Reasoning about "Hard" Cases in Talmudic Law

Steven S. Weiner
Massachusetts Institute of Technology
Dept. of Electrical Engineering and Computer Science
Harvard Law School

Introduction

Anne Gardner [1] has drawn an important distinction between "hard" and "easy" cases of law. Easy cases are defined as situations whose verdict would not be disputed by knowledgeable, rational lawyers; hard cases, on the other hand, pose questions about whose answers lawyers might intelligently disagree. Gardner posits that reasoning about easy cases generally involves combining legal predicates which are clearly defined in terms of each other and ultimately in terms of the descriptors given for the case. While "easy" should not be confused with trivial or routine, researchers have made considerable progress in modelling easy-case reasoning, by employing mechanisms such as Modus Ponens chaining of "if-then" rules [1,2], frame-based representations of some commonsense knowledge [1,2], simple deductive use of precedents [2], and Gardner's own use of augmented transition networks (ATN's) for temporal modelling in offer and acceptance cases [1]. Analyzing hard cases, suggests Gardner, typically requires the application of open-textured legal predicates: predicates whose scopes are not precisely defined in terms of case descriptors, but are instead defined by rules that have some latitude and by a set of positive and negative specific examples. Because straightforward deduction on the facts of a hard case is not possible, using precedents to reason by induction and analogy becomes critical. Research in this difficult area of cognitive modelling has progressed more slowly, but several interesting ideas about case-based reasoning have nevertheless surfaced, including: organizing precedents into discrimination trees [3], explanation-based generalization [4], and refining concepts using hypothesis and experiment [5].

Randall Davis [6] has discussed how "depth" is an important property that may be used to classify expert systems, and intelligent programs in general. The depth of a system, as Davis explains, is the extent to which its programs contain not only rules and algorithms for mapping conclusions onto input scenarios, but also a representation of the *underlying causes* linking facts and consequents. Davis argues that causal models and not just empirical relations are essential if a program is to have any chance of recovering gracefully and learning from error on the inevitable occasions when its hard-coded heuristics yield faulty predictions. On the other hand, building such depth into a system can be quite a daunting challenge, since it typically demands of the would-be system designer both a clear causal model of the domain as well as a practical control scheme for when to reason carefully from the underlying model and when to simply apply more shallow rules. Work to date on automating case-based reasoning has spanned a wide range of the "depth" spectrum, from heavily inductive jurimetrics work and the chemistry program Meta-Dendral [7] to the cause-dominated models of DeJong[5] and Doyle[8]. In a subtle domain like law, the payoff of depth's robustness as well as the difficulty of codifying useful and deep models of legal principles both loom large.

In Talmud study one constantly works with "hard" cases. A typical Talmudic exercise is to consider a set of seemingly inconsistent rules and/or cases and to carefully refine the rule definitions so as to resolve the conflicts. Subtle, creative analysis of predicates and precedents is necessary to generate plausible rule definitions that satisfy all of the given data. As we shall see, the Talmud student frequently comes across examples of both "deep" and "shallow" reasoning applied to hard cases. There thus appears to be an intriguing analog in human legal systems to Davis' spectrum for automated systems. This paper will present several examples of Talmudic reasoning to illustrate this depth spectrum in Talmudic scholarship and in related AI paradigms. Working with ideas based on Levi's exposition on the American legal system [9], the paper attempts to rationalize the basis for this diversity of approaches within Talmudic law, and concludes with the implications of this human spectrum for contemporary Artificial Intelligence research.

Hard Cases from the Talmud

In his extensive, thirteenth century glosses on the Babylonian Talmud [10], Rabbi Asher ben Yechiel comes to grips with the confusing, "hard" Talmudic issue of liability for indirect damage. Rabbi Asher cites and summarizes an assortment of accepted precedents representing positive and negative instances of liability (see Figure 1), and tries to formulate a liability rule that satisfies this data.

