

Advancing Mixture Models for Least Squares Optimization

Tim Pfeifer, Sven Lange and Peter Protzel

Abstract—Gaussian mixtures are a powerful and widely used tool to model non-Gaussian estimation problems. They are able to describe measurement errors that follow arbitrary distributions and can represent ambiguity in assignment tasks like point set registration or tracking. However, using them with common least squares solvers is still difficult. Existing approaches are either approximations of the true mixture or prone to convergence issues due to their strong nonlinearity. We propose a novel least squares representation of a Gaussian mixture, which is an exact and almost linear model of the corresponding log-likelihood. Our approach provides an efficient, accurate and flexible model for many probabilistic estimation problems and can be used as cost function for least squares solvers. We demonstrate its superior performance in various Monte Carlo experiments, including different kinds of point set registration. Our implementation is available as open source code for the state-of-the-art solvers Ceres and GTSAM.

Index Terms—Probabilistic Inference, Probability and Statistical Methods, Sensor Fusion, Localization, SLAM

I. INTRODUCTION

WOULD it not be great to have one probabilistic tool to describe arbitrary distributions that arise from simultaneous localization and mapping (SLAM), tracking or even point set registration? We show how to include Gaussian mixture models into the typical nonlinear least squares solvers that are used for these problems today.

State estimation in robotics is based on deterministic models, that describe the physical relation between measurements and states. Due to the absence of perfect measurements, probabilistic models are required additionally to describe the related uncertainty. For these models, the Gaussian distribution is the first choice, since it allows efficient least squares optimization.

However, sensor information is often corrupted by outliers that do not follow a Gaussian: Vision can blur, odometry can slip and GNSS can be affected by reflections and multipath. Even if the sensor noise itself is Gaussian, ambiguity can lead to multimodal densities for SLAM or point set registration problems. Without better probabilistic models, state estimation is likely to fail in these cases.

Manuscript received: August 25, 2020; Revised December 1, 2020; Accepted March 1, 2021. This paper was recommended for publication by Editor Allison Okamura upon evaluation of the Associate Editor and Reviewers' comments. This work was supported by "Bundesministerium für Wirtschaft und Energie" (German Federal Ministry for Economic Affairs and Energy).

The authors are with the Faculty of Electrical Engineering and Information Technology, Technische Universität Chemnitz, Germany, (e-mail: tim.pfeifer@etit.tu-chemnitz.de; sven.lange@etit.tu-chemnitz.de; peter.protzel@etit.tu-chemnitz.de)

Digital Object Identifier (DOI): 10.1109/LRA.2021.3067307




Fig. 1. Optimization of a Gaussian mixture's negative log-likelihood. Starting from 100 different initial positions, a nonlinear least squares solver is applied to find the minimum of three cost function variants that represent the same mixture, as described in Section IV-A. The scattered points represent the results after each individual optimization. While the proposed Max-Sum-Mixture converges fully to the global minimum, Sum-Mixture fails to converge for almost all initial points. By intersecting the cost surface, the lower plot shows why some points of Max-Mixture get stuck: It has a second local minimum. Please notice, that both plots share the same x-axis.

One of the most popular models for robust estimation is the Gaussian mixture model (GMM), which is composed of a sum of multiple weighted Gaussian components. In difference to the also common M-estimators, GMMs are preferred when outliers follow a systematic pattern or ambiguity has to be represented. They are flexible enough to represent asymmetry, heavy tails and skewness, but close enough to a single Gaussian to allow sufficient convergence using nonlinear least squares.

There already exist algorithms to include GMMs in nonlinear least squares optimization, like Max-Mixture (MM) [1] or the approach proposed in [2], in this work referred as Sum-Mixture (SM). However, both of them come with individual drawbacks in terms of model accuracy or convergence as we explain in Section II.

In this paper, we introduce a novel representation for Gaussian mixtures in least squares problems, such as state estimation using factor graphs or point set registration. Based on Sum-Mixture, we derive a hybrid formulation of Max-Mixture and Sum-Mixture using robust numerical methods in Section III. It offers the combined advantages of both, without their individual disadvantages.

In addition, we provide:

- A discussion how to reduce the impact of nonlinearity on the estimation process in Section III-B.
- A comparison against Max-Mixture and Sum-Mixture in extensive Monte Carlo experiments, as shown in Figure 1, in Section IV-A.
- A novel method to apply the proposed model to point set registration in Section IV-B.
- A detailed analysis of the reasons why state-of-the-art approaches fail under certain conditions in Section V.

Furthermore, we provide open source implementations¹ of our approach – as well as SM and MM – for two of the most common optimization frameworks in robotics: Ceres [3] and GTSAM [4].

II. PRIOR WORK

While the field of robust state estimation is remarkably broad, we focus on a brief overview of different robust probabilistic models for nonlinear least squares (NLS) and factor graph optimization. The common goal of all algorithms is to reduce the huge impact of non-Gaussian outliers to least squares problems.

M-estimators are one of the standard solutions for robust estimation. They describe a set of loss functions that weight large residuals down. Frameworks like GTSAM or Ceres bring a variety of M-estimators with them, like the historical Huber M-estimator [5] or the more recent Dynamic Covariance Scaling (DCS) [6]. Due to their symmetry, they are not suited to describe arbitrary multimodal distributions that arise from ambiguity or effects like reflected GNSS signals. Also, their performance depends strongly on the choice of specific parameters, which often leads to manual tuning. In more recent work like [7] and [8], the parametrization issue was resolved by self-tuning algorithms based on expectation-maximization. But still, they are limited by their symmetry constraint.

