

Inspection planning under execution uncertainty

Shmuel David Alpert¹, Kiril Solovey¹, Itzik Klein², Oren Salzman¹

Abstract—Autonomous inspection tasks require path-planning algorithms to efficiently gather observations from *points of interest* (POIs). However, localization errors in urban environments introduce execution uncertainty, posing challenges to successfully completing such tasks. Existing inspection-planning algorithms do not explicitly address this uncertainty, which can hinder their performance. To overcome this, we introduce *IRIS-under uncertainty* (IRIS-U²), an inspection-planning algorithm that provides statistical assurances regarding coverage, path length, and collision probability. Our approach builds upon IRIS —our framework for *deterministic*, highly efficient, and provably asymptotically-optimal framework. This extension adapts IRIS to uncertain settings using a refined search procedure that estimates POI coverage probabilities through Monte Carlo (MC) sampling. We demonstrate IRIS-U² through a case study on bridge inspections, achieving improved expected coverage, reduced collision probability, and increasingly precise statistical guarantees as MC samples grow. Additionally, we explore bounded suboptimal solutions to reduce computation time while preserving statistical assurances.

I. INTRODUCTION

We consider the problem of planning in an offline phase a collision-free path for a robot to inspect a set of *points of interest* (POIs) using onboard sensors. This can be challenging, especially in urban environments, where dynamics uncertainty and localization errors (e.g., inaccuracies in location estimates) increase the task’s complexity. In particular, localization can be a significant source of uncertainty in urban environments, leading to missed POIs and compromising the efficiency and accuracy of inspection missions.

One application that motivates this work is the inspection of bridges using *unmanned aerial vehicles* (UAVs) [1]. Almost 40% of the bridges in the United States of America exceed their 50-year design life [2], and regular inspections are critical to ensuring bridge safety. UAVs can efficiently inspect bridge structures via visual assessment at close range without involving human inspectors or expensive under-bridge inspection units [3]. In these scenarios, the UAV typically carries a camera for POI inspection and a navigation system that combines data from various sensors like a *global navigation satellite system* (GNSS) and inertial sensors. Yet, GNSS signal obstruction by the bridge can cause location inaccuracies, heavily relying on inertial measurements (which also suffer from inaccuracies) and potentially compromising inspection effectiveness (see Fig. 1).

Several approaches have been proposed to perform inspection planning without accounting for uncertainty [1], [4], [5]. These methods can be used when the uncertainty is low by compensating for lack of GNSS by upgrading the navigation




Fig. 1: A UAV executing a bridge-inspection task, following a command path pre-computed offline (dark blue). The UAV relies on GNSS satellite signals for navigation but the bridge obstructs one satellite signal during a particular segment of its path (light blue). This leads to a substantial deviation from the command path, which hinders the UAV’s ability to inspect the POIs and may result in a collision with the bridge.

system to use RF tools such as WIFI [6], Bluetooth [7] or tactical-grade inertial sensors [8]. However, these solutions require a supportive communication infrastructure or more expensive and heavy sensors.

As we detail in Sec. II, uncertainty has been extensively considered when planning the motion of a robot without accounting for inspection considerations. This has been done using various methods that allow to account for position uncertainty by e.g., using Monte Carlo localization [9], particle filters [10], or Bayesian filtering to estimate and propagate uncertainty in the robot’s position and environment [11]. However, uncertainty in inspection tasks introduces further complications as it requires reasoning about an exceptionally large search space that captures the interaction between the POI locations, the locations from which POIs can be inspected, and the order of visitation of the latter locations. As a result, the computational complexity of inspection planning can be large, even in the absence of localization uncertainty [5], [12].

In our previous work, we introduced the IRIS algorithm, a highly effective approach for deterministic inspection planning [12]. Importantly, IRIS guarantees asymptotic convergence to optimal solutions while being several orders of magnitude faster than previous work providing similar guarantees.

Unfortunately, extending IRIS to the uncertain setting is highly nontrivial. Roughly speaking, the efficiency of IRIS is due to a novel graph-search algorithm that searches in the space of possible paths, where the search is guided by reasoning about which POIs have been seen along any given path. Unfortunately, in our context, execution uncertainty can lead to the robot deviating from its proposed path, which can lead to inspecting a different set of POIs from what it initially intended.

Contribution. We present IRIS-under uncertainty (IRIS-U²)—an offline inspection planning algorithm that systematically

¹ Technion–Israel Institute of Technology, Haifa, Israel.

² The Hatter Dept. of Marine Technologies, Charney School of Marine Sciences, University of Haifa, Haifa, Israel.

accounts for execution uncertainty. Our algorithm combines the capabilities of (i) IRIS, which efficiently explores the space of high-quality inspection plans in the deterministic setting, together with (ii) *Monte Carlo* (MC) sampling to reason about uncertainty via POI inspection probabilities. Importantly, our approach does not merely utilize those two components in a decoupled manner, by, e.g., interleaving between planning and uncertainty estimation [13], but rather uses the uncertainty model to (i) obtain statistical bounds (on the number of POIs inspected by each path, collision probability and length) and (ii) uses these statistical bounds to guide the exploration of the search space. To compute and integrate those probabilities within a search-based approach in a computationally efficient manner, we develop novel mechanisms for extending, subsuming, and dominating nodes. The computed probabilities then serve as optimization objectives for the algorithm.

On the theoretical side, IRIS-U² estimates with a certain *confidence level* (CL) the desired performance criteria (i.e., path length, coverage, and collision probabilities) within some *confidence interval* (CI) depending on the number of the MC samples used [14]. However, when selecting an execution path from multiple optional paths, straightforward statistical analysis may be biased toward false negatives and either requires offline preprocessing to provide guarantees or should be used as a guideline (see details in Appendix C). We choose the latter approach and outline a procedure to set a user-defined parameter for providing a CI, which becomes tighter as the number of samples increases. Additionally, we highlight the potential benefits of using a bounded sub-optimal solution in certain situations to reduce computation time while still providing guarantees through the CI boundaries.

We demonstrate the effectiveness of IRIS-U² through a simulated case study of structural inspections of bridges using a UAV in an urban environment. Our results show that IRIS-U² is able to achieve a desired level of coverage, while also reducing collision probability and tightening the CI lower boundaries as the number of MC samples increases.

The rest of this paper is organized as follows. In Sec. II, we review relevant related work. In Sec. III, we formulate the problem of offline path planning for inspection tasks under execution uncertainty. In Sec. IV, we describe the algorithmic background and in Sec. V we describe our proposed approach followed by a theoretical analysis presented in Sec. VI. Then in Sec. VII and VIII we present the results of our simulated experiments. Finally, in Sec. IX, we summarize our main contributions and discuss future work.

II. RELATED WORK

A. Motion planning under uncertainty

Various approaches have been proposed to address localization uncertainty for safe motion planning (which only considers a motion plan between states without inspection points). For instance, minimum-distance collision-free paths can be computed while bounding the uncertainty by, e.g., a self-localization error ellipsoid, or uncertainty corridor [15]–[17]. Another approach [18], uses a mixed-observability Markov decision process approach to account for a-priori probabilistic

sensor availability and path execution error propagation. Englot et al. [19] suggest an RRT*-based method which minimizes the maximum uncertainty of any segment along a path. Sampling-based planners that explicitly reason about the evolution of the robot’s belief have also been proposed [20]–[23].

B. Inspection planning

Several approaches have been proposed for *offline* inspection planning without accounting for uncertainty. Some algorithms decompose the region containing the POIs into sub-regions and solve for each sub-region separately [24], while others divide the problem into two NP-hard problems: (i) solving the art gallery problem to compute a small set of viewpoints that collectively observe all the POIs, and (ii) solving the traveling salesman problem to compute the shortest path to visit this set of viewpoints [13], [25]–[27]. Another common approach is to simultaneously compute both inspection points (i.e., points from which POIs are inspected) and the trajectory to visit these inspection points using sampling-based techniques [5], [28]. However, these approaches do not consider the execution uncertainty, which can lead to missed POIs due to differences between the path executed and the path planned.

Alternative approaches consider the *online* setting where a planner has to run online, deciding on a next step as the exploration or inspection advances [29], [30]. Bircher et al. [29] consider the setting where no uncertainty exists. By choosing different objective functions, their approach can be used for either the exploration of unknown environments or inspection of a given surface manifold in both a known and an unknown volume. Perhaps most closely related to our work is the approach by Papachristos et al. [30] where uncertainty is considered during execution via a receding-horizon technique. Their work starts by computing a path, then optimizes uncertainty, iterating to minimize localization and mapping uncertainty. However, as uncertainty optimization relies heavily on online updates (via tracked landmarks) of the robot’s belief it cannot be used in an offline planning phase.

Finally, we note the related problem of *Coverage Planning* [31] concerned with a robot (or a robot team) physically covering a region of interest such as in autonomous cleaning [32] lawn mowing [33] or surveillance [34], [35]. There has been some research that incorporates uncertainty in coverage planning, albeit not in a manner directly transferable to our problem [36]. For example, Yang et al. [37] consider the coverage problem for road sweepers, which typically have a large turning radius and a small coverage radius, thus making the coverage path highly sensitive to mapping uncertainty near obstacles. They suggest a planning framework for processing the environment’s map to reduce uncertainty and allow for better physical coverage near workspace obstacles. This contrasts with our setting, where we account for execution uncertainty rather than obstacle uncertainty. Abreu et al. [38] consider the coverage problem for autonomous underwater vehicles executing lawnmower-type paths. Unfortunately, this work is specifically tailored for the case of lawn-mower patterns in the workspace and cannot be applied to our general setting. Lastly, Galceran et al. [39] present a coverage path-planning approach to minimize a robot’s position uncertainty

along a planned path while taking into account area coverage performance. However, similar to the work by Abreu et al. [38], it targets a specific type of paths encountered by bathymetric mapping applications and cannot be applied to our setting.

III. PROBLEM DEFINITION

In this section, we provide a formal definition of the inspection-planning problem under execution uncertainty. We start by introducing basic definitions and notations, and discuss the uncertainty considerations (Sec. III-A). Next, we formally describe the inspection-planning problem in the deterministic (uncertainty-free) regime (Sec. III-B). Then, we introduce the inspection-planning problem under execution uncertainty which will be the focus of this work (Sec. III-C).

A. Basic definitions and notations

We have a holonomic robot \mathcal{R} operating in a workspace $\mathcal{W} \subset \mathbb{R}^3$ amidst a set of (known) obstacles $\mathcal{W}_{\text{obs}} \subset \mathcal{W}$.¹ A *configuration* q is a d -dimensional vector uniquely describing the robot's pose (position and orientation) and let $\mathcal{X} \subset \mathbb{R}^d$ denote the robot's *configuration space*. Let $\text{Shape} : \mathcal{X} \rightarrow 2^{\mathcal{W}}$ be a function mapping a configuration q to the workspace region occupied by \mathcal{R} when placed at q (here, $2^{\mathcal{W}}$ is the power set of \mathcal{W}). We say that $q \in \mathcal{X}$ is *collision free* if $\text{Shape}(q) \cap \mathcal{W}_{\text{obs}} = \emptyset$.

A path $\pi = (q_0, q_1, \dots) \subset \mathcal{X}$ is a sequence of configurations called *milestones* connected by straight-line edges $\overline{q_i q_{i+1}}$. A path π is *collision free* if every configuration along the path (i.e., either one of the milestones or along the edges connecting milestones) is collision-free. We use the binary function $\mathcal{C}(\pi)$ to express the *collision state* of a path π where $\mathcal{C}(\pi) = 0$ and $\mathcal{C}(\pi) = 1$ correspond to π being collision free and in-collision, respectively. Finally, we use $\ell(\pi)$ to denote the length of a given path π .

When a path π is computed in an *offline phase* to be later executed by the robot, we refer to it as a *command path*. Unfortunately, when following the command path, the system typically deviates from π due to different sources of uncertainty. In particular, we assume that the robot operates under two sources of uncertainty: (i) *control uncertainty* and (ii) *sensor uncertainty*. (A simple toy example demonstrating these concepts is detailed in Sec. V.) The control uncertainty results from various sources, such as a mismatch between the robot model used during planning and the real robot model, and disturbances in the environment (e.g., wind gusts). Sensor uncertainty (which is usually the main source of uncertainty in an urban environment) corresponds to navigation-model parameters that determine the robot's configuration which is not accurately known. Those parameters need to be modeled as random variables and may include biases of the inertial sensors or GNSS error terms [40]. To counteract errors resulting from those two sources of uncertainty, a *localization algorithm*, e.g., Kalman or particle filter, is invoked to compute the estimated robot location [41].

¹The assumption that the robot is holonomic is realistic as our motivating application is UAV inspection in which the UAV typically flies in low speed in order to accurately inspect the relevant region of interest.

To this end, we assume that we have access to distributions \mathcal{D}_{unc} from which the parameters governing the execution uncertainty \mathcal{N}_{unc} are drawn (see example in Sec. V-C)². In addition, we assume to have access to the initial true location (i.e., there is no uncertainty in the initial configuration of any execution path regardless of the command path provided). Finally, we assume that we have access to a black-box simulator, or motion model, \mathcal{M} that, given a command-path π , the uncertainty parameters \mathcal{N}_{unc} and the initial location, outputs the path that the robot will pursue starting from the initial location while following π under \mathcal{N}_{unc} . Note that after \mathcal{N}_{unc} is drawn, the model is deterministic. We mention that our approach is general and can be applied to any type of control and sensor uncertainty, as well as the accompanying estimation algorithm, so long that they can be faithfully simulated.

B. Inspection planning (without execution uncertainty)

In the inspection-planning problem, we receive as input a set of POI $\mathcal{I} = \{\iota_1, \dots, \iota_k\} \subset \mathcal{W}$ which should be inspected using some on-board inspection sensors (e.g., a camera). We model the inspection sensors as a mapping $\mathcal{S} : \mathcal{X} \rightarrow 2^{\mathcal{I}}$ such that $\mathcal{S}(q)$ denotes the subset of \mathcal{I} that are inspected from a configuration $q \in \mathcal{X}$. By a slight abuse of notation, we define $\mathcal{S}(\pi) := \bigcup_i \mathcal{S}(q_i)$ to be the POI that can be inspected by traversing the path $\pi = (q_0, q_1, \dots)$. For simplicity, we only inspect POIs along milestones (rather than edges). We start with a simplified setting of the inspection problem which involves no uncertainty on the side of control or sensing. That is, a robot performing the inspection can precisely follow a path during execution time and its location is known exactly. Such a setting was solved by IRIS (see Sec. IV-A).

Problem 1 (Deterministic problem): In the *inspection-planning problem* we wish to compute in an offline phase a path π that maximizes its coverage $|\mathcal{S}(\pi)|$. Out of all such paths we wish to compute the paths whose length $\ell(\pi)$ is minimal.

