

Rules and Precedents as Complementary Warrants

L. Karl Branting

Department of Computer Science
University of Wyoming
Laramie, Wyoming 82071-3682
karl@master.uwyo.edu

Bruce W. Porter

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712
porter@cs.utexas.edu

Abstract

This paper describes a model of the complementarity of rules and precedents in the classification task. Under this model, precedents assist rule-based reasoning by operationalizing abstract rule antecedents. Conversely, rules assist case-based reasoning through case elaboration, the process of inferring case facts in order to increase the similarity between cases, and term reformulation, the process of replacing a term whose precedents only weakly match a case with terms whose precedents strongly match the case. Fully exploiting this complementarity requires a control strategy characterized by *impartiality*, the absence of arbitrary ordering restrictions on the use of rules and precedents. An impartial control strategy was implemented in GREBE in the domain of Texas worker's compensation law. In a preliminary evaluation, GREBE's performance was found to be as good or slightly better than the performance of law students on the same task.

The Complementarity of Rules and Precedents for Classification

In a variety of domains, such as law, both general rules and specific precedents are useful for performing **classification** — the task of assigning a given input, or case, to a category and explaining the assignment. This section explains the complementarity of rules and precedents for performing classification and the disadvantages of arbitrarily restricting the order in which they can be combined.

A case is classified as belonging to a particular category by relating its description to the criteria for category membership. The justifications, or **warrants** (Toulmin, 1958), that can relate a case to a category, can vary widely in the generality of their antecedents. For example, consider warrants for classifying a case into the legal category "negligence." A rule, such as "An action is negligent if the actor fails to use reasonable care and the failure is the proximate cause of an injury," has very general antecedent terms (*e.g.*, "breach of reasonable care"). Conversely, a precedent, such as "Dr. Jones was negligent because he failed to

count sponges during surgery and as a result left a sponge in Smith," has very specific antecedent terms (*e.g.*, "failure to count sponges"). Both types of warrants have been used by classification systems to relate cases to categories.

The Role of Precedents

Classification systems have used precedents to help match the antecedents of rules with cases. Completing this match is difficult when the terms in the antecedent are **open-textured**, *i.e.*, when there is significant uncertainty whether they match specific facts (Gardner, 1984; McCarty and Sridharan, 1982). This problem results from the "generality gap" separating abstract terms from specific facts (Porter et al., 1990).

Precedents of an open-textured term, *i.e.*, past cases to which the term applied, can be used to bridge this gap. Unlike rule antecedents, the antecedents of precedents are at the same level of generality as cases, so no generality gap exists between precedents and new cases. Precedents therefore reduce the problem of matching specific case facts with open-textured terms to the problem of matching two sets of specific facts.

For example, an injured employee's entitlement to worker's compensation depends on whether he was injured during an activity "in furtherance of employment." Determining whether any particular case should be classified as a compensable injury therefore requires matching the specific facts of the case (*e.g.*, John was injured in an automobile accident while driving to his office) to the open-textured term "activity in furtherance of employment." The gap in generality between the case description and the abstract term makes this match problematical. However, completing this match may be much easier if there are precedents of the term "activity in furtherance of employment" (*e.g.*, Mary's injury was not compensable because it occurred while she was driving to work, which is not an activity in furtherance of employment; Bill's injury was compensable because it occurred while he was driving to a house to deliver a pizza, an activity in furtherance of employment). In this case, John's driving to his office closely matches Mary's driving to work, so

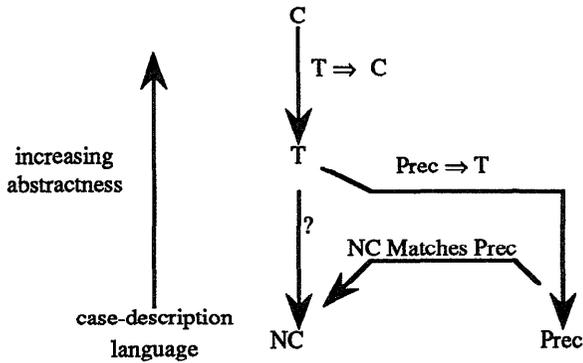


Figure 1: The role of a precedent, *Prec*, in classifying a new case, *NC*, into the category *C*. *Prec* is a precedent of the open-textured term *T*. Matching *NC* to *Prec* circumvents the “generality gap” separating *T* from *NC*.

