Published in the Workshop on Progress and Open Problems in Motion Planning at the IEEE/RSJ Conf. on Intelligent Robots and Systems

Look Before You Leap: Predictive Sensing and Opportunistic Navigation

D. K. Grady*, M. Moll*, C. Hegde†, A. C. Sankaranarayanan†, R. G. Baraniuk†, and L. E. Kavraki*
*Dept. of Computer Science, Rice University, Houston, TX 77005, USA. Email: {dkg1,mmoll,kavraki}@rice.edu
†Dept. of Elec. & Computer Engineering, Rice University, Houston, TX 77005, USA. Email: {ch3,saswin,richb}@rice.edu

***Abstract*— This paper describes a novel method for identifying multiple targets with multiple robots in a partially known environment. Two main issues are addressed. The first relates to the use of motion planning algorithms to determine whether robots can reach "good" positions that offer the most informative measurements. The second concerns the use of predictive sensing to decide where sensor measurements should be taken. The problem is formulated similar to a next-best-view problem with differential constraints on the robots' motion, with additional layers of complexity due to visual occlusions as well as navigational obstacles. We propose a new distributed sensing strategy that exploits the structure of image manifolds to predict the utility of the measurements at a given position. This information is encoded in a cost map that guides a motion planning algorithm. Coordination among robots is achieved by incorporating additional information in each robot's cost map. A range of simulations indicates that our approach outperforms current approaches and demonstrates the advantages of predictive sensing and accounting for reachability constraints.**

I. Introduction

Mobile robots often have sophisticated algorithms to extract information from sensor measurements. It is nevertheless a challenging problem to decide *when* and *where* sensor measurements should be taken and *how* robots can navigate to "good" positions that offer the most informative measurements. The following scenario illustrates the problem we are addressing in this paper. Suppose an Unmanned Aerial Vehicle (UAV) has identified several possible target locations. The UAV can send the locations to a team of mobile ground robots, who already have a partial map of the environment. The ground robots are able to sense the targets at a higher resolution than the UAV, thereby resolving any ambiguities. The ground robots proceed to navigate to good viewpoints, exchange information whenever they are within communication range, and coordinate their actions. The ground robots' goal is to collectively build models of the targets or classify the targets.

The scenario faced by the ground robots bears a strong resemblance to the *next best view problem* (NBV) [1]. In general this problem can be formulated as: given a series of previous measurements of an object, what is the position from which the most informative next measurement can be taken? The informativeness depends on both the sensing modality and the high-level task (model building, object classification, etc.). However, to the best of our knowledge, this problem has not been considered for *multiple, mobile* sensors simultaneously trying to identify *multiple* targets.

This work was supported in part by the US Army Research Laboratory and the US Army Research Office under grant number W911NF-09-1-0383.

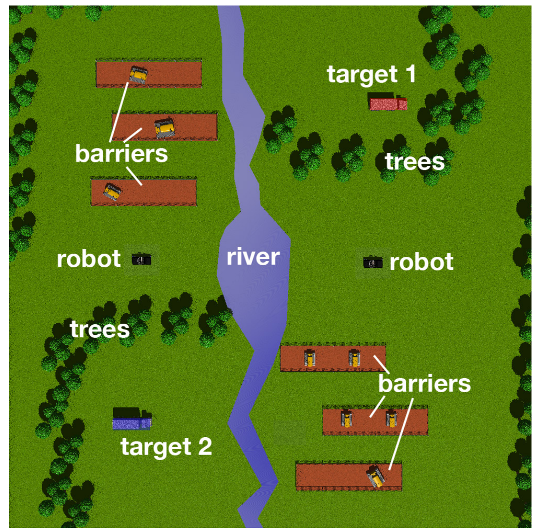

Fig. 1. **Next best view for multiple targets and multiple car-like robots.** Both predicting measurements and reasoning about reachability are essential for efficient target classification.

There is a significant amount of work on multi-robot target tracking and target localization (see, e.g., [2], [3]), but there the best viewing angle for target *classification* and *model building* is typically not considered.

Figure 1 illustrates the importance of predictive sensing and reachability for the NBV problem: neither robot can see either target completely or reach all viewing angles, but together they can collect enough information that views of either target from any angle can be synthesized. The targets are not statically assigned to robots, but, rather, the robots use local communication to coordinate opportunistically. Task allocation is performed dynamically and implicitly through continuously updated *cost maps*: grid decompositions of the workspace such that each grid cell encodes the expected cost of obtaining a measurement from a position in that cell. Traditionally, the cost to reach the NBV has been ignored, or has been arbitrarily set to be proportional to the distance from the current position. For robots with differential constraints, such as car-like robots, this is problematic. The robots cannot simply use a reactive, control-based approach to follow a gradient based on the cost map. Furthermore, the NBV may be unreachable and positions near the NBV may not be informative at all (due to, e.g., occlusion). On the other hand, there may exist many reachable positions that are only marginally less informative than the best view. Computing the next best *reachable* view is generally undecidable, and so instead we will use an approximate solution from an online replanning framework to select the next view.

