

A Geometric Perspective on Fusing Gaussian Distributions on Lie Groups


PREPRINT*

 **Yixiao Ge**

Systems Theory and Robotics Group
School of Engineering
Australian National University
ACT, 2601, Australia
Yixiao.Ge@anu.edu.au

 **Pieter van Goor**

Systems Theory and Robotics Group
School of Engineering
Australian National University
ACT, 2601, Australia
Pieter.vanGoor@anu.edu.au

 **Robert Mahony**

Systems Theory and Robotics Group
School of Engineering
Australian National University
ACT, 2601, Australia
Robert.Mahony@anu.edu.au

May 2, 2024

ABSTRACT

Stochastic inference on Lie groups plays a key role in state estimation problems such as; inertial navigation, visual inertial odometry, pose estimation in virtual reality, etc. A key problem is fusing independent concentrated Gaussian distributions defined at different reference points on the group. In this paper we approximate distributions at different points in the group in a single set of exponential coordinates and then use classical Gaussian fusion to obtain the fused *posteriori* in those coordinates. We consider several approximations including the exact Jacobian of the change of coordinate map, first and second order Taylor's expansions of the Jacobian, and parallel transport with and without curvature correction associated with the underlying geometry of the Lie group. Preliminary results on $\mathbf{SO}(3)$ demonstrate that a novel approximation using parallel transport with curvature correction achieves similar accuracy to the state-of-the-art optimisation based algorithms at a fraction of the computational cost.

1 Introduction

The rising interest in modern robotic and avionic systems over the past 20 years has led to increased interest in state estimation for systems defined on differentiable manifolds, and particularly on Lie groups and homogeneous spaces. In the Euclidean space setting, the *de facto* approach to sensor fusion is that used by the (extended) Kalman filter [1]: the prior probability of the state estimate and the measurement likelihood are linearised and combined through Bayesian fusion. This approach is not directly compatible with the nonlinear structure of a smooth manifold and has motivated a significant body of work to adapt Bayesian fusion methodologies to the manifold setting. The Invariant Extended Kalman Filter (IEKF) [2] and Equivariant Filter (EqF) [3] both contain an update step, where they pose the fusion problem in one set of exponential coordinates centered at the current state estimate, and linearise the probability functions to apply Bayesian fusion. In a separate work [4], Wang and Chirikjian present the concept of a concentrated Gaussian distribution, which can be used to model fusion of distributions on Lie groups with exponential coordinates centered at multiple different group elements. This formulation is used by [5, 6], where the authors propose an optimization algorithm to fit the posterior distribution using Bayes rule on the concentrated Gaussians directly. In [7], this same fusion problem on Lie groups is solved numerically by truncating the Baker-Campbell-Hausdorff (BCH) formula with different numbers of terms. Recently, in [8, 9], the authors present a new scheme that models the covariance as a tensor object on the tangent space, and uses parallel transport to compensate for the intrinsic nonlinearity of the underlying manifold.

In this paper we revisit the problem of stochastic fusion of concentrated Gaussians on Lie groups. We extend the concept of concentrated Gaussian to allow the mean of the Gaussian to be separate from the group element at which the exponential coordinates for the distribution are centred. This allows us to treat Gaussian fusion on the tangent space without requiring computationally expensive optimisation procedures. However, it is important to compensate for the change of coordinates associated with defining a covariance of a distribution at a non-zero mean in a given set of coordinates. We consider several approximations including computing the exact Jacobian of the change of exponential coordinates map, first and second order Taylor's expansions of the Jacobian, and a parallel transport method. We also propose a novel method that does parallel transport with a curvature correction associated with the underlying geometry of the Lie group. This is particularly of interest since it can be implemented using only the matrix exponential function for which efficient implementations are readily available. We

