Social Aspects of Computing

Rob Kling Editor

The Potential of Artificial Intelligence to Help Solve the Crisis in Our Legal System

Donald H. Berman and Carole D. Hafner

> The laws that govern affluent clients and large institutions are numerous, intricate and applied by highly sophisticated practitioners. In this section of society, rules proliferate, lawsuits abound, and the cost of legal services grows much faster than the cost of living. For the bulk of the population, however, the situation is very different. Access to the courts may be open in principle. In practice, however, most people find their legal rights severely compromised by the cost of legal services, the baffling complications of existing rules and procedures, and the long, frustrating delays involved in bringing proceedings to a conclusion . . . There is far too much law for those who can afford it and far too little for those who cannot. No one can be satisfied with this state of affairs.
>
> Derek Bok [5]

The American legal system[1] is widely viewed as being in a state of crisis, plagued by excessive costs, long delays, and inconsistency leading to a growing lack of public confidence. One reason for this is the vast amount of information that must be collected and integrated in order for the legal system to function properly. In many traditional areas of law, evolving legal doctrines have led to uncertainty and increased litigation at a high cost to both individuals and society. And in discretionary areas such as sentencing, alimony awards, and welfare administration, evidence has shown a high degree of inconsistency in legal decision making, leading to public dissatisfaction and a growing demand for "determinate" rules.

In this article, we consider the potential of artificial intelligence to contribute to a more fair and efficient legal system. First, using the example of a middle income home buyer who was misled by the statements of a real estate broker, we show how a predictive expert system could help each side assess its legal position. If expert systems were reasonably accurate predictors, some disputes could be voluntarily settled that are now resolved by costly litigation, and many others could be settled more quickly. We then consider the process of discretionary decision making, using the example of a judge sentencing a criminal. We describe how diagnostic expert systems developed in the medical domain could be adapted to criminal sentencing, and describe a process by which this technology could be used—first to build a consensus on sentencing norms, and then to make those norms accessible.

In the ideal case, legal decisions are made after lengthy study and debate, recorded in published justifications, and later scrutinized in depth by other legal experts. In contrast to this ideal, most day-to-day legal decisions are made by municipal and state court judges, police officers, prosecuting attorneys, insurance claims adjusters, welfare administrators, social workers, and lawyers advising their clients on whether to settle or litigate. These decisions must often be made under severe pressures of limited time, money, and information. Expert systems can provide decision makers with tools to better understand, evaluate and disseminate their decisions. At the same time, it is important to reiterate that expert systems should not and cannot replace human judgement in the legal decision making process.

PREDICTIVE LEGAL EXPERT SYSTEMS

To assess the potential of artificial intelligence (AI) in helping to streamline the dispute resolution process, we analyze the hypothetical case of an unfortunate middle income home buyer who responds to a real estate broker's advertisement. The buyer complains to his attorney that during the broker's showing of the house the buyer asked whether "the house had water problems" and the broker replied, "Not to my knowledge." Several weeks after acquiring the house for $150,000 a heavy rain storm resulted in significant water accumulation in the basement. The buyer discovers that the flooding occurred because the land on which his house is located has serious drainage problems. To prevent a reoc-

[1] Including legislatures, government agencies, the courts, private law firms, and the criminal justice system.

currence will cost over $20,000. After discovering that the seller had moved away, the buyer called the broker who said that since "he knew nothing about the potential water problems he felt no obligation to contribute to the cost of the repairs." After an hour's consultation with his attorney the buyer learns that his case is "complicated," the chances of obtaining redress "uncertain" and the cost of pursuing the claim will be at least $5,000. In a neighboring law office the broker hears from his attorney that the "law is in a state of flux." Therefore, even though the broker's statement was truthful, the buyer stands "a chance" and that to defend a law suit may cost "about $5,000."

Though many citizens have difficulty understanding why these parties cannot obtain a clear answer to what appears to be a relatively simple legal question, members of the legal community well understand that three potentiating factors have greatly complicated our dispute resolution processes. First, many new areas of law have emerged as society called upon its legal institutions to protect the environment, the rights of consumers, voters, minorities, the indigent, and criminals. Fifty years ago an attorney could have given this buyer a rather straightforward answer—no legal relief unless the broker lied. Today, the skilled lawyer knows that consumer protection statutes and evolving judicial doctrines might provide the basis of recovery against an honest, but uninformed broker.

Secondly, framers of legal rules have often abandoned clear directives in favor of open textured rules. Under traditional common law fraud principles the broker would not have been liable unless he knew his statement was false. Under a more modern view the broker is required "to conduct a reasonably competent and diligent inspection of the residence." [9] Litigation to determine what the broker "knew" is generally less complex than judicial attempts to fix standards for "a reasonably competent and diligent inspection" and then to decide whether the broker satisfied these amorphous norms of professional behavior.

Thirdly, procedural rules were changed from a system where the pleadings of the parties narrowly framed the issues for trial to a system where lengthy and costly pre-trial discovery is required before the contours of the litigation become known [32]. Under the earlier procedural system that relied heavily on pleadings the buyer would have alleged what transpired in considerable detail and then a trial would take place to see who was telling the truth. Today, the pleadings do not apprise the parties of the nature of the case and, therefore, lawyers must examine witnesses and documents before trial in order to learn what the case is about. In this particular case the lawyers, in addition to asking witnesses what transpired, might conduct a pre-trial examination of expert witnesses to obtain their opinions on the kind of inspection this broker should have conducted. For all practical purposes litigants now must pay for two lengthy trials, the first to find out what the witnesses will say and the second to hear them repeat it. Granted, such a system minimizes surprise but at a significant increase in cost.

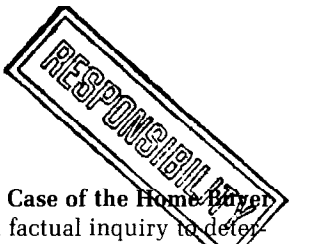

A Lawyer's Approach to the Case of the Home Buyer

Lawyers initially start with a factual inquiry to determine which broad area of law applies to the problem before them. After a few minutes the lawyer learns whether the problem involves a criminal or civil matter and shortly learns that he or she will need to apply legal principles applicable to the sale of houses. And in a case such as this the lawyer, relying on training and experience, would seek rather precise descriptions of what the broker said and what the client observed.

