

AI and LEGAL REASONING

Report of a panel chaired by

Edwina L. Rissland
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003

Abstract

This paper presents a summary of the responses of a panel to issues on AI and legal reasoning. The panel consisted of: Edwina L. Rissland, Chair (University of Massachusetts), Kevin D. Ashley (University of Massachusetts), Michael G. Dyer (UCLA), Anne v.d.L. Gardner (Stanford), L. Thorne McCarty (Rutgers), and Donald A. Waterman (RAND). Among the issues addressed by the panel were:

1. What are the characteristics of the legal domain that make it interesting or amenable to AI approaches - what is special about it;
2. The open-textured nature of legal concepts and the implications this has for using AI-techniques, especially knowledge representation;
3. The complementarity of rule-based and case-based reasoning - how cases are used, especially when the rules "run out";
4. The pervasive role of analogy in legal reasoning;
5. The special role played by hypotheticals in the legal domain and how *hypos* help with argumentation and strategic case planning;
6. The interleaving of justification, explanation, and argumentation;
7. How common law systems can be seen to be systems which learn from cases;
8. The appropriateness and feasibility of intelligent aids for practicing litigators and other legal experts;
9. Implications for other domains - like medicine - that use case-based reasoning;
10. Methodological and other issues.

For each issue considered, the comments of the panelists are summarized.

1. The Challenge and Special Characteristics of the Legal Domain

The legal domain presents some very interesting challenges to the AI researcher. While it is a domain which has established standards for deriving new truths (e.g., *stare decisis* or the doctrine of precedent), it is more of a "scruffy" domain than a "neat" one, despite its orderly, rule-like surface veneer. It is very much an experience-based example-driven field. Legal reasoning is also heavily intertwined with natural language processing and common sense reasoning and therefore inherits all the hard problems that these imply.

Several panelists emphasize that legal reasoning and argumentation take special skills and that learning to think like a lawyer requires considerably more than rote memorization of a large number of cases, a daunting task in itself. For instance, Dyer says:

Modeling what a lawyer does is more complex than modeling experts in technical/scientific domains. First, all of these complex conceptualizations are expressed in natural language, so modeling the comprehension ability of the lawyer requires solving the natural language problem. For example, giving legal advice often starts with hearing a "story" where the client was one of the actors. So legal advice often presupposes a story understanding capability. In other expert systems, the natural language problem can be largely finessed since the task (e.g., disease diagnosis, reconfiguring hardware, analyzing dipmeters) rarely involves the communication of complex conceptualizations. Second, the planning in this domain is complex. This is a world where actors obtain agents, where they counterplan, retaliate, attempt to set up subsumption states, etc. Third, Anglo-American jurisprudence, at least, makes use of a large body of concrete EPISODES. Becoming a lawyer requires learning how to index, generalize and apply these concrete episodes in service of reasoning and argumentation strategies. Fourth, law is simply an attempt to formalize COMMON SENSE reasoning and notions of justice, morality and fair play.

Rissland points out the analogies between legal and mathematical reasoning. In her analysis, learning to be an expert mathematician or mathematics student involves a lot more than simply learning definitions, theorems and proofs. There is a wealth of other knowledge and skills needed: for instance, there is a large corpus of examples and heuristics, and there is the important skill of being able to generate examples [Rissland, 1978]. Similar remarks can be made about the legal domain [Levi, 1949; Llewellyn, 1930]. In particular, cases in law play the analogous role to examples in other fields. However, in addition to the usual questions about examples - such as their function and generation - the law is a domain which is largely case or example-based and thus examples are truly central.

The law also exhibits many of the traits discussed by Kuhn with regard to scientific disciplines [Kuhn, 1970]. In particular, the dialectic between proposing an idea or rule and the testing of it through concrete cases – so well discussed by Lakatos [1976] and Polya [1973] – can also be found in the law and other fields [Rissland, 1984c].

Waterman notes several special characteristics of the law:

The law combines many different kinds of reasoning processes: rule-based, case-based, analogical, hypothetical.

It has the unique property of being pseudo-formalized, i.e., there exists a large body of formal rules that purport to define and regulate activity in the domain. However, these rules are often deliberately ambiguous, contradictory and incomplete. Many of the actual rules used in legal reasoning are rules about how to access and reason with the "formal" rules. The problem this creates is the naive notion (for some) that because a body of rules and regulations exists, all one has to do is translate them into executable code to create a legal reasoning program.

