Decision-theoretic Planning under Uncertainty for Active Cooperative Perception

Matthijs T.J. Spaan and **Pedro U. Lima**
Institute for Systems and Robotics
Instituto Superior Técnico
Av. Rovisco Pais, 1, 1049-001
Lisbon, Portugal
{mtjspaan,pal}@isr.ist.utl.pt

Abstract

This extended abstract summarizes our ideas on developing decision-theoretic planning approaches for active perception in networked robot systems.

Introduction

In our work we consider networked robot systems (NRS), in which robots interact with each other as well as with sensors present in the environment to accomplish certain tasks. For instance, in urban pedestrian areas, we can consider a group of robots assisting humans. The primary task of the robots could be to identify persons in need of assistance, and subsequently help them, for instance by guiding to a desired location. The pedestrian area in which the robots operate is equipped with surveillance cameras providing the robot with more information. Implementing such a system requires addressing many scientific and technological challenges, but in our work, we focus on one particular problem, namely *active cooperative perception*.

Active Cooperative Perception

In our context, cooperative perception refers to the fusion of sensory information between the fixed surveillance cameras and each robot, with as goal maximizing the amount and quality of perceptual information available to the system. Active perception means that an agent considers the effects of its actions on its sensors, and in particular it tries to improve their performance. This can mean selecting sensory actions, for instance pointing a pan-and-tilt camera or choosing to execute an expensive vision algorithm; or to influence a robot's path planning, e.g., given two routes to get to a desired location, take the more informative one. Combining the two concepts, active cooperative perception is the problem of active perception involving multiple sensors and multiple cooperating decision makers.

In our work, we consider decision-theoretic approaches to active cooperative perception, in particular Partially Observable Markov Decision Processes (POMDPs) (Kaelbling et al., 1998). POMDPs provide an elegant way to model the interaction of an active sensor with its environment. Based on prior knowledge of the sensor's model and the environment dynamics, we can compute policies that tell the active sensor how to act, based on the observations it receives. As we are essentially dealing with multiple decision makers, it could also be beneficial to consider modeling (a subset of) sensors as a decentralized POMDP (Dec-POMDP). The fact that sensors and robots are embedded in an environment that is physically distributed allows for applying decentralized planning under uncertainty methods that exploit such locality of interaction (Oliehoek et al., 2008).

In a cooperative perception framework, an important task encoded by the (Dec-)POMDP could be to reduce the uncertainty in its view of the environment as much as possible. Entropy can be used as a suitable measure for uncertainty. However, using a POMDP solution, we can tackle more elaborate scenarios, for instance in which we prioritize the tracking of certain objects. In particular, POMDPs inherently trade off task completion and information gathering. Sensory actions might also include other sensors, as we can reason explicitly about communicating with other sensors. For instance, a fixed sensor could ask a mobile sensor to examine a certain location.

As an example, Figure 1 illustrates how a simple task in which a robot needs to meet a person can be modeled as a factored POMDP. Observations about the person's location are provided by ceiling-mounted cameras, while the robot's localization uses its laser-range finder.

Applications

We have been applying these ideas in several problems. First, we considered the problem of dynamic sensor selection in camera networks (Spaan and Lima, 2009). Given the large resource demands of imaging sensors in terms of bandwidth and computing power, processing image streams of many cameras simultaneously might not be feasible. We proposed a decision-theoretic approach modeled as a POMDP, which selects k sensors to consider in the next time frame, incorporating all observations made in the past. We showed how we can model this problem as a POMDP, and how we can encode objectives such as maximizing coverage or improving localization uncertainty, and we successfully applied our techniques in our testbed with 10 cameras (Barbosa et al., 2009).

To our knowledge not many papers have been published

ICAPS 2010 POMDP Practitioners Workshop, May 12, 2010, Toronto, Canada.

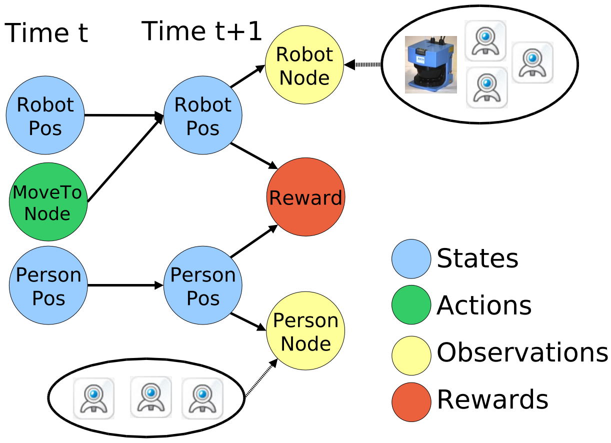

(a) Model.

