

Perspectives on Artificial Intelligence and Legal Reasoning

L. Thorne McCarty
Department of Computer Science
and
Faculty of Law
Rutgers University
New Brunswick, New Jersey 08903, USA
mccarty@cs.rutgers.edu

1 Introduction

Although most of the work on Artificial Intelligence and Law today is oriented towards the development of practical systems, there is a small group of researchers who are primarily interested in theoretical questions: How much of legal reasoning can be reduced to reasoning with rules? Is this rule-based component significant, or trivial? How is it possible to reason with cases at all? Are legal concepts just like ordinary common sense concepts, or do they have special characteristics? Is it possible to develop a computational theory of legal argument? The researchers who have investigated these questions include: Anne Gardner [Gar87]; Edwina Rissland and her students, Kevin Ashley [RA87, Ash88] and David Skalak [RS89, Ska89]; Michael Dyer and his student, Seth Goldman [GDF87]; Karl Branting [Bra89]; and Keith Belairs [Bel89]. In addition, researchers such as Richard Susskind [Sus87] and J.C. Smith [SD87], who have primarily built practical systems, have also been deeply concerned with the jurisprudential foundations of the field.

In this paper, I will describe my work on the *TAXMAN* Project, which for many years has pursued the goal of a computational theory of legal argument. More specifically, I am interested in *appellate* legal argument, rather than argument at trial, and I am trying to understand the role of *prototypical reasoning* in this process. Prototypes are ubiquitous in legal argument, as many researchers have noted [Ris85]: They appear as actual decided cases, which serve as precedents, and they appear as hypothetical cases, which serve to expose anomalies in an adversary's position. But what *theory* of legal argu-

ment would justify this prominent role for prototypical reasoning? Why, for example, does the citation of a hypothetical case have such a powerful effect?

In Section 2, below, I will discuss my early work with Sridharan on the *TAXMAN* Project, in which we attempted to answer some of these questions, and I will point out the difficulties we encountered. In Section 3, I will explain how my current research is intended to remedy these deficiencies. In Section 4, I will compare my approach to related work in the field.

2 Early Research on TAXMAN

More than ten years ago, Sridharan and I proposed a theory of legal reasoning in hard cases [McC80b, MS81, MS82]. We began by emphasizing the following three points, which should be familiar to most lawyers:

1. Legal concepts cannot be adequately represented by definitions that state necessary and sufficient conditions. Instead, legal concepts are incurably "open-textured".
2. Legal rules are not static, but dynamic. As they are applied to new situations, they are constantly modified to "fit" the new "facts". Thus the important process in legal reasoning is not theory application, but theory construction.
3. In this process of theory construction, there is no single "right answer". However, there are plausible arguments, of varying degrees of persuasiveness, for

each alternative version of the rule in each new factual situation.

The first of these points has been thoroughly discussed by Anne Gardner [Gar87], and seems to be generally accepted by researchers in AI and Law. The second point is less common, but it is related to the constructive approach to legal decisions proposed by Herbert Fiedler [Fie86] and Tom Gordon [Gor89], and to the rule-based representation of open-texture in law proposed by Trevor Bench-Capon and Marek Sergot [BCS87]. The third point, of course, has been thoroughly debated by legal philosophers for many years as part of the response to Ronald Dworkin's thesis [Dwo75]. Sridharan and I adopted this third point primarily as a methodological guideline: Since lawyers are more likely to agree on what counts as a plausible argument in a case than to agree on the appropriate outcome, we decided that it would be more fruitful to develop a theory of legal argument than to develop a theory of correct legal decisions.

This was the framework in which we worked. The specific theory we proposed was based on a representation of legal concepts by means of *prototypes and deformations*. Legal concepts have three components, we suggested: (1) an (optional) invariant component providing necessary conditions; (2) a set of exemplars providing sufficient conditions; and (3) a set of transformations that express various relationships among the exemplars. These three components are then refined further, for most concepts, so that one or more of the exemplars is designated as a prototype and the remaining exemplars are represented by a set of transformations, or deformations, of the prototypes. In this model, the transformations induce a partial order on the set of exemplars corresponding to the typicality gradient observed by psychologists in the study of human categorization [RL78, SM81], and the application of a concept to a new factual situation automatically modifies the definition of the concept itself, as required by Levi's classical account of legal reasoning [Lev42]. This was our response to the first two points noted above. In addition, in response to the third point, we were able to show that the arguments of lawyers and judges in a series of early corporate tax cases could be explained very well by the theory of prototypes and deformations. Our principal example was *Eisner v. Macomber*, 252 U.S. 189 (1920), an early stock dividend case, in which the arguments of Justice Pitney and Justice Brandeis took the form of a sequence of transformations from precedent cases through hypothetical cases to the factual situation of *Macomber*. It is important to note that these "explanations" were hand simulations. The *TAXMAN II* theory was partially implemented by Donna Nagel in

her thesis [Nag87], but a full implementation was never attempted.