Most lawyers understand the basic principles of fraud—had the broker *expressly* and *falsely* represented that the property had no water problems then the plaintiff's chances of recovery are great. Conversely, had the broker said nothing and the buyer had seen water in the basement then his chances of victory would be nil because the law is unlikely to protect a buyer who failed totally to use his common sense. But once having sailed into the right ocean the lawyer must draw on two kinds of specialized knowledge to navigate into the narrow channels of the inner harbor. First, he must determine what published legal rules will apply to the precise facts he discovers. Secondly, after having assembled the facts and the relevant legal principles he must call upon his experience in determining whether a court, after having considered these facts and legal principles, will find in favor of his client.

As the lawyer hears the facts from the client certain legal principles are called to mind and as the lawyer delves further into the law he then becomes aware of the relevance of certain facts. In our simple example, the lawyer might initially conclude that the doctrine of *caveat emptor* ("let the buyer beware") bars recovery. But as he studies the cases he will find that the buyer might recover if the broker knew of potential water problems in the basement. Then, as the lawyer gains more factual information, he might learn that the client noticed a water line, a fact that might invoke a legal rule that would exculpate the broker since the existence of a water line should have alerted the buyer to the potential flooding. But since legal rules change, the buyer's lawyer might anticipate a court altering the legal rules to hold that since brokers are experts in real estate they have a duty to explain the meaning of the water line to a customer who lacks this kind of specialized knowledge.

The process of *finding the law* (including statutes, prior court cases, administrative rulings and procedural requirements) may involve searching a database of millions of potentially relevant documents.[2] The proliferation of legal documents is a major cause of the growing cost of legal services noted by Harvard President (and former law professor) Derek Bok. Furthermore, as the chance of missing a relevant document increases, the legal status of a case becomes increasingly difficult to

[2] The database for LEXIS, the online legal retrieval system, currently contains more than 10 million documents, and is growing at the rate of 2.5 million documents per year.

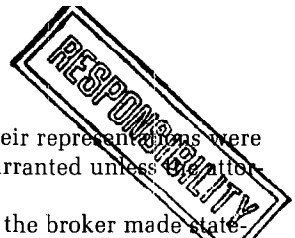

determine. Although still in its infancy, there is growing interest in using artificial intelligence techniques to develop conceptual legal retrieval systems [15, 16] to assist the lawyer in more effectively utilizing this vast database.

Once having discovered the facts and found the relevant sources of law, the lawyer must predict the outcome. Unlike an engineer whose knowledge of rules and facts engenders a confident prediction that a bridge will stand, lawyers operate within an indeterminate system. Fact indeterminacy always clouds the advice one gives to clients—a jury may believe either the buyer or the broker. Or a jury may believe the broker but find for the plaintiff because they sympathize with his plight [12].

Nor can attorneys ever be sure that a court will decide that brokers who either speak truthfully or do not speak at all are liable for defects they might have discovered had they made a diligent search, because accepted modes of legal reasoning support either side of this controversy. No court would be deemed unreasonable to insist that brokers only be held liable for their intentional misrepresentations. Conversely, a court might respond to plausible arguments that the broker should be held liable for innocent but careless representations. First, his expertise enables him to spot potential defects that would elude the eye of the inexperienced purchaser. Secondly, as businessmen brokers can better absorb the losses of the few individual buyers who sustain injury by spreading these losses evenly over all their customers.

While operating within an indeterminate system, lawyers nonetheless must regularly forecast the results of judicial proceedings. If the broker's attorney offered to settle the case for $10,000, the buyer's counsel would have to assess that offer. Few empirical studies have examined the way lawyers actually decide whether to accept settlements or proceed to litigation. Yet one can surmise that in our case of the home buyer a competent and ethical lawyer, presented with the offer described, would consider the following factors:

1. **Costs of litigation.** If litigation costs were equal to or greater than $10,000 settlement is clearly indicated given damages of only $20,000.
2. **Amount of damages.** If it appears that the plaintiff may have difficulty in proving damages in excess of $10,000 then settlement is clearly indicated.
3. **Disbelief of client.** Juries might not believe the buyer. After taking the depositions of the parties, attorneys predict how juries will judge the credibility of their clients. In making this prediction they will consider the demeanor of the witnesses. And given no clear indications that either party is lying, the attorney might conclude that a juror, moved by sympathy, would more readily accept the word of a first time home buyer than accept the credibility of a broker working for a large national realty company.
4. **Application of an unfavorable rule of law.** If, after researching the law, the attorney finds that in his jurisdiction the court presently imposes liability on only brokers who knew their representations were false then settlement is warranted unless the attorney believes that
   a) a jury will believe that the broker made statements about the condition of the property AND that he knew those statements to be false, or
   b) the court will change the law to hold brokers liable for false statements honestly made or will require brokers to make a diligent inspection of the premises.
5. **Skill of the Attorney.** Some lawyers have reputations for obtaining optimum results for their clients while others may be known to be less skillful. Lawyers confident of their litigating prowess or contemptuous of the abilities of opposing counsel often reject settlements that more timid counsel would readily accept.
6. **Judge and Jury.** Experienced lawyers realize that certain judges and juries from certain localities are prone to favor a certain class of plaintiffs or defendants.

The Role of Legal Expert Systems

Thirty years ago Layman Allen laid much of foundation for contemporary attempts to represent legal knowledge when he proposed encoding statutory law into logic as a way of reducing the syntactic ambiguity that is inherent in most statutes [2]. A decade ago computer scientists began their efforts to represent legal knowledge and model legal reasoning [15, 21, 23]. More recently, researchers have made serious efforts to translate legal rules into logic-based computer programs [30]. And now efforts are emerging to represent the unique factual and legal issues presented by individual cases [28]. Though these efforts may prove useful to lawyers and legal assistants,[3] they will not significantly impact the litigation morass that currently engulfs our courts because these systems treat the law as a set of abstract rules rather than a system manipulated by expert litigators. Lawyers do not make decisions based on what the rules "say." Rather, they base their decisions on how decision makers will apply those rules to specific factual situations.

In seminal research based on modeling the decision making of legal experts, Peterson and Waterman carefully examined how lawyers and insurance claims adjusters evaluate product liability claims [26]. From interviews with these experts they framed computer understandable rules that capture the way lawyers actually determine the settlement value of a case. For our case of the home buyer, some rules gathered from experienced litigators might read as shown in Figure 1. These rules can then be organized into decision trees (illustrated in Figure 2) that replicate the processes used by attorneys in estimating the settlement value of a case.

[3] Such systems could aid an administrator to determine whether a claimant is entitled to a retirement benefit because he became 65 years old but could not resolve a contested issue involving the question of whether a certain illness resulted in a compensable disability.

Peterson and Waterman recognize the limits of a strictly rule-based approach:

> But by themselves, the [rule-based] systems do not provide a general conceptual structure that can help us understand legal decision making. The rules are too specific, while the chains of reasoning are ad hoc products of the facts in particular cases.
>
> Peterson and Waterman [26, p. 632.]

