

Integrating Generalizations with Exemplar-Based Reasoning

L. Karl Branting
Department of Computer Sciences
University of Texas

ABSTRACT

Knowledge represented as generalizations is insufficient for problem solving in many domains, such as legal reasoning, because of a gap between the language of case-descriptions and the language in which generalizations are expressed, and because of the graded structure of domain categories. Exemplar-based representation addresses these problems, but accurate assessment of similarity between an exemplar of a category and a new case requires reasoning both with general domain theory and with the explanation of the exemplar's membership in the category. GREBE is a system that integrates generalizations and exemplars in a cooperative manner. Exemplar-based explanations are used to bridge the gap between case-descriptions and generalizations, and domain theory in the form of general rules and specific explanations is used to explain the equivalence of new cases to exemplars.

INTRODUCTION

In many important domains, knowledge expressed as generalizations is insufficient for such important tasks as determining membership in domain categories and evaluating domain predicates. One reason for the insufficiency of generalizations in such domains is that there may be a “gap” between the language in which cases are described and the language in which generalizations are expressed [PBH89]. A second reason is that domain categories may exhibit a gradient of centrality or typicality [Bar85] which generalizations expressed as rules are ill-suited to represent [SM81].

Both of these factors are illustrated by the domain of legal reasoning. Determining the legal consequences of a set of facts may require determining whether a surgeon acted with “reasonable care,” a killer acted with “malice,” or an employee was acting “in furtherance of employment,” since these terms appear in general legal rules for determining guilt and liability. However, the terms “reasonable care,” “malice,” and “in furtherance of employment” do not appear in case descriptions, and the domain theory provides no general rules for determining whether such terms are satisfied under the facts of a given case [vdLG84].

Graded structure is illustrated by the category “activities in furtherance of employment.” An employee working on an assembly line is clearly acting in furtherance of his employment, but what about an employee carrying equipment from his car to the shop, driving from home to work, or shaving in preparation for work? It is problematical to determine at what point an activity is sufficiently remote from work that it is no longer a category instance. This example illustrates a gradient of possible cases—from clear category instances through unclear cases to clear noninstances—that cannot easily be expressed by any single general rule.

These problems are addressed by an approach to knowledge representation in which full descriptions of known instances, or *exemplars*, of various categories are retained and each new case is analyzed by comparing it to the exemplars that it most closely resembles. Examples of this approach include MEDIATOR [Sim85], Protos [PBH89], Kibler and Aha's systems [KA87], and in

BRANTING

the legal reasoning community, HYPO [RA87]. Exemplar-based systems typically use a feature vector representation of cases and assess the degree of similarity between cases by calculating the weighted sum (or product) of features of the exemplar matched by the new case.

An exemplar-based representation makes it possible to reason about categories for which there are insufficient generalizations and is well suited for concepts with graded structure, since there is a range of possible degrees of match with an exemplar. However, determining category membership exclusively as a weighed function of shared features has been criticized on the grounds that it neglects the generalization-based domain theory in which exemplars are embedded [MM85]. Murphy and Medin point out that the relative feature weights that determine degree of similarity depend on the context and task. They conclude that exemplar-based categorization requires knowledge of the relations among features and of the explanatory principles that connect exemplars to the categories of which they are members.

The use of general domain theory to assist in the assessment of similarity between cases was investigated in Protos [PBH89], a learning apprentice for heuristic classification in the domain of clinical audiology. Categories are represented in Protos by exemplars embedded in a network of causal and associational rules derived from explanations of category membership. The similarity of a new case to an exemplar is evaluated by attempting to construct an explanation of featural equivalence between the cases from these rules. Protos demonstrated that use of general domain theory to assist in similarity assessment could lead to high levels of performance in audiology.

Protos is nevertheless inadequate for more complex domains such as legal reasoning. Protos is limited to a feature-vector representation of cases that is unsuited to the complex narratives that constitute the facts of legal cases. In addition, Protos can only apply generalization-based reasoning in assessing similarity between individual case features, and can only apply exemplar-based reasoning to its top-level goal or to infer a new case feature necessary to match an exemplar feature. More complex domains such as legal reasoning require the ability to choose between and combine exemplar-based and generalization-based reasoning more flexibly so that each technique can be used in support of the other.

