Session Overview
Planning

Nicholas Roy[1] and Roland Siegwart[2]

[1] MIT
[2] ETZ Zürich

When we discuss autonomous robots, we think of robots that move around, interacting with people and making changes in the world. The problem of actually choosing motor commands to achieve high level goals — such as moving to a desired destination or answering a query from a human — typically involves planning. Planning is of course one of the central questions of artificial intelligence, and the planning field has moved a long way from the early days when planning meant searching for a sequence of abstract actions that satisfied some symbolic predicate. Robots can now learn their own representations through statistical inference procedures, they can now reason using different representations and they can reason in worlds where action can have stochastic outcomes.

However, despite the successes of robots that use machine learning and statistical inference in such different areas as mapping, speech recognition, computer vision, etc., there remain open questions to be addressed before we will see ubiquitous, useful, mobile robots, and some of the most interesting problems are in the planning domain. Consider a mobile robot deployed in some populated environment such as the home. A human operator typically drives the robot around in order to collect sensor data. This data is then used to build a "good" map that is largely static. The robot planner then computes good paths through this map, assuming that the map is correct and complete. The planning system rarely has the ability to reason about the robot's position within the map and how different plans may lead to better or worse localization. The planner almost never has ability to reason about the quality of the map itself and plan to gather more data in order to get a better map. In contrast, a planner that can reason about how much it knows about the world, and can plan to learn more when necessary, is likely to be a much more robust and general system.

What is becoming clear as robots become increasingly sophisticated is that there are three key issues in planning for mobile robots. Firstly, robots must be able to reason about uncertainty at all levels, both in the current state but also in the current representation. Secondly, robots must be able to plan in extremely high dimensional spaces. Thirdly, robots must be able to plan in populated and dynamic spaces. Each of these three issues is addressed by the papers in the planning section of ISRR; these papers shed new light on these problems and provides new tools for autonomous robots.

Planning with uncertainty Pineau & Gordon's paper describes the PEMA algorithm for solving large Partially Observable Markov Decision Processes (POMDPs), in which planning decisions are made with respect to the full probability distribution over the state space. POMDPs in particular have been considered computationally intractable for any real world problems, but this paper demonstrates that good approximation techniques can be used to generate plans that lead to overall more robust performance for robots in uncertain worlds. Additionally, the PEMA algorithm addresses a fairly important problem of how a planning algorithm should reason about its model. PEMA uses sampled beliefs, or probability distributions, in the planning process; PEMA demonstrates an approach to choosing these samples intelligently, improving the overall plan.

Planning in high-dimensional spaces In order to find plans in high-dimensional problems, conventional discretization techniques have been superseded by techniques that sample configurations from the world and then retain only those samples that are useful configurations. Hsu, Latombe and Kurniawati address some important questions at the heart of stochastic sampling planners, in particular why these techniques work well, and they describe theoretically why some variants of the sampling techniques have not represented improvements. The critical issue is to recognize that the the closer the sampling measure is to the desired plan, the better the performance. Most sampling techniques are a long way from achieving this goal, but this paper points the way to developing even more efficient planners.

Planning in populated worlds Finally, Alami et al.'s presentation on planning in human environments highlighted the need to start building human models into autonomous systems. For example, being able to deal with unpredictable people safely is a critical issue, and one of the results in this paper describes a motion planning algorithm with the objective of safety around people. Additionally, knowing how to behave reasonably around people in highly ambiguous situations is also essential.

It is worth pointing out that all three of these topics are highly related. Planning under uncertainty inevitably leads to planning in high-dimensional information spaces. Planning around people inevitably requires planning under uncertainty. These ideas will be essential for furthering the field of autonomous robots.