Estimating Value of Localization Assistance for Path-Planning Robots

Sarah Keren[1]

***Abstract*— Navigating robots often need to plan their behavior based on partial information about their environment and on probabilistic localization estimations. In some settings, it may be possible to perform information acquisition actions that may reduce uncertainty about the robot's location but that may be costly and limited. We aim to support the decision of whether it is worthwhile to execute such actions with respect to the agent's objectives. For this, we offer several heuristic measures that provide estimations of the expected cost reduction the extra information will yield. Evaluation of a set of robotic navigation settings demonstrates the ability of our measures to estimate the expected benefit of such actions.**

I. INTRODUCTION

In many real-world settings, autonomous agents, and robots in particular, need to plan their behavior based on partial information about the environment. A common paradigm for control involves a modular approach that decomposes the system into modules for perception, state estimation, planning, and control [1], [2]. This decomposition increases flexibility by allowing each module to be designed, developed, and operated independently. However, since these modules typically communicate through a predefined and often narrow interface, this modular structure may also compromise robustness due to inadequate interfaces that this structure imposes.

Our focus here is on the often inadequate interface between the state estimation component, which captures the uncertainty regarding the state of the environment and the agent, and the path planner, responsible for choosing the trajectories to be followed by the agent. Typically, the state estimation component maintains the robot's *belief*, representing a probability distribution over possible world states. In a mapped environment, the belief represents the uncertainty associated with the robot's location within the map, referred to as *localization uncertainty*, which is often due to limited sensing and computational resources. In contrast, many common *path planners*, that prescribe a sequence of robot configurations a robot should follow to reach some goal configuration, accept as input a single discrete representation of the robot's state.

A typical approach to the interface between the two components is to provide the planner with the mean of the state estimation, thus abstracting location uncertainty from the planner and other downstream modules. This approach to planning relies on the robot's ability to recover from potential failure and replan its behavior, which may be very costly [3]. An alternative approach is contingent planning, which attempts to account for the different outcomes that may occur. This includes POMDP planners [**?**] which account for uncertainty by producing a contingent plan over the robot's belief. This approach provides more comprehensive policies but tends to be complex and offers limited applicability to real-world settings.

In this work, we suggest a novel interface between a probabilistic state estimation module and a discrete planner that exploits the planner in assessing the effect localization uncertainly will have on performance. This analysis, which is based on comparing the plans the robot would follow for different possible states in its belief, provides not only an efficient way to assess a robot's ability to achieve its goal but also to assess whether it is worthwhile to reduce its uncertainty by providing it with localization information. This allows us to support settings in which performing such information-sharing operations is costly or limited. For example, our analysis can help an underwater vehicle decide if to rise to the surface to observe its surrounding. It can also help a wheeled robot decide if to request costly assistance from a drone that can fly to its estimated location and reveal its exact location on the map.

Choosing when to acquire new information and communicate existing information is a fundamental problem in single-agent and multi-agent AI systems [4]. This is an especially challenging problem when communication is limited and costly and each agent has a different view of the state of the world. Solution approaches vary between those that distinguish between the information acquisition process and the decision on how to act [5] to those that apply a decision-theoretic approach that views information acquisition and communication as part of the actions an agent can perform, thus supporting a unified analysis that balances the benefits of communication against its costs [6].

In this work, we adopt the notion of *Value of Assistance* (VOA) [7] which is used to quantify the expected increase in an agent's expected value (or reduction in cost) as a result of obtaining the information. This measure is based on the notion of *value of information* (VOI) proposed by [8], [4] to represent the expected increase in an agent's expected value (or reduction in cost) as a result of obtaining the information. This idea has been applied in many settings [9], including its adaptation to a human-robot collaborative setting where it is used to quantitatively measure the helpfulness of a robotic partner for a given task at hand [10], [11]. Recently, VOA is used to quantify the effect localization information will have on a robot that relies only on internal sensors for localization and is tasked with following a fixed sequence of waypoints [7]. In this setting, localization uncertainty,

The author is with the Faculty of Computer Science, Technion - Israel Institute of Technology, Haifa 32000, Israel [3]`sarahk@cs. }technion.ac.il`

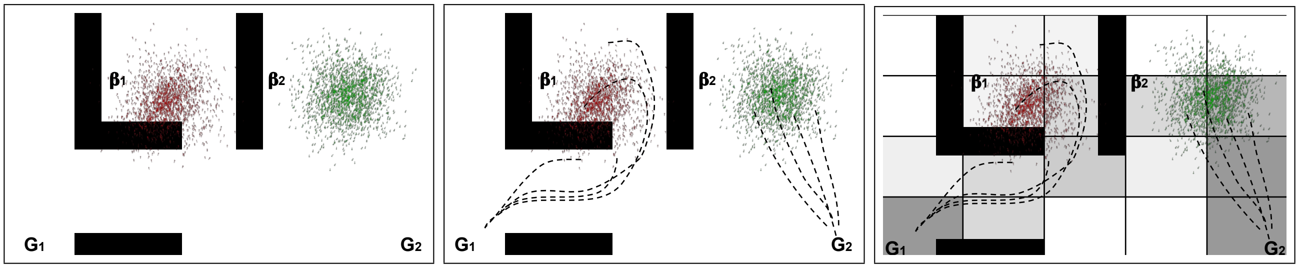

Fig. 1: Demonstrating ways to estimate the expected benefit of providing localization information.

modeled as a Gaussian process, grows as the agent progresses in the environment. When receiving localization information the agent attempts to return to its intended location by following a straight line. We extend VOA to support goal-directed re-planning agents that use a path planner to decide how to act and can re-plan their path based on newly acquired information.

