Robotics: Science and Systems 2026
Sydney, Australia, July 13–17, 2026

Provably Guaranteed Polytopic Uncertainty Quantification for SLAM

Guangyang Zeng[1,2], Yulong Gao[3], Yuan Shen[1], Lingpeng Chen[1], Haoying Li[1], Guodong Shi[4], and Junfeng Wu[1,2]
[1]School of Data Science, The Chinese University of Hong Kong, Shenzhen
[2]School of Artificial Intelligence, The Chinese University of Hong Kong, Shenzhen
[3]Department of Electrical and Electronic Engineering, Imperial College London
[4]School of Aerospace, Mechanical and Mechatronic Engineering, The University of Sydney

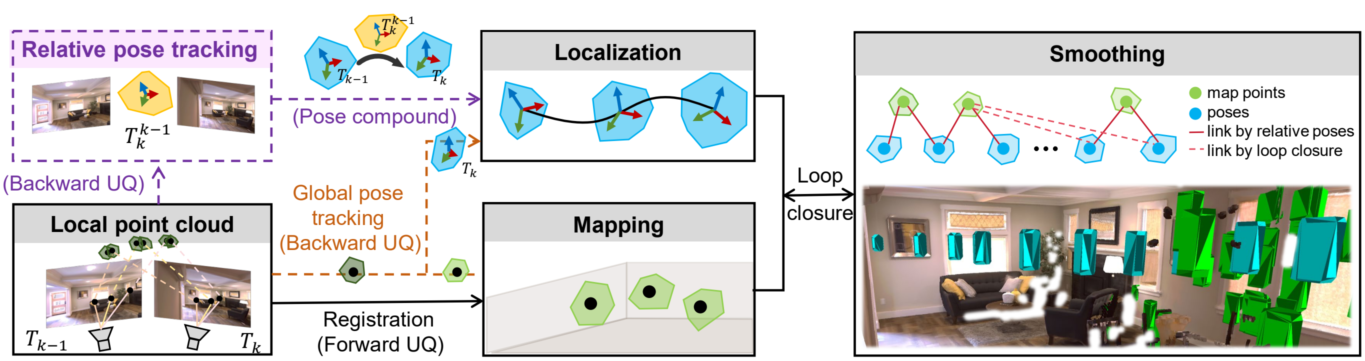

Figure 1: Our guaranteed polytopic SLAM uncertainty quantification (UQ) framework. For global robot localization, there are two alternatives: (i) relative pose tracking followed by pose compound (marked as dashed purple lines), and (ii) global pose tracking (marked as dashed orange lines).

***Abstract*—In safety-critical robotics applications, guaranteed and practical uncertainty quantification (UQ) in perception is vital. Many existing works either offer no formal containment guarantee, rely on restrictive modeling assumptions, or focus only on pose estimation rather than a complete SLAM pipeline. This paper presents provably guaranteed UQ algorithms for 3D-3D landmark-based SLAM. The algorithms consist of three basic UQ modules: forward UQ for mapping, backward UQ for pose tracking, and pose compound. Each module produces a certified uncertainty set; when the input uncertainty bounds are deterministic, the output sets inherit deterministic guarantees, i.e., they provably contain the true poses and landmarks. Specifically, we use polytopes to represent uncertainty sets, enabling tractable computations and a unified treatment of pose uncertainty. To enhance algorithms' practical usability, we incorporate conformal prediction to calibrate measurement uncertainty from data with prescribed probability. Simulations and experiments demonstrate that the proposed algorithms provide both strong theoretical guarantees and practical usability. The code is open-sourced at https://github.com/LIAS-CUHKSZ/Polytopic-SLAM-Uncertainty-Quantification.**

I. Introduction

Simultaneous localization and mapping (SLAM) has found widespread applications in robotics. In recent years, substantial progress has been made, with research primarily improving accuracy [1], robustness [2], and efficiency [3], among other aspects. Most SLAM algorithms calculate optimal point estimates for the robot state and the environment map. However, for *safety-critical* applications such as autonomous driving and unmanned aerial vehicle navigation, point estimates alone are insufficient. Reliable *uncertainty quantification* (UQ) is needed so that downstream planning and control can explicitly account for perception risk and make safety-aware decisions.

Probabilistic UQ typically assumes measurement errors follow a Gaussian distribution. Approaches in this domain often estimate parameters via maximum likelihood and approximate uncertainty covariance using the inverse Fisher information matrix (FIM) or Hessian [4, 5, 6]. However, these methods face several limitations: real-world noise is often non-Gaussian, particularly in dynamic scenes [7, 8], and reliance on local approximations like the FIM can underestimate uncertainty [9]. Furthermore, the resulting covariance ellipsoids offer no formal guarantee of containing the true parameter.

On the other hand, *deterministic UQ*, or *set-membership estimation (SME)*, models uncertainty using hard bounds to identify the set of all parameters consistent with measurements. This approach offers provable guarantees of containing the true parameter without assuming specific error distributions, gaining traction in robotics through interval-based formulations [10, 11, 12]. Recently, conformal prediction (CP) has been introduced to calibrate measurement error bounds statistically [13, 14]. However, computing the resulting feasible parameter set exactly remains difficult due to non-convexity. While tractable approximations exist for pose estimation [15, 16, 7], extending these methods to a complete SLAM pipeline is still an open challenge.

Detailed related work is in Sec.I of suppl. material.

Contributions. In our work, we develop a novel framework that bridges distribution-free calibration with SME to enable tractable, *provably guaranteed* UQ for 3D-3D landmark-based SLAM. Unlike heuristic approaches, our method guarantees that estimated sets contain the true parameter with a prescribed probability. A key novelty is the framework's duality and modularity: it seamlessly transitions between probabilistic confidence sets (via CP) and strict deterministic enclosures by simply interchanging the bounding module.

Technically, as illustrated in Figure 1, we realize this via three primitives: forward UQ (mapping), backward UQ (pose tracking), and pose compound. Each primitive presents distinct challenges, and the associated uncertainties entangle mutually due to the propagation between landmarks and poses. We resolve them with tractable algorithms and then integrate them for full SLAM UQ. In particular, by implementing these primitives using polytopic sets, we leverage mature tools to ensure computational tractability via linear constraints, while enabling a unified representation of rotation and translation uncertainty. Simulations and experiments demonstrate that our algorithms contain the true parameters with certified guarantees, achieve tighter uncertainty sets than recent guaranteed pose UQ methods, and apply to complete SLAM pipelines.

Notations: Let $\mathcal{X}$ and $\mathcal{Y}$ be two sets, $\mathcal{X}+\mathcal{Y}$ and $\mathcal{X}\cdot\mathcal{Y}$ denote their Minkowski sum and Minkowski product, respectively. The closed Euclidean ball centered at point $p_0 \in \mathbb{R}^n$ with radius $r \geq 0$ is denoted as $\mathcal{B}(p_0, r)$. For a set $\mathcal{X}$ in a Euclidean space, its convex hull is denoted by $\mathrm{conv}(\mathcal{X})$. Let $2^{\mathcal{X}}$ be the collection of all subsets of $\mathcal{X}$. A set-valued function $f : \mathcal{X} \to 2^{\mathcal{Y}}$ will be rewritten as $f : \mathcal{X} \rightrightarrows Y$ to emphasize its set-valuedness. The image of $\mathcal{C} \subset X$ under f is defined as $f(\mathcal{C}) = \bigcup_{x\in\mathcal{C}} f(x)$. The preimage of $\mathcal{C}$ is defined as $f^{-1}(\mathcal{C}) = \{x : f(x)\cap\mathcal{C} \neq \emptyset\}$. For two vectors x and y, $\langle x, y\rangle$ represents their inner product, $\mathrm{col}(x, y) = [x^\top, y^\top]^\top$ for short, and $x \leq y$ means the element-wise inequality.