Figure 1

case	direct?	immediate?	definite?	liable?
destroys friend's IOU note	yes	yes	yes	yes
reveals friend's property to thieves	yes	yes	yes	yes
judge errs negligently	yes	yes	yes	yes
mingles crops in friend's field, rendering them forbidden	yes	yes	yes	yes
opens gate releasing friend's animal, which goes out and dies	no	no	no	no
gives a lit torch to a child, who burns a field	no	no	no	no
sics dog on friend	yes	no	no	no
places poison before friend's beast	no	yes	no	no
removes cushions from underneath falling person or property	no	yes	yes	no

Rabbi Asher's approach is to first circumscribe the space of potential rules by assuming that for the purposes of distinguishing these cases, liability is a positive function of three binary variables: the immediacy, definiteness, and directness of the resultant damage relative to the questionable act. He then considers the values that each of his accepted cases has for these three attributes, and by looking at the verdict in each instance seeks to regressively derive the attribute values which must be necessary for liability. This regressive reasoning process is illustrated using "version space" terminology [4] in Figure 2. Version spaces involve the use of positive and negative instances to gradually prune a tree or graph-like space of possible attribute-verdict relations. Granting Rabbi Asher's assumptions that these three attributes are indeed the relevant discriminators for these cases, and that the attribute values are as he claims they are, it follows from version space analysis that there are only two tenable rules which satisfy all the given data: namely, either liability requires positive values for all three attributes, or else the pair "directness and definiteness" is a sufficient condition for liability (Rabbi Asher favored the latter rule).

Figure 2

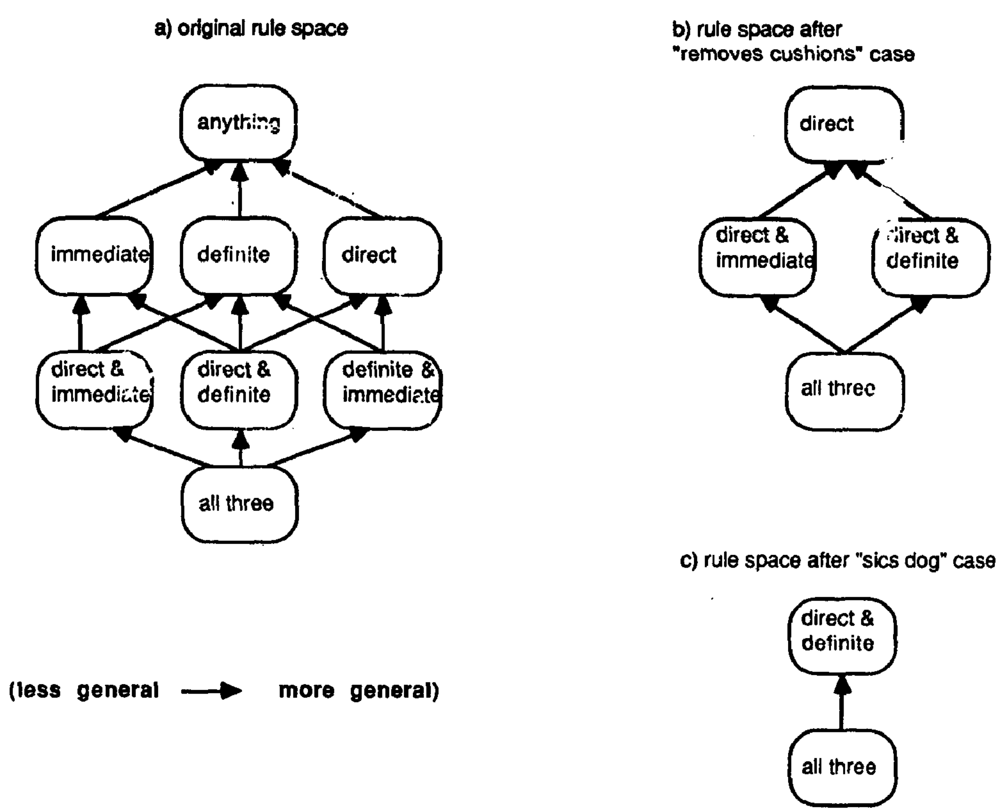

Rabbi Asher's case-based approach to the hard question of liability for indirect damage bears strong resemblance not only to Mitchell's version spaces technique but also to Schank and Kolodner's method of "discrimination trees" [3]. Discrimination trees organize precedents like the leaves of a tree, where each internal node of the tree represents a branch point on a potentially significant attribute. When a convincingly large number of leaves under a particular node are found to all share some additional common quality, the attributes associated with that node are inductively said to "predict" the common quality. For Rabbi Asher, the discriminating attributes are directness, definiteness, and immediacy, while the common quality for which predictors are sought is liability.