Switchable constraints [9] introduced weights for each residual that can be jointly optimized within the least squares problem. In this way, the optimizer can down-weight outliers individually. Albeit it is formulated different, the authors of [6] showed that its solution is equivalent to an M-estimator.

Gaussian mixtures are quite flexible models but hard to integrate into nonlinear least squares. So the authors of [10] proposed the idea of preselecting the dominant mode and using it as an approximation of the full GMM. Since the optimization problem was constructed fully Gaussian afterwards, it is just a limited workaround and no reselection is possible. Olson et al. went one step further in [1] and proposed a better approximation – a local maximum of multiple Gaussians. In contrast to the explicit selection in [10], the Max-Mixture approach selects the dominant mode locally and implicitly during the optimization. MM is widely adopted in SLAM and factor graph optimization as an efficient and robust model. However, it is still a limited approximation of the true density, especially GMMs with overlapping components are not represented well.

With [2], Rosen et al. proposed a very general approach to represent almost arbitrary distributions in NLS problems. This

idea had a fundamental importance for later work [11], where we firstly implemented an exact GMM for NLS. As we show in this publication, the Sum-Mixture model from [11] is fundamentally ill-posed beside some special cases. Nevertheless, both papers [1] and [2] inspired the novel approach that is presented in this work.

Besides least-squares optimization, the authors of [12] applied Gaussian mixtures to solve probabilistic point set registration. They used the full Gaussian mixture, but approximated the product of probability densities by a sum of them, which violates fundamental assumptions of probability theory.

The authors of [13] and [14] introduced novel algorithms to represent ambiguity, which includes equivalents of Gaussian mixture models. However, they require different kinds of specialized solvers and come with a lot of overhead.

We, instead, present a novel Gaussian-mixture implementation for common nonlinear least squares solvers, that can be applied to a variety of robotic applications. Our approach combines the well-behaved convergence of Max-Mixture with the exact representation of Sum-Mixture without adopting their particular weaknesses.

III. GAUSSIAN MIXTURE MODELS FOR LEAST SQUARES

To introduce our novel Gaussian mixture representation for least squares, we take a look at the Gaussian case before we extend it to mixtures.

Central idea of probabilistic state estimation is to find the most likely set of states \mathcal{X} based on a set of perceived measurements \mathcal{Z} . The likelihood maximization is done by minimizing the negative log-likelihood, which leads to the common least squares formulation:

$$\hat{\mathcal{X}} = \operatorname{argmax}_{\mathcal{X}} P(\mathcal{X}|\mathcal{Z}) \quad (1)$$

$$= \operatorname{argmin}_{\mathcal{X}} -\log P(\mathcal{X}|\mathcal{Z}) \quad (2)$$

$$= \operatorname{argmin}_{\mathcal{X}} -\sum_n \log P(\mathbf{r}(\mathbf{x}_n, \mathbf{z}_n)) \quad (3)$$

$$= \operatorname{argmin}_{\mathcal{X}} \frac{1}{2} \sum_n \|\mathbf{r}(\mathbf{x}_n, \mathbf{z}_n)\|_{\Sigma}^2 + C \quad (4)$$

Please note, that we use two notations for the squared Mahalanobis distance, which are

$$\|\mathbf{r}\|_{\Sigma}^2 = \|\mathcal{I}^{\frac{1}{2}}\mathbf{r}\|^2, \quad (5)$$

where Σ is the measurement covariance matrix and $\mathcal{I}^{\frac{1}{2}}$ is the square-root information matrix. Residual term \mathbf{r} describes a nonlinear difference between a subset of states \mathbf{x}_n and measurements \mathbf{z}_n . Constant C is the Gaussian log-normalization and does not depend on the states \mathcal{X} , so it can be omitted. The remaining pure quadratic form is the direct result of assuming Gaussian distributed noise for the residual \mathbf{r} . Since a quadratic problem is sensitive to any kind of outlier, the residual term is often wrapped into a loss function ρ which brings us to robust least squares:

$$\hat{\mathcal{X}} = \operatorname{argmin}_{\mathcal{X}} \frac{1}{2} \sum_n \|\rho(\mathbf{r}(\mathbf{x}_n, \mathbf{z}_n))\|^2 \quad (6)$$

¹Available under www.mytuc.org/mix

TABLE I
GAUSSIAN MIXTURE LOSS FUNCTIONS FOR LEAST SQUARES

Algorithm	Max-Mixture [1]	Sum-Mixture [2]	Max-Sum-Mixture
Notation	ρ_m	ρ_s	ρ_{ms}
Cost Function	$\left[\sqrt{\frac{\mathcal{I}_k^{\frac{1}{2}}(\mathbf{r} - \boldsymbol{\mu}_k)}{-2 \log \left(\frac{w_k \det \mathcal{I}_k^{\frac{1}{2}}}{\gamma_m} \right)}} \right]$	$\sqrt{-2e_k - 2 \log \left(\frac{1}{\gamma_s} \sum_{l=1}^L s_l \exp(e_l - e_k) \right)}$	$\left[\sqrt{\frac{\mathcal{I}_k^{\frac{1}{2}}(\mathbf{r} - \boldsymbol{\mu}_k)}{-2 \log \left(\frac{1}{\gamma_{ms}} \sum_{l=1}^L s_l \exp(e_l - e_k) \right)}} \right]$
Normalization	$\gamma_m \geq \max_l w_l \det \mathcal{I}_l^{\frac{1}{2}}$	$\gamma_s \geq \sum_l w_l \det \mathcal{I}_l^{\frac{1}{2}}$	$\gamma_{ms} = L \cdot \max_l w_l \det \mathcal{I}_l^{\frac{1}{2}} + \delta$