II. Preliminaries

A. Rigid Transformation

In 3D Euclidean space, a rigid transformation includes a translation $t \in \mathbb{R}^3$ and a rotation $R \in \mathrm{SO}(3)$. It preserves the distances between points and the orientation of rigid bodies. The group $\mathrm{SO}(3)$ forms a smooth manifold, and the tangent space at the identity is the Lie algebra, denoted as $\mathfrak{so}(3)$, which is comprised of all 3×3 skew-symmetric matrices. Every 3×3 skew-symmetric matrix can be generated from a vector—say $s \in \mathbb{R}^3$—via the *hat* operator $s^\wedge \in \mathfrak{so}(3)$. The exponential map $\exp : \mathfrak{so}(3) \to \mathrm{SO}(3)$ coincides with standard matrix exponential. We let $\mathrm{Exp} : \mathbb{R}^3 \to \mathrm{SO}(3)$ be the composition of the hat operator followed by exp, that is, $\mathrm{Exp}(s) = \exp(s^\wedge)$. The inverse operation of Exp is denoted as $\mathrm{Log} : \mathrm{SO}(3) \to \{s \in \mathbb{R}^3 \mid \|s\| \in [0, \pi)\}$. The geodesic distance between two rotations is defined as

$$\mathrm{dis}(R_1, R_2) = \cos^{-1}\left(\frac{\mathrm{tr}(R_1 R_2^\top) - 1}{2}\right) \in [0, \pi].$$

B. Conformal Prediction

CP is a lightweight statistical tool for UQ [17, 13]. Let $s_0, s_1, \ldots, s_L$ be $L + 1$ *exchangeable* random variables[1], and $C(s_1, \ldots, s_L) = \mathrm{Quantile}_{1-\delta}(s_1, \ldots, s_L, \infty)$, which is the $(1-\delta)$-quantile over the empirical distribution of $s_1, \ldots, s_L, \infty$. Let $s_{L+1} = \infty$, and without loss of generality, assume $s_1, \ldots, s_L$ are sorted in non-decreasing order. Then, we have

$$C(s_1, \ldots, s_L) = s_p,\ p = \lceil (1-\delta)(L+1)\rceil, \tag{1}$$

where $\lceil\cdot\rceil$ is the ceiling function.

Lemma 1 (**Marginal coverage [17, 13]**). *With the choice of C in* (1)*, it holds that* $\mathrm{Prob}(s_0 \leq C(s_1, \ldots, s_L)) \geq 1 - \delta$.

C. Polytopes

Given $A = [a_1, \ldots, a_M]^\top \in \mathbb{R}^{M\times n}$ and $b = [b_1, \ldots, b_M]^\top \in \mathbb{R}^M$, define the set

$$\mathcal{P}(A, b) := \{x \in \mathbb{R}^n \mid Ax \leq b\}. \tag{2}$$

If $\mathcal{P}(A, b)$ is bounded, we call it a polytope. Next, we introduce some definitions and polytope operations.

Definition 1 (Polytope diameter)**.** *For a polytope $\mathcal{P} \subset \mathbb{R}^n$, its diameter $d(\mathcal{P})$ is defined as* $d(\mathcal{P}) = \max_{p_1,p_2\in\mathcal{P}} \|p_1 - p_2\|$.

Definition 2 (Chebyshev ball and Chebyshev center)**.** *For a polytope $\mathcal{P}$, the Chebyshev ball is the unique minimal enclosing ball. Its center is denoted* cent $(\mathcal{P})$*, and its radius $r(\mathcal{P})$ is given by* $\min_{x_0,r}\{r : \|x_0 - x\| \leq r, \forall x \in \mathcal{P}\}$.

Definition 3. *Given three polytopes $\mathcal{P}, \mathcal{P}_1, \mathcal{P}_2 \subset \mathbb{R}^n$, a linear map $H \in \mathbb{R}^{m\times n}$, and a vector $q \in \mathbb{R}^m$, define*

- *(Affine Map)* $H\mathcal{P} + q := \{Hx + q \mid x \in \mathcal{P}\}$*;*
- *(Minkowski Sum)* $\mathcal{P}_1 + \mathcal{P}_2 := \{x + y \mid x \in \mathcal{P}_1, y \in \mathcal{P}_2\}$*;*
- *(Intersection)* $\mathcal{P}_1 \cap \mathcal{P}_2 := \{x \mid x \in \mathcal{P}_1 \text{ and } x \in \mathcal{P}_2\}$*;*
- *(Set projection to subspace $\mathbb{R}^d$)* $\pi_{\mathbb{R}^d}(\mathcal{P}) = \{x \in \mathbb{R}^d \mid \exists y \in \mathbb{R}^{n-d}, \mathrm{col}(x, y) \in \mathcal{P}\}$.

These operations can be defined for arbitrary sets, however, within the context of polytopes, they are closed, i.e., the result is again a polytope. Thorough preliminaries on the above three aspects are provided in Sec.II of suppl. material.

III. Basic Uncertainty Quantifications

We study three fundamental UQ primitives that form the core building blocks for Sec.IV.

Consider a rigid transformation $T = \begin{bmatrix} R & t \\ 0_3^\top & 1\end{bmatrix}$ of SE(3), which acts on points in $\mathbb{R}^3$ via a left group action ϕ :

[1]The random variables $s_0, \ldots, s_L$ are said to be exchangeable if the joint distribution of $s_0 \ldots, s_L$ is the same to that of $s_{\sigma(0)}, \ldots, s_{\sigma(L)}$ for any permutation σ on indices.

$\mathrm{SE}(3) \times \mathbb{R}^3 \to \mathbb{R}^3$. We denote $x(T) := \mathrm{col}(\mathrm{vec}(R), t) \in \mathbb{R}^{12}$ as the vectorization of T. In robotics, if T maps a local frame to a reference frame and p is a point in the local frame, then $\phi(T, p) = Rp + t$ gives the point in the reference frame. Define the map $\phi_p : \mathrm{SE}(3) \to \mathbb{R}^3, \quad T \mapsto \phi(T, p)$. The basic problems to be discussed involve the forward and backward transfer of uncertainty between the argument of ϕ_p (elements of $\mathrm{SE}(3)$) and its image in $\mathbb{R}^3$. The forward UQ aims to propagate uncertainty from $\mathrm{SE}(3)$ to $\mathbb{R}^3$, whereas the backward UQ infers the uncertainty on $\mathrm{SE}(3)$ from that on $\mathbb{R}^3$. We will also investigate the fundamental UQ problem for pose compound. Before proceeding, we introduce the following definition. For a given set $\mathcal{S}$ of points, define a set-valued function $\phi_{\mathcal{S}} : \mathrm{SE}(3) \rightrightarrows \mathbb{R}^3$ by

$$\phi_{\mathcal{S}}(T) = \{\phi_p(T) : p \in \mathcal{S}\}. \tag{3}$$

A. Forward Uncertainty Quantification

The forward UQ problem is interpreted as follows: given a fixed, known and bounded set $\mathcal{S}$, for a set $\mathcal{T} \subset \mathrm{SE}(3)$, find the image $\phi_{\mathcal{S}}(\mathcal{T})$. For tractability and computational issues, we restrict $\mathcal{S}$ to a polytope $\mathcal{P}(A_1, b_1) = \{p \in \mathbb{R}^3 \mid A_1 p \leq b_1\}$ and focus on the rigid transformation uncertainty form as $\mathcal{T} = \{T \in \mathrm{SE}(3) \mid Hx(T) \leq d\}$. In the SLAM context in Sec.IV, p denotes a landmark's local coordinates, with uncertainty $\mathcal{P}(A_1, b_1)$ provided by CP, and the set $\mathcal{T}$ for a robot pose can be obtained by backward UQ in each iteration.

The forward UQ derived from $\mathcal{T}$ and $\mathcal{P}(A_1, b_1)$ is to compute the set

$$\mathcal{Q} := \{q = \phi(T, p) \mid T \in \mathcal{T}, p \in \mathcal{P}(A_1, b_1)\}. \tag{4}$$

However, due to the bilinear property of ϕ and non-convex $\mathrm{SO}(3)$ constraint, the set $\mathcal{Q}$ is hard to characterize. In what follows, we instead compute a tight and computationally efficient polytopic approximation $\mathcal{P}(A_2, b_2)$ of $\mathcal{Q}$.

Fo computational efficiency, we fix the matrix $A_2 = [a_1, \ldots, a_M]^\top$ *a prior*, where each a_m is a *unit normal*. One way to select A_2 is from the regular 3D polytope with M sides. Now the problem of computing $\mathcal{P}(A_2, b_2)$ boils down to computing the smallest b_2 giving the tightest approximation. We need to solve M optimization problems in form of

$$\begin{aligned} [b_2]_m = \max_{q,T,p} \quad & a_m^\top q && \text{(5a)}\\ \text{s.t.} \quad & Hx(T) \leq d, \ A_1 p \leq b_1, && \text{(5b)}\\ & q = \phi(T, p), \ T \in \mathrm{SE}(3). && \text{(5c)} \end{aligned}$$

This problem can be simplified by only considering the vertices of $\mathcal{P}(A_1, b_1)$ for p, as shown in the following lemma.