John’s driving is probably not in furtherance of his employment. John’s injury is therefore probably not compensable.

Figure 1 illustrates the role of precedents in classification. A domain theory consisting of the rule $T \Rightarrow C$ and the precedent $Prec \Rightarrow T$ lacks rules connecting the terms of new case *NC* to the open-textured term *T*. However, a match between the new case and *Prec* permits an inference path to be constructed from the new case to category *C*. In this way, precedents often make it possible to classify cases that could not otherwise be classified.

The importance of precedents for the classification task is evidenced by the limitations of classification systems that do not use them. In particular, a number of systems use only rules, such as the Latent Damage System (Capper and Susskind, 1988), the Legal Decision-making System (Peterson and Waterman, 1985), the British Nationality Act Program (Sergot et al., 1986), and TA (Schlobohm, 1985). Without knowledge of precedents, these systems cannot determine whether the open-textured terms in rule antecedents match case descriptions, but must instead leave it to their users to make these determinations.

The Role of Rules

Despite the importance of precedents in performing classification, they too are often inadequate for the classification task when used alone. Because of the specificity and detail of case descriptions, few pairs of distinct cases have identical facts. Therefore, matching a new case to a precedent typically requires rules

to establish their equivalence. Two types of rule-based inference are useful: term reformulation and case elaboration.

Term reformulation. Rules can improve matching by reformulating an open-textured term as another term for which there are precedents that match the case more closely than do precedents of the original term. Suppose, for example, that a domain theory consists of the rule $T_1 \Rightarrow C$, precedent $Prec_1$ of T_1 , and precedent $Prec_2$ of T_2 . Suppose that new case *NC* matches $Prec_1$ weakly but matches $Prec_2$ strongly. The only explanation for classifying *NC* into the category *C* involves a weak match between *NC* and $Prec_1$. However, adding the rule $T_2 \Rightarrow T_1$ to the domain theory permits the goal T_1 to be reformulated as T_2 . This leads to a stronger alternative explanation involving a strong match between *NC* and $Prec_2$. See Figure 2.

More generally, term reformulation can be used to replace an open-textured term with a combination of terms. If some precedent of each of the new terms matches the new case, the system can explain the classification by combining the individual reasoning steps into a single explanation.

Previous researchers have acknowledged the importance of term reformulation by reporting limitations in systems that are unable to perform term reformulation. For example, Rissland and Skalak, working in the legal domain of the home-office deduction, reported that the TAX-HYPO system was limited by its inability to “...combine the analyses of individual predicates to generate an argument that takes into account the statute as a whole” (Rissland and Skalak, 1989a).¹ Similarly, Koton, working in the domain of cardiac disorders, found that her case-based system could not classify several cases involving “multiple noninteracting diagnoses” when the particular combination of diagnoses had not been seen before (Koton, 1988). Both researchers reported a similar need: the ability to combine solutions from several steps of case-based reasoning into a single explanation. Term reformulation improves matching by meeting this need.

Case elaboration. A second way in which rules can improve matching a case with a precedent is by inferring facts that are not explicitly stated in the case description. For example, a new case can match a precedent even if it lacks some of the terms in the precedent’s antecedent, provided that the missing terms can be inferred. Similarly, differing case terms can be

¹These limitations of TAX-HYPO were a motivation for the development of CABARET, a system that uses an approach to integration of rules and precedents that represents an alternative to the model proposed here. The distinguishing features of this approach include the use of an agenda-based control mechanism in which “heuristic control rules direct and interleave the two modes of reasoning by posting and prioritizing tasks for each to do” (Rissland and Skalak, 1989b) and the use of *dimensional analysis* (Ashley, 1988) within the case-based reasoning component.

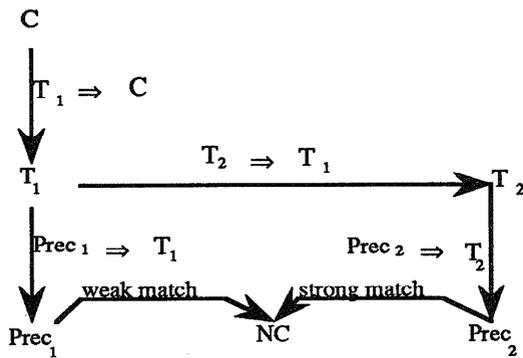


Figure 2: The role of term reformulation in classifying a new case, *NC*, into the category *C*. The rule $T_2 \Rightarrow T_1$ permits a reformulation step between the terms T_1 and T_2 . This leads to a stronger explanation of *NC*'s classification because *NC* matches *Prec₂* more strongly than *Prec₁*.

matched if both are manifestations of the same abstract term, *e.g.*, if both are symptoms caused by the same underlying physiological state or both have the same generalization.

