

Gaussian Variational Inference With Non-Gaussian Factors for State Estimation: A UWB Localization Case Study

Andrew Stirling¹, Mykola Lukashchuk¹, Dmitry Bagaev¹, Wouter Kouw², *Member, IEEE*,
and James R. Forbes¹, *Member, IEEE*

Abstract—This letter extends the exactly sparse Gaussian variational inference (ESGVI) algorithm for state estimation in two complementary directions. First, ESGVI is generalized to operate on matrix Lie groups, enabling the estimation of states with orientation components while respecting the underlying group structure. Second, factors are introduced to accommodate heavy-tailed and skewed noise distributions, as commonly encountered in ultra-wideband (UWB) localization due to non-line-of-sight (NLOS) and multipath effects. Both extensions are shown to integrate naturally within the ESGVI framework while preserving its sparse and derivative-free structure. The proposed approach is validated in a UWB localization experiment with NLOS-rich measurements, demonstrating improved accuracy and comparable consistency. Finally, a PYTHON implementation within a factor-graph-based estimation framework is made open-source to support broader research use.

Index Terms—Localization, probabilistic inference, sensor fusion, range sensing.

I. INTRODUCTION

AUTONOMOUS robots must estimate their position, orientation, and other key pieces of information to operate effectively. This collection of dynamic variables, known as the state of the robot, fully describes the system at a given time. Consistent and accurate state estimation is essential for high-level decision-making and the support of downstream tasks such as guidance and control. Within state estimation, approaches to localization can be broadly split into the filtering and batch approaches [1]. Filtering methods use data up to the current timestep to estimate the robot's state, whereas batch methods use all available data over a given time period. The latter forms the foundation of this work. Taking a Bayesian perspective, the

goal of state estimation is to compute the state $\mathbf{x} \in \mathbb{R}^{n_x}$, given a set of measurements $\mathbf{y} \in \mathbb{R}^{n_y}$, that is, the full posterior $p(\mathbf{x}|\mathbf{y})$ is computed, given a prior $p(\mathbf{x})$. Using Bayes' rule, the relationship

$$p(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} = \frac{p(\mathbf{x}, \mathbf{y})}{p(\mathbf{y})}, \quad (1)$$

is exploited to compute the posterior. For nonlinear measurement models $p(\mathbf{y}|\mathbf{x})$, the posterior is *not* a Gaussian probability density function (PDF). Currently, the majority of the batch state estimation literature is based on solving a *maximum-a-posteriori* (MAP) problem, which involves computing the maximum, that is the mode, of the Bayesian posterior. An alternative to MAP is variational inference [2], which approximates the posterior by fitting a distribution that minimizes the Kullback–Leibler (KL) divergence. In [3], exactly sparse Gaussian variational inference (ESGVI) was introduced as a tractable method for computing Gaussian approximations of the posterior in large-scale nonlinear batch state estimation problems, by exploiting the sparsity associated with robotics problems. By accounting for the full shape of the posterior distribution, ESGVI was shown to outperform MAP in nonlinear state estimation problems. However, in its current form, ESGVI operates on vector-valued states and, although the use of a robust cost function was suggested as a future line of inquiry in [3], it has only been evaluated with Gaussian-distributed sensor models. In practice, sensor measurements can deviate significantly from Gaussian assumptions, producing outliers that can severely degrade estimation performance. This issue is particularly evident in ultra-wideband (UWB) localization, where ranging measurements are increasingly used in robotics due to their long-range and low-cost characteristics [4]. Nevertheless, UWB measurements are often corrupted by non-line-of-sight (NLOS) effects and multipath propagation, leading to heavy-tailed and asymmetric noise distributions that must be accounted for [5]. To address these limitations, ESGVI is extended to operate on matrix Lie group (MLG) states, a modest but necessary step, as MLGs naturally capture the geometry of states commonly found in robot navigation problems [6]. In addition, non-Gaussian measurement factors are incorporated, providing a more flexible way to model heavy-tailed and skewed noise distributions. While such factors have been explored in related estimation frameworks [7], the contribution here lies in demonstrating their integration within the ESGVI setting and evaluating their impact in practice. Although the proposed extensions are general, their effectiveness is illustrated through a specific UWB example for concreteness and ease of interpretation. Building on [3], the contributions of this letter are as follows:

Received 31 August 2025; accepted 9 December 2025. Date of publication 12 January 2026; date of current version 26 January 2026. This article was recommended for publication by Associate Editor D. Belter and Editor S. Behnke upon evaluation of the reviewers' comments. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) Alliance Grant program, in part by Denso Corporation, in part by ARA Robotics, and in part by NSERC-FRQ NOVA Program. (*Corresponding author: Andrew Stirling.*)

Andrew Stirling and James R. Forbes are with the Department of Mechanical Engineering, McGill University, Montreal, QC H3A 0G4, Canada (e-mail: andrew.stirling2@mail.mcgill.ca; james.richard.forbes@mcgill.ca).

Mykola Lukashchuk, Dmitry Bagaev, and Wouter Kouw are with the Department of Electrical Engineering, TU Eindhoven, 5612 AZ Eindhoven, The Netherlands (e-mail: m.lukashchuk@tue.nl; d.v.bagaev@tue.nl; w.m.kouw@tue.nl).

Digital Object Identifier 10.1109/LRA.2026.3653370

- Extending ESGVI to support state estimation on matrix Lie groups.
- Introducing non-Gaussian measurement factors into the ESGVI framework, tailored for asymmetric and skewed noise distributions.
- Evaluating ESGVI in a real-world, outlier rich, UWB localization experiment, demonstrating greater accuracy and comparable consistency.
- Releasing an open-source PYTHON implementation of ESGVI within a factor-graph-based estimation framework at https://github.com/decargroup/gvi_ws.

In this letter bold uppercase letters such as \mathbf{A} denote matrices, bold lowercase letters $\mathbf{a} \in \mathbb{R}^n$ denote column vectors, while lowercase letters $a \in \mathbb{R}$ represent scalar values.

II. PROBLEM STATEMENT

Consider a robot in an indoor environment estimating its state using UWB range measurements y_r , from a set of fixed UWB anchors positioned throughout the room. When the signal path is unobstructed and in line-of-sight (LOS), the range error e_r , defined as the difference between the measured and true distance, is typically small. In contrast, when obstacles block the direct path, creating non-line-of-sight (NLOS) conditions, signals may reflect from surrounding surfaces, or even pass through obstacles. These multipath and NLOS measurements bias the ranges resulting in a skewed error distribution, as shown in Fig. 1.