Note. In practice, one may be interested in minimizing mission completion time or energy consumption (and not path length which is a first-order approximation for these metrics). Optimizing for these metrics is slightly more complex and is left for future work.

Prob. 1 can be defined for the *continuous setting* (i.e., when we consider all paths in \mathcal{X} between a given start configuration q_{start} and goal configuration q_{goal}) or for the *discrete setting* where we restrict the set of available paths to those defined via a given *roadmap*. Here, a roadmap $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a graph embedded in the configuration space \mathcal{X} such that each vertex is associated with a configuration and each edge with a local path connecting close-by configurations. Roadmaps are commonly used in motion-planning algorithms (see, e.g., [44], [45]) and, as we will see, will be the focus of this work as well.

²Having access to \mathcal{D}_{unc} is a common assumption, see, e.g., [42], [43].

C. Inspection-planning problem under execution uncertainty

In the inspection-planning problem under execution uncertainty, we calculate a command-path π in an offline stage. To account for execution uncertainty, the following definitions extend the notion of path length, path coverage, and collision state to be their expected values:

Definition 1: The *expected* collision, coverage, and length probabilities are defined as:

$$\bar{C}(\pi) := E_{\mathcal{N}_{\text{unc}} \sim \mathcal{D}_{\text{unc}}} [\mathcal{C}(\mathcal{M}(\pi, \mathcal{N}_{\text{unc}}))], \quad (1a)$$

$$|\bar{\mathcal{S}}(\pi)| := E_{\mathcal{N}_{\text{unc}} \sim \mathcal{D}_{\text{unc}}} [|\mathcal{S}(\mathcal{M}(\pi, \mathcal{N}_{\text{unc}}))|], \quad (1b)$$

$$\bar{\ell}(\pi) := E_{\mathcal{N}_{\text{unc}} \sim \mathcal{D}_{\text{unc}}} [\ell(\mathcal{M}(\pi, \mathcal{N}_{\text{unc}}))], \quad (1c)$$

respectively. Here, $\bar{C}(\pi) \in [0, 1]$, where $\bar{C}(\pi) = 0$ and $\bar{C}(\pi) = 1$ correspond to π being guaranteed to be collision-free and in-collision, respectively.

We are now ready to formally introduce the optimal inspection-planning problem under execution uncertainty.

Problem 2 (Optimal problem): In the *optimal inspection-planning problem under execution uncertainty* we are given a user-provided threshold $\rho_{\text{coll}} \in [0, 1]$ and we wish to compute in an offline phase a command-path π such that its expected execution collision probability is below ρ_{coll} (i.e., $\bar{C}(\pi) \leq \rho_{\text{coll}}$), and which maximizes the expected coverage $|\bar{\mathcal{S}}(\pi)|$. Of all such paths, we wish to choose the one whose expected length $\bar{\ell}(\pi)$ is minimal.³

Finally, we introduce a relaxation of the above problem to reduce its computational cost.

Problem 3 (Sub-optimal problem): Let π^* be the solution to the optimal inspection-planning problem under uncertainty. In addition, let $\varepsilon \geq 0$, and $\kappa \in (0, 1]$ be user-provided approximation factors with respect to path length and coverage, respectively. Then, in the *sub-optimal inspection-planning problem under execution uncertainty* we wish to compute in an offline phase a command path π such that:

$$\bar{C}(\pi) \leq \rho_{\text{coll}}, \quad (2a)$$

$$|\bar{\mathcal{S}}(\pi)| \geq \kappa \cdot |\mathcal{S}(\pi^*)|, \quad (2b)$$

$$\bar{\ell}(\pi) \leq (1 + \varepsilon) \cdot \ell(\pi^*). \quad (2c)$$

Notice that, by setting $\kappa = 1$ and $\varepsilon = 0$ the sub-optimal Prob. 3 is equivalent to Prob. 2. As we are only given a black-box model of \mathcal{M} , it is infeasible to directly compute the expected values for coverage, collision probability, and length, $|\bar{\mathcal{S}}(\pi)|$, $\bar{C}(\pi)$ and $\bar{\ell}(\pi)$, respectively. As we will see, our approach will be to solve Prob. 3 using estimates of these values.

IV. ALGORITHMIC BACKGROUND

In this section, we provide algorithmic background. We begin by describing IRIS [12], a state-of-the-art algorithm

for solving the continuous inspection-planning problem in the deterministic regime (Prob. 1). We then continue to outline the statistical methods we will use. Throughout the text we assume familiarity with the A* algorithm [46].

A. Incremental Random Inspection-roadmap Search (IRIS)

IRIS solves Prob. 1 by incrementally constructing a sequence of increasingly dense graphs, or roadmaps, embedded in \mathcal{X} and computes an inspection plan over the roadmaps as they are constructed. The roadmap $\mathcal{G} = (\mathcal{V}, E)$ is a Rapidly-exploring Random Graph (RRG) [47] rooted at the start configuration (though other types of graphs, such as PRM*, can be used as well). For simplicity, when describing IRIS below (and IRIS-U² later on), we focus on the behavior of the algorithm for a given roadmap. More information on how to construct such roadmaps can be found in [12].

Let $\mathcal{I}_{\mathcal{G}} := \{\iota \in \mathcal{I} \mid \exists v \in \mathcal{V} \text{ s.t. } \iota \in \mathcal{S}(v)\}$ be the set of all inspection points that can be inspected from some roadmap vertex. To compute an inspection plan, IRIS considers the *inspection graph* $\mathcal{G}_{\mathcal{S}} = (\mathcal{V}_{\mathcal{S}}, \mathcal{E}_{\mathcal{S}})$ induced by the roadmap \mathcal{G} . Here, vertices are pairs comprised of a vertex $u \in \mathcal{V}$ in the roadmap \mathcal{G} and subsets of $\mathcal{I}_{\mathcal{G}}$. Namely, $\mathcal{V}_{\mathcal{S}} = \mathcal{V} \times 2^{\mathcal{I}_{\mathcal{G}}}$, and note that $|\mathcal{V}_{\mathcal{S}}| = O(|\mathcal{V}| \cdot 2^{|\mathcal{I}_{\mathcal{G}}|})$. An edge $e \in \mathcal{E}_{\mathcal{S}}$ between vertices (u, \mathcal{I}_u) and (v, \mathcal{I}_v) exists if $(u, v) \in \mathcal{E}$ and $\mathcal{I}_u \cup \mathcal{S}(v) = \mathcal{I}_v$. The cost of such an edge is simply the length of the edge $(u, v) \in \mathcal{E}$, namely $\ell(u, v)$. The graph $\mathcal{G}_{\mathcal{S}}$ has the property that a shortest path in the inspection graph corresponds to an optimal inspection path π^* over \mathcal{G} . However, the size of $\mathcal{G}_{\mathcal{S}}$ is exponential in the number of POIs $|\mathcal{I}_{\mathcal{G}}|$. Thus, to reduce the runtime complexity IRIS uses a search algorithm that approximates π^* , which allows to prune the search space of paths over $\mathcal{G}_{\mathcal{S}}$.

Specifically, the approach for pruning the search space used by IRIS is done through the notion of *approximate dominance*, which allows to only consider paths that can significantly improve the quality (either in terms of length or the set of points inspected) of a given path. In particular, let π, π' be two paths in \mathcal{G} that start and end at the same vertices and let $\varepsilon \geq 0$ and $\kappa \in (0, 1]$ be some approximation parameters. We say that π (ε, κ) -dominates π' if $\ell(\pi) \leq (1 + \varepsilon) \cdot \ell(\pi')$ and $|\mathcal{S}(\pi)| \geq \kappa \cdot |\mathcal{S}(\pi) \cup \mathcal{S}(\pi')|$. If π indeed (ε, κ) -dominates π' then π' can potentially be pruned. However, if we prune away approximate-dominated paths, we need to efficiently account for all paths that were pruned away in order to bound the quality of the solution obtained. This is done through the notion of *potentially-achievable paths* described below.

The search algorithm used by IRIS employs an A*-like search over $\mathcal{G}_{\mathcal{S}}$, where each node in the search tree is associated with a *path pair* (PP) corresponding to a vertex v in $\mathcal{G}_{\mathcal{S}}$ (rather than only a path as in A*). Here, a PP is a tuple $(\pi, \tilde{\pi})$, where π and $\tilde{\pi}$ are the so-called *achievable path* (AP) and *potentially achievable path* (PAP), respectively. The AP π represents a realizable path in $\mathcal{G}_{\mathcal{S}}$, from the start vertex to v , and is associated with two scalars corresponding to the path's length and coverage, respectively. The PAP $\tilde{\pi}$ is a pair of scalars $\ell, \bar{\mathcal{S}}$ representing length and coverage, respectively, which are used to bound the quality of any achievable paths to v represented

³ At first glance, one may be tempted to always provide a collision probability of zero. However, this comes at a computational cost—placing unnecessarily tight constraints may yield longer running times and, in extreme cases, problem infeasibility. See Sec. VI for more details.

by a specific PP. Note that $\tilde{\pi}$ does not imply necessarily that there exists any path π' from the start to v such that $\ell(\pi') = \ell(\tilde{\pi})$ and $\mathcal{S}(\pi) = \mathcal{S}(\tilde{\pi})$. It merely states that such a path *could* exist.

A PP is said to be (ε, κ) -bounded if (i) the length of the AP is no more than $(1 + \varepsilon)$ times the length of the PAP and (ii) the coverage of the AP is at least κ percent of the coverage of the PAP.

The search algorithm starts with a path pair rooted at the start vertex v_{start} where both the AP and the PAP represent the trivial paths that only contain v_{start} (i.e., the scalars associated with the length of the AP and the PAP are zero and the scalars associated with the coverage of the AP and the PAP is $\mathcal{S}(v_{\text{start}})$). It operates in a manner similar to A*, with an OPEN list and a CLOSED set to track nodes that have not and have been considered, respectively. Each iteration begins with popping a node from the OPEN list and checking if an inspection path has been found. If this is not the case, the popped node from the OPEN list is inserted into the CLOSED set and the iteration continues.

The next step is *extending* this node and testing whether its successors are *dominated* by an existing node. If this is the case, the node is discarded. Otherwise, IRIS tests whether the node can be *subsumed* by or subsume another node.

These three core operations (extending, dominating, and subsuming) are key to the efficiency of IRIS. When extending a node, a path pair $\text{PP}_u = (\pi_u, \tilde{\pi}_u)$ to some vertex u is extended by an edge $e = (u, v) \in \mathcal{E}_S$ to create the path pair $\text{PP}_v = (\pi_v, \tilde{\pi}_v)$. The length and coverage of the AP of π_v are $\ell(\pi_u) + \ell(e)$ and $\mathcal{S}(\pi_u) \cup \mathcal{S}(v)$, respectively. Similarly, the length and coverage of the PAP $\tilde{\pi}_v$ are $\ell(\tilde{\pi}_u) + \ell(e)$ and $\mathcal{S}(\tilde{\pi}_u) \cup \mathcal{S}(v)$, respectively. When testing domination, two path pairs $\text{PP}_{u,1}$ and $\text{PP}_{u,2}$ to some vertex $u \in \mathcal{V}_S$ are considered. $\text{PP}_{u,1}$ is said to dominate $\text{PP}_{u,2}$ if both $\ell(\pi_{u,1}) \leq \ell(\pi_{u,2})$ and $\mathcal{S}(\pi_{u,2}) \subseteq \mathcal{S}(\pi_{u,1})$. In such a case, $\text{PP}_{u,1}$ is preferable to $\text{PP}_{u,2}$ and so $\text{PP}_{u,2}$ can be discarded. However, in many settings two path pairs will not dominate each other but their respective path lengths and coverage will be similar. To avoid maintaining and extending such similar path pairs, IRIS uses the subsuming operation. This operation is denoted by $\text{PP}_{u,1} \oplus \text{PP}_{u,2}$, which creates a new path pair $\text{PP}_{u,3}$ whose AP's length and coverage are identical to those of $\text{PP}_{u,1}$. The PAP's length of $\text{PP}_{u,3}$ is the minimum PAP's length of PP_1, PP_2 , and the PAP's coverage is the union of the coverage of the PAPs of PP_1 and PP_2 . Subsuming is only performed as long as the resultant PP is (ε, κ) -bounded which allows to guarantee bounds on the solution quality. For additional details, see [12].

The algorithm's asymptotic convergence to an optimal solution is achieved through a process of iterative roadmap densification and parameter tightening. This iterative approach involves systematically refining the roadmap \mathcal{G} (i.e., adding vertices and edges) while progressively reducing the parameters ε and κ . Roadmap densification ensures that the algorithm considers larger sets of configurations, leading to more accurate and refined solutions. Meanwhile, the tightening of parameters focuses the algorithm on increasingly promising paths within the roadmap.

B. Monte-Carlo methods & confidence intervals

Monte Carlo methods are a family of statistical techniques used to simulate and analyze complex systems that involve randomness. These methods involve generating multiple random samples and using them to estimate the value of a process being studied. Due to the finite number of samples, there is uncertainty regarding the true value of the process. To quantify this uncertainty, a common approach is CI and CL [14], [48].

A CI is a range of values likely to contain the true value of a population's parameter (such as its mean) with a certain confidence levels. For example, in our setting this could be the expected path length. The size of the CI reflects the uncertainty around the estimated value and is influenced by the number of samples used. As the number of samples increases, the accuracy of the estimate improves (i.e., CL increases) and the CI decreases. See additional background on the statistical tools we use in Appendix A.

V. METHOD

In this section, we present our method called *IRIS under uncertainty*, or IRIS-U², to solve Prob. 3. This is done by extending the algorithmic framework of IRIS to consider execution uncertainty within the inspection-planning algorithm. We start with a general description of our algorithmic approach (Sec. V-A), and then describe how the operations used in IRIS are modified to account for localization uncertainty (Sec. V-B). This is followed by a toy scenario that is used to demonstrate the key newly-introduced definitions and operations (Sec V-C) We conclude by describing how those updated operations are used by IRIS-U² to compute an inspection path (Sec. V-D).

A. IRIS-U²—Algorithmic approach

A naïve approach to address execution uncertainty is to penalize paths with high localization uncertainty. We describe one such approach as a baseline in Sec. VII-B. As we will see, while highly efficient in collision avoidance, even minor deviations from the command path due to execution uncertainty can lead to discrepancies between the intended POI coverage of the command path and the actual path taken during execution, particularly when obstacles are present. Thus, instead of reasoning about localization uncertainty, in IRIS-U², we directly consider and maximize POI coverage.