The authors then propose a general conceptual structure for evaluating civil liability claims, based on their experience in interviewing litigators and building models of product liability. In this model, the evaluation of a claim is broken down into steps, which include:

1. **Analysis of loss:** What damages can be claimed? In our case of the home buyer, this would include direct economic losses such as the cost of repairs and damage to personal property, if any. In the product liability cases studied by Peterson and Waterman, general damages due to pain and suffering, limitation of activity, etc., could also be claimed.
2. **Analysis of liability:** What is the probability of establishing that a defendant is legally liable? This may depend on the rules of the jurisdiction as well as the facts. For example, in some jurisdictions it may be necessary to show that the broker actually knew about the defect, while in others it may only be necessary to show that the broker could have discovered it.
3. **Analysis of responsibility:** What proportion of the responsibility should be assigned to the plaintiff for his own carelessness? This also depends on the rules of the jurisdiction, since some allow damages to be reduced by a theory of comparative negligence.
4. **Analysis of characteristics:** Characteristics of the litigants, judges, and attorneys are considered by legal experts when they evaluate a claim.
5. **Analysis of context:** Strategic considerations such as the timing of the case, the plaintiff's immediate need for money, and the effect on the defendant of publicity, etc.

R1. IF the defect is material
AND the cost of repair is greater than $1,000
THEN increase the economic loss by the cost of repair

R2. IF the defect is material
AND there is damage to personal property greater than $500
THEN increase the economic loss by the damage to personal property

R3. IF the economic loss is greater than $2,500
AND the defect existed at the time of the sale
AND the broker knew of the defect
THEN increase the loss factor by 70 percent of the economic loss

R4. IF the buyer is the first occupant
THEN increase the sympathy factor by 5 percent of the economic loss (assumption that jurors have greater sympathy for buyers of new homes who expect a dwelling with fewer defects.)

R5. IF the buyer is a first time home buyer
THEN increase the sympathy factor by 10 percent of the economic loss (assumption that jurors will feel less chance that first time home buyer knows what to look for.)

R6. IF the broker is associated with a national chain
THEN increase the sympathy factor by 10 percent of the economic loss (assumption that big companies can more easily absorb losses than an individual home buyer.)

R7. IF there were visible signs of damage
AND the buyer was not a first time home buyer
THEN decrease the sympathy factor by 30 percent of the economic loss (assumption that jurors will not be sympathetic to a buyer who could have discovered the defect himself.)

FIGURE 1. Sample Rules for a Predictive Expert System

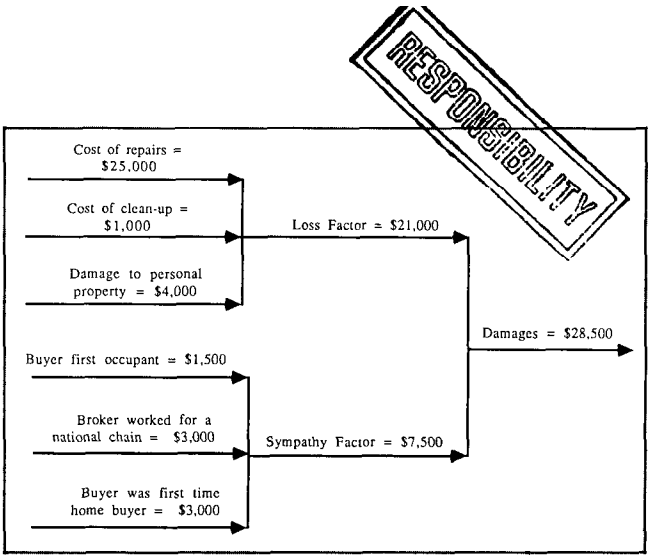

FIGURE 2. Decision Tree for a Predictive Expert System

The value of a case is estimated according to the following formula:

$$\text{VALUE} = \text{LOSS} * \text{LIABILITY} * \text{RESPONSIBILITY} * \text{CHARACTERISTICS} * \text{CONTEXT}$$

Figure 3 shows the decision tree for our case of the home buyer, re-structured according to this more general model. In this example, we will assume that the loss was $30,000, the probability of a plaintiff's verdict has been evaluated as .8, the plaintiff's responsibility for the loss has been evaluated as .3, the adjustment for characteristics was 1.1 (more favorable to plaintiffs than to defendant) and the adjustment for context was .9 (e.g., timing and strategy reduce the case value slightly), then the resulting case value would be:

$$\text{VALUE} = \$30{,}000 * .8 * .7 * 1.1 * .9$$

$$= \$16{,}632$$

We can see how this approach helps to explain the complex effect of a single factor: the degree to which the defect was discoverable. If the defect was not discoverable at all, then the question of liability (Step 2) would probably be answered in the negative since the

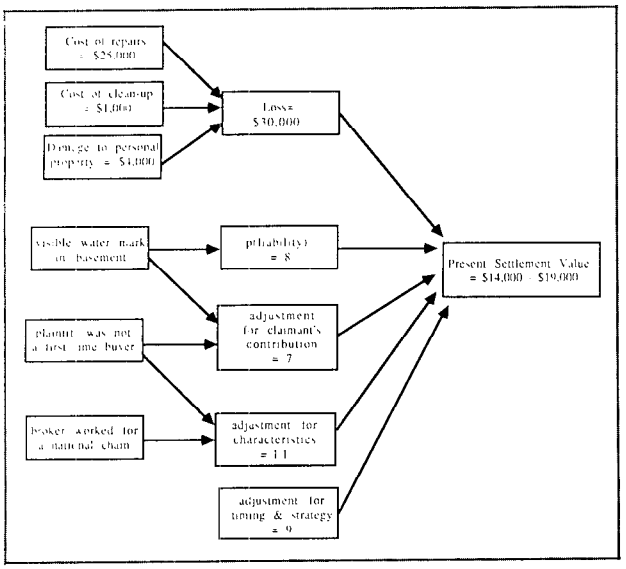

FIGURE 3. Decision Tree for a Settlement Value

broker was not at fault. If the defect was very easily discoverable then the buyer's carelessness (Step 3) would probably reduce the value of the claim. However, the answer to Step 3 would depend on both the discoverability of the defect and the experience (or inexperience) of the buyer, since a very inexperienced person might not recognize a defect that an average home buyer would easily discover.

