

Rules and Cases for Legal Reasoning: Notes on Some Neglected Aspects

Anne v.d.L. Gardner

286 Selby Lane, Atherton, California 94027
gardner@cs.stanford.edu

Abstract

Rules and cases are essential elements in legal reasoning, but computational models have barely begun to reflect the complexities of their roles. Based on experience with a real case, this paper identifies four areas that deserve attention from anyone concerned with understanding the processes of a general legal reasoner. These are (1) combining rules that were adopted for differing purposes but that all have application to the problem at hand; (2) allowing for argument over the logical structure of rules, and managing to reason with them even when unsure what the logical structure is; (3) allowing cases to be used mainly for their facts and outcome, mainly for their reasoning, or mainly for the rules they lay down, and employing each technique when appropriate; and (4) extending the legal sources that are treated as cases. The paper does not propose solutions but merely attempts, by way of examples, to suggest significant research areas.

Introduction

It is widely agreed that a realistic computational model of legal reasoning must use both rules and cases. Working out how to combine them has been the subject of a number of programs, including GREBE (Branting 1991a, 1991b, 1994), CABARET (Rissland and Skalak 1991; Skalak and Rissland 1992), TAXMAN II (McCarty and Sridharan 1981; McCarty 1989, 1995), and my own dissertation project (Gardner 1987).

The choice of directions for developing this work is goal-dependent. One may want mainly to build tools useful for lawyers; to understand human cognitive processing; or to understand the forms of legal analysis, argument, and decision-making. Toward the last goal at least, it is important to observe what moves take place in actual legal reasoning. Our models so far have mostly been based on abstractions—inspired, for instance, by the descriptions in standard works of jurisprudence like Levi (1949) and Hart (1961). According to my recent experience with a large, rather technical, real-life case, the proportion of the reasoning that existing computational models can account for is disappointingly small. This is not only for the expected reasons, such as inability to handle general natural language and commonsense knowledge (on the AI side) or arguments from purpose and from principle (on the legal side). There are also less familiar features, some funda-

mental. This paper identifies a few features, of varying importance, that are fairly easy to pick out and describe.

Background: The *Alaska* Case

United States v. Alaska was a lawsuit over the ownership of lands just off Alaska's north coast. The areas in dispute are potentially valuable for oil; the government that owns them is the government that gets to decide whether to open the lands for offshore oil exploration and, if opened, to enter leases with oil companies and collect royalties from them.

Being a suit between a state and the federal government, the *Alaska* case was initially filed in the Supreme Court as an original jurisdiction case under Article III, sec. 2, clause 2, of the Constitution. The Court, as it often does in such cases, appointed a Special Master to hear the case and report back to it. The Master's report (565 pages, covering six main groups of questions) was submitted in March 1996. The Supreme Court, after hearing oral argument by the parties on the parts of the report to which they took exception, issued its decision in June 1997 (117 S. Ct. 1888). Alaska's three exceptions were overruled; the United States' one exception was sustained. My work in all this was with the Special Master, Professor J. Keith Mann of the Stanford Law School.

Combining Rule Sets

Programs that work with statutes usually deal with the terms of a single enactment, for example the British Nationality Act (Sergot et al. 1986) or the Latent Damage Act (Susskind 1989). When cases are used as well, the rule-based part may be limited to a single statutory section, such as the Internal Revenue Code section on home office deductions (Rissland and Skalak 1991).

In *United States v. Alaska*, the range of rules was much broader. The basic statute was the 1953 Submerged Lands Act (43 U.S.C. §§ 1301–1315), which says that each state owns the submerged lands in a three-mile belt measured outward from its coastline. Secondary was a 1958 treaty, the Convention on the Territorial Sea and the Contiguous Zone (15 U.S.T. 1606), which defines the baselines from which nations are to measure their territorial seas. Linking the two was a Supreme Court decision saying, roughly, that

AAAI Spring Symposium Series, 1998:
Multimodal Reasoning

“coastline” in the Submerged Lands Act should be interpreted to mean the same thing as “baseline” in the Convention. (*United States v. California*, 381 U.S. 139 (1965).) So far so good. It is not hard to imagine a rule base that includes both the Submerged Lands Act and the Convention. A bit more detail is given in Gardner (1989).