For example, the Texas Court of Civil Appeals used case elaboration to match the facts of *Vaughn v. Highlands Underwriters Ins. Co.*, 445 S.W.2d 234 (1969) with the facts of an earlier precedent, *Janak v. Texas Employer's Ins. Co.*, 381 S.W.2d 176 (1964). *Janak* involved an accident that occurred during a deviation from the direct route to the drill site where the Draplia, the driver, and Janak, the passenger, worked on an oil drilling crew. The purpose of the deviation was to get ice to cool the crew's drinking water. Because of the hot and humid environment of the drill site, ice water was "reasonably essential" for the drilling operation.

One respect in which the facts of *Vaughn* differed from those of *Janak* was that the employee in *Vaughn* was injured while driving to get food rather than ice water. However, the court reasoned that in view of *Vaughn*'s 15 hour shift, getting food to eat during his shift was as important to his job activities as ice water was to the crew in *Janak*. Thus, the ice in *Janak* and the food in *Vaughn* match because both reduced physiological needs that would have interfered with the performance of the employee's job duties, *i.e.*, both were "reasonably essential" for their job duties.

Both Protos (Porter et al., 1990) and Casey (Koton, 1988) used rules for case elaboration. An ablation study of Protos demonstrated that case-elaboration made a significant contribution to Protos's ability to

classify cases (Mallory, 1989).

Rules and Precedents Should Be Treated Impartially

Some systems that perform classification use both rules and precedents. However, these systems do not always produce optimal explanations for the classification of cases because they restrict the order in which rules and precedents can be used. We characterize these methods as: (1) *Precedents-first systems* that use precedents strictly before rules, and (2) *Rules-first systems* that use rules strictly before precedents.

Precedents-first systems, such as Taxman-II (McCarty and Sridharan, 1982), Protos, and Casey apply precedents only to the top-level classification term and use rules only to improve the match between the new case and the precedents. Such systems can perform case elaboration, but not term reformulation. As a result, they cannot replace a classification term with an equivalent term, or boolean combination of terms, to improve the match between the new case and existing precedents.

Rules-first systems, such as Prolexs (Oskamp et al., 1989) and Gardner's system (Gardner, 1984), are unable to perform rule-based reasoning to assist in matching. In these systems, case-based reasoning can follow, but can never precede, rule-based reasoning in an explanation. This strict ordering permits term reformulation but not case elaboration. Consequently, the match between a precedent and a new case may be incomplete, reducing both classification accuracy and explanation quality.

Our classification system, GREBE, is unlike these systems in that it does not impose a rigid ordering on the use of precedents and rules. Instead, GREBE both permits precedents to operationalize rule antecedents and permits rules to assist in case matching. Moreover, if both a rule and a precedent are applicable to a goal, the choice is based on the strength of the resulting explanation, rather than on an arbitrary ordering scheme. A control strategy that permits rules and precedents to be used in this complementary fashion imposes no arbitrary ordering restrictions. Such a control strategy is said to be *impartial*.

GREBE: Impartiality in Legal Analysis

GREBE (GeneratoR of Exemplar-Based Explanations) is a system for *legal analysis* (Meldman, 1975), the task of determining whether a legal classification applies to a new case and explaining this determination. The top-level components of GREBE consist of an explanation generator, a memorandum generator, and a knowledge base of rules and cases. When the explanation generator receives a new case and a proposed classification for that case, it attempts to construct an explanation for the case's classification by back-chaining, using any combination of rules and precedents. If successful, the explanation generator outputs

one or more explanations. The memorandum generator orders these explanations by a heuristic estimate of their strengths. It then generates a natural-language equivalent of the explanations using a library of templates.

The algorithm of GREBE's explanation generator, a simplified version of which is set forth below, treats rules and precedents impartially.