A second challenge arises from the fact that some state components, such as position, admit a natural Euclidean representation, whereas others are inherently directional. For example, the vehicle's orientation is periodic, with 0 (rad) and 2π (rad) denoting the same heading. Treating rotational variables in Euclidean space \mathbb{R}^{n_x} neglects their non-commutative group structure [8], leading to singularities and distorted uncertainty representations [9]. To address this, the vehicle state is represented as a matrix Lie group (MLG), that is, $\mathbf{X}_k \in SE(2)$, which combines both translational and rotational components as [1], [6],

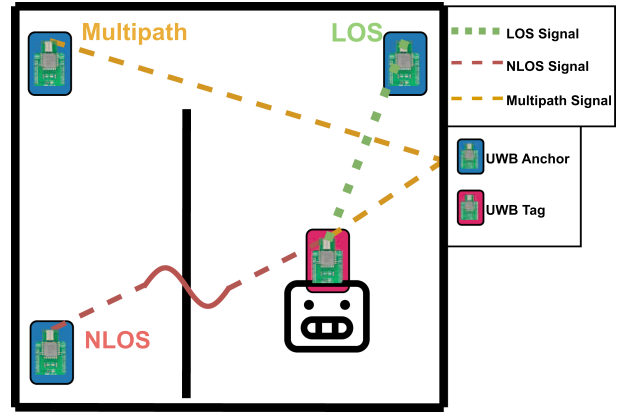
$$\mathbf{X}_k = \begin{bmatrix} \mathbf{C}_{ab} & \mathbf{r}_a^{zw} \\ \mathbf{0} & 1 \end{bmatrix} \in SE(2), \quad (2)$$

where $\mathbf{C}_{ab} \in SO(2)$ is the rotation matrix mapping the body frame \mathcal{F}_b to the absolute frame \mathcal{F}_a . The notation \mathbf{r}_a^{zw} represents the position of the point z on the robot relative to the origin w of the absolute frame, resolved in the frame \mathcal{F}_a . A broader overview of the required MLG operations used throughout this work is provided in Appendix A. By representing the state as a MLG, the underlying manifold structure is respected, and directional uncertainty is properly captured. The objective is to leverage this representation to estimate the robot state as accurately as possible under skewed measurement noise conditions. Standard estimation approaches assume Gaussian sensor likelihoods and thus struggle in this setting, often leading to bias or even numerical instability. A robust solution must instead incorporate skewed likelihood functions while still producing a Gaussian distribution for the state estimate.


III. BACKGROUND AND RELATED WORK

A. Probabilistic Model Specification

This work focuses on state-space models, that is, models that contain explicit dynamics and observation functions. For



(a) Illustration of LOS, NLOS and multipath signal propagation between fixed UWB anchors and a mobile UWB tag attached to the robot.



(b) Simulated noise distribution modeling NLOS and multipath effects, where 25% of Gaussian range measurements are corrupted with values drawn uniformly between 1 and 6 standard deviations. Various probabilistic models are fitted and compared for estimation performance in Section V-A.

Fig. 1. Depiction of obstacle-induced signal propagation effects that lead to errors in range-based localization.

simplicity, consider the vector-valued state \mathbf{x}_k at time index $k = 0$, with an initial state distribution of

$$p(\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_0 | \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0). \quad (3)$$

While attention is restricted to vector states here for clarity, in later sections this framework will be generalized to accommodate the MLG state, \mathbf{X}_k . The state evolves over time according to a stochastic state transition with the Markov property, i.e., only dependent on the previous state [10, Property 4.1]. It is modelled as a Gaussian distribution,

$$p(\mathbf{x}_k | \mathbf{x}_{k-1}, \mathbf{u}_{k-1}) = \mathcal{N}(\mathbf{x}_k | \mathbf{f}(\mathbf{x}_{k-1}, \mathbf{u}_{k-1}), \mathbf{Q}_k), \quad (4)$$

where \mathbf{u}_{k-1} represents the input applied at the previous step, $\mathbf{f}(\cdot)$ is the nonlinear dynamics function, and \mathbf{Q}_k is the discrete-time process noise covariance. We implicitly assume that stochastic perturbations to the state are independent over time. For the sensor measurement likelihood, non-Gaussian likelihoods can be incorporated when the noise characteristics are known. In particular, for range measurements y_k , the likelihood is modeled

using a Skew-Laplace distribution of the form [7],

$$p(y_k|\mathbf{x}_k) = \mathcal{SL}(y_k|g(\mathbf{x}_k), \sigma, \lambda), \quad (5)$$

$$= \frac{1}{2\sigma\alpha} \exp\left(\lambda\left(\frac{y_k - g(\mathbf{x}_k)}{\sigma^2}\right) - \alpha\left|\frac{y_k - g(\mathbf{x}_k)}{\sigma}\right|\right), \quad (6)$$

where $g(\cdot)$ denotes the predicted range as a function of the state, and the parameters, $\sigma \in \mathbb{R}^+$, $\lambda \in \mathbb{R}$, denoting the distribution scale, and skewness, where $\alpha = \sqrt{1 + (\lambda/\sigma)^2}$. The parameters are estimated separately using maximum likelihood. Combining the prior, process, and measurement models together, the probabilistic model of the discrete-time series of $k = 1, \dots, K$ is,

$$p(\mathbf{y}, \mathbf{x}|\mathbf{u}) = p(\mathbf{x}_0) \prod_{k=1}^K p(\mathbf{x}_k|\mathbf{x}_{k-1}, \mathbf{u}_{k-1}) p(y_k|\mathbf{x}_k). \quad (7)$$

The next section discusses the proposed method for obtaining the marginal posterior over the trajectory $\mathbf{x} = \{\mathbf{x}_0, \dots, \mathbf{x}_K\}$.

B. State of the Art

A common approach to batch state estimation is to assume a Gaussian measurement model and find the most probable state vector under the state posterior distribution (i.e., MAP estimation),

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}, \mathbf{u}). \quad (8)$$

By assuming Gaussian noise and minimizing the negative logarithm of the posterior, the MAP estimate can be formulated as a nonlinear least-squares (NLS) problem [1, Section 3.1.2]. Expressed in this sense, the estimator minimizes a sum-of-squared-errors cost function, which makes it highly sensitive to outliers [1, §5.4]. To address the sensitivity of MAP-based NLS estimators to outliers, a variety of strategies have been explored, particularly in UWB-based state estimation. For example, Omni-Swarm [11] introduced a visual-inertial-UWB estimation pipeline formulated as a MAP-based NLS optimization problem, where UWB outliers were rejected using fixed elevation-angle criteria. Distributed approaches [12], [13], [14] extend this concept to multi-agent settings, also solving MAP NLS problems while rejecting inconsistent range measurements through averaging, geometric or max velocity constraints. More recently, learning-based methods have been employed to detect and discard NLOS UWB measurements. In [15] and [16], convolutional neural network (CNN) and long short-term memory (LSTM) architectures were proposed to classify and reject suspected NLOS observations. However, since no model can guarantee perfect outlier identification, residual contamination typically persists in the dataset, introducing skewed measurement noise.

