

# GAPSLAM: Blending Gaussian Approximation and Particle Filters for Real-Time Non-Gaussian SLAM

Qiangqiang Huang<sup>1</sup> and John J. Leonard<sup>1</sup>

**Abstract**—Inferring the posterior distribution in SLAM is critical for evaluating the uncertainty in localization and mapping, as well as supporting subsequent planning tasks aiming to reduce uncertainty for safe navigation. However, real-time full posterior inference techniques, such as Gaussian approximation and particle filters, either lack expressiveness for representing non-Gaussian posteriors or suffer from performance degeneracy when estimating high-dimensional posteriors. Inspired by the complementary strengths of Gaussian approximation and particle filters—scalability and non-Gaussian estimation, respectively—we blend these two approaches to infer marginal posteriors in SLAM. Specifically, Gaussian approximation provides robot pose distributions on which particle filters are conditioned to sample landmark marginals. In return, the maximum a posteriori point among these samples can be used to reset linearization points in the nonlinear optimization solver of the Gaussian approximation, facilitating the pursuit of global optima. We demonstrate the scalability, generalizability, and accuracy of our algorithm for real-time full posterior inference on realworld range-only SLAM and object-based bearing-only SLAM datasets.

## I. INTRODUCTION

Simultaneous localization and mapping (SLAM) is a foundational technique for autonomous driving, AR/VR, and robotic navigation. While the primary goal of SLAM is to find a point estimate of robot paths and maps, inferring the posterior density encountered in SLAM is also essential. This is because full posterior inference reveals how uncertain the robot path and map could be, and supports the robot in planning how to reduce the uncertainty for safe navigation [1].

Standard *real-time* methods for full posterior inference in SLAM research either compute a Gaussian approximation, centered on the point estimate, to the posterior [2], or represent the posterior using samples in the framework of Rao-Blackwellized particle filtering (RBPF) [3]. Constructing the Gaussian approximation enjoys great advantages in scalability for tackling high-dimensional posteriors, owing to efficient nonlinear optimization solvers for SLAM [4, 5]. However, the Gaussian approximation is inherently incapable of describing highly non-Gaussian/multi-modal posteriors, which often appear in realworld SLAM problems due to nonlinear measurement models. While the RBPF can capture non-Gaussian features of the posterior via samples, it suffers from scalability issues incurred by particle degeneracy and depletion [6].

<sup>1</sup>Computer Science and Artificial Intelligence Lab, Massachusetts Institute of Technology, Cambridge, MA 02139, USA{hqq, jleonard}@mit.edu

This work was supported by ONR grant N00014-18-1-2832, ONR Neuroautonomy MURI grant N00014-19-1-2571 and the MIT Portugal Program. The project page is available at <https://github.com/doublestrong/gapslam>.

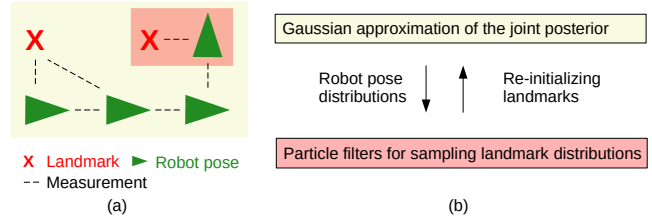


Fig. 1. Illustration of our method for inferring robot pose and landmark distributions: (a) a SLAM example, where the robot moves along poses in green and makes measurements to landmarks in red, and (b) our method, which blends Gaussian approximation in yellow and particle filters in pink. The Gaussian approximation, centered on the maximum a posteriori (MAP) estimate, provides distributions of robot poses on which the particle filters are conditioned to draw samples that represent landmark distributions. If a sample attains a higher probability than the MAP estimate when evaluating the posterior, landmarks in the Gaussian solver will be re-initialized by that sample.

By blending advantages of the **Gaussian Approximation** and **Particle filters** on scalability and non-Gaussian estimation, respectively, we present a novel algorithm, **GAPSLAM**, to infer marginal posteriors encountered in SLAM. The contributions of the paper include:

- 1) An adaptive modeling strategy whereby only marginals with great uncertainty are sampled by particle filters, while others are represented by the Gaussian approximation.
- 2) An uncertainty-aware re-initialization technique that leverages particle filters to reset linearization points in the nonlinear optimization solver.
- 3) Range-only and object-based bearing-only SLAM experiments that demonstrate the scalability, generalizability, and accuracy of GAPSLAM, as well as its ability to precisely describe the evolution of non-Gaussian posteriors in real-time.

## II. RELATED WORK

We briefly review some work on full posterior inference for SLAM with a focus on non-Gaussian representations of the posterior. Note that the non-Gaussian representation can be parametric models (e.g., sum of Gaussians) or non-parametric (i.e., samples).

The framework of RBPF has been leveraged in a big class of SLAM algorithms (e.g., FastSLAM) [3, 7, 8], where the posterior of robot poses is represented by a set of particles, and each particle is attached with parametric models [3, 8] or samples [7] of the conditional of the map. Although our work exploits the same conditional independence relation as FastSLAM, i.e., landmarks are conditionally independent given robot poses, our work swaps representations of the robot path and map in RBPF, resulting in a parametric (Gaussian) model to describe the belief about robot poses and a set of particle filters to independently draw samples

from the posterior of each landmark. Thus the issue of losing diversity in robot path particles no longer exists, warranting the scalability of our work. We will compare to [8] in our experiments section. mm-iSAM [9] and NF-iSAM [10] are recent non-Gaussian inference algorithms that leverage the Bayes tree algorithm [4] to exploit more conditional independence relations. As solvers for general factor graphs, they can tackle both nonlinear measurement models and multi-modal data association, while our work is only focused on nonlinear measurement models. Another class of approaches purely relies on Monte Carlo techniques to directly draw samples from the joint posterior [11–13]. These methods, in general, do not suit real-time online applications due to the exorbitant computational cost but can serve as reference solutions to full posterior inference.

Our work is inspired by some works [7, 8, 14] dedicated to bearing-only (point-object-based) or range-only SLAM. Existing approaches for these two problems can be categorized as batch optimizations [15], delayed initializations of landmarks [14, 16], and undelayed initializations of landmarks [7, 8, 17, 18]. Our work falls into the undelayed category. Our experiments show that we can solve both of these problems in a unified framework and demonstrate the evolution of landmark posteriors since time step zero. Furthermore, we clarify that, while our object-based SLAM system requires similar input (RGB images and object detections) as the systems in [19–22], we do not estimate occupied areas of objects and focus solely on inferring how distributions of object locations evolve.