Unfortunately, we cannot directly compute the expected values for coverage, collision probability, and length, $|\bar{\mathcal{S}}(\pi)|$, $\bar{\mathcal{C}}(\pi)$ and $\bar{\ell}(\pi)$, respectively. Thus, we tackle Prob. 3 using the estimated values $\hat{\mathcal{S}}(\pi)$, $\hat{\mathcal{C}}(\pi)$ and $\hat{\ell}(\pi)$ instead of the expected values $|\bar{\mathcal{S}}(\pi)|$, $\bar{\mathcal{C}}(\pi)$ and $\bar{\ell}(\pi)$ such that the command path π , satisfies:

$$\hat{\mathcal{C}}(\pi) \leq \rho_{\text{coll}}, \quad (3a)$$

$$|\bar{\mathcal{S}}(\pi)| \geq \kappa \cdot |\hat{\mathcal{S}}(\pi^*)|, \quad (3b)$$

$$\bar{\ell}(\pi) \leq (1 + \varepsilon) \cdot \hat{\ell}(\pi^*). \quad (3c)$$

Here, the estimated coverage $\hat{\mathcal{S}}(\pi)$, the estimated collision probability $\hat{\mathcal{C}}(\pi)$ and the estimated path length $\hat{\ell}(\pi)$ are computed by simulating $m \geq 1$ different executions (with respect to uncertainty) of π (see details below).

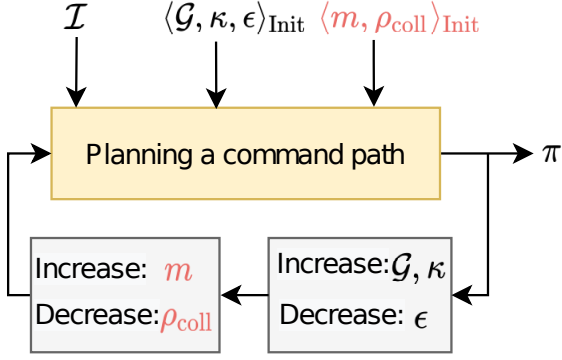


Fig. 2: Planning a command path π to inspect set of POI \mathcal{I} given the input of IRIS— \mathcal{G} , κ and ϵ , and the additional parameters of IRIS-U²— m the ρ_{coll} .

Similar to IRIS, IRIS-U² solves the inspection-planning problem by sampling an initial roadmap. It then iteratively (i) plans a command path on this roadmap and (ii) densifies the roadmap and refines the algorithm’s parameters. Importantly, the focus of this work is on the path-planning part of the algorithmic framework wherein a command path is computed for a given roadmap. This is visualized in Fig. 2 where the path-planning part is highlighted. For completeness, we reiterate that graph refinement is done by continuing to grow the RRG as described in Sec. IV.

Specifically, we start by initializing our algorithm by sampling m different parameters $\mathcal{N}_{\text{unc}}^1, \dots, \mathcal{N}_{\text{unc}}^m$ from \mathcal{D}_{unc} . For any command path π considered by the algorithm, we will estimate $\bar{\ell}(\pi)$ and $\bar{S}(\pi)$ by simulating m execution paths π_1^e, \dots, π_m^e using the m motion models. Namely,

$$\pi_i^e := \mathcal{M}(\pi, \mathcal{N}_{\text{unc}}^i). \quad (4)$$

To estimate the expected coverage, let $\xi_{i,j} \in \{0, 1\}$ be a variable that will be set to one if path π_i^e covers the j ’th POI. Namely,

$$\xi_{i,j} := \begin{cases} 1 & \iota_j \in \mathcal{S}(\pi_i^e), \\ 0 & \text{else.} \end{cases} \quad (5)$$

Then, we define for the command path π the *inspection probability vector* (IPV):

$$\text{IPV}(\pi) := \{\hat{p}_1^\pi, \dots, \hat{p}_k^\pi\}. \quad (6)$$

where:

$$\hat{p}_j^\pi := \frac{1}{m} \sum_{i=1}^m \xi_{i,j}, \quad (7)$$

is the estimated probability that the j ’th POI is viewed when executing the command path π . Finally, the estimated expected coverage is defined as:

$$|\hat{S}(\pi)| := \sum_{j=1}^k \hat{p}_j^\pi. \quad (8)$$

To simplify the theoretical analysis of our algorithm (see Section VI) and its implementation (see Equation (18)), which

depends on the interaction between the probabilities $\{\hat{p}_j^\pi\}$, we make the following assumption:

Assumption 1: The probability \hat{p}_j^π of inspecting each POI j is independent of other POIs.

As we will see in the experimental results (Section VIII), the algorithm empirically achieves the theoretical guarantees obtained from the analysis in practice, even when the assumption is relaxed in a realistic setting. We leave the theoretical implications of relaxing the assumption for future work.

To estimate the collision probability $\hat{C}(\pi)$ of a path π , IRIS-U² maintains for the command path π a *collision vector* (CV):

$$\text{CV}(\pi) := \{\zeta_1^\pi, \zeta_2^\pi, \dots\}. \quad (9)$$

Here, ζ_j^π indicates whether the path π_j^e was found to be in a collision. Namely:

$$\zeta_j^\pi := \begin{cases} 1 & \pi_j^e \text{ is in collision,} \\ 0 & \text{else.} \end{cases} \quad (10)$$

Subsequently, the estimated collision probability $\hat{C}(\pi)$ of a path π is defined as:

$$\hat{C}(\pi) := \frac{1}{m} \sum_{j=1}^m \zeta_j^\pi. \quad (11)$$

Similarly, the estimated expected path length $\bar{\ell}(\pi)$ is defined as:

$$\hat{\ell}(\pi) := \frac{1}{m} \sum_{j=1}^m \ell(\pi_j^e). \quad (12)$$

B. IRIS-U²—Modified search operations

Recall that, while IRIS maintains the POIs inspected using a set representation, IRIS-U² maintains an inspection probability vector (IPV) for each path. In addition, IRIS-U² maintains estimations of the expected path length and expected collision probability of its command path. This requires to modify node operations used by IRIS’s search algorithm to account for the uncertainty values. Next we will explain how to modify the operations used in IRIS-U²’s A*-like search that were originally used in IRIS. We start by formally defining nodes in Sec. V-B1 and detail the node extension, collision, and domination operations in Sec. V-B2, V-B3 and V-B4, respectively. Finally, we describe the node subsuming operation in Sec. V-B5 and termination criteria in Sec. V-B6.

1) *Node definition:* A node n in our search algorithm is a tuple:

$$n = \langle u, \pi_u, \Pi_u^e, \hat{\text{IPV}}_u, \hat{\ell}_u, \hat{\text{IPV}}_u, \bar{\ell}_u, \hat{C}_u \rangle, \quad (13)$$

where $u \in \mathcal{V}$ is a roadmap vertex, π_u is a path from the start vertex $v_{\text{start}} \in \mathcal{G}$ to u , $\Pi_u^e = \{\pi_{1,u}^e, \dots, \pi_{m,u}^e\}$ are m simulated execution paths calculated using Eq. (4), $\hat{\text{IPV}}_u = \{\hat{p}_1^{\pi_u}, \dots, \hat{p}_k^{\pi_u}\}$ is the estimated IPV (see Eq. (6)), and $\hat{\ell}_u$ is the estimated path length (see Eq. (12)) with respect to the command path π_u . In addition, $\hat{\text{IPV}}_u = \{\hat{p}_1^{\pi_u}, \dots, \hat{p}_k^{\pi_u}\}$ and $\bar{\ell}_u$ are the estimated IPV and estimated length of the PAP. Finally, \hat{C}_u is the estimated collision probability of π_u (see Eq. (11)).

We define the *initial node* as:

$$n_{\text{init}} := \langle v_{\text{start}}, \{v_{\text{start}}\}, \Pi_{v_{\text{start}}}^e, \text{IPV}_{v_{\text{start}}}, 0, \text{IPV}_{v_{\text{start}}}, 0, 0 \rangle, \quad (14)$$

where the command path π of n_{init} consists of the trivial path starting and ending at v_{start} whose estimated length and collision are initialized to zero (we assume that v_{start} corresponds to a collision-free configuration). In addition, the m MC starting configurations are set to be v_{start} (namely, $\Pi_{v_{\text{start}}}^e = \{v_{\text{start}}, \dots, v_{\text{start}}\}$), which induce the initial inspection probability vector $\text{IPV}_{v_{\text{start}}} = \{p_{\text{start},1}, \dots, p_{\text{start},k}\}$. Here $p_{\text{start},j} = 1$ if $\iota_j \in S(v_{\text{start}})$ and zero otherwise.

Finally, analogously to IRIS, the IPV and estimated length of the PAP are initialized to be the same as the AP.

2) *Node extension*: Let $u, v \in \mathcal{V}$ be two roadmap vertices such that $(u, v) \in \mathcal{E}$. and let $n_u = \langle u, \pi_u, \Pi_u^e, \text{IPV}_u, \hat{\ell}_u, \text{IPV}_u, \tilde{\ell}_u, \hat{C}_u \rangle$ be a node in the search algorithm associated with vertex u . We define the operation of *extending* n_u by the edge (u, v) as creating a new node $n_v = \langle v, \pi_v, \Pi_v^e, \text{IPV}_v, \hat{\ell}_v, \text{IPV}_v, \tilde{\ell}_v, \hat{C}_v \rangle$ such that:

- π_v is the result of concatenating π_u with the path $\pi_{u \rightarrow v}$ from u to v . Namely,

$$\pi_v := \pi_u \circ \pi_{u \rightarrow v}. \quad (15)$$

- $\Pi_v^e := \{\pi_{1,v}^e \dots \pi_{m,v}^e\}$ is a set of execution paths such that $\pi_{j,v}^e := \mathcal{M}(\pi_v, \mathcal{N}_{\text{unc}}^j)$. Notice that this can be efficiently computed by denoting:

$$\pi_{j,u \rightarrow v}^e := \mathcal{M}(\pi_{u \rightarrow v}, \mathcal{N}_{\text{unc}}^j, \pi_{j,u}^e), \quad (16)$$

and setting:

$$\pi_{j,v}^e := \pi_{j,u}^e \circ \pi_{j,u \rightarrow v}^e. \quad (17)$$

- $\text{IPV}_v := \{\hat{p}_1^{\pi_v}, \dots, \hat{p}_k^{\pi_v}\}$ is obtained by leveraging the assumptions, that (i) POI inspections are conducted at vertices (see Sec. III-A) and that (ii) each inspection is independent (see Assumption. 1).

$$\forall j \in [1, k] \hat{p}_j^{\pi_v} := 1 - (1 - \hat{p}_j^{\pi_{u \rightarrow v}}) \cdot (1 - \hat{p}_j^{\pi_u}). \quad (18)$$

- $\hat{\ell}_v$ is the estimated path length of π_v . Notice that this can be efficiently computed by setting:

$$\hat{\ell}_v := \hat{\ell}_u + \hat{\ell}(\pi_{u \rightarrow v}). \quad (19)$$

- The estimated IPV of the PAP is updated such that:

$$\forall j \in [1, k] \tilde{p}_j^{\pi_v} := 1 - (1 - \hat{p}_j^{\pi_{u \rightarrow v}}) \cdot (1 - \hat{p}_j^{\pi_u}). \quad (20)$$

and the estimated path length of the PAP is:

$$\tilde{\ell}_v := \tilde{\ell}_u + \hat{\ell}(\pi_{u \rightarrow v}). \quad (21)$$

- \hat{C}_v is the collision probability of node v , such that:

$$\hat{C}_v := 1 - (1 - \hat{C}_u) \cdot (1 - \hat{C}_{u \rightarrow v}). \quad (22)$$

3) *Node collision*: Recall that the collision probability of a node estimates the probability that the command path associated with n will intersect an obstacle. Now, let $n_u = \langle u, \pi_u, \Pi_u^e, \text{IPV}_u, \hat{\ell}_u, \text{IPV}_u, \tilde{\ell}_u, \hat{C}_u \rangle$ be a node in the search algorithm associated with vertex u . Then, given a user-defined threshold $\rho_{\text{coll}} \in [0, 1]$ a node n will be considered in *collision* (and hence pruned by the search) if its collision probability \hat{C}_u satisfies $\hat{C}_u \geq \rho_{\text{coll}}$.

4) *Node domination*: As in many A*-like algorithms, node domination is used to prune away nodes that cannot improve the solution compared to other nodes expanded by the algorithm. We introduce a similar notion that accounts for both converge (via IPV's) and path length.

Specifically, let n_1 and n_2 be two nodes that start at $v_{\text{start}} \in \mathcal{V}$ and end at the same vertex $u \in \mathcal{V}$. Here, we assume that for each $i \in \{1, 2\}$,

$$n_i = \langle u, \pi_{u,i}, \Pi_{u,i}^e, \text{IPV}_{u,i}, \hat{\ell}_{u,i}, \text{IPV}_{u,i}, \tilde{\ell}_{u,i}, \hat{C}_u \rangle. \quad (23)$$

Then, we say that n_1 *dominates* n_2 if n_1 's PAP is strictly better than n_2 's. Namely, if

$$\forall j \in [1 \dots k], \hat{p}_j^{\pi_1} \geq \hat{p}_j^{\pi_2} \quad \text{and} \quad \tilde{\ell}(\pi_1) \leq \tilde{\ell}(\pi_2). \quad (24)$$

5) (ε, κ) -*bounded nodes & node subsuming* : Similar to IRIS, we need to ensure that the PAP bounds the AP given the user-provided parameters $\varepsilon > 0$ and $\kappa \in [0, 1]$. Specifically, let $n_u = \langle u, \pi_u, \Pi_u^e, \text{IPV}_u, \hat{\ell}_u, \text{IPV}_u, \tilde{\ell}_u, \hat{C}_u \rangle$ be a node such that $\text{IPV}(\pi_u) := \{\hat{p}_1^{\pi_u}, \dots, \hat{p}_k^{\pi_u}\}$ and $\text{IPV}(\pi_u) := \{\tilde{p}_1^{\pi_u}, \dots, \tilde{p}_k^{\pi_u}\}$ are the IPV's of n 's AP and PAP, respectively. Similarly, let $\hat{\ell}_u$ and $\tilde{\ell}_u$ be the estimated lengths of n 's AP and PAP, respectively. We say that n is (ε, κ) -*bounded* if:

$$\sum_{j=1}^{j=k} \hat{p}_j^{\pi_u} \geq \kappa \cdot \sum_{j=1}^{j=k} \tilde{p}_j^{\pi_u} \quad \text{and} \quad \hat{\ell}_u \leq (1 + \varepsilon) \cdot \tilde{\ell}_u. \quad (25)$$

Similar to IRIS, (ε, κ) -bounded nodes will be used together with node subsuming to reduce the number of paths considered by the search while retaining bounds on path quality. Specifically we define node subsuming as follows: Let n_1 and n_2 be two nodes both starting at the same vertex and ending at the same vertex u such that:

$$n_i = \langle u, \pi_{u,i}, \Pi_{u,i}^e, \text{IPV}_{u,i}, \hat{\ell}_{u,i}, \text{IPV}_{u,i}, \tilde{\ell}_{u,i}, \hat{C}_u \rangle. \quad (26)$$

Then, the operation of n_1 *subsuming* n_2 (denoted as $n_1 \oplus n_2$) will create the new node (that will be used to replace n_1 and prune n_2):

$$n_1 \oplus n_2 := \langle u, \pi_{u,3}, \Pi_{u,3}^e, \text{IPV}_{u,3}, \hat{\ell}_{u,3}, \text{IPV}_{u,3}, \tilde{\ell}_{u,3}, \hat{C}_u \rangle. \quad (27)$$

Here, the components of n_3 's AP are identical to n_1 's. Namely, n_3 's components associated with the AP are defined




Fig. 3: Toy scenario for running example with one POI (red dot), two obstacles (red rectangles), and four vertices (black dots) connected with six edges (black dashed lines).

as follows:

$$\begin{aligned}\pi_{u,3} &= \pi_{u,1}, \\ \Pi_{u,3}^e &= \Pi_{u,1}^e, \\ \text{IPV}_{u,3} &= \text{IPV}_{u,1}, \\ \hat{\ell}_{u,3} &= \hat{\ell}_{u,1}, \\ \hat{C}_{u,3} &= \hat{C}_{u,1},\end{aligned}$$

and n_3 's components associated with the PAP are defined as follows:

$$\begin{aligned}\tilde{\ell}_{u,3} &:= \min(\tilde{\ell}_{u,1}, \tilde{\ell}_{u,2}), \\ \hat{p}_j^{\pi_{u,3}} &:= \max(\hat{p}_j^{\pi_{u,1}}, \hat{p}_j^{\pi_{u,2}}).\end{aligned}$$

6) *Termination criteria*: To terminate the search in IRIS-U² given a node n_v we check whether n_v satisfies:

$$\sum_{j=1}^{j=k} \hat{p}_j^{\pi_v} \geq k \cdot \kappa. \quad (28)$$

Here $\{\hat{p}_1^{\pi_v}, \dots, \hat{p}_k^{\pi_v}\}$ is the IPV of n_v 's AP and k is the number of POIs.