To more systematically account for these residual effects, M-estimators incorporate robust cost functions directly within the NLS formulation, allowing the optimization to down-weight the influence of large residuals rather than reject measurements outright. Various robust loss functions have been proposed [17], and in the context of UWB localization, several noise parameterizations have been studied. For example, in [18] a Huber loss was integrated into the NLS optimization, while [4] proposed a Cauchy loss, however these symmetric losses ignore the asymmetry present in NLOS noise. To capture skewed characteristics, [5], [19] used an asymmetric two-piece distribution. A similar asymmetric Cauchy distribution is used in this letter

defined by

$$p(y_k|\mathbf{x}_k) = \begin{cases} \alpha \frac{1}{1 + \left(\frac{r(y_k, \mathbf{x}_k)}{c^-}\right)^2}, & r(\cdot, \cdot) < 0, \\ \alpha \frac{1}{1 + \left(\frac{r(y_k, \mathbf{x}_k)}{c^+}\right)^2}, & r(\cdot, \cdot) \geq 0, \end{cases} \quad (9)$$

where $r(\cdot, \cdot)$ is the residual, c^- and c^+ are scale parameters, and $\alpha = 2/[\pi(c^- + c^+)]$ normalizes the distribution. The asymmetric Cauchy loss provides outlier robustness while explicitly modeling skewed measurement noise. A secondary flexible alternative for modeling skewed distributions is to employ a Gaussian mixture measurement model [20],

$$p(y_k|\mathbf{x}_k) = \sum_{j=1} w_j \mathcal{N}(r(y_k, \mathbf{x}_k)|\psi_j), \quad (10)$$

with weights $w_j \geq 0$, $\sum_j w_j = 1$ and $\psi_j = (\mu_j, \Sigma_j)$ [21]. Component weights and parameters are calibrated separately, or solved for in parallel with the state \mathbf{x} , in an Expectation Maximization (EM) framework [2, §9.3]. The Gaussian mixture measurement model of (10) is incorporated into the joint distribution of (7), and MAP estimation (8) produces a state estimate. Although Gaussian mixture models (GMMs) are flexible enough to capture skewed noise distributions, they require additional parameters and, while incorporation into a NLS framework is possible, it demands careful handling [21]. Moreover, a MAP estimator provides only a point estimate of the state, with uncertainty obtained using an ad-hoc Laplace approximation [1, §4.3.1]. Alternatively, while the methods discussed thus far rely on MAP-based optimization, ESGVI offers a fundamentally different probabilistic framework that estimates both the state and its uncertainty directly through variational inference. Performing parameter learning in parallel with ESGVI [22] has been shown to improve robustness, as jointly estimating the covariance with an appropriate prior implicitly rejects outliers. While robustness was identified as an open direction in [3], measurement models specific to a non-Gaussian distribution, such as the Skew-Laplace in (6) have not yet been studied directly within the ESGVI framework. Consequently, in this work, comparisons are made against robust MAP approaches rather than existing ESGVI methods. In the following, ESGVI is shown to naturally incorporate robust, non-Gaussian measurement models with matrix Lie group states, enabling principled estimation under heavy-tailed or asymmetric noise. This is illustrated using a Skew-Laplace model for UWB localization, capturing NLOS-induced asymmetry.

IV. GAUSSIAN VARIATIONAL INFERENCE

As introduced previously, an alternative to MAP estimation is Gaussian variational inference (GVI), which fits a Gaussian approximation $q(\mathbf{x}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, to the posterior $p(\mathbf{x}|\mathbf{y}, \mathbf{u})$, by minimizing the Kullback-Leibler (KL) divergence [2],

$$\text{KL}(q||p) = \mathbb{E}_q[\ln q(\mathbf{x}) - \ln p(\mathbf{x}|\mathbf{y}, \mathbf{u})]. \quad (11)$$

Expanding and dropping terms independent of $q(\mathbf{x})$, defines the loss functional

$$V(q) = \mathbb{E}_q[\phi(\mathbf{x})] + \frac{1}{2} \ln(|\boldsymbol{\Sigma}^{-1}|), \quad (12)$$

where $\phi(\mathbf{x}) = -\ln p(\mathbf{x}, \mathbf{y}, \mathbf{u})$ is the negative log joint state and measurement likelihood, and $|\boldsymbol{\Sigma}^{-1}|$ is the determinant of the information matrix of the variational approximation, $q(\mathbf{x})$. In the work of [3], an exactly sparse Gaussian variational inference (ESGVI) optimization scheme was defined for the vector state

\mathbf{x} , minimizing $V(q)$. The iterative update to the mean $\boldsymbol{\mu}$, and information matrix $\boldsymbol{\Sigma}^{-1}$, are defined according to

$$(\boldsymbol{\Sigma}^{-1})^{(i+1)} = \mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^\top \partial \mathbf{x}} \right], \quad (13a)$$

$$(\boldsymbol{\Sigma}^{-1})^{(i+1)} \delta \boldsymbol{\mu} = -\mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^\top} \right], \quad (13b)$$

$$\boldsymbol{\mu}^{(i+1)} = \boldsymbol{\mu}^{(i)} + \delta \boldsymbol{\mu}, \quad (13c)$$

where $\delta \boldsymbol{\mu}$ is computed by solving (13b). By making use of the factoring properties of the probabilistic model of (7), the expectations can be expressed over the marginal $q_k(\mathbf{x}_k)$. Defining the projection, $\mathbf{x}_k = \mathbf{P}_k \mathbf{x}$, the expectations become,

$$\mathbb{E}_q [\phi(\mathbf{x})] = \sum_{k=1}^K \mathbb{E}_{q_k} [\phi_k(\mathbf{x}_k)], \quad (14a)$$

$$\mathbb{E}_q \left[\frac{\partial \phi(\mathbf{x})}{\partial \mathbf{x}^\top} \right] = \sum_{k=1}^K \mathbf{P}_k^\top \mathbb{E}_{q_k} \left[\frac{\partial}{\partial \mathbf{x}_k^\top} \phi_k(\mathbf{x}_k) \right], \quad (14b)$$