This new cost function (6) is composed of the residual function \mathbf{r} and the loss function ρ . The loss can be derived directly from negative log-likelihood of the residual (3) to apply different probabilistic models:

$$\rho(\mathbf{r}) = \sqrt{-2 \log(P(\mathbf{r}))} \quad (7)$$

The Mahalanobis distance (5) is a special case of this loss function. Please note, that we define ρ inside the square of (6) to allow asymmetric loss functions, which is slightly different to the literature. The main contribution of this paper can be condensed down to the answer to a simple question: *How can we represent the log-likelihood of a GMM using the loss function $\rho(\mathbf{r})$?*

A. Max-Sum-Mixture – A Novel Approach

For a better understanding of the underlying mathematical structure, we introduce a compressed notation for the probability density of a Gaussian mixture with L components:

$$P(\mathbf{r}) \propto \sum_{l=1}^L s_l \exp(e_l(\mathbf{r})) \quad (8)$$

The constant scaling s contains the weighting w of a component l and the square root information matrix $\mathcal{I}^{\frac{1}{2}}$:

$$s_l = w_l \det \mathcal{I}_l^{\frac{1}{2}} \quad (9)$$

The exponent e is defined as half of the squared Mahalanobis distance between the residual \mathbf{r} and the mean $\boldsymbol{\mu}$:

$$e_l(\mathbf{r}) = -\frac{1}{2} \left\| \mathcal{I}_l^{\frac{1}{2}}(\mathbf{r} - \boldsymbol{\mu}_l) \right\|^2 \quad (10)$$

The log-likelihood of (8) is the scaled version of the log-sum-exp function

$$\text{sLSE}(\mathbf{s}, \mathbf{e}) = \log(P(\mathbf{r})) = \log \left(\sum_{l=1}^L s_l \exp(e_l) \right), \quad (11)$$

with the set of scalings $\mathbf{s} = \{s_1, \dots, s_L\}$ and exponents $\mathbf{e} = \{e_1, \dots, e_L\}$. While it is not possible to push the logarithm

inside the sum, it is indeed possible to pull one particular exponent e_k outside:

$$\begin{aligned} -\text{sLSE}(\mathbf{s}, \mathbf{e}) &= -\log \left(\sum_{l=1}^L s_l \exp(e_l) \right) \\ &= -\log \left(\sum_{l=1}^L \frac{s_l \exp(e_l)}{\exp(e_k)} \cdot \exp(e_k) \right) \\ &= \underbrace{-e_k}_{\text{linear}} - \underbrace{\log \left(\sum_{l=1}^L s_l \exp(e_l - e_k) \right)}_{\text{nonlinear}} \quad (12) \end{aligned}$$

This is a common approach to formulate a numerical robust log-sum-exp function, but we apply it algebraically to split the sum in a linear and a remaining nonlinear part. Since e_k can be chosen freely, we choose the locally dominant component to get a big linear term. So the index k of e_k is defined by

$$k = \underset{l}{\operatorname{argmax}} (s_l \exp(e_l)). \quad (13)$$

Interestingly, this part is identical with the Max-Mixture formulation from [1]. In difference to MM, we keep the remaining nonlinear term in the least squares problem. Therefore, we redefine the set of exponents as

$$\tilde{\mathbf{e}} = \{e_1 - e_k, \dots, e_L - e_k\}, \quad (14)$$

and apply the generalization approach of [2] by introducing a normalization constant γ_{ms} :

$$-\text{sLSE}(\mathbf{s}, \tilde{\mathbf{e}}) \propto \frac{1}{2} \left\| \sqrt{-2 \text{sLSE} \left(\frac{\mathbf{s}}{\gamma_{ms}}, \tilde{\mathbf{e}} \right)} \right\|^2 \quad (15)$$

The normalization guarantees that the expression inside the square root is always positive. As shown in detail in Appendix A, it can be set to

$$\gamma_{ms} = L \cdot \max_l s_l + \delta \text{ with } \delta \geq 0. \quad (16)$$

Damping parameter δ can be used to reduce the nonlinearity of (15) – its optimal choice is discussed the next section. Now, everything can be put together to formulate our main contribution – the square root of the negative log-likelihood of a GMM:

$$\rho_{ms}(\mathbf{r}) = \left[\sqrt{\frac{\mathcal{I}_k^{\frac{1}{2}}(\mathbf{r} - \boldsymbol{\mu}_k)}{-2 \log \left(\frac{1}{\gamma_{ms}} \sum_{l=1}^L s_l \exp(e_l - e_k) \right)}} \right] \quad (17)$$

Please note that this is a vector expression with a dimensionality of $\dim(\mathbf{r}) + 1$. The related Jacobian is provided in Appendix B.

B. Damping the Nonlinearity

The previously introduced damping factor δ seems superfluous in the first place, but it can be used to control the impact of nonlinearities. Therefore, its worth thinking about the best choice of δ .

The effect of the damping factor on the optimization itself is rather small, since it only shifts the cost surface. However, as elaborated in [15, p. 22], shifting the cost can decrease the rate of convergence. Additionally, a high δ decreases the relative cost change, which can lead to an early termination of the optimization. Therefore, δ should not be too high.

The damping factor affects also the estimated covariance $\hat{\Sigma}_{\mathbf{x}}$ of a state \mathbf{x} , that can be recovered from the optimization problem. Covariance recovery can be done by inverting the Hessian \mathbf{H} that is approximated by the squared Jacobian \mathbf{J} :

$$\hat{\Sigma}_{\mathcal{X}} = \mathbf{H}^{-1} \approx (\mathbf{J}^T \mathbf{J})^{-1} \quad (18)$$

We refer to this approximated Hessian as pseudo-Hessian. In Figure 2, we illustrate the result of this approximation for different choices of the damping parameter. For a high δ , the pseudo-Hessian of MSM becomes similar to the pseudo-Hessian of MM, which is slightly too big – and therefore slightly overconfident as shown in Figure 3. As smaller the damping gets, as more overconfident the pseudo-Hessian becomes. So, δ should not be too low.