### III. PROPOSED APPROACH

We focus on landmark-based SLAM. Let  $\mathbf{X}_t := (X_0, X_1, \dots, X_t)$  be robot pose variables up to time step  $t$  where  $X_i \in \text{SE}(d)$ . Landmark variables are denoted by  $\mathbf{L}_n := (L_1, L_2, \dots, L_n)$  where each element denotes the location or pose of a landmark (e.g., ultra-wideband tags, 3D points of visual features, and daily objects). In the SLAM problems, we consider two types of measurements: i) odometry and ii) landmark measurements (e.g., distances, directions, and relative poses). Let  $\mathbf{o}_t := (o_1, o_2, \dots, o_t)$  be odometry measurements of which each is modeled by the likelihood function  $p(o_i|X_{i-1}, X_i)$ . Without loss of generality, for the ease of notation in the method formulation, we assume the robot makes at most a landmark measurement at a time. Let  $\mathbf{z}_t := (z_0, z_1, \dots, z_t)$  be all landmark measurements in which each is modeled by the likelihood function  $p(z_i|X_i, L_{d_i})$ , where  $d_i \in \{1, 2, \dots, n\}$  is the landmark index associated with measurement  $z_i$ . Let  $\mathbf{d}_t := (d_0, d_1, \dots, d_t)$  be data associations for all landmark measurements thus the posterior of robot and landmark variables at time step  $t$  can be formulated by

$$p(\mathbf{X}_t, \mathbf{L}_n | \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \quad (1)$$

$$\stackrel{\text{Bayes}}{\propto} \prod_{i=0}^t p(z_i|X_i, L_{d_i}) \prod_{i=1}^t p(o_i|X_{i-1}, X_i) p(\mathbf{X}, \mathbf{L}), \quad (2)$$

where  $p(\mathbf{X}, \mathbf{L})$  denotes the prior.

Our goal is to infer marginal posteriors. Our method represents marginals either by samples or parametric models, depending on the uncertainty of the marginal. This hybrid

density modeling aims to achieve a balance between computational efficiency and expressiveness. Our method is inspired by a typical observation about the posterior: oftentimes, only a few of the landmarks have very uncertain marginal posteriors, while marginals of most variables can be well described by Gaussian distributions. For example, a noisy range-only or bearing-only measurement between a robot pose and a landmark can simulate either spherical or conic landmark distributions in 3D, while accumulated measurements may well constrain the landmark to a concentrated and (almost) unimodal distribution.

Specifically, we use particle filters to estimate very uncertain landmarks for the expressiveness of sample-based density representations. Meanwhile, we maintain a Gaussian approximation of the posterior for computation efficiency. Combining these two techniques leads to a win-win situation over time: i) the Gaussian approximation provides particle filters with a parametric model of the smoothed robot path, reducing the sampling complexity to landmarks, and ii) particle filters afford the Gaussian approximation statistically probable values of landmarks, which can be used to explicitly reset linearization points in nonlinear optimization solvers, thereby benefiting the pursuit of global optima. We describe the details of our method in the following sections.

#### A. Gaussian Approximation

We denote the Gaussian approximation (GA) of the posterior at time step  $t$  as  $g(\mathbf{X}_t, \mathbf{L}_n | \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t)$ . We categorize all landmarks into a set of non-Gaussian landmarks  $L_{\mathcal{N}_t}$  and a set of Gaussian landmarks  $L_{\mathcal{G}_t} = \mathbf{L}_n \setminus L_{\mathcal{N}_t}$ , where  $\mathcal{N}_t \cup \mathcal{G}_t = \{1, 2, \dots, n\}$  and  $\mathcal{N}_t \cap \mathcal{G}_t = \emptyset$  (see Sec. III-C for the approach to identify Gaussian landmarks). We use the Gaussian approximation to represent the posterior of robot poses and the Gaussian landmarks, resulting in

$$p(\mathbf{X}_t, \mathbf{L}_n | \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \quad (3)$$

$$\stackrel{\text{GA}}{\approx} p(L_{\mathcal{N}_t} | \mathbf{X}_t, L_{\mathcal{G}_t}, \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) g(\mathbf{X}_t, L_{\mathcal{G}_t} | \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t), \quad (4)$$

$$\stackrel{\text{CIR}}{\equiv} p(L_{\mathcal{N}_t} | \mathbf{X}_t, \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) g(\mathbf{X}_t, L_{\mathcal{G}_t} | \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t), \quad (5)$$

$$\stackrel{\text{CIR}}{\equiv} \prod_{i \in \mathcal{N}_t} p(L_i | X_{\mathcal{V}_i}, z_{\mathcal{V}_i}, d_{\mathcal{V}_i}) g(\mathbf{X}_t, L_{\mathcal{G}_t} | \mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \quad (6)$$

where  $X_{\mathcal{V}_i}$  denotes the set of robot poses that observed landmark  $L_i$  (i.e.,  $\mathcal{V}_i = \{j \in [0, t] | d_j = i\}$  denotes the time steps observing landmark  $L_i$ ). The reduction from (4) to (5) exploits the conditional independence relation (CIR) that landmarks are conditionally independent given a robot path. Note that while we use the CIR in a similar way to typical Rao-Blackwellized particle filters (RBPFs) for SLAM (e.g., FastSLAM 2.0 [3]), our probabilistic modeling and inference methods differ from these RBPFs. The Gaussian approximation can be computed via the Laplace approximation [23]: 1) finding the MAP estimate as the mean and 2) estimating the inverse of the Hessian of the negative logarithm of the posterior at the MAP estimate as the covariance. In practice, many nonlinear local optimization solvers, such as GTSAM, are capable of providing efficient out-of-the-box solutions to the Gaussian approximation, although they are subject to local optima of the MAP estimate. In this paper, we refer to

them as Gaussian solvers.

$$p(\mathbf{X}, \mathbf{L}|\mathbf{z}) \quad (7)$$

$$\stackrel{\text{GA}}{\approx} p(L_{\mathcal{N}}|\mathbf{X}, L_G, \mathbf{z})g(\mathbf{X}, L_G|\mathbf{z}), \quad (8)$$

$$\stackrel{\text{CIR}}{\approx} \prod_{i \in \mathcal{N}} p(L_i|X_{\mathcal{V}_i}, \mathbf{z})g(\mathbf{X}, L_G|\mathbf{z}) \quad (9)$$

### B. Marginal Posterior of non-Gaussian Landmarks

We will use particles to represent the marginal posterior of any non-Gaussian landmark  $L_i \in L_{\mathcal{N}_i}$ . Alg. 1 summarizes how we draw the particles. Integrating out all variables except  $L_i$  in (6), we can formulate the marginal posterior by

$$p(L_i|\mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \quad (10)$$

$$\stackrel{\text{GA}}{\approx} \int_{X_{\mathcal{V}_i}} p(L_i|X_{\mathcal{V}_i}, z_{\mathcal{V}_i}, d_{\mathcal{V}_i})g(X_{\mathcal{V}_i}|\mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \quad (11)$$

$$\stackrel{\text{MC}}{\approx} \frac{1}{K} \sum_{k=1}^K p(L_i|X_{\mathcal{V}_i} = x_{\mathcal{V}_i}^{(k)}, z_{\mathcal{V}_i}, d_{\mathcal{V}_i}), \quad (12)$$

where  $\{x_{\mathcal{V}_i}^{(k)}\}_{k=1}^K$  are  $K$  i.i.d. samples of part of the robot path drawn from the Gaussian marginal  $g(X_{\mathcal{V}_i}|\mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t)$  (line 1 in Alg. 1). Monte Carlo (MC) integration is applied in (12). The conditional of  $L_i$  in (12) is a result of eliminating  $L_i$  from factors adjacent to  $L_i$  (line 2 in Alg. 1), thus

$$p(L_i|X_{\mathcal{V}_i}, z_{\mathcal{V}_i}, d_{\mathcal{V}_i}) = \frac{\psi(L_i, X_{\mathcal{V}_i})}{\int_{L_i} \psi(L_i, X_{\mathcal{V}_i})}, \quad (13)$$

where

$$\psi(L_i, X_{\mathcal{V}_i}) = \prod_{j \in \mathcal{V}_i} p(z_j|X_j, L_i)p(L_i). \quad (14)$$

