on how to represent facts without begging the legal conclusions. It emphasizes common-sense knowledge and the speech act approach.

Chapter 6, "Representing and Using Legal Knowledge" contains a very detailed discussion of legal rules and how the program encodes and reasons with them.

Chapter 7, "Program Performance" gives an extended presentation of how Gardner's program handles the exam question in figure 1, including a march through the two output graphs, how they were computed, and what they mean. It discusses how the program performed on five other bar-exam-type questions, including one in which the program spotted an issue that a standard review book missed.

Chapter 8, "Conclusion" discusses what it all means and where to go from here. In particular, it examines the big questions concerning the use of exemplars and means of generating arguments.

Discussion

Gardner is right on point in placing importance on such paradigms as open texture, hard/easy cases, and core/penumbra distinctions. I think they are central to legal reasoning and open up valuable linkages to other AI disciplines as well as to legal philosophy. For instance, the core/penumbra distinction of H. L. A. Hart has close ties to example-based reasoning, the use of prototypes, and machine learning.

Gardner's model gives some real computational flesh to the philosophical skeleton of such matters as the hard/easy distinction. Not only does it provide a model of how to make the distinction but it also allows the reader to examine and criticize the philosophical approaches. Philosophical discussions rarely provide enough detail for this to be done. Thus, Gardner's work is a prime example of how an AI model can enlighten and inform work

Gardner's work is a prime example of how an AI model can enlighten and inform work on central issues in the law.

on central issues in the law.

Regarding her model of how to distinguish hard and easy questions, I am concerned that with any sort of realistic case base, there will always be contrary, opposite cases and thus, under her heuristics (especially the third that looks for oppositely decided cases), many if not all predicate resolution questions will appear hard. Even in a legal area as stable as contracts, there is a diversity of cases, which is reflected in casebooks and the *Restatement* itself. As things now stand, her approach might raise too many false alarms. One way to remedy this is to recognize that not all (opposite) cases are created equal or better yet, that not all (contrary) arguments are equally robust. This could give rise to a refined heuristic to be applied to tentatively hard questions: When there are opposite cases, evaluate the cases or arguments and dismiss those that are weak; if one still has an oppositely decided case, the issue is hard. Of course, this requires reasoning with cases or, more difficult, reasoning about arguments. This is not easy but I think it is vital, and I think Gardner would agree. I fear that shying away from the more case-based and argument-based aspects of legal reasoning is to put more faith in and emphasis on rules and generalized patterns than is appropriate, even for an issue-spotter task.

Being more of a rule-skeptic, I feel that one can turn just about any case into a hard one, for instance by discrediting the applicability or validity of a rule. I see cases as playing a much more central role than Gardner. However, using cases in a more central way and not just as annotations or existential embodiments of concepts would require deep changes in Gardner's program.

I also feel that cases need to be represented in more detail than is possible with generalized fact patterns; while there certainly are stereotypical situations, not all are so. Such a generalized level of representation under-emphasizes important details of individual cases that are critical for indexing, analogy, and other aspects of case-based reasoning. It would not be hard for Gardner to use a more fact-based

level of representation for cases, since she already has developed a representation to handle this level of specificity in describing the facts of the problem, which can be thought of as a new case.

Gardner's book is beautifully written. I have used it, and prior to its publication the thesis form of it for three years in my seminar at the Harvard Law School on "AI and Legal Reasoning" and in my graduate computer science seminars at the University of Massachusetts. Both groups of students found it accessible. In particular, the introduction and the explication of legal philosophy (chapters 2 and 3) are absolutely first rate. They hit the highlights without drowning the reader in a jurisprudential bog. By writing such a useful survey and also developing an excellent bibliography, Gardner has done us all, both legal and computer science types alike, a great favor. Without trivializing, she manages to convey the central issues and history concerning rules and the various jurisprudential paradigms.

The main criticism from both law and computer science students is that the representation chapters (chapters 5 and 6) are hard going. I differ from my students a bit on this point. While the chapters are somewhat uneven and tell us more than we might have ever wanted to know, the content of these chapters is essential for anyone attempting to replicate Gardner's approach. I'd opt for even more detail in the form of listings of rules and general fact patterns, perhaps in an appendix. There is no question that chapter 7, which presents the extended example, is very important and lets the reader pull it all together.