But of the four main sections of the Master’s report that applied these rules—each to a different geographical feature that might or might not form part of the coastline—only two were able to use just the rules mentioned and the related cases. The other two sections involved possible exceptions to the rules, that is, arguments that the usual definitions should not be applied in this particular situation. These arguments were based on other rules entirely.

For example, one part of the case involved a mile-long pier built out into the ocean. Is the three-mile limit to be measured from the mainland or from the end of the pier? Under the Convention, the Master’s report found, the end of the pier should be used. The United States, arguing for an exception, invoked sources including the statute authorizing construction of the Trans-Alaska pipeline (43 U.S.C. §§ 1651–1655), the statute giving the Army Corps of Engineers authority over structures built in navigable waters (33 U.S.C. § 403), regulations issued by the Army Corps of Engineers (33 C.F.R. § 209.120 (1976)), and some judge-made rules about what happens when government employees fail to follow government regulations.

This example is not extraordinary. Altogether the Master’s report contains citations to over forty different statutes. It could not have been known in advance just what statutes should form part of the rule base.

The question raised here is not just how to make a broader rule base available to programs. What reasoning is involved in figuring out how disparate rule sets fit together? Is it true that introducing a new rule set is usually associated with arguing for an exception to the main rule? If so, does this tell us anything new about defeasible reasoning?

Interpretation of Rules

There are some important operations in legal reasoning that are not covered by the general concepts of applying rules and analogizing cases. One is the interpretation of a rule, in the sense of reasoning about the meaning of a universally quantified proposition. Such reasoning may be needed as a step separate from reasoning about whether the antecedents of a rule are satisfied by the particulars of a case at hand. One problem of rule interpretation, familiar from the work of Allen and Saxon (e.g., 1987, 1991), arises from ambiguity in the natural-language counterparts of logical operators.

The need for rule interpretation goes further, however; it may involve clarifying the relations among domain con-

cepts. Some examples come from the Convention’s definition of a bay. (Where a bay is found, a line drawn across its mouth counts as part of the coastline.) The definition contains two sentences, with the first stating some general conditions and the second imposing a geometrical test based on the area of a semicircle:

Article 7(2). For the purpose of these articles, a bay is a well-marked indentation whose penetration is in such proportion to the width of its mouth as to contain landlocked waters and constitute more than a mere curvature of the coast. An indentation shall not, however, be regarded as a bay unless its area is as large as, or larger than, that of the semi-circle whose diameter is a line drawn across the mouth of that indentation.

The interpretation problems do not come from the rule-plus-exception structure of the English, which can uncontroversially be flattened to

*if general-conditions(x) and semicircle-test(x)
then bay(x) .*

Rather, the problems in the *Alaska* case were (1) does the semicircle test subsume the general conditions? and (2) if not, what is the logical structure of the general conditions? Dependent on the answers was the ownership of lands under a water body called southern Harrison Bay. It was agreed that southern Harrison Bay met the semicircle test.

Alaska’s first argument was, in effect, that

if semicircle-test(x) then general-conditions(x) ,

thus reducing the rule to

if semicircle-test(x) then bay(x) .

This is a startling position because it has no warrant in the syntax of the English rule; but it is not preposterous, because scholars have raised the same question of interpretation (see Report, pp. 182–83). In Alaska’s view, the interpretation was warranted by the drafting history of the definition. The Master’s report, after reviewing the history, disagreed.

The lesson here is that even where a rule has an authoritative text with no surface structural ambiguity, programs still need to leave room for argument over what logical expression correctly translates the rule. Had the Master’s report found that the drafting history supported Alaska’s argument, there would then have been a need for metalevel reasoning about whether preferring the history to the syntax is legitimate (for some legal sources, see Report, p. 186, n. 11). The latter point goes beyond the suggestion in Rissland and Skalak (1991) that one may always argue for dropping an antecedent from a rule.