*Under review

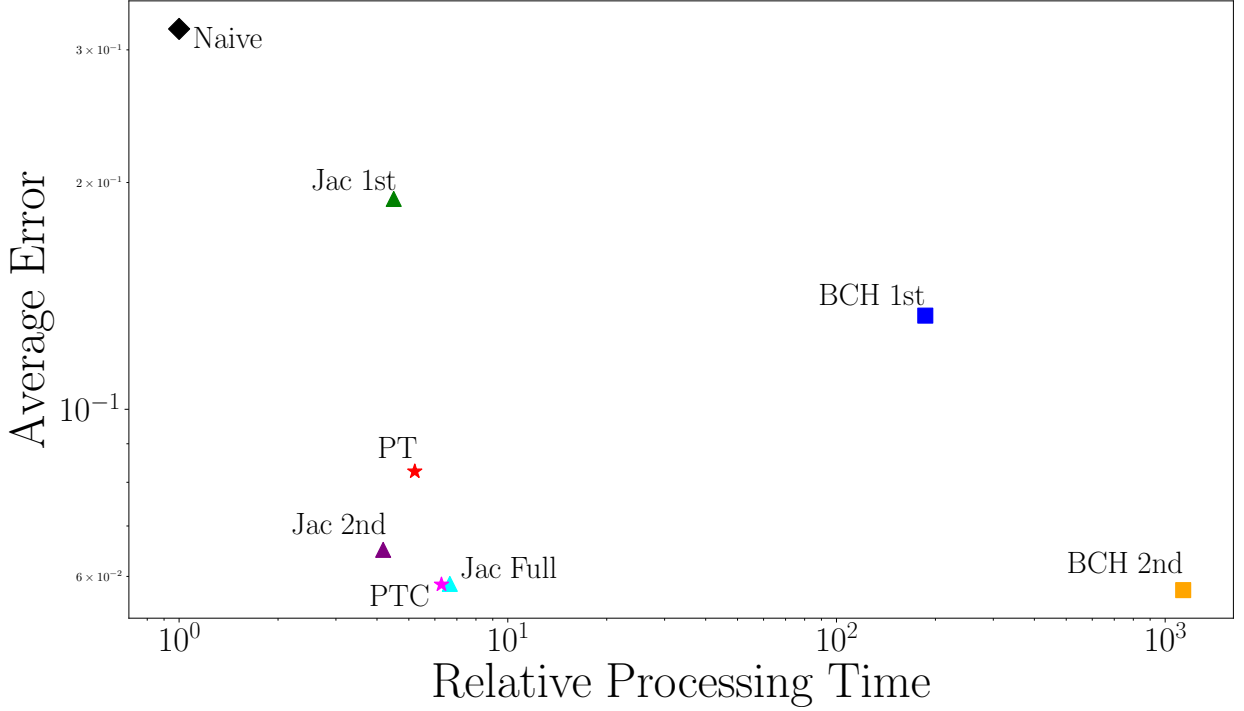


Figure 1: Comparative study of the average error (12) and the relative processing time of each fusion method. Relative processing time is measured as a ratio of algorithm run-time with respect to run-time for the naive fusion (\diamond) algorithm. “Jac Full(\triangle)” refers to using the analytic form of the Jacobian of the transition functions. “Jac 1st(\triangle)” and “Jac 2nd(\triangle)” refer to first and second order Taylor approximations of the Jacobian. “PT(\star)” refers to parallel transport and “PTC(\star)” to parallel transport with curvature correction. “BCH 1st(\square)” and “BCH 2nd(\square)” refer to the Baker-Campbell-Hausdorff approximation methods proposed in [7].

compare these five approximations with the BCH-based optimisation algorithms that are considered the state-of-art for fusion of concentrated Gaussians [7]. Preliminary results on $\mathbf{SO}(3)$ demonstrate that using parallel transport with curvature correction achieves similar accuracy to the state-of-the-art optimisation based algorithms at a fraction of the computational cost, as shown in Fig 1.

2 Preliminaries

2.1 General Lie groups

Let \mathbf{G} be a general Lie group with dimension n , associated with Lie algebra \mathfrak{g} . Let id denote the identity element of \mathbf{G} . Given arbitrary $X, Y \in \mathbf{G}$, the left and right translations are denoted by L_X and R_X , and are defined by

$$L_X(Y) := XY, \quad R_X(Y) := YX.$$

The Lie algebra \mathfrak{g} is isomorphic to a vector space \mathbb{R}^n with the same dimension. We use the wedge $(\cdot)^\wedge : \mathbb{R}^n \rightarrow \mathfrak{g}$ and vee $(\cdot)^\vee : \mathfrak{g} \rightarrow \mathbb{R}^n$ operators to map between the Lie algebra and vector space. The Adjoint map for the group \mathbf{G} , $\text{Ad}_X : \mathfrak{g} \rightarrow \mathfrak{g}$ is defined by

$$\text{Ad}_X[u^\wedge] = \text{DL}_X \circ \text{DR}_{X^{-1}} [u^\wedge],$$

for every $X \in \mathbf{G}$ and $u^\wedge \in \mathfrak{g}$, where DL_X , and DR_X denote the differentials of the left, and right translation, respectively. Given particular wedge and vee maps, the Adjoint matrix is defined as the map $\text{Ad}_X^\vee : \mathbb{R}^n \rightarrow \mathbb{R}^n$

$$\text{Ad}_X^\vee u = (\text{Ad}_X u^\wedge)^\vee.$$

The adjoint map for the Lie algebra $\text{ad}_{u^\wedge} : \mathfrak{g} \rightarrow \mathfrak{g}$ is given by

$$\text{ad}_{u^\wedge} v^\wedge = [u^\wedge, v^\wedge],$$

and is equivalent to the Lie bracket. Given particular wedge and vee maps, the adjoint matrix $\text{ad}_u^\vee : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is defined to be

$$\text{ad}_u^\vee v = (u^\wedge v^\wedge - v^\wedge u^\wedge)^\vee = [u^\wedge, v^\wedge]^\vee.$$

Let $\exp_{\mathbf{G}} : \mathfrak{g} \rightarrow \mathbf{G}$ denote the exponential map from the Lie algebra element to the group element. For matrix Lie groups such as $\mathbf{SO}(3)$, $\mathbf{SE}(3)$, this map is simply the matrix exponential. Let $\mathbf{G}' \subset \mathbf{G}$ be the subset of \mathbf{G} where the exponential map is invertible, one can then define the logarithm map $\log_{\mathbf{G}} : \mathbf{G}' \rightarrow \mathfrak{g}$ and $\log_{\mathbf{G}}^\vee : \mathbf{G}' \rightarrow \mathbb{R}^n$. For simplicity, we will suppress the wedge ‘ \wedge ’ operator in the exponential map throughout this paper.

