The Knowledge Engineering Review

Critical note—artificial intelligence and legal reasoning*

RICHARD E. SUSSKIND

Ernst & Whinney, Becket House, 1 Lambeth Palace Road, London SE1 7EU, UK

I

Solving a legal problem is a complex business. First, the lawyer—or whoever is doing the solving—has to assemble and marshall the facts of his case in as orderly a fashion as his knowledge of the relevant law will allow. Secondly, he must confront the range of pertinent legal sources and decompose these into manageable and applicable units. Thirdly, he will subsume the particular facts of his case under the more general terms of his decomposed units. And, finally, he draws inferences about the applicable legal provisions and their impact on the facts. Of course, assembly, decomposition, subsumption and inference may not occur in linear progression. What is more, problems can occur at each stage and interim evaluations and judgments must be made.

How do human beings manage this? Can computers assist them at any or all of the stages identified? These are two crucial questions for workers in the field of artificial intelligence and legal reasoning. The field itself can be said to date from 1970, when Bruce Buchanan, one of the fathers of expert systems, together with Thomas Headrick, a legal colleague, wrote a paper in the *Stanford University Law Review*. That article, "Some Speculation About Artificial Intelligence and Legal Reasoning", set the scene and laid down the challenge.

Pre-eminent in responding was L. Thorne McCarty (then at Stanford but now at Rutgers University), whose contribution to the field since then has not been paralleled. He placed the next major landmark in the published literature in 1977 in the prestigious *Harvard Law Review*. Since then, and largely inspired by his contribution, around 30 sustained research projects have been conducted in the field. While commercially viable expert systems in law remain exceptionally thin on the ground, the first book on the more general area of AI and law has emerged. Anne von der Leith Gardner's *An Artificial Intelligence Approach to Legal Reasoning* is that book. It is the first in a projected series of texts on artificial intelligence and legal reasoning and is under the joint editorship of the aforementioned McCarty together with Edwina Rissland (who has herself worked extensively on AI and analogical legal reasoning). Gardner's book is a revised version of her doctoral thesis submitted at Stanford University in 1984. The main body of text extends to just under 200 pages, divided into eight chapters, and throughout there are peppered clear and well produced illustrations. At the end, there is a comprehensive and useful bibliography, followed by a detailed index.

Dr Gardner was well equipped to write such a book. She holds a doctorate in computer science as well as a law degree (both from Stanford University). Capitalizing on her interdisciplinary background, she states that the object of her book is "to create a model of the legal reasoning process that makes sense from both jurisprudential and AI perspectives" (p. 1). The non-lawyer may not be familiar with the subject of jurisprudence. Broadly speaking, jurisprudence is legal theory or the philosophy of law. Its exponents strive for systematic and general understanding both of the law itself and of its administration in society. Amongst many other issues, jurisprudential debate focusses on the nature of law, the process of legal reasoning, the structure of legal rules and

* Book review. Anne von der Leith Gardner, *An Artificial Intelligence Approach to Legal Reasoning*, A Bradford Book, The MIT Press, London, 1987.

systems, the relationship of logic to the law, and the interpretation of statute and case law. So any computer system that seeks to solve legal problems—to engage in legal reasoning—necessarily makes a whole set of jurisprudential assumptions: it will presuppose some account (often no more than implicitly) of the nature of law and of legal reasoning. Gardner, then, is absolutely correct in perceiving the relevance of jurisprudence for AI and the law. Her enduring interest throughout the book is to enable a computer system not simply to solve legal problems but to "recognize the issues a problem raises and to distinguish between those it has enough information to resolve and those on which competent human judgments might differ" (p. 4). In adopting this perspective, Gardner ambitiously launches herself into the heart of legal philosophy. She sustains her commitment to the relevance of jurisprudence throughout her study but, for reasons I shall suggest shortly, I fear that seasoned legal philosophers will find her theorizing a little thin in places.

To demonstrate her findings, the author concentrates on one particular branch of the law: the formation of contracts by offer and acceptance. This makes her job all the more difficult as this legal domain is dominated not by authoritatively articulated legislation but by a rich body of judicial precedents. Gardner is not typical in operating on case law, but she is right in recognizing that it is unrealistic to focus on statute at the expense of precedent. There is a lesson here for all workers in the field.