The Impact of Predictive Expert Systems

Expert systems that predict the outcome of litigation with a fair degree of accuracy, if widely available, would have a profound effect on our legal system. If a court offered litigants the option of choosing a computer-predicted settlement, many lawsuits would end. Two factors would induce litigants to accept such settlements. First, many parties to litigation are risk averse, unable to absorb total defeat, thereby pressured to accept a predictable settlement that is below their desired goal. More importantly, greatly reduced attorneys' fees would make such settlements very attractive to all parties. Not only would a litigant avoid possible defeat, he or she would also avoid the high costs of litigation. However, what is desirable in theory may be difficult to attain in practice.

Researchers who have studied the implementation of information systems in organizations and in society at large [19, 20] have found that institutional impediments may retard the development and acceptance of computerized systems, no matter how beneficial they may be in the ideal case. In the real world, a complex interplay of socio-economic, political and interpersonal incentives determine whether and how computer systems are developed and used. Morrison [25], in a study of the attitudes of private law firms toward legal expert systems, found these institutional barriers to be a major cause of the slow growth of legal expert systems in the private sector.

In law, for example, it is clear that the current practice of time billing, as contrasted with value billing, provides a disincentive for attorneys to use legal expert systems which would reduce their number of billable hours. On the other hand, the client's desire for faster and less expensive legal service provides countervailing pressure to automate, and if legal expert systems can be shown to lead to better decision making this should provide an additional incentive for clients. However, this straightforward analysis may still be too naive; an expert system that expedites the settlement of cases would not necessarily be welcomed by all branches of a client corporation. For corporate managers who see their current assignments as short term, a settlement which is chargeable to this year's budget may be less desirable than a more expensive settlement negotiated at a later time, when the manager has moved on to a new assignment.

Institutional arrangements may also result in expert systems being developed and used for the benefit of one group of clients but not for their opponents. For example, it is not surprising that Peterson and Waterman's research on expert systems for product liability was funded by the insurance industry. Industries which undertake large amounts of repetitive litigation will have strong incentives to develop predictive expert systems for their own use. It is much less likely that expert systems will soon be used by the plaintiffs' attorneys who represent individual clients in suits against insurance companies.

Another environment where legal expert systems are likely to gain acceptance can be found in institutions providing free legal services to the poor. In legal aid clinics and the public defense bar, where attorneys on fixed salaries process a myriad of cases that have similar characteristics, there are strong incentives and few disincentives for using automation. In Project Pericles at Harvard Law School, an expert system was developed for representing indigent tenants, which led to an increase in both the quantity and effectiveness of service available to the tenants, causing problems for landlords who were required to pay increased hourly fees to their attorneys.

Even in the ideal case where predictive expert systems are developed and used based solely on their contribution to a more efficient legal system, a troubling paradox emerges if *all* litigants were to accept computer-predicted settlements. Accurate predictions result from encapsulating the thinking of experts who litigate cases. Without litigation the experts would have no additional data to input—leaving the law frozen. Fortunately, the universal acceptance of such settlements is unlikely. Since an expert system could not have rules that account for all possibilities, some litigants would find the computer-predicted settlement sufficiently unsatisfactory to justify the expenses and hazards of litigation. From such newly litigated cases the experts would develop new rules to update the system's performance. In addition, changing social values would lead some litigants to reject settlements based on relatively static rules. For example, a home

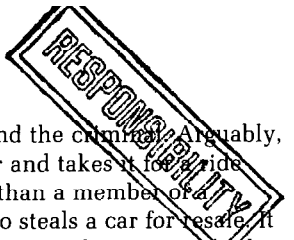

buyer might reject the computer-predicted settlement, recognizing the growing tendency toward consumer protection, and hoping this may lead the court to rule that brokers are liable even when they were not aware of the defect.

NORMATIVE LEGAL EXPERT SYSTEMS

Each year our criminal justice system sentences hundreds of thousands of criminals; our family court system determines the economic future of the spouses and children in over one million divorced families [6], and administrators of social welfare programs decide what benefits millions of clients should receive. Unfortunately, many citizens similarly situated receive greatly disparate treatment from discretionary decision makers.

For example, a welfare regulation might grant an extra allowance to a person who is "substantially incapable of performing normal household duties."[4] To one administrator, this might include a person who cannot carry groceries from the store; to another it might only apply to a person who is completely bedridden. Such discrepancies in interpretation may be uncovered during an official appeal, or they may go undetected indefinitely.

When a marriage ends in divorce, the legal system is often called upon to determine the division of property and the award of alimony and child support. These awards vary greatly depending on how a judge categorizes and values property; and how he or she measures the incomes and needs of spouses and children. Some judges may allow the custodial parent to remain in the marital home until the youngest child reaches the age of majority while others require the home to be sold within a short time after the divorce. Some judges favor lengthy alimony awards while others routinely deny such requests, favoring rehabilitative alimony that lasts only a few years. Needless to say such disparate awards create dissatisfaction among litigants [35]. Indeed, the potential for inequity is so great that over a decade ago it was suggested that maintenance awards be "calculated on a computer" [8].

Criminal sentencing, in particular, is receiving increasing attention from the legal community and the public. Studies reveal that some judges place the great majority of criminals on probation while others employ probation in a small number of cases [29], and anecdotal reports have documented numerous cases of nearly-identical crimes being committed by individuals with very similar backgrounds, in which one criminal was released while the other received a long prison sentence. Critics believe that inconsistent sentencing undermines efforts to both deter and rehabilitate criminals [31].

Public dissatisfaction with inconsistent sentencing has led to a growing demand for *determinate* or fixed sentences for certain crimes; however, this solution also engenders unfairness by failing to consider the unique characteristics of the crime and the criminal. Arguably, a 15-year-old who steals a car and takes it for a ride should be treated differently than a member of a professional car theft ring who steals a car for resale. It is clear that there is a need for agreed-upon standards or norms that would provide a framework for more consistent decision making but would still allow individual circumstances to be taken into account.

One approach to developing such norms would be to maintain a database of all decisions, which could be queried to find the typical treatment of particular types of cases. Judges and administrators would then be able to find out what happened in similar cases before making a decision. This approach has been pursued by the Law and Computers Project at the University of British Columbia. A Sentencing Data Base has been created which summarizes more than 40,000 sentencing decisions of the Provincial and Supreme Courts of British Columbia during the past few years. A very simple interactive interface asks the user to enter information about five factors:

1. the crime the defendant committed,
2. the age group of the defendant,[5]
3. the sex of the defendant,
4. whether or not the defendant is married,
5. whether or not the defendant has previously been convicted of an indictable (i.e., serious) offense.