The directional derivative of $\exp_{\mathbf{G}}$ at a point $u \in \mathfrak{g}$ in the direction of $w \in \mathfrak{g}$ is given by [10, Theorem 1.7]

$$\left. \frac{d}{dt} \exp_{\mathbf{G}}(u + tw) \right|_{t=0} = \text{DL}_{\exp_{\mathbf{G}}(u)} \left\{ \frac{\text{I} - \exp_{\mathbf{G}}(-\text{ad}_u)}{\text{ad}_u}(w) \right\}.$$

The differential $D_u \exp_{\mathbf{G}} : \mathfrak{g} \rightarrow \text{T}_{\exp_{\mathbf{G}}(u)} \mathbf{G}$ of $\exp_{\mathbf{G}}$ at u can be derived immediately:

$$D_u \exp_{\mathbf{G}} := \text{DL}_{\exp_{\mathbf{G}}(u)} \circ \frac{\text{I} - \exp(-\text{ad}_u)}{\text{ad}_u}.$$

This map is transcendental and is computed using asymptotic expansions for a general Lie group, although for specific Lie groups such as $\mathbf{SO}(3)$ and $\mathbf{SE}(3)$, algebraic forms exist in terms of known trigonometric transcendental functions \sin and \cos [11]. By identifying $\text{T}_{\exp_{\mathbf{G}}(u)} \mathbf{G}$ with \mathbb{R}^n via left trivialization and ‘ \vee ’ operator, one can define the *left trivialised* Jacobian map $J_u : \mathbb{R}^n \rightarrow \mathbb{R}^n$ as

$$J_u := \text{DL}_{\exp_{\mathbf{G}}(u)}^{-1} \circ D_u \exp_{\mathbf{G}}(u) = \frac{\text{I} - \exp(-\text{ad}_u^\vee)}{\text{ad}_u^\vee}, \quad (1)$$

for any $u \in \mathfrak{g}$ in the domain of $\exp_{\mathbf{G}}$. The left trivialised Jacobian is referred to as the *right* Jacobian in much of the key literature [11]. The *left* Jacobian can be obtained analogously by right trivialisation in the construction (1).

For two functions $A, B : \mathbb{R}^m \rightarrow \mathbb{R}^n$ then we write $A(u) = B(u) + \mathbf{o}(|u|^k)$ if $\lim_{|u| \rightarrow 0} (A(u) - B(u))/|u|^k = 0$. Note that this definition is invariant under a change of coordinates.

2.2 Connection, curvature and parallel transport

For an arbitrary manifold, an affine connection is a geometric structure that is additional to the underlying differential structure. On Lie groups, however, there is a *canonical connection* (also known as the Cartan-Schouten (0)-connection), which is defined as the only affine connection that is left-invariant, torsion-free, and has geodesics given by the Lie exponential [12]. In the case of compact Lie groups, such as $\mathbf{SO}(3)$, the canonical connection coincides with the Levi-Civita connection associated with the bi-invariant metric. In the rest of this paper, we will consider only the Cartan-Schouten (0)-connection as the default choice of affine connection on Lie groups.

Let $X, Y \in \mathfrak{X}(\mathbf{G})$ be two left-invariant vector fields on \mathbf{G} , corresponding to $x, y \in \mathfrak{g}$ respectively. For the (0)-connection, the covariant derivative $\nabla_X Y$ is given by

$$\nabla_X Y = \frac{1}{2} [X, Y]. \quad (2)$$

The Riemann curvature tensor $\mathcal{R} : \mathfrak{X}(\mathbf{G}) \times \mathfrak{X}(\mathbf{G}) \times \mathfrak{X}(\mathbf{G}) \rightarrow \mathfrak{X}(\mathbf{G})$ is defined by

$$\mathcal{R}(X, Y)Z = \nabla_X \nabla_Y Z - \nabla_Y \nabla_X Z - \nabla_{[X, Y]} Z.$$

Using the Jacobi identity of the Lie bracket, it is straightforward to verify that on a Lie group \mathbf{G} with (0)-connection (2) the Riemann curvature tensor is given by

$$\mathcal{R}(X, Y)Z = -\frac{1}{4} [[X, Y], Z], \quad (3)$$

where $X, Y, Z \in \mathfrak{X}(\mathbf{G})$ are left-invariant vector fields.

For the (0)-connection, the curve $\exp_{\mathbf{G}}(tu)$ for $u \in \mathfrak{g}$ arbitrary, is a geodesic. Given $w \in \mathfrak{g}$, the parallel translation of w along $\exp_{\mathbf{G}}(tu)$ is given by

$$\mathbf{P}_{tu}[w] = \text{DL}_{\exp_{\mathbf{G}}(tu)} \circ \text{Ad}_{\exp_{\mathbf{G}}(-\frac{1}{2}u)}(w) \in \text{T}_{\exp_{\mathbf{G}}(tu)} \mathbf{G}.$$

By identifying $\text{T}_{\exp_{\mathbf{G}}(tu)} \mathbf{G}$ with \mathbb{R}^n via the left trivialisation and the ‘ \vee ’ operator, one can define the following *left trivialised parallel transport* map $\mathbf{P}_u^\vee : \mathbb{R}^n \rightarrow \mathbb{R}^n$

$$\mathbf{P}_u^\vee := \text{Ad}_{\exp_{\mathbf{G}}(-\frac{1}{2}u)}^\vee = \exp(-\frac{1}{2} \text{ad}_u^\vee) \in \mathbb{R}^{n \times n}.$$

3 Problem Formulation

In this section, we recall the concept of a concentrated Gaussian distribution on Lie groups [4] and its extended version [8, 9]. We go on to formulate the fusion problem on Lie groups considered in this paper.