For example, it is sometimes necessary to use domain generalizations to reformulate a goal into subgoals, some of which are amenable to generalization-based reasoning and others of which require exemplar-based reasoning. Similarly, determining whether a feature of an exemplar is present in a new case can require both generalizations and exemplar-based reasoning.

This paper describes an approach to flexible integration of generalization-based and exemplar-based reasoning. This approach is applicable to complex cases that do not lend themselves to feature-vector representation and permits reasoning steps used in assessing similarity to be reused in subsequent cases.

OVERVIEW OF GREBE

GREBE (Generator of Recursive Exemplar-Based Explanations) is a system that uses knowledge in the form both of generalizations and category exemplars to determine the classification of new cases. GREBE integrates generalization-based knowledge with exemplars in two ways. First, exemplar-based reasoning is used to help evaluate antecedents of legal or common-sense rules for which there are no applicable generalizations. Second, general domain rules and specific explanations of category membership by exemplars are used in the assessment of similarity between cases. In this manner, generalization-based reasoning and exemplar-based reasoning are treated as complementary processes, each of which is necessary for the success of the other.

BRANTING

GREBE uses a semantic network representation of cases in which individual facts correspond to relation/unit/value triples and the facts of an entire case correspond to a labeled graph.

When GREBE is queried about whether a certain conclusion applies to a case, it tries to find the conclusion in the case description. If it is unable to do so, it attempts to construct either of the following types of explanations of the conclusion:

- *Generalization-based explanation.* The conclusion is the consequent of a general domain rule all the antecedents of which are themselves explained. This form of explanation is similar to “explanation as proof” [Moo88] [KC85].
- *Exemplar-based explanation.* The conclusion is justified by the similarity between the new case and the relevant aspects of an exemplar to which the conclusion applied.

In exemplar-based explanation, not all of the facts of an exemplar case are necessarily relevant to a given result. For example, if several explanations apply to an exemplar, it may be that only a subset of the facts are relevant to each explanation. The facts of a case that are used to explain a given result are the exemplar’s *critical facts* with respect to the result. The critical facts of an exemplar form a labeled subgraph of the graph that represents all the facts of the exemplar.

The critical facts of an exemplar are necessarily quite specific. Even though it is the pattern of relations and not the particular individuals in a precedent that must be matched in a new case, new cases nevertheless seldom contain exactly the same pattern of relations as any precedent. The solution to the problem posed by the specificity of exemplars is to use generalization-based and exemplar-based explanations to explain how individual critical facts are matched in a new case. This permits multiple sources of knowledge to be exploited in order to explain the equivalence of a new case to an exemplar.

GREBE assesses the degree of similarity between a new case and an exemplar with respect to a given conclusion by attempting to map the subgraph representing the critical facts of the exemplar onto the new case.¹ A best-first search is performed among possible mappings between the critical facts of the exemplar and the new case, using fewest unmatched triples as the evaluation function. GREBE is then called recursively to attempt to infer any facts missing from the new case that are needed for a perfect match.

Missing facts can be inferred by reusing the explanations from previous exemplars. These explanations may be either generalization-based explanations or exemplar-based explanations. For example, if in a previous exemplar a common-sense rule was used to infer a given relation, this common-sense rule is available to infer the same relation in subsequent cases in which the rule’s antecedents are met. Similarly, if there is an exemplar of the relation, then the relation can be inferred in any new case that shares the critical facts of the exemplar.

GREBE’s knowledge base currently contains rules and a small (but growing) collection of exemplar cases concerning the compensability under Texas worker’s compensation law of injuries to workers traveling outside of the work place. GREBE’s rules fall into two distinct categories: legal rules and common-sense rules. Legal rules are rules that are explicitly stated in statutes or judicial opinions. Common-sense rules represent reasoning in judicial opinions that is implicit because it is too obvious (to humans) to need pointing out. An example is the inference that if an activity is a duty of employment, then each step of that activity is a duty of employment as well.