The system then displays a histogram showing the frequency of various categories of sentences, for all cases in the database that match on the five factors. It is also possible to retrieve a list of the matching cases including the exact disposition, the judge, the location, and the plea. Judges in the Provincial Court, Supreme Court, and Court of Appeals of British Columbia have been using the Sentencing Data Base since late 1987, and many find it quite useful. In addition, both prosecutors and defense attorneys have used the system and found it useful.

The frequency analysis technique applied in the British Columbia sentencing project is a valuable contribution to legal decision making. If a judge is confronted by a crime of breaking and entering that seems to be of ordinary severity, he or she will not *inadvertently* give a sentence that is extremely harsh or extremely lenient. If a judge is confronted by a crime that is very heinous, but where the criminal is a first time offender, he or she can look up the most severe sentence previously given for this crime to a first time offender. However, there are three major limitations to this technique, where a rule based approach, used in addition to databases, may be useful.

1. A system that analyzes a few factors such as age, sex, and prior convictions does not approach the true complexity of judicial decision making. In making decisions, a large number of factors with complex interactions must be taken into account in order to determine a fair sentence: the nature of the

[4] This example, taken from the British welfare law, was discussed by Bench-Capon and Sergot [3].

[5] The system recognizes four age groups: under 21, 21–29, 30–59, and over 59.

crime including mitigating and aggravating factors; the individual characteristics of the defendant including age, education, family status, reputation in the community, etc.; a complete review of the defendant's prior criminal record; and the existence of any extenuating circumstances. An attempt to model these hundreds of factors and their interactions would preclude the use of simple statistics. Even if one could enumerate them, it would be virtually impossible to collect this information for all prior cases. Existing records often do not contain this information but even if the careful perusal of these records would yield the information, society lacks the resources necessary to encode this massive amount of data.

2. If our current decision makers are behaving inconsistently, as the anecdotal evidence seems to suggest, then the mere reporting of past behavior does not ensure a meaningful pattern for judges to base their future decisions on. The mode of sentences for a particular crime may represent an accidental congruence of light sentences from some judges and severe sentences from others for a collection of crimes and defendants with little in common other than the formal charge.
3. Databases cannot explain *why* a particular sentence was given, while the goal of expert systems is to model the reasoning of human experts, not only at the performance level but also at the cognitive level. One of the most important attributes of expert systems is their ability to explain why a particular analysis or recommendation was produced (in terms of their internal rule-based model). The creation and use of an expert system puts legal decision making under the microscope of precise computer specification. In areas of discretionary decision making that are believed to suffer from undesirable inconsistency, AI offers an exciting opportunity to bring more rational approaches into being.

It is possible that AI research could result in expert systems that would help judges produce an acceptable level of evenhandedness without resorting to a rigid model of fixed sentencing. The expert system would be normative rather than predictive, providing guidance based on a complex model encompassing characteristics of the crime, mitigating and aggravating circumstances, and the characteristics of the individual defendant. Although such a system might influence a judge's decisions, he or she would always be free to reject its recommendations, since no matter how many factors were included in the expert system's model, there would always be cases whose unique circumstances would justify a different result.

Modeling Judgemental Decisions: Sentencing as Diagnosis

Expert systems have achieved a high level of performance in a number of specialized domains, including mineral prospecting, computer configuration [22] and chemical analysis [11]. But perhaps the largest amount of significant expertise to be represented in expert systems is in the area of medical diagnosis and therapy [7, 10, 27, 34]. For example, according to Hayes-Roth et al. [18], the CADUCEUS system in 1982 possessed approximately 100,000 associations representing more knowledge of internal medicine than any human, and was able to correctly diagnose complex test cases missed by human experts.

Medical expert systems are based on a model of heuristic reasoning that relates symptoms of illness to suggested therapy, based on intermediate concepts of *disease states* (e.g., infectious disease organisms, conditions such as alcoholism, heart disease, etc.). The physician diagnoses the patient's symptoms by making a judgement about the disease state(s) that caused them, determines the general type of therapy required (antibiotics, surgery, special diet, etc.) and finally prescribes a specific treatment.

This model can be adapted to the sentencing of criminals by viewing the criminal acts as symptoms and the sentences as therapy. The general types of therapy correspond to various sentencing goals: specific deterrence, general deterrence, isolation of dangerous individuals, and rehabilitation [17, 24]. These in turn lead to specific recommendation: probation, a fine, a prison term, or a diversionary program such as community service or treatment for drug abuse. A relevant question that can only be answered through empirical research is: do judges use some notion corresponding to disease states to mediate between the description of a criminal act and the determination of an appropriate sentence?

The MYCIN system, [7] one of the first and best known medical expert systems, defined a basic approach to heuristic reasoning that appears to be very promising for modeling legal judgements:

1. Knowledge is represented by production rules, applied by a domain-independent inference engine. (This is now a standard approach in expert systems.)
2. A conceptual hierarchy of objects, attributes, and values is used to structure the problem, and is integrated with the inference engine. Thus, the sophistication of the rule formalism used in MYCIN is far beyond that normally found in expert system shells. Figure 4 shows a hypothetical context tree that might be applied in evaluating a criminal for sentencing. Using a top-down approach, first the nature of the crime would be considered, including each offense that the defendant is charged with, the defendant's actions leading to that charge, the victims of the crime, and the surrounding circumstances. Then, the characteristics of the defendant would be considered, including personal data and information on previous crimes. Finally, the existence of mitigating and aggravating circumstances would be considered.
3. Weights are attached to all knowledge elements in the database to indicate how strongly they are believed, and rules are also weighted to show how