Lemma 2. *Let $\mathcal{V} = \{v_i\}_{i=1}^N$ be the vertex set of $\mathcal{P}(A_1, b_1)$, where N is the vertex number. Problem* (5) *is equivalent to*

$$\begin{aligned} [b_2]_m = \max_{v_i \in \mathcal{V}} \max_{q,T} \quad & a_m^\top q && \text{(6a)}\\ \text{s.t.} \quad & Hx(T) \leq d, \ q = \phi(T, v_i), && \text{(6b)}\\ & T \in \mathrm{SE}(3). && \text{(6c)} \end{aligned}$$

Proofs of all lemmas, theorems, and propositions in the main manuscript are provided in Sec.IV in suppl. material.

In Lemma 2, we enumerate all vertices $v_i \in \mathcal{V}$ and solve the inner problem in (6) at each iteration. The maximum objective value over $\mathcal{V}$ is selected as $[b_2]_m$. However, the inner problem is still non-convex due to the $\mathrm{SE}(3)$ constraint. We can relax it to an SDP and obtain a certifiable upper bound. Let $y = \mathrm{col}(x(T), 1)$ and $X = yy^\top$. Then, the relaxed SDP is

$$\begin{aligned} [\hat{b}_2]_m = \max_{v_i \in \mathcal{V}} \max_{X \in \mathbb{S}^{13}} \quad & \mathrm{tr}(Q_{m,0} X) && \text{(7a)}\\ \text{s.t.} \quad & X \succeq 0, \ [X]_{13,13} = 1, && \text{(7b)}\\ & \mathrm{tr}(Q_i X) = 0, i = 1, \ldots, 6, && \text{(7c)}\\ & \mathrm{tr}(F_i X) \leq 0, i = 1, \ldots, d_T, && \text{(7d)} \end{aligned}$$

where d_T is the dimension of d. The data matrices $Q_{m,0}, Q_i, F_i$ are given in Sec.III of suppl. material. The time complexity to solve (7) using interior-point algorithms is $O(n^{3.5} N d_T + n^{2.5} N d_T^2 + n^{0.5} N d_T^3)$ [18] ($n = 13$ in our problem). Since the data matrices are sparse and the constraint numbers in polytopic description is typically moderate in our problem, the SDP can be efficiently solved by off-the-shelf solvers such as `MOSEK`. To tighten the SDP relaxation, the automatic method in [19] can also be adopted, which augments the formulation with automatically generated redundant constraints.

Finally, we conclude with the following *guaranteed* result of the whole procedure. The forward propagation algorithm is summarized in Algorithm 1 in suppl. material.

Theorem 1. *The set $\mathcal{P}(A_2, \hat{b}_2)$ is a guaranteed approximation for $\mathcal{Q}$ in* (4)*, that is, $\mathcal{Q} \subseteq \mathcal{P}(A_2, \hat{b}_2)$.*

B. Backward Uncertainty Quantification

The backward UQ problem can be interpreted as follows: given a fixed bounded set $\mathcal{S}$, for a target subset $\mathcal{Q} \subset \mathbb{R}^3$ find the preimage $\phi_{\mathcal{S}}^{-1}(\mathcal{Q})$ in the manifold of $\mathrm{SE}(3)$. The following theorem provides an upper bound on this preimage by relating it to the preimage of a single-valued map ϕ_s for some $s \in \mathcal{S}$.

Theorem 2. *For any $s \in \mathcal{S}$, $\phi_{\mathcal{S}}^{-1}(\mathcal{Q}) \subseteq \phi_s^{-1}(\mathcal{Q}^+)$ for all $\mathcal{Q}^+ \supseteq \mathcal{Q} + \mathcal{B}(0, d(\mathcal{S}))$. In particular, when $s = \mathrm{cent}(\mathcal{S})$, $\phi_{\mathcal{S}}^{-1}(\mathcal{Q}) \subseteq \phi_s^{-1}(\mathcal{Q}^+)$ for all $\mathcal{Q}^+ \supseteq \mathcal{Q} + \mathcal{B}(0, r(\mathcal{S}))$.*

When $\mathcal{S}$ and $\mathcal{Q}$ are arbitrary sets, the preimage is generally difficult to characterize. Although Theorem 2 provides an upper bound on this preimage, obtaining a closed-form expression remains impractical. For further analysis, we restrict $\mathcal{Q}$ to be a polytope. A useful lemma is

Lemma 3. *Give a polytope $\mathcal{P}(A, b)$, for any real number $\epsilon > 0$, it holds that $\mathcal{P}(A, b) + \mathcal{B}(0, \epsilon) \subset \mathcal{P}(A, b + \epsilon\mathbf{1})$.*

An immediate benefit of restricting $\mathcal{Q}$ to a polytope represented in H-form as $\mathcal{P}(A_2, b_2)$ is that the set $\mathcal{Q} + \mathcal{B}(0, r(\mathcal{S}))$ can be bounded by a polytope $\mathcal{P}(A_2, b_2 + r(\mathcal{S})\mathbf{1})$ as established in Lemma 3 by considering the Chebyshev ball of $\mathcal{S}$. Moreover, by Theorem 2, fixing $s = \mathrm{cent}(\mathcal{S})$ renders ϕ_s linear, ensuring that the preimage of a polytope under ϕ_s remains a polytope. In addition, if $\mathcal{S}$ itself is constrained to be a polytope $\mathcal{P}(A_1, b_1)$, locating its Chebyshev ball is computationally efficient. In the SLAM context in Sec.IV, $\mathcal{P}(A_1, b_1)$ for local point p can be provided by CP, and $\mathcal{P}(A_2, b_2)$ for global map coordinate q can be obtained by previous forward UQ rounds. The backward UQ from $\mathcal{P}(A_1, b_1)$ and $\mathcal{P}(A_2, b_2)$ is to compute the set

$$\mathcal{T} := \{T \in \mathrm{SE}(3) \mid \exists p \in \mathcal{P}(A_1, b_1), q \in \mathcal{P}(A_2, b_2) \text{ such that } q = \phi(T, p)\}. \tag{8}$$

From Theorem 2 and Lemma 3, by setting s =cent $(P(A_1, b_1))$, for any $T \in \phi_s^{-1}(\mathcal{P}(A_2, b_2 + r(\mathcal{P}(A_1, b_1)\mathbf{1}))$, we have $A_2\phi(T, s) \leq b_2 + r(\mathcal{P}(A_1, b_1)\mathbf{1}$. It follows that

$$\mathcal{T}_r = \{T \in \mathrm{SE}(3) \mid Hx(T) \leq d\}, \tag{9}$$

where $H = [A_2(s^\top \otimes I_3), A_2]$ and $d = b_2 + r(\mathcal{P}(A_1, b_1)\mathbf{1}$.

Theorem 2 also allows for a looser bound based on the polytope diameter, which is computationally simpler as it bypasses the computation of the Chebyshev center. Choosing arbitrary $s \in \mathcal{P}(A_1, b_1)$, and the preimage $\phi_s^{-1}(\mathcal{P}(A_2, b_2 + d(\mathcal{P}(A_1, b_1)\mathbf{1}))$ is given by

$$\mathcal{T}_d = \{T \in \mathrm{SE}(3) \mid Hx(T) \leq d\}, \tag{10}$$

where $H = [A_2(s^\top \otimes I_3), A_2]$ and $d = b_2 + d(\mathcal{P}(A_1, b_1)\mathbf{1}$.

In summary, we conclude with the following result on the *guarantee* of the whole procedure.

Theorem 3. *Both $\mathcal{T}_r$ in (9) and $\mathcal{T}_d$ in (10) are guaranteed approximations for $\mathcal{T}$ in (8), that is, $\mathcal{T} \subset \mathcal{T}_r, \mathcal{T}_d$. Moreover, $\mathcal{T}_r$ is the unique minimum set that*

1) *can be written as $\mathrm{SE}(3) \cap \mathcal{P}(H, \cdot)$ with template $H = [A_2(s^\top \otimes I_3), A_2], s = \mathrm{cent}\ (\mathcal{P}(A_1, b_1))$;*
2) *encloses $\mathcal{T}$;*
3) *is enclosed by any other set satisfying 1) and 2).*

Our backward UQ is reported in Algorithm 2 in suppl. material.

C. Rigid Transformation Uncertainty Compound

The rigid transformation uncertainty compound problem is formulated as follows: find a set $\mathcal{T}_3$, if there exists one, of $\mathrm{SE}(3)$ for uncertainty sets $\mathcal{T}_1, \mathcal{T}_2 \subset \mathrm{SE}(3)$ such that

$$\phi_{\mathcal{S}}(\mathcal{T}_3) = \phi_{\phi_{\mathcal{S}}(\mathcal{T}_2)}(\mathcal{T}_1) \tag{11}$$

holds for any bounded point set $\mathcal{S}$. The problem is well-defined only if $\mathcal{T}_3$ exists and is unique. The composition of the group action ϕ on $\mathbb{R}^3$ is compatible with the group multiplication in $\mathrm{SE}(3)$. Specifically, for $T_1, T_2 \in \mathrm{SE}(3)$, $\phi_{\phi_p(T_2)}(T_1) = \phi_p(T_1T_2)$. Consequently, the compound of these uncertainties can be fully characterized using the Minkowski product, establishing that the problem is well-defined.