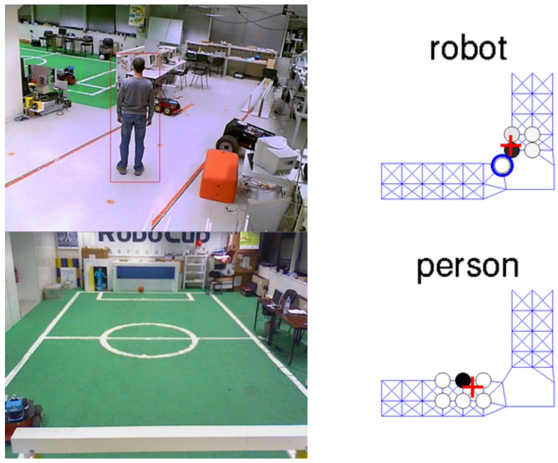

(b) Left: camera images, right: belief.

Figure 1: A factored POMDP model for a simple robot-meets-person task (a) and a snapshot during policy execution (b).

applying closed-loop non-myopic POMDP solutions to sensor selection problems. Although it is known that many related problems can be formulated as POMDPs, the complexity of solving continuous-state POMDPs in closed form has obstructed their solution. In our case, we tackle this problem by discretizing the state space, however, the final output of the system is still a continuous localization estimate. Related is the work of Krishnamurthy and Djonin (2007), who study so-called threshold policies for POMDPs for sensor scheduling. A crucial difference is that they consider reward functions that are not linear in the belief state, for instance based on the entropy of the belief state. In this case, the POMDP is nonstandard, and the optimal value function is no longer piecewise linear and convex. By defining a standard reward function over states, we remain in the standard POMDP setting, for which many results are known and successful approximate algorithms have been developed. Ji et al. (2007) present a POMDP formulation of the problem of multiaspect sensing on a single platform. A key difference is that we learn an observation model from data, instead of computing it from a physical model. Learning the model from data is potentially more reliable, as it will take into account limitations of the sensor or the event detection algorithm.

Our current work combines classification decisions with tracking uncertainty. We consider an environment with mobile and fixed sensors, where mobile sensors increase the observability of the system at a particular location but, as a counterpart, have an associated cost of moving. Furthermore, robot-mounted sensors will take time to arrive a particular destination, and are a scarce resource: given a limited number of robots available, the system needs to choose carefully when to send a robot where. The system must tackle at the same time robot and person localization as well as event detection and characterization, to have enough information to drive the robot to check an event at any location accurately.

Finally, in related works we exploited POMDP representations of tasks for multirobot task assignment (Spaan et al., 2010), as well as MDP models for sensor-aware path planning problems (Pahliani et al., 2009).

Acknowledgments This work was supported by the European Project FP6-2005-IST-6-045062-URUS, and was funded by Fundação para a Ciência e a Tecnologia (ISR/IST pluriannual funding) through the PIDDAC Program funds as well as by project PTDC/EEA-ACR/73266/2006.

References

M. Barbosa, A. Bernardino, D. Figueira, J. Gaspar, N. Gonçalves, P. U. Lima, P. Moreno, A. Pahliani, J. Santos-Victor, M. T. J. Spaan, and J. Sequeira. ISRobotNet: A testbed for sensor and robot network systems. In *Proc. of International Conference on Intelligent Robots and Systems*, pages 2827–2833, 2009.

S. Ji, R. Parr, and L. Carin. Non-myopic multi-aspect sensing with partially observable Markov decision processes. *IEEE Trans. Signal Processing*, 55(6):2720–2730, 2007.

L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.

V. Krishnamurthy and D. Djonin. Structured threshold policies for dynamic sensor scheduling–a partially observed Markov decision process approach. *IEEE Transactions on Signal Processing*, 55(10):4938–4957, 2007.

F. A. Oliehoek, M. T. J. Spaan, S. Whiteson, and N. Vlassis. Exploiting locality of interaction in factored Dec-POMDPs. In *Proc. of Int. Conference on Autonomous Agents and Multi Agent Systems*, pages 517–524, 2008.

A. Pahliani, M. T. J. Spaan, and P. U. Lima. Decision-theoretic robot guidance for active cooperative perception. In *Proc. of International Conference on Intelligent Robots and Systems*, pages 4837–4842, 2009.

M. T. J. Spaan and P. U. Lima. A decision-theoretic approach to dynamic sensor selection in camera networks. In *Int. Conf. on Automated Planning and Scheduling*, pages 279–304, 2009.

M. T. J. Spaan, N. Gonçalves, and J. Sequeira. Multirobot coordination by auctioning POMDPs. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2010. To appear.