The contributions of this work are as follows. We propose a novel formulation for multi-robot, multi-target NBV for car-like robots. The sensing strategy in this work allows robots to predict which measurements are expected to be informative. Often sensor measurements can be avoided, which leads to substantial savings in power usage, bandwidth, and computation. By exploiting reachability we are more likely to find short informative paths rather than potentially long paths to the "best" view point. Finally, the cost maps used in this approach enable multiple robots to discover informative views of multiple targets.

II. Overview of Our Approach

The primary focus is on the interaction between sensing and planning. Each robot executes a sense/plan/act loop, where it simultaneously senses, plans for the next cycle, and executes the plan computed in the previous cycle. At the heart of this sense/plan/act cycle is a geometric NBV algorithm that exploits relationships among images to propose novel views to sense. In stark contrast to conventional NBV algorithms where candidate views are suggested, we characterize the informativeness of viewpoints using a cost map. This choice is to accommodate reachability constraints (candidate NBV might be unreachable) as well as differential constraints in car-like robots. We therefore need to use a planning algorithm to produce dynamically feasible motions. We assume the robots have a map of the environment, and are able to self-localize on this map.

To achieve our goals of simultaneous sensing and navigation, our framework has two main parts: an *offline* model building effort where we build manifold models for each of the targets that we are interested in identifying, and an *online* real-time processing step on each of the robots wherein the sensed images are analyzed and navigation plans are made.

Offline: The offline model building process builds a manifold model for each target; manifold models offer a natural construct to represent and process multi-view image data efficiently. In particular, we are interested in quantifying informativeness of new views of a target given a set of images of the target. We achieve this by using the concept of a transport operator that links images on the manifold based on a prediction model. We characterize *informativeness* of a viewpoint using the size of the neighborhood on the image manifold that is predicted by the image obtained at the viewpoint. The key goal of our offline computations is to build a framework wherein we characterize informativeness of views. As we will see later, the NBV problem can be mapped as an elegant maximization of our notion of informativeness of viewpoints.

Online: The online processing part of our framework has two distinct processing stages: (a) an image analysis step where-in we use advances in object detection and pose estimation to locate target(s) of interest in the sensed image and suggest potential NBVs using cost maps, and (b) a path planning process that coordinates the movement of the robots to optimize some desired objective with reachability and inertial constraints. We assume that each robot has a camera with a limited field of view that is aimed at a target's expected location, and that through standard image analysis techniques the background is subtracted to extract an image of a target when a target is within view. The image analysis process does not simply pick the NBV according to some metric, but, instead, computes a score for each grid cell given a grid decomposition of the workspace. The grid-based scoring function is called a cost map and is computed at each time step by processing the aggregate set of measurements acquired by the mobile robots. The computation of cost maps is ostensibly expensive, especially in complex scenarios with multiple targets and robots. However, we have developed a technique to efficiently perform the computation by exploiting the intrinsic structure of the set of acquired measurements.

The planning process exploits this information by biasing the growth of a search tree containing feasible paths towards areas that have a low cost using a sampling-based planner [4]. Such planners have been demonstrated to be effective in finding valid paths for constrained, dynamic systems [5], [6] and can be guided in their search for good path by cost maps. This planning process accomplishes the goal of finding informative viewpoints that are also reachable. Whenever robots are within communication range, robots not only exchange any information about the targets that they have acquired so far, but also communicate their plans for the next cycle. We recently proposed a motion coordination framework for second-order car-like robots that enables them to safely operate in the same space [7]. This framework is used in this work. The plans from each neighboring robot are transformed into another cost map. The robots then use a composite cost map, formed by combining their own cost map, as well as cost maps for the neighbors' plans, to formulate a plan that balances finding an informative path with staying out of the way of the neighboring robots. This simple scheme allows robots to distribute naturally around an arbitrary number of targets without a complex task allocation scheme or negotiation process.

References

[1] J. E. Banta, L. R. Wong, C. Dumont, and M. A. Abidi, "A next-best-view system for autonomous 3-D object reconstruction," *IEEE Trans. Systems, Man and Cybernetics, Part A*, vol. 30, no. 5, pp. 589–598, 2000.

[2] L. E. Parker, "Distributed algorithms for multi-robot observation of multiple moving targets," *Auton. Robots*, vol. 12, pp. 231–255, 2002.

[3] B. Grocholsky, J. Keller, V. Kumar, and G. Pappas, "Cooperative air and ground surveillance," *IEEE Robotics Automation Magazine*, vol. 13, no. 3, pp. 16–25, Sept. 2006.

[4] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press, 2005.

[5] K. E. Bekris and L. E. Kavraki, "Greedy but safe replanning under kinodynamic constraints," in *IEEE Intl. Conf. on Robotics and Automation*, 2007, pp. 704–710.

[6] K. E. Bekris, K. I. Tsianos, and L. E. Kavraki, "Safe and distributed kinodynamic replanning for vehicular networks," *ACM/Springer Mobile Networks and Applications*, vol. 14, no. 3, pp. 292–308, 2009.

[7] D. Grady, K. E. Bekris, and L. E. Kavraki, "Asynchronous distributed motion planning with safety guarantees under second-order dynamics," in *Workshop on the Algorithmic Foundations of Robotics*, 2010.