Once it is decided that *general-conditions(SoHarBay)* must be tested, the second set of problems arises. Part of the difficulty comes from the usual source, namely vague or open-textured predicates such as “well-marked indentation,” “penetration,” “landlocked waters,” and “more than a

mere curvature of the coast.” (Another predicate, “the width of its mouth,” presented no problem in this case because the value was agreed to be about 12 nautical miles.) But beyond this, the relationship among the predicates is unclear: how do we write the rule

if $p_1(x)$ and . . . $p_n(x)$ then general-conditions(x) ?

From the English, one might think that “well-marked indentation” and “more than a mere curvature of the coast” were separate requirements. As used in the drafting history, however, they seem synonymous; and indeed the general conditions were criticized as circular during the drafting (Report, p. 191). In addition, the relationship between “penetration” and “landlocked waters” is unclear. Are they independent requirements, or does the value of one determine the value of the other? Supreme Court precedents seem to lead to the first conclusion; the syntax, to the other (see Report, pp. 199–200).

In the Master’s report, the outcome of this second exercise in rule interpretation was in effect to conclude that we do not know the logical form of the first English sentence in the definition of a bay. This conclusion triggered a case-based approach to testing the general conditions, using the various predicates as factors rather than as a neat conjunction of preconditions.

Modes of Reasoning with Cases

Programs that work with legal cases usually apply the same algorithm no matter what the source of the case. Human reasoners, however, adapt their style to the situation. If the

court hearing the current case has issued some recent decisions that are more or less on point, an extremely elaborate analysis may be called for, spelling out every point of similarity and dissimilarity and attending to every nuance in what the court said. In other situations one may use a precedent only for its facts and its result, ignoring the reasoning. In still others the important thing may be the pattern of reasoning displayed in the precedent rather than close factual similarity. The Master’s report contains examples of all of these. Here are illustrations of the latter two.

Fact-based Comparison

Continuing with the example of southern Harrison Bay, the most difficult feature to reach a conclusion about was whether the waters were landlocked. As shown in figure 1, the area has two arms, where both parties agreed the waters were landlocked, and a middle area, on which they disagreed. The most relevant precedents were other indentations with two arms, sometimes called double-headed bays. Five such precedents were available (covering one formation in Norway, one in California, and three in Alaska including the northwestern part of Harrison Bay). Given maps of each precedent area, together with the decision on whether it formed a single bay, it was possible to order the precedents along a numerical scale (Report, pp. 216–26). Southern Harrison Bay fell within the landlocked range; the Master’s report recommended accordingly; and the parties filed no exceptions.

This portion of the report seems unusual in making so little use of the reasoning in past cases and in producing a



Figure 1. A section of the north coast of Alaska. Disputed areas include the southern part of Harrison Bay (east of the Eskimo Islands), the islands around Prudhoe Bay, and the Arco pier. The northwestern part of Harrison Bay was agreed to form a bay.

basically geometrical solution to a legal problem. There are several explanatory factors. First, the precedents on double-headed bays did not contain much usable reasoning: only one was the subject of a judicial opinion, and even that one, from the International Court of Justice, did little more than announce its conclusion that the Norwegian Svaerholthavet had “the character of a bay” (see Report, p. 207).

Such guidance as there was came instead from a United States Supreme Court decision in which the geographical facts had less resemblance to southern Harrison Bay. In that decision (on whether Long Island Sound and the adjacent Block Island Sound formed a bay), the Court identified some requirements for landlocked waters:

We agree with the general proposition that the term “landlocked” “implies both that there shall be land in all but one direction and also that it should be close enough at all points to provide [a seaman] with shelter from all but that one direction.” [*Rhode Island and New York Boundary Case*, 469 U.S. 504, 525 (1985)]

In the report, an elaboration of “land in all but one direction” led to the numerical scale used to decide whether southern Harrison Bay was landlocked. The question whether the land was “close enough at all points” was essentially answered by article 7(4) of the Convention, which permits a line across the mouth of a bay to be as much as 24 miles long. As a final simplification, the evidence on the extent of shelter for a mariner was limited to the information available from two-dimensional maps. The last move seems to be standard, both in the legal cases and also, no doubt for practical reasons, among geographers who need to draw boundaries.

