

A Hessian for Gaussian Mixture Likelihoods in Nonlinear Least Squares

Vassili Korotkine[✉], Mitchell Cohen[✉], and James Richard Forbes[✉], *Member, IEEE*

Abstract—This letter proposes a novel Hessian approximation for Maximum a Posteriori estimation problems in robotics involving Gaussian mixture likelihoods. Previous approaches manipulate the Gaussian mixture likelihood into a form that allows the problem to be represented as a nonlinear least squares (NLS) problem. The resulting Hessian approximation used within NLS solvers from these approaches neglects certain nonlinearities. The proposed Hessian approximation is derived by setting the Hessians of the Gaussian mixture component errors to zero, which is the same starting point as for the Gauss-Newton Hessian approximation for NLS, and using the chain rule to account for additional nonlinearities. The proposed Hessian approximation results in improved convergence speed and uncertainty characterization for simulated experiments, and similar performance to the state of the art on real-world experiments. A method to maintain compatibility with existing solvers, such as *ceres*, is also presented.

Index Terms—Localization, optimization and optimal control, probabilistic inference, sensor fusion, SLAM.

I. INTRODUCTION AND RELATED WORK

ESTIMATING the state of a system from noisy and incomplete sensor data is a central task for autonomous systems. To describe inherent uncertainty in the measurements and state, probabilistic tools are required. Gaussian measurement likelihoods are commonly used in state estimation, allowing estimation problems to be easily formulated as instances of nonlinear least squares (NLS) optimization [1, §4.3]. The NLS formulations for state estimation stem from considering the negative log-likelihood (NLL) of the Gaussian, allowing a cancellation between the logarithm and the Gaussian exponent. Furthermore, the structure of NLS problems allows for the use of efficient solvers such as *ceres* [2], whereas more general optimizers or non-parametric methods [3] may be more computationally demanding. Therefore, the NLS structure, which for many problems arises from Gaussianity of the error model, is highly desirable. However, in practice, many sensor models are highly non-Gaussian, as in underwater acoustic positioning [4], or subsea Simultaneous Localization and Mapping (SLAM) [5].

Manuscript received 4 April 2024; accepted 14 July 2024. Date of publication 23 July 2024; date of current version 1 August 2024. This article was recommended for publication by Associate Editor H. Yang and Editor J. Civera upon evaluation of the reviewers' comments. This work was supported by the McGill University as well as the Natural Sciences and Engineering Research Council of Canada (NSERC) Alliance Grant in collaboration with Denso Corporation. (Corresponding author: Vassili Korotkine.)

The authors are with the Department of Mechanical Engineering, McGill University, Montreal, QC H3A 0C3, Canada (e-mail: vassili.korotkine@mail.mcgill.ca; mitchell.cohen3@mail.mcgill.ca; james.richard.forbes@mcgill.ca).

Accompanying software and supplementary material can be found at https://github.com/decargroup/hessian_sum_mixtures.

Digital Object Identifier 10.1109/LRA.2024.3432350

Even when the sensor characteristics are well-modeled by Gaussian distributions, multimodal error distributions arise in real-world situations involving ambiguity. Examples include simultaneous localization and mapping (SLAM) with unknown data associations [6], 3D multi-object tracking [7], and loop-closure ambiguity in pose-graph optimization [8]. In such cases, Gaussian measurement likelihoods are insufficient, and realizing robust and safe autonomy requires more accurate modeling of the underlying distributions. A popular parametric model used to represent multimodal distributions is the Gaussian mixture model (GMM), composed of a weighted sum of Gaussian components [9, §2.3.9]. Extensions of the *incremental smoothing and mapping* framework iSAM2, such as multi-hypothesis iSAM [10], have been developed to represent ambiguity in state estimation problems. However, these approaches are incompatible with standard NLS optimization methods and require a dedicated solver. Hence, other research has focused on incorporating GMM measurement likelihoods for use in standard NLS optimization frameworks. In the GMM case, the NLL does not simplify since the logarithm cannot cancel with the Gaussian exponents, meaning that GMMs are not directly compatible with NLS solvers. A common method to overcome this is the *Max-Mixture*, introduced by Olson et al. [11], which approximates the summation over Gaussians with a maximum operator, reducing the problem back to a single Gaussian. This technique allows the problem to be cast into an NLS form at the cost of introducing additional local minima. Rosen et al. [12] introduce the extension of Robust Incremental least-Squares Estimation (RISE) to non-Gaussian models, which allows for error terms with arbitrary distributions to be cast into instances of NLS minimization. The application of the framework in [12] to GMMs error terms is presented in Pfeifer et al. [13] and called the *Sum-Mixture* method. While the approach in [12] works well for its intended scope, its application to GMMs results in issues when iterative optimization methods are used. Common iterative methods for solving NLS problems, such as Gauss-Newton, rely on an approximation of the Hessian of the cost function. Accuracy of the Hessian approximation manifests itself in convergence of the algorithm, since it is used to compute the step in NLS algorithms. Furthermore, the Hessian approximation can affect consistency, since the Laplace approximation [9, Sec. 4.4] uses the Hessian to provide a covariance on the estimate computed using the NLS algorithm.

When the Sum-Mixture approach is used in conjunction with such optimization techniques, the resulting Hessian approximation is inaccurate, leading to degraded performance.




Fig. 1. Comparison of the Hessian approximation for a Gaussian mixture factor consisting of two components, both centered around zero but with different covariances. For a scalar cost and state variable, the Hessian is also scalar. When substituted into a local optimization method, a better Hessian approximation results in better convergence. In particular, the proposed Hessian (orange) is much closer visually to the exact Hessian (green) than the Max-Mixture (blue) and Max-Sum-Mixture (yellow) approaches.

This inaccuracy is pointed out by Pfeifer et al. in [13], but not explained theoretically. Pfeifer et al. [13] then propose an alternative to the Sum-Mixture approach, termed the *Max-Sum-Mixture* approach, which is a hybrid between the Max-Mixture and Sum-Mixture approaches. The dominant component of the mixture is factored out, creating a Max-Mixture error term, while the remaining non-dominant factor is treated using the approach of Rosen et al. [12]. The Hessian contribution for the non-dominant factor is inaccurate in a similar fashion as that of the Sum-Mixture, and improving on this aspect is the focus of this letter.

The contribution of this letter is a novel Hessian approximation for NLS problems involving GMM terms. The proposed Hessian approximation is derived by using the Gauss-Newton Hessian approximation for the *component errors* corresponding to each component of the GMM, and using the chain rule to take into account the additional nonlinearities. This approach is similar in spirit to how the robust loss nonlinearity is treated when deriving the Triggs correction proposed in [14] and used in the *ceres* library [2]. The proposed Hessian approximation has better mathematical properties in the sense that all mixture components contribute in the same way to the Hessian, without treating the dominant mixture component in a different framework than the rest, unlike other state-of-the-art approaches. Improvements to Hessian accuracy can therefore be expected in cases where multiple mixture components overlap. Furthermore, traditional NLS solvers do not allow the user to specify a separate Hessian, only requiring an error definition and its corresponding Jacobian. Therefore, this letter also proposes a method of defining an error and Jacobian that uses the proposed Hessian while maintaining compatibility with the NLS solver.

