Long-term Autonomous Mobile Manipulation under Uncertainty

Extended Abstract

Michael W. Lanighan
University of Massachusetts Amherst
lanighan@cs.umass.edu

Roderic A. Grupen
University of Massachusetts Amherst
grupen@cs.umass.edu

ABSTRACT

We describe a hierarchical belief space planning framework to achieve robust long-term autonomous mobile manipulation behavior under uncertainty. The approach relies on condensing belief distributions across different abstractions to simultaneously suppress uncertainty and mitigate risk at run-time. We evaluate this system in an experimental domain that requires a robot to monitor and clean a dynamic unstructured environment, executing hundreds of physical actions without failures that require human intervention. Results indicate that this framework provides a sound basis for cognitive robots in uncertain and dynamic environments.

KEYWORDS

Long-term Autonomy; Robot Manipulation; Failure Recovery; Cognitive Architecture/System

ACM Reference Format:

Michael W. Lanighan and Roderic A. Grupen. 2019. Long-term Autonomous Mobile Manipulation under Uncertainty. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION AND BACKGROUND

Despite advances in manipulation, locomotion, and perception, results from recent robot challenges indicate that robots that are competent in unstructured (*open*) environments remains an elusive goal [5, 9]. This is due to *uncertainty* present in unstructured environments, sensing, and actuation. Developing planning and execution frameworks able to make effective decisions while overcoming uncertainty and managing risk is, therefore, a major goal of robotics. Research in Long-Term Autonomy (LTA) focuses on developing technologies to address problems introduced by *open worlds* [12]. Open worlds are defined by a lack of structure, partial observability, and complex dynamic environments.

Generally speaking, studies in LTA do not address challenges that arise when the robot intentionally alters the environment. The best examples of long-term deployments of semi-autonomous mobile manipulators and the issues they encounter are the systems deployed at the DARPA Robotics Challenge. These systems often experienced failures that required external interventions within tens of actions due to a low mean-time between decision making failures [9]. In this work, we view "long-term deployments" as autonomous deployments that last on the order of hours during which 100s of decisions about actions that intentionally alter the environment are executed.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

Decision making under uncertainty falls into a class of problems known as Partially Observable Markov Decision Processes (POMDP) [7]. Finding optimal solutions to POMDPs is PSPACE complete, so researchers have developed a variety of approximate techniques. A common technique is to plan in *belief space*, where instead of reasoning over the states themselves, reasoning is performed on the probability distribution defined over the states [6, 8, 15–17]. Lanighan *et al.* demonstrated simple assembly tasks that are robust to perceptual and motor uncertainty using a hierarchy of Active Belief Planners [13]. With this approach a robot is able to manage stochastic actions in partially observable systems to avoid unrecoverable outcome states (that require external intervention). Such a system supports robust and reliable autonomous systems. In this abstract, we summarize preliminary work on evaluating the performance of long-term, autonomous deployments by measuring the number of actions between unrecoverable decision failures.

2 HIERARCHICAL ACTIVE BELIEF

We build upon the work of Lanighan *et al.* [13], using an approach that decomposes a task into multiple belief spaces where belief in lower levels informs decision making at higher levels. Each planner is defined by the set S of world states, the conditional transition probability T between states, an observation function O, the set of available actions A, the belief distribution b_k over states $\in S$ at time k, a reward function $r(b_k, A) \rightarrow \mathbb{R}$ parameterized by the belief distribution and actions, and a state abstraction function $Z(b_i) \rightarrow z$, where b_i is the belief of the preceding level. Z allows higher levels of the hierarchy to form observations from lower level belief distributions, causing information to stabilize as it progresses up the hierarchy.

In this work, we introduce a hierarchy of three levels to control uncertainty in the environment, over object identities, and task geometry. We accomplish this through an *environmental level*, an *intra-object level*, and an *inter-object levels.* Each level evaluates uncertainty present in its state and selects actions to manage uncertainty. Belief states at each level are maintained by recursive Bayesian filtering.

The *environmental level* actively manages uncertainty in the large scale, volumetric occupancy of space surrounding the robot. Gross and fine motor actions produce autonomous mobile manipulation strategies to actively manage uncertainty and risk by interacting with the unstructured environment in response to visual and tactile observations. The *environmental level* implemented in this work respectively considers trajectories that re-position the mobile base of the robot. To generate trajectories, a harmonic function path planner [3, 4] is used to find collision-free trajectories while simultaneously reducing uncertainty by optimizing information gain.