¹This process resembles the structure mapping of [Gen83]. It differs, however, in that the critical facts of an exemplar are a part of the domain theory and cannot be recognized *a priori* using syntactic criteria such as relationality or systematicity.

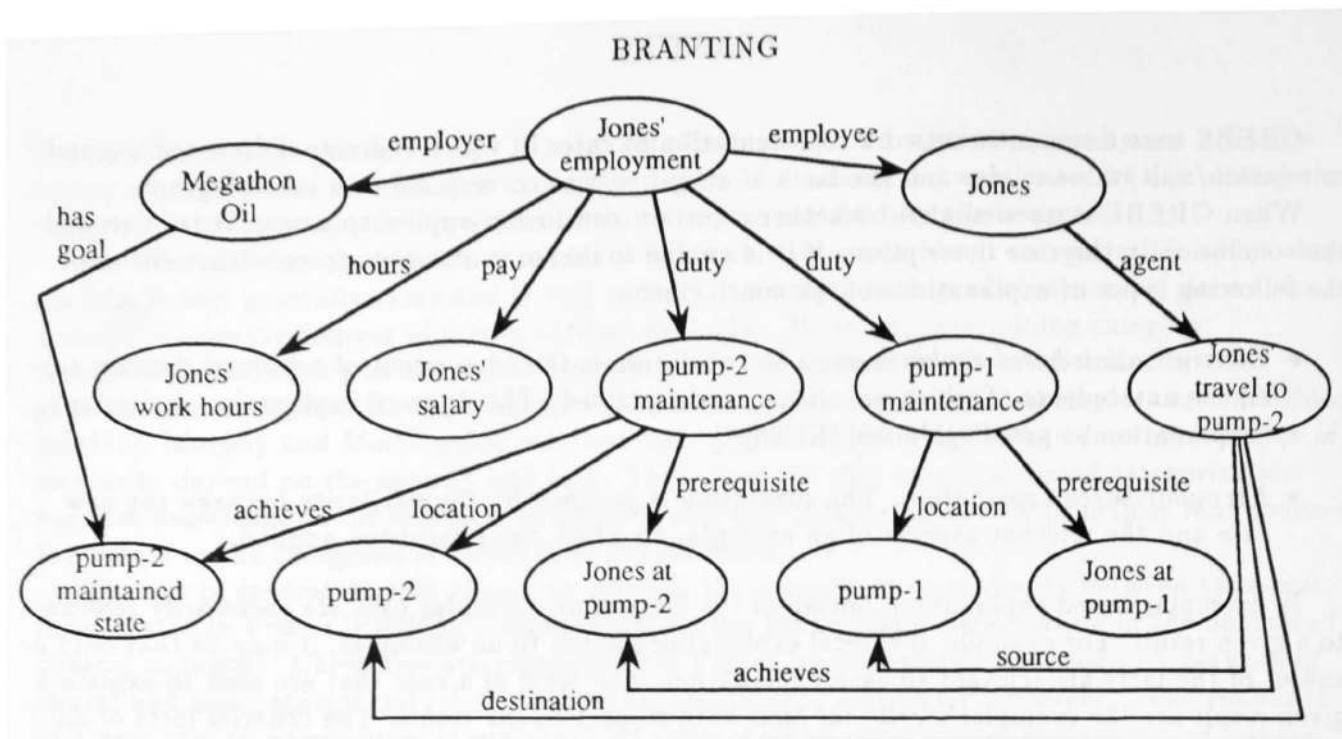


Figure 1: A partial representation of the facts of *Jones*.