3.1 Extended Concentrated Gaussian

For a random variable $g \in \mathbf{G}$, the classical construction of concentrated Gaussian distribution on Lie group [4] is given by

$$p(g; x, \Sigma) = \alpha \exp \left\{ -\frac{1}{2} [\log_{\mathbf{G}}^{\vee}(x^{-1}g)]^{\top} \Sigma^{-1} \log_{\mathbf{G}}^{\vee}(x^{-1}g) \right\},$$

where α is the normalising factor. The parameters $x \in \mathbf{G}$, $\Sigma \in \mathbb{S}_+(n)$ are the *group* mean and covariance respectively, which are defined such that [11]

$$\int_{\mathbf{G}} \log_{\mathbf{G}}^{\vee}(x^{-1}g) p(g) dg = \mathbf{0},$$

and

$$\Sigma = \int_{\mathbf{G}} [\log_{\mathbf{G}}^{\vee}(x^{-1}g)] [\log_{\mathbf{G}}^{\vee}(x^{-1}g)]^{\top} p(g) dg.$$

This construction is equivalent to defining a random variable on \mathbf{G} by

$$g = x \exp_{\mathbf{G}}(\varepsilon), \quad \varepsilon \sim \mathbf{N}(0, \Sigma)$$

where ε is a random variable associated with a normal distribution on \mathbb{R}^n .

In more recent works [8][9], this concept was extended to allow offset mean in the Lie algebra

$$p(g; x, \mu, \Sigma) = \tag{4}$$

$$\alpha \exp \left\{ -\frac{1}{2} [\log_{\mathbf{G}}^{\vee}(x^{-1}g) - \mu^{\vee}]^{\top} \Sigma^{-1} [\log_{\mathbf{G}}^{\vee}(x^{-1}g) - \mu^{\vee}] \right\}, \tag{5}$$

where $x \in \mathbf{G}$ is termed the *reference point*, $\mu \in \mathfrak{g}$ is termed the *mean*. The extended concentrated Gaussian distribution is equivalent to defining a random variable

$$g = x \exp_{\mathbf{G}}(\varepsilon), \quad \varepsilon \sim \mathbf{N}(\mu^{\vee}, \Sigma).$$

The extended concentrated Gaussian makes the role of the reference point as the origin of local coordinates on the group, as separate from the mean of the underlying distribution, clear. Both the classical and extended concentrated Gaussian distributions are approximations of the true distributions of a random variable on a Lie group after fusion. The key question is not whether they are the correct model, but rather how accurately they can represent real distributions.

3.2 Fusion problem

Consider the scenario where one has n independent unbiased estimates $g \sim \mathbf{N}_{x_i}(0, \Sigma_i) = p_i(g|x_i, \Sigma_i)$ for a random variable $g \in \mathbf{G}$. Each estimate is derived from independent data captured in the parameters $x_i \in \mathbf{G}$ and $\Sigma_i \in \mathbb{S}_+(n)$. We want to derive a fused estimate $g \sim \mathbf{N}_{x^+}(0, \Sigma^+) \approx p(g|x_1, \dots, x_n, \Sigma_1, \dots, \Sigma_n)$ based on all the available data. In classical Gaussian fusion the solution is the product of the Gaussians and can be written as a Gaussian. However, the product of concentrated Gaussians is not a concentrated Gaussian and the goal of this work is to find the parameters $x^+ \in \mathbf{G}$ and $\Sigma^+ \in \mathbb{S}_+(n)$ that best approximate the fused density.

4 Coordinate Representation of Concentrated Gaussian Distribution

In this section, we present the technical results including expressing the concentrated Gaussian in different reference coordinates using Jacobian mapping, and a novel approach to approximate the Jacobian with the curvature structure of the exponential map.

4.1 Changing reference

The extended concentrated Gaussian (4) is introduced in order to provide a model to express Gaussians around a reference that does not coincide with their mean. A detailed formulation is provided in the following lemma.

Lemma 4.1. *Given an extended concentrated Gaussian distribution $p(g) = \mathbf{N}_{x_1}(\mu_1, \Sigma_1)$ on \mathbf{G} and a point $x_2 \in \mathbf{G}$ then the concentrated Gaussian $q(g) = \mathbf{N}_{x_2}(\mu_2, \Sigma_2)$ with parameters*

$$\mu_2 = \log_{\mathbf{G}}(x_2^{-1}x_1 \exp_{\mathbf{G}}(\mu_1)) \tag{6}$$

$$\Sigma_2 = J_{\mu_2}^{-1} J_{\mu_1} \Sigma_1 J_{\mu_1}^{\top} J_{\mu_2}^{-\top} \tag{7}$$

minimises the Kullback-Leibler divergence of $q(g)$ with respect to $p(g)$ up to second-order linearisation error.

Proof. The Kullback-Leibler divergence between $p(g)$ and $q(g)$ is given by

$$\begin{aligned} \text{KL}(p||q) &= \mathbb{E}_p[\log(p) - \log(q)] \\ &= C_p + \frac{n}{2} \log(2\pi) + \frac{1}{2} \log \det(\Sigma_2) \\ &\quad + \frac{1}{2} \mathbb{E}_p[(\log_{\mathbf{G}}^{\vee}(x_2^{-1}g) - \mu_2^{\vee})^{\top} \Sigma_2^{-1} (\log_{\mathbf{G}}^{\vee}(x_2^{-1}g) - \mu_2^{\vee})], \end{aligned}$$

where C_p is the negative entropy of $p(g)$. The mean μ_2 in (6) can be derived by assuming that

$$x_1 \exp_{\mathbf{G}}(\mu_1) = x_2 \exp_{\mathbf{G}}(\mu_2).$$

Take the derivative of $\text{KL}(p||q)$ with respect to Σ_2 and the critical point is given by

$$\Sigma_2 = \mathbb{E}_p[(\log_{\mathbf{G}}^{\vee}(x_2^{-1}g) - \mu_2^{\vee}) (\log_{\mathbf{G}}^{\vee}(x_2^{-1}g) - \mu_2^{\vee})^{\top}]. \quad (8)$$