The proposed Hessian approximation is showcased in Fig. 1 for a 1D example, where it is visually more accurate compared to state of the art. Performance of the proposed method is validated on a toy problem, a simulated point-set registration problem [13], [15] and on a SLAM problem with unknown data associations on the “Lost in the Woods” dataset [16]. The implementation is open-sourced at <https://github.com/>

decargroup/hessian_sum_mixture and uses the *navlie* NLS solver [17].

The rest of this letter is organized as follows. Newton’s method and the Gauss-Newton Hessian approximation for NLS are presented in Section II. Section III presents Gaussian mixtures, the difficulties in applying NLS to Gaussian mixtures, as well as a review of the previous approaches to solving this problem. Section IV discusses the reasons for the Hessian inaccuracy of existing methods, and Section V presents the proposed approach. Section VI describes how to maintain compatibility with traditional NLS solvers while using the proposed Hessian. Section VII presents simulation and experimental results highlighting the benefits of the proposed approach.

II. THEORETICAL BACKGROUND

Popular optimization methods used for state estimation are reviewed to motivate the improvement of the Hessian approximation used. Furthermore, the link to NLL minimization of Gaussian probabilities is reviewed, while Section III presents the extension to GMMs.

A. Newton’s Method

Newton’s method is a local optimization method that iteratively minimizes a quadratic approximation to the loss function [1, §4.3.1]. Denoting the optimization variable as $\mathbf{x} \in \mathbb{R}^n$ and the loss function to be minimized as $J(\mathbf{x})$, Newton’s method begins with an initial guess $\mathbf{x}^{(0)}$, and iteratively updates the optimization variable using $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} + \alpha^{(i)} \Delta \mathbf{x}^{(i)}$ where $\alpha^{(i)}$ is a user-defined step size and $\Delta \mathbf{x}^{(i)}$ is the descent direction. Setting the gradient of the local quadratic approximation of the loss at $\mathbf{x}^{(i)}$ to zero results in the descent direction being given by [18, §2.2]

$$\frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \Delta \mathbf{x} = -\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}^\top}, \quad (1)$$

where $\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}} \in \mathbb{R}^{1 \times n}$ and $\frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \in \mathbb{R}^{n \times n}$ are the loss Jacobian and Hessian respectively, and are evaluated at the current iterate $\mathbf{x}^{(i)}$. The notation $\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}^\top} = \left(\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}} \right)^\top$ is used throughout this letter. The iterate $\mathbf{x}^{(i)}$ is updated and the process is iterated to convergence. The main drawbacks of this method are the need to compute the Hessian, as well as the requirement for the Hessian to be positive definite. The Hessian is only guaranteed to be positive semidefinite for convex losses [19, §3.1.4], which are infrequent in robotics applications.

B. Gauss-Newton Method

The Gauss-Newton method is an optimization approach applicable to problems that have a NLS structure. This structure is leveraged to form an approximate Hessian that is in turn used in the Newton step (1). The loss function is written as the sum of squared error terms,

$$J(\mathbf{x}) = \sum_{i=1}^{n_f} J_i(\mathbf{x}) = \sum_{i=1}^{n_f} \frac{1}{2} \mathbf{e}_i(\mathbf{x})^\top \mathbf{e}_i(\mathbf{x}), \quad (2)$$

where $J_i(\mathbf{x})$ denotes a single summand, n_f is the number of summands, and $\mathbf{e}_i(\mathbf{x}) \in \mathbb{R}^{n_{e,i}}$. The Jacobian and Hessian are computed using the chain rule,

$$\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}} = \sum_{i=1}^{n_f} \frac{\partial J(\mathbf{x})}{\partial \mathbf{e}_i} \frac{\partial \mathbf{e}_i}{\partial \mathbf{x}}, \quad (3)$$

$$\frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} = \sum_{i=1}^{n_f} \left(\frac{\partial \mathbf{e}_i^\top}{\partial \mathbf{x}} \frac{\partial \mathbf{e}_i}{\partial \mathbf{x}} + \sum_{j=1}^{n_{e,i}} e_{i,j} \frac{\partial^2 e_{i,j}}{\partial \mathbf{x} \partial \mathbf{x}^\top} \right), \quad (4)$$

where the argument \mathbf{x} has been dropped in writing \mathbf{e}_i . The Gauss-Newton approximation sets the second term of (4) to zero, such that

$$\frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \approx \mathbf{H}_{\text{GN}} = \sum_{i=1}^{n_f} \frac{\partial \mathbf{e}_i^\top}{\partial \mathbf{x}} \frac{\partial \mathbf{e}_i}{\partial \mathbf{x}}, \quad (5)$$

where \mathbf{H}_{GN} is the Hessian approximation made by Gauss-Newton. For an error term $\mathbf{e}_i(\mathbf{x})$ affine in \mathbf{x} , this approximation is exact, and the algorithm converges in a single iteration.

C. Maximum a Posteriori Estimation

Maximum a Posteriori (MAP) estimation seeks to estimate the system state $\mathbf{x} \in \mathbb{R}^n$ by maximizing the posterior probability [1, Sec. 3.1.2], which for Gaussian error models results in the optimization problem

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \sum_{i=1}^{n_f} -\log \mathcal{N}(\boldsymbol{\eta}_i(\mathbf{x}); \boldsymbol{\mu}_i, \mathbf{R}_i), \quad (6)$$

where $\mathcal{N}(\boldsymbol{\eta}_i; \boldsymbol{\mu}_i, \mathbf{R}_i)$ is a Gaussian probability density function (PDF) in the error $\boldsymbol{\eta}_i$ with mean $\boldsymbol{\mu}_i$ and covariance \mathbf{R}_i . The MAP estimation problem results in an NLS problem of the form (2), with the normalized error $\mathbf{e}_i(\mathbf{x})$ obtained through the change of variables

$$\mathbf{e}_i(\mathbf{x}) = \sqrt{\mathbf{R}_i^{-1}} (\boldsymbol{\eta}_i(\mathbf{x}) - \boldsymbol{\mu}_i). \quad (7)$$

For non-Gaussian error models, the Gaussian in (6) is replaced by an arbitrary PDF $p(\boldsymbol{\eta}_i)$ resulting in a more general summand expression $J_i(\mathbf{x}) = -\log p(\boldsymbol{\eta}_i) + c_i$ in (2). The constant c_i absorbs extraneous normalization constants from $-\log p(\boldsymbol{\eta}_i)$ that do not affect the optimization.