The critical point that I wish to emphasize is the nature of Rabbi Asher's approach vis-a-vis Davis' "depth" issue. Rabbi Asher presents no explicit discussion of the causality relationship between his attributes and liability, other than to assume that immediacy, definiteness, and directness - and no other attributes of the cases in point - may each influence liability in a positive way, if at all. However, no precise theory is suggested as to why any particular combination of those attributes ought to be a cause of liability. Rabbi Asher's thought process seems far more inductive, less cause-driven: he knows that some line must be drawn, and he feels that his three attributes are sensible variables to work with, so his next step is to best-fit the data. A similar state of affairs appears to prevail in the programs of Schank and in Mitchell's original LEX, in which all causal knowledge is embedded in the actual selection of potentially relevant attributes and attribute groups, and relatively shallow induction drives the subsequent learning process. (Mitchell himself lamented this shallowness in [4], and journeyed far to the other extreme with LEX2 and its explanation-based generalization which resembles theorem proving.) By way of contrast, in the next section we demonstrate that sometimes legal analysis is conducted with much deeper, cause-dominated reasoning.

Rabbi Chaim Soloveitchik of Brisk was a turn of the century Talmud scholar renowned for his exceptionally probing, analytical approach to case-based reasoning. The following example of his methodology sharply demonstrates the key role that "deep" induction can play in at least one scholar's vision of case-based reasoning.

The Talmud recognizes two basic property classes - real estate and movables. The two classes differ with respect to many property rules; in particular, real estate "can not be stolen" - i.e., wrongful occupation of land is viewed as temporary trespass and the land does not attain the legal status of stolen property. The Talmud prescribes special rules for stolen property, making the thief responsible for any damage (even accidental) to the goods, and nullifying the original owner's title to the stolen object if the object is substantially changed (thus protecting innocent buyers; the victim is still entitled to payment from the thief, of course). The Talmud's property laws cover slaves, and there is a general property principle that slaves are as real estate. Accordingly, it is not surprising to find Maimonides [11] writing that an original slave owner retains his title even if the slave is stolen and undergoes a change (like age or injury) which would be sufficient to nullify original title to a stolen cow. What comes as a shock in the light of all this is Maimonides' ruling that a thief who steals a slave does become responsible for accidental injuries, as is true of movables. How can Maimonides' paradoxical classification of slaves as part real estate and part movables be explained and justified?