GIVEN:

a new case NC , and
 the proposed classification of NC : $(Pred\ Arg_1 \dots Arg_n)^2$

DO:

If $(Pred\ Arg_1 \dots Arg_n)$ unifies with a proposition in NC
then return a trivial explanation
else

1. {rule-based reasoning}
For every rule R whose consequent unifies with $(Pred\ Arg_1 \dots Arg_n)$
if the antecedents $A_1 \dots A_m$ of R have explanations $E(A_1) \dots E(A_m)$
then return a rule-based explanation with warrant R and grounds $E(A_1) \dots E(A_m)$.
2. {case-based reasoning}³
 - (a) {precedent retrieval}
 Find PC , the precedent of $Pred$ most similar to NC .
 - (b) {structure matching}
 Find the best mapping, $M : PC \Rightarrow NC$, from PC to NC .
 - (c) {case elaboration}
 - i. **If** the match between PC and NC would be improved if propositions $p_1 \dots p_m$ were true in NC
then attempt to find explanations for $p_1 \dots p_m$. Let $E(p_i) \dots E(p_j)$ be the successful explanations and $p_k \dots p_l$ be the facts that could not be explained.
 - ii. **If** the match between NC and PC meets the acceptability threshold after case elaboration

²It is usually necessary to express a legal classification as a proposition, rather than simply a category, because legal predicates generally apply to tuples of case entities rather than to a case as a whole. For example, an employer may be liable to an employee for an injury arising from a particular activity. Thus, the predicate "worker's compensation liability" takes as arguments an employer, employee, injury, and activity. Since a single case may involve multiple employees, injuries, or activities, classifying a case as an instance of worker's compensation liability requires specifying the particular arguments to which the legal predicate applies.

³A discussion of GREBE's method of case-based reasoning is beyond the scope of this paper. See (Branting, 1991) for details.

then return a case-based explanation for $(Pred\ Arg_1 \dots Arg_n)$ with grounds $M : PC \Rightarrow NC$ and $E(p_i) \dots E(p_j)$ and qualification $p_k \dots p_l$.

This algorithm uses rules and precedents in a complementary fashion. In step 2(c)(i) of the algorithm, propositions that could improve the match if inferred become subgoals to which the explanation generator is recursively applied. This permits case elaboration because rules can apply to subgoals arising in case matching. Similarly, the explanation generator is recursively applied to the antecedents of any rule whose consequent unifies with the current goal in step 1 of the algorithm. This enables precedents to apply to rule antecedents, permitting term reformulation.

Moreover, the algorithm treats rules and precedents impartially because it imposes no arbitrary restrictions on the order in which they can be used. Instead, steps 1 and 2 of the algorithm are always performed together, so the explanation generator always attempts to apply both rule-based and case-based reasoning to every goal. If multiple explanations are generated, GREBE's memo generator orders the explanations by a heuristic estimate of their strength and presents the strongest to the user.

Impartiality Leads to Improved Explanations

This section illustrates how GREBE's impartial control strategy leads to improved explanations. Consider the following hypothetical case:

Joan and Donald were employed by the school district as teachers at a middle school and car-pooled together. Each workday, the driver of the car-pool was responsible for picking up some sandwiches on the way to work for both teachers to eat because there was no cafeteria at the school. On the day of the accident, Donald picked up Joan at her house and drove toward the school. Donald then deviated from the direct route to the school on his way to the sandwich shop. Before reaching the sandwich shop, Donald had an automobile accident in which Joan was injured. Does Joan have a claim for worker's compensation against the school district?

GREBE begins its analysis by searching for warrants for worker's compensation liability. Finding only statutory rules for this predicate, it chains through these rules until it reaches the goal of showing that the traveling was within the scope of Joan's employment. At this point it finds two distinct alternative explanations. The first uses a case-based explanation in which Joan's traveling is compared to the traveling of Draplia in the case of *Janak*, discussed above. The analogy to Draplia is weak, however, because, unlike Draplia, Joan was

not the driver, nor was she responsible for picking up the sandwiches.

The alternative explanation involves term reformulation using a common-law rule set forth in *Janak* that a passenger in a business car-pool is in the course of employment whenever the driver is in the course of employment. This rule is used to replace the goal of showing that Joan was within the scope of her employment with a new goal of showing that Donald was within the scope of his employment at the time of the accident. Donald's traveling is much more closely analogous to Draplia's traveling, so a strong argument can be made that Donald, and therefore Joan, was acting in furtherance of employment.