Define $\phi_1 : \mathbf{G} \rightarrow \mathbb{R}^n$ and $\phi_2 : \mathbf{G} \rightarrow \mathbb{R}^n$ as

$$\begin{aligned} \phi_1(g) &:= \log_{\mathbf{G}}^{\vee}(x_1^{-1}g) - \mu_1^{\vee}, \\ \phi_2(g) &:= \log_{\mathbf{G}}^{\vee}(x_2^{-1}g) - \mu_2^{\vee}. \end{aligned}$$

For an arbitrary $g \in \mathbf{G}$ one has

$$\phi_1(g) = \log_{\mathbf{G}}^{\vee}(x_1^{-1}x_2 \exp_{\mathbf{G}}(\phi_2(g) + \mu_2^{\vee})) - \mu_1^{\vee} \quad (9)$$

Taking the Taylor series of (9) at $\phi_2(g) = 0_{n \times 1}$ up to first order yields:

$$\begin{aligned} \phi_1(g) &\approx D \log_{\mathbf{G}}^{\vee}(\exp_{\mathbf{G}}(\mu_1)) \cdot DL_{x_1^{-1}x_2} \cdot D \exp_{\mathbf{G}}(\mu_1) [\phi_1(g)] \\ &= D \log_{\mathbf{G}}^{\vee}(\exp_{\mathbf{G}}(\mu_1)) \cdot DL_{\exp_{\mathbf{G}}(\mu_1)} \cdot DL_{(x_1 \exp_{\mathbf{G}}(\mu_1))^{-1}} \cdot \\ &\quad DL_{x_2 \exp_{\mathbf{G}}(\mu_2)} \cdot DL_{\exp_{\mathbf{G}}(-\mu_2)} \cdot D \exp_{\mathbf{G}}(\mu_2) [\phi_2(g)] \\ &= (DL_{\exp_{\mathbf{G}}(-\mu_1)} \cdot D \exp_{\mathbf{G}}(\mu_1))^{-1} \cdot \\ &\quad DL_{\exp_{\mathbf{G}}(-\mu_2)} \cdot D \exp_{\mathbf{G}}(\mu_2) [\phi_2(g)] \\ &= J_{\mu_1}^{-1} J_{\mu_2} [\phi_2(g)]. \end{aligned}$$

Note that due to the choice of μ_2 in (6), $D \log_{\mathbf{G}}^{\vee}$ in the Taylor series is evaluated at $\exp_{\mathbf{G}}(\mu_1)$. Substitute the result into (8):

$$\Sigma_2 = \mathbb{E}_p[\phi_2(g) \phi_2(g)^{\top}] \approx J_{\mu_2}^{-1} J_{\mu_1} \Sigma_1 J_{\mu_1}^{\top} J_{\mu_2}^{-\top},$$

which follows from the definition of $\Sigma_1 = \mathbb{E}_p[\phi_1(g) \phi_1(g)^{\top}]$. \square

Note that (6) is the exact coordinates of the mean in the new coordinates, and that (7) is the covariance conjugated by the Jacobian of the change-of-coordinates maps. That is, Lemma 4.1 can be thought of as transforming a Gaussian distribution under a non-linear change of coordinates.

Remark 4.2. *In the special case when $\mu_1 = 0$ or $\mu_2 = 0$, the covariance Σ_2 is given by*

$$\Sigma_2 = J_{\mu_2}^{-1} \Sigma_1 J_{\mu_2}^{-\top} \quad \text{or} \quad \Sigma_2 = J_{\mu_1} \Sigma_1 J_{\mu_1}^{\top},$$

respectively. Both cases can happen in the fusion problem, as discussed in Sec 5.

4.2 Approximation with Curvature

The result in Lemma 4.1 relies on computing the linear map $J_{\mu} : \mathbb{R}^n \rightarrow \mathbb{R}^n$. However, as presented in Sec 2.1, this map is transcendental and except in special cases, must be computed using approximations of infinite power series. Analytic formulae in terms of classical trigonometric functions such as \cos and \sin exist for a limited range of Lie groups such as $\mathbf{SO}(3)$ and $\mathbf{SE}(3)$. In this section we propose a method to approximate the Jacobian using geometric structure of the Lie group.

Theorem 4.3. *For any $u \in \mathfrak{g}$, the Jacobian map $J_u : \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined in (1) satisfies*

$$\begin{aligned} J_u[w^{\vee}] &= \mathbf{P}_u^{\vee} [w + \frac{1}{6} \mathcal{R}(u, w)u]^{\vee} + \mathbf{o}(|u^{\vee}|^3) \\ &\approx \text{Ad}_{\exp_{\mathbf{G}}(-\frac{1}{2}u)}^{\vee} \left(I + \frac{1}{24} \text{ad}_u^{\vee 2} \right) [w^{\vee}], \end{aligned}$$

for all $w \in \mathfrak{g}$.

Proof. Define the geodesic $\gamma(t) := \exp_{\mathbf{G}}(tu)$ with $t \in I$, where $I \subseteq \mathbb{R}$ is an open interval containing 0. Define the Jacobi field $J_w(t)$ to be the unique solution of the Jacobi equation [13]

$$D_t^2 J_w(t) + \mathcal{R}(J_w(t), \dot{\gamma}(t))\dot{\gamma}(t) = 0$$

for $J_w(0) = 0$ and $D_t J_w(0) = w$. For $t \in I$, one has [13, Theorem 3.1]

$$D_{tu} \exp_{\mathbf{G}}[w] = \frac{1}{t} J_w(t).$$

Consider the following map: $t \rightarrow (\mathbf{P}_{tu})^{-1} \left(\frac{1}{t} J_w(t) \right)$. Note that taking the Taylor expansion of this map around $t = 0$ is equivalent to studying the Taylor expansion of $(\mathbf{P}_u)^{-1} \circ D_u \exp_{\mathbf{G}}[w]$ around $u = 0$. Applying [14, Theorem A.2.9] yields

