UNIVERSITY OF MASSACHUSETTS AMHERST

Resource-Bounded Reasoning Lab

Contact Information

Computer Science Building
University of Massachusetts
Amherst, MA 01003-9264
(413) 545-4189 (voice)
(413) 545-1249 (fax)
http://rbr.cs.umass.edu

Principal Investigator

Prof. Shlomo Zilberstein
(413) 545-4189
shlomo@cs.umass.edu

Postdocs

Dr. Siddharth Srivastava
Dr. William Yeoh

Graduate Students

Alan Carlin
Akshat Kumar
Thomas Rossi
Xiaojian Wu

Undergraduates

Cathy Zhang

Grant Administrator

Michele Roberts
(413) 545-4389
michele@cs.umass.edu

Lab Location

326 CS Building
(413) 545-1985

Research Areas

- anytime algorithms
- automated planning
- autonomous systems
- bounded rationality
- coordination
- cost-sensitive learning
- decentralized planning
- decision theory
- heuristic search
- metareasoning
- models of rational agency
- multi-agent systems
- probabilistic reasoning
- reinforcement learning

The Resource-Bounded Reasoning Lab conducts research on the computational foundations of automated reasoning and planning. We are particularly interested in the design of autonomous agents that can cope with uncertainty and limited computational resources. In most practical settings, it is not feasible to find the optimal course of action, making it necessarily to resort to some form of approximate reasoning. This raises a simple fundamental question: what does it mean for an agent to be "rational" when it does not have enough knowledge or computational power to derive the best action? Our overall approach to this problem is based on probabilistic reasoning and decision-theoretic principles, used both to develop planning algorithms and to monitor their execution and maximize the value of computation. We study meta-level deliberation processes that reason explicitly about the cost of decision-making and can optimize the amount of "thinking" an agent does before taking action. This research spans both theoretical issues and the development of effective algorithms and applications. We also study novel models and algorithms for planning in situations involving multiple decision makers operating in either collaborative or adversarial domains.

CURRENT RESEARCH PROJECTS

Decision-Theoretic Foundations for Multi-Agent Systems

Developing new models (decentralized generalizations of MDPs) and algorithms for planning in situations that involve multiple decision makers, each having different partial information about the domain. In collaboration with the Multi-Agent Systems Lab (UMass). Sponsored by AFOSR and NSF.

Adaptive Optimization Techniques for Large Scale Stochastic Planning

Developing robust algorithms for planning under uncertainty in very large domains using approximate linear programming. Sponsored by AFOSR.

Foundations and Applications of Generalized Planning

Developing rich new plan representations that include branches and loops, and creating new ways to produce automatically algorithm-like plans that can apply to a large set of problem instances. In collaboration with Neil Immerman (UMass). Sponsored by NSF.

Resource-Bounded Reasoning Techniques for Autonomous Systems

Developing anytime planning algorithms that offer a tradeoff between computational resources and solution quality, and developing mixed human-robot control for semi-autonomous systems. In collaboration with INRIA and the University of Caen (France).

Online Planning for Decentralized POMDPs

Developing effective techniques for online coordination of multi-agent systems with limited communication. In collaboration with USTC (China).

RECENT PUBLICATIONS

S. Srivastava, N. Immerman, and S. Zilberstein. "A New Representation and Associated Algorithms for Generalized Planning." *Artificial Intelligence*, 175(2):615–647, 2011.

F. Wu, S. Zilberstein, and X. Chen. "Online Planning for Multi-Agent Systems with Bounded Communication." *Artificial Intelligence*, 175(2):487–511, 2011.

A. Carlin and S. Zilberstein. "Decentralized Monitoring of Anytime Decision Making." *Proc. of the 10th International Conf. on Autonomous Agents and Multiagent Systems*, Taipei, Taiwan, 2011.

C. Amato, D. S. Bernstein, and S. Zilberstein. "Optimizing Fixed-Size Stochastic Controllers for POMDPs and Decentralized POMDPs." *Autonomous Agents and Multi-Agent Systems*, 21(3):293–320, 2010.

A. Kumar and S. Zilberstein. "Convergent MAP Estimation for Graphical Models by Likelihood Maximization." *Proc. of Neural Information Processing Systems* (NIPS), Vancouver, Canada, 2010.

S. Srivastava, N. Immerman, S. Zilberstein. "Computing Applicability Conditions for Plans with Loops", *Proc. of the 20th International Conf. on Automated Planning and Scheduling*, Toronto, Canada, 2010.

M. Petrik and S. Zilberstein. "A Bilinear Programming Approach for Multiagent Planning", *Journal of Artificial Intelligence Research*, 35:235–274, 2009.

R. Becker, A. Carlin, V. Lesser, and S. Zilberstein. "Analyzing Myopic Approaches for Multiagent Communications." *Computational Intelligence*, 25(1):31–50, 2009.

A. Kumar and S. Zilberstein, "Event-Detecting Multi-Agent MDPs: Complexity and Constant-Factor Approximation", *Proc. of the 21st International Joint Conf. on Artificial Intelligence,* Pasadena, CA, 2009.

A. Kumar and S. Zilberstein, "Constraint-Based Dynamic Programming for Decentralized POMDPs with Structured Interactions", *Proc. of the 8th International Conf. on AAMAS*, Budapest, Hungary, 2009.

S. Seuken and S. Zilberstein. "Formal Models and Algorithms for Decentralized Decision Making Under Uncertainty", *Autonomous Agents and Multi-Agent Systems*, 17(2):190–250, 2008.