Ideally, we would like to know VOA before the robot starts to move. However, since the actual effect of the novel information is only known at execution time, we seek ways to estimate it. Our focus is on robots with localization uncertainty and on an offline estimation of the value of revealing their exact location before they begin execution. For this, we suggest novel ways to compute VOA based on the assumption that uncertainty is modeled as a *particle filter* [12], [13], where particles represent a possible location of the robot. Our analysis is based on sampling configurations from the belief and computing the plans the agent would follow for each sample (i.e., particles). Based on this computation we offer novel estimations that consider the plan's cost variance and the areas traversed by the plans.

Example 1. To demonstrate our suggested VOA measures, consider the setting depicted in Figure 1[left], in which two mobile agents (1 and 2) navigate towards some goal location G_i. Both agents have the same level of uncertainty, depicted by the scattered particles of β_i, each representing a possible robot location. One way to estimate VOA is to consider the variance of the length of the plans the robot would follow for a sampled set of particles, as depicted in Figure 1[center]. Another option is to overlay the grid with a heat-map representation of the grid Figure 1[right], where the heat of a cell represents the number of plans that traverse it, and to estimate VOA according to the heat variance.

In our example above, both suggested measures are aimed at accounting for agents with limited external sensing and thus associate a greater need for information and higher VOA to robot 1. However, it is important to note that such measures may have different meanings in different settings, depending on the agent's capabilities, the environment in which it operates, and its task. In particular, while obstacles in the environment might hinder the ability to reach the goal for a robot with limited sensors, they may facilitate localization and reduce the need for assistance by representing landmarks for an agent with sensing capabilities. This raises the need to develop different VOA measures and empirically examine their effect in the domain of interest.

The contributions of this work are threefold. First, we suggest a novel interface between probabilistic state estimators and discrete path planners that accounts for localization uncertainty. Secondly, we suggest novel and efficient ways to estimate the value of providing localization information for agents for which the belief is represented by a probability distribution that can be sampled. Lastly, we provide an empirical analysis that demonstrates the ability of our measures to estimate VOA and enhance performance.

References

[1] H. Kurniawati, D. Hsu, and W. S. Lee, "Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces." in *Robotics: Science and systems*, vol. 2008. Citeseer, 2008.

[2] L. P. Kaelbling and T. Lozano-Pérez, "Integrated task and motion planning in belief space," *The International Journal of Robotics Research*, vol. 32, no. 9-10, pp. 1194–1227, 2013.

[3] L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical task and motion planning in the now," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 1470–1477.

[4] R. Becker, A. Carlin, V. Lesser, and S. Zilberstein, "Analyzing myopic approaches for multi-agent communication," *Computational Intelligence*, vol. 25, no. 1, pp. 31–50, 2009.

[5] S. Gil, S. Kumar, D. Katabi, and D. Rus, "Adaptive communication in multi-robot systems using directionality of signal strength," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 946–968, 2015.

[6] C. V. Goldman and S. Zilberstein, "Optimizing information exchange in cooperative multi-agent systems," in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, 2003, pp. 137–144.

[7] A. Amuzig, D. Dovrat, and S. Keren, "Value of assistance for mobile agents," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.

[8] R. A. Howard, "Information value theory," *IEEE Transactions on systems science and cybernetics*, vol. 2, no. 1, pp. 22–26, 1966.

[9] R. J. Marcotte, X. Wang, D. Mehta, and E. Olson, "Optimizing multi-robot communication under bandwidth constraints," *Autonomous Robots*, vol. 44, no. 1, pp. 43–55, 2020.

[10] J. A. Shah, "Fluid coordination of human-robot teams," Ph.D. dissertation, Massachusetts Institute of Technology, 2011.

[11] R. G. Freedman, S. J. Levine, B. C. Williams, and S. Zilberstein, "Helpfulness as a key metric of human-robot collaboration," *arXiv preprint arXiv:2010.04914*, 2020.

[12] N. J. Gordon, D. J. Salmond, and A. F. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," in *IEE Proceedings F-radar and signal processing*, vol. 140, no. 2. IET, 1993, pp. 107–113.

[13] D. Fox, W. Burgard, F. Dellaert, and S. Thrun, "Monte carlo localization: Efficient position estimation for mobile robots," *AAAI/IAAI*, vol. 1999, no. 343-349, pp. 2–2, 1999.