Based on practical experience, almost any value will work for the damping factor. Nevertheless, we could achieve the best results with values of between 1 and 1000. So, we recommend using $\delta = 10$ and apply this for our evaluation. Figure 3 shows the resulting cost function compared to Max-Mixture [1] and Sum-Mixture [2].

C. Relation between Sum-Mixture and Max-Sum-Mixture

Since the proposed Max-Sum-Mixture has a substantial algebraic overlap with Sum-Mixture, we take a closer look at their mathematical relation. Both loss functions are provided together with MM in Table I and their Jacobians are given in Appendix B. For a fair evaluation, we apply the numerical robust log-sum-exp formulation for SM, even though this was not part of the original publication.

By comparing the scalar loss ρ_s of SM with the intermediate step (12) of MSM, it becomes clear that both describe the same log-likelihood and their squared costs are directly proportional:

$$\rho_s^2 \propto \boldsymbol{\rho}_{ms}^T \cdot \boldsymbol{\rho}_{ms} \quad (19)$$

They differ only in two points: their normalization constants are slightly different and MSM uses a vectorized loss function where the dominant component is pulled out.

Since the normalization constants barely influence the optimization, the major advantage of MSM stems from the vectorization. Separating the dominant linear part from the nonlinear in (17), exposes a linear representation of the

dominant mode to the optimizer. For SM on the other hand, the optimizer has to calculate its steps based on the Jacobian of a completely nonlinear cost term. Thus, the actual steps of MSM behave better than those of SM, even though their squared costs are almost equal. In consequence, the proposed MSM is a much more efficient implementation of the concepts of [2] than the original scalar formulation. We will verify this advantage in the next section.

IV. MONTE CARLO EVALUATION

In the following, we explain the motivation behind our simulative evaluation. Our experiments are designed to analyze the properties of the proposed Max-Sum-Mixture formulation compared to state-of-the-art approaches from [1] and [2]. A good probabilistic loss function should be an accurate and efficient representation of the corresponding distribution. Therefore, we stress all approaches for their convergence properties, their accuracy, their credibility and their computational efficiency.

We provide two types of synthetic Monte Carlo experiments to generalize our findings beyond specific applications. A plain optimization of a random set of Gaussian mixtures and a set of randomized point set registration problems. All experiments were executed on a standard AMD Ryzen 7 desktop PC.

A. Plain Optimization

We design the first synthetic experiment to test our approach for a well-behaved convergence. This includes the ability to represent the true mode of a Gaussian mixture and a cost surface that allows the optimizer to converge fast.

Hence, we create the simplest possible optimization problem that involves a GMM – a single linear residual wrapped into a nonlinear loss function:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \frac{1}{2} \|\boldsymbol{\rho}(\mathbf{x})\|^2 \quad (20)$$

The loss function represents the negative log-likelihood as described in (17). To allow a rough comparison with M-estimators, we also add DCS [6] to the symmetric scenarios. Starting from 100 different initial positions, least-squares optimization is applied to find the minimum of each cost function. By repeating the experiment for a set of 1000 different Gaussian mixtures, we run 100 000 optimizations in total.

Dataset Generation: To generate a variety of different models, we sample the mixture parameters randomly from a set of uniform distributions. All mixtures are composed of two components, the first one with a zero mean and a standard deviation drawn between 0.1 m and 1 m. The second component has a random mean from the range of ± 2 m and the standard deviation of the first, multiplied by a random factor between 2 and 10. A weight between 0.2 and 0.8 is used for the first component – the second weight is simply the difference to 1. For two-dimensional models, we draw the values for both dimensions independently. To guarantee that converging to the global optimum is always possible, we detect models with multiple minima and reject them, using a combination of grid-based sampling and local maxima detection. The initial positions are linearly spaced in a range of ± 4 m to cover the




Fig. 2. The influence of the damping factor δ on the squared cost and the approximated Hessian of our approach. Each line represents a Max-Sum-Mixture cost function with a different value of δ – as darker the line, as lower the damping. A higher damping decreases the effect of nonlinearities on the Hessian, but shifts its cost surface upwards.

relevant cost surface. We implement the evaluation as part of the libRSF [16], using the Ceres [3] version of the Levenberg-Marquardt algorithm.

Metrics: We evaluate the root-mean-square error (RMSE) as distance between the true mode of the GMM and the optimized variable after convergence. The true mode is found by using a grid search, followed by a local optimization. To check if an optimization converged to the right minimum, we apply an error threshold of 0.01 m and define the ratio of runs below this limit as convergence success rate. All results are listed together with the averaged run time and number of iterations in Table III.

B. Point Set Registration

Our motivation behind the second simulated evaluation stems from ego-motion estimation using sensors that measure sets of spatial points, such as radars or lidars. This task can be interpreted as a robust point set registration problem, as shown for example by [17] with radar sensors. We apply the concept from [18] to represent two set of points with Gaussian mixture models and register them in a distribution-to-distribution manner. By using a GMM representation, we can simply integrate a varying uncertainty of each target point (or sometimes referred to as landmark) within the two point sets – crucial for having a credible estimator. Furthermore, we can omit the usually required step of finding explicit correspondences compared to other point set registration algorithms.

However, with focus on credibility, we differ from the implementation in [17]. They used a sum approximation to align both GMM distributions which leads to robustness against outliers, but results in wrong covariances. Their approach increases furthermore the possible number of local minima – we refer to future work for a more detailed investigation.