I still believe that the *TAXMAN II* theory is qualitatively correct. But there were two major problems with our earlier work. First, the theory makes enormous demands on our knowledge representation language. To see this, it is sufficient to note that a transformation is a *syntactic* operation, and for such an operation to be meaningful it must correspond to the significant *semantic* relationships in the legal domain. However, the frame-based language in which we (re)implemented the *TAXMAN I* system [McC77, MS80, Sri81] did not have an adequate semantic foundation, and this meant that a full *TAXMAN II* implementation would have been entirely *ad hoc*. The second problem involves the theory of prototypes and deformations itself. What determines the choice of a prototype? What are the criteria for constructing transformations? It was clear that the set of transformations had to be tightly constrained, or else anything could be "transformed" into anything. But what was the source of these constraints? Much of my work since 1982 has been devoted to finding solutions to these two problems.

3 Current Research

My answer to the first problem identified above has been the development of a *Language for Legal Discourse*, or *LLD*, which is described in [McC89]. *LLD* has facilities for the representation of *states, events, actions*, and various modalities over actions such as *permission* and *obligation*. There are similar facilities in the *Event Calculus* of Kowalski and Sergot [KS86], although Kowalski seems to have taken a principled stance against the explicit representation of the deontic modalities [Kow89]. *LLD* also provides a systematic treatment of *sorts* and *subsorts* (e.g., an 'Actor' can be a 'Person' or a 'Corporation'), and it includes both *count terms* and *mass terms* (e.g., 'Person' is a count term and 'Stock' is a mass term). For both technical and philosophical reasons, the language is based on *intuitionistic logic* rather than classical logic. I have argued elsewhere that an intuitionistic semantics offers distinct advantages for a logic programming language [McC88a, McC88b], and these advantages are inherited by the action language and the deontic language in *LLD* [McC83, McC86b, MvdM92, McC94].

Why do I insist that *LLD* is a partial solution to the problems encountered in the *TAXMAN II* Project? First, it is no accident that the common sense categories embodied in the current version of the language are just

those categories that we need for an initial representation of corporate tax law: count terms, mass terms, states, events, actions, permissions, obligations. Other categories will surely be needed later, if we wish to develop a more sophisticated analysis of the tax code: purpose, intention, knowledge, belief are prime examples. More important than the substantive coverage of *LLD*, however, is the close correspondence between its surface syntax and its deep semantics. This correspondence is largely a result of the intuitionistic semantics of the language, and it addresses directly the first deficiency in our earlier work: Syntactic transformations now map directly onto significant semantic relationships.

The second problem is not completely solved by *LLD*, but the necessary tools are now available. I remarked above that the theory of prototypes and deformations requires a set of tight constraints on transformations. In our earlier papers [MS81, MS82], Sridharan and I noted that these constraints seem to be related to a sense of "conceptual coherence". But what does that mean? My conjecture now is that conceptual coherence can be explained, at least partially, by an analysis of the computational complexity of the inferences that we need to make in a language with the features of *LLD*. My research programme thus resembles, abstractly, the research programme of Marcus [Mar80] and Berwick [Ber85] for natural language grammars: Natural language grammars should be (i) easy to parse and (ii) easy to learn, Berwick suggests. Likewise, coherent concepts should be (i) *easy to compute with* and (ii) *easy to learn*. The problem, then, is to show how a representation of concepts using prototypes and deformations can have these properties.

So far, most of my work along these lines has been concerned with the proof theory for concepts represented by prototypes and deformations [McC93, MS94]. An overview appears in [MvdM91]. The technical idea is to construct definitions of concepts using only (intuitionistic) *definite rules*, and then to use *circumscription* [McC80a, McC86a] when we need to express indefinite information. To draw inferences in such a system, we construct a *prototypical proof*, which is complete but not sound for intuitionistic logic, and we achieve soundness by showing that the prototypical proof is preserved under the appropriate *transformations*. Intuitively, if the concept is "coherent" (in a certain context, and for a certain purpose), then it should be possible to compute the inferences that follow from the concept (in the specified context, and for the specified purpose) by applying only *local transformations* to prototypical proofs. Although much work remains to be done, I hope, in this way, to link some of the standard criteria for computational tractability to

the intuitive idea of conceptual coherence.