The law is in a constant state of change, so expert legal reasoning systems have to be easy to modify and update. Many other application domains have this characteristic, but perhaps not to quite the same extent.

Gardner not only joins others in pointing out the challenge for AI presented by the intertwining of natural language and common sense knowledge in legal reasoning but also she remarks on some differences concerning the nature of expertise in the law compared with other domains of expert systems work:

Law is an area for "expert systems" in the sense of involving professional knowledge, but it is unusual in that we have to expect the experts to disagree. (E.g., opposing counsel; judges writing majority and dissenting opinions.) Questions then arising are: (1) What kinds of things can they disagree on, and what kinds of disagreement are beyond the bounds of professional competence? (2) How can a reasoning system leave room for the appropriate kinds of disagreement?

Law is also an area deeply involving natural language and commonsense understanding. In particular: (1) Legal materials (such as statutes and other statements of legal rules) are written in natural language, and the problem of saying how they apply to a particular situation is in large part a problem of saying

what this language means. This aspect of law calls for a different approach to meaning from that taken in most AI programs. (See discussion of open texture below.) (2) The situations to which we might want to apply legal knowledge encompass practically the whole range of human activity – with a special focus on those human situations where something goes wrong. Thus a truly expert system would be able to represent and reason about such situations using a great deal of commonsense knowledge as well as technical knowledge. This last is true of some other kinds of expert systems too.

2. The Open-Textured Nature of Legal Concepts

Unlike certain other domains like mathematics or even medicine, concepts in the law tend to be "open-textured", that is, one cannot provide black and white definitions. Interpreting a legal concept in a new situation depends on past interpretations. To decide whether a legal predicate applies to a fact situation, one must usually find cases in which the predicate was applied and focus on similarities or dissimilarities.

Gardner, who addressed this issue in considerable detail in her recent doctoral thesis [Gardner, 1984, especially Sections 3.1-3.3], summarizes:

Legal language is open-textured in that (1) we cannot state necessary and sufficient conditions for the application of legal predicates and (2) neither can we take legal predicates as primitives in the sense that a program could simply recognize whether or not the predicate applies to some fact situation. (That is, a legal rule (Px implies Qx) and a fact Pa do not entail legal consequence Qa .)

This characterization of open texture covers several phenomena. One is the variable standard (e.g., "a reasonable time"). Another is vagueness (although legal rules do tend to avoid using obviously vague words – legislators write rules about people over 65, not about "old" people). The most important, however, is that legal conclusions that appear to follow deductively are in fact defeasible. That is, given the legal rule (Px implies Qx) and the fact Pa , the conclusion Qa is only a default conclusion.

Ashley points out that the law's handling of open-textured predicates is not neutral, that is, the decision about the concept is tied in with who the decider wants to win the case:

In deciding whether a legal predicate applies to a fact situation, a court usually also

determines who should win the issue or the case. The criteria for deciding whether the predicate applies are rarely completely specifiable in rules. The court must usually find cases in which the predicate was applied and focus on similarities or dissimilarities between the new fact situation and past cases. It must decide if the comparison cuts in favor of following the prior case or distinguishing it. This analysis is significantly different from trying to apply some abstract definition of the legal predicate. Any attempt to model legal reasoning must take account of this aspect of case-based reasoning.

3. The Complementarity of Case-Based and Rule-Based Reasoning

One way cases are used is to give concrete content to the concepts and abstractions in legal rules. Even if one believes that the law can be captured in rules – which many, particularly the legal realists, do not – one needs cases to flesh out the meaning and intent of the rules. As Gardner has put it, cases are what you use “when the rules run out”. Most scholars of the common law, would say that cases are what you use all the time – even if you have rules and know how to apply them. McCarty has long campaigned the importance of this point for AI work on legal reasoning; for instance see [McCarty, 1977; McCarty and Sridharan, 1981a].

Rissland restated the theme of the intertwining of case-based and rule-based reasoning, and that it is akin to the proofs-and-refutations dialectic between concept/conjecture and examples talked about by Lakatos. In fact, she suggests, as has Gardner, that one could carry this computationally further for handling the different types of questions discussed by Gardner [Gardner, 1984]: a rule-based approach for the “easy” or black-and-white questions and a case-based approach for the “hard” or gray-area questions. A rule-based reasoner should be able to call on the services of a case-based reasoner when the predicates get murky or the rules run out and a case-based reasoner should be able to call a rule-based reasoner for the well-understood, well-structured aspects of its problems.