USING GENERALIZATIONS AND EXEMPLARS TO ANALYZE A NEW CASE

Consider the following hypothetical case: Jones, a maintenance man employed by Megathon Oil Company, was involved in a one-car accident while driving in his own car from one pumping station where he had performed maintenance duties to a second pumping station where he planned to perform additional maintenance duties. Figure 1 shows a portion of the facts of *Jones*.²

If the system is queried whether Megathon is liable to Jones for his injuries, it is able to create a partial generalization-based explanation for Megathon's liability, shown in Figure 2. Megathon's liability to Jones is explained by statutory rule 1 and by the conclusions that Jones was employed by Megathon and that the injury was "sustained in the course" of the job. That the injury was "sustained in the course" of Jones' job follows under statutory rule 2 from the following conclusions: the injury occurred during Jones' travel to pump station 2, this travel was "in furtherance of" Jones' employment, and Jones' injury "originated in" the employment.

The gap between case descriptions and domain generalizations emerges in attempting to determine whether Jones' travel was in furtherance of his employment. There are no general rules for determining whether an activity is in furtherance of employment, so GREBE attempts to construct an exemplar-based explanation of this predicate. First, a promising exemplar of traveling in furtherance of employment, *Jecker v. Western Alliance Ins. Co.*, 369 S.W.2d 776 (Tex. 1963), is identified and retrieved.³ Then, the degree of similarity between *Jecker* and *Jones* is assessed by attempting to create an exemplar-based explanation using *Jecker* as the exemplar.

Creating an Exemplar-Based Explanation

In constructing an exemplar-based explanation, GREBE begins by mapping the criterial facts of *Jecker* with respect to travel in furtherance of employment onto *Jones*. Figure 3 shows a

²Italicized names refer to the case involving the person named, e.g., the *Jones* case, whereas unitalicized names refer to the person himself, e.g., Jones.

³Identification and retrieval of appropriate exemplars is performed in a manner similar to Protos' use of difference links [PBH89].

BRANTING




Figure 2: A partial generalization-based explanation of workers' compensation liability.

portion of the mapping from these criterial facts (which actually consist of 31 triples) onto the facts of *Jones*. Since it is the pattern of relationships, and not the particular individuals, of a case that are responsible for its legal consequences, a relation/unit/value triple in the exemplar is considered to match a triple in the new case if the relations are equal and if the unit and value mappings are consistent with those of other matched triples. In Figure 3, each of the triples of *Jecker* shown has a match in *Jones* under the mapping shown except that *Jecker*'s travel was an implied duty of his employment, whereas this relation is not given as part of the *Jones* case. GREBE therefore attempts to infer that *Jones* had a duty to travel.

GREBE is unable to construct a generalization-based explanation that *Jones* had a duty to travel, but finds that *Jecker* has an exemplar-based explanation that *Jecker*'s traveling was an implied duty of employment. This exemplar-based explanation was used by the court that decided *Jecker* as part of its explanation of the similarity between *Jecker* and an earlier exemplar in which traveling was an express duty.

GREBE fetches the criterial facts of *Jecker* with respect to this explanation (represented as 19 triples) and performs a mapping from this set of criterial facts onto *Jones*. Under the best mapping (a portion of which is shown in Figure 4) a criterial fact for *Jecker*'s implied duty to travel—that the traveling occurred during his work hours—is unmatched in *Jones*. However, GREBE constructs an explanation that *Jones*' travel occurred during his work hours (shown in Figure 5) by reusing a common-sense rule from an earlier exemplar, *Brown*. This rule provides that if an employee determines his own hours, then any time he spends performing job duties is, in effect, working hours.

GREBE returns an explanation structure showing that there is an actual or inferrable fact in *Jones* corresponding to each criterial fact of *Jecker* with respect to travel in furtherance of employment, with the sole exception that the maintenance site to which *Jones* was traveling was not under the employer's direct control. *Jones* is therefore strongly analogous to *Jecker*. The exemplar-based explanation that *Jones*' travel was in furtherance of his employment satisfies the antecedent of statutory rule 2 (in Figure 2) and completes the explanation that Megathon is liable to *Jones* for his accident.