C. Modified search operation—illustrative example

Consider the toy problem illustrated in Fig. 3. The roadmap \mathcal{G} contains four vertices A , B , C , and D which represent configurations of a point robot (namely, each configuration defines the (x, y) location of the point robot). Here, the single POI can be seen from any configuration as long as the straight-line connecting them does not intersect an obstacle.

We use a simple toy motion model \mathcal{M}_{toy} to define the execution localization uncertainty in which we assume that the position of every configuration along the command path is normally distributed around the position of the corresponding configuration. Specifically, parameters $\mathcal{N}_{\text{unc}} := (r, \theta)$ are drawn from the distribution \mathcal{D}_{unc} such that $r \sim |\mathcal{N}(0, 1)|$, and $\theta \sim \mathcal{N}(0, 2\pi)$. Now, given a command path π^c with $n > 1$ configurations $\pi^c = \{\langle x_1^c, y_1^c \rangle, \dots, \langle x_n^c, y_n^c \rangle\}$, the corresponding executed path (which is also a random variable) is $\pi_e = \{\langle x_1^e, y_1^e \rangle, \dots, \langle x_n^e, y_n^e \rangle\}$ such that $x_i^e = x_i^c + r \cdot \cos \theta$ and $y_i^e = y_i^c + r \cdot \sin \theta$ for $i > 1$.

Note that (i) as we assume that there is no uncertainty in the initial location, $\langle x_1^c, y_1^c \rangle = \langle x_1^e, y_1^e \rangle$ and that (ii) in contrast to the general setting, here uncertainty is only a function of the current configuration and *not* of the entire command path.

Finally, in the running example, we assume that the algorithm uses three MC planning samples ($m = 3$) and the start configuration is located at vertex A .

1) *Node definition—example*: For the running example in our toy problem, the first location of each MC planning sample is at v_{start} and the POI cannot be inspected from that location. Specifically, n_{init} is defined such that:

$$\begin{aligned}v_{\text{start}} &= A, \\ \Pi_{v_{\text{start}}}^e &= \{\langle -1, 1 \rangle; \langle -1, 1 \rangle; \langle -1, 1 \rangle\}, \\ \text{IPV}_{v_{\text{start}}} &= \{\hat{p}_1^{\pi_{v_{\text{start}}}}\} = \{0\}.\end{aligned}$$

2) *Node extension—example*: Extending the node n_{init} by edge $e(A, B)$ will result in a new node:

$$n_B = \langle B, \pi_B, \Pi_B^e, \text{IPV}_B, \hat{\ell}_B, \text{IPV}_B, \tilde{\ell}_B, \hat{C}_B \rangle,$$

where $\pi_B = \{A, B\}$ and Π_B^e is calculated using \mathcal{M}_{toy} to be:

$$\begin{aligned}\Pi_B^e &= \{\langle -1, 1 \rangle, \langle 0.6, 2.1 \rangle; \\ &\quad \langle -1, 1 \rangle, \langle 0.5, 1.2 \rangle; \\ &\quad \langle -1, 1 \rangle, \langle -0.2, 1.6 \rangle\}.\end{aligned}$$

As a result, since POI 1 can only be seen from $\Pi_{B,2}^e$ and $\Pi_{B,3}^e$, then $\hat{p}_1^{\pi_{A \rightarrow B}} \cong 0.67$. Using these values, we can calculate $\hat{p}_1^{\pi_B}$ and $\hat{p}_1^{\tilde{\pi}_B}$ as:

$$\begin{aligned}\hat{p}_1^{\pi_B} &= 1 - (1 - \hat{p}_1^{\pi_{A \rightarrow B}}) \cdot (1 - \hat{p}_1^{\pi_A}) = \\ &= 1 - (1 - 0.67) \cdot (1 - 0) = 0.67, \\ \hat{p}_1^{\tilde{\pi}_B} &= \max(\hat{p}_1^{\pi_{A \rightarrow B}}, \hat{p}_1^{\pi_A}) = 0.67.\end{aligned}$$

Thus, $\text{IPV}_B = \text{IPV}_B = 0.67$. As $\hat{\ell}(\pi_{u \rightarrow v}) \cong 1.48$ (using Π_B^e), we have that $\hat{\ell}_B = \tilde{\ell}_B = 0 + 1.48$. Finally, $\Pi_{B,1}^e$ collides with the middle obstacle, and $\hat{C}_B = \frac{1}{3}$.

3) *Node collision—example*: Assuming $\rho_{\text{coll}} = 0$ (i.e., we only allow collision-free paths), then $\hat{C}_B = \frac{1}{3} \geq \rho_{\text{coll}}$ is considered in collision and this extension is discarded.

4) *Node domination—example*: Assume that node $n_{D,1}$ represents the command path $A-D$ in the running example and its IPV and estimated length of the AP and the PAP are:

$$\begin{aligned}\text{IPV}_{D,1} &= \{0.67\}, \quad \hat{\ell}_{D,1} = 2.4 \\ \text{IPV}_{D,1} &= \{0.67\}, \quad \tilde{\ell}_{D,1} = 2.4.\end{aligned}$$

In addition, assume that later in the search node $n_{D,2}$ represents another command path $A-C-D$. This path also reached vertex D and its IPV and estimated length of the AP and the PAP are:

$$\begin{aligned}\text{IPV}_{D,2} &= \{0.77\}, \quad \hat{\ell}_{D,2} = 2.9 \\ \text{IPV}_{D,2} &= \{0.77\}, \quad \tilde{\ell}_{D,2} = 2.9.\end{aligned}$$

Here, the estimated length of 2.9 is the sum of the estimated lengths of $A-C$ and $C-D$. In addition, the IPV of the AP and PAP contains a probability of 0.33 to inspect the POI from vertex C and 0.67 to inspect it from vertex D . Thus, the value

of $\hat{\text{IPV}}_{D,2}$ equals $1 - (1 - 0.67) \cdot (1 - 0.33) \cong 0.77$ (using Eq. (18)). Similarly, the value of $\tilde{\text{IPV}}_{D,2}$ is also equal to 0.77. Notice that, the path lengths and IPV's of the AP and PAP are identical here. This will change shortly when we introduce node subsuming (Sec V-B5).

Here, despite that $\tilde{\ell}_{D,1} = 2.4 \leq 2.9 = \tilde{\ell}_{D,2}$ (namely, the path to $n_{D,1}$ is shorter than the path to $n_{D,2}$), we have that $\hat{p}_1^{\tilde{\pi}_{D,1}} = 0.67 \leq 0.77 = \hat{p}_1^{\tilde{\pi}_{D,2}}$ (namely, the path to $n_{D,1}$ has a smaller probability of inspecting the POI). Thus, $n_{D,1}$ does not dominates $n_{D,2}$.

5) *Node subsuming—example:* Consider the operation $n_{D,1} \oplus n_{D,2}$. This results in a new node $n_{D,3}$ such that its IPV and estimated length of the AP and the PAP are:

$$\begin{aligned} \hat{\text{IPV}}_{D,3} &= \{0.67\}, \\ \tilde{\text{IPV}}_{D,3} &= \max(\{0.67\}, \{0.77\}) = \{0.77\} \\ \hat{\ell}_{D,3} &= 2.4, \\ \tilde{\ell}_{D,3} &= \min(2.4, 2.9) = 2.4. \end{aligned}$$

Here, if we choose $\kappa = 0.85$ and $\varepsilon = 0$, then $n_{D,3}$ is (ε, κ) -bounded, namely:

$$\begin{aligned} \sum_{j=1}^{i=k} \hat{p}_j^{\pi_{D,3}} &= 0.67; \quad \kappa \cdot \sum_{j=1}^{i=k} \hat{p}_j^{\tilde{\pi}_{D,3}} = 0.85 \cdot 0.77 \cong 0.66; \\ \hat{\ell}_{D,3} &= 2.4; \quad (1 + \varepsilon) \cdot \tilde{\ell}_{D,3} = 2.4. \end{aligned}$$

and indeed both $0.67 \geq 0.66$ and $2.4 \geq 2.4$

6) *Termination criteria—example:* As a result of the subsume operation $n_{D,3} = n_{D,2} \oplus n_{D,1}$ we have that the IPV of $n_{D,3}$'s AP and PAP are:

$$\hat{\text{IPV}}_{D,3} = \tilde{\text{IPV}}_{D,3} = \{0.77\}.$$

Now, assume that $\kappa = 0.97$ and recall that we have one POI. As $0.77 < 1 \cdot 0.97$, the algorithm cannot terminate. However, assume we extend this path by returning to vertices C and D to obtain the command path $A-C-D-C-D-C-D$ in which we perform inspection three times from C and three times from D and recall that the coverage probability of C and D equals 0.33 and 0.67, respectively. Thus, the coverage of the command path $A-C-D-C-D-C-D$ equals:

$$1 - \underbrace{(1 - 0.33)^3}_{\text{Three times } C} \cdot \underbrace{(1 - 0.67)^3}_{\text{Three times } D} \cong 0.98.$$

Finally,

$$\sum_{j=1}^{j=k} \hat{p}_j^{\tilde{\pi}_{A-C-D-C-D-C-D}} = 0.98 \geq 1 \cdot 0.97,$$

and the algorithm terminates.

D. IRIS-U²—Algorithmic description

In the previous sections, we described how IRIS-U² modifies the search operations of IRIS to account for localization uncertainty. In this section, we complete the description of the algorithm. As we will see, despite these modifications, the high-level framework of IRIS remains the same, and subsequently,

Algorithm 1 IRIS-U²

Input: $\langle \mathcal{G}, m, \kappa, \varepsilon, \rho_{\text{coll}} \rangle$
Output: Command path π

```

1: Initialize  $n_{\text{init}}$  //See V-B1 & Eq. (14)
2: OPEN  $\leftarrow n_{\text{init}}$ , CLOSED  $\leftarrow \emptyset$ 
3: while OPEN  $\neq \emptyset$  do
4:    $n \leftarrow \text{OPEN.extract\_node\_with\_max\_coverage}()$ 
5:   CLOSED.insert( $n$ )
6:   if to_terminate( $n$ ) then //See V-B6 & Eq. (28)
7:     return  $n.\pi$  //command path
8:   for  $v \in \text{neighbour}(u, \mathcal{G})$  do //  $u$  is  $n$ 's vertex
9:      $n' \leftarrow n.\text{extend}(u, v)$  //See V-B2 & Eq.(15)–(22)
10:    if  $n.\text{is\_in\_collision}()$  then //See V-B3
11:      continue
12:    valid = True
13:    for  $n'' \in \text{CLOSED}$  with vertex  $v$  do
14:      if  $n''$  dominates  $n'$  then //See V-B4 & Eq. (24)
15:        valid = False
16:        break
17:    if !valid then //  $n'$  was dominated
18:      continue
19:    for  $n'' \in \text{OPEN}$  with vertex  $v$  do
20:      if  $n'' \oplus n'$  is  $(\varepsilon, \kappa)$ -bounded then //See V-B5 & Eq. (25)
21:        if  $n''$  has a better AP coverage than  $n'$  then
22:           $n'' \leftarrow n'' \oplus n'$  //See V-B5 & Eq. (27)
23:          valid = False
24:          break
25:    if !valid then //  $n'$  was subsumed
26:      continue
27:    for  $n'' \in \text{OPEN}$  with vertex  $v$  do
28:      if  $n' \oplus n''$  is  $(\varepsilon, \kappa)$ -bounded then //See V-B5 & Eq. (25)
29:        if  $n'$  has a better AP coverage than  $n''$  then
30:          OPEN.remove( $n''$ ) //  $n''$  was subsumed
31:           $n' \leftarrow n' \oplus n''$  //See V-B5 & Eq. (27)
32:    OPEN.insert( $n'$ )
33: return NULL

```

its original guarantees, such as asymptotic convergence to an optimal solution. This is done while also incorporating execution uncertainty and providing statistical guarantees (Sec. VI). To this end, we proceed to outline an iteration of IRIS-U² given a tuple $\langle \mathcal{G}, m, \kappa, \varepsilon, \rho_{\text{coll}} \rangle$ (see Fig. 2) whose pseudo-code is detailed in Alg. 1.

In particular, similar to IRIS's graph search, IRIS-U² uses a priority queue OPEN and a set CLOSED while ensuring that all nodes are always (ε, κ) -bounded. IRIS-U² starts with an empty CLOSED list and with an OPEN list initialized with the start node n_{init} . At each step, the search proceeds by iteratively popping a node n from the OPEN list whose PAP coverage IPV_u is maximal.⁴ Then, if n satisfies the termination criteria (see Eq. (28)) we terminate the search and return n 's command path (Lines 4-7). Otherwise, we create a new node n' by extending n along its neighboring edges (Lines 8-9). However, if the estimated probability of collision for node n' exceeds a certain threshold ρ_{coll} , then the newly created node n' is discarded (Lines 10-11).

