Efficient Task Dependent Localization through Submodularity

Shervin Javdani, J. Andrew Bagnell, Siddhartha S. Srinivasa
The Robotics Institute, Carnegie Mellon University

Uncertainty is a fundamental challenge in robotics, particularly for tasks of fine manipulation. To handle uncertainty, many works perform a sequence of uncertainty reducing actions prior to attempting the task [1]–[3]. Ideally, the selected actions reduce uncertainty enough to accomplish the task while optimizing a performance criterion like minimum energy or time. Computing the *optimal* such sequence can be formulated as a Partially Observable Markov Decision Process (POMDP) [4]. However, finding optimal solutions to POMDPs for manipulation tasks is often intractable due to the large state and action spaces.

Previous work on uncertainty reduction utilizes online planning within the POMDP framework, looking at locally reachable states during each decision step. In general, these methods limit the search to a low horizon [1], often using the *greedy* strategy of selecting actions with the highest expected benefit in one step [2]. This is out of necessity - computational time increases exponentially with the search depth. However, this simple greedy strategy often works surprisingly well.

In our own work, we addressed the efficient automatic construction of such a sequence when information gaining actions are guarded moves. Our primary insight was making an explicit connection to *adaptive submodularity* [5], enabling us to use an efficient lazy-greedy algorithm guaranteed to select a sequence competitive with the optimal solution. Our experiments confirmed that are methods provided accurate localization efficiently. See Fig. 1 for an example sequence which enabled a successful grasp of door handle, and Fig. 2 for an image summarizing the method.

Our methods developed so far naively reduce uncertainty without considering the underlying task. In actuality, a task and planner may not require the exact pose, but that all uncertainty lie within a particular distribution. Consider pushing a button we may only care that if the robotic end effector started at a fixed pose and moved forward, it would successfully push the button. By optimizing for this target directly, we can achieve our goal more efficiently.

Our current work has focused on optimizing for this criteria directly with an adaptive submodular maximization. We do so by drawing on recent developments for active learning with noisy observations [6]. See Fig. 3 for an image summarizing the method. Importantly, this method is also adaptive submodular, enabling an efficient lazy-greedy algorithm guaranteed to be near-optimal.

In this poster, I will present our ongoing work on adaptive submodular formulations for dealing with uncertainty in robotic manipulation. It will encompass both our previous work on quickly localizing the object [3], as well as our more recent developments on task-based uncertainty reduction. This poster will focus on the ideas and directions underlying these formulations, rather than implementations.

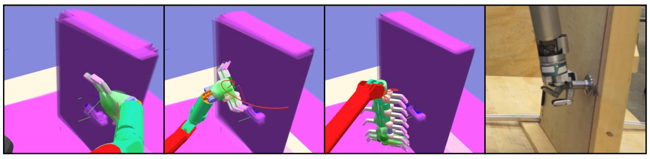

Fig. 1: We adaptively select a sequence of touch actions to reduce uncertainty. Here, we show actions selected by our Hypothesis Pruning method, enabling a successful grasp.

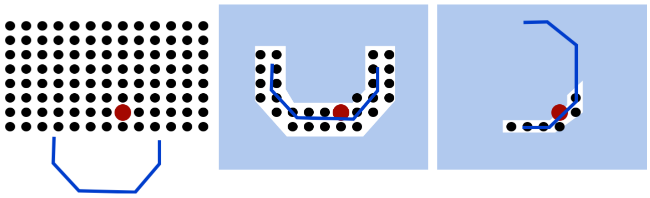

Fig. 2: We think of tactile localization as set cover problem, which is adaptive submodular [5]. Actions correspond to moving the hand until contact. Observations cover the hypotheses (black dots) which do not agree (blue area). True object shown in red. Our objective is to maximize coverage, or maximally rule out hypotheses.

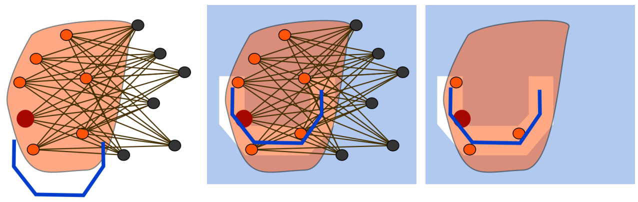

Fig. 3: Our previous approach pruned hypotheses without considering the task [3]. We may be able to accomplish a task successfully if all uncertainty is encapsulated in a region (orange).Optimizing for this criteria may be more efficient. We do so by considering an “edge” between all hypotheses inside the region and outside, and removing edges greedily.

References

[1] K. Hsiao, “Relatively robust grasping,” Ph.D. dissertation, Massachusetts Institute of Technology, 2009.
[2] P. Hebert, J. W. Burdick, T. Howard, N. Hudson, and J. Ma, “Action inference: The next best touch,” in *RSS Workshop on Mobile Manipulation*, 2012.
[3] S. Javdani, M. Klingensmith, J. A. D. Bagnell, N. Pollard, and S. Srinivasa, “Efficient touch based localization through submodularity,” in *IEEE ICRA*, 2013.
[4] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, vol. 101, pp. 99–134, 1998.
[5] D. Golovin and A. Krause, “Adaptive submodularity: Theory and applications in active learning and stochastic optimization,” *JAIR*, vol. 42, no. 1, pp. 427–486, 2011.
[6] D. Golovin, A. Krause, and D. Ray, “Near-optimal bayesian active learning with noisy observations,” in *NIPS*, 2010.