III. GAUSSIAN MIXTURE ERROR TERMS

A possible choice for the error model probability $p_i(\cdot)$ is a *Gaussian mixture model*. Gaussian mixture models are the sum of n_k weighted Gaussian components, allowing for a wide range of arbitrary probability distributions to be easily represented. The expression for a single Gaussian mixture term, with the term index i omitted for readability, is given by

$$p(\mathbf{x}) = \sum_{k=1}^{n_k} w_k \mathcal{N}(\boldsymbol{\eta}_k(\mathbf{x}); \boldsymbol{\mu}_k, \mathbf{R}_k), \quad (8)$$

where w_k is the weight of the k 'th Gaussian component. Note that here the sum is over components, whereas in Section II the sums considered are over problem error terms, of which only

a single one is now considered. To find an expression of the NLL for error models of the form (8), first let α_k be the weight normalized by covariance, written as $\alpha_k = w_k \det(\mathbf{R}_k)^{-1/2}$. Then, dropping state- and covariance-independent normalization constants and using the change of variables (7) yields the expression for the corresponding summand given by

$$J_{\text{GMM}}(\mathbf{x}) = -\log p(\mathbf{x}) + c \quad (9)$$

$$= -\log \sum_{k=1}^{n_k} \alpha_k \exp \left(-\frac{1}{2} \mathbf{e}_k(\mathbf{x})^\top \mathbf{e}_k(\mathbf{x}) \right). \quad (10)$$

The summand corresponding to a Gaussian mixture model NLL, given by (10), is not of the form of a standard NLS problem in (2). Unlike with Gaussian error terms, an NLS problem cannot be directly obtained, since the logarithm does not cancel the exponent. Hence, standard NLS optimization methods, such as Gauss-Newton, are not directly applicable to problems with Gaussian mixture error terms. Three existing approaches have been proposed to utilize Gaussian mixture terms within the framework of NLS, as discussed in the following sections.

A. Max-Mixtures

The Max-Mixture approach, first introduced in Olson et al. [11], replaces the summation in (10) with a max operator, such that the NLL is given by

$$-\log p_{\text{max}}(\mathbf{x}) = -\log \left(\max_k \alpha_k \exp \left(-\frac{1}{2} \mathbf{e}_k(\mathbf{x})^\top \mathbf{e}_k(\mathbf{x}) \right) \right). \quad (11)$$

The key is that the max operator and the logarithm may be swapped, allowing the logarithm to be pushed inside the max operator to yield the Max-Mixture error, given by

$$-\log p_{\text{max}}(\mathbf{x}) = -\log \alpha_{k^*} + \frac{1}{2} \mathbf{e}_{k^*}(\mathbf{x})^\top \mathbf{e}_{k^*}(\mathbf{x}), \quad (12)$$

where

$$k^* = \arg \max_k \alpha_k \exp \left(-\frac{1}{2} \mathbf{e}_k(\mathbf{x})^\top \mathbf{e}_k(\mathbf{x}) \right) \quad (13)$$

is the integer index of the dominant mixture component at any given \mathbf{x} . The Max-Mixture error in (12) is in the standard NLS form, and methods such as Gauss-Newton can be used to solve the problem.

B. Sum-Mixtures

A general method for treating non-Gaussian likelihoods in NLS is proposed in Rosen et al. [12]. For a general non-Gaussian NLL $-\log p_{\text{non-gauss}}(\mathbf{x})$, the corresponding NLS error $e_{\text{non-gauss}}(\mathbf{x})$ is derived by writing the relationship

$$J_{\text{SM}}(\mathbf{x}) = -\log p_{\text{non-gauss}}(\mathbf{x}) + c \quad (14)$$

$$= \frac{1}{2} \left(\sqrt{\log \gamma_{\text{SM}} - \log p_{\text{non-gauss}}(\mathbf{x})} \right)^2, \quad (15)$$

where γ_{SM} is a normalization constant chosen such that $\gamma_{\text{SM}} \geq p_{\text{non-gauss}}(\mathbf{x})$ to ensure that the square root argument is always positive. In this case, the constant c absorbs the constant offset

that is added by the normalization constant. The expression (15) allows the error definition

$$e_{\text{non-gauss}}(\mathbf{x}) = \sqrt{\log \gamma_{\text{SM}} - \log p_{\text{non-gauss}}(\mathbf{x})}. \quad (16)$$

The Sum-Mixture approach consists of applying this approach to the Gaussian mixture case,

$$\begin{aligned} J_{\text{SM}}(\mathbf{x}) &= \frac{1}{2} e_{\text{SM}}^2(\mathbf{x}) \\ &= \frac{1}{2} \left(\sqrt{2} \sqrt{-\log \frac{1}{\gamma_{\text{SM}}} \sum_{k=1}^{n_k} \alpha_k \exp\left(-\frac{1}{2} \mathbf{e}_k^T \mathbf{e}_k\right)} \right)^2. \end{aligned} \quad (17)$$

$$(18)$$

Ensuring positivity of the square root argument can be achieved by setting the normalization constant to $\gamma_{\text{SM}} = \sum_{k=1}^{n_k} \alpha_k$ [20].

C. Max-Sum-Mixtures

The Max-Sum-Mixture is introduced by Pfeifer et al. [13], and is a hybrid between the Max-Mixture and Sum-Mixture approaches. The dominant component of the mixture is factored out such that the corresponding summand is given by

$$\begin{aligned} J_{\text{MSM}}(\mathbf{x}) &= \frac{1}{2} \mathbf{e}_{k^*}^T \mathbf{e}_{k^*} \\ &\quad - \log \sum_{k=1}^{n_k} \alpha_k \exp\left(-\frac{1}{2} \mathbf{e}_k^T \mathbf{e}_k + \frac{1}{2} \mathbf{e}_{k^*}^T \mathbf{e}_{k^*}\right). \end{aligned} \quad (19)$$

The first term of (19) is analogous to the Max-Mixture case. The second term is treated using the Sum-Mixture approach,

$$e_2(\mathbf{x}) = \sqrt{2} \sqrt{-\log \frac{1}{\gamma_{\text{MSM}}} \sum_{k=1}^{n_k} \frac{\alpha_k}{\alpha_{k^*}} \exp(-\mathbf{e}_k^T \mathbf{e}_k + \mathbf{e}_{k^*}^T \mathbf{e}_{k^*})}, \quad (20)$$

where γ_{MSM} is chosen as $\gamma_{\text{MSM}} = n_k \max \frac{\alpha_k}{\alpha_{k^*}} + \delta$. The parameter $\delta > 0$ is a damping constant [20] that dampens the influence of this nonlinear term in the optimization.