If *Jones* were modified to provide that *Jones* was traveling home from a maintenance site rather than between two maintenance sites, GREBE would find two additional facts to be unmatched: there would be no employment activity performed at *Jones*' destination, and being at the destination would not be a prerequisite for an employment activity. GREBE's analysis is

BRANTING




Figure 3: The dashed vertical arrows indicate the mapping from a portion of the criterial facts of *Jecker* with respect to travel in furtherance of employment onto the facts of *Jones*. All the criterial facts shown are matched except for Jones' travel being an implied duty of employment.

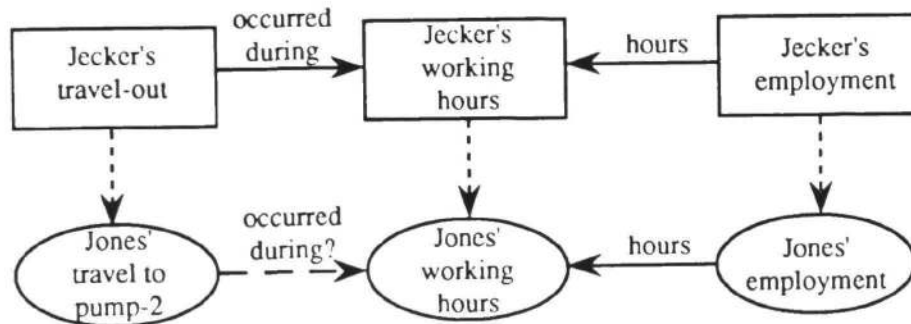


Figure 4: Under the mapping between a portion of the criterial facts of *Jecker* with respect to implied duties and facts of *Jones*, the fact that the travel occurred during working hours is unmatched.

consistent with the assessment of human experts that the original hypothetical is strongly analogous to *Jecker*, but that the modified hypothetical differs significantly from *Jecker* [Bra88].

Explanation Reuse

The accurate assessment of similarity between *Jones* and *Jecker* depended on reuse of the explanation that *Jecker*'s travel was in furtherance of his employment. This example illustrates that exemplar-based explanations can be reused in two ways.

First, knowledge that a particular set of facts of an exemplar is criterial for a given category is reused when a new case is classified on the basis of its match with those criterial facts. For example, knowing the criterial facts of *Jecker* for traveling in furtherance of employment means knowing that *Jecker* was acting in furtherance of his employment because he was driving from one location where he performed a job duty on a direct route to a second location where he intended to perform a job duty, and the driving was an implied duty under the employment contract. This portion of the explanation is reused in explaining that *Jones*' travel was in furtherance of his

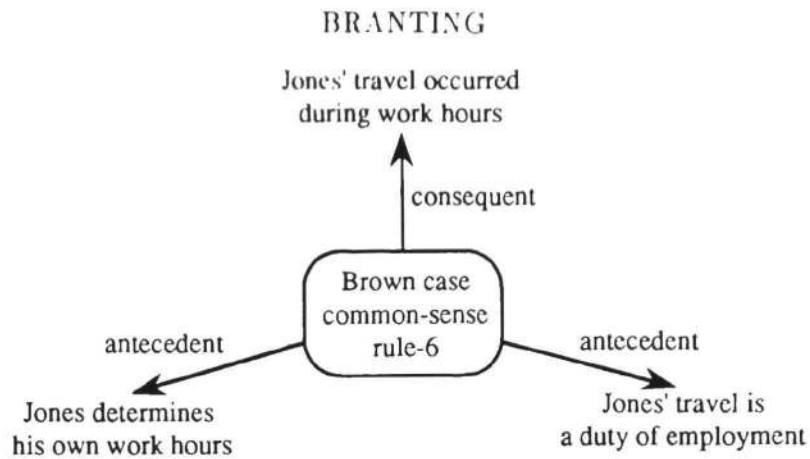


Figure 5: A generalization-based explanation that *Jones*' travel occurred during his work hours.

employment because of the similarity between these criterial facts and the facts of *Jones*.