In general, I wish that the text had been given a bit more structure (despite the abundance of section numbers) and that there were more schematics or figures. For instance, I wish more explicit enumeration were used inline in the text instead the phrases, "one way" and then "another way." This would help the reader avoid wondering what might have been missed because either it was not presented or not fully understood. I would have appreciated a computer-science-style presentation—perhaps a diagram or structured English—on the

hard/easy algorithm and the various details supporting it. While all the pieces are there, some minor editorial changes would make it easier for the reader to synthesize and keep track of the information. Such shortcomings would be easy to remedy in the next edition.

Gardner has done landmark research in the field of AI and legal reasoning. Notwithstanding our jurisprudential differences, I feel that Gardner has made a substantial contribution to AI and legal reasoning, which will have a major impact on both disciplines. Not only has she done groundbreaking work on certain topics such as issue-spotter questions, open-textured predicates, and the hard/easy paradigm, but she also points the way to further work on case-based reasoning and argumentation. This book is not simply a thesis bound between cloth covers. Rather it is the product of refinements by a researcher who has further mulled over and synthesized the results of the thesis. It displays a maturity of understanding seldom reached in a thesis itself. Her book is essential to anyone working in the field and invaluable to those who would like to do so. And that is an easy question.

The Emerging Field

The field of AI and legal reasoning is burgeoning. I think the field is at the beginning of a period of exciting growth, like AI and medicine about ten years ago. Over the past years, there has been a steady rise of interest and accomplishments and the pace is picking up.

In May, 1987, the first International Conference on AI and Law (ICAIL-1) was held. This conference drew approximately seventy paper submittals, out of which 30 were chosen for inclusion in the proceedings. The attendance surprised all the organizers; 175 attendees showed up, far in excess of the 60 to 90 expected, and a good fifth of these registered for tutorials. The attendees represented a broad cross-section which included academics and students from both disciplines, senior partners from well-known firms, and lawyers and experts from various government agencies.

The work presented at the conference covered many aspects of legal reasoning such as the following:

- Legal expert systems. Presentations included examples of useful applications of expert systems technology in such areas as tax and pension law, and discussion of the jurisprudential status of these applications.
- Legal retrieval. The discussion covered how to design conceptual knowledge bases using traditional methods, AI approaches, and even connectionist models.
- Modelling legal reasoning. Topics included specifying expert lawyer reasoning and case-based reasoning in areas as divergent as landlord-tenant and trade secrets law, using conceptual approaches from natural language understanding in case law, and developing ways of handling hard cases.
- Logic-based approaches. Conference participants discussed the continuing efforts to write statutes in structures with well-defined logical meanings and to formalize the deontic logic of permissions and obligations. They also described the use of logic programming for applications such as the drafting of statutes.

There were several presentations of work being undertaken as doctoral research in computer science that will provide future landmarks in the field. For instance, my colleague Kevin Ashley (1988), who is another J.D./Ph.D. researcher, has developed a detailed AI model of case-based reasoning and how to frame legal arguments, tasks high on Gardner's and my research agendas. Other students of mine are investigating how to combine case-based and rule-based reasoning, for instance, in statutory domains. A student of Michael Dyer is conducting doctoral research on understanding short natural language summaries of cases.

Plans for ICAIL-2 are now under development and the meeting is scheduled for 13-16 June 1989, in Vancouver, British Columbia, with Robert Franston and Joseph Smith as coconference chairs and myself as program chair. The call for papers is reproduced in the Announcements section of this issue of *AI Magazine*.

In the meantime, the conference committee of last May's meeting has

decided to establish an informal newsletter, to be edited by myself and distributed to the community once or twice a year, and AAAI has also formed a special interest group on AI and Law.

Conclusion

The field of AI and legal reasoning has reached critical mass and shows all the signs of being a source of of insightful research that should be of interest to both disciplines. From the point of view of a person whose primary allegiance is the law, such work should shed light on the workings of legal reasoning and on the validity of certain jurisprudential stances. From the point of view of AI, such work should provide advances in basic techniques and probing examples and test beds for existing methods.

Gardner's book is an important and thoughtful investigation of AI and legal reasoning. It is a landmark case of research that enlightens and informs both disciplines. If further work, and ultimately books, on this topic follow the precedent set by Gardner and others, this field is off to a vigorous start.

Selected Bibliography

The following is a sampling of both recent and classic papers on AI and legal reasoning. The Proceedings of the First International Conference on Artificial Intelligence and Law, at Northeastern University, Boston, in May 1987 in cooperation with ACM Sigart, provides an excellent collection of current research.