If n' was not discarded, then, we perform the following operations (here we assume that the vertex corresponding

⁴As described in the original exposition of IRIS, we can order the OPEN list either according to the PAP coverage or the AP coverage.

with n' is v):

- We start by discarding n' if there exists a node n'' in CLOSED that also reaches v and dominates n' (Lines 13-18).
- We continue by testing whether there exists a node n'' in OPEN that also reaches v that may subsume n' , be (ε, κ) -bounded (Lines 19-26) and has a better AP coverage than n' (Line 21). If so, n' is discarded and n'' is set to be $n'' \oplus n'$.
- Then, we test whether there exists a node n'' in OPEN that also reaches v that may be subsumed by n' while the resultant node being (ε, κ) -bounded (Lines 27-31), and the AP coverage of n' is better than n'' (Line 29). If so, n'' is removed from OPEN and n' is set to be $n' \oplus n''$.

Finally, if n' was not discarded, it is inserted into the OPEN list (Line 32).

Note. When there is no execution uncertainty, running IRIS-U² with $m = 1$ is identical to IRIS.

VI. IRIS-U²—STATISTICAL GUARANTEES

In this section, we detail in Sec. VI-A different statistical guarantees regarding a given command path (proofs are provided in Appendix B). Then, we discuss the implication for the command path computed by IRIS-U² in Sec. VI-B and provide guidelines on how to choose parameters for IRIS-U² given the statistical guarantees and the implications above. We emphasize that the theoretical guarantees obtained here are under Assumption 1, which states that POI inspection probabilities $\{\hat{p}_j^\pi\}$ are independent of one another.

A. Guarantees for a given command path

Consider a command path π and consider m MC simulated executions of π such that $\text{IPV}(\pi) := \{\hat{p}_1^\pi, \dots, \hat{p}_k^\pi\}$ is the associated inspection probability vector, $\hat{C}(\pi)$ is the associated estimated collision probability, and $\hat{\ell}(\pi)$ and $\hat{s}_{\hat{\ell}}(\pi)$ are the associated average and standard deviation of the path's length, respectively.

Lemma 6.1 (Executed path's expected coverage): For any desired CL of $1 - \alpha \in [0, 1]$, the expected coverage of an executed path following π , denoted by $|\bar{\mathcal{S}}(\pi)|^-$, is at least:

$$|\bar{\mathcal{S}}(\pi)|^- := \sum_{j=1}^{j=k} \hat{p}^-(\hat{p}_j^\pi, m, \alpha). \quad (29)$$

Here, the function \hat{p}^- is obtained from the Clopper-Pearson method [49] and is defined in Eq. (35) in Appendix A.

Lemma 6.2 (Executed path's collision probability): For any desired CL of $1 - \alpha \in [0, 1]$, the expected collision probability of an executed path following π , denoted by $\bar{\mathcal{C}}(\pi)$ is at most:

$$\bar{\mathcal{C}}(\pi)^+ := \hat{p}^+(\hat{C}(\pi), m, \alpha). \quad (30)$$

Here, the function \hat{p}^+ is defined in Eq. (35).

Lemma 6.3 (Executed path's expected length): For any desired CL of $1 - \alpha \in [0, 1]$, the expected length of an executed path following π , denoted by $\bar{\ell}$ is bounded such that:

$$\bar{\ell}(\pi) \in [\bar{X}^-(\hat{\ell}(\pi), m, \alpha), \bar{X}^+(\hat{\ell}(\pi), m, \alpha)]. \quad (31)$$

Here, \bar{X}^- and \bar{X}^+ are defined in Eq. (37).

Before stating our final Lemma, we introduce the following assumption:

Assumption 2: For any fixed values of m and α , the function $\hat{p}^-(\hat{p}_j^\pi, m, \alpha)$, which depends solely on \hat{p}_j^π , is both monotonically increasing and strictly convex.

Lemma 6.4 (Bounding executed path's sub-optimal coverage): Recall that for any desired CL of $1 - \alpha \in [0, 1]$, $|\bar{\mathcal{S}}(\pi)|^-$ is the lower bound value of the expected coverage of an executed path following π and can be expressed as $\sum_{j=1}^k \hat{p}^-(\hat{p}_j^\pi, m, \alpha)$ (see Lemma 6.1 and Eq. (35)).

If Assumption 2 holds, then minimizing $|\bar{\mathcal{S}}(\pi)|^-$ subject to (i) $\hat{p}_j^\pi \in [0, 1]$ for all j , and (ii) $\sum_{j=1}^k \hat{p}_j^\pi \geq \kappa \cdot k$ yields that $\forall j, \hat{p}_j^\pi = \kappa$. Consequently,

$$\min_{\text{IPV}(\pi)} |\bar{\mathcal{S}}(\pi)|^- = k \cdot \hat{p}^-(\kappa, m, \alpha). \quad (32)$$

As we will see, Assumption 2 is used to prove Lemma 6.4. Proving that Assumption 2 holds is non-trivial. However, in Appendix B we numerically demonstrate that it holds for all tested values.

B. Implication of statistical guarantees to IRIS-U²

One may be tempted to use the bounds on the expected coverage (Lemma 6.1) and the collision probability (Lemma 6.2) on the command path computed by IRIS-U². Indeed, these bounds hold if the estimations (e.g., the path's IPV) computed via MC simulated executions were computed *after* the command path was computed by IRIS-U² and not on the fly *while* the command path is computed by IRIS-U². That is, using these guarantees may lead to false negatives (e.g., estimation of a collision-free path despite the expectation of a collision occurring) since the command path is computed from a pool of multiple optional paths. A detailed illustrative example to explain this is provided in Appendix C.

To summarize, the different statistical guarantees provided in Sec. VI-A can be used if the command path output by IRIS-U² is simulated multiple times (an alternative is to use the Bonferroni correction, which is a multiple-comparison correction used when conducting multiple dependent or independent statistical tests simultaneously [50]). However, we can use them to understand the relationship between the system's parameters m , κ , and ρ_{coll} and as *guidelines* on how to choose them.

Lemma 6.4—implications: Recall that Lemma 6.4 states that for any desired CL, and regardless of the values of m and κ , the minimum lower bound on the executed path's coverage is $\min_{\text{IPV}(\pi)} |\bar{\mathcal{S}}(\pi)|^- = k \cdot \hat{p}^-(\kappa, m, \alpha)$. This allows us to provide guidelines on how to choose the algorithm's parameters m and κ according to the desired CL which is application specific. As an example, in Fig. 4a we plot $\hat{p}^-(\kappa, m, \alpha)$ for $\alpha = 0.05$ for various values of m and κ . Now, consider a user requirement that the POI coverage of the executed path will exceed 93% with a CL of 95% (i.e., $\alpha = 0.05$). This corresponds to choosing any point on the line of 93% which can be, e.g., $m = 70$ and $\kappa = 0.99$ or $m = 95$ and $\kappa = 0.98$.




Fig. 4: (a), (b) Values (in percentage) of $\hat{p}^-(\kappa, m, \alpha)$ and $\hat{p}^+(\rho_{\text{coll}}, m, \alpha)$ for $\alpha = 0.05$ as a function of the number of MC samples m (x -axis) and the coverage approximation factors κ and ρ_{coll} (y -axis), respectively.

Lemma 6.2—implications: Similar to Lemma 6.4, Lemma 6.2 can be used as a guideline on how to choose the algorithm’s parameters m and ρ_{coll} according to the desired CL. As an example, in Fig. 4b we plot $\hat{p}^+(\rho_{\text{coll}}, m, \alpha)$ for $\alpha = 0.05$ and various values of m and ρ_{coll} . Now, consider a user requirement that the executed path’s collision probability does not exceed 7% with a CL of 95% (i.e., $\alpha = 0.05$). This can be achieved by selecting a point on the 7% line, for instance, $m = 94$ and $\rho_{\text{coll}} = 0.02$.

Lemma 6.3—note: Both Lemma 6.4 and Lemma 6.2 gave clear guidelines on how to choose parameters for desired confidence levels. This was possible because there exists a bound on the best possible outcome (i.e., 100% coverage of POIs and 0% collision probability) and the system parameters κ and ρ_{coll} are defined with respect to these bounds. In contrast, there is no a-priory bound on path length and the parameter ε is only defined with respect to the (unknown) optimal length (i.e., $\hat{\ell}_u \leq (1 + \varepsilon) \cdot \ell_u$).

VII. ILLUSTRATIVE SCENARIO

In this section, we demonstrate the performance of IRIS- U^2 in a toy scenario using a simple motion model. We start (Sec. VII-A) by describing the setting and continue to describe the methods we will be comparing IRIS- U^2 with (Sec. VII-B). We finish with a discussion of the results and their implications (Sec. VII-C). All tests were run on an Intel(R) Core(TM) i7-4510U CPU @ 2.00GHz with 12GB of RAM. The implementation IRIS- U^2 algorithm is available at <https://github.com/CRL-Technion/IRIS-UU.git>.

A. Setting

Here we consider the toy scenario depicted in Fig. 5a in which we have a two-dimensional workspace that contains 10 obstacles (red rectangles) and 27 POIs (red points in seven groups of three). The robot is described by three degrees of freedom—its location (x, y) and heading ψ . We model its sensor as having a field-of-view of 94° and a range of $10[m]$.

We choose a roadmap \mathcal{G} with 27 vertices (where vertex 0 is the initial vertex) such that the configuration associated with every vertex is facing up towards the POIs. We note that in the toy scenario, we hand-pick the roadmap only for illustrative purposes. In practical settings the roadmap would be generated

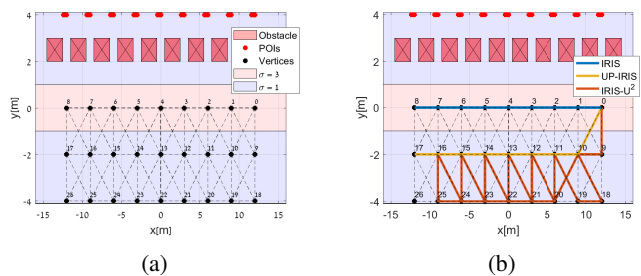


Fig. 5: (a) Illustrative scenario for empirical evaluation depicting a roadmap, (black points being vertices and dashed lines being edges), POIs (red points), obstacles (red rectangles) and regions with high and low uncertainty. (b) Command paths computed by IRIS (blue), UP-IRIS (yellow), and IRIS- U^2 (orange) with $m = 100$.

by the systematic approach of IRIS (as in the following section). The first three POIs can be seen from vertices 0, 9 and 18. Similarly, the next three POIs can be seen from vertices 1, 10 and 19 and so on.

Finally, we assume that the environment contains two types of regions corresponding to different levels of uncertainty (to be explained shortly): the first level (pink, containing vertices 0–8) with a standard deviation of $\sigma = 3[m]$ and the second level (light blue, containing vertices 9–26) with $\sigma = 1[m]$.

The motion model we use here, denoted as $\mathcal{M}_{\text{simple}}^{2d}$, is an extension of \mathcal{M}_{toy} (Section V-C) where the parameters $\mathcal{N}_{\text{unc}} := (r, \theta)$ are drawn from the distribution \mathcal{D}_{unc} such that $r = |\mathcal{N}(0, 1)|$ if the robot is located in low (pink) uncertainty region and $r = |\mathcal{N}(0, 3)|$ if the robot is located in high (light blue) uncertainty region. Additionally $\theta = \mathcal{N}(0, 2\pi)$. Just as in \mathcal{M}_{toy} , given a command path π^c with $n > 1$ configurations $\pi^c = \{(x_1^c, y_1^c), \dots, (x_n^c, y_n^c)\}$, the corresponding executed path is $\pi_e = \{(x_1^e, y_1^e), \dots, (x_n^e, y_n^e)\}$ such that $x_i^e = x_i^c + r \cdot \cos \theta$ and $y_i^e = y_i^c + r \cdot \sin \theta$ for $i > 1$.

B. Baselines

We consider two baselines—the original IRIS algorithm and a straw-man approach which we call uncertainty-penalizing IRIS (UP-IRIS). In UP-IRIS, the cost of an edge is its length added to a penalty factor which is proportional to the uncertainty associated with the edge. As the original IRIS minimizes path length, this modification will compute paths that are both short and have low uncertainty.

For each algorithm, we compute a command path π^c and compare performance in the execution phase. In preparation for the results, we highlight the distinction between MC samples used in planning and in execution. Planning MC samples are used by IRIS- U^2 to compute the command-path (referred to as m). In contrast, execution MC samples do not affect the command-path and are only used to evaluate the performance of the executed path and IRIS- U^2 ’s CI boundaries.

C. Results

To compare IRIS- U^2 with our baselines, IRIS and UP-IRIS, we ran the planning phase with $\varepsilon = 3$ and $\kappa = 0.99$. As

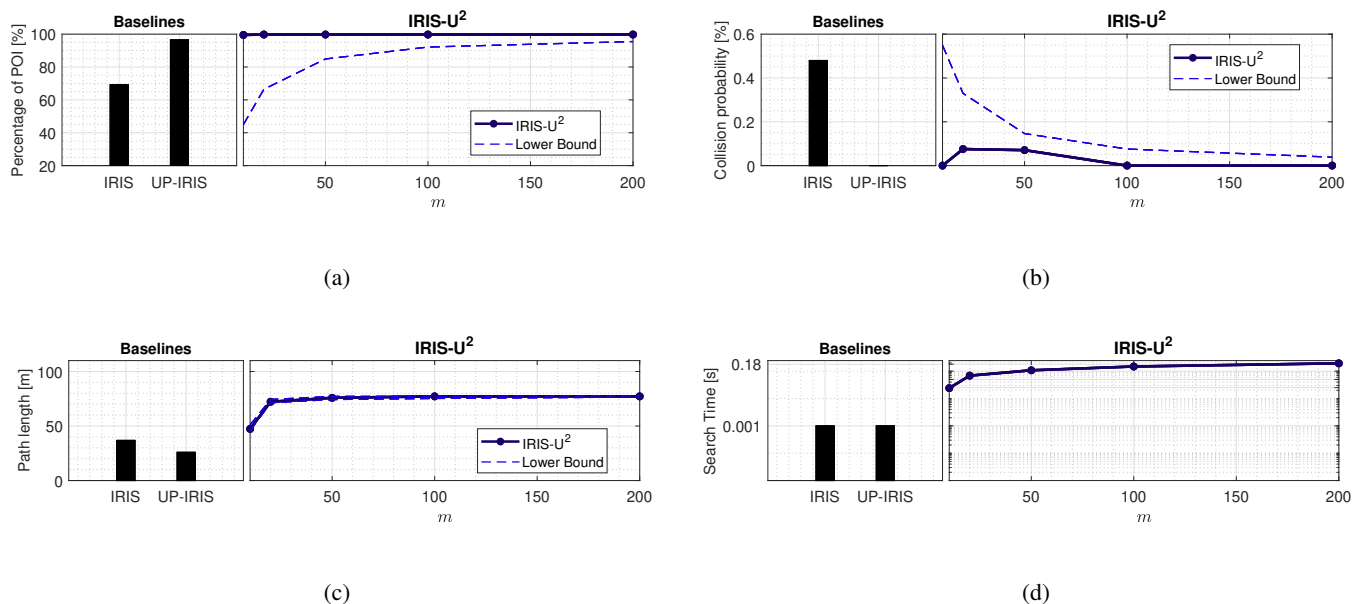


Fig. 6: (a)-(d) POI coverage, collision probability, path length, and search time as a function of m for IRIS-U² (right) and for the baselines (left) for the illustrative scenario. Each of figures (a)-(c) displays the average of 10,000 MC execution samples, with the corresponding CI of the expected performance as detailed in Lemma. 6.1, 6.2, and 6.3.

expected, IRIS (without accounting for uncertainty) computed a command path which is the straight line connecting vertices 0 and 8 (blue path in Fig. 5b) as it is the shortest path that results in full coverage (when ignoring uncertainty). UP-IRIS on the toy scenario computes a command path (depicted as the yellow path in Fig. 5b) which prioritizes regions with low localization uncertainty. This path moves towards vertex 10 and then proceeds in a straight line towards vertex 17 while staying within the low-uncertainty region. Similarly, the command path computed by IRIS-U² (represented by the red path in Fig. 5b) also prioritizes regions with low localization uncertainty, even though this is not its explicit objective. However, instead of simply traversing through the low-uncertainty region once, IRIS-U² revisits multiple vertices to ensure full coverage.