Reason-based Comparison

An interesting contrast to the Harrison Bay reasoning comes from a different part of the case, this one on the effect that near-shore islands have on drawing the coastline (see figure 1). The United States said that each island has its own coastline and its own three-mile limit. Alaska wanted to draw a single line as the coastline, running along the seaward side of any islands, and measure the three-mile limit from there. For these results the United States invoked the Convention, while Alaska argued on various grounds that the usual rules of the Convention should not apply.

One suggestion by Alaska was that the islands should be treated in the same way as islands along the most similar parts of the United States coast, notably in Louisiana and Mississippi. This would have been in line with the reasoning used for Harrison Bay. As the analysis was worked out in the Master’s report, however, the exact geography in past cases proved less important than the theory behind each decision on how to treat an island configuration. The

theories that were used were sometimes unclear, mutually inconsistent, and unpredictable even in application to the same area at different times. Consequently there was no basis for inferring how the islands off the north coast of Alaska would have been treated before the Convention took effect, and thus no basis for applying an exception to the Convention’s rules. This was so even though a 1985 Supreme Court case contained a statement that strongly supported Alaska’s position:

Prior to its ratification of the Convention on March 24, 1961, the United States had adopted a policy of enclosing as inland waters those areas between the mainland and off-lying islands that were so closely grouped that no entrance exceeded 10 geographical miles. [*Alabama and Mississippi Boundary Case*, 470 U.S. 93, 106 (1985)]

The Master’s report (pp. 53–54) took the statement to be nonbinding because (1) it was a statement of fact, not of law, and (2) it was not strictly necessary to the 1985 decision. The Supreme Court agreed and overruled Alaska’s exception.

The example highlights the importance of working on representations for the reasoning in legal cases, not just for the facts and the outcome. As for the differences from the Harrison Bay example—with respect to the role of the key sentence from the most important precedent, and with respect to the treatment of geographically similar cases—these are not inconsistencies. Rather, they result from differences in the available evidence, the available cases, and the arguments based on these that were or could have been made.

Sources of Cases

For a human reasoner, one of the most satisfying moments is finding a case that solves a puzzle or clinches an argument. Case-based programs may have the same goal. Achieving the goal, however, often requires going beyond the case base that seems natural for the problem at hand. Perhaps it is precisely because the normal stock of judicial opinions provides no answer that an issue becomes salient as a puzzle. Examples from the Master’s report of non-standard cases include the following:

1. To help settle the meaning of “permanent” in the Convention, a case from a domain having nothing to do with submerged lands but holding that eight years is long enough to count as permanent. (Report, p. 320.) The case was cited in one of the parties’ briefs; the brief writer might have found it from the legal reference *Words and Phrases*.

2. To help interpret an early boundary description, a case that had matching facts but that turned on another point. The statement relevant to our problem was thus dictum, and moreover it appeared in a concurring opinion.

(Report, p. 378.) The case was located through an early Supreme Court opinion (*Shively v. Bowlby*, 152 U.S. 1 (1894)), which provided a virtual treatise on the submerged lands decisions to that time.

3. To help establish the meaning of “high tide” in the Convention, a decree entered in a previous case. There had been no discussion of the question in the opinions leading up to the decree, but the decree itself equated “high tide” with “mean high water.” (Report, p. 234.) A Supreme Court decree also yielded one of the examples of a double-headed bay, and others provided some of the information on past treatment of islands as part of the coastline or not. The decrees are published in United States Reports, the same source as for Supreme Court opinions.