IV. INACCURACY OF THE NON-DOMINANT COMPONENT MAX-SUM-MIXTURE HESSIAN

NLS problems are defined by the choice of error. Different choices of error may be made, which in turn defines the Gauss-Newton Hessian approximation. Because of the Gaussian noise assumption in most robotics problems, the error choice is natural because the negative log-likelihood is quadratic. However, since GMMs do not yield a quadratic log-likelihood, there is no natural error choice. The Sum-Mixture is one such error choice but yields an inaccurate Hessian approximation. The inaccuracy is made clear by considering a problem with a single Gaussian mixture error term of the form (8), having a single mode. The cost function is given by $J(\mathbf{x}) = \frac{1}{2} \mathbf{e}(\mathbf{x})^T \mathbf{e}(\mathbf{x})$. When the Sum-Mixture formulation is applied to this problem, the cost function becomes

$$J(\mathbf{x}) = \frac{1}{2} e_{\text{SM}}(\mathbf{x})^2, \quad (21)$$

with $e_{\text{SM}}(\mathbf{x}) = \sqrt{\mathbf{e}^T(\mathbf{x}) \mathbf{e}(\mathbf{x})}$. Even for a Gaussian likelihood problem that is easily solved by Gauss-Newton in its original formulation, since the error $e_{\text{SM}}(\mathbf{x})$ is a scalar, the Sum-Mixture approach results in a rank one Hessian approximation (5) for the corresponding summand. For state dimension $n > 1$, the Hessian approximation is non-invertible and the Newton update (1) fails. While this is not an issue within the original application in [12], the Hessian approximation when applied to mixtures is inaccurate as demonstrated with this base case.

The fundamental reason for this inaccuracy is that higher-level terms in the NLS Hessian (4) are neglected. Due to the added nonlinearity $e_{\text{SM}}(\mathbf{x}) = \sqrt{\mathbf{e}^T(\mathbf{x}) \mathbf{e}(\mathbf{x})}$, terms get absorbed from the left term of (4) into the right one and end up neglected. Even for well-behaved component errors, the additional nonlinearity imposed by the LogSumExp expression inside the square root in (18), together with the square root itself, make it such that the Gauss-Newton approximation is inaccurate, as demonstrated by the unimodal single-factor case. The Max-Sum-Mixture approach mostly addresses the shortcomings of the Sum-Mixture method by extracting the dominant component.

However, the non-dominant components in the second term of (19) are still subjected to the same treatment. Their contribution to the overall Hessian is thus inaccurate, leading to the discrepancies shown in Fig. 1.

V. PROPOSED APPROACH: HESSIAN-SUM-MIXTURE

The proposed approach is termed the ‘‘Hessian-Sum-Mixture’’ (HSM) method and improves on previous formulations by proposing a Hessian that takes into account the nonlinearity of the LogSumExp expression in the GMM NLL (10), while keeping the Gauss-Newton approximation for each component. The overall loss function is exactly the same, up to a constant offset, as for the Sum-Mixture and Max-Sum-Mixture approaches. However, the proposed Hessian derivation does not depend on defining a Gauss-Newton error.

The proposed Hessian-Sum-Mixture Hessian approximation is derived by setting the second derivatives of the component errors to zero, and using the chain rule to take into account the LogSumExp nonlinearity and derive the cost Hessian.

Formally, by writing the negative LogSumExp nonlinearity as

$$\rho(f_1, \dots, f_{n_k}) = -\log \sum_{k=1}^{n_k} \alpha_k \exp(-f_k), \quad (22)$$

and setting the f_k components to the quadratic forms $f_k(\mathbf{x}) = \frac{1}{2} \mathbf{e}_k(\mathbf{x})^T \mathbf{e}_k(\mathbf{x})$, the negative Gaussian mixture log-likelihood may be rewritten as $-\log p(\mathbf{x}) = J(\mathbf{x}) + c$ with

$$J(\mathbf{x}) = \rho(f_1(\mathbf{e}_1(\mathbf{x})), \dots, f_k(\mathbf{e}_k(\mathbf{x})), \dots, f_{n_k}(\mathbf{e}_{n_k}(\mathbf{x}))). \quad (23)$$

Using the Gauss-Newton Hessian approximation for the mixture components yields

$$\frac{\partial^2 f_k(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^T} \approx \frac{\partial \mathbf{e}_k}{\partial \mathbf{x}^T} \frac{\partial \mathbf{e}_k}{\partial \mathbf{x}}. \quad (24)$$

Applying the chain rule to (23) and using the approximation (24) yields the factor Jacobian as

$$\frac{\partial J(\mathbf{x})}{\partial \mathbf{x}} = \sum_{k=1}^{n_k} \frac{\partial \rho}{\partial f_k} \frac{\partial f_k}{\partial \mathbf{x}}, \quad (25)$$

and the Hessian as

$$\frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} = \sum_{k=1}^{n_k} \left(\frac{\partial \rho}{\partial f_k} \frac{\partial^2 f_k}{\partial \mathbf{x} \partial \mathbf{x}^\top} + \frac{\partial f_k}{\partial \mathbf{x}^\top} \sum_{j=1}^{n_k} \frac{\partial^2 \rho}{\partial f_j \partial f_k} \frac{\partial f_j}{\partial \mathbf{x}} \right). \quad (26)$$

The partial derivatives of ρ are given by

$$\frac{\partial \rho}{\partial f_k} = \frac{\alpha_k \exp(-f_k)}{\sum_{i=1}^{n_k} \alpha_i \exp(-f_i)} \quad (27)$$

$$\begin{aligned} \frac{\partial^2 \rho}{\partial f_j \partial f_k} &= \frac{-\delta_{jk}(\alpha_k \exp(-f_k)) \sum_{i=1}^{n_k} \alpha_i \exp(-f_i)}{(\sum_{i=1}^{n_k} \alpha_i \exp(-f_i))^2} \\ &\quad + \frac{\alpha_k \alpha_j \exp(-f_k) \exp(-f_j)}{(\sum_{i=1}^{n_k} \alpha_i \exp(-f_i))^2}, \end{aligned} \quad (28)$$

where δ_{jk} is the Kronecker delta and the Jacobian of $f_k(\mathbf{x}) = \frac{1}{2} \mathbf{e}_k(\mathbf{x})^\top \mathbf{e}_k(\mathbf{x})$ is given by

$$\frac{\partial f_k(\mathbf{x})}{\partial \mathbf{x}} = \mathbf{e}_k(\mathbf{x})^\top \frac{\partial \mathbf{e}_k(\mathbf{x})}{\partial \mathbf{x}}. \quad (29)$$

