Mixed-Initiative Reasoning for Integrated Domain Modeling, Learning and Problem Solving

Mihai Boicu and Gheorghe Tecuci
Learning Agents Laboratory, Dept. of Computer Science, MS 4A5
George Mason University, 4400 University Drive, Fairfax, VA 22030-4444
(703) 993-4669, URL: lalab.gmu.edu/mihaib, Email: mboicu@gmu.edu

The main challenge addressed by this research is the knowledge acquisition bottleneck defined as the difficulty of creating and maintaining a knowledge base that represents a model of the expertise domain that exists in the mind of a domain expert (Buchanan and Wilkins, 1993). The mixed-initiative approach we are investigating, called Disciple (Tecuci, Boicu et al. 1999), relies on developing a very capable agent that can collaborate with the domain expert to develop its knowledge base. In this approach both the agent and the expert are accorded responsibility for those elements of knowledge engineering for which they have the most aptitude, and together they form a complete team for knowledge base development. The domain modeling and problem solving approach is based on task reduction, the knowledge base to be developed consisting of an ontology that defines the terms from the application domain, and a set of task reduction rules expressed with these terms. The ontology is based on the OKBC knowledge model (Chaudhri et al. 1998) which facilitates the import of ontological knowledge from the OKBC compliant knowledge repositories, such as Ontolingua (Farquhar et al. 1996; Boicu et al. 1999). Each task reduction rule is learned by the agent through a mixed-initiative multistrategy learning method, starting from a specific example E_1 provided by the expert. Such a rule is a complex IF-THEN structure that specifies a plausible space for the conditions under which the task from the IF part can be reduced to the tasks from the THEN part. This space is represented by a plausible upper bound condition which, as an approximation, is more general than the exact (but not yet known) condition E_h, and a plausible lower bound condition which, as an approximation, is less general than E_h. The rule may also include several except-when conditions (that should not hold in order for the rule to be applicable), "except-for" conditions (that specify negative exceptions of the rule) and "for" conditions (that specify positive exceptions). The rule also includes generalizations of natural language phrases used by the expert to describe the example E_1.

The main focus of our research is the development of a powerful and flexible mixed-initiative plausible reasoner that allows the expert to train the agent in a variety of ways, and in as natural a manner as possible, similar to the way the expert would train a human apprentice. This reasoner exploits the structure of the ontology and of the plausible task reduction rules to integrate the domain modeling, learning and problem solving processes involved in developing the KB of the agent. The goal is to develop a knowledge base that will allow the agent to exhibit the same problem solving competence as the domain expert. We call the set of all correct solutions generated with this "final" knowledge base the *Target Solution Space* (see Fig.1). However, the current knowledge base of the agent is incomplete and may be partially incorrect. Therefore, part of the *Target Solution Space* is not even included in the *Current Representation Space* of the agent which will have to be extended by introducing new terms in the ontology.

The plausible reasoner allows the agent to distinguish between four types of increasingly complex problem solving situations: routine, innovative, inventive and creative. This capability guides the interaction with the domain expert, leading to a cooperative problem solving process where the agent solves the more routine parts of the problem and the expert solves the more creative ones. In this process the agent will learn from the expert improving its knowledge base.

The *routine solutions* are those that satisfy the plausible lower bound conditions of the task reduction rules, and are very likely to be correct. The *innovative solutions* are those that satisfy

the plausible upper bound conditions. These solutions may or may not be correct and have to be checked by the expert who can accept or reject them. These situations will lead to a refinement of the involved task reduction rules. The *inventive solutions* are based on weaker forms of plausible reasoning (such as partial matching of the plausible conditions of the rules, and tasks similarity based on the structure of the ontology). An inventive task reduction step is based on an analysis of several rules, and is generally a novel decomposition. From inventive solutions the agent will learn new plausible task reduction rules. Finally, the *creative solutions* are those that cannot even be expressed in the current agent's representation

Figure 1: Types of solutions generated through mixed-initiative reasoning

language. These solutions must be provided by expert. They will lead both to an extension of the ontology, and to the learning of new rules. As a result of this learning process, the problem solving situations that were innovative for the agent gradually become routine, and those that were creative, gradually become inventive, then innovative and ultimately routine.
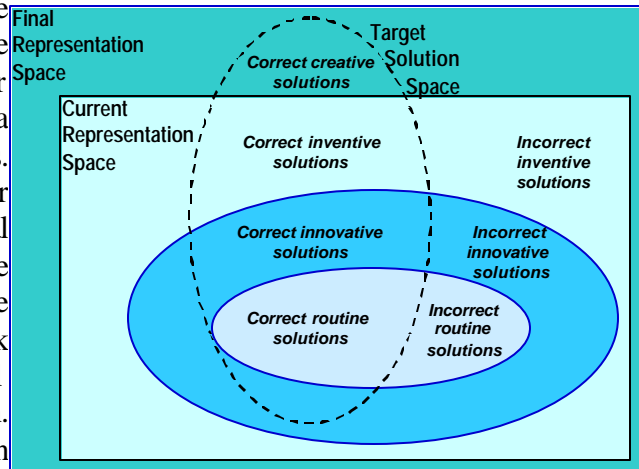
A very important feature of the mixed-initiative reasoner is that it fulfils multiple roles, supporting domain modeling, learning and problem solving, depending of the agent's knowledge. Initially, when the agent does not have much knowledge, the emphasis is on domain modeling where most of the problems require "creative" or "inventive" solutions. During this phase, the plausible reasoner supports the definition of the inventive solutions and the explanation-based learning of the rules. As the agent learns from the expert, it is increasingly able to propose routine and innovative solutions. During this phase the plausible reasoner supports solution generation and explanation-based rule refinement.