4. The Pervasive Role of Analogy

Ashley identifies several uses of analogy in legal argumentation and practice [Ashley, 1984] including:

1. Precedents that are analogous in some sense to the case being decided are used to focus the attention of the advocate on those issues and authorities which have a bearing upon the argument;
2. The precedent provides an example of what a reasonable argument in an analogous context looks like;

3. An existing classification from a precedent may be applied to analogous facts in the case being decided;

4. A new classification, analogous to a classification in a precedent, may be created and applied to the facts in the case being decided.

Gardner points out that if one takes analogy to refer to comparing facts of cases for the purpose of arguing that a new case should be decided in the same way as an old, similar case, then the crux of the issue is what makes some factual similarities and differences important and others irrelevant; this involves more than a simple comparison of facts.

McCarty has addressed this issue in his TAXMAN II research, especially concerning the famous tax case of *Eisner v. Macomber*. If one can imagine a sequence of intermediate hypotheticals connecting the two cases to be analogized then the sequence of intermediate transformations provide a key to the analogy. For a detailed analysis, see [McCarty and Sridharan, 1981b].

Dyer echoes some of McCarty's and Rissland's concerns, particularly on the utility of the idea of deformations and standard or prototype cases, with an added emphasis on the key role of memory in analogy:

Functioning successfully as a lawyer requires having indexed key cases in such a way that the comprehension of novel cases leads immediately to access of key cases. In a sense, legal reasoning and comprehension involves finding “islands” of legal prototypes from which well-known reasoning paths can be adapted to the novel situations and then traversed. Much of legal reasoning, therefore, is a memory indexing and search problem rather than a logical or “theorem-proving” problem.

5. The Role of Hypotheticals

Hypotheticals – that is, make-believe cases – serve many roles in legal reasoning. Several of the panelists, in particular, McCarty, Rissland, and Ashley, agreed on their importance. All three make strong use of hypotheticals in their models of legal reasoning and generate them in a somewhat similar framework (called “prototype plus deformation” by McCarty and “modification plus retrieval” by Rissland, who also taxonomizes examples into classes like “model”, standard reference, anomaly, etc.).

Rissland enumerates some of the uses that hypotes play in legal reasoning and argumentation [Rissland, 1984a, b]. For instance, some general functions of hypotes are:

1. hypotes remake experience
2. hypotes create experience
3. hypotes organize and cluster cases
4. hypotes tease out hidden assumptions

For instance, the second usage allows one to consider issues that have not yet been litigated, and may well be, (e.g., the recent case of the frozen embryo which was predeceased by its parents or a malpractice suit against an expert system). Thus hypos can contribute to strategic planning by providing what-if situations – especially contrary ones – to consider. This can lead to a debugging an argument or strategy ahead of time.

Ashley elaborates on three roles played by hypotheticals in the context of argumentation:

1. In a legal argument, a hypothetical can be used to isolate the weaknesses and strengths of an attorney's case. A hypo can be constructed to emphasize one aspect of a case that pushes strongly for a decision contrary to the proponent's, perhaps by making it more like a line of contrary cases or by exaggerating some common sense objection to following the proponent's classification. One way for the proponent to respond is to distinguish the hypo and his case. In so doing, the proponent is forced to emphasize the other strengths of his case that are not present in the hypo.

2. A hypothetical can be used to illuminate the consequences of a decision. A court may ask, "If we decide in favor of the proponent on this case, what will we do when a case with somewhat less sympathetic facts is posed?" It is an exercise in line drawing. If no defensible lines can be drawn, the court may choose not to decide in favor of the proponent, or at least not on the grounds argued for.

3. In both of the above, the hypothetical is used as a short hand for lines of cases, lines of reasoning, even whole arguments.

Such uses of hypotheticals in legal argument can be illustrated by excerpts from oral arguments before the United States Supreme Court in which the justices pose hypotheticals to make and elicit legal points and to control the presentation of a proponent's case (e.g., the questioning by the Justices in the recent civil rights case of *Gomez-Bethke v. U.S. Jaycees* (1983)) and to justify and explain their conclusions (e.g., see the opinion of Justice Pitney and the dissent of Justice Brandeis in the landmark corporate tax case of *Eisner v. Macomber* (1920)).

6. Interleaving of Justification, Explanation and Argumentation

Waterman reminds us that explanation (explaining how you reached a conclusion) and justification (providing a convincing argument that the method used to reach the conclusion is valid) aren't as easy in law as they are in

some other domains which have clear underlying (deep) models of the mechanisms involved, mechanisms that can provide predictive and explicative power. He is joined by McCarty in this concern.