Fig. 3. Comparison of cost functions and approximated Hessians of different mixture representations against the true Gaussian mixture. The dotted line represents the negative log-likelihood of the underlying GMM and its true Hessian. Solid lines represent the different approaches to describe this GMM for least squares. The Hessians in the lower plot are approximated by squared Jacobians. At the cost minimum, Max-Sum-Mixture and Max-Mixture are close to the true Hessian, while the Hessian of Sum-Mixture is singular.

Theoretical Background: In general, point set registration is the task of finding the transformation which aligns a current or moving point set \mathcal{M} , so that it overlaps with a reference or fixed set \mathcal{F} . We use the index i for a point measurement \mathbf{m}_i within \mathcal{M} and j for a point in \mathcal{F} respectively. For finding the best transformation $\mathbf{x} = [\mathbf{x}^{\text{tran}}, \mathbf{x}^{\text{rot}}]^T$, we can maximize the joint distribution

$$P(\mathcal{M}|\mathcal{F}, \mathbf{x}) = \prod_{i=1}^{|\mathcal{M}|} P(\mathbf{m}_i|\mathcal{F}, \mathbf{x}). \quad (21)$$

Following the idea of [18] and [12], we model the reference set as GMM, which results in the following distribution for one point \mathbf{m}_i of \mathcal{M} :

$$P(\mathbf{m}_i|\mathcal{F}) = \sum_{j=1}^{|\mathcal{F}|} w \cdot \mathcal{N}(\mathbf{m}_i|\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j) \quad \text{with } w = \frac{1}{|\mathcal{F}|} \quad (22)$$

Taking into account the transformation and rewriting the normalization terms leads to the likelihood function

$$P(\mathbf{m}_i|\mathcal{F}, \mathbf{x}) \propto \sum_{j=1}^{|\mathcal{F}|} s_j \exp\left(-\frac{1}{2} \|\mathbf{r}_i - \boldsymbol{\mu}_j\|_{\boldsymbol{\Sigma}_j + \tilde{\boldsymbol{\Sigma}}_i}^2\right). \quad (23)$$

The residual \mathbf{r}_i represents the transformation of point \mathbf{m}_i into the reference frame with the current estimate of \mathbf{x} :

$$\mathbf{r}_i(\mathbf{x}) = \mathbf{R}(\mathbf{x}^{\text{rot}})\mathbf{m}_i + \mathbf{x}^{\text{tran}} \quad (24)$$

Also, the covariance matrix $\tilde{\boldsymbol{\Sigma}}_i = \mathbf{R}\boldsymbol{\Sigma}_i\mathbf{R}^T$ of \mathbf{m}_i is transformed into the current frame, using the current estimate's rotation matrix $\mathbf{R}(\mathbf{x}^{\text{rot}})$. The resulting least squares problem is

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \sum_{i=1}^{|\mathcal{M}|} \frac{1}{2} \|\boldsymbol{\rho}(\mathbf{r}_i(\mathbf{x}))\|^2, \quad (25)$$




Fig. 4. Example of one random point set registration problem with otherwise unknown correspondence information for better visual association. The transformation of the ground truth points, before adding simulated measurement noise to both sets, was 0.5 m in x-direction and 15° in rotation. Specific properties of the described experiment like a measurement noise based on polar coordinates and clustered landmarks are apparent.

TABLE II
PARAMETERS FOR THE POINT SET REGISTRATION EXPERIMENTS.

Experiment	2D	3D
Num. of Configurations	100	100
Landmark Generation	$\mathcal{U}(x \pm 5 \text{ m})$ $\mathcal{U}(y \pm 5 \text{ m})$	$\mathcal{U}(r (10 \pm 1) \text{ m})$ $\mathcal{U}(\theta \pm \pi)$ $\mathcal{U}(\phi \pm 1/2\pi)$
Clustering (size)	40 % (3)	40 % (3)
Cluster spread	$\mathcal{N}(x, y 0, 0.1^2)$	$\mathcal{N}(x, y, z 0, 0.1^2)$
Runs per Configuration	1000	1000
Transformation sampling	$\mathcal{U}(x, y \pm 0.5 \text{ m})$ $\mathcal{U}(\alpha \pm 15^\circ)$	$\mathcal{U}(x, y, z \pm 0.5 \text{ m})$ $\mathcal{U}(\alpha_{x,y,z} \pm 5^\circ)$
Measurement noise	$\mathcal{N}(r 0, (0.2 \text{ m})^2)$ $\mathcal{N}(\theta 0, (3^\circ)^2)$	$\mathcal{N}(r 0, (0.2 \text{ m})^2)$ $\mathcal{N}(\theta, \phi 0, (3^\circ)^2)$

where the loss ρ describes the negative log-likelihood of the GMM (23). Note, that we can simply add an uncertain Gaussian component to the fixed frame, representing outliers, to have a robust version of the point set registration problem. We use this robust version in our last experiment with the Manhattan dataset. Again, Levenberg-Marquardt is applied to solve the optimization, in this experiments based on the GTSAM framework [4].

Dataset Generation: To evaluate the different mixture models applied to the point set registration problem, we simulate 1000 different measurement pairs on 100 different landmark configurations. Each consisting of 10 uniformly distributed landmarks in a square of (10×10) m. Following the typical characteristics of an automotive radar sensor, 40 % of the landmarks are duplicated two times based on a standard deviation of 0.1 m to generate clusters. Now, we assign a measurement noise (in polar space) with standard deviation

of 3° and 0.2 m to each of the 18 landmarks. Based on this, we sample a new measurement pair for each Monte Carlo run. The transformation between the measurements follows a uniform distribution in the range of ± 0.5 m in x- and y-direction and $\pm 15^\circ$ for the angle. The random set of 1000 different transformations is the same for each of the 100 landmark configurations.