We use importance sampling to draw samples representing  $p(L_i|X_{\mathcal{V}_i} = x_{\mathcal{V}_i}^{(k)}, z_{\mathcal{V}_i}, d_{\mathcal{V}_i})$ . We design the proposal distribution of landmark  $L_i$  as the sum-mixture of binary factors, as seen in

$$q(L_i; x_{\mathcal{V}_i}^{(k)}) = \frac{1}{|\mathcal{V}_i|} \sum_{j \in \mathcal{V}_i} p(L_i|X_j = x_j^{(k)}, z_j). \quad (15)$$

This proposal<sup>1</sup> was chosen for two reasons: i) covering more area in the space of the landmark variable and ii) it is efficient to draw landmark samples  $\{l_i^{(m,k)}\}_{m=1}^M$  from the proposal  $q(L_i)$  by sampling a multinomial distribution and binary factors independently. With proposal samples  $\{l_i^{(m,k)}\}_{m=1}^M$  in hand, the normalized weight of each sample can be computed by

$$w^{(m,k)} \propto \frac{p(l_i^{(m,k)}|x_{\mathcal{V}_i}^{(k)}, z_{\mathcal{V}_i}, d_{\mathcal{V}_i})}{q(l_i^{(m,k)}; x_{\mathcal{V}_i}^{(k)})} \propto \frac{\psi(l_i^{(m,k)}, x_{\mathcal{V}_i}^{(k)})}{q(l_i^{(m,k)}; x_{\mathcal{V}_i}^{(k)})}, \quad (16)$$

where  $\sum_{m=1}^M w^{(m,k)} = 1$  (line 8 in Alg. 1). Finally we apply the re-sampling and regularization steps in [6, Alg. 6] to generate equally weighted samples of  $L_i$  (line 9 in Alg. 1).

<sup>1</sup>Although we sacrificed sample diversity, we adopted a lighter-weight proposal in our experiments to ensure computational efficiency. Specifically, we limited the number of components in the sum-mixture to 5 and chose these components randomly.

---

### Algorithm 1: LandmarkSampler

---

**Input:** Landmark  $l$ , Gaussian solver  $g$

- 1 Draw robot path samples  $\mathcal{X}$  from the Gaussian approx.
- 2  $\mathcal{F} = \{f(l, \mathbf{x})\} \leftarrow$  All factors adjacent to  $l$  // robot vars  $\mathbf{x}$
- 3  $\psi(l, \mathbf{x}) \leftarrow$  Product of all factors in  $\mathcal{F}$  // potential func. (14)
- 4  $q(l, \mathbf{x}) \leftarrow$  Sum of elements in  $\mathcal{F}$  // proposal density (15)
- 5 Initialize a set  $\mathcal{S}$  for storing landmark samples
- 6 **for** sample  $\hat{\mathbf{x}}$  in  $\mathcal{X}$  **do**
- 7     Draw landmark samples  $\mathcal{L}$  from density  $q(l, \mathbf{x} = \hat{\mathbf{x}})$
- 8     Weights  $\mathcal{W}$  of landmark samples by computing  $\psi/q$
- 9      $\mathcal{L} \leftarrow$  Resampling using  $\mathcal{L}$  and  $\mathcal{W}$
- 10    Add  $\mathcal{L}$  to  $\mathcal{S}$
- 11 **return**  $\mathcal{S}$  // samples of  $l$

---

### Algorithm 2: GAPSLAM

---

**Input:** New factor  $f$  of landmark  $l$ , Gaussian solver  $g$ , dictionary  $\mathcal{D}$  of non-Gaussian landmark samples

- 1 **if**  $l$  not in the Gaussian solver **then** // new landmark
- 2    Add  $l$  to non-Gaussian landmarks  $\mathcal{D}$  as a new key
- 3 Add  $f$  to the Gaussian solver
- 4 **if**  $l$  in non-Gaussian landmarks  $\mathcal{D}$  **then**
- 5    LandmarkReinitializer( $l, g, \mathcal{D}[l]$ ) // re-init. in the solver
- 6 Update the MAP estimate in the Gaussian solver
- 7 **if**  $l$  in non-Gaussian landmarks  $\mathcal{D}$  **then**
- 8     $\mathcal{D}[l] =$  LandmarkSampler( $l, g$ ) // update samples
- 9    Compute empirical covariance matrix  $C$  of  $\mathcal{D}[l]$
- 10    **if** the largest eigenvalue of  $C$  is small **then**
- 11     Delete  $l$  from  $\mathcal{D}$  // no longer non-Gaussian lmk
- 12 **return**  $g$  and  $\mathcal{D}$  // updated solver and dictionary

---

### Algorithm 3: LandmarkReinitializer

---

**Input:** Landmark  $l$ , Gaussian solver  $g$ , samples  $\mathcal{S}$  of landmark  $l$

- 1  $\psi(l, \mathbf{x}) \leftarrow$  Product of all factors adjacent to  $l$  // robot vars  $\mathbf{x}$
- 2  $\hat{l}, \hat{\mathbf{x}} \leftarrow$  Current estimate in the Gaussian solver  $g$
- 3  $l^* \leftarrow$  Value in  $\{\mathcal{S}, \hat{l}\}$  maximizes function  $\psi(l, \mathbf{x} = \hat{\mathbf{x}})$
- 4 Reset the value of  $l$  in the Gaussian solver to  $l^*$
- 5 **return**  $g$  // updated solver

---

Thus the marginal posterior in (12) can be approximated by samples, as seen in

$$p(L_i|\mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \stackrel{\text{MC}}{\approx} \tilde{p}(L_i|\mathbf{z}_t, \mathbf{o}_t, \mathbf{d}_t) \quad (17)$$

$$= \frac{1}{K} \sum_{k=1}^K \sum_{m=1}^M \delta(L_i - l_i^{(m,k)}). \quad (18)$$

To summarize, for each of the  $K$  robot path samples, we draw  $M$  landmark samples. The delta functions (or other kernels) of all  $KM$  landmark samples can be used to approximate the marginal posterior of the landmark.

### C. Updates with new odometry and landmark measurements

The previous sections describe how the Gaussian approximation helps to compute the particle-based belief of landmarks. Now we discuss how these particles provide the Gaussian approximation with linearization points, thus leading to the uncertainty-aware re-initialization in the Gaussian approximation. We will begin by analyzing the incremental update by one time step.

