Book Reviews

An Artificial Intelligence Approach to Legal Reasoning

Thomas C. Arcidiacono

An Artificial Intelligence Approach to Legal Reasoning by Anne von der Lieth Gardner (Cambridge, Massachusetts: The MIT Press, 1987, 225 pages, $22.50), is a revision of Anne von der Lieth Gardner's dissertation at Stanford, and is the first in a planned series examining aspects of artificial intelligence and legal reasoning. As both a computer scientist and a lawyer, Gardner understands the importance of participation by scholars from both fields in future research. Her work is directed at two groups of readers: those with technical knowledge of AI programming techniques, and those trained in law.

The first part of the book is intended to be an introduction to computational jurisprudence for both groups. It identifies issues critical to the purpose, behavior, knowledge sources, knowledge structures, and reasoning processes of expert legal systems. The second part implements a simple prototype system for a well-defined area of contract law and is more appropriate for experienced developers of knowledge-based systems.

Law is a domain in which the experts are supposed to disagree, and lawyers must be able to argue either side of a case. A judge or juror must decide which argument is "best." A knowledge based legal reasoning program can only guide analysis and identification of technically defensible positions in a case. However, it should also be able to distinguish between questions that are "easy" to decide, and those demanding human analysis. These two ideas form the basis of the prototype's behavior, making it somewhat different from knowledge based systems in most other expert domains.

According to Gardner, legal reasoning systems are further distinguished by their knowledge sources and knowledge structures. She reviews the evolution of legal thought in the context of knowledge engineering, raises several critical issues, and draws conclusions about how legal knowledge must be used and represented in programs. In the human world, legal knowledge is represented in cases, and statutes. Although not all areas of law use both sources, she concludes that expert legal systems need both types of knowledge, plus some additional "common sense knowledge" to guide analysis effectively.

Gardner views statutes as rules defining legal states and their consequences. Although they are convenient starting points for legal analysis, they are usually insufficient for making wise legal decisions. Most litigation involves questions about whether the rules have been followed, what the rules actually mean, and sometimes, which set of rules should be used.

Cases contain written arguments about how to answer these questions under specific circumstances, along with their final interpretation by the juror. Lawyers can use similar cases as examples to guide their formulation of arguments in future disputes. Cases are used as precedents for deciding which rules to use in a given situation, and how to apply them. They can be used to annotate and clarify rules that conflict in some context, or whose relevance might be disputed. They can even change the way rules are applied to similar factual situations in the future.

In these respects, cases embody much more legal knowledge than statutes. Most of their representational power results directly from inherent ambiguities in natural language text. This problem of "open texture" causes questions that are too hard for legal expert systems to answer.

Gardner concludes that rules can be used to make legal decisions when cases are available as examples of how to clarify, modify, and extend the rule base. This approach is taken by her prototype.

Abundant legal knowledge already exists in the form of cases, but this knowledge is beyond the ability of present natural language processing programs to interpret. Gardner's approach has been to let humans identify the salient facts or predicates provided by a case, and then encode them in formal structures by hand. A serious loss and perhaps contamination of information can result, which is considered unavoidable due to the scope of the study.

Perhaps to avoid a similar problem with the translation of statutes to rules, the author has selected contract law for the study domain because it is already described using rules. Although it evolved through case decisions (as have many other areas of law), it was translated into rules by the legal profession to assist human jurors. Although it is a good choice to illustrate the author's ideas about the ways cases and rules can be used, it has some real limitations in areas of law based primarily on cases. Construction of a similar prototype for those domains would require either a translation of "open text" statutes to a format better suited to rules, or the generation of rules entirely from existing cases. Neither process is currently practical according to the author, so the reader must consequently wonder what other alternative uses, sources, or transformation techniques if any, will be needed.

Also of interest is the idea that an expert legal system must be able to recognize deductive limitations of specific rules for specific problems. While reasoning about a problem, the system must decide which rules, if any, it can apply. If it cannot do so with simple predicate matching, the facts of one or more cases can be used as examples to support or refute the use of the rule. If cases are found, deduction can continue, possibly

along multiple alternative paths. Otherwise the rules could be applied in a way that is inconsistent with their intended purpose. The system must recognize when a rule is insufficient, and how to select appropriate examples from the body of cases. Both tasks may involve more sophisticated reasoning and searching techniques. Gardner discusses some of these techniques later in the context of her prototype but does not implement them.