A second component of conceptual coherence, according to this research programme, is *learnability*. I have not yet pursued this approach in any detail, but there are hints about how to proceed in William Cohen's dissertation at Rutgers [Coh90a]. Cohen shows that an overly general theory (expressed in Horn clauses) can be specialized by inductive learning from examples, but only if the specialized theory takes certain restricted syntactic forms. Without syntactic restrictions, the specialized theory is not "*PAC-learnable*" according to Valiant's test [Val84]; with syntactic restrictions, Cohen shows that it is possible to learn, e.g., the best opening bids in the game of bridge [Coh90b]. It is interesting to note that one of Cohen's learnable classes (the class of "*k-complete prefixes*") can be viewed as a type of prototypical definition. There is thus some hope that the conditions for tractable prototypical proofs and the conditions for *PAC-learnable* concepts will coincide to some extent.

So far, in these studies of prototypical proofs and *PAC-learnable* concepts, the knowledge representation language has been restricted to the simple first-order non-modal case. In fact, the language in [MvdM91] and [Coh90a] is restricted to Horn clauses. It is a major challenge to extend these ideas to the full complexity of my *Language for Legal Discourse*. However, the complexity of *LLD* actually *strengthens* the preceding arguments about conceptual coherence, assuming it is possible to work out the technical details. Why is this so? As I pointed out in [McC89], deductive inference in a modal logic is notoriously difficult [Wal87, Ohl88], but prototypical proofs are relatively simple. We might thus expect the relationship between tractability and conceptual coherence to be especially pronounced in a language with the expressive power of *LLD*.

4 Related Work

In broad terms, I would summarize my approach to appellate legal argument as follows: The task for a lawyer or a judge in a "hard case" is to construct a *theory* of the disputed legal rules and legal concepts that produces the desired legal result, and then to persuade the relevant audience that this theory is preferable to any theories offered by an opponent. Empirically, legal theories seem to take the form of *prototypes and deformations*, and one important component of a persuasive argument is an appeal to the *coherence* of the theory thus constructed. Therefore, to obtain a deeper understanding of the phenomenon of legal argument, we need to explain, in com-

putational terms, why one theory constructed using prototypes and deformations is more (or less) coherent than another. These considerations motivate the research programme outlined in Section 3 above.

Viewed in this way, work by other AI researchers on legal argument can be seen as complementary to my own research: For example, Kevin Ashley's use of *dimensions* in *HYPO* [Ash88] resembles my use of deformations in *TAXMAN II*, but Ashley treats these dimensions as fixed at the time of argument rather than constructible. Karl Branting, in his work on *GREBE* [Bra89], integrates rule-based and case-based reasoning, and generates arguments about the open-textured predicates that appear in the legal rules, but these arguments are based on rough similarity metrics. Taken together, the work of Ashley and Branting seems likely to lead to useful techniques for organizing legal data bases, but not likely to lead to genuine insights into the nature of legal argument. Seth Goldman's work on *STARE* [GDF87] uses an *episodic memory* to index contract cases in a way that resembles my use of prototypes and deformations to index corporate tax cases, but Goldman emphasizes (in his published papers) the analysis of new cases rather than the construction of new arguments. Goldman's work also seems likely to lead to useful techniques for organizing legal data bases. Keith Bellairs, in his work on *BRAMBLE* [Bel89], proposes a sophisticated theory of analogical reasoning, in which the strength of an analogy, and hence the persuasiveness of an argument, is a function of the conceptual context in which the analogy occurs. Finally, Skalak and Rissland, in their work on *CABARET* [SR92], emphasize the intimate connection between cases and arguments, and show how different patterns of argument can influence the selection of different supporting cases. Philosophically, this latter work is probably the most congruent with my own approach, but it lacks a deep conceptual model of the legal domain (the IRS home office deduction). And without such a model, I have argued, it is not entirely clear what the computational approach contributes to the theory, over and above the usual verbal formulation.