$$(\mathbf{P}_u)^{-1} \circ D_u \exp_{\mathbf{G}}[w] = w + \frac{1}{6} \mathcal{R}(u, w)u + \mathbf{o}(|u|^3).$$

where the order \mathbf{o} approximation is understood over a basis in the Lie-algebra \mathfrak{g} . Take the first order approximation,

$$\begin{aligned} J_u[w^\vee] &= \text{DL}_{\exp_{\mathbf{G}}(-u)} \circ D_u \exp_{\mathbf{G}}[w]^\vee \\ &= \mathbf{P}_u^\vee \left[w + \frac{1}{6} \mathcal{R}(u, w)u \right]^\vee + \mathbf{o}(|u^\vee|^3) \\ &\approx \mathbf{P}_u^\vee \left[w - \frac{1}{24} [[u, w], u] \right]^\vee && \text{(Equation 3)} \\ &= \mathbf{P}_u^\vee \left[w + \frac{1}{24} \text{ad}_u^2(w) \right]^\vee && \text{(Anti-symmetry of } [\cdot, \cdot] \text{)} \\ &= \mathbf{P}_u^\vee \left(\mathbf{I} + \frac{1}{24} \text{ad}_u^{\vee 2} \right) [w^\vee] && \text{(Bilinearity of } [\cdot, \cdot] \text{)} \\ &= \text{Ad}_{\exp_{\mathbf{G}}(-\frac{1}{2}u)}^\vee \left(\mathbf{I} + \frac{1}{24} \text{ad}_u^{\vee 2} \right) [w^\vee]. && \text{(Definition of } \mathbf{P}_u^\vee \text{)} \end{aligned}$$

This completes the proof. \square

We show that this approximation captures enough of the necessary information of the Jacobian of the exponential map to achieve good fusion results at a low computational cost.

5 Fusion on Lie groups

In this section, we propose a general methodology to fuse multiple concentrated Gaussians on Lie groups. The proposed methodology has 3 steps. The first step is to compute a reference point $\hat{x} \in \mathbf{G}$. In the second step, one of the approximation methods is used to express the independent concentrated Gaussians provided as data as extended concentrated Gaussians with respect to the chosen reference point. In these coordinates, classical Gaussian fusion is applied. The last step is to rewrite the fused extended concentrated Gaussian as a concentrated Gaussian around the group element corresponding to its mean.

Step 1: Reference

The goal of the first step of the methodology is to choose a reference point $\hat{x} \in \mathbf{G}$ as close to the true mean of the fused distribution as possible for the least reasonable computational cost. This point will be used as the reference point for the approximation of the independent concentrated Gaussians. The closer \hat{x} is to the correct group-mean, the less approximation error will be incurred before the full fusion process is undertaken. However, spending excessive computation at this point is wasted since the independent concentrated Gaussians are defined at different points on the Lie group anyway and, as long as \hat{x} is roughly central to the data, the particular choice of reference point will make little difference to the approximation.

Choosing a reference point is common to most Lie-group fusion algorithms. In [9], and in general for Extended Kalman Filters (EKF), the mean of the prior distribution is used as reference. For the fusion of Gaussians, the mean x_1 of the first independent distribution can be used as reference. In [7], an initial estimate \hat{x} is derived by using the naive fusion method. Such choice can be iterated to achieve better accuracy at a higher computational cost. The authors in [6] use an optimization process to compute \hat{x} .

Naive Fusion: We use a simple algorithm to choose \hat{x} that will also act as benchmark for the comparison study in Section 6. Consider n independent estimates $\mathbf{N}_{x_i}(0, \Sigma_i)$ for $i = 1, \dots, n$ of a random variable $g \in \mathbf{G}$. Consider exponential coordinates on the Lie group \mathbf{G} around the origin. Approximate

$$\mathbf{N}_{x_i}(0, \Sigma_i) \approx \mathbf{N}_{\text{id}}(\log_{\mathbf{G}}(x_i), \Sigma_i)$$

by an extended concentrated Gaussian in origin coordinates without any consideration of the change of coordinates on the covariance Σ_i . The distributions are now Gaussian in a single set of coordinates (the Lie algebra) and classical fusion is used to estimate the mean of the distribution

$$\hat{\Sigma} = \left(\sum_{i=1}^n \Sigma_i^{-1} \right)^{-1}, \quad \hat{\mu}^\vee = \hat{\Sigma} \sum_{i=1}^n \Sigma_i^{-1} \log_{\mathbf{G}}^\vee(x_i).$$

Set $\hat{x} = \exp_{\mathbf{G}}(\hat{\mu})$, the final distribution is $\mathbf{N}_{\hat{x}}(0, \hat{\Sigma})$.

Step 2: Fusion

Consider independent estimates $p(g|x_i, \Sigma_i) = \mathbf{N}_{x_i}(0, \Sigma_i)$ for $i = 1, \dots, n$. We approximate each distribution by an extended concentrated Gaussian

$$p(g|x_i, \Sigma_i) \approx \mathbf{N}_{\hat{x}}(\mu_i, \hat{\Sigma}_i)$$

where $\mu_i = \log_{\mathbf{G}}(\hat{x}^{-1}x_i)$ and $\hat{\Sigma}_i$ is given by the chosen approximation scheme.

Full Jacobian: The most direct approximation is provided by applying Lemma 4.1:

$$\hat{\Sigma}_i = J_{\mu_i}^{-1} \Sigma_i J_{\mu_i}^{-\top} \quad (10)$$

where an analytic expression for the inverse Jacobian $J_{\mu_i}^{-1}$ can be computed.