Using only the first of the terms in the summand of (26) guarantees a positive definite Hessian approximation since $\frac{\partial \rho}{\partial f_k}$ is guaranteed to be positive. The second term has no such guarantee. The final proposed Hessian approximation is

$$\frac{\partial^2 J(\mathbf{x})}{\partial \mathbf{x} \partial \mathbf{x}^\top} \approx \mathbf{H}_{\text{HSM}} = \sum_{k=1}^{n_k} \frac{\partial \rho}{\partial f_k} \frac{\partial^2 f_k}{\partial \mathbf{x} \partial \mathbf{x}^\top}, \quad (30)$$

where $\frac{\partial^2 f_k}{\partial \mathbf{x} \partial \mathbf{x}^\top}$ is approximated using (24). Examining (27) allows a somewhat intuitive interpretation of (30). The $\frac{\partial \rho}{\partial f_k}$ term is the relative strength of each component at the evaluation point, while $\frac{\partial^2 f_k}{\partial \mathbf{x} \partial \mathbf{x}^\top}$ is the Hessian of that component. The Hessian approximation \mathbf{H}_{HSM} in (30) corresponding to each summand is substituted directly into Newton's method (1), instead of the Gauss-Newton approximation in (5).

The key difference with respect to Max-Sum-Mixture is that the dominant and non-dominant components are all treated in the same manner. In the Max-Sum-Mixture the dominant component has a full-rank Hessian contribution, while the non-dominant components have a rank one Hessian contribution that is inaccurate as detailed in Section IV. In situations where the components overlap significantly, the Hessian is thus more accurate in the proposed method.

VI. NONLINEAR LEAST SQUARES COMPATIBILITY

Most NLS solvers do not support a separate Hessian definition such that the user may only specify an error and error Jacobian, which then are input to the GN approximation (5). This may be circumvented by engineering an "error" and "error Jacobian" that result in the same descent direction as using Newton's method (1) with the HSM Hessian (30), while maintaining

a similar cost function value to the Sum-Mixture (15). First, defining the following quantities

$$\mathbf{e}_{\text{solver},1}^\top = \left[\sqrt{\frac{\partial \rho}{\partial f_1}} \mathbf{e}_1^\top \quad \cdots \quad \sqrt{\frac{\partial \rho}{\partial f_K}} \mathbf{e}_K^\top \right], \quad (31)$$

$$\mathbf{J}_{\text{solver},1}^\top = \left[\sqrt{\frac{\partial \rho}{\partial f_1}} \frac{\partial \mathbf{e}_1}{\partial \mathbf{x}^\top} \quad \cdots \quad \sqrt{\frac{\partial \rho}{\partial f_K}} \frac{\partial \mathbf{e}_K}{\partial \mathbf{x}^\top} \right], \quad (32)$$

and substituting into the GN Hessian approximation (5) results in exactly the Hessian-Sum-Mixture Hessian (26) as well as the loss Jacobian (25). However, the resulting cost function $\|\mathbf{e}_{\text{solver},1}\|_2^2$ is not of the same form as (10), which can cause issues with methods that use the cost function to guide the optimization such as Levenberg-Marquardt (LM) [18]. To rectify this, an additional error term is required to make the loss into the same form as (10). The proposed error and Jacobian to be input into the solver are thus

$$\mathbf{e}_{\text{solver}}^\top = \left[\mathbf{e}_{\text{solver},1}^\top \quad \sqrt{2(\gamma_{\text{HSM}} + \Delta J)} \right], \quad (33)$$

$$\mathbf{J}_{\text{solver}}^\top = \left[\mathbf{J}_{\text{solver},1}^\top \quad 0 \right], \quad (34)$$

where ΔJ is the difference between the desired loss J_{GMM} (10) and the squared norm of $\mathbf{e}_{\text{solver},1}$,

$$\Delta J = J_{\text{GMM}} - \frac{1}{2} \|\mathbf{e}_{\text{solver},1}\|_2^2, \quad (35)$$

and γ_{HSM} is a normalization constant that ensures positivity of the square root in (33). The evaluated cost $\frac{1}{2} \|\mathbf{e}_{\text{solver}}\|_2^2$ is thus $\frac{1}{2} \|\mathbf{e}_{\text{solver}}\|_2^2 = J_{\text{GMM}} + \gamma_{\text{HSM}}$. Note that since $\mathbf{J}_{\text{solver}}$ is not the true Jacobian of the error $\mathbf{e}_{\text{solver}}$, caution must be taken if using automatic differentiation. The normalization constant γ_{HSM} is taken to be the lower bound on ΔJ to ensure positivity of the square root argument. By taking the logarithm and exponent of the second term in (35), and using the properties of exponents, (35) becomes

$$\Delta J = -\log \sum_{k=1}^{n_k} (\alpha_k \exp(S_k)), \quad (36)$$

$$S_k = \sum_{j=1}^{n_k} \frac{\alpha_j \exp(-f_j)}{\sum_{i=1}^{n_k} \alpha_i \exp(-f_i)} (f_j - f_k), \quad (37)$$

with $f_k = \frac{1}{2} \mathbf{e}_k^\top \mathbf{e}_k$. A sequence of algebraic manipulations and using the fact that $t \exp(-t) \leq \exp(-1)$ yields $S_k \leq \frac{1}{\alpha_k} \sum_{j=1}^{n_k} \alpha_j$, and thus

$$\Delta J \geq -\log \sum_{k=1}^{n_k} \left(\alpha_k \exp \left(\frac{1}{\alpha_k} \sum_{j=1}^{n_k} \alpha_j \right) \right). \quad (38)$$

The negative of the lower bound (38) is used as γ_{HSM} .

VII. SIMULATION AND EXPERIMENTS

To show the benefits of the proposed approach, the performance of the Max-Mixture (MM), Max-Sum-Mixture (MSM), Sum-Mixture (SM), the Hessian Sum-Mixture (HSM), are all evaluated on two simulated experiments and one real dataset. A benefit of HSM is in treating all of the components in the same way, unlike MM or MSM that privilege the dominant

component. The simulated examples thus focus on cases where components overlap. The considered simulation examples are similar to the ones provided in [13], and the real-world experiment is a SLAM problem with unknown data associations. In some of the tested examples, particularly the 2D toy example case, the Sum-Mixture approach fails when using standard Gauss-Newton due to a rank-deficient Hessian. To provide a fair comparison, Levenberg-Marquardt [18] (LM) is used instead of Gauss-Newton for all approaches. LM adds a multiple of the identity to the approximate Hessian, written as $\mathbf{H}^{\text{LM}} = \mathbf{H} + \lambda \mathbf{I}$, where λ is a damping constant for the LM method, \mathbf{H} is the original Hessian approximation given by (5) for the Max-Mixture, Sum-Mixture, and Max-Sum-Mixture, and by (30) for the Hessian Sum Mixture approach. This ensures that the Hessian approximation is always full rank. The LM damping constant λ is set as described in [18].