The book contains an excellent critique of previous computational approaches to legal reasoning. It chronologically reviews the evolution of different techniques, their results, and their limitations. Failure to achieve significant advances was in most cases attributed to insufficient knowledge. This review might have been more helpful as part of the introduction to computational legal reasoning because it clearly defines the field's origin and current status. The discussion helps put things in perspective for developers seeking to avoid earlier mistakes, and for legal scholars trying to get a handle on the direction of computational research. A review of early design methodology could also have given the reader a better understanding of the author's motivations, and helped to clarify the discussion of legal reasoning paradigms.

The prototype's purpose was to provide an example of how the author's design choices could be applied to a realistic legal reasoning domain. Analysis of contract "offer and acceptance" problems by first year law students was the domain chosen, for reasons that included the availability of a well-defined legal rule base, cases for analysis, and a minimum of "common sense" knowledge to identify and resolve predicates. Given the facts of a dispute, the program attempts to identify "easy" decisions, otherwise analyzing possible arguments and representing them with decision trees. It also tries to identify aspects of the arguments that are "hard."

The program uses an ATN to represent legal states such as "offer" and "acceptance." Several different types of rules are used to move between states. Rule predicates are satisfied by applying "common sense knowledge" and MRS structured case facts. Solutions are obtained by pruning a space of possible paths through the network. Description of the program implementation is highly technical and provides enough detail for a reader to functionally duplicate the system. Operation of the system is traced through a test problem to show each important type of activity.

The choice of a formalism with which to represent knowledge and problems in the prototype had an affect on the implementation of inferencing and control. MRS was selected to temporally represent events, states, objects, and physical measurements. Conceptual differences between these types of information were discussed but the reasons for using MRS to represent them were unclear. The author even points out that limitations in the ability of MRS to represent inheritence and generalization increases the system's dependence on explicit rules. Some justification is given indirectly in subsequent discussions of implementation detail.

One of the few problems with the book is that it was written as a doctoral thesis and tends to wade through detail without providing the kind of overview that some readers might need. Perhaps less time should have been spent on design methodology in a book with so potentially diverse an audience. As an alternative, each chapter could have been partitioned into a general overview of the material, followed by a more detailed treatment. This would have allowed readers to focus on areas of particular interest without interrupting the progression of ideas.

Throughout the discussion of legal theory there should have been clearer explanations of some basic legal terms used. There also should have been more diagrams illustrating the dynamics of system-knowledge structure interaction. There are enough figures showing structural relationships between different types of program objects, and showing examples of input and output.

Certain problems are not addressed. For example, "hard questions" can often have nothing to do with open texture. Decisions can involve philosophical, political, emotional, or intangible factors entirely unrelated to either statutes or cases. There was no analysis of how these might be included in the model, nor was there an adequate discussion of how conflicting case decisions might be analyzed in light of them. In addition, the explanation of "common sense" knowledge seemed too vague for generalities to be made about its availability or importance within a model of legal knowledge.

The text fails to emphasise that not all law involves disputes. In fact, it could be argued that a primary objective of law is the avoidance of dispute. There should have been more speculation about how the features of the prototype could be applied to the problem of "legality advisement." For example, given a context, what should be done to result in a particular legal situtation? If such a system is used to examine possible legal pitfalls, would there be additional representation difficulties in order to predict their affect on related legal situtations? Could a similar system use a "critic," or a learning process to find grounds for appeal after an argument fails? Some of these questions are beyond the scope of the book; hopefully they will be addressed by forthcoming works in the series.

Many of the major issues for developing computational law seem to have been identified by this work, if not analyzed in detail. The prototype reflects this analysis, and serves as an effective instructional tool for the reader. Aside from some minor limitations, the book is an exceptionally good and thought provoking introduction to a potentially complex interdisciplinary field.

The book is valuable to AI researchers, to legal scholars, and to anyone interested in computational legal reasoning. The text explicitly asks or implicitly raises many questions for further study and attempts to answer the most fundamental of them. Researchers seeking design or implementation resources will find the work particularly useful, and an extensive bibliography supplements the literature review. The series editors have achieved their objective by providing a foundation on which computational studies of legal reasoning can build.