From time step  $t$  to time step  $t+1$ , the robot receives a new odometry reading  $o_{t+1}$  and a new landmark measurement  $z_{t+1}$ . It is easy to create a new robot pose  $X_{t+1}$  and absorb the odometry  $o_{t+1}$  to form the intermediate Gaussian

$g(\mathbf{X}_{t+1}, \mathbf{L}_n | \mathbf{z}_t, \mathbf{o}_{t+1}, \mathbf{d}_t)$ . However, fusing the measurement  $z_{t+1}$  is a more involved process. Alg. 2 summarizes steps for tackling the landmark measurement. Assuming that the measurement  $z_{t+1}$  comes from a landmark with index  $d_{t+1}$ <sup>2</sup>, several cases are considered:

- 1) If the observed landmark  $L_{d_{t+1}}$  is not new or a non-Gaussian landmark (lines 3 and 6 in Alg. 2), we simply add the likelihood model of  $z_{t+1}$  (i.e., factor in the algorithm) to the Gaussian solver to update the Gaussian approximation without making any further changes.
- 2) If the observed landmark  $L_{d_{t+1}}$  is new (line 2 in Alg. 2), i.e.  $d_{t+1} = n + 1$ , we update the landmark set to  $\mathbf{L}_{n+1} = \mathbf{L}_n \cup \{L_{n+1}\}$  and indices of non-Gaussian landmarks to  $\mathcal{N}_{t+1} = \mathcal{N}_t \cup \{n + 1\}$ . We then add the likelihood model of measurement  $z_{t+1}$  to the Gaussian solver. Finally, we can simulate equally-weighted samples of the landmark using this single measurement and randomly select a point from the samples as the initial value of the landmark.
- 3) If the observed landmark  $L_{d_{t+1}}$  is not new and is in non-Gaussian landmarks  $\mathcal{N}_t$ , this is the challenging case we will focus on.

In the third case, our goal is to use the new measurement  $z_{t+1}$  to determine whether we should explicitly re-initialize the non-Gaussian landmark in the Gaussian solver (line 5 in Alg. 2). Furthermore, if our belief of the landmark converges to a less uncertain situation after fusing measurement  $z_{t+1}$ , we remove the landmark from the non-Gaussian set  $\mathcal{N}_t$  (line 11 in Alg. 2).

Algorithm 3 summarizes how we re-initialize a non-Gaussian landmark  $L_i$ . As described in Sec. III-B and Alg. 1, we already have samples of the landmark at hand, which can be passed to Alg. 3 as input. The union of the samples and the current estimate of the landmark forms a set of candidate points for re-initialization. The point in the set that maximizes the posterior will be used to re-initialize the landmark (lines 3 and 4 in Alg. 3). During the evaluation of the posterior, all variables other than the landmark are fixed by the current mean of the Gaussian approximation. This reduces the evaluation to only the product of factors that are adjacent to the landmark, simplifying the computation. The re-initialization can be formulated as

$$l_i^* = \operatorname{argmax}_{l_i \in \{\mathcal{S}, \hat{l}_i\}} \psi(L_i = l_i, X_{\mathcal{V}_i} = \hat{x}_{\mathcal{V}_i}), \quad (19)$$

where the potential  $\psi(\cdot)$  is the same as (14),  $\mathcal{S} = \{l_i^{(j)}\}_{j=1}^M$  denotes current samples of the landmark, and  $\hat{l}_i$  and  $\hat{x}_{\mathcal{V}_i}$  denote current estimates of the landmark and robot poses in the Gaussian solver.

After performing the re-initialization, we compute the Gaussian approximation for time step  $t + 1$  (line 6 in Alg. 2),  $g(\mathbf{X}_{t+1}, \mathbf{L}_n | \mathbf{z}_{t+1}, \mathbf{o}_{t+1}, \mathbf{d}_{t+1})$ , where the new measurement  $z_{t+1}$  has been incorporated. Given the new Gaussian approximation, we update samples of the marginal posterior

<sup>2</sup>Data association and creation of new landmarks are left to the specific application in Sec. IV-B.3.

of  $L_i$  (line 8 in Alg. 2), which are cached for potential re-initialization in the future. Following [7], we compute an empirical covariance matrix of these samples as well as its largest eigenvalue. When the largest eigenvalue falls below a threshold, it indicates that the belief of landmark  $L_i$  has become sufficiently certain, and  $L_i$  is removed from the non-Gaussian landmark set  $\mathcal{N}_{t+1}$  (line 11 in Alg. 2). As a result,  $L_i$  will no longer be estimated by particle filters and will only be involved in updates in the Gaussian solver. Note that currently, we set the threshold of the eigenvalue with predetermined values. One can explore automatic approaches such as monitoring the convergence rate of the largest eigenvalue or performing normality tests on samples.

#### IV. EXPERIMENTS AND RESULTS

We implemented Algorithms 1-3 in Python by blending our in-house code of particle filters and the Gaussian solver provided by the GTSAM library. In addition, we applied explicit regularization to the Gaussian solver by adding large-covariance priors on new non-Gaussian landmarks. These prior factors were not involved in the sampling or re-initialization and were removed once the corresponding landmarks were removed from non-Gaussian landmarks. We perform two sets of experiments: 1) we validate our solutions for distribution and point estimations using range-only SLAM experiments, and 2) we demonstrate the generalizability and scalability of our method using object-based bearing-only SLAM experiments. All experiments were conducted on a laptop with a 2.30GHz Intel Core i7-10875H CPU running Ubuntu 20.04.4 LTS.

##### A. Range-only SLAM

1) *Datasets and methods for comparison:* The performance of GAPSLAM is evaluated by comparing it with other methods on a range-only SLAM dataset, as range measurements can easily lead to highly non-Gaussian/multi-modal posteriors. We use the Plaza 1 dataset [24] which provides time-stamped range and odometry measurements collected by a mobile robot in a planar environment. Ranges between the robot and four landmarks were measured by an ultra-wide band ranging system so each range measurement was tagged with the identity of the landmark. Two methods were used for comparison: 1) NSFG [13], which directly draws samples from the joint posterior using nested sampling methods [25, 26] and is considered as reference solutions at the expense of computational burden, and 2) RBPF-SOG [8], which models landmark conditionals as sum-of-Gaussians (SOGs) in the RBPF framework, updates the SOGs using multi-hypothesis EKFs, and attains real-time operation. We present two sets of results: 1) marginal posteriors for demonstrating how the Gaussian approximation helps particle filters to draw landmark samples, and 2) point estimates for validating the use of particle filters in re-initializing the Gaussian approximation.