Numerical results concerning these command paths in the execution phase are depicted in Fig. 6. IRIS and UP-IRIS achieved an average coverage of 69% and 96%, respectively. Both being lower than the at least 99% achieved by IRIS-U² for all values of m . Additionally, the collision probability of IRIS's path was found to be 47% as opposed to the collision-free execution path of UP-IRIS and IRIS-U² (for $m \geq 100$).

When looking at the performance of IRIS-U² as a function of the number of MC samples m , one can see (as to be expected) that the expected POI covered increases up to 99% and that the collision probability decreases down to 0% as m increases (Fig. 6a and 6b, respectively). This comes at the price of longer paths and longer computation times (Fig. 6c and 6d, respectively).

Note that in all cases, the CI bounds hold empirically. This is important as our analysis relies on the fact that the probability of inspecting each POI is independent of other POIs (Assumption 1), which may not necessarily hold.

VIII. BRIDGE SCENARIO

In the following set of experiments, we consider a realistic bridge-inspection scenario to assess the performance of IRIS-U². The evaluation consists of two parts wherein both parts use a well-established approximation of the full motion model, which we call the *simplified model*, in the planning stage. In the first part, we use the same simplified motion model to evaluate the command path in the execution phase. This experiment demonstrates our statistical bounds (which hold in the setting that the planning-stage model is identical to the execution-stage model). In the second part, we use a more accurate model to evaluate the command path in the execution phase, which is much more computationally demanding (see discussion below). This part aims to demonstrate the performance of our algorithm on a more realistic motion model and explores the implications of having a simplified (but computationally efficient) model in the planning stage, and a more accurate model (but also more computationally demanding) in the execution stage (both of which accumulate errors in the GNSS-denied region). In our setting, the mismatch arises due to computational considerations (see discussion below). We first describe below the scenario and then proceed to the experimental results for each of the motion models.

A. Setting

We consider a UAV with six degrees of freedom corresponding to its location $(x, y, z) \in \mathbb{R}^3$ and its orientation $(\phi, \theta, \psi) \in \mathbb{R}^3$. We model the sensor as having a field-of-view of 94° and a range of $10m$. We use the 3D model of a bridge⁵, as depicted in Fig. 7 and set the values $\varepsilon = 3$ and $\kappa = 0.9$. Following [12]

⁵Model taken from <https://github.com/UNC-Robotics/IRIS>




Fig. 7: Bridge scenario with 533 POIs (marked in red). Search is performed on a roadmap \mathcal{G} (marked in black).

the roadmap \mathcal{G} was generated using an RRG [47] with 100 vertices and 1224 edges.

a) *Exact motion model:* We make use of a highly realistic model, which we denote by $\mathcal{M}_{\text{exact}}$ which requires fusing measurements from the GNSS and the inertial navigation system (INS) via an EKF [40], known as inertial navigation system GNSS-INS fusion. Our implementation⁶ of $\mathcal{M}_{\text{exact}}$ which uses an EKF is based on the quadcopter dynamics of ArduCopter⁷, an open-sourced quadrotor system with an adaptation of the EKF fusion. However, $\mathcal{M}_{\text{exact}}$ demands significant computational resources, rendering it impractical during the planning stage. In particular, in our implementation, the exact model is typically two orders of magnitude slower than the simplified model. Consequently, we restrict the utilization of this exact model solely to the execution phase, facilitating the evaluation of our approach. In the planning stage, we rely on the following simplified model.

b) *Simplified motion model:* During the planning stage, we consider a computationally efficient and well-established approximation model from [40], which we denote by $\mathcal{M}_{\text{simplified}}$. Indeed, our experimental results below demonstrate that $\mathcal{M}_{\text{simplified}}$ serves as a good proxy for $\mathcal{M}_{\text{exact}}$ as the predicted behavior during the planning stage correlates with the execution behavior. Specifically, in our model $\mathcal{M}_{\text{simplified}}$, we assume that outside the bridge, the GNSS-INS system can achieve an accuracy with an error of 1σ (in meters) around a vertex location. In contrast, beneath the bridge, where GNSS signal reception is typically compromised, we assume accumulating uncertainty over time. For this region, we adopt a model [40] that assumes that the robot moves at a constant speed in a straight line between nodes. Specifically, the location uncertainty of an execution path, $\langle x_1^e, y_1^e, z_1^e \rangle; \langle x_2^e, y_2^e, z_2^e \rangle; \dots$, when following the command path, $\langle x_1^c, y_1^c, z_1^c \rangle; \langle x_2^c, y_2^c, z_2^c \rangle; \dots$, can be expressed as follows:

$$[x_i^e, y_i^e, z_i^e] = [x_i^c, y_i^c, z_i^c] + \frac{1}{2} C_b \mathbf{b}_a t^2 + \frac{1}{6} C_b (\mathbf{b}_g \times \mathbf{g}) t^3. \quad (33)$$

⁶See <https://github.com/CRL-Technion/Simulator-IRIS-UU.git>.

⁷See <https://wilselby.com/research/arducopter/>.

Here, $\mathbf{g} = (0, 0, -g)$ is the gravity vector, \mathbf{b}_a and \mathbf{b}_g are the accelerometer and gyro biases. Additionally, C_b denotes the rotation matrix transforming from the body frame to the inertial frame, and t denotes the continuous time spent in GNSS outage regions (beneath the bridge).

B. Evaluation

Similar to Sec. VII, we compare IRIS-U² with our two baselines IRIS and UP-IRIS. However, for each of the three algorithms, we consider the two different motion models in the execution phase. Thus, we use ALG-SEM to refer to the setting where an algorithm $\text{ALG} \in \{\text{IRIS-U}^2, \text{IRIS}, \text{UP-IRIS}\}$ uses $\mathcal{M}_{\text{simplified}}$ both in the planning and in the execution phase (here, ‘SEM’ refers to Simplified Execution Model). Similarly, we use ALG-EEM to refer to the setting where an algorithm $\text{ALG} \in \{\text{IRIS-U}^2, \text{IRIS}, \text{UP-IRIS}\}$ uses $\mathcal{M}_{\text{simplified}}$ in the planning phase while using $\mathcal{M}_{\text{exact}}$ in the execution phase (here, ‘EEM’ refers to Exact Execution Model).

Fig. 8 contains the coverage, collision probability, path length and search time for the different algorithms.

IRIS achieved an average coverage of 72%, which is significantly lower than the 91% coverage (more than the value of $\kappa = 0.9$) achieved by IRIS-U² even when using only $m = 10$. Additionally, the path computed by IRIS was found to be in collision in 60% of the time as opposed to IRIS-U² whose path was found to be less than 4% collision in all tested execution paths.

In the case of UP-IRIS, despite its shorter calculation time and the fact that the command path has a low collision probability, it achieved an average coverage of only 71%. This highlights the critical significance of considering uncertainty not only in terms of localization uncertainty, which impacts collision probability and path length but also in projecting its effects on the primary objective of inspecting the POIs.

Finally, note that all statistical guarantees stated in Sec. VI empirically hold. We reiterate that the guarantees do not necessarily follow from our analysis in Section VI as we cannot assure that Assumption 1 holds.

2) *Planning and execution with different models:* One may expect that paths computed when planning with a simplified model can be very poor in quality (coverage, collision probability and length) when evaluated with an exact motion model. Despite $\mathcal{M}_{\text{simplified}}$ being an approximation of $\mathcal{M}_{\text{exact}}$, results across all algorithms differ only slightly.

More importantly, as we use a different execution model from the one used when planning, CI bounds are not guaranteed to hold. However, we can see empirically that all results IRIS-U²-EEM falls within the confidence bounds computed for IRIS-U²-SEM.

In addition, in Fig. 8e and 8f we plot the evolution of the location error for one scenario during the mission’s execution stage for the simplified and exact motion models, respectively. We note that for the SEM setting, all algorithms exhibit similar maximal errors. Returning to the improved coverage of our algorithm, this demonstrates our motivating insight that inspection uncertainty should be optimized rather than location

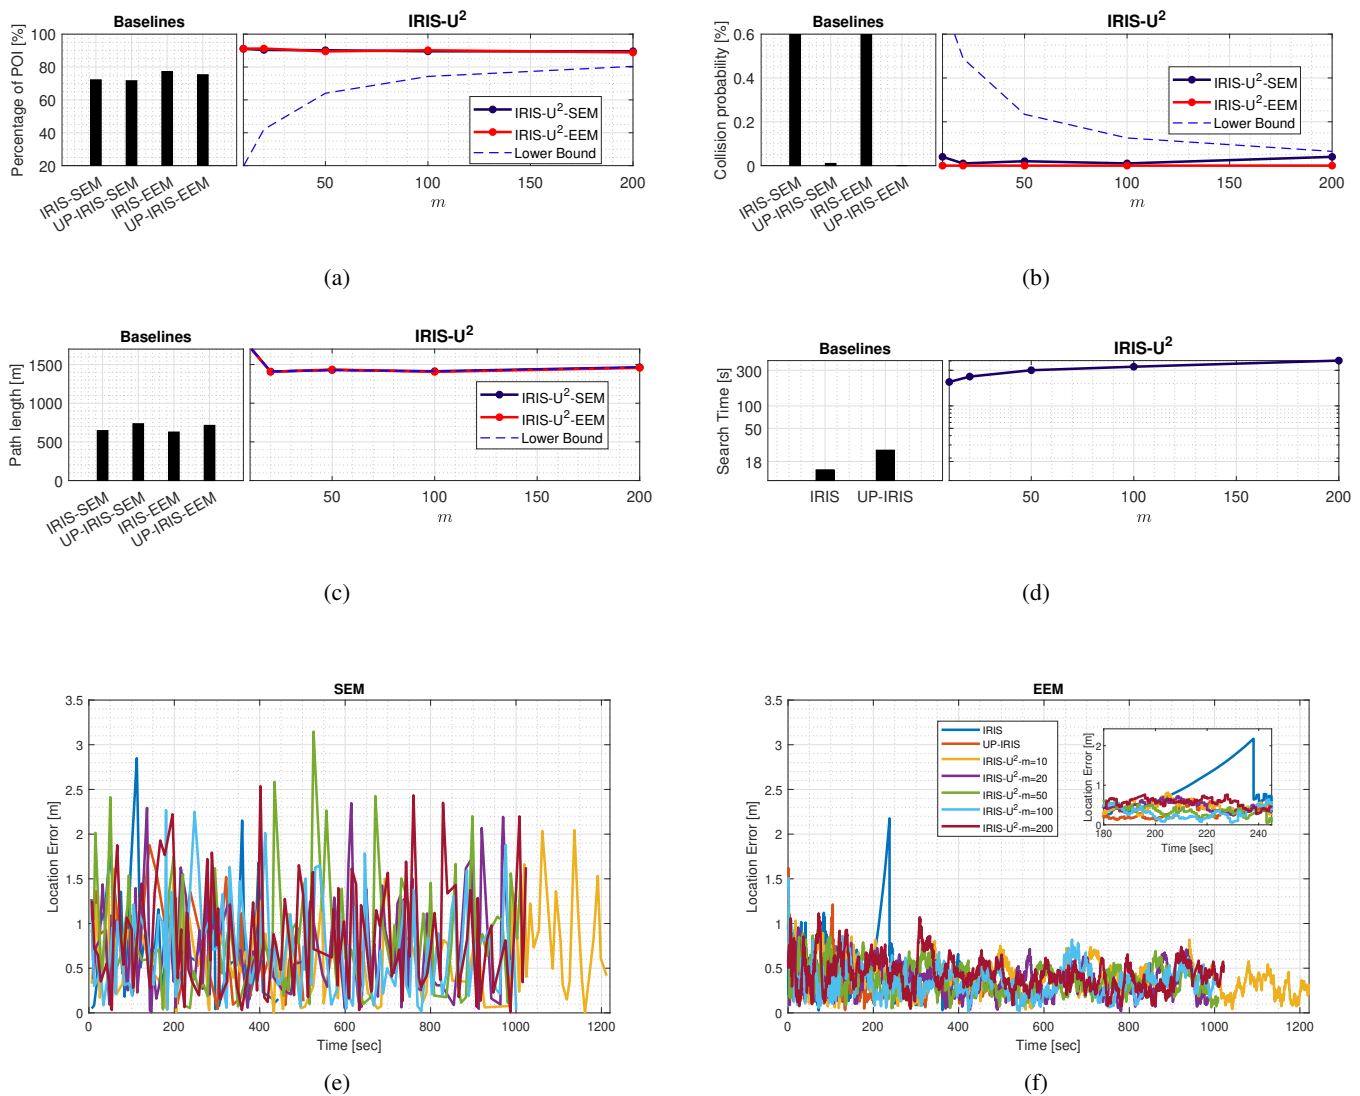


Fig. 8: Experimental results for the bridge scenario. (a)-(d) POI coverage, collision probability, path length, and planning search time as a function of m for IRIS-U² (right) both with the baseline performance of IRIS and UP-IRIS (left). Each of figures (a)-(c) displays the average of 100 MC execution samples (blue dots), corresponding to the CI of the expected performance (red line) as detailed in Lemma. 6.1, 6.2, and 6.3. (e) and (f) show the evolution of the location error for one scenario during the mission's execution stage for the simplified and exact motion model, respectively.

uncertainty. For the EEM model, we can see that (i) IRIS has a peak in its location uncertainty (highlighted in the box inside the figure) as the algorithm plans a path that goes under the bridge, significantly degrading the coverage and that (ii) similar to the SEM setting, even though UP-IRIS does not have a higher location uncertainty when compared to IRIS-U², its coverage is not as good, emphasizing our approach of tracking inspection uncertainty.