For the AI scientist, the meat of Gardner's book is to be found in Chapters 5, 6 and 7. There she describes how the implementation of her legal reasoning program actually works. The system itself was written in Maclisp, augmented by MRS's database and knowledge base facilities, and runs on a DECSYSTEM-20. The computational model of legal reasoning she develops required: (1) a method of representing the facts of a legal case (Chapter 5); (2) a representation formalism for the expression of legal knowledge (Chapter 6); and an inference mechanism for applying the knowledge held within the system to any problem data presented to it (also Chapter 6). Chapter 7 brings it all together by showing, through examples, how the system addresses concrete legal problems.

II

If one chooses to design a legal reasoning program that requires the facts of cases entered in some formal or semi-formal representation, then, as Gardner's book illustrates clearly, the order is a tall one. The language in question must accommodate actions, events, states, objects, substances, speech acts and conversation, as well as the more general notions of time and space. Using ideas associated with natural language processing and frame-based representation, and sensitive to the idiosyncrasies of legal reasoning, Gardner tackled the problem in two stages (p. 89). In the first instance, by using a simple propositional, KRL-type syntax, she was able easily to identify what basic notions would need to be described. Thereafter, she adopted a richer and more expressive language—predicate logic-based MRS.

Not being frame-based, MRS lacks the features of slots, fillers and hierarchies that Gardner argues are required in this context. This she overcomes—essentially thereby reproducing a frame-based formalism—by introducing CSK rules: her system's very rough approximation to common sense (non-legal) knowledge. These rules allow for knowledge about the objects and relations expressed in MRS, and support a selection of inferences. For instance, CSK rules provide for movement through a hierarchy of terms ranging from the very general (such as "event") to the very particular (such as "beginning of a transfer by mail"). In sum, then, Gardner concocts a sophisticated and expressive quasi-frame-based language in which to represent problems. Yet sophistication can be a double-edged sword, as is evidenced by her daunting formalization of a "sale of a carload of salt" (pp. 99, 112 and 114).

The program is designed to process at four different levels, associated with each of which is a different knowledge representation formalism. To date, only the top three have been implemented. The fourth will hold knowledge sufficient to solve "questions that lawyers could reasonably disagree on, [and] will require a great deal of further research" (p. 119). With this very last remark, I am in full agreement. The tripartite division itself runs as follows. First, there are the fundamental legal

states and events, such as "a contract exists" and "offer" respectively. This knowledge is held in an "augmented transition network" (the most widely used technique for processing natural language). The second level of knowledge is definitional in nature, expressed as very complex "if-then rules", such as the one laying out the conditions under which an offer can be said to have been made (p. 142). The third knowledge type, residing below the network and rule levels, is "knowledge of examples of the kinds of things that are and are not covered by the predicates used in the rules" (p. 155). Gardner has in mind here knowledge needed when legal rules alone do not cover any particular case.

To understand why four levels of processing are needed, a departure to the world of jurisprudence is required. Gardner's jurisprudential analysis is held in Chapters 2 and 3. Although reasoning with rules is central to legal problem-solving, the resolution of some cases requires both reflection on and response to principles (such as "no man may profit from his own wrong") as well as regard to the underlying purpose of enactments and decisions. Moreover, doubt over legal rules can result from the "open texture" of their language. If "vehicles" are prohibited in the park, for instance, we may be confident that the intrusion of a Porsche would be a transgression but less sure of the position regarding planes or skateboards. Further difficulties that arise from legal rules are their tendencies to conflict or to suffer from gaps or to change through time.

Not all legal problems, however, are fettered by considerations of the type just noted and in reflection of this the jurisprudential literature speaks of the distinction between "clear" and "hard" cases, a division that Gardner adopts and refers to widely. (Note, in passing, this distinction is one in which jurisprudents have no grounds for feeling pride. It is commonly referred to, but is inadequate in several respects; chief of which is its neglect of the prospective as well as the retrospective senses in which cases can meaningfully be said to be clear or hard.) For Gardner, a case is hard when experts themselves genuinely disagree over the legal position, whereas clear or "easy" cases are those in which there is no such disagreement. The problem then identified is that there is no clear dividing line between such cases as just classified. Should problem-solvers when confronted with legal cases assume every case to be *prima facie* hard or clear? Legal practice suggests the latter, while legal theory implies the former. For AI purposes, Gardner's key question is how systems can recognize when they are confronted with a problem which they lack the knowledge to solve. To tackle this, Gardner articulates heuristics that determine whether cases are easy or hard and therefore can or cannot be solved by the system.