2) *Results on marginal posteriors:* We present the posteriors of early time steps, which are expected to be strongly non-Gaussian (e.g., multi-modal). Fig. 2 shows samples drawn by different methods to represent the posteriors. GAPSLAM shows good consistency with the reference samples by NSFG, while RBPF-SOG overestimates the uncertainty in the posteriors, as indicated by spurious modes of landmark

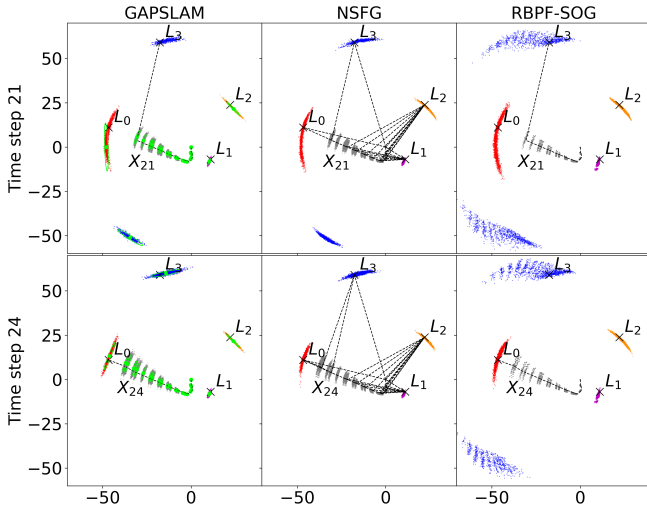


Fig. 2. Samples from marginal posteriors by three methods at two time steps. The robot moves from (0,0) to pose  $X_{21}$  or  $X_{24}$ , observing 4 landmarks  $L_{0-3}$ . Black lines and markers indicate the ground truth. Measurements are shown as dashed lines (accumulated measurements are shown in the middle column). Gray dots indicate robot position samples while colored dots denote landmark samples. Green ellipses in GAPSLAM represent confidence intervals ( $2\sigma$ ) of the Gaussian approximation.

$L_3$  samples (blue dots in the last column of Fig. 2). The excessive number of modes is an expected result of tracking multiple hypotheses in SOGs. Although we applied the strategy suggested in [8] to prune hypotheses with negligible weights, as shown by the decreasing blue line in the bottom of Fig. 3b, the averaged number of modes per landmark is still 11 at time step 21.

In order to explain how GAPSLAM draws the samples in Fig. 2 and when the re-initialization occurs, we can start by looking at the green confidence intervals, which represent the Gaussian approximation. Samples of robot positions and Gaussian landmarks are drawn from the Gaussian. At time step 21,  $L_0$  and  $L_3$  are non-Gaussian landmarks, and we use the method described in Sec. III-B and Algorithm 1 to draw their samples. Following (19), these samples are then used for re-initialization across time steps 21 and 24, which is reflected by the increased computation time for re-initialization in Fig. 3a. Because the set of Gaussian landmarks extends from  $(L_1, L_2)$  at time step 21 to all landmarks at time step 24, the number of non-Gaussian landmarks drops to zero, as shown in the bottom of Fig. 3b.

Fig. 3a illustrates the computation time of different methods for updating their probabilistic models. Note that the computation time for drawing dense samples ( $K = 200, M = 100$  in (18)) of landmarks shown in Fig. 2 is separately shown in the top of Fig. 3b. There are two reasons for excluding it from Fig. 3a: 1) sampling landmark posteriors can be implemented as a parallel process, and 2) RBPf-SOG does not need landmark samples to update the SOGs, and GAPSLAM only needs sparse samples to assist the re-initialization to the Gaussian approximation. For the sparse samples (i.e., ‘update samples’ in Fig. 3a), we draw landmark samples given only the current mean of the robot path (i.e.,  $K = 1$  in (18)). While sampling landmark posteriors consumes slightly more computation time than updating models in GAPSLAM before time step 24, the

TABLE I. Point estimates for the full sequence of the Plaza1 dataset. The RMSE is presented by the mean  $\pm$  standard deviation across 50 runs with different random seeds which lead to different random initial values in the Gaussian solver. The case ‘GAP. w/o reinit.’ serves as an ablation study and is configured by disabling the non-Gaussian landmarks set and operations for re-initialization and sampling, which is essentially just GTSAM with the explicit regularization treatment (large-covariance priors). The RMSE in this case is computed on 29 out of 50 runs (i.e., 58% success rate) that were not ceased by the GTSAM error, indeterminate linear system. The GTSAM error is expected because, in this problem, bad initial values may incur near-singular linear systems.

Metric	Odom.	GAPSLAM	RBPf-SOG	GAP. w/o reinit.
RMSE (cm)	639.6	<b>34.4 <math>\pm</math> 0.0</b>	56.0 $\pm$ 5.4	34.5 $\pm$ 0.0
Success rate (%)	-	<b>100</b>	<b>100</b>	58

sampling still supports an update frequency of at least 30 Hz. After time step 24, there are no non-Gaussian landmarks, so GAPSLAM is just the Gaussian solver with additional functions for sampling Gaussian marginals, which enjoys significantly faster speeds.

Fig. 3c presents the performance evaluation of point estimates and distribution estimation using the root mean square error (RMSE) and maximum mean discrepancy (MMD) [27], respectively. We compute the MMD between samples by GAPSLAM (or RBPf-SOG) and the reference samples by NSFG. The MMD is a distance between two distributions represented by samples, so a lower MMD here indicates a more accurate set of samples for representing posteriors. The results show that GAPSLAM significantly outperforms RBPf-SOG in terms of MMD, especially after time step 21 when the spurious modes in SOGs become more prominent, providing quantitative evidence of the superior accuracy of GAPSLAM in sampling posteriors. Additionally, it is worth noting that the RMSE of GAPSLAM only slightly outperforms that of RBPf-SOG at these early time steps. This indicates that MMD is a more suitable evaluation metric than RMSE for *full posterior inference*, especially in highly non-Gaussian settings.

3) *Results on point estimates*: Table I presents the errors of point estimates for the entire Plaza1 sequence. RBPf-SOG incurs less accurate and consistent results due to the particle depletion issue. GAPSLAM without re-initialization is essentially the Gaussian solver with additional priors on landmarks. However, only 58% of runs without re-initialization return solutions despite having RMSE values comparable to those of GAPSLAM. The superior performance of GAPSLAM supports the effectiveness of the sample-based, uncertainty-aware re-initialization.

## B. Object-based bearing-only SLAM

1) *Datasets and real-world experiments*: We estimate 6DOF camera poses and a map of object locations, using visual odometry and object detections from RGB videos. We aim to demonstrate the scalability of GAPSLAM to three-dimensional (3D) environments and its ability to fuse other types of measurements such as bearing-only. We test GAPSLAM using RGB data from the RGB-D Scenes Dataset v2 [28] and a video collected by a monocular camera in our office area. For visual odometry, we use camera poses of key frames computed by ORB-SLAM3 [29], while object classes and masks in the key frames are detected using the Detic detector [30]. Camera poses are optimized using both visual odometry and measurements to objects.

