

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

Mihai Boicu and Gheorghe Tecuci

Learning Agents Laboratory, Department of Computer Science, MS 4A5  
George Mason University, 4400 University Drive, Fairfax, VA 22030-4444  
(703) 993-4669, URL: lalab.gmu.edu, 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. The mixed-initiative approach we are investigating, called Disciple (Tecuci et al. 1999; Boicu et al., 2000), 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 paradigm. The knowledge base to be developed consisting of an OKBC-type ontology that defines the terms from the application domain, and a set of plausible task reduction rules expressed with these terms.

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 knowledge base of the agent. The goal is to develop a knowledge base that will allow the agent to exhibit a problem solving competence that is similar with that of the domain expert. We call the set of all correct solutions generated with this "final" knowledge base the Target Solution Space. 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.

A very important feature of the mixed-initiative reasoner is that it fulfills 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.

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: synergism in cooperative problem solving, synergism between teaching and learning, and synergism between different learning strategies.

## Acknowledgments

This work was supported by AFOSR and DARPA through the grants F49620-97-1-0188 and F49620-00-1-0072.

## References

- Boicu, M., Tecuci, G., Marcu, D., Bowman, M., Ciucu F., and Levcovici C. 2000. Disciple-COA: From agent Programming to Agent Teaching. In *Machine Learning: Proceedings of the Seventeenth International Conference*, San Francisco, Calif.: Morgan Kaufmann Pub.
- 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. of the Sixteenth National Conference on Artificial Intelligence*, 250-257, Menlo Park, Calif.: AAAI Press.