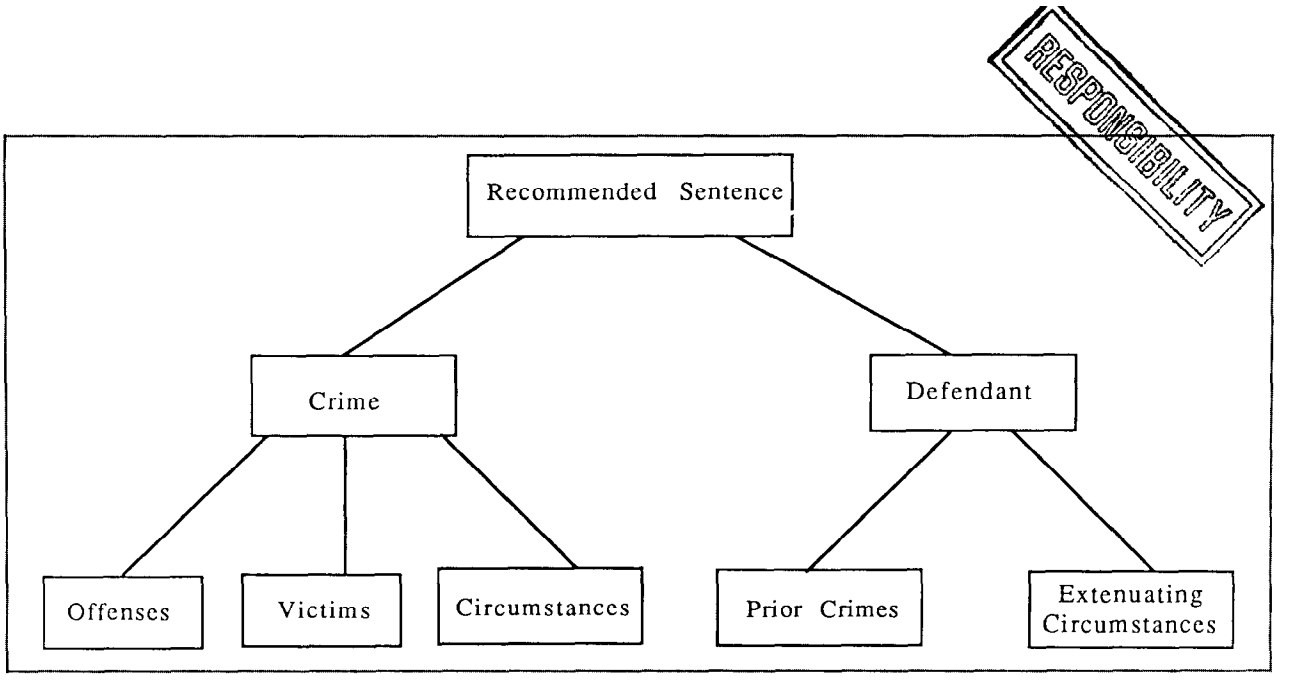

FIGURE 4. Context Tree for a Diagnostic Model of Criminal Sentencing

important they are. The system provides an algorithm for combining weights to measure the overall strength of each conclusion. The ability to represent varying degrees to which attributes apply (e.g., degrees of violence of an assault) is critical in the legal domain.

Figure 5, although greatly oversimplified, shows the kinds of rules that might be included in such an expert system. The sample rules are intended to help in weighing the evidence for and against the hypothesis, "D is dangerous to society." For example, rule R1 in Figure 5 says the "danger factor" associated with illegal use of firearms is .5 and rule R2 says the danger factor associated with possession or sale of drugs is .5. Using the MYCIN rule for combining evidence [33], the "danger factor" associated with both of these offenses would be .75.

4. Rules are allowed to explicitly refer to the strength of belief in the factors appearing in their premises; thus, a jail term may be indicated when the danger factor exceeds a particular value, e.g.:

   IF there is strongly suggestive evidence (.7) that D is dangerous to society
   THEN the recommended sentence for D should include incarceration 1.0.

Experiments with medical expert systems have shown several aspects of the MYCIN model that need to be improved. One aspect is the reliance on top-down control, which frequently leads to a line of questioning that has nothing to do with the current case. Another aspect is the fact that symptoms tend to occur in meaningful clusters, so that their weights are not really independent. The CADUCEUS system [27] includes an initial *consultation* in which the user can enter the symptoms of the current case in whatever order he or she wants. CADUCEUS also employs a more complex model of medical knowledge, which takes account of the way certain combinations of symptoms provide strong evidence for particular conclusions.

Is the MYCIN model an appropriate tool for modeling the way judges evaluate criminals for sentencing? Will the described limitations of MYCIN also be relevant in the legal domain? Or will other information structures that have not yet been considered be suggested by a study of judicial decision making? As AI researchers attempt to answer these questions, they will contribute to a better understanding of how our legal system works. Whether or not expert systems are actually used in the courtroom, the exercise has potential value as a tool for jurisprudential analysis.

The Impact of Normative Expert Systems

Even if computers did not exist it would make sense for legal decision makers to carefully analyze how and why they make their decisions. To that end, some courts have formal mechanisms by which judges gather to share information. In addition, many informal mechanisms permit decision makers to exchange views about the laws they administer. However, the growing complexity of the legal system and the time pressure on decision makers makes it difficult to maintain norms through merely gathering to share information.

One solution which is gaining popularity is the development of statutory guidelines. When legislators mandate the creation of guidelines for sentencing criminals or dividing up marital property they strive to incorporate societal wisdom on how to best resolve a certain class of case. But written guidelines are of limited utility when compared to expert systems' ability to model certain kinds of decision-making problems. Most importantly, the computer can handle a far greater number of interacting rules than a person looking at a printed page; thus, the practical limitations of legislative drafting preclude the listing of the myriad of factors that ideally should be taken into account. In addi-

R1. IF one of the offenses was illegal use of firearms
THEN D is dangerous to society .5

R2. IF one of the offenses was possession of illegal drugs
OR one of the offenses was sale of illegal drugs
THEN D is dangerous to society .5

R3. If one of the offenses involved bodily harm
AND D was acquainted with the victim
THEN D is dangerous to society .4

R4. IF none of the offenses involved bodily harm
AND D was acquainted with the victim
THEN D is dangerous to society −.2

R5. IF one of the offenses involved bodily harm
AND D was NOT acquainted with the victim
THEN D is dangerous to society .8

R6. IF the circumstances included the influence of alcohol or drugs
AND D's crime was driving-related
AND D's crime did not involve bodily harm
THEN D is dangerous to society .25

R7. IF the circumstances included the influence of alcohol or drugs
AND D's crime involved bodily harm
THEN D is dangerous to society .7

R8. IF D has previously spent time in prison
THEN D is dangerous to society .5

R9. IF the number of D's previous crimes = 0
THEN D is dangerous to society −.5

R10. IF D is regularly employed
THEN D is dangerous to society −.5

FIGURE 5. Sample Rules for a Normative Expert System

tion, the legislative process makes fine tuning a system most difficult while expert systems can be easily adjusted to correct errors and omissions.

Normative expert systems would never replace human judgement. Rather, they would provide decision makers with the tools to better understand, evaluate, and disseminate their ideas. No computer would ever sentence a defendant, make a child support award, or grant a housing allowance. But a computer could advise a decision maker on what factors and issues should be considered, and how his or her colleagues would be likely to view a case with similar facts.

In spite of the potential of normative expert systems to assist discretionary decision makers, there are significant barriers to their development and use. First, decision makers who are accustomed to viewing themselves as independent, judgemental thinkers may view expert systems as mechanizing or "de-skilling" their jobs. To the extent that this perception exists, the systems are likely to be strongly resisted. The best way to overcome this barrier is to make the decision makers the "owners" of the knowledge base, so that they will see it as a tool that they have created for themselves. However, most legal decision makers are not accustomed to using computers, so there is a "comfort" barrier that must be overcome, as well as the natural skepticism of professionals that computers could be an effective tool in their domain.