Waterman asks us to consider the question, "What are the fundamental models in the legal domain and what constitutes a real "justification" of an answer?"

Of course, the legal community in its answer would surely mention the doctrine of precedent and the use of cases as fundamentals in justification and argumentation. The AI community would emphasize representation and models of process. Ashley, McCarty and Rissland all cite the important role of hypotheticals in justification and argumentation. Ashley and Rissland go further to also say that cases and hypotheticals can carry a large part of the burden of explanation as well. Thus Waterman's question can be re-posed to ask for the details of *how* cases are used in justification, argumentation, and explanation.

Dyer points out that this interleaving has methodological consequences: process models of legal expertise must be designed as a whole, where beliefs, goals, planning, and arguing are integrated.

7. The Law as a Learning System

On this issue there is some disagreement. Rissland can view the law, with all its cases, rules, statutes, etc., as a learning system in that the law responds to its environment which presents it cases to be dealt with and in response changes its case base, rules, statutes, etc., as well as its ways of dealing with its environment. This view is elaborated somewhat by Ashley:

Experience in the legal system is accumulated in two places: in its database of cases and in the implicit rules that develop for comparing fact situations in a particular domain. These rules are learned by law students in learning to "think like lawyers." They are evidenced by the choices that courts make in applying precedents but they are not authoritatively set forth in any restatement, statute or case. The rules make it possible to apply the precedents to new cases. They are how the law learns.

McCarty sees the law as a "theory construction" system that is constructing new concepts.

Dyer feels strongly that it is the lawyers who learn:

The law is NOT a learning system; lawyers ARE. Lawyers index legal episodes in episodic memory. As new situations arise, new judgements are made and episodes are reorganized in memory. This is just how the law student learns: by being presented with a case, applying whatever legal and common

sense intuitions are currently available, indexing the case along recognized abstract legal issues, and then modifying and reindexing as new cases (with alternate outcomes and opinions) are encountered.

8. Intelligent Aids for Practicing Litigators

The panel as a whole voiced somewhat cautious optimism that it would eventually be possible to develop a lawyer's workbench which would include tools ranging from standard retrieval tools like the existing WESTLAW and LEXIS full text retrieval systems, document generation aids, scheduling and calendar managers, to tools needing more intelligence like briefing assistants and interpretive analysis programs which could understand cases. This caution is based upon the nearly common experience of how long it takes to develop a program that can handle a few cases or problems, let alone the plethora occurring in real practice.

McCarty sees two main categories of intelligent applications programs: conceptual legal retrieval systems and legal analysis and planning systems. He elaborates:

"...the most critical task in the development of an intelligent legal information system, either for document retrieval or for expert advice, is the construction of a conceptual model of the relevant legal domain.

...these models will not be easy to formulate and the corresponding information systems will not be easy to construct. For the near term, then, the critical problem will be to select an appropriate level of conceptual detail, and an appropriate level of system complexity." [McCarty, 1983a, p. 286].

9. Implications and Relevance for Other Case-Based Domains

Most of the panelists did not express much enthusiasm for useful cross-fertilization between law and other domains, like medicine, except for Rissland who believes that there is something shared in that both domains use cases – albeit in different ways – especially in areas where the diseases are new or not well understood (i.e., where there might not be rules). All feel that the details of such cross-discipline enlightenment are vague at this point.

Dyer puts it:

I believe that modeling lawyers is a much more complex task than modeling doctors. This is mainly because the expert systems approach has restricted itself to those areas where immediate success is possible. If we admit that medicine is not simply rule-based diagnosis over symptoms with weighted and propagated certainty factors, then we open up the harder issues in medicine. These issues

are those shared with law: namely, how do doctors (lawyers) read and understand medical (legal) texts? What must doctors (lawyers) know about the physiological (socio-cultural) processes of the body (everyday affairs) in order to be able to understand, plan and give advice?

