Knowledge Base Revision through Exception-driven Discovery and Learning

Seok Won Lee and Gheorghe Tecuci

Learning Agents Laboratory, Computer Science Department, MSN 4A5, George Mason University, Fairfax, VA 22030-4444, USA
{swlee, tecuci}@gmu.edu

We are currently witnessing a trend toward an architectural separation of a knowledge base (KB) into an ontology and a set of rules. The ontology is a description of the concepts and relationships from the application domain; the rules are problem solving procedures expressed with the terms from the ontology. Moreover, terminological standardization taking place in more and more domains has led to the development of domain ontologies. These two developments raise the prospect of reusing existing ontologies when building a new knowledge based system. For instance, the Disciple approach for building a knowledge based agent relies on importing ontologies from existing repositories of knowledge, and on teaching the agent how to perform various tasks, in a way that resembles how an expert would teach a human apprentice when solving problems in cooperation (Tecuci, 1998; Tecuci et al. 1999). In Disciple, the ontology serves as the generalization hierarchy for learning, an example being basically generalized to a rule by replacing its objects with more general objects from the ontology. However, the learning works well only if the ontology contains all the concepts needed to represent the application domain. We make the assumption that an ontology built from previously developed KBs will contain useful concepts, but it is incomplete and will not contain the more subtle distinctions needed for competent and efficient problem solving in a particular domain. These missing concepts will manifest themselves as *exceptions* to the learned problem solving rules. A negative exception of the rule is a negative example that is covered by the rule and any specialization (within the current ontology) of the rule that would uncover the exception would result also in uncovering of positive examples of the rule. Similarly, a positive exception of the rule is a positive example that is not covered by the rule and any generalization of the rule that would cover the exception would also result in covering of negative examples of the rule.

We are enhancing Disciple by developing a mixed-initiative multistrategy approach to KB revision that will result in an extended and domain-adapted ontology, as well as a set of rules with fewer (if any) exceptions. We are developing two classes of KB revision methods, a class of local methods and a class of global methods. The local methods focus on one rule with its exceptions at a time, in conjunction with the current ontology. Some of the local methods use analogical transfer of discriminating features from some objects to other objects in the positive examples of a rule, by considering the similarities between the positive examples and their dissimilarities with the negative exceptions. Other local methods use explanation-based techniques, and similarities between the current rule and other rules, to discover or elicit from the expert discriminating features in the form of explanations of why a negative exception of a rule is not a correct problem solving episode. These local methods work well when the ontology already contains the definitions of the discriminating features, but the descriptions of some of the objects from the ontology are incomplete with respect to those features. The methods perform a local extension of the ontology that leads to a refinement of a rule and a removal of some of its exceptions.

Other local methods do not immediately extend the ontology, but suggest characterizations of new concepts that would remove the exceptions. An example of such a characterization is the following: a concept that covers the maximum number of objects from a set of objects and the minimum number of objects from another set of objects. The global methods analyze the alternative concept characterizations suggested by the local methods and attempt to discover a reduced set of concept characterizations that would remove exceptions from a set of rules. In this process, they use various specialization and generalization operators that combine the concept characterizations into a reduced set. They then interact with the domain expert to identify which of the most useful characterizations correspond to meaningful concepts or features in the application domain. The global methods lead to the extension of the ontology with definitions of new objects and features.

A significant part of our research effort is also devoted to the evaluation of the developed approach by measuring the effectiveness of the exception-driven discovery and learning with respect to the number of exceptions removed, the impact of the discovered knowledge on the remaining rules, the knowledge acquisition effort required from the domain expert, and the effect of the discovered knowledge on the agent's performance on a set of tasks.

In conclusion, we are developing a suite of methods that continuously extend the ontology and revise the rules in the KB through discovery, learning and an interaction with a domain expert in order to achieve competent and efficient problem solving in a particular domain.

References

Tecuci, G. 1998. *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies.* Academic Press.

Tecuci, G., Boicu, M., Wright, K., Lee, S.W., Marcu, D., and Bowman, M. 1999. An Integrated Shell and Methodology for Rapid Development of Knowledge-Based Agents. In *Proc. AAAI-99*, Menlo Park, CA: AAAI Press.