4. To help settle whether a body of water qualified as a bay under the Convention, the minutes of a meeting of a committee of federal officials, deciding that a similar neighboring body of water was a bay. (Report, p. 225.) The minutes were a document introduced into evidence—as were, again, many of the items pertaining to past treatment of islands. This suggests, at least for a case in which much of the evidence is documentary, that the line between precedent cases and the facts of a current case is less sharp than usually supposed: building the case base may require processing some of the evidence.

Conclusion

In this paper I have tried to identify some significant differences between human legal reasoning and the computational models we have so far. If the paper is successful, it will have suggested some fruitful directions for future research.

References

- Allen, L. E., and Saxon, C. S. 1987. Some Problems in Designing Expert Systems To Aid Legal Reasoning. In Proceedings, First International Conference on Artificial Intelligence and Law, Boston, 94–103. New York: Association for Computing Machinery.
- Allen, L. E., and Saxon, C. S. 1991. More IA Needed in AI: Interpretation Assistance for Coping with the Problem of Multiple Structural Interpretations. In Proceedings, Third International Conference on Artificial Intelligence and Law, Oxford, 53–61. New York: Association for Computing Machinery.
- Branting, L. K. 1991a. Integrating Rules and Precedents for Classification and Explanation: Automating Legal Analysis. Ph.D. diss., University of Texas at Austin.
- Branting, L. K. 1991b. Reasoning with Portions of Precedents. In Proceedings, Third International Conference on Artificial Intelligence and Law, Oxford, 145–154. New York: Association for Computing Machinery.
- Branting, L. K. 1994. A Computational Model of Ratio Decidendi. *Artificial Intelligence and Law* 2: 1–31.
- Gardner, A. v.d.L. 1987. *An Artificial Intelligence Approach to Legal Reasoning*. Cambridge: MIT Press.
- Gardner, A. v.d.L. 1989. Representing Developing Legal Doctrine. In Proceedings, Second International Conference on Artificial Intelligence and Law, Vancouver, 16–22. New York: Association for Computing Machinery. Also in A. Martino, ed., *Expert Systems in Law*, 201–209. Elsevier Science Publishers, 1992.
- Hart, H. L. A. 1961. *The Concept of Law*. Oxford: Clarendon Press.
- Levi, E. H. 1949. *An Introduction to Legal Reasoning*. Chicago: University of Chicago Press.
- McCarty, L. T. 1989. A Language for Legal Discourse: I. Basic Features. In Proceedings, Second International Conference on Artificial Intelligence and Law, Vancouver, 180–188. New York: Association for Computing Machinery.
- McCarty, L. T. 1995. An Implementation of *Eisner v. Macomber*. In Proceedings, Fifth International Conference on Artificial Intelligence and Law, College Park, Maryland, 276–286. New York: Association for Computing Machinery.
- McCarty, L. T., and Sridharan, N. S. 1981. The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations. In Proceedings, Seventh International Joint Conference on Artificial Intelligence, Vancouver, 246–253.
- Rissland, E. L., and Skalak, D. B. 1991. CABARET: Rule Interpretation in a Hybrid Architecture. *International Journal of Man-Machine Studies* 34: 839–887.
- Sergot, M. J.; Sadri, F.; Kowalski, R. A.; Kriwaczek, F.; Hammond, P.; and Cory, H. T. 1986. The British Nationality Act as a Logic Program. *Communications of the ACM* 29: 370–386.
- Skalak, D. B., and Rissland, E. L. 1992. Arguments and Cases: An Inevitable Intertwining. *Artificial Intelligence and Law* 1: 3–44.
- Susskind, R. E. 1989. The Latent Damage System: A Jurisprudential Analysis. In Proceedings, Second International Conference on Artificial Intelligence and Law, Vancouver, 23–32. New York: Association for Computing Machinery.
- United States v. Alaska* (No. 84, Original), Report of Special Master J. Keith Mann (1996), *exceptions overruled in part and sustained in part*, 117 S. Ct. 1888 (1997). A limited number of copies of the report are available from Anne Gardner. Major law libraries may have the report either in hard copy or on microfiche.
- Words and Phrases*. 1940–. St. Paul, Minn.: West Publishing Co. 46 vols.