Second, if the use of an expert system would change the distribution of decisions, this might be unacceptable to the administering organization or to society. For example, many social entitlement programs such as workers compensation, medicaid, and social security disability, have statutory rules that determine eligibility. According to some experts,[6] the workers compensation rules are often misapplied to the detriment of the workers—resulting in a much lower cost to the government. If this view is correct, a legal expert system that accurately embodied these rules could result in a significant increase in benefits paid out, a cost which society might not be willing to pay. On the other hand, if social security disability benefits, which are determined by the states but paid by Washington, are being administered too liberally, state officials might resist the use of an expert system which would lower the overall benefits paid to their constituents.

The relationship between computer information systems and the development of norms has been studied by Kling [19], whose research shows that when norms are not taken into account, development projects are likely to fail. Kling presents several examples, including a project to upgrade the U.S. World Wide Military Command and Control System (WWMCCS), which was criticized in a 1981 report by the General Accounting Office (GAO) as follows: "DOD, despite dozens of large-scale studies, has failed to make meaningful progress . . ." [19, p. 323]. Kling traces the failure of the WWMCCS modernization to the need for field commanders in the three services to agree on precisely what data they need. He notes that different tactical battle doctrines held by officers in the three services result in different requirements for an information system, and suggests that until this problem is addressed the project cannot succeed.

Kling's research, which is focused on the organizational response to computer systems, reveals how the absence of shared norms can be an obstacle to achieving the primary goal of developing a useful computer system. On the other hand, in certain areas of law such as sentencing and child support, the development of norms is a primary goal; the use of an expert system is justified by its contribution to that objective. In either case, however, the need to agree on the behavior of the computer brings normative issues dramatically into focus. It is this focusing effect which is the primary benefit claimed here for legal expert systems.

There are potential dangers as well as benefits in the use of normative expert systems. Critics of artificial intelligence such as Weizenbaum [36] have warned about over-reliance on computers; for example, a decision maker might be reluctant to vary from the computer's advice even when the facts so warranted. This is not a new problem for the legal system; the same danger exists with presumptive sentencing when the legislature suggests, but does not require, judges to sentence within certain parameters. In the same way that a judge might blindly follow a legislative guideline, a

[6] James Rowan, personal communication.

judge might ignore a directive to treat the output of a computer as merely advisory. Furthermore, it is possible that expert systems will carry greater implicit authority than legislative guidelines, due to their ability to respond to more factual variations and also due to the dynamic quality of computer programs. Although expert systems, if properly constructed, have the ability to explain the reasoning on which their recommendations are based, it is all too easy to just accept the recommendations without examining the reasoning behind them.

Is this a realistic scenario? Perhaps in the early 1970's when Weizenbaum was building the ELIZA program, it was more realistic than it will be in the 1990's. In a seminar on Artificial Intelligence and Law taught by the authors at Northeastern University, we have observed a growing understanding on the part of law students who take the course on the capabilities and limitations of computers. It is not unreasonable to suppose that, as computers become commonplace in schools and colleges, as youngsters grow up playing "adventure" games, and as computer programs for various tasks (for example, tax preparation and estate planning) become commonplace, expert systems will become a part of everyday life whose uses and dangers are generally recognized, just as electricity and automobiles are today. The dangers of any technology must be weighed against its benefits. The question for expert systems is whether they can contribute to making the legal system more fair and less costly than it is today.

Of all the issues surrounding the use of normative expert systems, perhaps the most significant is the possibility for review of discretionary decisions in areas where such review is now difficult to obtain. Most jurisdictions do provide for the review of criminal sentences and alimony awards. But as a practical matter such reviews rarely succeed because appellate judges routinely hold that such decisions are within the discretion of the trier of fact. Such non-reviewability is virtually dictated by the unavailability of norms. Rarely can an appellate judge find a sentence or alimony award unreasonable because standards of reasonableness do not exist. The process of knowledge engineering that would be required in order to build normative expert systems would have the desirable side effect of helping to develop such standards.

CONCLUSIONS

Weizenbaum [36] argues that computers ought not to make judicial decisions, even though they could conceivably do so. We strongly agree, and in another report [4] discuss a number of reasons why legal decisions cannot be explained by any rule-based model of law, however complex. The legal decision maker not only applies existing rules, but is often called upon to creatively modify rules, to choose among competing rules, and to create new rules [14].

But we must recognize the practical benefits as well as the theoretical limitations of technology:

a. In legal disputes where money is at stake, the cost of the dispute resolution process must be allowed to enter the moral equation. If a perfectly just award must be shared 50–50 with a lawyer, an approximate result computed by an expert system of which the petitioner keeps 90 percent may be preferable.
b. Where discretion has been shown to foster gross inconsistency and shelter discriminatory practices, it may be an improvement to have an expert system generating a range of standard alternatives, with a requirement for written justification when decisions fall outside the range. The important questions to ask are: Would more people get a fairer result? and: Would more people understand WHY they got the result they did?
c. When citizens governed by a system of legal rules cannot ascertain the content of those rules, we must question whether we have a system of law, as that term is commonly understood [13]. Given the expanded complexity and amount of legal materials, the cost of "finding the law" has, for many citizens, rendered the law unfathomable. By making the law more accessible, conceptual retrieval systems have the potential of enhancing the legitimacy of our legal system.

The specter of a computer sentencing a criminal or forcing parties in a dispute to accept a predetermined settlement is often raised as a potential negative result of developing expert systems for law. But such things are happening even without expert systems—the emergence of no-fault insurance, proposed caps on malpractice awards, the trend toward determinate sentencing, and formal child support guidelines are all indications that, in reaction to the excessive cost and complexity of the legal system, citizens are willing to resort to an algorithmic determination of their rights. Algorithms that are written in the statute books and applied mechanically by human administrators are the paper equivalents of allowing a computer to make the decisions. Expert systems can be used in the service of such algorithmic decision making, but they can also be used to organize and present the relevant facts and issues in the service of human decision makers. The predictive and normative expert systems described in this article are intended to serve the latter function.