As GREBE reports, the stronger explanation identifies the business car-pool passenger rule from *Janak* and shows how its antecedents are satisfied:⁴

...

The trip to the sandwich shop was an activity in furtherance of Joan's employment.

This conclusion follows from the rule of *Janak v. Texas Employer's Ins. Co.*, 381 S.W.2d 176 (1964) that a passenger in a business car-pool is in the course of employment whenever the driver is in the course of employment if:

...

- iv. Joan was a passenger in the trip to the sandwich shop pursuant to a business car-pool.

This conclusion follows from the very strong analogy between the given case and the facts of *Janak v. Texas Employer's Ins. Co.*, 381 S.W.2d 176 (1964) that were relevant to the conclusion that *Janak* was a passenger in the deviation to Runge pursuant to a business car-pool.

- v. The trip to the sandwich shop was an activity in furtherance of Donald's employment.

Having shown that Joan's status depends on whether Donald's traveling was in furtherance of his employment, GREBE's analysis turns to the latter question. Donald's traveling matches both Draplia's driving in *Janak* and ordinary commuting.

Two conflicting explanations can be made concerning whether the trip to the sandwich shop was an activity in furtherance of Donald's employment. The stronger explanation is that:

The trip to the sandwich shop was an activity in furtherance of Donald's employment.

This conclusion follows from the very strong analogy between the given case and the facts of *Janak v. Texas Employer's Ins. Co.*, 381 S.W.2d 176 (1964) that were relevant to the conclusion that the deviation to Runge was an activity in furtherance of Draplia's employment.

⁴Each of the following excerpts, printed with typewriter font, is verbatim text from GREBE's memorandum generator. See (Branting, 1991) for details on the generator.

An important intermediate conclusion in the reasoning of *Janak* was that the object of Draplia's deviation, ice water, was *reasonably essential* for oil drilling. However, the facts of the hypothetical case don't state that sandwiches were "reasonably essential" for teaching. GREBE therefore attempts case elaboration, *i.e.*, attempts to improve the match by inferring this fact. GREBE's knowledge base contains two precedents for "reasonably essential": ice water was found to be reasonably essential under the facts of *Janak*, and food was found to be reasonably essential under the facts of *Vaughn*. The facts of *Vaughn* relevant to the predicate "reasonably essential" are closer to the facts of the hypothetical case than are the facts of *Janak* relevant to the same predicate, so GREBE uses *Vaughn* to support the conclusion that sandwiches were reasonably essential for teaching. This analogy is supported by the following inference:

Sandwiches being at the middle school was reasonably essential for teaching children.

This conclusion follows from the very strong analogy between the given case and the facts of *Vaughn v. Highlands Underwriters Ins. Co.*, 445 S.W.2d 234 (1969) that were relevant to the conclusion that *Vaughn's* having food was reasonably essential for *Vaughn* transporting sulfur.

Consider the effect of altering the hypothetical case so that the accident occurs after Donald purchased the sandwiches and was back on a direct route to the middle school. As in the previous case, both the business car-pool passenger rule and case-based reasoning are applicable to the goal of showing that Joan's traveling was in furtherance of her employment. Unlike the previous case, however, the business car-pool passenger rule does not lead to the stronger explanation. This is because the accident occurred after the deviation from the direct route to the school in the second hypothetical case, whereas in *Janak* and in the previous hypothetical case, it occurred during the deviation. As a result, the match between Donald's traveling and Draplia's traveling in *Janak* is weaker than in the previous cases. GREBE's strongest explanation involves instead, a direct match of Joan's traveling to ordinary commuting.

The stronger explanation in Car-pool Case number 2 is that the school district is not liable under worker compensation to Joan for Joan's injury because her trip to the Middle School was not an activity in furtherance of Joan's employment. This conclusion follows from the very strong analogy between the given case and the facts that are relevant to the conclusion that ordinary commuting to work is not an activity in furtherance of a typical employee's employment as held in *American General Ins. v. Coleman*, 157 Tx. 377, 303 S.W.2d 370.