The metrics used to evaluate the quality of the obtained solution are root-mean-squared error (RMSE), defined as $\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N \|\hat{\mathbf{x}}_i \ominus \mathbf{x}_i\|_2^2}$, where $\hat{\mathbf{x}}_i$ and \mathbf{x}_i are, respectively, the estimated and ground truth states at timestep i . The \ominus operator represents a generalized “subtraction” operator, required as some examples considered involve states defined on Lie groups. Following the definitions in [21], a left perturbation is used for all Lie group operations in this letter. Additionally, for states that involve rotations, the difference between the states is split into a rotational and translational component, yielding a rotational RMSE (deg) and a translational RMSE (m).

Consistency, the ability of the estimator to characterize its error uncertainty, is measured using average normalized estimation error squared (ANEES) [22]. A perfectly consistent estimator has ANEES equal to one. The main drawback of this metric is that it characterizes the posterior state belief as Gaussian, which is never completely true in practice.

Lastly, the average number of iterations and runtime are used to characterize how quickly the estimator converges to the solution. A smaller number of iterations indicates better descent directions in (1), which is expected if a better Hessian approximation is used. The solver exits once the step size of the LM algorithm is below 1×10^{-8} or a cap of 200 iterations is reached.

A. Toy Example

The first simulated example consists of optimizing a problem with a Gaussian mixture model error term of the form

$$J(\mathbf{x}) = -\log \sum_{k=1}^{n_k} w_k \mathcal{N}(\mathbf{x}; \tilde{\mathbf{x}}_k, \mathbf{R}_k). \quad (39)$$

The toy example experiments in this section use four mixture components, such that $n_k = 4$. To quantify the performance of all algorithms considered, Monte-Carlo trials are run, where in each trial, random mixture parameters are generated and run for many different initial conditions. Both 1D and 2D examples are considered, where in the 1D case, the design variable is a scalar, and in the 2D case, $\mathbf{x} \in \mathbb{R}^2$. In all trials, the first mixture component is chosen with weight $w_1 \sim \mathcal{U}(0.2, 0.8)$, where $\mathcal{U}(a, b)$ denotes the uniform distribution over the interval

TABLE I
RESULTS OF THE TOY EXAMPLE MONTE CARLO EXPERIMENT

Dims.	Method	RMSE (m)	Avg Iter.	Time (s)	Succ. Rate [%]
1D	MM	4.07e-01	1.4	4.24e-03	24.1
	SM	1.84e-02	26.1	4.12e-02	98.9
	MSM	1.66e-02	18.6	3.10e-02	99.0
	HSM [†]	1.67e-02	8.8	1.50e-02	99.0
2D	MM	7.16e-01	1.5	4.42e-03	30.4
	SM	5.30e-02	27.9	4.61e-02	97.5
	MSM	4.83e-02	12.9	2.31e-02	97.8
	HSM [†]	4.89e-02	9.1	1.57e-02	97.8

Bold indicates best performance.

$[a, b]$. The weights for the other mixture components component are chosen as $w_k = (1 - w_1)/(n_k - 1)$. The mean for the first mixture component is set to $\tilde{\mathbf{x}}_1 = \mathbf{0}$, while each component of the mean of the other mixture components is chosen as $\tilde{x}_2^i \sim \mathcal{U}(-2, 2)$. The covariance for the first component is set to $\mathbf{R}_1 = \sigma_1^2 \mathbf{I}$, where $\sigma_1^2 \sim \mathcal{U}(0.4, 1)$. The second component covariance is set to a multiple of the first component such that $\mathbf{R}_2 = m \mathbf{R}_1$, where $m \sim \mathcal{U}(4, 10)$. For each generated set of mixture parameters, the optimization is run from 100 different initial points, chosen uniformly on a grid on $[-4, 4]$ for each axis in the problem dimension. Additionally, 1000 different mixture parameters are generated, resulting in a total of 100,000 Monte-Carlo trials for both the 1D example and 2D example. The ground truth is determined by sampling on a grid combined with local optimization to find the global optimum. The result is marked as successful if it converges to the ground truth within a given threshold, set to 0.01 meters for all experiments. The results for both the 1D and 2D Monte-Carlo experiments are shown in Table I. In all tables, the proposed method is denoted by [†]. The posterior in this case is extremely non-Gaussian, thus the ANEES is not reported. All methods achieve similar error values, with Sum-Mixture presenting a small improvement over the other methods. Note however that this is due to its interaction with LM, since for the 2D case particularly, the Gauss-Newton version of Sum-Mixture results in a non-invertible Hessian approximation and the method does not work. Nevertheless, the HSM approach still achieves a lower number of iterations than SM or MSM. The Max-Mixture approach achieves the fewest iterations but has a very low success rate.

B. Point-Set Registration

The point-set registration example follows the structure of the example in [13], which itself builds on the formulation of [15]. Given an uncertain source set of points $\mathcal{M} = \{\mathbf{m}_i\}_{i=1}^M$, as well as an uncertain reference set of points $\mathcal{T} = \{\mathbf{p}_j\}_{j=1}^F$ with $\mathbf{m}_i, \mathbf{p}_j \in \mathbb{R}^N$, the point-set registration task seeks to find the rigid-body transformation $\mathbf{T}_{ts}^* \in SE(N)$ that best aligns \mathcal{M} to \mathcal{T} , where N is either 2 or 3. The transformation is found by forming point-to-point residuals, \mathbf{e}_{ij} ,

$$\mathbf{e}_{ij}(\mathbf{p}_j, \mathbf{m}_i, \mathbf{C}_{ts}, \mathbf{r}_t^{st}) = \mathbf{p}_j - \mathbf{C}_{ts} \mathbf{m}_i - \mathbf{r}_t^{st}, \quad (40)$$

where $\mathbf{C}_{ts} \in SO(N)$ and $\mathbf{r}_t^{st} \in \mathbb{R}^N$ are the rotation and translation components of \mathbf{T}_{ts} , respectively. These errors are assumed

TABLE II
RESULTS OF THE POINT-SET REGISTRATION EXPERIMENT

Dims.	Method	RMSE (deg)	RMSE (m)	ANEES	Avg Iter.	Time (s)
2D	MM	1.59	0.03	2.92	13.54	1.75
	SM	1.15	0.02	0.05	55.05	9.50
	MSM	1.15	0.02	1.61	25.79	4.27
	HSM [†]	1.15	0.02	1.44	22.56	3.66
3D	MM	1.10	0.03	1.81	14.71	3.06
	SM	0.95	0.02	0.03	71.08	20.53
	MSM	0.95	0.02	1.34	22.81	6.02
	HSM [†]	0.95	0.02	1.26	21.24	5.51

Bold indicates best performance.