Approximate Jacobian: If no analytic version of the inverse Jacobian $J_{\mu_i}^{-1}$ is available, we can approximate this by Taylor's expansions. We consider both first and second order approximations [6].

Parallel Transport and Curvature: As discussed in Section 4.2 the Jacobian can be approximated by parallel transport $J_{\mu_i}^{-1} \approx \mathbf{P}_{\mu_i}^{-1}$ or recalling Theorem 4.3 by parallel transport and curvature

$$J_{\mu_i}^{-1} \approx \left(\mathbf{I} + \frac{1}{24} \text{ad}_{\mu_i}^2 \right)^{-1} \mathbf{P}_{\mu_i}^{-1} \approx \left(\mathbf{I} - \frac{1}{24} \text{ad}_{\mu_i}^2 \right) \mathbf{P}_{\mu_i}^{-1}.$$

Once the Gaussians are written in the same coordinates the distributions are fused using classical Gaussian fusion.

$$\Sigma_\diamond = \left(\sum_{i=1}^n \hat{\Sigma}_i^{-1} \right)^{-1}, \quad \mu^{+\vee} = \Sigma_\diamond \left(\sum_{i=1}^n \hat{\Sigma}_i^{-1} \mu_i^\vee \right). \quad (11)$$

Under the assumption that the distributions are independent, the fused estimate (μ^+, Σ_\diamond) is optimal with respect to multiple criteria such as the weighted least squares error, minimum covariance estimation error and maximum-likelihood estimation [15].

Step 3: Reset

The outcome of the fusion step is an extended concentrated Gaussian with non-zero mean $\mathbf{N}_{\hat{x}}(\mu^+, \Sigma_\diamond)$. If a concentrated Gaussian is required, which is the normal case for most filtering algorithms and estimators, then the extended concentrated Gaussian estimate must be transformed into a concentrated Gaussian around a new group mean. The problem is equivalent to finding $x_+ \in \mathbf{G}$ and $\Sigma_+ \in \mathbb{S}_+(n)$ such that

$$\mathbf{N}_{\hat{x}}(\mu^+, \Sigma_\diamond) \approx \mathbf{N}_{\hat{x}^+}(0, \Sigma^+).$$

This is also a direct application of Lemma 4.1. The new reference point \hat{x}^+ and the covariance Σ^+ are given by

$$\hat{x}^+ = \hat{x} \exp_{\mathbf{G}}(\mu^+),$$

$$\Sigma^+ = J_{\mu^+} \Sigma_\diamond J_{\mu^+}^\top.$$

The implementation of the reset will require either a full analytic version of the Jacobian to be computed, or an approximation to be used based on one of the methods discussed in Step 2. The new distribution $\mathbf{N}_{\hat{x}^+}(0, \Sigma^+)$ is a zero-mean concentrated Gaussian.

6 Simulation

To numerically evaluate the proposed methods we undertake a simulation on the special orthogonal group $\mathbf{SO}(3)$. This group provides an intuitive example for which the formula are relatively straightforward to understand. Following the experimental study in [7] we consider two zero-mean concentrated Gaussians on $\mathbf{SO}(3)$, denoted by $p_1(g) = \mathbf{N}_{x_1}(0, \Sigma_1)$ and $p_2(g) = \mathbf{N}_{x_2}(0, \Sigma_2)$. The parameters are chosen as

$$x_1 = \exp_{\mathbf{G}} \left(\frac{\gamma}{\sqrt{3}} \begin{bmatrix} 1.0 \\ 1.0 \\ -1.0 \end{bmatrix} \right), \quad x_2 = \exp_{\mathbf{G}} \left(\frac{\gamma}{\sqrt{2}} \begin{bmatrix} 1.0 \\ -1.0 \\ 0.0 \end{bmatrix} \right)$$

and

$$\Sigma_1 = \xi \begin{bmatrix} 1.0 & 0 & 0 \\ 0 & 0.75 & 0 \\ 0 & 0 & 0.5 \end{bmatrix}, \quad \Sigma_2 = \xi \begin{bmatrix} 0.5 & 0 & 0 \\ 0 & 1.0 & 0 \\ 0 & 0 & 0.75 \end{bmatrix}.$$

The scalars γ , and $\xi \in \mathbb{R}_+$ control the distance between means and the concentration of the covariance. As γ increases and ξ decreases the fusion becomes more non-linear.

To evaluate the performance, we use the cost function proposed in [7]. Define $p_{12}(g) := \frac{p_1(g)p_2(g)}{\int_G p_1(h)p_2(h)dh}$ to be the fused probability density of $p_1(g)$ and $p_2(g)$. We use the ℓ_1 metric:

$$C := |p^+ - p_{12}|_{\ell_1} = \int_G |p^+(g) - p_{12}(g)| dg, \quad (12)$$

where $p^+(g) = \mathbf{N}_{\hat{x}^+}(0, \Sigma^+)$ is the estimated distribution using different methods. This metric evaluates the difference between the approximated concentrated Gaussian distribution and the underlying fused density, and can be interpreted as the *total variation distance* between them.

It is implemented by uniform sampling over a bounded domain on $\mathfrak{so}(3)$.

We run the simulation with different combinations of γ and ξ , where both parameters are varied from 0.1 to 1.8. Each result is averaged over a Monte Carlo simulation with 500 runs where the covariance matrices are rotated by random rotation matrices.