$$\mathbb{E}_q \left[\frac{\partial^2 \phi(\mathbf{x})}{\partial \mathbf{x}^\top \partial \mathbf{x}} \right] = \sum_{k=1}^K \mathbf{P}_k^\top \mathbb{E}_{q_k} \left[\frac{\partial^2}{\partial \mathbf{x}_k^\top \partial \mathbf{x}_k} \phi_k(\mathbf{x}_k) \right] \mathbf{P}_k, \quad (14c)$$

which reduces the computation to smaller, factor-specific subsets. Next, using Stein's lemma [1, Section 2.2.16], it was demonstrated how taking derivatives of $\phi_k(\cdot)$ can be avoided [3],

$$\mathbb{E}_{q_k} \left[\frac{\partial}{\partial \mathbf{x}_k^\top} \phi_k(\mathbf{x}_k) \right] = \boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E}_{q_k} [(\mathbf{x}_k - \boldsymbol{\mu}_k) \phi_k(\mathbf{x}_k)], \quad (15a)$$

$$\begin{aligned} \mathbb{E}_{q_k} \left[\frac{\partial^2}{\partial \mathbf{x}_k^\top \partial \mathbf{x}_k} \phi_k(\mathbf{x}_k) \right] &= -\boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E}_{q_k} [\phi_k(\mathbf{x}_k)] \\ &+ \boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E}_{q_k} [(\mathbf{x}_k - \boldsymbol{\mu}_k)(\mathbf{x}_k - \boldsymbol{\mu}_k)^\top \phi_k(\mathbf{x}_k)] \boldsymbol{\Sigma}_{kk}^{-1}. \end{aligned} \quad (15b)$$

Lastly, these marginal expectations can be approximated using a cubature method [10, §6.3], that is

$$\mathbb{E}_{q_k} [\phi_k(\mathbf{x}_k)] \approx \sum_{l=1}^L w_{k,l} \phi_k(\mathbf{x}_{k,l}), \quad (16)$$

where $w_{k,l}$ are the weights, and $\mathbf{x}_{k,l}$ are the sigma points. The remaining expectations in (15a) and (15b) are approximated in the same manner.

A. ESGVI With Matrix Lie Group States

Up to this point, the discussion has been restricted to vector-valued states \mathbf{x}_k , for simplicity. However, as introduced previously, some robotics states are more naturally expressed as MLG states \mathbf{X}_k . As such the variational algorithm (13) must be reformulated to work with MLG states. To do so, uncertainty must be defined in a way that is compatible with the underlying group structure [8]. A Gaussian distribution associated with a MLG state is defined by perturbing the mean element $\bar{\mathbf{X}}$ with a zero-mean Gaussian random variable $\delta \boldsymbol{\xi}$. Formally,

$$\mathbf{X} = \bar{\mathbf{X}} \oplus \delta \boldsymbol{\xi}, \quad \delta \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad (17)$$

where the \oplus notation is borrowed from [6]. A concise overview of the required MLG operations, including the definitions of \oplus and its inverse \ominus , can be found in Appendix A. Having defined distributions associated with MLGs, the computational structure of the ESGVI algorithm can be reformulated for MLG-valued

states. Although Stein's lemma was originally used in the Euclidean setting to rewrite the gradients of (15a), and (15b) in terms of expectations, its application remains valid when considering a perturbation $\delta \boldsymbol{\xi}$ relative to $\bar{\mathbf{X}}$, as in (17). Specifically, computations are performed relative to the mean state $\bar{\mathbf{X}}_k$ using $\delta \boldsymbol{\xi}_k$, which is Euclidean and therefore admits the same calculus rules as the standard vector space. In this way, the information update in (13a) and the gradient computation in (13b) proceed as before. The marginal MLG normal distribution parametrizes uncertainty through $\delta \boldsymbol{\xi}_k$, meaning $q_k(\delta \boldsymbol{\xi}_k)$, preserving the expectation formulation. Additionally, the Euclidean difference $(\mathbf{x}_k - \boldsymbol{\mu}_k)$, is replaced with $(\mathbf{X}_k \ominus \bar{\mathbf{X}}_k)$. The expectations in (15a), and (15b) are then

$$\mathbb{E}_{q_k} \left[\frac{\partial \phi_k(\mathbf{X}_k)}{\partial \delta \boldsymbol{\xi}_k^\top} \right] = \boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E}_{q_k} [(\mathbf{X}_k \ominus \bar{\mathbf{X}}_k) \phi_k(\mathbf{X}_k)], \quad (18a)$$

$$\begin{aligned} \mathbb{E}_{q_k} \left[\frac{\partial^2 \phi_k(\mathbf{X}_k)}{\partial \delta \boldsymbol{\xi}_k^\top \partial \delta \boldsymbol{\xi}_k} \right] &= -\boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E}_{q_k} [\phi_k(\mathbf{X}_k)] \\ &+ \boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E}_{q_k} [(\mathbf{X}_k \ominus \bar{\mathbf{X}}_k)(\mathbf{X}_k \ominus \bar{\mathbf{X}}_k)^\top \phi_k(\mathbf{X}_k)] \boldsymbol{\Sigma}_{kk}^{-1}. \end{aligned} \quad (18b)$$

Although the expressions appear structurally similar to their Euclidean counterparts, their interpretation on MLGs requires further clarification, provided in Appendix A. Numerically, the expectations of (18) are evaluated using a cubature method [10, §6.3] such that

$$\mathbb{E}_{q_k} [\phi_k(\mathbf{X}_k)] \approx \sum_{\ell=1}^L w_k^\ell \phi_k(\boldsymbol{\chi}_k^\ell), \quad (19a)$$

$$\mathbb{E}_{q_k} [(\mathbf{X}_k \ominus \bar{\mathbf{X}}_k) \phi_k(\mathbf{X}_k)] \approx \sum_{\ell=1}^L w_k^\ell (\boldsymbol{\chi}_k^\ell \ominus \bar{\mathbf{X}}_k) \phi_k(\boldsymbol{\chi}_k^\ell), \quad (19b)$$

where the matrix-valued expectation associated with (18b) is omitted for brevity. The group sigma points $\boldsymbol{\chi}_k^\ell$ are computed via (17) as

$$\boldsymbol{\chi}_k^\ell = \bar{\mathbf{X}}_k \oplus \delta \boldsymbol{\xi}_k^\ell, \quad (20)$$

where the perturbations are constructed as $\delta \boldsymbol{\xi}_k^\ell = \sqrt{\boldsymbol{\Sigma}_{kk}} \boldsymbol{\alpha}^\ell$. Here, the unit sigma points $\boldsymbol{\alpha}^\ell$, and their corresponding weights w^ℓ , are generated by the same cubature method used in the Euclidean case of (16). Now that (18a), and (18b) have been generalized for the MLG state, the reformulation of (13c) follows naturally from (17). Instead of updating the Euclidean mean with a vector addition, the MLG mean $\bar{\mathbf{X}}^{(i)}$, is updated by mapping $\delta \boldsymbol{\mu}$ back to the group using the exponential map,

$$\bar{\mathbf{X}}^{(i+1)} = \bar{\mathbf{X}}^{(i)} \oplus \delta \boldsymbol{\mu}. \quad (21)$$