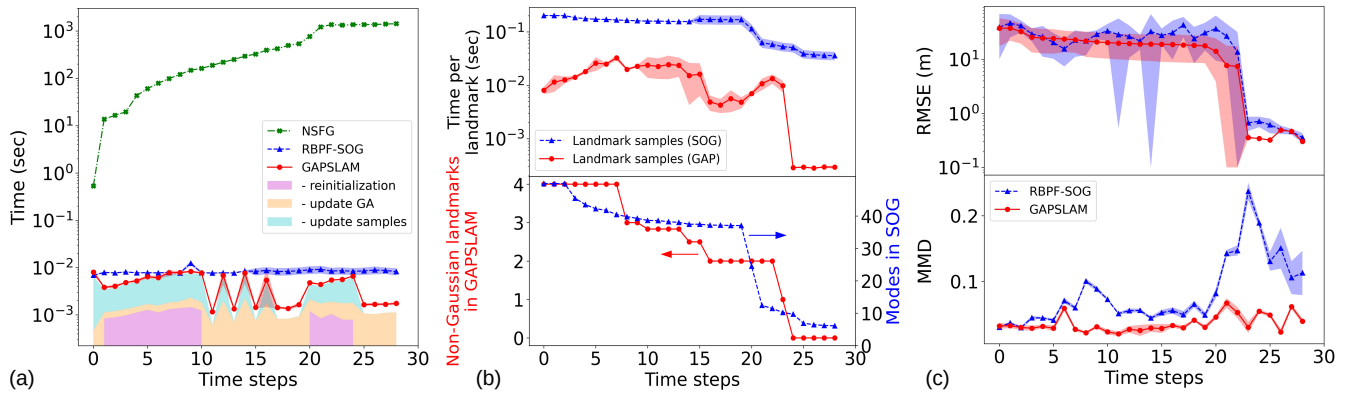


Fig. 3. Performance comparison at early time steps in the Plaza1: (a) computational time for updating probabilistic models, (b) computational time for posterior samples of landmarks, the number of non-Gaussian landmarks in GAPSLAM, and the averaged number of modes per landmark in RBPF-SOG, and (c) the root mean square error (RMSE) of point estimates and the maximum mean discrepancy (MMD) [27] of landmark samples. The MMD is a distance between two distributions represented by samples. We compute the MMD between the reference samples and samples by GAPSLAM (or RBPF-SOG), so a lower MMD means samples that describe the posterior better. Error bands indicate the mean and standard deviation across 6 runs with different random seeds.

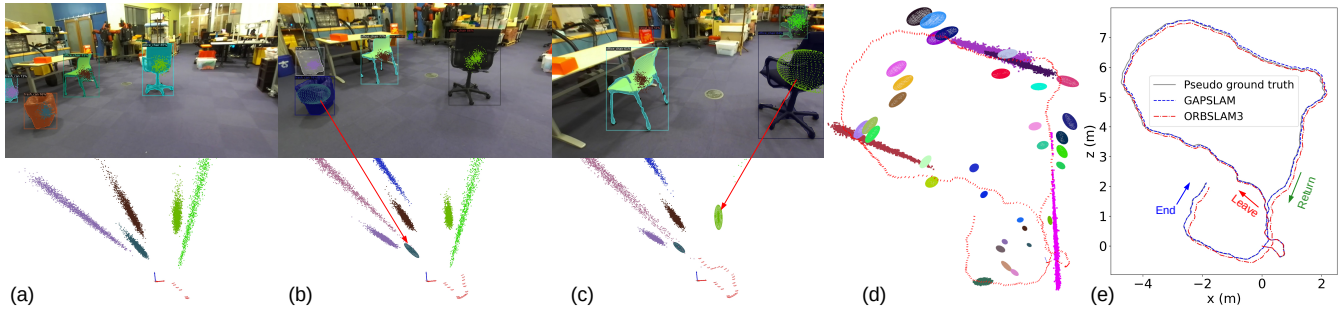


Fig. 4. Estimates made by GAPSLAM for the realworld RGB sequence: (a,b,c) samples or confidence intervals ( $3\sigma$ ) of object locations in 3D and their reprojection in images at key frames 10, 20, and 30, (d) final estimates of camera poses and samples and confidence intervals of object locations (we show samples of only 5 non-Gaussian landmarks to avoid clutter), and (e) a comparison of trajectory estimates with the pseudo ground truth. ORBSLAM3 did not close the big loop between ‘leave’ and ‘return’ due to the large change ( $\sim 180$  degrees) in viewing angles.

2) *Pre-processing object detections*: We treat the center  $z_k$  of an object mask on an image as a projection factor that models the reprojection error  $h(x_i, l_j; K) - z_k$ , where  $h(\cdot)$  denotes the reprojection of 3D object center  $l_j$  in the image with pose  $x_i$  and camera intrinsics  $K$ . This measurement  $z_k$  is either associated with an existing landmark or yields a new landmark, depending on the semantic and geometric information of the map. Thus wrong semantic labels may affect data association or spawn spurious landmarks. On the other hand, partial views of objects, which are often caused by occlusions or objects on image edges, lead to 2D mask centers that significantly deviate from the object center in 3D. To address these issues, we pre-process object detections as follows:

- Rejecting object detections that are potentially occluded by other detections on the image. Specifically, for any pair of detections (say  $i$  and  $j$ ) whose bounding boxes intersect, one of the detections (say  $i$ ) will be discarded if any of the following conditions is satisfied:
  - The object mask of detection  $i$  is smaller than a fraction (e.g., 10% in our experiments) of the mask of detection  $j$ . This condition is for rejecting very small masks.
  - The area of mask  $i$  within the bounding-box intersection is smaller than that of mask  $j$  and the area ratio between the intersection and the mask

$i$  is greater than a fraction (e.g., 10% in our experiments). This condition is for rejecting object detections that may be occluded by large bounding-box intersections.

- Increasing uncertainty in noise models of object mask centers that are close to image edges. Let us denote the relative position of the center of an object mask on an image as  $(r_x, r_y) \in [-1, 1]^2$  so  $r_x = 0$  and  $r_y = 0$  indicate a mask center located at the image center and  $|r_x|$  or  $|r_y|$  that approaches 1 indicates a mask center that is close to image edges. Thus we define standard deviations in the noise model as

$$\sigma_x = \delta_x(1 - |r_x|) + \max(\delta_x, \delta_y)|r_x|, \quad (20)$$

$$\sigma_y = \delta_y(1 - |r_y|) + \max(\delta_x, \delta_y)|r_y| \quad (21)$$

where  $\delta_x$  and  $\delta_y$  denote standard deviations of the mask in pixels.

- Rejecting any object detection which cannot be associated with landmarks in the same class via the maximum likelihood (ML) data association [31, 32] but passes the data association with landmarks that belong to different classes. This step intends discarding detections with wrong object classes. See the following section about the data association.

3) *Data association and creation of new landmarks*: We choose to implement the ML data association with some

additional treatment for exploiting semantic information in object detections and posterior samples in GAPSLAM. For a mask center  $z_k$  labeled by an object class and observed from pose  $x_i$ , we first compute its Mahalanobis distances to reprojections of any landmark  $l_j$  in the same class, as defined in

$$D_{kj}^2 = \|h(\hat{\mathbf{x}}) - z_k\|_C^2 \quad (22)$$

where  $\mathbf{x} = (x_i, l_j)$  and

$$C = \frac{\partial h}{\partial \mathbf{x}} \Big|_{\hat{\mathbf{x}}} \Sigma \frac{\partial h}{\partial \mathbf{x}} \Big|_{\hat{\mathbf{x}}}^T + \Gamma. \quad (23)$$

$\Sigma$  is the covariance between  $x_i$  and  $l_j$  in the Gaussian approximation of the posterior.  $\Gamma$  is the covariance of measurement noise model (e.g., (20) and (21)). The measurement is associated with  $l_j$  if  $D_{kj}^2$  is the smallest among landmarks in the same class and passes the chi-square test,  $D_{kj}^2 < \chi_{d,\alpha}^2$ , where  $d$  is the dimension of the measurement and  $\alpha$  is the desired confidence level.