In addition to the 2D dataset, a 3D variant with 6 degrees of freedom was generated in the same style, but with 20 landmarks to compensate for the higher degree of freedom. Parameters for both experiments are summarized in Table II, where the variables r , θ , and (ϕ) are used for radius, azimuth, and (elevation) to define parameters in polar and (spherical) coordinates. Please note, that r means the radius in this context, while \mathbf{r} is the residual function. Furthermore, we use Cartesian coordinates with the nomenclature x, y, z and an Euler angle representation with $\alpha_{x,y,z}$ for rotations around the corresponding axes.

Since real world applications have to deal with outliers and sets of not fully overlapping points, we include an additional evaluation based on the Manhattan dataset M3500b from [19]. A number of 300 landmarks are added uniformly into the world with the same clustering and noise properties as before. Instead of random transformations, the dataset's trajectory is used – excluding loop closure transitions. Landmark measurements are simulated by a sensor with 10 m range and 320° FoV, resulting in an average outlier rate of 13 % and a maximum of 32 %.

Metrics: As it is common practice, we apply the RMSE separately for the translational and rotational errors to measure accuracy. Aside from accuracy, we also want to measure, how good the estimator calculates its own uncertainty – also referred to as the estimator's credibility. Therefore, the average normalized estimation error squared (ANEES) is used as originally described in [20, Chp. 5.4], but normalized with the problem's dimension as stated in [21]. If the estimator is perfectly credible, the ANEES value equals 1. Since the optimization problem is likely to have multiple local minima, we expect that none of the evaluated algorithms will be fully credible. Nevertheless, we use the ANEES as a comparative metric to evaluate which algorithm is the most credible under this rough conditions.

V. DISCUSSION

Based on the Monte Carlo experiments, we discuss the performance of our approach in the following section – for the numbers see Table III and Table IV.

In terms of accuracy, we have to distinguish between plain optimization and point set registration. In our plain optimization evaluation, we can see that the accuracy depends strongly on a well-behaved convergence. Due to its linearity, *Max-Mixture* converges perfectly for symmetric mixture models. However, it fails in almost 40 % of the asymmetric cases, where it converges to a false local minimum. Although we only include GMMs with a distinct global optimum, the MM approximation is prone to have multiple ones as shown in Figure 3. This demonstrates that albeit it is a close approximation, it behaves differently from the exact model. *Sum-Mixture* on the other hand is an

TABLE III
RESULTS OF THE SINGLE COST FUNCTION EVALUATION.

	Alg.	RMSE [m]	Converg. Succ. Rate [%]	Average Iterations	Average Runtime [μ s]
1D sym.	DCS	$7.48e-14$	100.0	4.2	26
	MM	$9.64e-11$	100.0	4.0	25
	SM	$4.75e-9$	100.0	29.0	155
	MSM	$8.83e-5$	100.0	7.0	44
1D asym.	MM	0.72	60.1	3.5	23
	SM	$1.65e-5$	100.0	29.3	142
	MSM	$9.47e-5$	100.0	9.6	58
2D sym.	DCS	$1.34e-13$	100.0	4	25
	MM	$3.38e-7$	100.0	3.9	26
	SM	2.19	6.8	75.7	415
	MSM	$5.11e-4$	100.0	3.6	27
2D asym.	MM	1.09	56.8	3.1	23
	SM	2.34	3.8	75.2	407
	MSM	$5.98e-4$	100.0	8.0	51

TABLE IV
RESULTS OF THE POINT SET REGISTRATION EXPERIMENTS.

	Alg.	RMSE [m]	RMSE [$^{\circ}$]	ANEES	Average Iterations	Average Runtime [ms]
2D	MM	0.117	1.48	1.99	3.7	2.7
	SM	0.100	1.29	0.13	22.0	7.3
	MSM	0.098	1.23	1.32	4.5	3.3
3D	MM	0.140	1.98	2.98	3.8	6.0
	SM	0.228	2.21	0.69	26.8	31.1
	MSM	0.122	1.71	2.18	4.4	8.5
M3500b (2D)	MM	0.156	2.10	2.09	3.4	3.7
	SM	0.178	2.22	0.40	19.4	7.4
	MSM	0.132	1.29	1.37	5.2	4.0

exact model and performs well for one-dimensional problems. For two-dimensional problems, however, it fails as shown in Figure 1 on the first page. A convergence success rate below 10% indicates that the underlying optimization problem is seriously ill-posed. This is caused by the problematic Jacobian that drops to zero, as shown in Figure 3. We cannot discuss this deeply here, but these singularities slow down the convergence in 1D and prevent it in 2D. The proposed *Max-Sum-Mixture* has none of these issues. It converges fast to the correct minimum in all evaluated cases and achieves a consistently low RMSE. The error is slightly higher than for well-converging cases of SM/MM, because of the early termination due to the remaining constant part of the cost function (see Section III-B). However, with a magnitude of 10^{-5} , this effect is usually negligible for practical applications. The M-estimator DCS shows the highest accuracy through all algorithms, which is expectable, since it is perfectly quadratic in the final phase of convergence. Still, it is not applicable for asymmetric problems.

