

Research Statement

Pradeep Varakantham
School of Information Systems, Singapore Management University
Tel: (65) 6828-0519; Email: pradeepv@smu.edu.sg
Updated on 10 (Day) 02 (Month) 2012 (Year)

Background

My research goal is to build large-scale intelligent systems (both single- and multi-agent) that make decisions or provide decision support in the face of uncertainty in complex, real world environments. I foresee an integration of such systems in many critical facets of human life:

- (a) decision support agents for improving operational efficiency of a taxi fleet,
- (b) cyber defense agents to counter different types of computer virus,
- (c) intelligent assistants in hospitals to offices,
- (d) rescue-agents in large scale disaster response,
- (e) sensor-agents tracking weather phenomena in earth observing sensor webs,
- (f) multi-agent coordination in large scale logistics (like humanitarian missions, national day parades etc.)

and others. My long-term plan is to make this integration successful, primarily by understanding the fundamental principles of decision making in the presence of uncertainty and applying these principles in building single- and multi-agent systems.

Two major obstacles in building such systems are choosing the right models to represent the underlying uncertainty (about the world and humans) and reasoning with these models. Uncertainty is the only certainty there is, when dealing with complex, real-world environments. It can arise due to numerous factors: (a) non-deterministic outcomes of actions; (b) limited or faulty sensors; (c) durations of actions (d) incomplete model of the environment; (e) uncertainty about human behavior; and (e) uncertainty in the actions or motivations of other agents (when interacting with other agents). Thus, providing models expressive enough to represent decision problems accurately in the context of one or more of these types of uncertainties is a key problem that needs to be addressed. Choosing an expressive model however leads to the second major obstacle: the more expressive the model, the higher the complexity involved in solving the model and the less expressive the model (not representing all the uncertainties), the more difficult it is to capture the domain accurately (required if these systems are to make decisions that affect our daily life). Due to the significant computational complexity involved in solving these models optimally, efficiency can often be achieved only by sacrificing solution quality.

This leads to a fundamental tradeoff between expressivity of the model, efficiency of techniques and being able to provide guarantees on solution quality. My research

focuses on providing approximations that efficiently solve accurate representations of complex domains while providing theoretical or empirical guarantees on solution quality. Thus far, in my research, I have addressed this tradeoff in the context of uncertain domains that have uncertainty arising due to a combination of the following factors: (a) non-deterministic outcomes to actions; (b) faulty sensors; (c) uncertain action durations; (d) uncertainty regarding the actions of other agents.

Current Research

Single-agent domains: Within the single agent context, my contributions can be divided into two categories based on the problem studied:

(i) ***Scheduling with uncertainty:*** In [1], we considered the problem of assisting a busy user, who is invariably juggling multiple concurrent streams of tasks (or workflows) of varying importance and urgency, in managing his/her workload of pending tasks. There is uncertainty with respect to the duration of a pending task as well as the amount of follow-on task load that may be generated as a result of executing the task. In [2], we have considered a more general problem of scheduling with multiple resources in the face of durational uncertainty. Our goal here was to provide robust schedules without violating resource constraints.

(ii) ***Planning with uncertainty:*** In [3], the focus here was to provide quality bounded planning policies efficiently in domains where there is uncertainty due to non-deterministic outcomes of actions, noisy observations and unknown starting state and the goal is to maximize expected reward. In [4], we have considered high stakes domains (such as financial planning, gambling, and rescue scenarios) where expected utility (and not expected reward) needs to be maximized. In [5], we are considering domains where there is an additional uncertainty about when the observations are received. In [9] we are focused on large-scale organizations and the structure in the decision problem corresponding to those organizations.

Multi-agent domains: Most of my earlier and current research deals with multi-agent problem domains with heterogeneous teams of agents. Distributed POMDPs are important to model domains that have multiple coordinating agents, which are required to take decisions in the presence of uncertainty due to nondeterministic outcomes of actions, noisy observations and unknown starting state. I have addressed this issue by providing a combination of techniques that allow for quality bounded and efficient strategy computation in distributed POMDPs. My current work has focused on an approach called TREMOR [6, 7, 8, 9], which scales to large-scale teams of agents while providing good quality solutions. The promise of this approach lies in intelligently searching in the space of single agent policies to obtain a good joint policy.