These examples illustrate how impartiality leads to improved explanations. A system that was restricted to applying precedents to its top level goal couldn't produce the strongest explanation in the first hypothetical case. The strongest explanation in the first hypothetical case required term reformulation, *i.e.*, using a rule to replace the goal of showing that Joan's traveling was in furtherance of employment with the goal of showing that the driver's traveling was in furtherance of employment. Nor could this explanation have been produced by a system that was incapable of case elaboration. Matching the hypothetical case to the *Janak* case required inferring a fact—that sandwiches were reasonably essential for the activity of teaching—that was not given in the facts of the case. Similarly, a system that was restricted to applying rules to its top level goal couldn't have produced the strongest explanation in the second hypothetical case. The strongest explanation in that case required applying a precedent to the top level goal. Only an impartial system has the flexibility to produce both explanations.

The next section presents some preliminary empirical results tending to establish the utility of an impartial control strategy.

Empirical Evaluation of GREBE

This section describes an experimental evaluation of GREBE in which GREBE's analysis of 18 worker's compensation cases was compared to the analysis of the same cases by law students. The purpose of the evaluation was primarily to assess the overall effectiveness of GREBE in addressing this task, but a secondary purpose was to demonstrate that analysis of worker's compensation cases is a challenging task, even for humans with legal training.

The experiment was conducted as follows. Five students at the University of Texas Law School, responding to an advertisement for several "short legal-research projects," were each presented with a different group of three or four related hypothetical worker's compensation cases. For each hypothetical case, the students were asked to perform whatever research was necessary to determine the applicable legal warrants, construct the strongest explanations for and against worker's compensation liability based on those warrants, and set forth the explanations in a short memorandum. The subjects were asked to record the length of time they spent on each problem. Three of the subjects were second-year Juris Doctor students, and two had foreign law degrees and were enrolled in the Masters of Comparative Law program. Each student was paid for his or her participation.

The memoranda produced by the students, together with GREBE's analysis of the same 18 hypothetical cases, were then given to a domain expert, an attorney who is a recognized authority on Texas worker's compensation law. The domain expert was asked to grade all of the analyses, applying the following criteria:

Problem Solver	Issues	War.	Expl.	Overall
Students	1.94	1.78	1.78	1.78
GREBE	2.11	1.88	2.22	2.00

Table 1: Average grades for analyses of 18 worker's compensation hypothetical cases by students and GREBE broken down by issues, warrants, explanations, and overall grade. Letter grades have been converted into their numeric equivalent on a 4-point scale.

1. **Issues.** Does the memorandum correctly identify the relevant legal issues?
2. **Warrants.** Have the legal warrants (rules and precedents) applicable to the issues been identified?
3. **Explanations.** Are the explanations contained in the memorandum complete, sound, and persuasive?

The domain expert was asked to use the three criteria to assign each analysis an overall grade.

The results of the domain expert's grading for the students and GREBE are summarized in Table 1. Analyzing the worker's compensation cases was clearly a challenging task for the law students. The average student solution time was 2.77 hours (an average of 5.12 hours for the first problem in each set), and the overall grade on the analyses they produced was only 1.78, equivalent to a C-. GREBE's analyses almost always received a slightly higher grade, receiving an average overall grade of 2.0. GREBE received 4 B's, 9 C's, and 4 F's, whereas the students received 1 B, 12 C's, and 5 F's.

In only one problem did GREBE receive a lower grade than a student. In that case GREBE received an F for not identifying the issue that the injured worker might have been an independent contractor rather than an employee, whereas the student identified this issue. GREBE failed to identify this issue because its knowledge base simply has no information about the distinction between employees and independent contractors. In two other cases, both a student and GREBE were down-graded for failing to identify potentially relevant warrants that were outside of GREBE's knowledge base.

The greatest difference between GREBE and the students was in grades for explanation quality, where GREBE's average grade was 2.22 and the student's average grade was 1.78, a difference of .34. By contrast, there was only a difference of .1 between GREBE's grade for identifying the correct warrants (1.88) and the average grade for the students (1.78). This suggests that both the students and GREBE were hindered by incomplete knowledge of the domain, but that GREBE was somewhat better at assembling its knowledge into explanations. That GREBE received a grade of B on 4 of the problems indicates that GREBE is capable of producing sound and informative analyses on problems

for which the knowledge base is adequate.