Such concerns are also shared by McCarty who points out that the shallow reasoning exhibited by certain medical diagnosis programs, like MYCIN, is not at all like, or sufficient for, the deep reasoning (for instance, with respect to its representation and causal models) required in the legal domain. McCarty recaps some of his recent comments from [McCarty, 1983a] where he asked, "If shallow rule-based systems have been so successful in the medical field, why do I insist that they have such serious limitations in the legal field?" McCarty's answer to this paradox lies in the differences of the nature of the rules involved. In medicine, the rules are empirical, associative, probabilistic rules of thumb, which are used cumulatively and which do not reflect any deep causal models, say of bacterial disease. Whereas in medicine, this might be enough to get the job done, in the law it is not except for certain discretionary legal issues like "reasonable care". Interestingly, rule-based treatment of discretionary judgements (e.g., the worthwhileness of a civil suit) has already met with some success [Waterman and Peterson, 1981].

10. Methodological and Other Issues

There is the usual two way relationship between AI and the law. The law is a field which might be able to profitably use AI concepts and technology to further its own aims such as developing jurisprudential analyses of legal reasoning or developing intelligent litigation aids and the law provides an excellent task domain for AI research.

Dyer addresses whether these two approaches entail different standards and whether this matters:

For a person not interested in modeling intelligence, but in supplying the lawyer with smart legal aids, the approach may be different. I believe that this particular subdiscipline for AI is too young to decide that one (or the other) approach should not be tried, but I believe that more is to be gained ultimately by modeling the mind in the domain of law than modeling law using currently known insights about intelligence (artificial or otherwise).

Gardner feels that it might be a difference more of degree than kind:

At one end, there are people trying to build systems useful to lawyers, or trying to find out how far present AI techniques are

applicable in law. The results here may look like lots of law and not much that's new for AI.

At the other end, for the purpose of studying some AI problem, you could create a microworld and call it law. The results might have as little to do with a lawyer's view of law as Winston's "Shakespeare world" has to do with an English professor's view of Shakespeare.

But the reason I said the distinction was a matter of degree was that, to some extent, we can't help working with microworlds. That is, every study of law or legal reasoning, AI or otherwise, has to work with an abstracted version of its domain. It's most important to try to be clear about what simplifications one is making, and to consider from what point of view these simplifications are reasonable.

Several panelists feel that progress in the area of AI and law requires further progress in fundamental AI and cognitive science research on issues like representation. This is a point emphasized increasingly by McCarty:

Successful work in law and AI — both theory and application — requires much better "conceptual models" of the legal domain than the field of AI can currently provide "off-the-shelf".

Part of McCarty's response to this problem is his work on "permissions and obligations" which addresses the fundamental "Hohfeldian" mechanisms of rights, privileges, duties, etc. underlying our legal system. Other work, like that of deBessonnet [1984] on legal representation primitives, also contributes to this basic research.

Other issues of concern to the panel included sociological issues, for instance, the impact of intelligent or (semi-)automated legal services and how this will effect society's relationship with the legal system (e.g., will legal services become accessible to more people, will people become more "do it yourself"). What will be the nature of the symbiosis between man and machine in the legal domain. Of course, there was also the issue of the law applied to AI programs, for instance, potential tort, or even criminal, actions involving AI programs (e.g., a medical malpractice suit against a medical expert system). Many of these issues are addressed in the panel on "AI and Social Responsibility", chaired by Maggie Boden, and reported on in these proceedings.

Bibliography

Ashley, K. D. (1984). "Reasoning by Analogy: A Survey of Selected A.I. Research with Implications for Legal Expert Systems". In *Proceedings First Annual Conference on Law and Technology*, University of Houston, Houston TX. To be included in a collection to be published by West Pub. Co. in 1985.

August, S. E. and Dyer, M. G. (1985). "Understanding Analogies in Editorials." UCLA AI Lab Memo.

Birnbaum, L. (1982). "Argument Molecules: A Functional Representation of Argument Structure". In *Proceedings AAAI-82*. Pittsburgh, PA.

deBessonnet, C. G. (1984). *CCLIPS, The Civil Code Legal Information Processing System*. Law and AI Project, Louisiana State Law Institute, Baton Rouge, LA.

Corbin, A. L. (1961). *The Judicial Process Revisited*. *Yale Law Journal* 71, 195-201.

Dyer, M. G., & Flowers, M. (1984). "Toward Automating Legal Expertise." In *Proceedings of the First Law and Technology Conference*. University of Houston, Houston, Texas.

Flowers, M. (1982). "On Being Contradictory". In *Proceedings of AAAI-82*. Pittsburgh, PA.

Flowers, M. and Dyer, M. G. (1984). "Really Arguing with your Computer in Natural Language." In *Proceedings of NCC-84*. Las Vegas.