Lately, we have also employed a multi-agent learning algorithm for performing active malware analysis [11]. This is an approach that is applicable at large in cyber security problems.

Future research

For agents and multi-agent systems to finally break out in the real-world, in a very fundamental sense, they must conquer uncertainty and reason with human models. In the future, I would like to build upon my research towards understanding the reasoning process in ever more realistic environments. There are three key research directions, I would like to pursue in the future:

(a) **Objective:** Most existing research has focused on providing expected optimal solutions. However, there are other meaningful properties of the objective functions in real world domains that need to be handled. For instance,

1. **Robustness:** In domains such as disaster response, it is desired that the decisions taken err on the side of caution when faced with uncertainty.
2. **Stability:** In systems that provide decision support to humans, it is desirable that the guidance provided changes minimally over time so as not to confuse the human. The research theme here will be to develop approaches that will optimize over one or more of these criteria. As part of addressing interesting practical problems in this direction, we have developed an algorithm for providing decision support in a taxi fleet, such that the revenue for the taxi drivers is improved, while at the same time reducing the congestion or starvation of taxis.

(b) **Complexity of the constraints:** Constraints such as resource boundedness, are a regular in many real domains, increase the complexity of the decision problem considerably. Providing scalable approaches that handle such constraints on the decision problems will be the key aspect of this thread of research.

(c) **Reason with human behavior models:** For agents to operate in the real world, the decision making process has to account for humans or in other words boundedly rational agents. My goal here is to build boundedly rational models for humans from behavioral data sets and use these models in computing strategies for agents.

I believe that understanding the process of decision making in these critical settings and utilizing this knowledge towards **building practical** intelligent agent/multi-agent systems will result in a smooth transition of intelligent systems into our daily life.

Selected Publications and Outputs

- [1] P. Varakantham and S. F. Smith. *Linear Relaxation Techniques for Task Management in Uncertain Settings*. In Proceedings of the Eighteenth International Conference on Automated Planning and Scheduling (ICAPS), 2008.
- [2] P. Varakantham, R. Maheswaran, T. Gupta and M. Tambe. *Towards efficient computation of quality bounded solutions*. in POMDPs. In Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI), 2007.
- [3] N. Fu, P. Varakantham and H.C. Lau. *Towards Finding Robust Execution Strategies for RCPSP with Durational Uncertainty*. In Proceedings of the Twentieth International Conference on Automated Planning and Scheduling (ICAPS), 2010
- [4] J. Marecki and P. Varakantham. *Risk sensitive Planning in Partially Observable Domains*. In Proceedings of the International Conference on Autonomous Agents and Multi Agent Systems, 2010 (to appear).
- [5] P. Varakantham and J. Marecki. *Planning with delayed observations in partially observable domains*. In Proceedings of the Eleventh International Conference on Autonomous Agents and Multi-Agent Systems, 2012.
- [6] P. Varakantham, J. Y. Kwak, M. Taylor, J. Marecki, P. Scerri and M. Tambe. Exploiting Coordination Locales in Distributed POMDPs via Social Model Shaping. In Proceedings of the Nineteenth International Conference on Automated Planning and Scheduling (ICAPS), 2009.
- [7] P. Velagapudi, P. Varakantham, P. Scerri. Towards distributed planning in very large scale teams. In Proceedings of the Tenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2011.
- [8] P. Varakantham. Improved planning in Distributed POMDPs with TREMOR. Submitted to the International Joint Conference on Artificial Intelligence (IJCAI), 2010.
- [9] P. Varakantham, N. Schurr, A. Carlin and C. Amato. "Adaptive Decision Support for Structured Organizations: A Case for OrgPOMDPs", Proceedings of the Tenth International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)
- [10] P. Varakantham, W. Yeoh, P. Velagapudi, K. Sycara and P. Scerri. "Prioritized Model Shaping for Solving DEC-POMDPs", Proceedings of the Eleventh International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS)
- [11] S. Williamson, P. Varakantham, D. Gao and O. Chen Hui. "Active Malware Analysis using Stochastic Games". Proceedings of the Eleventh International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS).