

Geometric Data Fusion for Collaborative Attitude Estimation

Yixiao Ge* Behzad Zamani** Pieter van Goor*
Jochen Trumpf* Robert Mahony*

* *School of Engineering, Australian National University, Canberra, Australia (e-mail: {first name}.{last name}@anu.edu.au).*

** *Department of Electrical and Electronic Engineering, University of Melbourne, Melbourne, Australia (e-mail: Behzad.Zamani@unimelb.edu.au)*

Abstract: In this paper, we consider collaborative attitude estimation problem of a multi-agent system. The agents are equipped with sensors that provide directional measurements and relative attitude measurements. We present a bottom-up approach where each agent runs an extended Kalman filter (EKF) locally using directional measurements and augments this with relative attitude measurements provided by neighbouring agents. The covariance estimates of the relative attitude measurements are geometrically corrected to compensate for relative attitude between the agent that makes the measurement and the agent that uses the measurement before being fused with the local estimate using convex combination ellipsoid (CCE) method to avoid data incest. Simulations are undertaken to numerically evaluate the performance of the proposed algorithm.

Keywords: data fusion, geometric approaches, stochastic modelling, fusion, covariance matrices

1. INTRODUCTION

The problem of collaborative state estimation over sensor networks has drawn significant attention in the past 20 years (Roumeliotis and Bekey, 2002). In this problem, different agents share measurements and state-estimates to improve overall state estimation. Sharing data in this way, however, introduces the possibility of data incest (Julier and Uhlmann, 1997). To see this, consider a network of individual estimators each estimating some states while communicating with other nodes on the network. Information received from other agents will depend on information transmitted by the agent itself in preceding communications, potentially reinforcing its own hypothesis and increasing the risk of overconfidence in the resulting state estimates (Julier and Uhlmann, 1997).

To overcome these challenges there are two main approaches: a top down approach where the full state estimation is formulated as a joint estimation problem and then the computation is decentralised to each node (Carrillo-Arce et al. (2013); Luft et al. (2018); Zhu and Ren (2020)), and the bottom up approach, where each agents runs an independent estimator locally and fuses data from other agents taking precautions to avoid data incest (Olfati-Saber, 2005, 2006). The key enabling step in the bottom up approach is a methodology to provide safe fusion of correlated data into a local agent state estimation such that it retains the common uncertainty of the original random variables. This problem has been studied since the 60s (Schweppe, 1968; Kahan, 1968; Bertsekas and Rhodes, 1971). In more recent work (Julier and Uhlmann (1997), Julier and Uhlmann (2017)) Julier and Uhlmann proposed the Covariance Intersection (CI) algorithm which restricts

the fusion problem to a family of convex combinations of the inverse covariance matrices and is the most commonly used data fusion method in multi-agent problems. The CI algorithm, however, is known to be too conservative in certain situations (Julier and Uhlmann, 2007). The Inverse Covariance Intersection (ICI) method (Noack et al., 2017) computes the maximum possible common information shared by the estimates to be fused, and is known to be less conservative than the CI method. An alternative solution is the Convex Combination Ellipsoid (CCE) fusion method which arises from the set-theoretic fusion technique (Schweppe, 1968; Ros et al., 2002). The CCE method shares a similar structure with CI, however it improves the tightness of the fusion result while avoiding unnecessary uncertainties as the byproduct of the fusion process (Zamani et al., 2023).

All these fusion algorithms are originally formulated in global Euclidean space, and there have been many attempts to adapt these classical methods to systems that live on smooth manifolds, particularly Lie groups. One popular approach is to consider the fusion problem as finding the optimal mode of the posterior distribution by solving an optimization problem (Bourmaud et al., 2016; Barfoot, 2017). In (Wolfe and Mashner, 2011; Chirikjian and Kobilarov, 2014), the authors solved the fusion problem by posing a set of algebraic equations using the Baker-Campbell-Hausdorff formula. Recent work in equivariant filter theory (Mahony et al., 2020) and geometric extended Kalman filtering (Mueller et al., 2017; Ge et al., 2023) has provided a strong understanding of the geometry of filtering and data fusion. In particular, these works demonstrate that it matters in which coordinates the generative noise

model for a measurement is posed and provides formulae and methodology to transform covariance from one set of coordinates to another (Ge et al., 2023).

In this paper, we consider a bottom up approach to the problem of collaborative attitude estimation, where each node estimates its own attitude as well as taking relative measurements of other nodes. The problem is posed on the Lie group $\mathbf{SO}(3)$ representing the attitude of a single agent of interest, termed the *ego-agent*. The information used are local directional measurements, angular velocity, and a noisy relative attitude measurement of the ego-agent as observed by a neighbouring *altruistic-agent* along with the altruistic agent’s own state estimate (estimated attitude and its error covariance). This relative attitude measurement, combined with the altruistic agent’s state estimate, is effectively an attitude measurement of the ego-agent, and can be fused into the ego-agent’s state estimation, at the appropriate point, in the filter algorithm. However, in a collaborative communication scenario, the altruistic agent’s state estimate is itself dependent on shared information from the ego-agent, and this relative pose measurement should not be treated as an independent measurement. Furthermore, the attitude measurement is observed from the perspective of the altruistic agent and is written in these coordinates. The covariance of the measurement must be transformed into the ego-state coordinates to avoid incorrect inference. The contribution of the paper is to combine the geometric modification of the covariance into the correct coordinates with the CCE fusion method to obtain a high-performance bottom up collaborative state estimation scheme for multi-agent attitude estimation. We provide a Monte-Carlo simulation to demonstrate the performance of the proposed estimation algorithm together with the geometric modifications.

2. PRELIMINARIES

2.1 Special orthogonal group $\mathbf{SO}(3)$

Attitude of an agent is represented as a rotation matrix R in the special orthogonal group $R \in \mathbf{SO}(3)$. The identity element of $\mathbf{SO}(3)$, denoted id , is the identity matrix. Given arbitrary $X, Y \in \mathbf{SO}(3)$, the left and right translations are denoted by $L_X(Y) := XY$ and $R_X(Y) := YX$. The Lie algebra $\mathfrak{so}(3)$ of $\mathbf{SO}(3)$ consists of all skew-symmetric matrices of the form

$$u^\wedge = \begin{pmatrix} 0 & -u_3 & u_2 \\ u_3 & 0 & -u_1 \\ -u_2 & u_1 & 0 \end{pmatrix},$$

and is isomorphic to the vector space \mathbb{R}^3 . We use the wedge $(\cdot)^\wedge : \mathbb{R}^3 \rightarrow \mathfrak{so}(3)$ and vee $(\cdot)^\vee : \mathfrak{so}(3) \rightarrow \mathbb{R}^3$ operators to map between the Lie algebra and vector space. The Adjoint map for the group $\mathbf{SO}(3)$, $\text{Ad}_X : \mathfrak{so}(3) \rightarrow \mathfrak{so}(3)$ is defined by

$$\text{Ad}_X[u^\wedge] = Xu^\wedge X^\top,$$

for every $X \in \mathbf{SO}(3)$ and $u^\wedge \in \mathfrak{so}(3)$. Given particular wedge and vee maps, the Adjoint matrix is defined as the map $\text{Ad}_X^\vee : \mathbb{R}^3 \rightarrow \mathbb{R}^3$

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

The adjoint map for the Lie algebra $\text{ad}_{u^\wedge} : \mathfrak{so}(3) \rightarrow \mathfrak{so}(3)$ is given by

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

We define the adjoint matrix $\text{ad}_u^\vee : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ to be:

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

Let $\exp_{\mathbf{G}} : \mathfrak{so}(3) \rightarrow \mathbf{SO}(3)$ denote the matrix exponential (\mathbf{G} denotes the $\mathbf{SO}(3)$ group) and note that this is a one-parameter Lie-subgroup. In order to improve the analogy to the fusion literature that is usually written in \mathbb{R}^n coordinates we will use the \boxplus (‘boxplus’) operator for the exponential map

$$X \boxplus u = X \exp_{\mathbf{G}}(u^\wedge),$$

for $X \in \mathbf{SO}(3)$ to represent the state and $u \in \mathbb{R}^3$ to represent a certain noise process. Let $\mathbf{SO}^\circ(3) \subset \mathbf{SO}(3)$ be the subset of $\mathbf{SO}(3)$ where the exponential map is invertible and note that $\mathbf{SO}^\circ(3)$ is almost all of $\mathbf{SO}(3)$, excluding only those points with a rotation of π radians. The logarithm map $\log_{\mathbf{G}} : \mathbf{SO}^\circ(3) \rightarrow \mathfrak{so}(3)$ and $\log_{\mathbf{G}}^\vee : \mathbf{SO}^\circ(3) \rightarrow \mathbb{R}^3$ is well defined on $\mathbf{SO}^\circ(3)$.

The Jacobian $\mathbf{J}_{u^\wedge} : \mathfrak{so}(3) \rightarrow \mathfrak{so}(3)$ is defined to be the left-trivialised directional derivative of $\exp_{\mathbf{G}} : \mathfrak{so}(3) \rightarrow \mathbf{SO}(3)$ at a point $u^\wedge \in \mathfrak{so}(3)$ on $\mathbf{SO}(3)$. Given an arbitrary $w^\wedge \in \mathfrak{so}(3)$, it satisfies

$$\mathbf{J}_{u^\wedge}[w^\wedge] = \text{DL}_{\exp_{\mathbf{G}}(-u^\wedge)} \cdot \text{D}\exp_{\mathbf{G}}(u^\wedge)[w^\wedge],$$