Here, the \oplus notation is slightly abused, as each component $\bar{\mathbf{X}}_k^{(i)}$ is updated by the projected increment $\mathbf{P}_k \delta \boldsymbol{\mu}$, rather than the batch vector. This update ensures the new mean, $\bar{\mathbf{X}}_k^{(i+1)}$ remains in the group, while incorporating the computed variational variables, $\boldsymbol{\Sigma}_{kk}^{-1}$ and $\delta \boldsymbol{\mu}_k$. Recall from the earlier discussion in Section III-A that the assumed form of $\phi_{k,\text{meas}}(\cdot)$ was based on the Skew-Laplace density function,

$$\begin{aligned} \phi_{k,\text{meas}}(\mathbf{X}_k) &= -\ln p(y_k | \mathbf{X}_k), \\ &= -\left(\lambda \frac{y_k - g(\mathbf{X}_k)}{\sigma^2} - \alpha \left| \frac{y_k - g(\mathbf{X}_k)}{\sigma} \right| \right), \end{aligned} \quad (22)$$

where normalization constants are safely dropped due to the optimization requiring only derivatives of $\phi_k(\mathbf{X}_k)$. Similarly,

the Gaussian process model is represented as,

$$\begin{aligned} \phi_{k,\text{proc}}(\mathbf{X}_k) &= -\ln p(\mathbf{X}_k | \mathbf{X}_{k-1}, \mathbf{u}_{k-1}), \\ &= \frac{1}{2}(\mathbf{X}_k \ominus \mathbf{f}(\mathbf{X}_{k-1}, \mathbf{u}_{k-1}))^\top \mathbf{Q}_{k-1}^{-1}(\mathbf{X}_k \ominus \mathbf{f}(\mathbf{X}_{k-1}, \mathbf{u}_{k-1})). \end{aligned} \quad (23)$$

The flexibility of $\phi_k(\mathbf{X}_k)$ to represent any likelihood model, including Huber or Cauchy, was noted in [3], and such factors are incorporated into ESGVI using the same procedure described above. Bringing these components together, the generalized form of (13) for MLGs is reformulated as,

$$(\boldsymbol{\Sigma}^{-1})^{(i+1)} = \mathbb{E}_{q^{(i)}} \left[\frac{\partial^2 \phi(\mathbf{X})}{\partial \delta \boldsymbol{\xi}^\top \partial \delta \boldsymbol{\xi}} \right], \quad (24a)$$

$$(\boldsymbol{\Sigma}^{-1})^{(i+1)} \delta \boldsymbol{\mu} = -\mathbb{E}_{q^{(i)}} \left[\frac{\partial \phi(\mathbf{X})}{\partial \delta \boldsymbol{\xi}^\top} \right], \quad (24b)$$

$$\bar{\mathbf{X}}^{(i+1)} = \bar{\mathbf{X}}^{(i)} \oplus \delta \boldsymbol{\mu}, \quad (24c)$$

where the gradients correspond to the expectations defined in (18). This completes the extension of the ESGVI update rules to MLG states.

V. SIMULATION AND EXPERIMENTAL RESULTS

The proposed method is validated in simulation and in experiments, and compared against two robust on-manifold MAP baselines, one employing the asymmetric Cauchy loss (MAP-C) shown in (9) and the other a GMM (MAP-GMM) to model skewed measurements. All estimators are implemented in PYTHON within a common factor-graph-based framework. No comparison is made against a vector-valued ESGVI formulation, as the advantages of MLG-based state representations, such as improved consistency and avoidance of singularities, are well established in the literature [8], [9]. The following sections present results first on simulated data, which allows controlled evaluation of performance, and then on real-world experiments to demonstrate practical applicability.

A. Simulation Results

The considered simulation comprises a planar robot equipped with wheel odometry, a gyroscope, and multiple range sensors capable of ranging to fixed landmarks of known position. The robot's motion in discrete-time is modeled as

$$\mathbf{X}_k = \mathbf{X}_{k1} \oplus_r (\Delta t(\mathbf{u}_{k1} + \mathbf{w}_{k1})), \quad \mathbf{w}_{k1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_{k1}), \quad (25)$$

where $\mathbf{X}_k \in SE(2)$ denotes the robot pose at time step k . The input \mathbf{u}_{k-1} is a stacked vector of angular and linear velocity measurements obtained from the gyroscope and wheel odometry. Since these measurements are expressed in the robot's body frame, the operator \oplus_r , is used to update the pose accordingly. Process noise \mathbf{w}_{k-1} , is modeled as zero-mean Gaussian with covariance \mathbf{Q}_{k-1} , and accounts for uncertainty in the motion estimate. To better reflect unmodeled dynamics such as wheel slip, the lateral velocity noise component is deliberately inflated. The robot's environment is cluttered, creating an alternating LOS and NLOS scenario. These skewed range measurements are generated as

$$y_k = \|\mathbf{r}_a^{\ell_i w} - \mathbf{C}_{ab} \mathbf{r}_b^{p_j z} - \mathbf{r}_a^{z w}\|_2 + \nu_k, \quad \nu_k \sim \mathcal{D}, \quad (26)$$

where $\mathbf{r}_a^{\ell_i w}$ denotes the i th landmark position relative to point w resolved in the absolute frame \mathcal{F}_a , $\mathbf{C}_{ab} \in SO(2)$ is the rotation matrix mapping the body frame \mathcal{F}_b to the absolute frame \mathcal{F}_a ,

TABLE I
ESTIMATION PERFORMANCE OVER 50 MONTE CARLO TRIALS

Method	RMSE (rad)	RMSE (m)	aNEES
MAP-C	0.026	0.050	0.916
MAP-GMM	0.026	0.046	0.990
ESGVI	0.025	0.046	1.019