Returning to Gardner's four proposed levels of processing and remembering that the fourth has not been implemented, the system's capabilities can be regarded in two ways. On the one hand, the program's three levels seek to identify the hard questions in any cases brought before it; while on the other hand, the three levels pinpoint the easy questions and answer them. This reminds Gardner of one major skill expected of law students—"issue spotting".

The first two levels of processing are chiefly deductive in nature. Level 1 processes the events relating to a legal problem through the augmented transition network mentioned earlier. This process seems to be essentially classificatory in nature and in very straightforward cases may, on the basis of its characterization, yield a conclusion: for example, that there is a contract (crudely) because X made an offer and Y accepted it. In less straightforward cases, classification cannot easily be made, in which event level 2 applies its definitional rules to recognize instances of the more general categories held in the network, that is, to determine if a given arc of the network can be followed. At this level also, conflict of rules is tackled, in that competing rules dictating diverging conclusions are treated as giving rise to a hard question.

Level 3 is there to handle problems in respect of which rules have "run out". For Gardner, the problem here is really that of the "open texture" of legal rules—because rules are couched in natural language and the precise scope of their general terms is inherently uncertain, their application in some cases is clear, but in others decidedly not so. Open texture, in the use of Gardner's system, gives rise to most hard cases. Of course, the phenomenon of open texture does not render all cases hard cases, and Gardner allows for this via level 3.

Ultimately, the human expertise that Gardner tries to encapsulate in heuristic form is the ability

to assess, in advance (in a sense), whether a case does indeed raise easy or hard questions.For this will indicate if a conclusion may be inferred without further ado or whether human judgment is required. She makes a brave attempt, but even after reading her illustrative examples in Chapter 7, I was left wondering about the generalizability of her ideas and their applicability in other, far more extensive, branches of law. It would be nice to see a working system.

III

If adverse criticism were to be made of Gardner's book, this could be done in a jurisprudential dimension. Although she adeptly identifies some of the jurisprudential issues in this field, her usage of legal theory as a source of guidance is a touch haphazard. While wholeheartedly endorsing her idea of deploying jurisprudence, I have misgivings over her method. Given the enormity of the jurisprudential literature, it is surely methodologically appropriate that scholars who rely heavily on findings of legal theory clarify on what basis, and in accordance with what criteria, they choose their philosophical source materials. Remember that jurisprudence is best used to develop AI models and *not* to justify retrospectively pre-philosophical ideas and thoughts. Unfortunately, Gardner tells us nothing of how or why she selected the jurisprudential writings upon which she relies.

In legal theory, there are a great many diverse schools of thought. She rarely cites those that are at variance with the position she adopts, a posture, it should be added, that itself by no means commands universal support. She simply does not mention the works of Bentham, Kelsen and Raz, unquestionably three of the greatest legal philosophers of all time. All breathtakingly prolific men, in each of their writings can be found arguments that challenge some of Gardner's most basic assumptions. She does not confront their objections. Nor does she refer to several modern theorists—MacCormick, Wasserstrom, and Gottlieb—who have philosophized soundly about precisely the issues of logic and law that Gardner addresses. From a jurisprudential perspective these are significant omissions. Her jurisprudence is confined to a few, major theorists with whom all students are familiar. Scholars of jurisprudence would ask for more.

The legal theorist would also want to know why Gardner confines her jurisprudential attention to inference and the use of knowledge yet does not bring legal philosophy to bear on the question of the representation of legal knowledge. The jurisprudential issues of rule structure, legal systems, individuation, legal science and legal epistemology, for instance, are as relevant for AI and law as are the topics of legal logic and clear and hard cases. Again, there is no explanation for this discrepancy. One major shortcoming of the book, then, is the uncertainty surrounding Gardner's selection—and neglect—of works and issues of jurisprudence.

Leaving jurisprudence behind, where does Gardner's research fit into the more particular field of AI and law? In the world of artificial intelligence and legal reasoning, I would like to suggest that there are two major schools of thought. These may be called the "pragmatist" and "purist" schools. The overriding aim of pragmatists is to develop and implement working computer systems that can actually assist in the solving of legal problems. For purists, in contrast, completion even of a modest prototype is not always necessary for success, for their major goals are to clarify the nature of legal reasoning, of human and artificial intelligence, and of computational models of law: live systems are not necessary for this. On the whole, pragmatists tend to be found in the commercial world, while purists inhabit research establishments.