Recently, several articles have appeared in the American law review literature that are consistent with the research programme outlined in Section 3. Most closely related are a series of papers by Steven Winter [Win88, Win89], who specifically applies George Lakoff's theory of prototypes [Lak87] to a variety of legal issues. Winter's general approach is explained in [Win89], and the law of "standing" is analyzed and criticized within this framework in [Win88]. Also closely related is Clark Cunningham's linguistic analysis of "search" under the Fourth

Amendment [Cun88]. Both of these authors emphasize the role of coherence in legal argument, although their examples are primarily negative, that is, they use linguistic and cognitive theories to show that "standing" and "search" as expounded by the courts are *incoherent* concepts. Even the Critical Legal Studies movement, which often seems interested only in "trashing" legal doctrine, has explained some of its positions in cognitive terms: See, for example, Mark Kelman's concluding chapter in [Kel87], entitled "Toward a Cognitive Theory of Legitimation."

There are hints, in these papers and elsewhere, that the next successor to "Law and Sociology", "Law and Economics", "Law and Literature", etc., will be a field called: "Law and Cognition". If so, the work of AI researchers will be in demand. We will be needed to keep this new field on a rigorous (and computationally sound) path.

References

- [Ash88] K.D. Ashley. *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. PhD thesis, University of Massachusetts, 1988.
- [BCS87] T.J.M. Bench-Capon and M.J. Sergot. Towards a rule-based representation of open texture in law. In C. Walter, editor, *Computing Power and Legal Language*. Greenwood/Quorum Press, 1987.
- [Bel89] K. Bellairs. *Contextual Relevance in Analogical Reasoning: A Model of Legal Argument*. PhD thesis, University of Minnesota, 1989.
- [Ber85] R.C. Berwick. *The Acquisition of Syntactic Knowledge*. MIT Press, 1985.
- [Bra89] L.K. Branting. Representing and reusing explanations of legal precedents. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 103-110. ACM Press, June 1989.
- [Coh90a] W.W. Cohen. *Concept Learning Using Explanation Based Generalization as an Abstraction Method*. PhD thesis, Rutgers University, 1990.
- [Coh90b] W.W. Cohen. Learning from textbook knowledge: A case study. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 743-748, 1990.

- [Cun88] C.D. Cunningham. A linguistic analysis of the meanings of "search" in the fourth amendment: A search for common sense. *Iowa Law Review*, 73:541-609, 1988.
- [Dwo75] R. Dworkin. Hard cases. *Harvard Law Review*, 88:1057-1109, 1975.
- [Fie86] H. Fiedler. Expert systems as a tool for drafting legal decisions. In A.A. Martino and F. Socci Natali, editors, *Automated Analysis of Legal Texts: Logic, Informatics, Law*, pages 607-612. Elsevier North-Holland, 1986.
- [Gar87] A.v.d.L. Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. MIT Press, 1987.
- [GDF87] S.R. Goldman, M.G. Dyer, and M. Flowers. Precedent-based legal reasoning and knowledge acquisition in contract law: A process model. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, pages 210-221. ACM Press, May 1987.
- [Gor89] T.F. Gordon. Issue spotting in a system for searching interpretation spaces. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 157-164. ACM Press, June 1989.
- [Kel87] M. Kelman. *A Guide to Critical Legal Studies*. Harvard University Press, 1987.
- [Kow89] R.A. Kowalski. The treatment of negation in logic programs for representing legislation. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 11-15. ACM Press, June 1989.
- [KS86] R.A. Kowalski and M.J. Sergot. A logic-based calculus of events. *New Generation Computing*, 4(1):67-95, 1986.
- [Lak87] G. Lakoff. *Women, Fire, and Dangerous Things*. University of Chicago Press, 1987.
- [Lev42] E.H. Levi. *An Introduction to Legal Reasoning*. University of Chicago Press, 1942.
- [Mar80] M.P. Marcus. *A Theory of Syntactic Recognition for Natural Language*. MIT Press, 1980.
- [McC77] L.T. McCarty. Reflections on TAXMAN: An experiment in artificial intelligence and legal reasoning. *Harvard Law Review*, 90:837-93, 1977.
- [McC80a] J. McCarthy. Circumscription: A form of non-monotonic reasoning. *Artificial Intelligence*, 13:27-39, 1980.
- [McC80b] L.T. McCarty. The TAXMAN project: Towards a cognitive theory of legal argument. In B. Niblett, editor, *Computer Science and Law: An Advanced Course*, pages 23-43. Cambridge University Press, 1980.
- [McC83] L.T. McCarty. Permissions and obligations. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 287-294, 1983.
- [McC86a] J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28:89-116, 1986.
- [McC86b] L.T. McCarty. Permissions and obligations: An informal introduction. In A.A. Martino and F. Socci Natali, editors, *Automated Analysis of Legal Texts: Logic, Informatics, Law*, pages 307-337. Elsevier North-Holland, 1986. Also available as Rutgers Technical Report LRP-TR-19.
- [McC88a] L.T. McCarty. Clausal intuitionistic logic. I. Fixed-point semantics. *Journal of Logic Programming*, 5(1):1-31, 1988.
- [McC88b] L.T. McCarty. Clausal intuitionistic logic. II. Tableau proof procedures. *Journal of Logic Programming*, 5(2):93-132, 1988.
- [McC89] L.T. McCarty. A language for legal discourse. I. Basic features. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 180-189. ACM Press, June 1989.
- [McC93] L.T. McCarty. Proving inductive properties of PROLOG programs in second-order intuitionistic logic. In *Proceedings, Tenth International Conference on Logic Programming*, pages 44-63. MIT Press, 1993.
- [McC94] L.T. McCarty. Modalities over actions. In *Principles of Knowledge Representation and*