This evaluation is tentative for two reasons. First, the accuracy of the evaluation is uncertain because it rests upon the judgment of a single domain expert. The paucity of objective standards for evaluating legal analysis makes it desirable to base an evaluation upon the judgment of multiple domain experts. A second reason that the evaluation is tentative is that it fails to isolate the contributions of the various components of the GREBE system. An ablation study is needed to determine relative importance of, *e.g.*, GREBE's impartial control strategy, relational case-description language, precedent-retrieval algorithms, and choice of natural-language templates.

Conclusion

This paper has described a model of the complementarity of rules and precedents in the classification task. Under this model, precedents assist rule-based reasoning by operationalizing abstract rule antecedents. Conversely, rules assist case-based reasoning through case elaboration, the process of inferring case facts in order to increase the similarity between cases, and term reformulation, the process of replacing a term whose precedents only weakly match a case with terms whose precedents strongly match the case. Fully exploiting this complementarity requires a control strategy characterized by impartiality, the absence of arbitrary ordering restrictions on the use of rules and precedents.

An impartial control strategy was implemented in GREBE in the domain of Texas worker's compensation law. In a preliminary evaluation, GREBE's performance was found to be as good or slightly better than the performance of law students in the same task. While the contribution of GREBE's control strategy to its overall performance was not isolated in the evaluation, GREBE's strong performance in comparison to law students suggests that impartiality can contribute to effective integration of rules and precedents for classification.

Acknowledgements

Support for this research was provided by a grant from the National Science Foundation (IRI-8620052), and by contributions from Hughes Research Laboratories, GTE Research Laboratories, and the Cray Foundation.

References

- Ashley, K. D. (1988). *Modelling Legal Argument: Reasoning with Cases and Hypotheticals*. PhD thesis, The University of Massachusetts.
- Branting, L. K. (1991). *Integrating Rules and Precedents for Classification and Explanation: Automating Legal Analysis*. PhD thesis, University of Texas at Austin.
- Capper, P. N. and Susskind, R. E. (1988). *Latent Damage Law — The Expert System*. Butterworths.
- Gardner, A. (1984). *An Artificial Intelligence Approach to Legal Reasoning*. PhD thesis, Stanford University.
- Koton, P. (1988). *Using Experience in Learning and Problem Solving*. PhD thesis, Massachusetts Institute of Technology. Department of Electrical Engineering and Computer Science.
- Mallory, R. S. (1989). Sources of classification accuracy in Protos. Technical Report AI-TR89-118, Artificial Intelligence Laboratory, Department of Computer Sciences, University of Texas at Austin.
- McCarty, L. T. and Sridharan, N. S. (1982). A computational theory of legal argument. Technical Report LRP-TR-13, Laboratory for Computer Science Research, Rutgers University.
- Meldman, J. A. (1975). *A Preliminary Study in Computer-Aided Legal Analysis*. PhD thesis, Massachusetts Institute of Technology.
- Oskamp, A., Walker, R. F., Schrickx, J. A., and van den Berg, P. H. (1989). Prolexs, divide and rule: a legal application. In *Proceedings of the Second International Conference on Artificial Intelligence and Law*, Vancouver, B.C.
- Peterson, M. and Waterman, D. (1985). Rule-based models of legal expertise. In Walters, C., editor, *Computing Power and Legal Reasoning*, pages 627-659. West Publishing Company, Minneapolis, Minnesota.
- Porter, B. W., Bareiss, E. R., and Holte, R. C. (1990). Concept learning and heuristic classification in weak-theory domains. *Artificial Intelligence Journal*, 45(1-2).
- Rissland, E. L. and Skalak, D. B. (1989a). Case-based reasoning in a rule-governed domain. In *Proceedings of the Fifth IEEE Conference on Artificial Intelligence Applications*. Institute of Electrical and Electronic Engineers, Inc.
- Rissland, E. L. and Skalak, D. B. (1989b). Combining case-based and rule-based reasoning: A heuristic approach. In *Eleventh International Joint Conference on Artificial Intelligence*, pages 524-530, Detroit, Michigan.
- Schlobohm, D. (1985). TA—a prolog program which analyzes income tax issues under section 318(a) of the internal revenue code. In Walters, C., editor, *Computing Power and Legal Reasoning*, pages 765-815. West Publishing Company, Minneapolis, Minnesota.
- Sergot, M. J., Sadre, 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.
- Toulmin, S. E. (1958). *The Uses of Argument*. Cambridge University Press.