We present the main results demonstrating the performance of different approximation methods. The Naive method is implemented directly. Approximate Jacobian and parallel transport methods proposed in Section 5 are implemented as described with the same Jacobian approximation used in the reset. For comparison, we include the BCH-based methods (first and second order) proposed in [7]. We use the standard minimizer in the SciPy library to implement the optimisations required by these methods. The BCH methods do not require reset. All methods except the naive one use the same naive posterior as the initial guess \hat{x} .




Figure 2: Estimation error using different approximation methods with fixed $\gamma = 1.0$ or $\xi = 1.0$.

Fig 1 plots the average error against relative processing time of each method, averaged over all the parameter values of ξ and γ . To account for the dependence on the computer hardware used, we show the ratio of processing time for each algorithm as compared to the naive fusion. Since all algorithms require the naive fusion as a first estimate of the reference point, this ratio is always greater than one. In Fig 2, we present the average estimation error of different methods when either γ or ξ is fixed at 1.0 and the other parameter changes. Fig 3 shows how the error in each method varies with different combinations of the parameters γ and ξ .

In general, the naive fusion method has the worst performance among all the methods, as expected. The first order methods (BCH 1st order and Jac 1st order) perform marginally better but still show significant approximation error, and especially so




Figure 3: Heatmaps showing the estimation error with different approximation methods. The full Jacobian and PTC perform identically, while second-order BCH method has slightly better performance for small γ .

for large values of γ . Interestingly, the first order BCH method is also highly sensitive to small ξ . The parallel transport method can also be thought of as a first order method. It performs better than the other first order methods but incurs more performance error than the remaining methods.

The last four methods all use higher order information in the Jacobian approximation (higher order terms in the BCH expansion for the BCH method). Clearly, the full Jacobian outperforms the second order Taylor approximation, although the difference is less significant than the difference with the first order Jacobian. The full order Jacobian, parallel transport with curvature (PTC) and BCH 2nd achieve very similar average error. The combination of parallel transport and curvature clearly captures the major nonlinearities in the full Jacobian without requiring computation of the analytic form of the Jacobian. This is particularly of interest for Lie groups where a closed-form expression of the Jacobian is not available; the parallel transport only requires computing a matrix exponential, which has an efficient implementation in many linear algebra programming libraries. The second order BCH method achieves the lowest average error in our simulations. However, the BCH-based methods do not admit closed-form solutions and can only be implemented with an optimization process, resulting in much higher computational cost (> 100 times greater than the PTC method), as shown in Fig 1. These results are based on applying a standard optimisation routine, however, even a tailored optimisation algorithm is unlikely to significantly reduce this computational discrepancy.

7 Conclusion

This paper presents a general design methodology for fusion of independent concentrated Gaussian distributions on Lie groups from a geometric perspective. It is shown that by transforming a collection of distributions into a single, unified set of coordinates, the fusion problem can be solved using the same methodology as in the Euclidean case. In the simulation, we present the effectiveness of the different methods proposed, and, in particular, show that the parallel transport with curvature correction can achieve good performance at a low computational cost, while using only functions that are available in most linear algebra libraries.

References

- [1] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, Mar. 1960. 1
- [2] A. Barrau and S. Bonnabel, "The Invariant Extended Kalman Filter as a Stable Observer," *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1797–1812, Apr. 2017. 1
- [3] R. Mahony, P. van Goor, and T. Hamel, "Observer Design for Nonlinear Systems with Equivariance," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. 1, pp. 221–252, May 2022. 1
- [4] Y. Wang and G. Chirikjian, "Error propagation on the Euclidean group with applications to manipulator kinematics," *IEEE Transactions on Robotics*, vol. 22, no. 4, pp. 591–602, Aug. 2006. 1, 4
- [5] G. Bourmaud, R. Mégret, A. Giremus, and Y. Berthoumieu, "From Intrinsic Optimization to Iterated Extended Kalman Filtering on Lie Groups," *Journal of Mathematical Imaging and Vision*, vol. 55, no. 3, pp. 284–303, Jul. 2016. 1
- [6] T. D. Barfoot and P. T. Furgale, "Associating Uncertainty With Three-Dimensional Poses for Use in Estimation Problems," *IEEE Transactions on Robotics*, vol. 30, no. 3, pp. 679–693, Jun. 2014. 1, 6, 7
- [7] K. C. Wolfe and M. Mashner, "Bayesian Fusion on Lie Groups," *Journal of Algebraic Statistics*, vol. 2, no. 1, Apr. 2011. 1, 2, 6, 7, 8
- [8] Y. Ge, P. van Goor, and R. Mahony, "Equivariant Filter Design for Discrete-time Systems," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, Dec. 2022, pp. 1243–1250. 1, 4
- [9] Y. Ge, P. Van Goor, and R. Mahony, "A Note on the Extended Kalman Filter on a Manifold," in *2023 62nd IEEE Conference on Decision and Control (CDC)*, Dec. 2023, pp. 7687–7694. 1, 4, 6
- [10] S. Helgason, *Differential Geometry, Lie Groups, and Symmetric Spaces*. Academic Press, Feb. 1979. 3
- [11] G. S. Chirikjian, *Stochastic Models, Information Theory, and Lie Groups, Volume 2: Analytic Methods and Modern Applications*. Springer Science & Business Media, Nov. 2011. 3, 4
- [12] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry, Volume 2*. John Wiley & Sons, Feb. 1996. 3
- [13] S. Lang, *Fundamentals of Differential Geometry*. Springer Science & Business Media, Dec. 2012. 6
- [14] S. Waldmann, "Geometric wave equations," *arXiv preprint arXiv:1208.4706*, 2012. 6
- [15] F. C. Schwegpe, *Uncertain Dynamic Systems*. Prentice-Hall, 1973. 7