Gaussian $\mathbf{e}_{ij} \sim \mathcal{N}(\mathbf{e}_{ij}; \mathbf{0}, \mathbf{R}_{ij})$ with covariance

$$\mathbf{R}_{ij} = \mathbf{C}_{ts} \Sigma_m \mathbf{C}_{ts}^\top + \Sigma_f, \quad (41)$$

where Σ_m, Σ_f are the noise covariances on point measurements of points in \mathcal{M} and \mathcal{T} , respectively. Similarly to [13], to incorporate data association ambiguity, the likelihood of a measurement \mathbf{m}_i is given by the Gaussian mixture

$$p(\mathbf{m}_i | \mathbf{C}_{ts}, \mathbf{r}_t^{st}, \mathcal{T}) = \sum_{j=1}^M w_j \mathcal{N}(\mathbf{e}_{ij}(\mathbf{p}_j, \mathbf{m}_i, \mathbf{C}_{ts}, \mathbf{r}_t^{st}); \mathbf{0}, \mathbf{R}_{ij}), \quad (42)$$

with uniform weights $w_j = \frac{1}{M}$. 100 different landmark configurations are generated, with 100 different noisy point-cloud pairs generated for each. The experiments are run for a planar 2D case where $N = 2$, and a 3D case where $N = 3$. For the 2D (3D) case, each landmark configuration is simulated by first generating 15 (20) landmarks where each position component is generated according to $\mathcal{U} \sim (-5, 5)$, followed by duplicating 30% of them 4 times in a Gaussian with 0.11 covariance around the original. 100 measurement pairs are then generated for each landmark. The uncorrupted measurements are generated using the ground truth transformation \mathbf{T}_{ts} . The groundtruth transformation itself is randomly sampled as $\mathbf{T}_{ts} = \exp(\xi_{ts}^\wedge)$, where $(\cdot)^\wedge : \mathbb{R}^m \rightarrow \mathfrak{g}$, with $m = 3$ for $SE(2)$ and $m = 6$ for $SE(3)$, and $\exp(\cdot) : \mathfrak{g} \rightarrow G$ is the corresponding exponential map for $SE(2)$ or $SE(3)$. The Lie group and its corresponding Lie algebra are denoted G and \mathfrak{g} respectively. In all presented experiments, $\xi_{ts}^\top = [\xi^\phi^\top \ \xi^r^\top]$ with $\xi^\phi \sim \mathcal{U}(-\frac{15}{180}, \frac{15}{180})$, and $\xi^r \sim \mathcal{U}(-0.5, 0.5)$. The source and target point covariances Σ_m, Σ_f are each generated randomly as $\Sigma = \mathbf{C} \mathbf{D} \mathbf{C}^\top$, where \mathbf{D} is a diagonal matrix with entries generated from $\mathcal{U}(0.1, 0.6)$ and $\mathbf{C} \in SO(N)$ is a random direction-cosine-matrix generated as $\mathbf{C} = \exp(\xi^\wedge)$. Each entry in ξ is generated as $\xi_i \sim \mathcal{U}(-\pi, \pi)$, and the $(\cdot)^\wedge$ and $\exp(\cdot)$ operators are overloaded for $SO(N)$. Since the state is now defined on a Lie group, the Jacobian and Hessian of the loss with respect to the state are now replaced with their Lie group counterparts [21]. The results are presented in Table II. In this case, the proposed HSM approach slightly improves ANEES. This is expected since the Hessian approximation influences the resulting Laplace approximation and uncertainty characterization. The RMSE is similar for all approaches, with the exception of Max-Mixture, which has a higher RMSE, although it converges the fastest.

C. Lost in the Woods Dataset

The “Lost in the Woods” dataset [16] consists of a wheeled robot navigating a forest of plastic tubes, which are treated as landmarks. The robot receives wheel odometry measurements providing forward velocity and angular velocity measurements. The robot is equipped with a laser rangefinder that provides range-bearing measurements to the landmarks. The task is to estimate the robot poses, $\mathcal{T} = \{\mathbf{T}_1^{\text{rob}}, \dots, \mathbf{T}_K^{\text{rob}}\}$, with $\mathbf{T}_i^{\text{rob}} \in SE(2)$, $i = 1, \dots, K$, and the landmark positions $\mathcal{L} = \{\ell_1, \dots, \ell_{n_\ell}\}$, $\ell_i \in \mathbb{R}^2$, $i = 1, \dots, n_\ell$, given the odometry and range-bearing measurements. The process model consists of nonholonomic vehicle kinematics with the wheel odometry forward and angular velocity inputs detailed in [23]. The range-bearing measurements of landmark j observed from pose k are in the form $\mathbf{y}_{jk} = \mathbf{g}(\mathbf{T}_k^{\text{rob}}, \ell_j) + \mathbf{v}_{jk}$, where $\mathbf{v}_{jk} \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_{jk})$ and $\mathbf{g}(\cdot, \cdot)$ is the range-bearing measurement model, also detailed in [23]. While the dataset provides the data-association variables (i.e., which landmark each measurement corresponds to), to evaluate the performance of the algorithms, the challenging example of *unknown data associations* is considered. When the data associations are unknown, a multimodal measurement likelihood may be used, as in [6], where each measurement could have been produced from any landmark in \mathcal{L} , such that

$$p(\mathbf{y}_k | \mathbf{T}_k^{\text{rob}}, \mathcal{L}) = \sum_{j=1}^{n_\ell} w_i \mathcal{N}(\mathbf{y}_k - \mathbf{g}(\mathbf{T}_k^{\text{rob}}, \ell_j); \mathbf{0}, \mathbf{R}_{jk}). \quad (43)$$

The optimization problem to be solved then consists of process model residuals corresponding to the nonholonomic vehicle dynamics, as well as multimodal range-bearing error terms of the form (43) for each measurement. To simplify the problem, it is assumed that the groundtruth number of landmarks are known, so that the number of components in each mixture (43) is known *a priori*. Additionally, it is assumed that an initial guess for each landmark position is available. This is similar to the setup of the “oracle baseline” described in [24]. The dataset provides odometry and landmark measurements at 10 Hz. To reduce the problem size, the landmark measurements are subsampled to 1 Hz.

To provide an initial guess of the robot states and the landmark states, robot poses are initialized through dead-reckoning the wheel odometry measurements, while the landmarks are initialized by inverting the measurement model for the first measurement of each landmark. Note that the initialization of the landmark positions utilizes the data association labels, but the actual optimization solves for the associations implicitly using the multimodal likelihoods for each measurement.

The dataset of length 1200 seconds is split into subsequences of length T_{subseq} seconds each, such that there is no overlap between subsequences. The performance metrics computed for each subsequence to each subsequence are averaged to provide mean metrics. For each subsequence length, two scenarios are considered that correspond to different limits r_{max} on the range of the range-bearing measurements used.