Lemma 4. *For two sets $\mathcal{T}_1, \mathcal{T}_2 \subset \mathrm{SE}(3)$, $\mathcal{T}_3 \subset \mathrm{SE}(3)$ satisfies (11) for any bounded set $\mathcal{S}$ if and only if $\mathcal{T}_3 = \mathcal{T}_1 \cdot \mathcal{T}_2$.*

By Lemma 4, the problem of compounding uncertainties in rigid transformations can be conveniently reformulated as a set product problem. To facilitate further analysis, let us restrict our attention to the polytopic uncertainty sets $\mathcal{T}_1 = \{T \in \mathrm{SE}(3) \mid A_1x(T) \leq b_1\}$ on T_1 and $\mathcal{T}_2 = \{T \in \mathrm{SE}(3) \mid A_2x(T) \leq b_2\}$ on T_2. Assume that $\mathcal{P}(A_1, b_1)$ and $\mathcal{P}(A_2, b_2)$ are given. The polytopic structure is not generally preserved under the group product. We propose two methods—one direct and one indirect—to compute enclosures for $\mathcal{T}_1 \cdot \mathcal{T}_2$.

1) Direct Computation Method: In this part, we are interested in developing an algorithm to compute a tight polytopic set $\mathcal{P}(A_3, b_3) = \{x \in \mathbb{R}^{12} \mid A_3x \leq b_3\}$ on $x(T_3)$. Similar to the forward UQ in Sec.III-A, we fix the template $A_3 = [a_1, \ldots, a_M]^\top$ *a prior* to improve the computational efficiency. Here each a_m is a *unit normal*. Then, the problem to be solved is

$$\begin{aligned}
[b_3]_m = \max_{T_1, T_2, T_3} \quad & a_m^\top x(T_3) && \text{(12a)}\\
\text{s.t.} \quad & T_3 = T_1T_2,\ A_1x(T_1) \leq b_1, && \text{(12b)}\\
& A_2x(T_2) \leq b_2,\ T_1, T_2 \in \mathrm{SE}(3). && \text{(12c)}
\end{aligned}$$

This bilinear optimization problem is hard to be globally solved. Let $y = \mathrm{col}(x(T_1), x(T_2), 1)$ and $X = yy^\top$. We have the following relaxed SDP giving a certifiable upper bound.

$$\begin{aligned}
[\hat{b}_3]_m = \max_{X \in \mathbb{S}^{25}} \quad & \mathrm{tr}(Q_{m,0}X) && \text{(13a)}\\
\text{s.t.} \quad & X \succeq 0,\ [X]_{25,25} = 1, && \text{(13b)}\\
& \mathrm{tr}(Q_iX) = 0, i = 1, \ldots, 12, && \text{(13c)}\\
& \mathrm{tr}(F_iX) \leq 0, i = 1, \ldots, b_{1,T}, && \text{(13d)}\\
& \mathrm{tr}(G_iX) \leq 0, i = 1, \ldots, b_{2,T}, && \text{(13e)}
\end{aligned}$$

where $b_{1,T}$ and $b_{2,T}$ are the dimensions of b_1 and b_2, respectively, and the data matrices $Q_{m,0}, Q_i, F_i, G_i$ are given in Sec.III of suppl. material. The time complexity to solve the SDP using interior-point algorithms is $O(n^{3.5}b_{1,T} + n^{2.5}b_{1,T}^2 + n^{0.5}b_{1,T}^3 + n^{3.5}b_{2,T} + n^{2.5}b_{2,T}^2 + n^{0.5}b_{2,T}^3)$ [18] ($n = 25$ in this problem). Similarly, this SDP can be tightened using the method in [19] and solved by off-the-shelf solvers like `MOSEK`.

Theorem 4. *The set $\{T_3 \in \mathrm{SE}(3) \mid A_3x(T_3) \leq \hat{b}_3\}$ is a guaranteed approximation for the original set $\mathcal{T}_3$.*

Our direct pose uncertainty compound algorithm is summarized in Algorithm 3 in suppl. material.

2) Indirect Computation Method: Since $x(T_3) \in \mathbb{R}^{12}$, its polytopic template A_3 can be complex (i.e., M is large), and solving M SDPs (13) may become prohibitive. We instead propose an efficient indirect method. This approach decomposes the polytope $\mathcal{P}(A, b)$ via set projection into translation $\mathcal{P}(A_t, b_t)$ and rotation $\mathcal{P}(A_r, b_r)$ (where $r = \mathrm{vec}(R)$), and then further relax the rotation uncertainty.

We begin with over-approximating $\mathcal{P}(A_r, b_r)$ using the polytopic set $\mathcal{X} := \{\mathrm{Exp}(\theta w) \mid w \in \mathbb{S}^2, \theta \in [0, \theta']\}$ (defined by one-dimensional uncertainty θ). The following proposition characterizes this set.

Proposition 1. *Consider the polytope $\mathcal{P}(A_r, b_r)$ where $A_r = [a_1, \ldots, a_M]^\top$ and $b_r = [b_1, \ldots, b_M]^\top$. Then, the set $\bar{R}\mathcal{X}$ is a guaranteed approximation for the set $\{R \in \mathrm{SO}(3) \mid$

$\mathrm{vec}(R) \in \mathcal{P}(A_r, b_r)\}$, *where* $\mathcal{X} := \{\mathrm{Exp}(\theta w) \mid w \in \mathbb{S}^2, \theta \in [0, \theta']\}$. *Here,* $\bar{R} \in \mathrm{SO}(3)$ *is obtained from* $\bar{r}$ *via inverse vectorization and singular value decomposition-based* $\mathrm{SO}(3)$ *projection [20] and* $\theta' = \cos^{-1}\left(\max\{\frac{c^*-1}{2}, -1\}\right)$. *In particular,* $\bar{r}$ *is the Chebyshev center of* $\mathcal{P}(A_r, b_r)$*, and* c^* *is obtained by solving the problem*

$$
\begin{aligned}
c^* = \max_{X \in \mathbb{S}^{10}} \quad & \mathrm{tr}(Q_0 X) && \text{(14a)} \\
\text{s.t.} \quad & X \succeq 0, \; [X]_{10,10} = 1, && \text{(14b)} \\
& \mathrm{tr}(Q_i X) = 0, i = 1, \ldots, 6, && \text{(14c)} \\
& \mathrm{tr}(F_i X) \leq 0, i = 1, \ldots, M. && \text{(14d)}
\end{aligned}
$$

The data matrices Q_0, Q_i, F_i *are provided in Sec.III in suppl. material.*

Now let us consider the pose compound $T_3 = T_1 T_2$, as defined before. Suppose that the uncertainties on the pose T_i are characterized by the set $\bar{R}_i \mathcal{X}_i$ on the rotation R_i and the set $\mathcal{P}(A_{t_i}, b_{t_i})$ on the translation t_i, where $\mathcal{X}_i = \{\mathrm{Exp}(\theta_i w) \mid w \in \mathbb{S}^2, \theta_i \in [0, \theta_i']\}$. According to the pose compound operation, we have that $R_3 = R_1 R_2$ and $t_3 = R_1 t_2 + t_1$. From a set-valued perspective, we have

$$R_3 \in (\bar{R}_1 \mathcal{X}_1 \bar{R}_2) \cdot \mathcal{X}_2, \; t_3 \in \bar{R}_1 \mathcal{X}_1 \cdot \mathcal{P}(A_{t_2}, b_{t_2}) + \mathcal{P}(A_{t_1}, b_{t_1}).$$

Next we develop an efficient method to compute tight uncertainty sets $\bar{R}_3 \mathcal{X}_3 = \{\bar{R}_3 \mathrm{Exp}(\theta_3 w) \mid w \in \mathbb{S}^2, \theta_3 \in [0, \theta_3']\}$ on R_3 and $\mathcal{P}(A_{t_3}, b_{t_3}) = \{t \in \mathbb{R}^3 \mid A_{t_3} t \leq b_{t_3}\}$ on t_3. The following lemma gives the exact characterization of $\bar{R}_3 \mathcal{X}_3$.