A version of this plausible reasoner has been implemented in Disciple-COA, which has been developed as part of the DARPA's High Performance Knowledge Bases program to solve the Course Of Action challenge problem (Tecuci, Boicu et al. 2000), and has been evaluated in two intensive studies. The first was a two week annual DARPA evaluation where Disciple-COA demonstrated a very high rate of knowledge acquisition and the best problem solving performance among all the developed COA critiquers. The second study was a one week knowledge acquisition experiment at the US Army Battle Command Battle Lab which demonstrated that domain experts that do not have knowledge engineering experience can quickly be trained to extend the knowledge base of Disciple-COA.

In summary, this mixed-initiative reasoner allows the achievement of several levels of synergism between the expert that has the knowledge to be formalized and the agent that is able to formalize it. At the highest level there is the synergism in solving complex problems, where the agent contributes routine and innovative problem solving steps and the expert contributes inventive and creative ones. At the next level down, there is the synergism between teaching and learning, where the expert helps the agent to understand the problem solving steps contributed by him or her, and the agent learns general problem solving rules that will allow it to apply similar steps in future problem solving situations. Finally, at the lowest level, there is the synergism between different learning strategies employed by the agent to learn from the expert in situations in which no single strategy learning method would be sufficient.

References

Alphatech, Inc. *HPKB Course of Action Challenge Problem Specification*, Burlington, MA, 1999

Boicu M., Tecuci G., Bowman M., Marcu D., Lee S.W. and Wright K., "A Problem-Oriented Approach to Ontology Creation and Maintenance," in *Proceedings of the Sixteenth National Conference on Artificial Intelligence Workshop on Ontology Management*, July 18-19, Orlando, Florida, AAAI Press, Menlo Park, CA. 1999.

Boicu M., Wright K., Marcu D., Lee S.W., Bowman M. and Tecuci G., "The Disciple Integrated Shell and Methodology for Rapid Development of Knowledge-Based Agents," in *Proceedings of the Sixteenth National Conference on Artificial Intelligence Workshop on Ontology Management*, July 18-19, Orlando, Florida, AAAI Press, Menlo Park, CA. 1999.

Buchanan, B.G. and Wilkins, D.C. (eds.), *Readings in Knowledge Acquisition and Learning: Automating the Construction and Improvement of Expert Systems*, Morgan Kaufmann, CA, 1993

Chaudhri, V. K., Farquhar, A., Fikes, R., Park, P. D., and Rice, J. P., OKBC: A Programmatic Foundation for Knowledge Base Interoperability. In *Proc. AAAI-98*, pp. 600 – 607, Menlo Park, CA: AAAI Press, 1998.

Cohen P., Schrag R., Jones E., Pease A., Lin A., Starr B., Gunning D., and Burke M., The DARPA High-Performance Knowledge Bases Project, *AI Magazine*, 19(4),25-49, 1998.

Farquhar, A., Fikes, R., and Rice, J., The Ontolingua Server: a Tool for Collaborative Ontology Construction. In *Proceedings of the Knowledge Acquisition for Knowledge-Based Systems Workshop,* Banff, Alberta, Canada, 1996.

Tecuci, G., Boicu, M., Wright, K., Lee, S.W., Marcu, D. and Bowman, M. An Integrated Shell and Methodology for Rapid Development of Knowledge-Based Agents, in *Proceedings of the Sixteenth National Conference on Artificial Intelligence Workshop on Ontology Management*, July 18-19, Orlando, Florida, AAAI Press, Menlo Park, CA. 1999.

Tecuci G., Boicu M., Wright K. and Lee S.W., "Mixed-Initiative Development of Knowledge Bases," in *Proceedings of the Sixteenth National Conference on Artificial Intelligence Workshop on Mixed-Initiative Intelligence,* Orlando, Florida, AAAI Press, Menlo Park, CA. 1999.

Tecuci G., Boicu M., Wright K., Lee S.W., Marcu D. and Bowman M., "A Tutoring Based Approach to the Development of Intelligent Agents," to appear in Teodorescu, H.N., Mlynek, D., Kandel, A. and Zimmermann, H.J. (editors). *Intelligent Systems and Interfaces,* Kluwer Academic Press. To appear.

Tecuci, G., Boicu, M., Marcu, D., Bowman, M., Ciucu F., and Levcovici C., Rapid Development of a High Performance Knowledge Base for Course of Action Critiquing, Submitted to the *Twelfth Innovative Applications of AI Conference,* Austin, TX, 2000.

Wright K., Boicu M., Lee S.W. and Tecuci G., Building agents from shared ontologies through apprenticeship multistrategy learning, *The Fifteenth National Conference on Artificial Intelligence, AAAI-98,* Student Abstract Poster Program, Madison, Wisconsin, Morgan Kaufmann, 1998.