A Toolbox for Goal-driven Knowledge Acquisition

Derek Sleeman (DSLEEMAN@CSD.ABDN.AC.UK)
Simon White (SWHITE@CSD.ABDN.AC.UK)
Department of Computing Science, King's College,
University of Aberdeen, Aberdeen, AB24 3UE, Scotland, UK.

Introduction

Many tools and techniques have been developed for the systematic acquisition of domain knowledge, including knowledge elicitation methods for use with a human expert; machine learning algorithms that infer knowledge from data; and knowledge base refinement tools that refine existing knowledge bases. As the number and sophistication of these knowledge acquisition tools increases, it becomes increasingly difficult for users to choose between them for particular applications, especially when more than one is needed. We argue the importance of driving this whole process by the epistemological requirements of the problem solver(s) which have been selected to solve a particular task. To support this approach, we introduce a toolbox which includes an advisory system coupled to several knowledge acquisition tools and problem solvers.

Approach

The focus of this project is a toolbox called MUSKRAT (Multistrategy Knowledge Refinement and Acquisition Toolbox) which is able to reason about the suitability of knowledge acquisition (KA) tools for capturing the right knowledge content and form. Given a choice of problem-solver, the system compares the currently available knowledge with the problem solver's epistemological requirements. A shortfall at this stage defines a knowledge acquisition process for which advice to the user must be generated. We therefore incorporate knowledge-level descriptions (Newell, 1982) of KA tools and problem solvers so that MUSKRAT can provide such advice. Although knowledge-level descriptions have been applied to learning tasks, they have not been applied, to our knowledge, to the KA tool selection task. Furthermore, since the whole process is directed towards the successful execution of a problem solver, we are providing a mechanism for *goal-driven knowledge acquisition*.

We are illustrating our ideas by building a prototype in the domain of planning and cooking a meal. This domain is not only challenging in itself, but also carries heightened significance through an analogy with just-in-time manufacturing. For our implementation, we have identified three distinct problem solvers: a constraint satisfier, which composes a set of menus consistent with some given constraints; a design analyst, which analyses a set of menus subject to some preferences; and a scheduler, which, given some time and resource constraints, plans the preparation of a given menu. This current implementation is provided with the necessary knowledge to perform these three tasks (Sleeman & White, 1996). In future we will experiment with knowledge bases which are inconsistent and/or incomplete, thus forcing MUSKRAT to apply a range of KA tools to produce knowledge bases in the formats required by the several problem solvers.

For the representation of knowledge, we are using the Common Knowledge Representation Language (CKRL), which emerged as the generic knowledge description language of the ESPRIT Machine Learning Toolbox project (Kodratoff et al., 1992). CKRL was devised specifically for enabling the interchange of knowledge among heterogeneous knowledge acquisition components.

Summary

We have proposed a computational schema for knowledge acquisition and problem solving. We are currently implementing a prototype in the domain of meal preparation which incorporates three problem solvers and a number of KA tools. The research is relevant to Cognitive Science as it promises greater efficacy in knowledge acquisition for problem solving, but principally because it embodies a goal-directed model of the problem solving/KA process.

Acknowledgements

This research is supported by an EPSRC studentship.

References

Kodratoff, Y., Sleeman, D., Usynski, M., Causse, K., Craw, S., (1992). Building a Machine Learning Toolbox. In L. Steels & B. LePape (Eds.), *Enhancing the Knowledge Engineering Process - Contributions from ESPRIT,* North Holland: Elsevier Science Publishers.

Newell, A. (1982). The Knowledge Level. *Artificial Intelligence,* 18(1), 87-127.

Sleeman, D., & White, S. (1996). *A Multistrategy Knowledge Refinement and Acquisition Toolbox: Revisited* (Tech. Rep. AUCS/TR9605). Aberdeen, Scotland, UK: University of Aberdeen, Department of Computing Science.