In difference to the artificial example of plain optimization, all approaches are able to converge in our point set registration experiment. This is possible because the reasons of failure – singularities of SM or minima of MM – occur just locally in one of many cost terms. To give an example: A cost function with a Jacobian of zero, does not contribute to the optimization if the other derivatives are non-zero – the algorithm can still converge. Nevertheless, the individual weaknesses of each

approach degrade the accuracy and the proposed MSM provides the smallest RMSE in translation as well as in rotation. SM achieves a similar accuracy in the 2D Monte Carlo scenario, but for 3D it has twice the error of MSM. MM is more consistent than SM through the different scenarios, but its RMSE is about 15–20% higher than our proposed approach. The comparatively low rotational error of MSM in the Manhattan world scenario supports these findings. Since variant M3500b is characterized by a large initial orientation error, MSM can profit from its well-behaved convergence.

Analyzing the credibility of the point set registration problem, shows again the consequences of the problematic Jacobian of SM – a strongly underconfident result. It is worth mentioning, that the ANEES is an inverse proportional value, so an ANEES of 0.13 means the estimated covariance is about 8 times higher than the actual squared error. Hence, it is clear that in the 2D case, the covariance estimated by SM is way too high for practical applications. MSM is just slightly overconfident with an ANEES of 1.3, which is a good result if we take the multiple minima into account that occur in point set registration. The estimated uncertainty of MM is similar to MSM (as shown by Figure 3) and since its error is higher, it is more optimistic (overconfident). For 3D, the results are slightly shifted: Due to the increased number of local minima, all algorithms tend to have a higher ANEES. Since the ANEES of SM was too low in the 2D case, it now has the “best” ANEES for 3D. However, we have to mention that this is just the result of two undesirable effects that compensate each other – multiple minima and singular Jacobians. If any measure is taken to improve the convergence, like a better initialization or graduated non-convexity [22], it will become worse again. Therefore, we are still convinced that our approach is able to deliver the best credibility if unmodeled properties like multiple minima are taken into account.

By comparing the run time of all GMM-based algorithms, we can see that MM is the most rapidly converging one – which is not surprising due to its perfect linearity. Interestingly, DCS is not noticeably faster, even though the evaluation of multiple GMM components comes with a bigger overhead. The proposed MSM is just slightly slower across all experiments, with a comparable number of iterations. SM however, is far behind and converges between 2 and 10 times slower. Its nonlinearity forces the optimizer to iterate often with many small or even unsuccessful steps.

In summary, we can say that the proposed Max-Sum-Mixture algorithm is the most consistent one in our evaluation. It shows good convergence properties in the single cost function tests and the best accuracy combined with a good credibility in the point set registration experiment.

VI. CONCLUSION

In this work, we proposed a novel approach to implement Gaussian mixture models for robust least squares optimization. These models provide the flexibility to describe almost arbitrary distributions in state estimation or to represent the ambiguity of assignment problems. We demonstrated how to derive a well-posed loss function that represents the log-likelihood of a GMM mostly linear and without approximations.

By comparing against state-of-the-art methods, we showed the improved convergence properties of our Max-Sum-Mixture approach. It was able to converge fast and accurate to the global minimum for a broad range of models, as our extensive Monte Carlo evaluation confirmed. With point set registration, a practical application supported this results and showed also the superior combination of accuracy, credibility and computational efficiency that our approach provides.

We think that many state estimation problems in robotics can benefit from a well-converging Gaussian mixture formulation. Therefore, we want to emphasize again the available open source implementation of the proposed algorithm.

APPENDIX A

PROOF OF THE NORMALIZATION TERM

In this section, we prove that our choice (16) of the normalization constant γ_{ms} guarantees a negative log-likelihood above zero.

Proof: Define $\gamma_{ms}, s_l \in \mathbb{R}^{\geq 0}$ and $e_l \in \mathbb{R}^{\leq 0}$, where $s_k \exp(e_k) = \max_l s_l \exp(e_l)$ as defined in (13) and $\gamma_{ms} = L \cdot \max_l s_l$ as defined in (16).

$$\begin{aligned}
s_k \exp(e_k) &= \max_l s_l \exp(e_l) \\
s_k \exp(e_k) &\geq \frac{\sum_{l=1}^L s_l}{L} \cdot \exp(e_l) \\
1 &\geq \frac{\sum_{l=1}^L s_l \exp(e_l)}{L \cdot s_k \exp(e_k)} \\
1 &\geq \frac{\sum_{l=1}^L s_l \exp(e_l)}{L \cdot \max_l(s) \exp(e_k)} \\
\log(1) &\geq \log\left(\frac{\sum_{l=1}^L s_l \exp(e_l)}{L \cdot \max_l(s) \exp(e_k)}\right) \\
0 &\leq -\log\left(\frac{\sum_{l=1}^L s_l \exp(e_l)}{L \cdot \max_l(s) \exp(e_k)}\right) \\
0 &\leq -\log\left(\sum_{l=1}^L \frac{s_l \exp(e_l)}{\gamma_{ms} \exp(e_k)}\right) \quad \blacksquare
\end{aligned}$$

APPENDIX B

JACOBIANS

Here, we provide the Jacobian matrices for all cost functions that are given in Table I.

For *Max-Mixture* [1] the Jacobian is

$$\frac{\partial \rho_m}{\partial \mathbf{r}} = \begin{bmatrix} \mathcal{I}_k^{\frac{1}{2}} \\ \mathbf{0}_{1,N} \end{bmatrix},$$

with $\mathbf{0}_{1,N}$ as the $[1 \times N]$ zero matrix.

For *Sum-Mixture* [2] the Jacobian is

$$\frac{\partial \rho_s}{\partial \mathbf{r}} = \frac{\sum_{l=1}^L (\hat{s}_l \exp(e_l) \cdot \mathcal{I}_l(\mathbf{r} - \boldsymbol{\mu}_l))}{\sum_{l=1}^L \hat{s}_l \exp(e_l) \cdot \sqrt{-2 \text{sLSE}(\hat{\mathbf{s}}, \mathbf{e})}},$$

with $\hat{\mathbf{s}} = \{\hat{s}_1 \dots \hat{s}_L\}$ and $\hat{s}_l = \frac{s_l}{\gamma_s}$.

