Skill Bootstrapping: Hierarchical Skill Learning For High-level Planning

James MacGlashan (jmac1@cs.umbc.edu)
Marie desJardins (mariedj@cs.umbc.edu)

Skill Bootstrapping (SB)

• SB is a proposed framework and set of algorithms to allow an agent to start with a set of low-level primitive actions, and then abstract its actions allowing it to tractably solve complex problems.

• An agents starts by solving simple problems using planning, and then identifies *skills* that could be used as generalized and adaptive macro-actions during planning.

• Skills are learned using RL with function approximation, and learning is bootstrapped with previous successful plan traces that have previously solved the same kind of problem.

Motivating Properties

• Planning for new problems first, and seeding learning with the successful traces, improves the quality of newly created skill policies.

• Learning generalized and parameterized skills allows a single skill to decompose into many possible action sequences.

• Using a reactive skill policy means planning is required at only the abstract-level.

• Viable skill structures can be revealed by analyzing the successful plan traces.

• Using a top-level planner allows skills to be shared in many different hierarchical structures that may each solve very different problems.

Research Progress

This is a preliminary proposal for dissertation work by James MacGlashan, a Ph.D. student attending the University of Maryland Baltimore County. Currently, many design issues still need to be solved. We present here the preliminary ideas about how to create the Skill Bootstrapping architecture.

Related Work

Macro-Action Learning

• Botea, A., Enzenberger, M., Muller, M., & Schaeffer, J. (2005). Macro-FF: Improving AI planning with automatically learned macro-operators. Journal of Artificial Intelligence Research, 24, 581–621.

• Newton, M., Levine, J., & Fox, M. (2005). Genetically evolved macro-actions in AI planning problems. Proceedings of the 24th UK Planning and Scheduling SIG, 163–172.

• Coles, A., Fox, M., & Smith, A. (2007). Online identification of useful macro-actions for planning. Proceedings of the International Conference on Automated Planning and Scheduling.

RL Structure Learning

• Mehta, N., Ray, S., Tadepalli, P., & Dietterich, T. (2008). Automatic discovery and transfer of MAXQ hierarchies. Proceedings of the 25th International Conference on Machine Learning (pp. 648–655).

• Simsek, O., & Barto, A. (2007). Betweenness centrality as a basis for forming skills (Technical Report). University of Massachusetts, Department of Computer Science.

• G.D. Konidaris and A.G. Barto. Skill Discovery in Continuous Reinforcement Learning Domains using Skill Chaining. Technical Report UM-CS-2008-24, Department of Computer Science, University of Massachusetts Amherst, July 2008.

A Possible Scenario

An agent has been placed in a room with objects it can interact with to solve a number of different problems. With primitive actions consisting of basic movement of the agent's components, the agent has developed many abstract skills such as walking, and grabbing.

A High-Level Goal and Problem

There is a banana hanging from the ceiling of the room that the agent needs to grab, and there are various boxes scattered around the room.

A Possible Top-level Plan and Skill Decomposition

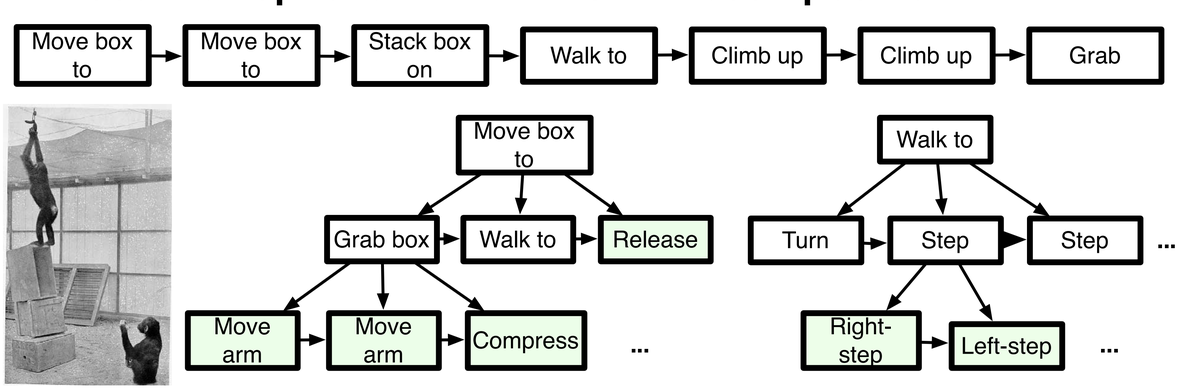

Architecture