Finally, it is important to note that location uncertainty is generally larger in SEM, as the EEM model incorporates a filtering mechanism that smooths GNSS measurement noise, which can vary depending on the dynamics.

3) *Sensitivity analysis of parameters to runtime*: To obtain better execution coverage and fewer collisions, IRIS-U² typically requires longer planning times (see Fig. 8d). However,

as described in Sec. VI-B, for a desired confidence level, we can choose between several parameters. We plot in Fig. 9 the search time for different values of κ and m . Roughly speaking, increasing m (and thus decreasing κ) reduces computation time. This is because higher values of κ do not allow IRIS-U² to subsume nodes and the computational price of maintaining more nodes is typically larger than using more MC samples. However, after a certain number of MC samples is reached, this trend is reversed. This trade-off is dramatic, for example, when considering a desired value of $\hat{p}^- = 0.85$ (i.e., at least 85% POIs will be covered for a CL of $1 - \alpha = 0.95$), the planning times range from 1,912 seconds for $\langle m, \kappa \rangle = \langle 71, 0.934 \rangle$ to 533 seconds for $\langle m, \kappa \rangle = \langle 371, 0.885 \rangle$.

4) *Summary*: We summarize this section with a high-level comparison of the different approaches, also visualized in the




Fig. 9: Search time in the planning phase in the bridge scenario using the simple motion model $\mathcal{M}_{\text{simple}}^{3d}$ as a function of m for several values of \hat{p}_{desired} . Here, for each value of m , κ is the smallest value for which $\hat{p}^-(\kappa, m, \alpha) \geq \hat{p}_{\text{desired}}$ (see Fig. 4).

accompanying [video](#). As IRIS does not account for uncertainty, the executed path tends to miss POIs and may even collide with the bridge. UP-IRIS, on the other hand, prioritizes low-uncertainty regions and hence typically yields a collision-free executed path, albeit it can still miss POIs. Finally, the planned path of IRIS-U² is typically longer, often traversing edges several times to ensure that no POIs are missed. This is done while ensuring that the path is collision free, even in regions with high uncertainty.

IX. CONCLUSION AND FUTURE WORK

In this study, we proposed IRIS-U², an extension of the IRIS offline path-planning algorithm that considers execution uncertainty through the use of MC sampling. Our empirical results demonstrate that IRIS-U² provides better performance under uncertainty in terms of coverage and collision while providing statistical guarantees via confidence intervals. In addition, we provide a guideline on how to choose parameters to reduce the computation time based on the statistical guarantees.

However, as we discussed in Sec. VI-B, the guarantees on the execution path could be affected by a false-negative bias, particularly when there are multiple optional paths in the planning process. As a promising direction, we suggest to explore the Bonferroni correction method [50], and combine it with the information about the number of optional paths considered during planning. By doing so, we can strengthen our CI bounds and enhance the robustness and reliability of our guarantees.

Additionally, recall that we made two key assumptions in this work: (i) that inspecting different POIs is i.i.d (Assumption 1) and (ii) that we wish to minimize path length and not energy consumption or mission completion time (see note following Prob. 1). In future work, we plan to relax both assumptions.

Another direction for future research involves exploring alternative methods to decrease search time. One promising approach is to reduce the number of MC samples by implementing alternative sampling methods such as *Latin hypercube sampling* (LHS) [51], which better distributes samples across the parameter space. The number of samples can be further reduced by leveraging information from the covariance matrix of the EKF (see, e.g., [13]), particularly when navigation sensors uniformly cover the entire uncertainty region.

ACKNOWLEDGEMENTS

This research was supported in part by the Technion Autonomous Systems Program (TASP), the Israeli Ministry of Science & Technology grants No. 3-17385 and in part by the United States-Israel Binational Science Foundation (BSF) grants no. 2019703 and 2021643.

REFERENCES

- [1] A. Bircher, M. Kamel, K. Alexis, M. Burri, P. Oettershagen, S. Omari, T. Mantel, and R. Siegwart, "Three-dimensional coverage path planning via viewpoint resampling and tour optimization for aerial robots," *Autonomous Robots*, vol. 40, no. 6, pp. 1059–1078, 2016.
- [2] B. McGuire, R. Atadero, C. Clevenger, and M. Ozbek, "Bridge information modeling for inspection and evaluation," *Journal of Bridge Engineering*, vol. 21, no. 4, p. 04015076, 2016.
- [3] B. Chan, H. Guan, J. Jo, and M. Blumenstein, "Towards UAV-based bridge inspection systems: A review and an application perspective," *Structural Monitoring and Maintenance*, vol. 2, no. 3, pp. 283–300, 2015.
- [4] M. Fu, A. Kuntz, O. Salzman, and R. Alterovitz, "Toward asymptotically-optimal inspection planning via efficient near-optimal graph search," *Robotics science and systems: online proceedings*, vol. 2019, 2019.
- [5] A. Bircher, K. Alexis, U. Schwesinger, S. Omari, M. Burri, and R. Siegwart, "An incremental sampling-based approach to inspection planning: the rapidly exploring random tree of trees," *Robotica*, vol. 35, no. 6, pp. 1327–1340, 2017.
- [6] Y. Sun, M. Liu, and M. Q.-H. Meng, "Wifi signal strength-based robot indoor localization," in *2014 IEEE International Conference on Information and Automation (ICIA)*. IEEE, 2014, pp. 250–256.
- [7] M. E. Rida, F. Liu, Y. Jari, A. A. A. Algawhari, and A. Askourih, "Indoor location position based on bluetooth signal strength," in *2015 2nd International Conference on Information Science and Control Engineering*. IEEE, 2015, pp. 769–773.
- [8] I. Klein, S. Filin, and T. Toledo, "Vehicle constraints enhancement for supporting ins navigation in urban environments," *NAVIGATION, Journal of the Institute of Navigation*, vol. 58, no. 1, pp. 7–15, 2011.
- [9] L. Janson, E. Schmerling, and M. Pavone, "Monte carlo motion planning for robot trajectory optimization under uncertainty," in *International Symposium of Robotics Research (ISRR)*, vol. 3. Springer, 2015, pp. 343–361.
- [10] N. A. Melchior and R. G. Simmons, "Particle RRT for path planning with uncertainty," in *International Conference on Robotics and Automation (ICRA)*, 2007, pp. 1617–1624.
- [11] J. van den Berg, S. Patil, and R. Alterovitz, "Motion planning under uncertainty using iterative local optimization in belief space," *Int. J. Robotics Res.*, vol. 31, no. 11, pp. 1263–1278, 2012.
- [12] M. Fu, A. Kuntz, O. Salzman, and R. Alterovitz, "Asymptotically optimal inspection planning via efficient near-optimal search on sampled roadmaps," *Int. J. Robotics Res.*, vol. 42, no. 4-5, pp. 150–175, 2023.
- [13] C. Papachristos, M. Kamel, M. Popović, S. Khattak, A. Bircher, H. Oleynikova, T. Dang, F. Mascarić, K. Alexis, and R. Siegwart, "Autonomous exploration and inspection path planning for aerial robots using the robot operating system," in *Robot Operating System (ROS)*. Springer, 2019, pp. 67–111.
- [14] A. Hazra, "Using the confidence interval confidently," *Journal of thoracic disease*, vol. 9, no. 10, p. 4125, 2017.
- [15] R. Peppy and A. Lambert, "Safe path planning in an uncertain-configuration space using rrt," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 5376–5381.
- [16] R. Alami and T. Simeon, "Planning robust motion strategies for a mobile robot," in *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*. IEEE, 1994, pp. 1312–1318.
- [17] S. Candido and S. Hutchinson, "Minimum uncertainty robot path planning using a pomdp approach," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 1408–1413.
- [18] J.-A. Delamer, Y. Watanabe, and C. P. Carvalho Chanel, "Solving path planning problems in urban environments based on a priori sensors availabilities and execution error propagation," in *AIAA Scitech 2019 Forum*, 2019, p. 2202.
- [19] B. Englot, T. Shan, S. D. Bopardikar, and A. Speranzon, "Sampling-based min-max uncertainty path planning," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 6863–6870.

- [20] A. Wu, T. Lew, K. Solovey, E. Schmerling, and M. Pavone, "Robust-RRT: Probabilistically-complete motion planning for uncertain nonlinear systems," in *International Foundation of Robotics Research*, 2022.
- [21] D. Zheng and P. Tsiotras, "Ibtt: Informed batch belief trees for motion planning under uncertainty," *arXiv preprint arXiv:2304.10984*, 2023.
- [22] Q. H. Ho, Z. N. Sunberg, and M. Lahijanian, "Gaussian belief trees for chance constrained asymptotically optimal motion planning," in *International Conference on Robotics and Automation*. IEEE, 2022, pp. 11 029–11 035.
- [23] A. R. Pedram, R. Funada, and T. Tanaka, "Gaussian belief space path planning for minimum sensing navigation," *IEEE Trans. Robotics*, vol. 39, no. 3, pp. 2040–2059, 2023.
- [24] E. Galceran and M. Carreras, "A survey on coverage path planning for robotics," *Robotics and Autonomous systems*, vol. 61, no. 12, pp. 1258–1276, 2013.
- [25] T. Danner and L. E. Kavraki, "Randomized planning for short inspection paths," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 971–976.
- [26] B. Englot and F. Hover, "Inspection planning for sensor coverage of 3d marine structures," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 4412–4417.
- [27] B. J. Englot and F. S. Hover, "Sampling-based coverage path planning for inspection of complex structures," in *Twenty-Second International Conference on Automated Planning and Scheduling*, 2012.
- [28] G. Papadopoulos, H. Kurniawati, and N. M. Patrikalakis, "Asymptotically optimal inspection planning using systems with differential constraints," in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 4126–4133.
- [29] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon path planning for 3d exploration and surface inspection," *Auton. Robots*, vol. 42, no. 2, pp. 291–306, 2018.
- [30] C. Papachristos, F. Mascariich, S. Khattak, T. Dang, and K. Alexis, "Localization uncertainty-aware autonomous exploration and mapping with aerial robots using receding horizon path-planning," *Auton. Robots*, vol. 43, no. 8, pp. 2131–2161, 2019.
- [31] C. S. Tan, R. Mohd-Mokhtar, and M. R. Arshad, "A comprehensive review of coverage path planning in robotics using classical and heuristic algorithms," *IEEE Access*, vol. 9, pp. 119 310–119 342, 2021.
- [32] H. H. Viet, V.-H. Dang, M. N. U. Laskar, and T. Chung, "BA*: an online complete coverage algorithm for cleaning robots," *Applied intelligence*, vol. 39, pp. 217–235, 2013.
- [33] I. A. Hameed, "Coverage path planning software for autonomous robotic lawn mower using dubins' curve," in *2017 IEEE International Conference on Real-time Computing and Robotics (RCAR)*, 2017, pp. 517–522.
- [34] T. Kusnur and M. Likhachev, "Complete, decomposition-free coverage path planning," in *International Conference on Automation Science and Engineering*, 2022, pp. 1431–1437.
- [35] T. Kusnur, S. Mukherjee, D. M. Saxena, T. Fukami, T. Koyama, O. Salzman, and M. Likhachev, "A planning framework for persistent, multi-UAV coverage with global deconfliction," in *Field and Service Robotics*, vol. 16, 2019, pp. 459–474.
- [36] R. Schirmer, P. Biber, and C. Stachniss, "Coverage path planning in belief space," in *International Conference on Robotics and Automation (ICRA)*, 2019, pp. 7604–7610.
- [37] N. Yang, W. Zhang, and W. Yu, "Coverage path planning for autonomous road sweepers in obstacle-cluttered environments," in *Conference on Control Technology and Applications (CCTA)*, 2022, pp. 491–497.
- [38] N. Abreu, N. Cruz, and A. Matos, "Accounting for uncertainty in search operations using auvs," in *Underwater Technology (UT)*, 2017, pp. 1–8.
- [39] E. Galceran, S. Nagappa, M. Carreras, P. Ridao, and A. Palomer, "Uncertainty-driven survey path planning for bathymetric mapping," in *International Conference on Intelligent Robots and Systems*, 2013, pp. 6006–6012.
- [40] P. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems, Second Edition*, 03 2013.
- [41] S. Thrun, W. Burgard, and D. Fox, "Probabilistic robotics," 2005.
- [42] J. N. Gross, Y. Gu, and M. B. Rhudy, "Robust UAV relative navigation with dgps, ins, and peer-to-peer radio ranging," *IEEE Transactions on Automation Science and Engineering*, vol. 12, no. 3, pp. 935–944, 2015.
- [43] M. Khaghani and J. Skaloud, "Autonomous vehicle dynamic model-based navigation for small UAVs," *NAVIGATION: Journal of the Institute of Navigation*, vol. 63, no. 3, pp. 345–358, 2016.
- [44] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006.
- [45] O. Salzman, "Sampling-based robot motion planning," *Commun. ACM*, vol. 62, no. 10, pp. 54–63, 2019.
- [46] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Trans. Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [47] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robotics Research*, vol. 30, no. 7, pp. 846–894, Jun. 2011.
- [48] J. Neyman, "Outline of a theory of statistical estimation based on the classical theory of probability," *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, vol. 236, no. 767, pp. 333–380, 1937.
- [49] D. Habtzghi, C. Midha, and A. Das, "Modified clopper-pearson confidence interval for binomial proportion," *J. Stat. Theory Appl.*, vol. 13, no. 4, pp. 296–310, 2014.
- [50] E. W. Weisstein, "Bonferroni correction," <https://mathworld.wolfram.com/>, 2004.
- [51] W.-L. Loh, "On latin hypercube sampling," *The annals of statistics*, vol. 24, no. 5, pp. 2058–2080, 1996.
- [52] L. Knüsel, "Computation of the chi-square and poisson distribution," *SIAM Journal on Scientific and Statistical Computing*, vol. 7, no. 3, pp. 1022–1036, 1986.
- [53] P. A. Games, "An improved t table for simultaneous control on g contrasts," *Journal of the American Statistical Association*, vol. 72, no. 359, pp. 531–534, 1977.
- [54] M. Grasmair, "Basic properties of convex functions," *Department of Mathematics, Norwegian University of Science and Technology*, 2016.