Lemma 5. *Let* $\theta_3' = \theta_1' + \theta_2'$ *and* $\bar{R}_3 = \bar{R}_1 \bar{R}_2$. *Then, the uncertainty set* $\bar{R}_3 \mathcal{X}_3$ *is equivalent to* $(\bar{R}_1 \mathcal{X}_1 \bar{R}_2) \cdot \mathcal{X}_2$*, i.e.,* $\bar{R}_3 \mathcal{X}_3 = (\bar{R}_1 \mathcal{X}_1 \bar{R}_2) \cdot \mathcal{X}_2$.

The following lemma gives a polytopic overapproximation of the uncertainty set $\mathcal{P}(A_{t_3}, b_{t_3})$ over t_3.

Lemma 6. *Given a matrix* $A = [a_1, \ldots, a_M]^\top$ *where each* a_m *is a normal vector, the set* $\mathcal{P}(A_{t_3}, b_{t_3}) = \mathcal{P}(A \bar{R}_1^\top, b) + \mathcal{P}(A_{t_1}, b_{t_1})$ *is a guaranteed approximation for the uncertainty set on* t_3*, i.e.,* $\bar{R}_1 \mathcal{X}_1 \cdot \mathcal{P}(A_{t_2}, b_{t_2}) + \mathcal{P}(A_{t_1}, b_{t_1})$. *Here,*

$$[b]_m = \max_{v \in \mathcal{V}_2} \left\{ \max_{\theta_1 \in [0, \theta_1'], w \in \mathbb{S}^2} \langle \mathrm{Exp}(\theta_1 w) a_m, v \rangle \right\}, \tag{15}$$

where $\mathcal{V}_2$ *is the vertex set of* $\mathcal{P}(A_{t_2}, b_{t_2})$.

As illustrated in Figure 2, the inner maximization in (15) has a closed-form solution:

$$
\begin{aligned}
&\max_{\theta_1 \in [0, \theta_1'], w \in \mathbb{S}^2} \langle \mathrm{Exp}(\theta_1 w) a_m, v \rangle \\
&= \begin{cases} \|v\|, \text{ if } \angle(a_m, v) \leq \theta_1', \\ \|v\| \cos(\angle(a_m, v) - \theta_1'), \text{ if } \angle(a_m, v) > \theta_1', \end{cases}
\end{aligned} \tag{16}
$$

where $\angle(a_m, v)$ denotes the angle between a_m and v. This enables efficient set computation.

The above two lemmas ensures a *guaranteed* result for the uncertainty set on T_3 via pose compound.

Theorem 5. *The set* $\{T_3 \in \mathrm{SE}(3) \mid R_3 \in \bar{R}_3 \mathcal{X}_3, t_3 \in \mathcal{P}(A_{t_3}, b_{t_3})\}$ *is a guaranteed approximation for the set* $\mathcal{T}_3$.

Our indirect pose uncertainty compound algorithm is summarized in Algorithm 4 in suppl. material.

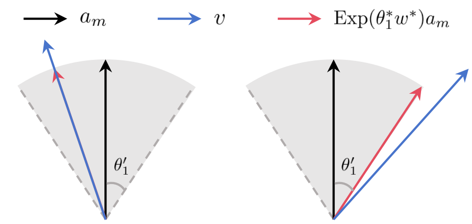

Figure 2: Illustration of Eq. (16). When $\angle(a_m, v) \leq \theta_1'$, the optimal rotation $\mathrm{Exp}(\theta_1^* w^*)$ aligns a_m to the direction of v; when $\angle(a_m, v) > \theta_1'$, the optimal rotation $\mathrm{Exp}(\theta_1^* w^*)$ rotates a_m along the optimal axis w^* by $\theta_1^* = \theta_1'$ such that the rotated vector has an angle of $\angle(a_m, v) - \theta_1'$ relative to v.

IV. Uncertainty Quantification for SLAM

In the section, we construct a guaranteed SLAM UQ pipeline based on the basic UQ primitives in Sec.III. Here we focus on 3D-3D landmark-based SLAM (joint pose-landmark estimation) [21, 22] as a representative formulation.

A. Algorithm Overview

An overview of our algorithm is shown in Figure 1. The robot is assumed to possess metric-scale 3D perception capabilities, enabled by sensors such as stereo cameras, RGB-D cameras, or LiDARs, for local point cloud generation. The proposed algorithm consists of the following key components:

1) *Local point cloud generation*: extracting point clouds in the local frame from the raw sensor data.
2) *Localization*: estimating the global poses of the robot.
3) *Mapping*: estimating the global landmark positions.
4) *Smoothing*: leveraging loop closures to correct the estimated robot's trajectory and map landmarks.

As illustrated in Figure 1, the localization module can be implemented as a **relative framework**, which estimates relative poses between consecutive frames and compound it with the previous global localization, or a **global framework**, which executes direct pose tracking by establishing correspondences between the local point cloud and the global map.

Next, we will utilize CP and the basic UQ in Sec.III to illustrate how to quantify uncertainty in each component.

B. Uncertainty Quantification for Local Point Cloud

We employ CP to provide polytopic uncertainty sets for local point clouds. When a new frame arrives, the robot generates a local point cloud using onboard sensors. Let p_i be the coordinates of the i-th local 3D point. The objective is to characterize its uncertainty by constructing a polytopic uncertainty set $\mathcal{P}(A_i, b_i) = \{p_i \mid A_i p_i \leq b_i\}$. The derivation of ellipsoidal uncertainty sets for points directly measured by LiDARs or RGB-D cameras using CP has been studied in [7], which can be easily adapted for polytopic uncertainties. In Sec.V of the suppl. material, we present a CP-based approach for uncertainty modeling of 3D points reconstructed from stereo image measurements.

C. Localization

1) Relative pose tracking followed by pose compound: When a new frame arrives, we perform relative pose tracking, i.e., estimating the relative pose between the k-th frame and the previous frame, denoted by T_k^{k-1}. After feature matching, we have point correspondences $\{p_{k,i}, q_{k,i}\}_{i=1}^{n_k}$, where $p_{k,i}$ and $q_{k,i}$ are coordinates in the k-th frame and the previous frame, respectively. They have the relationship $q_{k,i} = \phi(T_k^{k-1}, p_{k,i})$. The uncertainty sets $\mathcal{P}(A_{p_{k,i}}, b_{p_{k,i}})$ for $p_{k,i}$ and $\mathcal{P}(A_{q_{k,i}}, b_{q_{k,i}})$ for $q_{k,i}$ are provided by CP. Applying the backward UQ in Sec.III-B gives $\mathcal{T}_k^{k-1} = \{T_k^{k-1} \in \mathrm{SE}(3) \mid H_k^{k-1}\mathrm{vec}(T_k^{k-1}) \leq d_k^{k-1}\}$. Note that $T_k = T_{k-1}T_k^{k-1}$. The uncertainty set $\mathcal{T}_{k-1} = \{T_{k-1} \in \mathrm{SE}(3) \mid H_{k-1}\mathrm{vec}(T_{k-1}) \leq d_{k-1}\}$ for T_{k-1} is stored by the previous frame localization. Quantifying the uncertainty of T_k requires UQ through pose compound. We can apply either direct or indirect rigid transformation uncertainty compound in Sec.III-C to compute the set $\mathcal{T}_k = \{T_k \in \mathrm{SE}(3) \mid H_k\mathrm{vec}(T_k) \leq d_k\}$ for T_k.

2) Global pose tracking: As an alternative, we can directly estimate the global pose by tracking points in the global map. After feature matching between consecutive frames, we can obtain point correspondences $\{p_{k,i}, q_{k,i}\}_{i=1}^{n_k}$, where $p_{k,i}$ and $q_{k,i}$ are coordinates in the k-th frame and the global frame, respectively. They have the relationship $q_{k,i} = \phi(T_k, p_{k,i})$. The uncertainty sets $\mathcal{P}(A_{p_{k,i}}, b_{p_{k,i}})$ for $p_{k,i}$ and $\mathcal{P}(A_{q_{k,i}}, b_{q_{k,i}})$ for $q_{k,i}$ are provided by CP and the mapping function, respectively. We can utilize our backward UQ in Sec.III-B for global pose estimation and compute $\mathcal{T}_k = \{T_k \in \mathrm{SE}(3) \mid H_k\mathrm{vec}(T_k) \leq d_k\}$.