The second way that exemplar-based explanations can be reused is that explanations of any inferred criterial facts can be reused in subsequent exemplar-based explanations. The exemplar-based explanation for *Jecker*'s implied duty to travel was reused to infer that *Jones* had an implied duty to travel. Similarly, the generalization-based explanation that the travel in *Brown* occurred during work hours was reused to infer that *Jones*' travel occurred during his work hours.

Integrating Multiple Explanations

The *Jones* case illustrates how generalization- and exemplar-based explanations drawn from various sources can be integrated by GREBE into a single explanation for the classification of a new case. The explanation of Megathon's liability for *Jones*' injury combines two statutory rules, exemplar-based explanations involving two different aspects of *Jecker*, and a common-sense rule taken from a second exemplar, *Brown*. Exemplar-based explanation was necessary to satisfy an antecedent of a generalization-based explanation for workers' compensation liability. The assessment of similarity between *Jecker* and *Jones*, in turn, required both generalization-based reasoning and additional exemplar-based reasoning.

Such an explanation could not be produced by a system limited exclusively to generalization- or exemplar-based reasoning. Neither could it be produced by Protos, which is unable to apply generalization-based reasoning to its top-level rule or exemplar-based reasoning to rule antecedents.

CONCLUSION

GREBE integrates generalizations with exemplars in a manner that compensates for the weakness of each form of knowledge representation. Exemplars help bridge the gap between case descriptions and the language of generalizations, and aid in the representation of graded concepts. General domain theory and specific exemplar-based explanations are necessary for accurate assessment of similarity between complex and superficially dissimilar cases.

GREBE represents an advance over previous exemplar-based systems in that it retains and reuses the explanations of category exemplars, uses exemplar-based reasoning recursively to assist in assessment of similarity, and allows both generalization- and exemplar-based reasoning to be freely combined.

BRANTING

ACKNOWLEDGEMENTS

Support for this research was provided by the Army Research Office under grant number ARO DAAG29-84-K-0060.

REFERENCES

- [Bar85] Lawrence W. Barsalou. Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(4):629–649, October 1985.
- [Bra88] L. Karl Branting. Protocol analysis of five worker's compensation problems. Unpublished protocol analysis of problem solving by attorneys at the Colorado Court of Appeals, September 1988.
- [Gen83] Dedre Genter. Structure mapping: A theoretical framework for analogy. *Cognitive Science*, 7(2):155–170, April-June 1983.
- [KA87] Dennis Kibler and David W. Aha. Learning representative exemplars of concepts: an initial case study. In *Proceedings of the 4th International Workshop on Machine Learning*, pages 1–11, 1987.
- [KC85] Smadar Kedar-Cabelli. Purpose-directed analogy. In *Proceedings of the 7th Annual Conference of the Cognitive Science Society*, 1985.
- [MM85] George L. Murphy and Douglas L. Medin. The role of theories in conceptual coherence. *Psychological Review*, pages 289–316, 1985.
- [Moo88] Raymond Mooney. *A General Explanation Based Learning Mechanism and its Application to Narrative Understanding*. PhD thesis, University of Illinois, 1988.
- [PBH89] Bruce W. Porter, E. Ray Bareiss, and Robert C. Holte. Knowledge acquisition and heuristic classification in weak-theory domains. Technical Report AI-TR89-96, Artificial Intelligence Laboratory, Department of Computer Sciences, University of Texas at Austin, February 1989.
- [RA87] Edwina Rissland and Kevin Ashley. Hypo: A case-based reasoning system. Project memo 18. Department of Computer and Information Sciences, University of Massachusetts, 1987.
- [Sim85] Robert L. Simpson. *A Computer Model of Case-based Reasoning in Problem Solving: An Investigation in the Domain of Dispute Mediation*. PhD thesis, Georgia Institute of Technology, 1985.
- [SM81] Edward E. Smith and Douglas L. Medin. *Categories and Concepts*. Harvard University Press, 1981.
- [vdLG84] Anne van der Lieth Gardner. *An Artificial Intelligence Approach to Legal Reasoning*. PhD thesis, Stanford University, 1984.