$\mathbf{r}_b^{p_j z}$ is the known position of the j th range sensor relative to the gyroscope at z resolved in the robot's body \mathcal{F}_b , and similarly $\mathbf{r}_a^{z w}$ is the position of the robot's gyroscope relative to the origin of the absolute frame, resolved in \mathcal{F}_a . The skewed distribution \mathcal{D} is generated by sampling from a zero-mean Gaussian, from which 25% of measurements are additively perturbed by values drawn uniformly between 1σ and 6σ , where the resulting distribution is illustrated in Fig. 1(b). The objective is to estimate the full trajectory $\mathbf{X} = \{\mathbf{X}_0, \dots, \mathbf{X}_K\}$. The proposed ESGVI algorithm is compared against two robust MAP-based approaches, one using an asymmetric Cauchy loss (MAP-C) and the other a GMM (MAP-GMM) of range measurement noise. Both are formulated to be compatible with MLG states for a fair comparison. Performance is evaluated using root-mean-squared error (RMSE),

$$\text{RMSE} = \sqrt{\frac{1}{NK} \sum_{i=1}^N \sum_{k=1}^K \|\hat{\mathbf{X}}_{i,k} \ominus \mathbf{X}_{i,k}\|_2^2}, \quad (27)$$

where $\hat{\mathbf{X}}_{i,k}$, and $\mathbf{X}_{i,k}$ are the estimated and ground-truth states, for trajectory i , at time t_k , with RMSE reported separately for rotation (rad) and translation (m). Consistency, meaning the ability of the estimator to correctly quantify its own uncertainty, is measured using the average normalized estimation error squared (aNEES) [23, Section 5.4],

$$\text{aNEES} = \frac{1}{NK n_x} \sum_{i=1}^N \sum_{k=1}^K (\hat{\mathbf{X}}_{i,k} \ominus \mathbf{X}_{i,k})^\top \boldsymbol{\Sigma}_{i,kk}^{-1} (\hat{\mathbf{X}}_{i,k} \ominus \mathbf{X}_{i,k}), \quad (28)$$

where a perfectly consistent estimator has an aNEES equal to one, due to the normalization in (28) by the dimension of the state n_x . Performance comparisons are based on 50 Monte Carlo trials over a trajectory of 400 poses, with measurements and initial conditions regenerated randomly in each trial, with odometry being produced at 100 (Hz), and range measurements occurring at 10 (Hz). In all trials, the same GMM and Skew-Laplace parameters are used, fit using 5,000 samples as shown in Fig. 1(b). All algorithms are initialized with dead-reckoned gyroscope and wheel odometry measurements, though ESGVI can optionally be "warm-started" using a MAP-based solution to accelerate convergence. Finally, throughout this section, ESGVI uses a third-order Gauss-Hermite cubature method [10, Section 6.3] to evaluate the expectations in (19). As seen in Table I, all three estimators have similar performance. The GMM model is the most consistent, while the RMSE is comparable between the ESGVI and the MAP estimators. ESGVI is slightly overconfident compared to the MAP estimators, which follows the results of [2] as minimizing the KL-divergence, $\text{KL}(q||p)$ can result in a Gaussian that is too confident, underestimating the variance.




Fig. 2. Experimental setup showing the Clearpath Husky platform equipped with four UWB tags operating in a cluttered environment with four fixed UWB anchors. UWB modules are circled in blue, and obstacles are outlined in red.




Fig. 3. Example trajectory (Trajectory 6) driven by the Husky through a cluttered environment, comparing estimator outputs against ground-truth data recorded by a motion capture system.

B. Experimental Setup

Experiments were conducted using the Clearpath Husky vehicle, equipped with four DWM1000 Decawave UWB tags, as shown in Fig. 2. These tags measure ranges to a set of fixed UWB anchors positioned throughout the environment, which are then used for localization. UWB ranging relies on measuring the time-of-flight of radio signals [4] and is therefore sensitive to non-line-of-sight (NLOS) propagation and multipath effects, which can introduce bias in the measured ranges, as illustrated in Fig. 1. Such conditions are common in indoor environments, where obstacles, clutter, and frequent foot traffic create dynamic obstructions that lead to skewed measurements. To reproduce these effects in a controlled manner, obstacles made of metal, wood, and foam were placed near selected UWB anchors, as shown in Fig. 2.

The Clearpath Husky was driven along ten independent trajectories designed to randomly expose the UWB range measurements to alternating LOS and NLOS conditions. An example of the sixth trajectory is shown in Fig. 3. Ground-truth poses were recorded at 120 (Hz) using a Vicon motion capture system and used for estimator performance evaluation. The robot was operated at a nominal forward velocity of 0.4 (m/s), resulting in trajectories of approximately 60 seconds each in duration

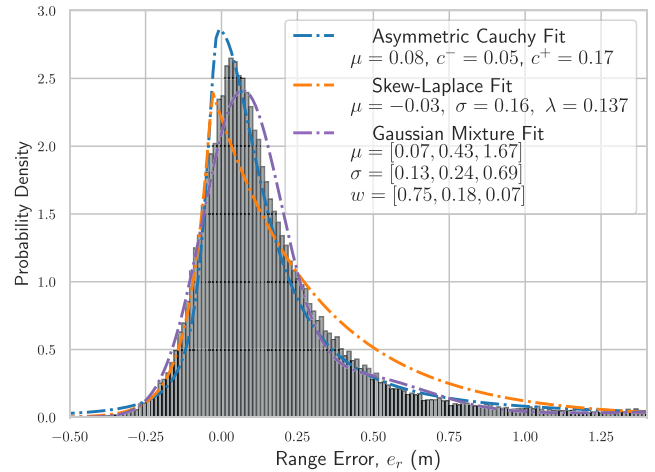


Fig. 4. Calibrated UWB range error histogram across all trajectories with fitted noise models.

within a 40 (m^2) indoor test area. UWB range measurements were collected at 30 (Hz) per tag, yielding over 75,000 range observations across all ten trajectories. In addition to NLOS effects, UWB measurements are affected by other error sources, including antenna delays, irregular radiation patterns, and clock skews between UWB tags. To account for these factors, the calibration procedure described in [4] was applied. In particular, the antenna-delay calibration was performed, and the data-driven model relating received signal power to attitude-dependent range bias was reused to isolate for NLOS-induced bias effects. Range measurement errors were then computed with respect to the ground-truth poses, across all trajectories. The resulting error distribution, shown in Fig. 4, clearly demonstrates the influence of obstacles through its pronounced skewed tails. Based on this distribution, the asymmetric Cauchy, GMM, and Skew-Laplace models were fitted. The GMM utilized three mixture components, consistent with the methodology of [20], where UWB errors were modeled in a similar fashion. All three models effectively capture the skewness arising from the prevalence of NLOS measurements, with the Skew-Laplace model exhibiting the most conservative behavior for highly biased ranges, while the asymmetric Cauchy and GMM models more closely align with the empirical histogram. Notably, the experimental error distribution closely mirrors the simulated case shown in Fig. 1(b), reinforcing the validity of the simulation assumptions.