Flowers, M., McGuire, R. and Birnbaum, L., (1982). "Adversary Arguments and the Logic of Personal Attacks." In Lehnert and Ringle (eds.) *Strategies for Natural Language Processing*. LEA Press, NJ.

Gardner, Anne v.d.L. (1984). *An Artificial Intelligence Approach to Legal Reasoning*. Doctoral Dissertation, Department of Computer Science, Stanford University.

—— (1983). "The Design of a Legal Analysis Program." In *Proceedings AAAI-83*. Washington, DC.

Goldman, S. R., Dyer, M. G. and Flowers, M., (1985). "Precedence-Based Legal Reasoning in Contract Law." UCLA AI Lab Memo.

Hohfeld, W. N. (1919). "Fundamental Legal Conceptions as Applied in Judicial Reasoning." Yale U. Press, New Haven, CT, Ed., W. W. Cook.

Kuhn, T. S. (1970). *The Structure of Scientific Revolutions*. Second Edition. University of Chicago Press.

- Lakatos, I. (1976). *Proofs and Refutations*. Cambridge University Press, London.
- Levi, E. H. (1949). *An Introduction to Legal Reasoning*. University of Chicago Press.
- Llwylyn, K. N. (1930). *The Bramble Bush*. Ocean Publications.
- _____. (1960). *The Common Law Tradition: Deciding Appeals*. Little Brown, Boston.
- McCarty, L. T. (1983a). "Intelligent Legal Information Systems: Problems and Prospects." *Rutgers Computer and Technology Law Journal* 9(2):265-294.
- _____. (1983b). "Permissions and Obligations." In *Proceedings IJCAI-83*. Karlsruhe, W. Germany.
- _____. (1977). "Reflections on TAXMAN: An Experiment in Artificial Intelligence and Legal Reasoning." *Harvard Law Review* 90:837-93.
- McCarty, L. T. and Sridharan, N. S. (1981a). "The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations." In *Proceedings IJCAI-81*.
- _____. (1981b). *A Computational Theory of Legal Argument*. Technical Report LRP-TR-13, Laboratory for Computer Science Research, Rutgers University.
- McGuire, R., Birnbaum, L., and Flowers, M. (1981). "Opportunistic Processing in Arguments." In *Proceedings IJCAI-81*. Vancouver, B.C.
- Polya, G. (1973). *Induction and Analogy in Mathematics*, Volumes 1 & 2. Princeton University Press, New Jersey.
- Peterson, M. A. (1983). *New Tools for Reducing Civil Litigation Expenses*. Memo R-3013-ICJ. The Institute for Civil Justice, The RAND Corporation, Santa Monica, CA.
- Reichman, R. (1981). "Modeling Informal Debates." In *Proceedings IJCAI-81*. Vancouver, B.C., August.
- Rissland, E. L. (1984a). "Argument Moves and Hypotheticals." In *Proceedings First Annual Conference on Law and Technology*. University of Houston, Houston TX. To be included in a collection to be published by West Pub. Co. in 1985.
- _____. (1981). *Example Generation*. Technical Report 81-24, Department of Computer and Information Science, University of Massachusetts, Amherst, MA.
- _____. (1982). "Examples in the Legal Domain: Hypotheticals in Contract Law." In *Proceedings Fourth Annual Cognitive Science Society Conference*. Univ. of Michigan, Ann Arbor.
- _____. (1983). "Examples in Legal Reasoning: Legal Hypotheticals." In *Proceedings IJCAI-83*. Karlsruhe, W. Germany.
- _____. (1984b). "Hypothetically Speaking: Experience and Reasoning in the Law". To appear in *Cognitive Science*. An earlier version appeared in *Proceedings First Annual Conference on Theoretical Issues in Conceptual Information Processing*, Georgia Tech, March.
- _____. (1984c). "The Ubiquitous Dialectic". In *Proceedings Sixth European Conference on Artificial Intelligence (ECAI-84)*, held in Pisa, Italy, September, published by Elsevier (North-Holland).
- _____. (1978). "Understanding Understanding Mathematics". *Cognitive Science*, Vol. 2, No. 4.
- Rissland, E. L., Valcarce, E. M., and Ashley, K. D. (1984). "Explaining and Arguing with Examples". In *Proceedings AAAI-84*, Austin, Texas.
- Waterman, D. A., and Peterson, M. A. (1981). *Models of Legal Decision Making*. Memo R-2717-ICJ. The Institute for Civil Justice, The RAND Corporation, Santa Monica, CA.