The *intra-object level* of the three-level hierarchy manages belief over identities of objects in the environment relative to a set of object models. Objects are modeled using Aspect Transition Graphs (ATGs), a multi-graph that describes how actions create changes between multi-modal constellations of observable features (aspects) of an object [10, 16]. Belief distributions are instantiated using a state abstraction function Z that samples belief of spatial occupancy from the environmental level. A generalized Hough transform [1] computes observation probabilities to known models. To control uncertainty among the object models, sensorimotor actions are selected to change the object-robot relationship in order to minimize the future expected entropy of the belief distribution.

The *inter-object level* of the hierarchy manages uncertainty in spatial precision of the task with regard to relative position of multiple objects. The state abstraction function samples the maximum likelihood state of the *intra-object level* distribution to "observe" the expected positions of objects that are relevant to the task at hand. This level recommends pick and place actions that rearrange objects in the environment relative to each other such that they maximize a reward function. Kullback-Leibler divergence D_{KL} [11] between the current belief distribution of the *inter-object level*, b_2, and a goal distribution, G, is used in this work.

To determine control authority in the hierarchy we introduce a *belief subsumption* arbitration mechanism inspired by subsumption architectures [2]. In the hierarchy, decisions at the higher levels are only informative given that lower level beliefs are confident. The highest-level planner with enough confidence in its state is chosen for execution. Authority for action at any level depends on avoiding unrecoverable states. The highest-level that can guarantee safe and productive future states is selected. In this way, belief actively condenses in the hierarchy from the bottom-up.

3 EVALUATION CRITERIA AND RESULTS

To test the system's capability for long term autonomy, we propose a set of mobile manipulation scenarios with the uBot-6 mobile manipulator [14] in unstructured environments where the robot must "tidy-up" a room. This scenario is analogous to tasks like clearing tables that is suited to our platform. In the scenario, the robot must search a room, identify objects that are out of place, and return them to a pre-specified position. During execution the object may be re-oriented or re-placed by a disruptive external agent. The robot must adapt to these perturbations and avoid states that lead probabilistically to future unrecoverable failure states. The approach is evaluated in terms of the number and frequency of external resets over several such deployments.

The robot was deployed in the evaluation environment for over four hours. Deployments were interrupted when the batteries dropped to unsafe levels or when external resets were required. The deployments are summarized in Table 1. The robot successfully put the target object away and never "put-away" a non-target/distractor object during the deployments. The robot encountered nine failures that required interventions during the deployments. Two failures were due to system-level (hardware) failures. The remaining failures occurred while manipulating objects in the environment as prescribed by the *intra-object level.* According to the belief state at the time of these failures, these were legal actions. However, due to uncertainties in the underlying controllers these actions failed during execution. This type of error is attributed to inadequate precision in the empirical object models and, thus, a compromised ability to predict future states.

#	Duration (h:m:s)	Actions	Failures	Interactions
1	1:18:29	37	2	5
2	0:41:54	28	2	4
3	0:16:40	12	2	2
4	0:02:20	2	1	0
5	0:32:23	19	0	5
6	0:45:47	29	1	7
7	0:48:05	27	1	12
Total	4:23:28	154	9	37

Table 1: Run time, number of actions executed, number of failures, and the number of interactions per deployment.

During the seven deployments, the experimenter perturbed experimental objects in the middle of execution a total of 37 times. These interactions re-positioned and re-oriented objects in the environment, and/or inserted or removed distractor or target objects. A single failure attributable to such a disturbance occurred in the first deployment. The cause of the failure was an erroneous implementation of the state abstraction function between the *environmental* and *intra-object* levels. Due to an implementation error, although the *environmental level* belief provided little support for volumetric occupancy, the *intra-object level* belief was not correctly updated. This caused the robot to grasp empty space, which required intervention to correct. This implementation error was fixed in subsequent deployments. This failure highlights the impact that a correctly implemented hierarchy has in such situations—the robot avoids executing actions that are not supported by history and observation.