Note that the quality of the Gaussian approximation affects the ML data association so, if the ML data association does not accept any landmark, we then perform another round of data association using landmark samples. For any landmark  $l_j$  in the same class as detection  $z_k$ , we reproject samples of  $l_j$  onto the image and count the percent  $p_{kj}$  of samples that fall in the bounding box of the detection. The detection is associated with  $l_j$  if  $p_{kj}$  is the highest among landmarks in the same class and higher than a threshold (e.g., 10% in our experiments).

If the sample-based data association does not accept any landmark either, we create a new landmark in the map.

4) *Results:* Fig. 4 shows qualitative results on our realworld sequence. We set the Detic detector to recognize only a certain number of classes: cup, cereal box, trash can, skateboard, office chair, football, bottle, traffic cone, and toy car since they are quite common in our office and we found Detic worked with good recall and precision for these classes. Figs. 4a, b, and c show how marginal posteriors of object locations evolve over key frames 10, 20, and 30, which occur before the ‘leave’ arrow on the path in Fig. 4e. The trash can in Fig. 4b and the chair in Fig. 4c have been identified as Gaussian landmarks by GAPSLAM so we draw confidence intervals ( $3\sigma$ ) of their locations using the Gaussian approximation in lieu of samples.

There are 415 camera poses and 100 landmarks being created in the end. Fig. 4d visualizes all camera poses and landmarks marginals in 3D. Fig. 4e compares the pseudo ground truth and the estimated trajectories by GAPSLAM and ORBSLAM3. ORBSLAM3 fails to close the big loop between ‘leave’ and ‘return’ since the camera returns from an opposite viewing angle, and the bag-of-word approach in ORBSLAM3 does not recognize the return. The trajectory by GAPSLAM is aligned with the pseudo ground truth much better since object detections introduce many loop closures along the path. We found object detections by Detic performed consistently well under big changes of viewing angles. In GAPSLAM, the big loop between ‘leave’ and ‘return’ is successfully closed when the camera returns to and recognizes chairs and trash cans that appeared at early key frames (Figs. 4a, b, and c). The pseudo ground truth is

TABLE II. Root mean square error (cm) of estimated paths under different settings. Each row below GAPSLAM serves as an ablation study that disables one (or three) of the functions in the object-based SLAM system.

Method	Sequences in RGB-D Scenes Dataset v2						Real world
	1	2	3	4	9	10	
ORBSLAM3	1.9	2.2	2.0	2.0	2.6	3.7	14.0
GAPSLAM	<b>0.9</b>	<b>1.0</b>	<b>0.7</b>	<b>1.2</b>	<b>0.7</b>	<b>2.6</b>	<b>5.2</b>
- no pre-processing	1.0	<b>1.0</b>	0.8	1.4	2.4	2.7	6.6
- no sample-based DA	<b>0.9</b>	<b>1.0</b>	<b>0.7</b>	<b>1.2</b>	0.8	<b>2.6</b>	5.8
- no reinitialization	<b>0.9</b>	<b>1.0</b>	0.8	1.3	<b>0.7</b>	<b>2.6</b>	6.6
- disable all above	1.0	2.2	0.9	1.3	4.2	2.7	9.5

TABLE III. Runtime profiling of the object-based SLAM system and specifics for different datasets.

Items	Sequences in RGB-D Scenes Dataset v2						Real world
	1	2	3	4	9	10	
Time/frame (ms)	43.2	28.6	29.7	32.1	23.9	19.7	63.3
- pre-processing	4.2	2.7	2.2	3.0	1.5	1.1	9.0
- data association	6.2	4.6	7.6	7.2	7.0	4.8	9.8
- GAPSLAM	32.8	21.3	21.2	21.9	15.4	13.9	44.4
- reinit.	3.2	2.4	1.6	1.4	1.0	1.3	6.1
- GA	8.8	6.1	7.6	8.8	4.8	3.7	13.9
- samples	13.6	8.0	6.8	6.2	4.7	5.5	17.1
Obj. det./frame (-)	5.1	5.0	5.2	5.0	2.8	3.0	5.3
Key frames (-)	99	95	108	109	96	87	415
# of landmarks (-)	8	7	8	8	4	3	100

the early portion of ORBSLAM3 results on a longer video, which circles our office three times along the same path, so the early portion has been corrected by loop closures introduced in the later portion.

Table II shows the RMSEs of estimated trajectories on the Scenes v2 sequences and our realworld sequence. GAPSLAM attains lowest errors in all the sequences. We attribute this accuracy to two reasons: 1) extra constraints introduced by object-based measurements, and 2) functions in the GAPSLAM algorithm. We disable some functions in the GAPSLAM algorithm to demonstrate their contributions to the estimation of the path. We find the pre-processing makes the biggest impact to the estimation while the re-initialization ranks the second and the sample-based data association ranks the third. Disabling all of them incurs significant estimation errors in some of the sequences (e.g., seqs. 9 and realworld).

Fig. 5 shows the runtime of our system for the realworld sequence. The system implemented in Python supports an update frequency of 6 Hz for the slowest key frames and, on average, affords an update frequency of 16 Hz. We are confident that the runtime can be greatly optimized by an improved implementation. Given that we only process key frames, GAPSLAM is able to support real-time operation. Table III shows the runtime profiling of our system and specifics of different datasets. The runtime on Scenes v2 sequences is expected to be faster since the sequences involve fewer poses and objects. The crux of averaged runtime is on steps for updating samples and the Gaussian approximation. However, Fig. 5 shows that updating samples only dominates the runtime in early key frames, when many non-Gaussian landmarks had just been created. Revisiting objects at later key frames reduces the uncertainty in the posteriors of object locations, so more and more objects are moved to the set of Gaussian landmarks. Thus the computation burden gradually transfers to updating the Gaussian approximation.

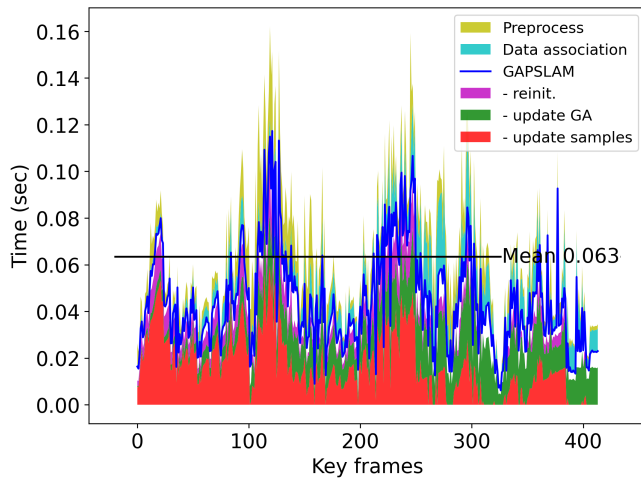


Fig. 5. Runtime profiling of the object-based SLAM system (in Python).