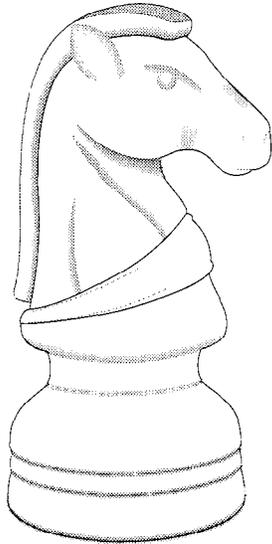
Allen, L. E., and Saxon, C. S. 1987. Some Problems in Designing Expert Systems to Aid Legal Reasoning. In Proceedings of the First International Conference on Artificial Intelligence and Law, 94-103.

Ashley, K. D. 1988. Modelling Legal Argument: Reasoning with Cases and Hypotheticals. Ph.D. diss., Dept. of Computer and Information Science, Univ. of Massachusetts.

Ashley, K. D., and Rissland, E. L. 1988. A Case-Based Approach to Modelling Legal Expertise. *IEEE Expert*, Fall 1988, 70-77.

Ashley, K. D., and Rissland, E. L. 1987. Compare and Contrast, A Test of Expertise. In Proceedings of the Sixth National Conference on Artificial Intelligence, pp. 273-278. Menlo Park, Calif.: American Association for Artificial Intelligence.

Have KEE® — Will Travel



RAPID PROTOTYPING of Robust, Extensible Knowledge-Based Systems

AMOS OSHRIN

**Knowledge Systems Consultant
4 Morning Sun Court • Mountain View, CA 94043**

(415) 969-9600

**KEE is a registered trademark
of IntelliCorp**

RESULTS, Not Just Reports

For free information, circle no. 145

Dworkin, R. 1977. *Taking Rights Seriously*. Cambridge, Mass.: Harvard University Press.

Dworkin, R. 1985. *A Matter of Principle*. Cambridge, Mass.: Harvard University Press.

Eisenberg, M. A., ed. 1982. *Gilbert Law Summaries: Contracts*, New York: Harcourt Brace Jovanovich.

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

Gilmore, G. 1974. *The Death of Contract*. Columbus, Ohio: Ohio State University Press.

Goldman, S. R.; Dyer, M. G.; and Flowers, M. 1987. Precedent-Based Legal Reasoning and Knowledge Acquisition in Contract Law: A Process Model. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, 210-221.

Hafner, C. D. 1987. Conceptual Organization of Case Law Knowledge Bases. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, 35-42.

Hart, H. L. A. 1961. *The Concept of Law*. Oxford: Clarendon Press.

Hart, H. L. A. 1983. *Essays in Jurisprudence and Philosophy*. Oxford: Clarendon.

Kuhn, T. S. 1970. *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.

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

Llewellyn, K. N. 1930. *The Bramble Bush*, Second Edition, 1960, Seventh Printing, 1981. Dobbs Ferry, N.Y.: Oceana Publications.

McCarty, L. T. 1985. Permissions and Obligations. In *Computing Power and Legal Reasoning*, ed. C. Walter, 573-594. St. Paul, Minn.: West Publishing.

McCarty, L. T., and Sridharan, N. S. 1982. A Computational Theory of Legal Argument, Technical Report, LRP-TR-13, Lab. for Computer Science Research, Rutgers Univ.

Meldman, J. A. 1977. A Structural Model for Computer-Aided Legal Analysis. *Rutgers Journal of Computers and Law* 6: 27-71.

Restatement of the Law, Second: Contracts, 3 vols. 1981. St. Paul, Minn.: Ameri-

can Law Institute.

Rissland, E. L., and Ashley, K. D. 1986. Hypotheticals as Heuristic Device. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 289-297. Menlo Park, Calif.: American Association for Artificial Intelligence.

Rissland, E. L. and Skalak, D. B. 1988. Case-Based Reasoning in a Rule-Governed Domain, Technical Report, CPTM 26, Dept. of Computer Science, Univ. Massachusetts, Amherst.

Schlobohm, D. A., and Waterman, D. A. 1987. Explanation for an Expert System that Performs Estate Planning. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, 18-27.

Shepard's Federal Citations 1988. Colorado Springs, Colorado: McGraw-Hill.

Sergot, M. J.; Sadri, F.; Kowalski, R. A.; Kriwaczek, F.; Hammond, P.; and Cory, H. T. 1986. The British Nationality Act as a Logic Program. *Communications of the ACM* 29: 370-383.

Waterman, D. A., and Peterson, M. 1981. Models of Legal Decisionmaking, Technical Report, R-2717-1CJ, Rand Corporation.