D. Mapping

Once an arriving frame has been localized, the mapping module is invoked to incorporate the newly observed local points into the global map. Let $p_{k,i}^{\mathrm{Ne}}$ and $q_{k,i}^{\mathrm{Ne}}$ be the coordinates of the i-th newly observed point in the k-th robot frame and the global frame, respectively, which is related by $q_{k,i}^{\mathrm{Ne}} = \phi(T_k, p_{k,i}^{\mathrm{Ne}})$. The uncertainty set $\mathcal{P}(A_{p_{k,i}}^{\mathrm{Ne}}, b_{p_{k,i}}^{\mathrm{Ne}})$ for $p_{k,i}^{\mathrm{Ne}}$ is provided by CP and the uncertainty set $\mathcal{T}_k = \{T_k \in \mathrm{SE}(3) \mid H_k\mathrm{vec}(T_k) \leq d_k\}$ for T_k is provided by the localization function. We can apply the forward propagation in Sec.III-A to compute the uncertainty set $\mathcal{P}(A_{q_{k,i}}^{\mathrm{Ne}}, b_{q_{k,i}}^{\mathrm{Ne}})$ for $q_{k,i}^{\mathrm{Ne}}$.

E. Smoothing

Upon loop closure detection, the trajectory and map are jointly optimized to mitigate accumulated drift and improve global consistency. As shown in Figure 1, the frames that involve in the loop closures may additionally link to more map points, denoted by dashed factors in the figure. Then, we jointly optimize all frame poses and map points within the loop closures. Specifically, we adopt an alternating optimization scheme: at each iteration, we first update the map points and subsequently refine the frame poses. For each map point, we incorporate all associated frames and execute the forward UQ in Sec.III-A to refine the map uncertainty set. To update each frame pose, we utilize all visible points and apply the backward UQ in Sec.III-B. Because both the pose and map point sets decrease monotonically and are bounded below by the empty set, the iteration is guaranteed to converge. As detailed in the supplementary pseudocode, we employ a maximum of three iterations or stop when the uncertainty set refinement becomes negligible. It is noteworthy that if an incorrect loop closure is introduced, the uncertainty sets may shrink incorrectly, possibly losing containment or even becoming empty. Note that empty uncertainty sets can indicate an incorrect loop closure and thus support rejection.

The pseudo code for the whole pipeline is provided in Sec.VI in suppl. material.

Remark 1. *Our algorithm relies on the exchangeability assumption of calibration and test data to construct statistically valid uncertainty sets via CP. In SLAM, this assumption holds when both datasets originate from the same underlying process, encompassing the sensor model, front-end pipeline, and environmental distribution. Outliers do not inherently violate exchangeability, provided the contamination mechanism remains consistent across both datasets; however, higher outlier rates typically yield looser uncertainty sets. As the core building blocks in our SLAM UQ pipeline, three primitives (forward UQ, backward UQ, and pose compound) require no specialized assumptions. Our theoretical derivation rigorously shows that if the input uncertainty sets are valid, then the outputs are guaranteed to contain the true values. While we integrate these primitives with CP here, they remain agnostic to the specific local point cloud calibration method used.*

V. SIMULATION AND EXPERIMENT

A. Basic Uncertainty Quantification Simulation

1) Forward uncertainty quantification: To test our forward UQ in Sec.III-A, six random trials are performed. In each trial, we randomly generate a pose $T \in \mathrm{SE}(3)$ and a point $p \in \mathbb{R}^3$. Uncertainty generation for T and p and more detailed settings can be found in Sec.VII in suppl. material. To verify the conservatism of our algorithm, we randomly sample 1000 points from the original set $\mathcal{Q}$ for $q = \phi(T, p)$ and test whether they fall inside our estimated set $\mathcal{P}(A_2, \hat{b}_2)$. The result is that all samples belong to the estimated set, and Figure 3 visualizes one of the trials, which validates our guaranteed UQ.

2) Backward uncertainty quantification: In this part, we evaluate our backward UQ in Sec.III-B and compare it with the pose UQ methods `CLOSURE` [7] and `GRCC` [15]. `CLOSURE` produces an inner approximation of the original set, whereas `GRCC` gives an outer approximation. Six random trials are conducted, where the detailed setting can be found in Sec.VII in the suppl. material. For `CLOSURE`, we adopt the default reasonable parameter configuration provided in [7], without tuning its hyperparameters in the random walk sampling. The conservatism test results are presented in Table I. All pose

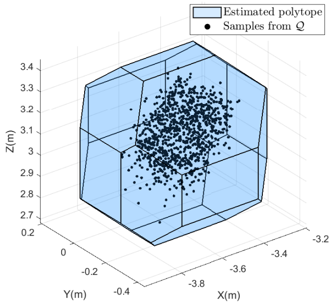

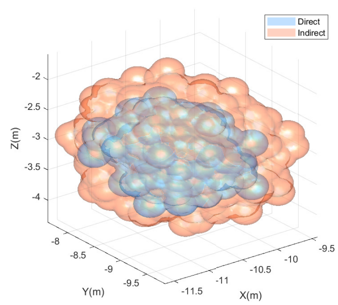

Figure 3: Forward propagation. Samples from $\mathcal{Q}$ all lie within our estimated polytope.

Figure 4: Sampling-based pose uncertainty visualization in pose compound.

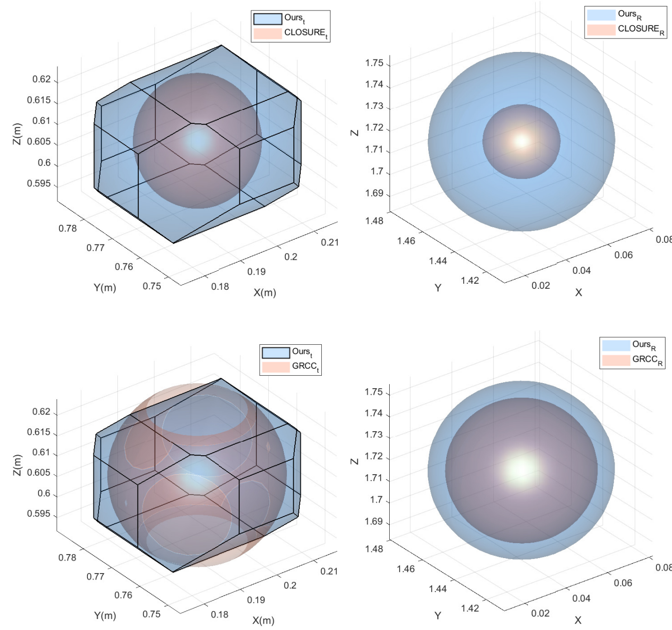

Figure 5: Visualization of uncertainties for translation (left) and rotation (right) in backward UQ.

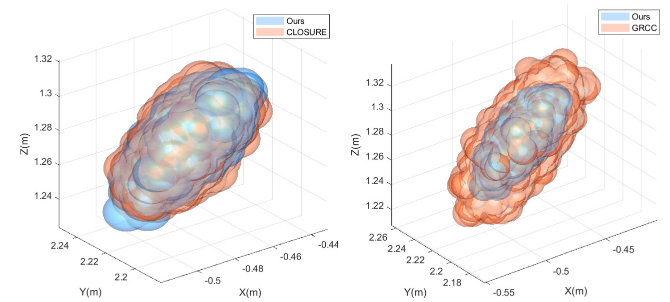

Figure 6: Sampling-based pose uncertainty visualization in backward UQ. A ball centered at $[1\ 1\ 1]^\top$ with radius 0.01m is transformed by 1000 poses sampled from the estimated pose set, and the union of the transformed balls is plotted.

samples drawn from the original set $\mathcal{T}$ lie within the sets given by our algorithm and GRCC, validating the guaranteed property of the two methods. However, since CLOSURE only gives an inner approximation of $\mathcal{T}$, some samples fall outside its estimated set. In addition, the conservatism of CLOSURE varies a lot in different random trails, resulting in inconsistent containment rates.

We note that CLOSURE and GRCC quantify translation and rotation uncertainties separately. The translation uncertainty is represented by a Euclidean ball $\{t \in \mathbb{R}^3 \mid \|t - \bar{t}\| \leq r_t\}$, and the rotation uncertainty by a geodesic ball $\{R \in \mathrm{SO}(3) \mid \mathrm{dis}(R, \bar{R}) \leq r_R\}$. To enable a direct comparison, we follow the convention used by the baselines by projecting our estimated polytopes $\mathcal{P}(H, d)$ onto the rotation and translation polytopes. We then enclose each rotation polytope within a geodesic ball; note that this representation tends to favor the baseline methods by changing our polytopic bounds. For visualization, we map rotation matrices to $\mathbb{R}^3$ via the Log operation and plot the rotation uncertainty as the Euclidean ball $\{s \in \mathbb{R}^3 \mid \|s - \mathrm{Log}(\bar{\mathrm{R}})\| \leq r_R\}$. The translation and rotation uncertainties in the first trial are visualized in Figure 5. Our projected results are more conservative than CLOSURE for both translation and rotation and more conservative than GRCC for rotation. For complete pose uncertainty visualization, we adopt a sampling-based method. Specifically, we draw 1000 poses from the estimated set, transform a ball of radius 0.01m using these poses, and visualize the union of the transformed balls. The ball center is set as $[1\ 0\ 0]^\top, [0\ 1\ 0]^\top, [0\ 0\ 1]^\top, [1\ 1\ 1]^\top$, respectively, and the results in the first trial with ball center $[1\ 1\ 1]^\top$ are shown in Figure 6. The full results are available in Sec.VII in suppl. material. Interestingly, unlike Figure 5, the pose uncertainty of our method is comparable to CLOSURE and tighter than GRCC. The reason may be attributed to the conservatism introduced by the comparison adaptation to ensure a direct comparison.