APPENDIX A STATISTICAL BACKGROUND

In this paper we use CI to evaluate two different types of quantities with respect to the behavior of the IRIS-U² algorithm. The first type is the probability of success for some event, such as the collision probability of a given path being above a certain value. The second type is the mean of a population, such as the expectation of a path's length. Calculating the CI for these two quantities is done differently and we now detail each method:

a) *Probability of success*: To calculate the CI of the success probability of a random variable, we use the Clopper-Pearson method [49]. This method evaluates the maximum likelihood of the probability and its CI assuming a binomial distribution given finite independent trials. Specifically, let X be a random variable whose true unknown probability of success is \bar{p} . Let x_1, \dots, x_m be the outcome of m samples drawn from X (i.e., $x_i \in \{0, 1\}$) and let $\hat{p} := \frac{1}{m} \sum_{i=1}^m x_i$ be the estimated success probability. Then, according to the Clopper-Pearson method for any $\alpha \in [0, 1]$, we can say with CL of $1 - \alpha$ that \bar{p} , the true unknown success probability of X , is within the following the CI:

$$\bar{p} \in [\hat{p}^-(\hat{p}, m, \alpha), \hat{p}^+(\hat{p}, m, \alpha)], \quad (34)$$

where,

$$\hat{p}^-(\hat{p}, m, \alpha) := \frac{\hat{p} \cdot F^{-1}\left(2\hat{p}, 2\lambda_{\hat{p}, m}^-, 1 - \frac{\alpha}{2}\right)}{\lambda_{\hat{p}, m}^- + \hat{p} \cdot F^{-1}\left(2\hat{p}, 2\lambda_{\hat{p}, m}^-, 1 - \frac{\alpha}{2}\right)}, \quad (35a)$$

$$\hat{p}^+(\hat{p}, m, \alpha) := \frac{\hat{p}' \cdot F^{-1}\left(2\hat{p}', 2\lambda_{\hat{p}, m}^+, \frac{\alpha}{2}\right)}{\lambda_{\hat{p}, m}^+ + \hat{p}' \cdot F^{-1}\left(2\hat{p}', 2\lambda_{\hat{p}, m}^+, \frac{\alpha}{2}\right)}. \quad (35b)$$

Here, F^{-1} is the inverse F-distribution function [52], $\lambda_{\hat{p}, m}^- := m - \hat{p} + 1$, $\lambda_{\hat{p}, m}^+ := m - \hat{p}$ and $\hat{p}' := \hat{p} + 1$.

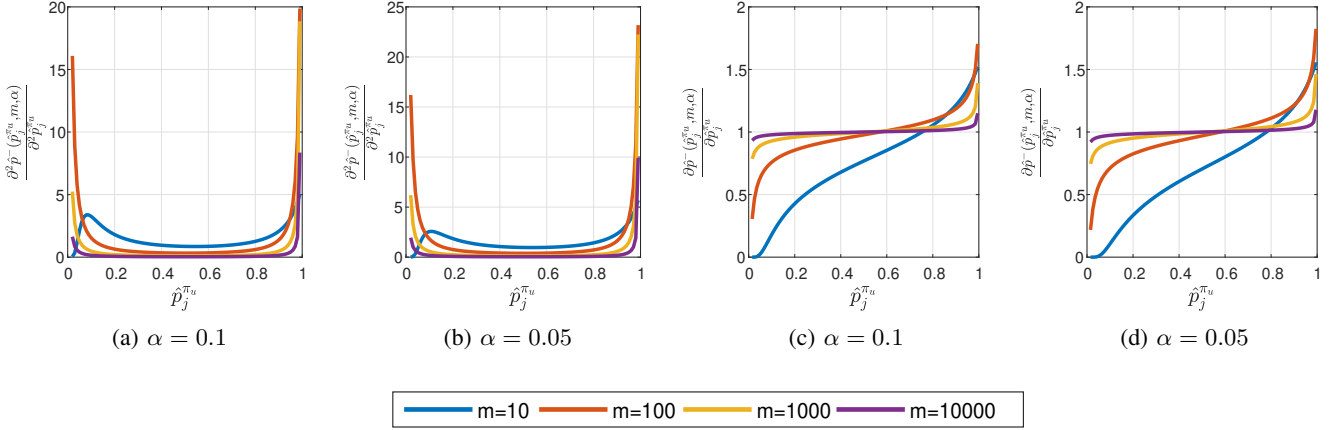


Fig. 10: Numerical demonstration that Eq. (43) (a and b) and Eq. (44) (c and d) hold for different values of α and m .

b) Mean of a population: To calculate the CI of the mean of a population, we follow Habtzghi et al. [49]. Specifically, let X be some random process with unknown mean value \bar{X} . In addition, let x_1, \dots, x_m be m samples drawn from X and let $\hat{X} := \frac{1}{m} \sum_{i=1}^m x_i$ and $\hat{s} := \sqrt{\frac{1}{m-1} \sum_{i=1}^m (x_i - \hat{x})^2}$ be their estimated mean and standard deviation, respectively. Then, for any $\alpha \in [0, 1]$ we can say with CL of $1 - \alpha$ that \bar{X} , the true unknown mean value of X , is within the following CI:

$$\bar{X} \in [\bar{X}^-(\hat{X}, m, \alpha), \bar{X}^+(\hat{X}, m, \alpha)]. \quad (36)$$

Where,

$$\bar{X}^-(\hat{X}, m, \alpha) := \hat{X} - t^* \frac{\hat{s}}{\sqrt{m}}, \quad (37a)$$

$$\bar{X}^+(\hat{X}, m, \alpha) := \hat{X} + t^* \frac{\hat{s}}{\sqrt{m}}. \quad (37b)$$

Here, t^* is computed according to standard t-tables [53].

APPENDIX B PROOFS

We provide proofs for our lemmas.

Proof of Lemma 6.1: We treat the executed path's coverage probability as a random variable and recall that \hat{p}_j^π is the estimated probability to inspect POI j computed by m independent samples of execution paths. Then, for any desired CL of $1 - \alpha$, the lower bound on the coverage is defined as $\hat{p}_j^{\pi-} = \hat{p}^-(\hat{p}_j^\pi, m, \alpha)$. Here, the function \hat{p}^- is defined in Eq. (35). Since we assume that the inspection of each POI is independent of the inspection outcomes of the other POIs (see Assumption. 1), we can add up the lower bounds of the individual POIs to obtain a lower bound on the executed path's coverage. Namely, we can say with a CL of at least $1 - \alpha$, a lower bound on the executed path's coverage is:

$$|\bar{S}(\pi)|^- := \sum_{j=1}^k \hat{p}_j^{\pi-} = \sum_{j=1}^k (\hat{p}_j^\pi, m, \alpha). \quad (38)$$

We now proceed to address the next four lemmas, whose proofs are straightforward. For Lemma 6.2, note that since we

can treat the executed path's collision probability as a random variable, Eq. (30) is an immediate application of the Clopper-Pearson method (see Eq. (35)). Similarly, for Lemma 6.3 we treat the executed path's expected length as random variable and Eq. (31) follows from Eq. (36). We proceed to the final proof.

Proof of Lemma 6.4: Considering that both values m and α are fixed, it would be convenient to define:

$$f(x) := \hat{p}^-(x, m, \alpha). \quad (39)$$

Assume that minimization of $\sum_{j=1}^k f(x_j)$ is achieved for some values x_1^*, \dots, x_k^* where $x_j^* \in [0, 1]$ and $\sum_{j=1}^k x_j^* \geq k \cdot \kappa$. We need to show that $\forall j, x_j^* = \kappa$. Finally, recall that $\kappa \in [0, 1]$. Thus, we distinguish between the cases where $\kappa = 1$ and $\kappa < 1$.

Case 1 ($\kappa = 1$): Here, $\sum_{j=1}^k x_j^* \geq k \cdot \kappa = k$. As $x_j^* \in [0, 1]$, it follows that $\forall j, x_j^* = 1$.

Case 2 ($\kappa < 1$): In this case, the proof will be done in two steps: In step 1, we show that minimization of $\sum_{j=1}^k f(x_j)$ is achieved when all values of x_j^* are equal. Namely, $\exists \kappa' \geq \kappa$ s.t. $\forall j, x_j^* = \kappa'$. In step 2, we show that this value is obtained for $\kappa' = \kappa$. These steps require that Assumption 2 holds. Namely, that $f(x)$ is convex and monotonically increasing.

Case 2, step 1: Let $\kappa' \in [\kappa, 1)$ be a constant such that $\sum_{j=1}^k x_j^* = k \cdot \kappa'$ (note that such κ' always exists). We will prove by contradiction that $\forall j, x_j^* = \kappa'$ which concludes this step.

W.l.o.g., $x_1^* > \kappa$, i.e., there exists $\delta_1 > 0$ such that $x_1^* = \kappa + \delta_1$. As $\sum_{j=1}^k x_j^* = k \cdot \kappa'$, there at least one $j \neq 1$ s.t. $x_j^* < \kappa'$. Additionally, w.l.o.g. $x_2^* < \kappa'$. Namely, there exists $\delta_2 > 0$ such that $x_2^* = \kappa' - \delta_2$.

Let $\delta = \min(\delta_1, \delta_2)$ and consider the following solution x'_1, \dots, x'_k to our minimization problem: $x'_1 = x_1^* - \delta$, $x'_2 = x_2^* + \delta$ and $x'_j = x_j^*$ for $2 < j \leq k$. Notice that this is a valid solution (i.e., $\forall j, x'_j \in [0, 1]$ and $\sum_{j=1}^k x'_j \geq k \cdot \kappa$). We will show that $\sum_{j=1}^k f(x'_j) < \sum_{j=1}^k f(x_j^*)$ which will lead to a contradiction that minimizing $\sum_{j=1}^k f(x_j)$ is achieved for x_1^*, \dots, x_k^* . ■

As $x'_j = x_j^*$ for $2 < j \leq k$, to show that $\sum_{j=1}^{j=k} f(x'_j) < \sum_{j=1}^{j=k} f(x_j^*)$, it suffices to prove that:

$$f(x'_1) + f(x'_2) < f(x_1^*) + f(x_2^*). \quad (40)$$

This will be done using Assumption 2 (i.e., that $f(\cdot)$ is a strictly convex function). Namely, that $\frac{\partial^2 f(x)}{\partial^2 x} > 0, \forall x \in [0, 1]$ which implies (see, e.g., [54]) that $\forall \lambda$ s.t. $0 < \lambda < 1$ we have:

$$f(\lambda x_1 + (1 - \lambda)x_2) < \lambda f(x_1) + (1 - \lambda)f(x_2). \quad (41)$$

Thus,

$$\begin{aligned} f(x'_1) + f(x'_2) &= f(x_1^* - \delta) + f(x_2^* + \delta) \\ &\stackrel{(a)}{=} f(x_1^* - \lambda(x_1^* - x_2^*)) + f(x_2^* + \lambda(x_1^* - x_2^*)) \\ &= f(\lambda x_2^* + (1 - \lambda)x_1^*) + f(\lambda x_1^* + (1 - \lambda)x_2^*) \\ &\stackrel{(b)}{<} (1 - \lambda)f(x_1^*) + \lambda f(x_2^*) + \\ &\quad \lambda f(x_1^*) + (1 - \lambda)f(x_2^*) \\ &= f(x_1^*) + f(x_2^*). \end{aligned} \quad (42)$$

In (a) we use the equation $\lambda = \delta/(x_1^* - x_2^*)$ and note that $0 < \lambda < 1$. (b) follows from Eq. (41).

Case 2, step 2: We will prove this step by contradiction. Assume that minimizing $\sum_{j=1}^{j=k} f(x_j)$ is achieved when $\sum_{j=1}^{j=k} x_j^* > k \cdot \kappa$. Namely, $\exists \delta > 0$ s.t. $\sum_{j=1}^{j=k} x_j^* = k \cdot (\kappa + \delta)$. Following step 1, $\forall j, x_j^* = \kappa + \delta$ and $\sum_{j=1}^{j=k} f(x_j^*) = k \cdot f(\kappa + \delta)$.

However, setting $x_j^* = \kappa$ is also a valid solution for which $\sum_{j=1}^{j=k} f(x_j^*) = k \cdot f(\kappa)$. Following Assumption 2, $f(\cdot)$ is a monotonically increasing function which leads to a contradiction. ■

As a supplement to the proof of Lemma 6.4 we numerically demonstrate that Assumption 2 holds. Specifically, recall that we used $f(x)$ to denote $\hat{p}^-(x, m, \alpha)$. Thus, to show that $f(\cdot)$ monotonically increases and is a strictly convex function we wish to show that $\forall \hat{p}_j^{\pi_u} \in [0, 1]$ both:

$$\frac{\partial \hat{p}^-(\hat{p}_j^{\pi_u}, m, \alpha)}{\partial \hat{p}_j^{\pi_u}} > 0, \quad (43)$$

and

$$\frac{\partial^2 \hat{p}^-(\hat{p}_j^{\pi_u}, m, \alpha)}{\partial^2 \hat{p}_j^{\pi_u}} > 0. \quad (44)$$

This is demonstrated in Fig. 10 for different values of m and α .

APPENDIX C

ILLUSTRATIVE EXAMPLE FOR POSSIBLE FALSE NEGATIVES

In this section, we present an illustrative example that serves to elucidate the notion of false negatives eluded to in Sec. VI. The purpose of this example is to shed light on how the use of statistical guarantees, such as those outlined in Sec. VI-A, can potentially lead to erroneous conclusions when applied to the output of the IRIS-U² algorithm under specific conditions.

Specifically, consider a scenario in which IRIS-U² is configured with $\rho_{\text{coll}} = 0$, meaning that if the algorithm outputs

a path π , it confidently asserts that $\hat{C}(\pi) = 0$, indicating that the path is collision-free.

Furthermore, assume that the algorithm uses $m = 120$. As established in Lemma 6.2, when considering a given path π and setting $\alpha = 0.05$, we find that $\bar{C}(\pi)^+ \approx 0.03$. This implies a 95% probability that if this path is executed 100 times, at most 3 of these executions will result in collisions.

Now, let's examine a situation where the false-negative bias comes into play — when there is more than one option for the command path. For instance, assume we have $k = 100$ paths, denoted as π_1, \dots, π_k , connecting the start and the goal where each of these paths has a true collision probability of $\bar{C}(\pi_i) = 0.04$. The probability that a specific path will be estimated to be collision-free is:

$$(1 - \bar{C}(\pi_i))^m = 0.96^{120} \approx 0.0075.$$

However, the false-negative bias effect lies in the probability that at least one of the k paths will be estimated as collision-free:

$$1 - (1 - (\bar{C}(\pi_i))^m)^k \approx 1 - (1 - 0.0075)^{100} \approx 0.52.$$

Namely, there is more than 50% chance that the algorithm will output a path π_i assumed to be collision-free whose true collision probability is 4% (which, of course, is larger than the upper bound of $\bar{C}(\pi_i)^+ \approx 0.03$ guaranteed with 95% confidence if Lemma 6.2 was wrongly used).