Exploiting prior knowledge in Intelligent Assistants - Combining relational models with hierarchies

Sriraam Natarajan, Prasad Tadepalli, and Alan Fern

School of EECS,
Oregon State University

There has been a growing interest in developing intelligent assistant systems that help users in a variety of tasks. The emphasis in these systems has been to provide a well-engineered domain-specific solution to the problem of reducing the users' cognitive load in their daily tasks. A decision-theoretic model was proposed recently [1] to formalize the general problem of assistantship as a partially observable Markov decision process (POMDP). In this formulation, there is a goal-oriented user and an assistant acting interactively in the environment. The goal of the assistant is to take actions that minimize the expected cost of completing the user's task. In most situations, however, the user's task or goal is not directly observable to the assistant, which makes the problem of quickly inferring the user's goals from observed actions critically important. To perform this goal inference, it is important to learn the user's policy as early as possible. In our previous work, we assumed that the user is reasonably rational to constrain his policy. Also, we assumed a flat user policy to perform effective inference.

In this work, we aim to use the combination of hierarchical and relational knowledge about the user's goal structure to constrain his policy. For instance, a user who submits a paper would decompose the goals into writing the abstract, the main paper, run experiments, compile the results and turn in the paper. Also, the user would use the same methodology irrespective of whether he turns in a paper to ICML or IJCAI. Similarly, the choice of whether he runs the experiments or writes the main section would be influenced by the closeness of deadline. We believe that an assistant equipped with such a relational hierarchical knowledge would be able to provide effective assistance to the user.

Our current work extends the assistantship model [1] to hierarchical and relational settings, building on the work in hierarchical reinforcement learning and statistical relational learning [3, 4]. We extend the assistantship framework by including parameterized task hierarchies and conditional relational influences as prior knowledge of the assistant. An example of parameterized task hierarchy is presented in Figure 1. We refer the reader to [2] for the semantics and execution of these hierarchies. We compile this knowledge into an underlying Dynamic Bayesian network and use Bayesian network inference algorithms to infer the distribution of user's goals given a sequence of their atomic actions. The DBN that is obtained for inferring the user's goal is similar to the ones used for plan recognition [5]. We estimate the parameters for the user's policy and influence relationships by observing the users' actions. Once the user's goal distribution is inferred, we determine an approximately optimal action by estimating the Q-values of different actions using rollouts and picking the action that has the least expected cost. We evaluate our relational hierarchical assistantship model in two

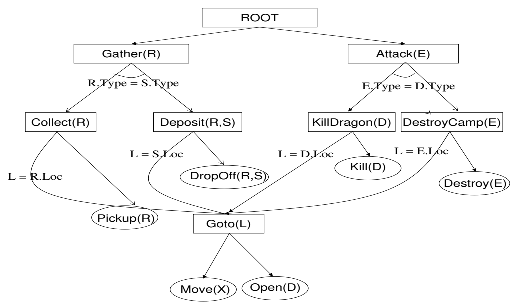

Fig. 1. Example of a task hierarchy of the user. The inner nodes indicate subtasks while the leaves are the primitive actions. The tasks are parameterized and the tasks at the higher level call the tasks at the lower level

different toy domains and compare it to a propositional flat model, propositional hierarchical model, and a relational flat model. Through simulations, we show that when the prior knowledge of the assistant matches the true behavior of the user, the relational hierarchical model provides superior assistance in terms of performing useful actions. The relational flat model and the propositional hierarchical model provide better assistance than the propositional flat model, but fall short of the performance of the relational hierarchical approach. We refer the user to [2] for a detailed discussion of the experimental setup and the results.

In our current work, we unrolled the observations and goal structure into a ground DBN. Though this is justified in many domains, the inference could be computationally expensive in many domains. An important future work is to develop faster inference methods that do not need full unrolling. To this effect, we are currently working on dynamic models that can avoid full grounding. Yet another important future work is to improve the action selection mechanism of our model and use methods that can exploit the goal structure of the user.

References

1. Alan Fern and Sriraam Natarajan and Kshitij Judah and Prasad Tadepalli: A Decision-Theoretic Model of Assistance, IJCAI, 2007
2. Sriraam Natarajan and Prasad Tadepalli and Alan Fern: A Relational Hierarchical Model for Decision-Theoretic Assistance, In proceedings of 17th Annual International Conference on Inductive Logic Programming - To appear, 2007
3. Tom Dietterich: Hierarchical Reinforcement Learning with the MAXQ Value Function Decomposition, JAIR, volume 13, 227-303, 2000,
4. Tom Dietterich and Lise Getoor and Kevin Murphy: Statistical Relational Learning and its Connections to Other Fields, ICML 2004 Workshop, Banff Canada
5. Hung Bui and Svetha Venkatesh and Geoff West: Policy recognition in the Abstract Hidden Markov Models, JAIR, volume 17, 2002