3) Pose compound: We perform six trials. In each trial, we randomly generate poses T_1 and T_2 and use the same uncertainty-generation procedure as in the forward UQ experiments. To assess conservatism, we draw 1000 samples from the original set $\mathcal{T}_3$ and verify that all samples lie within the pose uncertainty sets provided by both the direct method and the indirect method in Sec.III-C. We plot the translation and rotation uncertainties in the first trial in Figure 8. The direct algorithm yields a more conservative translation uncertainty but a tighter rotation uncertainty. Part of the higher rotation conservatism of the indirect method is attributed to its geodesic-ball approximation. We then use a sampling-based approach for complete pose uncertainty visualization. We draw 1000 poses from the estimated set and transform a ball of radius 0.2m. The ball center is set as $[10\ 0\ 0]^\top, [0\ 10\ 0]^\top, [0\ 0\ 10]^\top, [10\ 10\ 10]^\top$, respectively, and the results in the first trial with ball center $[10\ 10\ 10]^\top$ are shown in Figure 4. The full results are available in Sec.VII in suppl. material. We see that the direct algorithm achieves a tighter pose uncertainty.

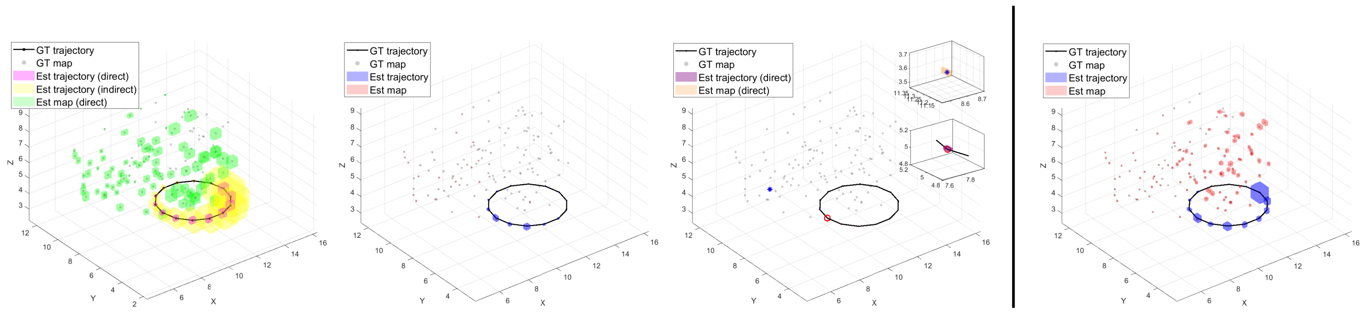

(a) Relative framework (b) Global framework (c) Loop closure (d) Global framework ablation

Figure 7: Guaranteed SLAM simulation results. In the relative framework, "direct" and "indirect" denote the direct pose compound and the indirect pose compound results, respectively, and the map are constructed based on direct pose compound results. The loop closure results are obtained by feeding the "direct" results into the smoothing module. The global framework ablation uses the setting that each frame registers all visible (instead of newly observed) landmarks into the global map.

Table I: Conservatism test for backward UQ: the percentage of 1000 samples from the original set $\mathcal{T}$ that lie within the estimated set.

Trial	1	2	3	4	5	6
CLOSURE	87.1%	59.6%	32.8%	71.3%	73.2%	10.4%
GRCC	100%	100%	100%	100%	100%	100%
Ours	100%	100%	100%	100%	100%	100%

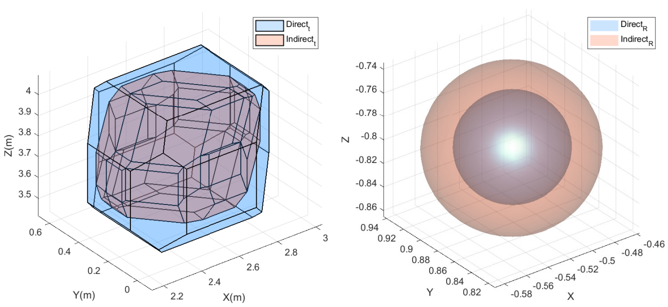

Figure 8: Visualization of uncertainties for translation (left) and rotation (right) in pose compound.

B. Numerical SLAM Simulation

We evaluate the proposed guaranteed SLAM UQ algorithms using synthetic data. We uniformly scatter 3D landmarks within a $50\text{m} \times 50\text{m} \times 50\text{m}$ workspace and randomly generate a circular robot trajectory. The robot's field of view is 60° horizontally, 60° vertically, and $[0.5, 5]\text{m}$ in depth. For each visible landmark, we impose box uncertainties on the visual measurements, where the box size is proportional to the landmark's distance from the robot. The results are shown in Figure 7. In the relative framework, uncertainties for both robot poses and landmark positions accumulate over time. Notably, the pose uncertainty in the indirect compound algorithm grows more rapidly than in the direct one, aligning with the results from the previous subsection. We note that the map and trajectory uncertainties in the global framework are much smaller than those in the relative framework. This phenomenon occurs because, during motion, the robot may observe previously registered landmarks with low uncertainty; these observations mitigate current pose uncertainties. To substantiate this explanation, we conduct an ablation study in which each frame registers all visible landmarks $p_{k,i}$, rather than only newly observed landmarks $p_{k,i}^{\text{Ne}}$, into the global map. Under this modification, previously registered landmarks are re-registered whenever they become visible in a new frame and thus may have a larger uncertainty set. The ablation results show that the estimated uncertainty increases significantly, especially when the trajectory returns to the starting point, which verifies our claim. Moreover, when a loop closure is detected and smoothing is performed, the uncertainties of both the trajectory and the map can be significantly reduced. Finally, we observe that all ground-truth poses and landmarks lie within their corresponding uncertainty sets.

We also evaluate the computational complexity of our algorithm. When the average number of points per frame is 19, the average runtime of each module (MATLAB implementation on a computer with 32 GB RAM and an Intel Core Ultra 7 265 CPU) is as follows: forward UQ (9.3s), backward UQ (0.0036s), direct pose compound (12s), indirect pose compound (0.5s), and loop-closure smoothing (29.55s per frame, 591s over 20 frames). The primary computational bottleneck in our framework arises from the repeated SDP optimizations required for forward UQ and direct pose compound. While our current MATLAB implementation is a research prototype not yet intended for real-time use, it remains usable and effective for offline applications like uncertainty-aware mapping.

For comparison with probabilistic counterparts, we implement a probabilistic baseline that computes the exact posterior in a linear-Gaussian setting without $\text{SE}(3)$ constraints. We conduct open-loop comparisons under both global and relative frameworks. Unlike our algorithm, this baseline provides no formal containment guarantee. We represent its uncertainty sets by covariance-induced ellipsoids at the 99% confidence level. In the global framework, with a moderate number of points per frame (~ 15), by frame 20, the baseline fails to

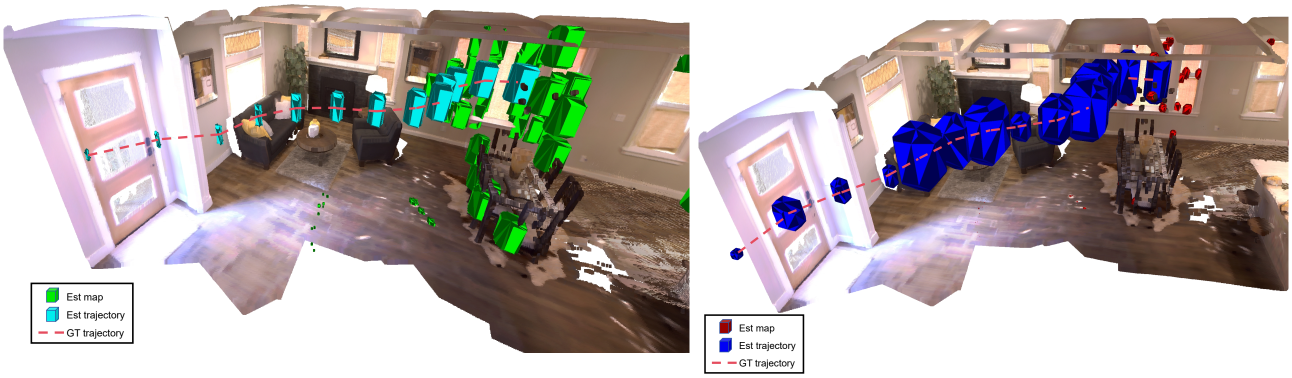

(a) Relative framework (direct)

(b) Global framework

Figure 9: Replica experiment results. The proposed relative framework using the direct pose compound method (left); the proposed global framework (right).

cover the ground truth in $9/20$ pose estimates and $26/120$ map points, whereas our method achieves full containment. The probabilistic baseline generally exhibits lower uncertainty than our method. However, since it ignores $\mathrm{SE}(3)$ constraints, in the relative framework, its uncertainty grows faster when only a few points are visible. For ~ 5 points per frame, by frame 30, the translation ellipsoid volume is 3.12, whereas our polytope volume is 1.2, showing about $3\times$ slower growth.