For each subsequence length and maximum range radius, the average RMSE and the average number of iterations are

TABLE III
AVERAGE RMSE (m) ON THE LOST IN THE WOODS SUBSEQUENCES

Max Radius (m)	2			4		
Subseq. Len (s)	20	30	40	20	30	40
MM	0.30	0.41	0.55	0.14	0.16	0.24
SM	0.44	0.51	0.61	0.40	0.48	0.56
MSM	0.30	0.41	0.54	0.14	0.16	0.24
HSM [†]	0.30	0.41	0.48	0.14	0.16	0.25

Bold indicates best performance.

TABLE IV
AVERAGE ITERATIONS ON THE LOST IN THE WOODS SUBSEQUENCES

Max Radius (m)	2			4		
Subseq. Len (s)	20	30	40	20	30	40
MM	18.7	22.3	27.5	16.7	18.4	24.1
SM	200.0	200.0	200.0	200.0	200.0	200.0
MSM	18.7	22.3	27.6	16.7	18.8	24.4
HSM [†]	18.7	22.3	27.0	16.7	18.4	23.9

Bold indicates best performance.

TABLE V
AVERAGE TIME ON THE LOST IN THE WOODS SUBSEQUENCES

Max Radius (m)	2			4		
Subseq. Len (s)	20	30	40	20	30	40
MM	1.45	3.08	5.22	2.88	5.55	10.46
SM	15.04	35.27	49.27	48.86	87.48	121.64
MSM	1.63	3.53	6.22	3.57	7.16	13.76
HSM	1.60	3.41	5.83	3.38	6.66	12.76

Bold indicates best performance.

summarized in Tables III and IV, respectively. With the exception of Sum-Mixture that fails due to the inaccurate Hessian, performance is similar in this experiment for all methods in terms of RMSE and iteration count. This is explained by landmarks being far to each other relative to the range-bearing sensor noise characteristics. The less the components overlap, the smaller performance improvement can be expected from HSM relative to previous methods as stated in Section I. There is a slight improvement for runtime for HSM compared to MSM in Table V. This difference must be interpreted with caution. Although reasonable efforts were made to optimize all of the approaches, the Python implementation may still not be representative of optimized real-time systems. Nevertheless, both methods require the computation of $\frac{\partial \rho}{\partial f_k}$, \mathbf{e}_k , $\frac{\partial \mathbf{e}_k}{\partial \mathbf{x}}$ for each component to compute error Jacobians at each iteration. Thus, it is recommended to consider the iteration count as an implementation-independent proxy for runtime.

VIII. CONCLUSION

This letter proposes a novel Hessian approximation for GMMs that is compatible with NLS solvers. Speed and consistency improvements are shown in simulation and similar performance to state of the art is demonstrated in experiment. While aimed at improving convergence, the proposed method remains a local optimization method, requiring a good enough initial guess for proper operation. Future work will consider global convergence properties of the solution.

REFERENCES

- [1] T. D. Barfoot, *State Estimation for Robotics*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [2] S. Agarwal, K. Mierle, and The Ceres Solver Team, "Ceres solver," Oct. 2023. [Online]. Available: <https://github.com/ceres-solver/ceres-solver>
- [3] A. Doucet et al., *Sequential Monte Carlo Methods in Practice*, vol. 1. Berlin, Germany: Springer, 2001.
- [4] M. Y. Cheung, D. Fourie, N. R. Rypkema, P. V. Teixeira, H. Schmidt, and J. Leonard, "Non-Gaussian SLAM utilizing synthetic aperture sonar," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 3457–3463.
- [5] P. Newman and J. Leonard, "Pure range-only sub-sea SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2003, pp. 1921–1926.
- [6] K. J. Doherty, D. P. Baxter, E. Schneeweiss, and J. J. Leonard, "Probabilistic data association via mixture models for robust semantic SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 1098–1104.
- [7] J. Pöschmann, T. Pfeifer, and P. Protzel, "Factor graph based 3D multi-object tracking in point clouds," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 10343–10350.
- [8] G. H. Lee, F. Fraundorfer, and M. Pollefeys, "Robust pose-graph loop-closures with expectation-maximization," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 556–563.
- [9] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [10] M. Hsiao and M. Kaess, "MH-iSAM2: Multi-hypothesis iSAM using bayes tree and hypo-tree," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 1274–1280.
- [11] E. Olson and P. Agarwal, "Inference on networks of mixtures for robust robot mapping," *Int. J. Robot. Res.*, vol. 32, no. 7, pp. 826–840, Jun. 2013.
- [12] D. M. Rosen, M. Kaess, and J. J. Leonard, "Robust incremental online inference over sparse factor graphs: Beyond the gaussian case," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2013, pp. 1025–1032.
- [13] 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.
- [14] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon, "Bundle Adjustment—A Modern Synthesis," in *Vision Algorithms: Theory and Practice*. Berlin, Germany: Springer, 2000, pp. 298–372.
- [15] B. Jian and B. C. Vemuri, "Robust point set registration using Gaussian mixture models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 8, pp. 1633–1645, Aug. 2011.
- [16] T. D. Barfoot, "AER1513 course notes, assignments, and data sets," University of Toronto, Institute for Aerospace Studies, 2011.
- [17] C. C. Cossette, M. Cohen, V. Korotkine, A. D. C. Bernal, M. A. Shalaby, and J. R. Forbes, "navlie: A python package for state estimation on lie groups," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2023, pp. 5282–5287.
- [18] K. Madsen, H. B. Nielsen, and O. Tingleff, "Methods for nonlinear least squares problems," 2004. [Online]. Available: <https://orbit.dtu.dk/en/publications/methods-for-non-linear-least-squares-problems-2nd-ed>
- [19] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [20] T. Pfeifer, "Adaptive estimation using Gaussian mixtures," Ph.D. dissertation, Dept. Elect. Eng. Inf. Technol., Chemnitz Univ. Technol., Chemnitz, Germany, Aug. 2023.
- [21] J. Solá, J. Deray, and D. Atchuthan, "A micro lie theory for state estimation in robotics," 2021. [Online]. Available: <http://arxiv.org/abs/1812.01537>
- [22] X. R. Li, Z. Zhao, and V. P. Jilkov, "Estimator's credibility and its measures," in *Proc. Triennial World Congress*, 2002, pp. 1–6.
- [23] C. C. Cossette, A. Walsh, and J. R. Forbes, "The complex-step derivative approximation on matrix lie groups," *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 906–913, Apr. 2020.
- [24] Y. Zhang, O. A. Severinsen, J. J. Leonard, L. Carlone, and K. Khosoussi, "Data-association-Free landmark-based SLAM," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 8349–8355.