C. Experimental Results

The experimental evaluation follows the same process and measurement models used in the simulation study presented in Section V-A. Specifically, the process dynamics remain as defined in (25), while the UWB measurement model follows (26), with parameters identified from the empirical error distributions aggregated across all ten trajectories shown in Fig. 4. Each estimator was initialized using dead-reckoned gyroscope and wheel odometry measurements prior to evaluation across all cluttered trajectories. The results, summarized in Fig. 5, demonstrate that ESGVI provides improved positional accuracy and comparable consistency and orientation accuracy relative to the asymmetric Cauchy MAP estimator. Notably, unlike in the simulation study, the GMM-based estimator performs




Fig. 5. Distribution of orientation RMSE, translation RMSE, and aNEES across 10 experimental trajectories for each estimator. The consistency bounds for the aNEES are computed using the median trajectory length.

less effectively in the experimental setting. This degradation is likely due to varying proportions of NLOS measurements across trajectories, which hinders the GMM’s ability to generalize, in contrast to the Monte Carlo simulations where NLOS measurements occurred at a fixed rate of 25%. From Fig. 3, it is evident that estimator performance deteriorates when the vehicle traverses NLOS regions caused by the metal, while other obstacles appear to have a less pronounced effect on the trajectory. On the other hand, both the asymmetric Cauchy MAP and ESGVI estimators exhibit comparable orientation RMSE, consistent with the findings in [3], which demonstrated that the performance advantages of the vector-valued ESGVI algorithm over MAP methods are primarily in translational accuracy rather than orientation estimation. This trend appears to extend to the MLG case as well. Regarding consistency, all estimators tend to be slightly overconfident, with the asymmetric Cauchy MAP and ESGVI exhibiting the most similar behavior. This overconfidence is expected because the metric is computed separately for each independent trajectory, and the proportion of NLOS measurements varies from one trajectory to another, making it a challenging dataset.

D. Limitations

Although ESGVI demonstrates improved accuracy, it is computationally more expensive than MAP-based methods, often an order of magnitude slower in practice. As noted in [3], ESGVI has the same theoretical asymptotic complexity as MAP, since both require a lower–diagonal–upper decomposition of the information matrix. However, the cubature-based computation of marginal expectations increases the constant factor, resulting in higher runtimes. In practice, this is further exacerbated by the lack of speed-oriented optimizations and the use of a simple backtracking line search.

While warm-starting ESGVI from a MAP solution improves convergence, the current implementation remains restricted to offline use and is therefore unsuitable for real-time applications. For a quantitative runtime evaluation, please refer to the relevant section of [3]. Nevertheless, ESGVI remains valuable in scenarios where robustness and improved estimation performance outweigh computational cost, and accelerating ESGVI remains an open direction for future work.

VI. CONCLUSION

In this letter, a more flexible and MLG compatible ESGVI method was presented. Building on the work of [3], it was shown how ESGVI can naturally incorporate skewed measurement noise models while supporting states defined on MLGs. The method was evaluated in an outlier-rich UWB localization experiment, demonstrating greater accuracy and comparable consistency relative to on-manifold robust MAP solutions using GMMs or asymmetric Cauchy losses to handle outliers. Although experimentally ESGVI was more computationally intensive than MAP, these results highlight its potential as a robust and flexible framework for state estimation in challenging environments.

APPENDIX

A. Matrix Lie Groups

Matrix Lie groups (MLGs) consist of elements that are matrices, typically denoted as $\mathbf{X} \in G$. To relate the Lie algebra, \mathfrak{g} , with its associated vector space, \mathbb{R}^g , two key mappings are defined: the *Hat* operator, which maps a vector $\boldsymbol{\tau} \in \mathbb{R}^g$ to its Lie algebra element $\boldsymbol{\tau}^\wedge \in \mathfrak{g}$, and the *Vee* operator, which performs the inverse mapping from $\boldsymbol{\tau}^\wedge$ back to $\boldsymbol{\tau}$, given by

$$\text{Hat} : \mathbb{R}^g \rightarrow \mathfrak{g}; \quad \boldsymbol{\tau} \mapsto \boldsymbol{\tau}^\wedge, \quad (29)$$

$$\text{Vee} : \mathfrak{g} \rightarrow \mathbb{R}^g; \quad \boldsymbol{\tau}^\wedge \mapsto (\boldsymbol{\tau}^\wedge)^\vee = \boldsymbol{\tau}. \quad (30)$$

The exponential and logarithmic maps describe the relationship between the Lie algebra \mathfrak{g} , and the MLG G , and are defined as

$$\exp : \mathfrak{g} \rightarrow G; \quad \boldsymbol{\tau}^\wedge \mapsto \mathbf{X} = \exp(\boldsymbol{\tau}^\wedge), \quad (31)$$

$$\log : G \rightarrow \mathfrak{g}; \quad \mathbf{X} \mapsto \boldsymbol{\tau}^\wedge = \log(\mathbf{X}). \quad (32)$$

The plus and minus operators are a concise means to describe operations involving elements of the MLG G and the Lie algebra \mathfrak{g} . Given that the composition of two group elements is not commutative, there are right and left versions of the plus and minus operators. The right versions are

$$\oplus_r : \mathbf{Y} = \mathbf{X} \oplus_r \delta \boldsymbol{\xi} \triangleq \mathbf{X} \exp(\delta \boldsymbol{\xi}^\wedge), \quad (33)$$

$$\ominus_r : \boldsymbol{\tau} = \mathbf{Y} \ominus_r \mathbf{X} \triangleq \log(\mathbf{X}^{-1} \mathbf{Y})^\vee, \quad (34)$$

and the left operations are

$$\oplus_{\ell} : \mathbf{Y} = \mathbf{X} \oplus_{\ell} \delta \boldsymbol{\xi} \triangleq \exp(\delta \boldsymbol{\xi}^{\wedge}) \mathbf{X}, \quad (35)$$

$$\ominus_{\ell} : \boldsymbol{\tau} = \mathbf{Y} \ominus_{\ell} \mathbf{X} \triangleq \log(\mathbf{Y} \mathbf{X}^{-1})^{\vee}. \quad (36)$$

If no subscript is indicated when writing either \oplus or \ominus , it is assumed that the expression holds for both directions, as long as the chosen direction remains consistent. For a more detailed review of matrix Lie groups, please refer to [6].