4 DISCUSSION

The hierarchy demonstrated in these deployments commits to control decisions only when belief in the various levels of abstraction have stabilized. By selecting which level of the hierarchy possesses control authority through *belief subsumption*, the system is able to manage uncertainty across the levels of the hierarchy. As result, the robot uses the three-level active belief hierarchy to interact with the unstructured domain within constraints derived from uncertainty in spatial occupancy, object identity, and inter-object precision and actively gathers information to meet safety and performance specifications. Despite repeated attempts by experimenters to frustrate the robot, it completed the specified 'tidy-task' over deployments lasting hours during which over 150 actions were executed. During these deployments nine failures requiring external interventions occurred. Although more experience with such an architecture and more experimental data is required, these preliminary results suggest that this kind of architecture—one that actively shapes posterior belief in hierarchically arranged planners across multiple levels of abstraction—can be an effective approach to systems that must accommodate extended periods of deployment with no human intervention.

REFERENCES

[1] Dana H Ballard. 1981. Generalizing the Hough transform to detect arbitrary shapes. *Pattern recognition* 13, 2 (1981), 111–122.

[2] Rodney Brooks. 1986. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation* 2, 1 (1986), 14–23.

[3] Christopher I Connolly, John B Burns, and R Weiss. 1990. Path planning using Laplace's equation. In *Robotics and Automation, 1990. Proceedings., 1990 IEEE International Conference on*. IEEE, 2102–2106.

[4] Christopher I Connolly and Roderic A Grupen. 1993. The applications of harmonic functions to robotics. *Journal of Robotic Systems* 10, 7 (1993), 931–946.

[5] Nikolaus Correll, Kostas E Bekris, Dmitry Berenson, Oliver Brock, Albert Causo, Kris Hauser, Kei Okada, Alberto Rodriguez, Joseph M Romano, and Peter R Wurman. 2018. Analysis and Observations From the First Amazon Picking Challenge. *IEEE Transactions on Automation Science and Engineering* 15, 1 (2018), 172–188.

[6] Dylan Hadfield-Menell, Edward Groshev, Rohan Chitnis, and Pieter Abbeel. 2015. Modular task and motion planning in belief space. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 4991–4998.

[7] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 101, 1 (1998), 99–134.

[8] Leslie Pack Kaelbling and Tomás Lozano-Pérez. 2013. Integrated task and motion planning in belief space. *The International Journal of Robotics Research* (2013).

[9] Eric Krotkov, Douglas Hackett, Larry Jackel, Michael Perschbacher, James Pippine, Jesse Strauss, Gill Pratt, and Christopher Orlowski. 2018. The DARPA Robotics Challenge Finals: Results and Perspectives. In *The DARPA Robotics Challenge Finals: Humanoid Robots To The Rescue*. Springer, 1–26.

[10] L. Ku, S. Sen, E. Learned-Miller, and R. Grupen. 2014. Aspect Transition Graph: An Affordance-Based Model. In *European Conference on Computer Vision, Workshop on Affordances: Visual Perception of Affordances and Functional Visual Primitives for Scene Analysis*.

[11] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86.

[12] Lars Kunze, Nick Hawes, Tom Duckett, Marc Hanheide, and Tomáš Krajník. 2018. Artificial intelligence for long-term robot autonomy: A survey. *IEEE Robotics and Automation Letters* 3, 4 (2018), 4023–4030.

[13] Michael W Lanighan, Takeshi Takahashi, and Roderic A Grupen. 2018. Planning Robust Manual Tasks in Hierarchical Belief Spaces.. In *Proceedings of the Twenty-Eighth International Conference on International Conference on Automated Planning and Scheduling (ICAPS)*. 459–467.

[14] Dirk Ruiken, Michael W Lanighan, and Roderic A Grupen. 2013. Postural Modes and Control for Dexterous Mobile Manipulation: the UMass uBot Concept. In *13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 721–726.

[15] Dirk Ruiken, Tiffany Q Liu, Takeshi Takahashi, and Roderic A Grupen. 2016. Reconfigurable Tasks in Belief-Space Planning. In *16th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*.

[16] Dirk Ruiken, Jay Ming Wong, Tiffany Q Liu, Mitchell Hebert, Takeshi Takahashi, Michael W Lanighan, and Roderic A Grupen. 2016. Affordance-based active belief recognition using visual and manual actions. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.

[17] Takeshi Takahashi, Michael W Lanighan, and Roderic A Grupen. 2017. Hybrid Task Planning Grounded in Belief: Constructing Physical Copies of Simple Structures. In *Proceedings of the Twenty-Seventh International Conference on International Conference on Automated Planning and Scheduling (ICAPS)*. AAAI Press.