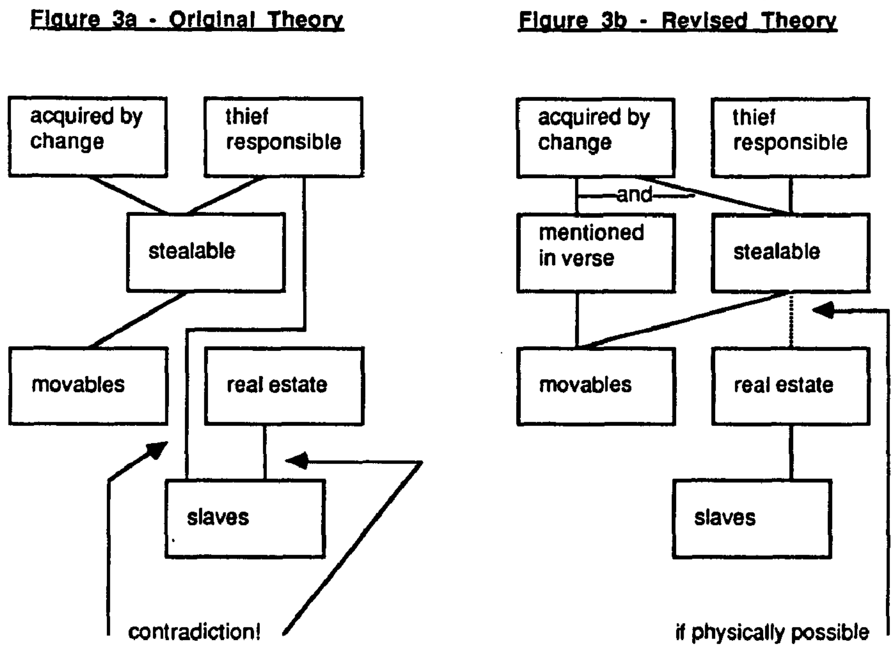

Rabbi Chaim of Brisk, in his commentary to Maimonides' code, suggests that Maimonides may well have seen the laws of thief responsibility and title nullification as emerging not from a single predicate "stolen property", but rather from two separate predicates. That is, "stolen property" is one predicate of which a consequence is thief responsibility, and "acquirable through change" is a distinct predicate for which "stolen property" is a necessary but not sufficient condition. Rabbi Chaim notes that the exemption of real estate from title nullification through change is explicitly derived by the Talmud from exegesis of a Biblical text. Perhaps that text excludes real estate only from the specific predicate of "acquirable through change" (possibly because naive buyers are not as likely for real estate as for movables), and that the exclusion of real estate from thief responsibility is caused not by fundamental exclusion from "stolen property" but rather because land cannot physically be appropriated but only trespassed upon.

Accordingly, since slaves are physically "movable", there would be no reason to exempt thieves from damage responsibility; since slaves are legally real estate, the exemption from title nullification would apply to them as well - precisely as Maimonides has written. To support this theory of separate causal models for "acquirable through change" and "thief responsibility", Rabbi Chaim's students identify precedents indicating that if means for physically stealing land can be devised - eg., moving boundary markers - then properties of stolen goods like "thief responsibility" do apply.

Rabbi Chaim's approach to the stolen slave problem is fairly characteristic of his general methodology: when a set of precedents appears to present contradictory evidence about the scope of some predicate, create a new predicate by proposing a new causal model for the verdicts involved, so that conflicting precedents may be associated with separate predicates. Finally, confirm the new theory by finding cases which lie within the scope of just one of the new predicates, and see if those cases exhibit the desired "split" behavior. This methodology has some similarity to research by DeJong [5]. In DeJong's work, a computer with a simplified model of laboratory chemistry gradually learns to refine concepts like permeability by starting with a basic theory - "liquids leave through holes or evaporation" - and proposing modifications - "or through membranes" - after seeing examples where solutions separated by membrane exhibited osmosis. DeJong's system tests its new theory of membranes by asking to see experiments where larger cross sections of membrane are inserted, and confirms its theory of permeable membranes causing osmosis by observing that the rate of osmosis increases. Related research has also been reported by Doyle [8], seeking to automate the learning of simple, causal models for devices like toasters and sinks. In Doyle's work, specific observations of things like toaster settings and bread shades are combined with simplified principles about how appliance components generally interact, to hypothesize functional dependencies that satisfy and explain all observed data. As in DeJong's system, observing new data inconsistent with current theories causes an attempted modification of beliefs such that all data are satisfied but that a plausible, causal theory remains. DeJong, Doyle, and Rabbi Chaim are thus good evidence for Davis' claim that paying attention to underlying causes is a critical factor in being able to recover from surprising examples with plausibly revised theories.

Analysis and Conclusions

In his discussion of the case-based reasoning process associated with American common law, Edward Levi writes [9]:

> Thus it cannot be said that the legal process is the application of known rules to diverse facts. Yet it is a system of rules; the rules are discovered in the process of determining similarity or difference... The problem for the law is: When will it be

> just to treat different cases as though they were the same? A working legal system must therefore be willing to pick out key similarities and to reason from them to the justice of applying a common classification... It could be suggested that reasoning is not involved at all; that is, that no new insight is arrived at through a comparison of cases. But reasoning appears to be involved... It seems better to say there is reasoning, but it is imperfect.