- Reasoning: Proceedings of the Fourth International Conference (KR94)*, pages 437–448. Morgan Kaufmann, 1994.
- [MS80] L.T. McCarty and N.S. Sridharan. The representation of an evolving system of legal concepts: I. Logical templates. In *Proceedings of the Third Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, pages 304–311, Victoria, British Columbia, 1980.
- [MS81] L.T. McCarty and N.S. Sridharan. The representation of an evolving system of legal concepts: II. Prototypes and deformations. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 246–53, 1981.
- [MS82] L.T. McCarty and N.S. Sridharan. A computational theory of legal argument. Technical Report LRP-TR-13, Computer Science Department, Rutgers University, 1982.
- [MS94] L.T. McCarty and L.A. Shklar. A PROLOG interpreter for first-order intuitionistic logic (abstract). In *Proceedings, 1994 International Logic Programming Symposium*, page 685. MIT Press, 1994.
- [MvdM91] L.T. McCarty and R. van der Meyden. Indefinite reasoning with definite rules. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, pages 890–896, 1991.
- [MvdM92] L.T. McCarty and R. van der Meyden. Reasoning about indefinite actions. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR92)*, pages 59–70. Morgan Kaufmann, 1992.
- [Nag87] D.J. Nagel. *Learning Concepts with a Prototype-Based Model for Concept Representation*. PhD thesis, Rutgers University, 1987.
- [Ohl88] H.J. Ohlbach. A resolution calculus for modal logics. In *Proceedings, Ninth International Conference on Automated Deduction*, pages 500–516, 1988.
- [RA87] E.L. Rissland and K.D. Ashley. A case-based system for trade secrets law. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, pages 60–66. ACM Press, May 1987.
- [Ris85] E.L. Rissland. AI and legal reasoning. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1254–1260, 1985. Report of a panel consisting of Rissland (chair), Ashley, Dyer, Gardner, McCarty and Waterman.
- [RL78] E. Rosch and B.B. Lloyd. *Cognition and Categorization*. Lawrence Erlbaum Associates, 1978.
- [RS89] E.L. Rissland and D.B. Skalak. Interpreting statutory predicates. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 46–53. ACM Press, June 1989.
- [SD87] J.C. Smith and C. Deedman. The application of expert systems technology to case-based law. In *Proceedings of the First International Conference on Artificial Intelligence and Law*, pages 84–93. ACM Press, May 1987.
- [Ska89] D.B. Skalak. Taking advantage of models for legal classification. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, pages 234–241. ACM Press, June 1989.
- [SM81] E. Smith and D. Medin. *Categories and Concepts*. Harvard University Press, 1981.
- [SR92] D.B. Skalak and E.L. Rissland. Arguments and cases: An inevitable intertwining. *Artificial Intelligence and Law*, 1:3–44, 1992.
- [Sri81] N.S. Sridharan. Representing knowledge in AIMDS. *Informatica e Diritto*, 7:201–221, 1981.
- [Sus87] R.E. Susskind. *Expert Systems in Law: A Jurisprudential Inquiry*. Oxford University Press, 1987.
- [Val84] L.G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.
- [Wal87] L.A. Wallen. Matrix proof methods for modal logics. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 917–923, 1987.

- [Win88] S.L. Winter. The metaphor of standing and the problem of self-governance. *Stanford Law Review*, 40:1371-1516, 1988.
- [Win89] S.L. Winter. Transcendental nonsense, metaphoric reasoning and the cognitive stakes for law. *University of Pennsylvania Law Review*, 137:1105-1237, 1989.