Not all pragmatists are qualified to build expert systems in law: some software developers have made attempts without any assistance from domain experts. While many purists are well qualified for the job, commercial exploitation of the technology is not the game they are in. The diverging orientation of the schools is further reflected in their reactions to the label of their field: pragmatists prefer not to be associated with "artificial intelligence", whereas this designation for the purists concisely encapsulates their aspirations.

The difference in tools used by pragmatists and purists is yet more evidence of the divide. The former opt for low-cost, "rule-based shells" employing classical logics and running on personal computers, while the latter invariably deploy vast-cost, AI "toolkits" or "frame-based" languages,

using extended logics (such as "deontic" logic), operating on mainframes or dedicated AI workstations.

Most workers in the field, of course, are not simply pragmatists or purists but are subtle combinations of both. The terms purism and pragmatism, in sociologists' terminology, are "ideal types". Inevitably, however, scholars incline more towards one camp than the other. Gardner is fairly clearly a purist and would benefit from regard to pragmatist teachings. Nowhere is this more apparent than in her approach to the user interface.

In using Gardner's system, the user is required to enter the facts of the case in the formal representation language that was discussed earlier. Without further human intervention, the system can then differentiate between those cases it can solve and those it cannot, between clear and hard cases. Therein lies its intelligence. The hitch is that any user of a system such as Gardner's must be thoroughly conversant with the quasi-computer language by means of which the details of any legal problem are imparted. In this, those who are assessing the commercial viability of AI systems for the law will find a problem. For it would be unrealistic to expect all potential users of such systems to immerse themselves in the niceties of processable grammars. Yet this is not a worry for the purist such as Gardner for whom sale and distribution of systems is not the prime concern. Her foci of interest are natural language processing, common sense reasoning, artificial intelligence, and the nature of legal problem-solving. There is the danger, however, that work in these spheres is too far dissociated from potentially useful systems.

The pragmatist would approach the question of the interface in a quite different manner. Ever-sensitive to market requirements, he would want any expert system in law to be very easy to use. To that end, he would design his system so that the facts of any case are elicited simply by asking the user to answer questions, input data and select from menus, all in a manner requiring of the user no more than primitive keyboard skills. Gardner has some robust purist objections to this (pp. 60–61), but the pragmatist posture is defensible in theoretical as well as practical terms. For the pragmatist can claim that his system simply requires more intelligent users. And these users can bridge the intelligence gap between pragmatist and purist systems. It may well be that this means users of pragmatist systems will need to be lawyers or legally informed persons: but such systems—in offering assistance to knowledgeable individuals who lack highly specific knowledge within a field with which they have broad familiarity—are undoubtedly commercially significant. The resultant systems can solve "clear cases of the expert domain"—cases that are easy for experts but usually hopelessly difficult for non-experts.

Over and above such pragmatist offerings, however, Gardner's type of work on the interface must go on: from continuation of such research will emerge successively more intelligent systems. There is room in the world of AI and law for both pragmatists and purists. Indeed, it is essential that contributions emerge from both camps. Equally crucial is that neither claims superiority over the other. Above all perhaps, collaboration between the two is desirable.

IV

In conclusion, *An Artificial Intelligence Approach to Legal Reasoning* is, I believe, a contribution predominantly to the world of computer science. It is, of course, essential reading for all workers in the specific field of artificial intelligence and legal reasoning, not least for the clear overview of some related projects in Chapter 4. It is also a valuable introduction to the legal domain for AI scientists involved with natural language processing, and, to a lesser extent, expert systems. For those of legal orientation, however, the book will be of less interest: legal theorists will find it does not contribute to but draws (a little arbitrarily from) jurisprudence while practising lawyers will be disappointed (or perhaps delighted!) by what seems to emerge as the remote possibility of AI systems in law. Commercially viable systems can be built today, but more along the lines of what I have termed the approach of the "pragmatist" school rather than that of the "purist" camp.

Gardner's book is at once well written, stimulating and yet too ambitious. It will require a polymath of some distinction to write a text on AI and law that will fully satisfy both the legal and

the computational fraternities. The problem here is rooted not simply in the amount of data, information and knowledge that would need to be at the fingertips of such a scholar. There is the further question of *perspective*: the history, styles, and epistemological bases of writings on law and computing are diverse. This raises the question of whether academic offerings in the field should not best be made in the tradition of one or other but not both disciplines. This would not be to deny the incontrovertible worth of collaborative, interdisciplinary work, but to acknowledge that you cannot please all people all of the time.