Levi, in probing the nature of hard-case reasoning, can not accept the idea that case law is purely arbitrary selection by each judge of some induction terms and an uninspired, shallow, best-fitting of some random data. On the other hand, Levi acknowledges that decisions are not always backed by a solid, logical theory explaining all precedents. This" imperfection" - or, as I would say, shallowness - in case law is not a "bug" according to Levi, but rather a "feature":

> Not only do new situations arise, but in addition peoples' wants change. The categories used in the legal process must be left ambiguous in order to permit the infusion of new ideas... Reasoning by example shows the decisive role which the common ideas of society and the distinctions made by experts can have in shaping the law.

In other words, the price to pay for requiring a solid, underlying theory is lack of flexibility in decision making. Finding superficial distinctions and drawing a new line to best-fit data is not difficult; developing a new, complete, consistent physics is. If a judge needed to devise in real time a creative, complete theory, Rabbi Chaim style, every time he felt impelled to classify a new case outside its most obvious grouping, more often than not the judge's hands would be tied.

It is interesting to note that Rabbi Chaim's reputation as a scholar is much greater within the Talmudic community than is his reputation as an actual judge. Instructive is the history of a theory which emerged from Rabbi Chaim's school concerning observance of the Passover holiday's Seder ritual. Rabbi Chaim and his students [12] presented a very ingenious and plausible explanation of a long-standing dispute between authorities relating to the drinking of wine during the Seder - an explanation which resolved a number of difficult passages in Talmudic lore but which carried as a logical consequence a fairly new and strict ruling forbidding substitute bevarages (such as grape juice) for that particular ritual. In the early twentieth century Europe of Rabbi Chaim's day, authoratative codifiers (see [13]) continued to permit beverage substitutions when necessary, relying on a relatively shallow (compared to Rabbi Chaim's theory) comparison to other ritual ceremonies. Athough Rabbi Chaim's approach elegantly explained a number of Talmudic puzzles, better theory evidently does not always mean better law. In more affluent, late twentieth century America, Rabbi Chaim's theory enjoys a new popularity. It seems reasonable to suggest that where the actual practice of law is concerned - as opposed to pure scholarship - adherence to a more flexible, if more shallow, legal reasoning process may sometimes be preferable, in the face of factors like the economics of wine.

We have seen that balancing flexibility and rigorous theory is quite in the spirit of Talmudic law as well as of other legal systems. The coexistence of these two approaches within human legal thought lends encouragement to the continuance of research along both "deep" and "shallow" fronts in the Artificial Intelligence world. The pragmatic, judicial value of shallow induction - flexibility - suggests further research on adaptively modifying inductive vocabulary, within those mechanical systems that pursue the practical, "shallow" approach.

References

[1] Gardner, A., "Overview of an AI Approach to Legal Reasoning". from *Computing Power and Legal Reasoning*, edited by C. Waters, West Publishing Co., 1985.

[2] Meldman, J., "A Preliminary Study in Computer-Aided Legal Analysis", TR-157, MIT Project MAC ,1975.

[3] Kolodner, J., "Maintaining Organization in a Dynamic Long-Term Memory", from *Cognitive Science*, vol. 7, 1983.

[4] Mitchell, T., "Learning and Problem Solving", from *Proceedings of IJCAI*, 1983.

[5] DeJong, G., "Towards a Model of Conceptual Knowledge Acquisition Through Directed Experimentation", from *Proceedings of IJCAI*, 1985.

[6] Davis, "Expert Systems: Where Are We? And Where Do We Go from Here?", AI Memo 665, MIT AI Lab, 1982.

[7] Buchanan, B., and Feigenbaum, E., "Dendral and Meta-Dendral: Their Applications Dimensions", from *Artificial Intelligence*, vol. 11, 1978.

[8] Doyle, R., "Hypothesizing and Refining Causal Models", AI Memo 811, MIT AI Lab, 1984.

[9] Levi, E., "An Introduction to Legal Reasoning", pp. 3-5, University of Chicago Press, 1949.

[10] Glosses of Rabbi Asher ben Yechiel to the Babylonian Talmud, Tractate Baba Batra, chapter 2.

[11] Maimonides' Code, Laws of Theft and Loss, Chapter 9.

[12] Novellae of Rabbi Yitzhak Zev, commentary to Maimonides' Code, Laws of Passover, Chapter 7.

[13] Mishna Berurah, Vol. Orach Chaim, Chapter 472, note 37.