The scalability of updating the Gaussian approximation can be improved via many strategies such as fix-lag smoothing [33].

## V. CONCLUSION

We presented a real-time algorithm, GAPSLAM, that precisely infers non-Gaussian/multi-modal marginal posteriors encountered in SLAM. Our experiments justified the efficacy of the adaptive modeling strategy for efficient full posterior inference, and the uncertainty-aware re-initialization technique for explicitly correcting linearization points in non-linear optimization solvers. Future work includes inferring the joint posterior of multiple variables, incorporating more complex likelihood of measurements (e.g., learned orientation distributions [34], multi-modal data association [35]), and applying GAPSLAM to realworld planning tasks that tackle the non-Gaussian belief of obstacle locations [36].

## REFERENCES

- [1] D. M. Rosen, K. J. Doherty, A. T. Espinoza, and J. J. Leonard, "Advances in inference and representation for simultaneous localization and mapping," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 4, no. 1, pp. 215–242, Jan. 2021.
- [2] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," Georgia Institute of Technology, Tech. Rep., 2012.
- [3] M. Montemero, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM 2.0: An improved particle filtering algorithm for simultaneous localization and mapping that provably converges," in *Proc. Int. Joint Conf. Artif. Intell.*, 2003, pp. 1151–1156.
- [4] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Int. J. Robot. Res.*, vol. 31, no. 2, pp. 216–235, 2012.
- [5] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 3607–3613.
- [6] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, 2002.
- [7] J. Blanco, J. González, and J.-A. Fernández-Madrigal, "A pure probabilistic approach to range-only SLAM," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2008, pp. 1436–1441.
- [8] J. Blanco, J.-A. Fernández-Madrigal, and J. González, "Efficient probabilistic range-only SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.* IEEE, 2008, pp. 1017–1022.
- [9] D. Fourie, J. J. Leonard, and M. Kaess, "A nonparametric belief solution to the Bayes tree," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.* IEEE, 2016, pp. 2189–2196.
- [10] Q. Huang *et al.*, "Incremental non-Gaussian inference for SLAM using normalizing flows," *IEEE Trans. Robot.*, pp. 1–18, 2022.

- [11] P. Torma, A. György, and C. Szepesvári, "A Markov-chain Monte Carlo approach to simultaneous localization and mapping," in *Proc. Int. Conf. Artif. Intell. Stat.*, 2010, pp. 852–859.
- [12] R. Shariff, A. György, and C. Szepesvári, "Exploiting symmetries to construct efficient MCMC algorithms with an application to SLAM," in *Artificial Intelligence and Statistics*. PMLR, 2015, pp. 866–874.
- [13] Q. Huang, A. Papalía, and J. J. Leonard, "Nested sampling for non-Gaussian inference in SLAM factor graphs," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 9232–9239, Oct. 2022.
- [14] A. J. Davison, "Real-time simultaneous localisation and mapping with a single camera," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 3, 2003, pp. 1403–1403.
- [15] P. Newman and J. Leonard, "Pure range-only sub-sea SLAM," in *Proc. IEEE Int. Conf. Robot. Autom.*, vol. 2, 2003, pp. 1921–1926.
- [16] T. Lemaire, S. Lacroix, and J. Sola, "A practical 3D bearing-only SLAM algorithm," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2005, pp. 2449–2454.
- [17] N. Kwok and G. Dissanayake, "An efficient multiple hypothesis filter for bearing-only SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2004, pp. 736–741.
- [18] J. Sola, A. Monin, M. Devy, and T. Lemaire, "Undelayed initialization in bearing only SLAM," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2005, pp. 2499–2504.
- [19] L. Nicholson, M. Milford, and N. Sünderhauf, "QuadricSLAM: Dual quadrics from object detections as landmarks in object-oriented SLAM," *IEEE Robot. Autom. Lett.*, vol. 4, no. 1, pp. 1–8, 2018.
- [20] S. Yang and S. Scherer, "CubeSLAM: Monocular 3-D object SLAM," *IEEE Trans. Robot.*, vol. 35, no. 4, pp. 925–938, 2019.
- [21] K. Ok, K. Liu, K. Frey, J. P. How, and N. Roy, "Robust object-based SLAM for high-speed autonomous navigation," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 669–675.
- [22] C. Rubino, M. Crocco, and A. Del Bue, "3D object localisation from multi-view image detections," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 6, pp. 1281–1294, 2017.
- [23] C. M. Bishop, *Pattern recognition and machine learning*. New York, USA: Springer, 2006, ch. 4, p. 213.
- [24] J. Djugash, B. Hammer, and S. Roth, "Navigating with ranging radios: Five data sets with ground truth," *J. Field Robot.*, vol. 26, no. 9, pp. 689–695, 2009.
- [25] J. Skilling, "Nested sampling for general Bayesian computation," *Bayesian Anal.*, vol. 1, no. 4, pp. 833–859, 2006.
- [26] J. S. Speagle, "dynesty: a dynamic nested sampling package for estimating Bayesian posteriors and evidences," *Monthly Notices Roy. Astronomical Soc.*, vol. 493, no. 3, pp. 3132–3158, 2020.
- [27] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 723–773, 2012.
- [28] K. Lai, L. Bo, and D. Fox, "Unsupervised feature learning for 3D scene labeling," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 3050–3057.
- [29] C. Campos, R. Elvira, J. J. Gómez, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM3: An accurate open-source library for visual, visual-inertial and multi-map SLAM," *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [30] X. Zhou, R. Girdhar, A. Joulin, P. Krähenbühl, and I. Misra, "Detecting twenty-thousand classes using image-level supervision," in *Proc. Eur. Conf. Comput. Vis.*, 2022, pp. 350–368.
- [31] M. Kaess and F. Dellaert, "Covariance recovery from a square root information matrix for data association," *Robotics and autonomous systems*, vol. 57, no. 12, pp. 1198–1210, 2009.
- [32] J. Neira and J. D. Tardós, "Data association in stochastic mapping using the joint compatibility test," *IEEE Trans. Robot. Autom.*, vol. 17, no. 6, pp. 890–897, 2001.
- [33] F. Dellaert and M. Kaess, "Factor graphs for robot perception," *Found. Trends Robot.*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [34] K. A. Murphy, C. Esteves, V. Jampani, S. Ramalingam, and A. Makadia, "Implicit-PDF: Non-parametric representation of probability distributions on the rotation manifold," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 7882–7893.
- [35] K. Doherty, D. Fourie, and J. J. Leonard, "Multimodal semantic slam with probabilistic data association," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 2419–2425.
- [36] W. Han, A. Jasour, and B. Williams, "Non-gaussian risk bounded trajectory optimization for stochastic nonlinear systems in uncertain environments," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022, pp. 11 044–11 050.