C. SLAM Experiment

We evaluate our method on the Replica dataset [23], which provides photo-realistic indoor environments. For any camera pose, it renders an RGB image and a dense depth map, which we back-project to generate a point cloud. To calibrate CP, we randomly generate three trajectories, each containing hundreds of frames. When a new frame arrives, we extract ORB features [22] from the current and previous frames and perform feature matching. To ensure tighter uncertainty bounds, we apply RANSAC-based outlier rejection. We reject outliers using the feature depths and the ground-truth relative pose provided by the dataset. For each 3D correspondence (p_l, q_l), we transform p_l using the relative pose (R, t) and define the nonconformity score as $s_l = \|q_l - Rp_l - t\|$. We set the target miscoverage level in CP to $\delta = 0.01$ to calibrate local point uncertainties. During testing, we use the same feature extraction and matching pipeline. Loop closure detection is not involved. To balance runtime and accuracy, we randomly retain 15 correspondences for each pair of consecutive frames. In our feature-matching setting, the global framework works in the same way as in the ablation study in Figure 7. As shown in Figure 9, the estimated trajectories and maps enclose their ground truths well. An interesting phenomenon is that the global framework achieves tighter map uncertainty but more conservative trajectory (translation) uncertainty than the relative framework. The reason is that although the marginally projected translation uncertainty of the global framework is larger, its joint pose uncertainty in the form $\{T \in \mathrm{SE}(3) \mid Hx(T) \leq d\}$ is smaller due to the coupling between translation and rotation. To verify this, we conduct a sampling-based pose uncertainty test. As shown in Figure S4 in the suppl. material, the poses sampled from the relative framework-generated uncertainty sets span a ball to a larger volume, indicating its larger pose uncertainty. Overall, the uncertainty sets of both frameworks expand relatively rapidly in the absence of loop-closure smoothing.

VI. Conclusion and Future Work

In this paper, we proposed SME-based UQ for 3D-3D SLAM pipelines. We identified three UQ primitives—forward UQ, backward UQ, and pose compound—and derived uncertainty sets with guarantees for each. By integrating them with CP, we obtain a complete SLAM UQ algorithm. Representing uncertainties as polytopes further enables tractable computation and a unified treatment of pose uncertainty.

Future work will focus on three directions. (i) Tightness: In our tests, the uncertainty sets expand relatively quickly. We plan to incorporate sliding-window bundle adjustment and develop tighter SDP formulations to slow this growth. We are also interested in providing a rigorous theoretical analysis of the uncertainty set growth rate. (ii) Real-time performance: While tractable, the current algorithms are still in an early stage; future iterations will target real-time performance by transitioning to C++ and incorporating parallel SDP computation, keyframe mechanisms, and customized solver designs. (iii) Practical utility: Our current implementation is limited to 3D-3D landmark-based SLAM. Extending the approach to more sensor modalities (e.g., IMUs and range sensors) and validating it in complex, large-scale (loop-closure) scenarios remain promising directions.

Acknowledgement

This work was supported in part by NSFC under Grants 62273288 and 62336005, in part by the Shenzhen Science and Technology Program under Grants JCYJ20241202124010014 and JCYJ20240813113609013, and in part by Australian Research Council (ARC) under Grants DP190103615, LP210200473, and DP230101014.

References

[1] Igor Cvišić, Ivan Marković, and Ivan Petrović. Soft2: Stereo visual odometry for road vehicles based on a point-to-epipolar-line metric. *IEEE Transactions on Robotics*, 39(1):273–288, 2022.

[2] Heng Yang, Jingnan Shi, and Luca Carlone. Teaser: Fast and certifiable point cloud registration. *IEEE Transactions on Robotics*, 37(2):314–333, 2020.

[3] Haoyu Han and Heng Yang. Building rome with convex optimization. In *Proceedings of Robotics: Science and Systems*, pages 1–19, 2025.

[4] Jakob Engel, Thomas Schöps, and Daniel Cremers. LSD-SLAM: Large-scale direct monocular SLAM. In *Proceedings of European Conference on Computer Vision*, pages 834–849, 2014.

[5] David M Rosen, Luca Carlone, Afonso S Bandeira, and John J Leonard. SE-Sync: A certifiably correct algorithm for synchronization over the special euclidean group. *The International Journal of Robotics Research*, 38(2-3):95–125, 2019.

[6] Joshua G Mangelson, Maani Ghaffari, Ram Vasudevan, and Ryan M Eustice. Characterizing the uncertainty of jointly distributed poses in the Lie algebra. *IEEE Transactions on Robotics*, 36(5):1371–1388, 2020.

[7] Yihuai Gao, Yukai Tang, Han Qi, and Heng Yang. Closure: Fast quantification of pose uncertainty sets. In *Proceedings of Robotics: Science and Systems*, pages 1–13, 2024.

[8] Devansh Ramgopal Agrawal, Ruichang Chen, and Dimitra Panagou. gatekeeper: Online safety verification and control for nonlinear systems in dynamic environments. *IEEE Transactions on Robotics*, 40:4358–4375, 2024.

[9] Richard Szeliski. *Computer Vision: Algorithms and Applications*. Springer Nature, 2022.

[10] Mohamed Mustafa, Alexandru Stancu, Nicolas Delanoue, and Eduard Codres. Guaranteed SLAM—an interval approach. *Robotics and Autonomous Systems*, 100:160–170, 2018.

[11] Raphael Voges and Bernardo Wagner. Interval-based visual-lidar sensor fusion. *IEEE Robotics and Automation Letters*, 6(2):1304–1311, 2021.

[12] Aaronkumar Ehambram, Raphael Voges, Claus Brenner, and Bernardo Wagner. Interval-based visual-inertial LiDAR SLAM with anchoring poses. In *Proceedings of International Conference on Robotics and Automation*, pages 7589–7596, 2022.

[13] Lars Lindemann, Yiqi Zhao, Xinyi Yu, George J Pappas, and Jyotirmoy V Deshmukh. Formal verification and control with conformal prediction: Practical safety guarantees for autonomous systems. *IEEE Control Systems*, 45(6):72–122, 2025.

[14] Anastasios N Angelopoulos, Stephen Bates, et al. Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning*, 16(4):494–591, 2023.

[15] Yukai Tang, Jean-Bernard Lasserre, and Heng Yang. Uncertainty quantification of set-membership estimation in control and perception: Revisiting the minimum enclosing ellipsoid. In *Proceedings of Annual Learning for Dynamics & Control Conference*, pages 286–298, 2024.

[16] Heng Yang and Marco Pavone. Object pose estimation with statistical guarantees: Conformal keypoint detection and geometric uncertainty propagation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8947–8958, 2023.

[17] Glenn Shafer and Vladimir Vovk. A tutorial on conformal prediction. *Journal of Machine Learning Research*, 9(3):371–421, 2008.

[18] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.

[19] Frederike Dümbgen, Connor Holmes, Ben Agro, and Timothy Barfoot. Toward globally optimal state estimation using automatically tightened semidefinite relaxations. *IEEE Transactions on Robotics*, 40:4338–4358, 2024.

[20] Guangyang Zeng, Yuan Shen, Ziyang Hong, Yuze Hong, Viorela Ila, Guodong Shi, and Junfeng Wu. Bias-eliminated PnP for stereo visual odometry: Provably consistent and large-scale localization. *IEEE Robotics and Automation Letters*, 10(11):11840–11847, 2025.

[21] Ji Zhang, Sanjiv Singh, et al. LOAM: Lidar odometry and mapping in real-time. In *Proceedings of Robotics: Science and Systems*, pages 1–9, 2014.

[22] Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap SLAM.

[23] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.