B. Stein's Lemma for Matrix Lie Groups

The expressions in (18) are not new results but rather simply Stein's identity written in Lie-group notation for the marginal MLG density. For the perturbation variable $\delta \boldsymbol{\xi}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{kk})$ and any smooth test function $f(\delta \boldsymbol{\xi}_k) = \phi_k(\bar{\mathbf{X}}_k \oplus \delta \boldsymbol{\xi}_k)$, Stein's lemma actually takes the form

$$\mathbb{E} \left[\frac{\partial f}{\partial \delta \boldsymbol{\xi}_k^{\top}} \right] = \boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E} [\delta \boldsymbol{\xi}_k f(\delta \boldsymbol{\xi}_k)], \quad (37a)$$

$$\mathbb{E} \left[\frac{\partial^2 f}{\partial \delta \boldsymbol{\xi}_k \partial \delta \boldsymbol{\xi}_k^{\top}} \right] = -\boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E} [f(\delta \boldsymbol{\xi}_k)] + \boldsymbol{\Sigma}_{kk}^{-1} \mathbb{E} [\delta \boldsymbol{\xi}_k \delta \boldsymbol{\xi}_k^{\top} f(\delta \boldsymbol{\xi}_k)] \boldsymbol{\Sigma}_{kk}^{-1}. \quad (37b)$$

Using the retraction $\mathbf{X}_k = \bar{\mathbf{X}}_k \oplus \delta \boldsymbol{\xi}_k$ and the corresponding difference $\mathbf{X}_k \ominus \bar{\mathbf{X}}_k$, these relations can be rewritten in the Lie-group form shown in (18). This rewriting is exact when log and exp acts as a local inverse, which holds within a neighborhood of the identity.

ACKNOWLEDGMENT

The authors would like to thank Kyle Biron-Gricken and Nicholas Dahdah for their assistance with the experimental setup, and Vassili Korotkine for helpful discussions on GMMs.

REFERENCES

- [1] T. D. Barfoot, *State Estimation for Robotics: Second Edition*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2024.
- [2] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [3] T. D. Barfoot, J. R. Forbes, and D. J. Yoon, "Exactly sparse Gaussian variational inference with application to derivative-free batch nonlinear state estimation," *Int. J. Robot. Res.*, vol. 39, no. 13, pp. 1473–1502, 2020.
- [4] M. A. Shalaby, C. C. Cossette, J. R. Forbes, and J. Le Ny, "Calibration and uncertainty characterization for ultra-wideband two-way-ranging measurements," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 4128–4134.
- [5] T. Kim, B. Yoon, and D. Lee, "UWB-based localization system considering antenna anisotropy and NLOS/multipath conditions," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2024, pp. 8810–8816.
- [6] J. Solà, J. Deray, and D. Atchuthan, "A micro Lie theory for state estimation in robotics," 2021, *arXiv:1812.01537*.
- [7] J. Vilà-Valls, F. Vincent, and P. Closas, "Decentralized information filtering under Skew-Laplace noise," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, 2019, pp. 291–295.
- [8] T. D. Barfoot and P. T. Furgale, "Associating uncertainty with three-dimensional poses for use in estimation problems," *IEEE Trans. Robot.*, vol. 30, no. 3, pp. 679–693, Jun. 2014.
- [9] A. W. Long, K. C. Wolfe, M. J. Mashner, and G. S. Chirikjian, "The banana distribution is Gaussian: A localization study with exponential coordinates," *Robot., Sci. Syst. VIII*, MIT Press, pp. 265–272, 2013.
- [10] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge, U.K.: Cambridge Univ. Press, 2013.
- [11] H. Xu et al., "Omni-swarm: A decentralized omnidirectional visual-Inertial-UWB state estimation system for aerial swarms," *IEEE Trans. Robot.*, vol. 38, no. 6, pp. 3374–3394, Dec. 2022.
- [12] A. Fishberg and J. P. How, "Multi-agent relative pose estimation with UWB and constrained communications," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2022, pp. 778–785.
- [13] A. Fishberg, B. J. Quiter, and J. P. How, "MURP: Multi-agent ultra-wideband relative pose estimation with constrained communications in 3D environments," *IEEE Robot. Automat. Lett.*, vol. 9, no. 11, pp. 10612–10619, Nov. 2024.
- [14] Y. Cao and G. Beltrame, "VIR-SLAM: Visual, inertial, and ranging SLAM for single and multi-robot systems," *Auton. Robots*, vol. 45, no. 6, pp. 905–917, 2021.
- [15] L. Qi, Y. Zhao, Q. Guo, M. Kaliuzhnyi, and E. Wang, "Deep learning-based NLOS identification of UWB signals," *IEEE Signal Process. Lett.*, vol. 32, pp. 2394–2398, 2025.
- [16] C. Jiang, J. Shen, S. Chen, Y. Chen, D. Liu, and Y. Bo, "UWB NLOS/LOS classification using deep learning method," *IEEE Commun. Lett.*, vol. 24, no. 10, pp. 2226–2230, Oct. 2020.
- [17] Z. Zhang, "Parameter estimation techniques: A tutorial with application to conic fitting," *Image Vis. Comput.*, vol. 15, no. 1, pp. 59–76, 1997.
- [18] B. Yang, J. Li, and H. Zhang, "UVIP: Robust UWB aided visual-inertial positioning system for complex indoor environments," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 5454–5460.
- [19] M. Kok, J. D. Hol, and T. B. Schön, "Indoor positioning using ultrawideband and inertial measurements," *IEEE Trans. Veh. Technol.*, vol. 64, no. 4, pp. 1293–1303, Apr. 2015.
- [20] W. Zhao, A. Goudar, M. Tang, X. Qiao, and A. P. Schoellig, "Uncertainty-aware gaussian mixture model for UWB time difference of arrival localization in cluttered environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2023, pp. 5266–5273.
- [21] T. Pfeifer, S. Lange, and P. Protzel, "Advancing mixture models for least squares optimization," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 3941–3948, Apr. 2021.
- [22] J. N. Wong, D. J. Yoon, A. P. Schoellig, and T. D. Barfoot, "Variational inference with parameter learning applied to vehicle trajectory estimation," *IEEE Robot. Automat. Lett.*, vol. 5, no. 4, pp. 5291–5298, Oct. 2020.
- [23] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation With Applications to Tracking and Navigation: Theory Algorithms and Software*. Hoboken, NJ, USA: Wiley, 2001.