REFERENCES

1. American Law Institute. Using Personal Computers for Decision-Making in Law Practice. ALI-ABA Invitational Conference, 1985.
2. Allen, L.E. Symbolic logic: A razor-edged tool for drafting and interpreting legal documents. *Yale Law J. 66*, (1957), 837.
3. Bench-Capon, T. and Sergot, M. *Toward a rule-based representation of open texture in law*. Deptartment of Computing, Imperial College of Science and Technology, London, England (Apr. 1985).
4. Berman, D.H., and Hafner, C.D. Indeterminacy: A challenge to logic-based models of legal reasoning. *3 Yearbook of Law Computers & Technology* (1987, to appear). Also published as Research Report NU-CCS-87-1, College of Computer Science, Northeastern University, Boston, Mass., 1987.
5. Bok, D. A flawed system. *Harvard Magazine 85*, 5 (1983), 38–40.
6. Bureau of the Census, U.S. Dept. of Commerce. *1981 Statistical Abstract of the United States*, 80.
7. Davis, R., Buchanan, B., and Shortliffe, E. Production rules as a

representation for a knowledge-based consultation program. *Artif. Intell. 8*, 1 (1977), 15–45.
8. *Desk Guide to the Uniform Marriage and Divorce Act*, (BNA, Washington, D.C., 1974). Remarks of Ms. Betty Berry, past National Coordinator of NOW's Marriage and Divorce Task Force.
9. *Easton v. Strassburger*, 152 Cal. App. 3d 90 (1984).
10. Fagan, L.M., et al. Representation of cynamic clinical knowledge: Measurement interpretation in the intensive case unit. *IJCAI-6* (1979), 260–262.
11. Feigenbaum et al. On generality and problem solving: A case study using the DENDRAL program." In *Machine Intelligence 6*, B. Meltzer and D. Michie, Eds. Edinburgh University Press, (1971), 165–190.
12. Frank, J. *Law and the Modern Mind* (1963), xxvii–xxviii.
13. Fuller, L. *The Morality of Law* 39 (1964).
14. Gardner, A.L. Overview of an artificial intelligence approach to legal reasoning. In Walter, C. Ed. *Computer Power and Legal Reasoning*. West Publishing Co., St. Paul, Minn. (1985), 247–74.
15. Hafner, C.D. *An Information Retrieval System Based on a Computer Model of Legal Knowledge* Ph.D. dissertation, Univ. of Michigan, 1978. Republished by UMI Research Press, Ann Arbor, Mich. (1981).
16. Hafner, C.D. Conceptual organization of case law knowledge bases. In *Proceedings of the 1st Intelligence Conference on Artificial Intelligence and Law*. ACM, New York, 1987, pp. 35–42.
17. Hart, H. The Aims of the Criminal Law. *23 Law & Cont. Prob.* 401 (1958).
18. Hayes-Roth, F., Waterman, D.A., and Lenat, D.B. *Building Expert Systems*. Addison-Wesley, Reading, Mass., 1983, 9.
19. Kling, R. Defining the boundaries of computing across complex organizations. In *Critical Issues in Information Systems*, R.J. Boland and R.A. Hirscheim, Eds. John Wiley, New York, 1987, 307–362.
20. Kraemer, K.L., Dickhoven, S., Tierney, S.F., and King, J.L. *Datawars: The Politics of Modeling in Federal Policymaking*. Columbia University Press, New York, 1987.
21. McCarty, L.T. Reflections on TAXMAN: An experiment in artificial intelligence and legal reasoning. *Harvard Law Review 90*, (Mar. 1977), 837.
22. McDermott, J. R1: A rule-based configurer of computer systems. In *Proceedings of the 1st National Conference on Artificial Intelligence (AAAI-80)*, pp. 269–271.
23. Meldman, J.A. A structural model for computer-aided legal analysis. *Rutgers Journal of Computer Law 6*, (1977).
24. Model Penal Code, Sec. 1.02(2) American Law Institute, 1962.
25. Morrison, R.W. Market realities of rule-based software for lawyers: Where the rubber meets the road. In *Proceedings of the 2nd International Conference on Artificial Intelligence and Law (ICAIL 89)*. (Vancouver, Canada, June 1989). ACM, New York, pp. 33–36.
26. Peterson, M.A., and Waterman, D.A. An expert systems approach to evaluating product liability cases. In Walter, C. Ed. *Computer Power and Legal Reasoning*. West Publishing Co., St. Paul, Minn. 1985, 629–59.
27. Pople, H.E., Jr. Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics. In *Artificial Intelligence in Medicine*. P. Szolovitz, Ed. Westview Press, Boulder Colo., 1981, 119–185.
28. Rissland, E.L., Valarce, E.M., and Ashley, K.D. Explaining and arguing with examples. In *Proceedings of the 1984 National Conference on Artificial Intelligence (AAAI-84)*.
29. Rubin, S. Disparity and equality of sentencing—A constitutional challenge. *Federal Rules Decision 40*, 55 (1965).
30. Sergot, M.J., et al. The British Nationality Act as a logic program. *Commun. ACM 29*, 5 (May 1986), 370–385.
31. Shane-Dubow, S., Brown, A., and Olsen, E. *Sentencing Reform in the United States*. National Institute of Justice, U.S. Department of Justice (1985).
32. Subrin, S.N. How equity conquered common law: The federal rules of civil procedure in historical prospective. Pa. La. Rev., Vol. 135 (Apr. 1987). Manuscript on file at Northeastern Univ. School of Law, Boston, Mass. 02115.
33. Van Melle, W., Scott, A.C., Bennett, J.S., and Peairs, M. The Emycin Manual. Report No. STAN-CS-81-885, Department of Computer Science, Stanford Univ., Stanford, Calif. (1981).
34. Weiss, S.M., et al. A model-based consultation system for the long-term management of glaucoma. *IJCAI-5* (1977), 826–832.
35. Weitzman, L.J. *The Divorce Revolution*. The Free Press, New York, 1985.
36. Weizenbaum, J. *Computer Power and Human Reason*. W. H. Freeman and Co., San Francisco, 1976.

CR Categories and Subject Descriptors: I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems—*law*; K.4.0 [**Computers and Society**]: General; K.4.3 [**Computers and Society**]: Organizational Impacts
General Terms: Legal Aspects
Additional Key Words and Phrases: Computers and law, social impact of artificial intelligence

ABOUT THE AUTHORS:

DONALD H. BERMAN graduated from Harvard Law School and now holds the title of Richardson Professor of Law at the Northeastern University School of Law in Boston. An expert in both tax law and legal expert systems, Berman formed the Center for Law and Computer Science at Northeastern in 1985 with coauthor Hafner.

CAROLE D. HAFNER is an associate professor at Northeastern University's College of Computer Science. Hafner holds a Ph.D. in computer and communication sciences from the University of Michigan. Her dissertation, published in 1978, was one of the first research projects to develop a computer model of legal knowledge.
Authors' Present Address: Center for Law and Computer Science, Northeastern University, Boston MA 02115.