The proposed *Max-Sum-Mixture* uses the modified scalings $\tilde{\mathbf{s}} = \{\tilde{s}_1 \dots \tilde{s}_L\}$ and exponents $\tilde{\mathbf{e}} = \{\tilde{e}_1 \dots \tilde{e}_L\}$, which are

defined in detail as $\tilde{s}_l = \frac{s_l}{\gamma_{ms}}$ and $\tilde{e}_l = e_l - e_k$. The Jacobian is

$$\frac{\partial \rho_{ms}}{\partial \mathbf{r}} = \begin{bmatrix} \mathcal{I}_k^{\frac{1}{2}} \\ \frac{\sum_{l=1}^L (\tilde{s}_l \exp(\tilde{e}_l) \cdot (\mathcal{I}_l(\mathbf{r} - \boldsymbol{\mu}_l) - \mathcal{I}_k(\mathbf{r} - \boldsymbol{\mu}_k)))}{\sum_{l=1}^L \tilde{s}_l \exp(\tilde{e}_l) \cdot \sqrt{-2 \text{sLSE}(\tilde{\mathbf{s}}, \tilde{\mathbf{e}})}} \end{bmatrix}.$$

REFERENCES

- [1] E. Olson and P. Agarwal, "Inference on networks of mixtures for robust robot mapping," *Int. Journal of Robotics Research*, vol. 32, no. 7, pp. 826–840, 2013.
- [2] D. M. Rosen, M. Kaess, and J. J. Leonard, "Robust incremental online inference over sparse factor graphs: Beyond the Gaussian case," in *Proc. of Int. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 1025–1032.
- [3] S. Agarwal, K. Mierle, and Others, "Ceres Solver," <http://ceres-solver.org>.
- [4] F. Dellaert and Others, "GTSAM," <http://research.cc.gatech.edu/borg/gtsam>.
- [5] P. J. Huber, "Robust estimation of a location parameter," *The Annals of Mathematical Statistics*, vol. 35, pp. 73–101, 1964.
- [6] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard, "Robust map optimization using dynamic covariance scaling," in *Proc. of Int. Conf. on Robotics and Automation (ICRA)*, 2013, pp. 62–69.
- [7] G. Agamennoni, P. Furgale, and R. Siegwart, "Self-tuning m-estimators," in *Proc. of Int. Conf. on Robotics and Automation (ICRA)*, 2015, pp. 4628–4635.
- [8] N. Chebrolu, T. Labe, O. Vysotska, J. Behley, and C. Stachniss, "Adaptive robust kernels for non-linear least squares problems," *Robotics and Automation Letters (RA-L)*, 2021.
- [9] N. Sunderhauf and P. Protzel, "Switchable constraints for robust pose graph slam," in *Proc. of Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 1879–1884.
- [10] M. Pflingstorn and A. Birk, "Simultaneous localization and mapping with multimodal probability distributions," *The Int. Journal of Robotics Research*, vol. 32, no. 2, pp. 143–171, 2012.
- [11] T. Pfeifer, P. Weissig, S. Lange, and P. Protzel, "Robust factor graph optimization – a comparison for sensor fusion applications," in *Proc. of Int. Conf. on Emerging Technologies and Factory Automation (ETFA)*, 2016, pp. 1–4.
- [12] T. Stoyanov, M. Magnusson, H. Andreasson, and A. J. Lilienthal, "Fast and accurate scan registration through minimization of the distance between compact 3D NDT representations," *The Int. Journal of Robotics Research*, vol. 31, no. 12, pp. 1377–1393, 2012.
- [13] M. Hsiao and M. Kaess, "Mh-isam2: Multi-hypothesis isam using bayes tree and hypo-tree," in *Proc. of Int. Conf. on Robotics and Automation (ICRA)*, Montreal, QC, Canada, 2019, pp. 1274–1280.
- [14] D. Fourie, J. Leonard, and M. Kaess, "A nonparametric belief solution to the bayes tree," in *Proc. of Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016, pp. 2189–2196.
- [15] K. Madsen, H. Nielsen, and O. Tingleff, "Methods for nonlinear least squares problems," Technical University of Denmark, Tech. Rep., 2004.
- [16] T. Pfeifer and Others, "libRSF," <https://github.com/TUC-ProAut/libRSF>.
- [17] M. Barjenbruch, D. Kellner, J. Klappstein, J. Dickmann, and K. Dietmayer, "Joint spatial- and Doppler-based ego-motion estimation for automotive radars," in *Proc. of Int. Symposium on Intelligent Vehicles (IV)*, 2015, pp. 839–844.
- [18] B. Jian and B. C. Vemuri, "Robust Point Set Registration Using Gaussian Mixture Models," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 33, no. 8, pp. 1633–1645, 2011.
- [19] L. Carlone and A. Censi, "From Angular Manifolds to the Integer Lattice: Guaranteed Orientation Estimation With Application to Pose Graph Optimization," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 475–492, 2014.
- [20] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software*, 1st ed. John Wiley & Sons, Inc., 2001.
- [21] X. R. Li, Z. Zhao, and V. P. Jilkov, "Estimator's credibility and its measures," in *Proc. of Int. Federation of Automatic Control (IFAC)*, 2002.
- [22] H. Yang, P. Antonante, V. Tzoumas, and L. Carlone, "Graduated non-convexity for robust spatial perception: From non-minimal solvers to global outlier rejection," *Robotics and Automation Letters (RA-L)*, pp. 1127–1134, 2019.