where the tangent space $\text{T}_{\exp_{\mathbf{G}}(u^\wedge)}\mathbf{SO}(3)$ is isomorphic to $\mathfrak{so}(3)$ via left trivialisation. Equivalently, $\text{D}\exp_{\mathbf{G}}(u^\wedge)[w^\wedge] = \text{DL}_{\exp_{\mathbf{G}}(u^\wedge)}\mathbf{J}_{u^\wedge}[w^\wedge] \in \text{T}_{\exp_{\mathbf{G}}(u^\wedge)}\mathbf{SO}(3)$. Its matrix form, denoted by $\mathbf{J}_u \in \mathbb{R}^{3 \times 3}$, is given by (Chirikjian, 2011)

$$\mathbf{J}_u := I_3 - \frac{1 - \cos \|u\|}{\|u\|^2} u^\wedge + \frac{\|u\| - \sin \|u\|}{\|u\|^3} u^\wedge^2.$$

For an arbitrary $u^\wedge \in \mathfrak{so}(3)$, the inverse of its Jacobian is given by

$$\mathbf{J}_u^{-1} := I_3 + \frac{1}{2}u^\wedge + \left(\frac{1}{\|u\|^2} - \frac{1 + \cos \|u\|}{2\|u\| \sin \|u\|} \right) u^\wedge^2.$$

2.2 Concentrated Gaussians on $\mathbf{SO}(3)$

We use the concept of a concentrated Gaussian to model distributions on $\mathbf{SO}(3)$. For a random variable $X \in \mathbf{SO}(3)$, the probability density function is defined as

$$p(X; \hat{X}, \mu, \Sigma) = \alpha e^{-\frac{1}{2}(\log_{\mathbf{G}}^\vee(\hat{X}^{-1}X) - \mu)^\top \Sigma^{-1}(\log_{\mathbf{G}}^\vee(\hat{X}^{-1}X) - \mu)}, \quad (1)$$

where α is a normalizing factor. The stochastic parameters are $\mu \in \mathbb{R}^3$, a mean vector in local coordinates, and $\Sigma \in \mathbb{S}_+(3)$ a positive-definite symmetric 3×3 covariance matrix parameter. The geometric parameter $\hat{X} \in \mathbf{SO}(3)$ is termed the reference point and plays the role of the origin of the local coordinates. We will term a concentrated Gaussian *zero-mean* if $\mu \equiv 0$. In this case the distribution corresponds to the classical concentrated Gaussian Chirikjian and Kobilarov (2014); Bourmaud et al. (2016) where one can think of the reference point $\hat{X} \in \mathbf{SO}(3)$ as a sort of ‘geometric’ mean. We will write $X \sim \mathbf{N}_{\hat{X}}(\mu, \Sigma)$ for the random variable $X \in \mathbf{SO}(3)$.

Lemma 2.1. Given an arbitrary concentrated Gaussian distribution $p(X) = \mathbf{N}_{X_1}(\mu_1, \Sigma_1)$ on $\mathbf{SO}(3)$, then the zero-mean concentrated Gaussian distribution $q(X) = \mathbf{N}_{X_2}(0, \Sigma_2)$ with parameters

$$X_2 = X_1 \exp_{\mathbf{G}}(\mu_1)$$

$$\Sigma_2 = \mathbf{J}_{\mu_1} \Sigma_1 \mathbf{J}_{\mu_1}^\top$$

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

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

$$\begin{aligned} \text{KL}(p||q) &= \mathbb{E}_p[\log(p) - \log(q)] \\ &= C_p + \frac{n}{2} \log(2\pi) + \frac{1}{2} \log \det(\Sigma_2) \\ &\quad + \frac{1}{2} \mathbb{E}_p[\log_{\mathbf{G}}^{\vee}(X_2^{-1}X)^\top \Sigma_2^{-1} \log_{\mathbf{G}}^{\vee}(X_2^{-1}X)], \end{aligned}$$

where C_p is the negative entropy of $p(X)$. Taking the derivative of $\text{KL}(p||q)$ with respect to Σ_2 in the direction u yields

$$D_{\Sigma_2} \text{KL}(p(X)||q(X))[u] = \frac{1}{2} \text{tr}(\Sigma_2^{-1}u - \Sigma_2^{-1} \mathbb{E}_p[\log_{\mathbf{G}}^{\vee}(X_2^{-1}X) \log_{\mathbf{G}}^{\vee}(X_2^{-1}X)^\top] \Sigma_2^{-1}u).$$

The critical point is given by

$$\Sigma_2 = \mathbb{E}_p[\log_{\mathbf{G}}^{\vee}(X_2^{-1}X) \log_{\mathbf{G}}^{\vee}(X_2^{-1}X)^\top]. \quad (2)$$

Define $\phi_1 : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ and $\phi_2 : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ as

$$\begin{aligned} \phi_1(X) &:= \log_{\mathbf{G}}^{\vee}(X_1^{-1}X) - \mu_1, \\ \phi_2(X) &:= \log_{\mathbf{G}}^{\vee}(\exp_{\mathbf{G}}(-\mu_1)X_1^{-1}X), \end{aligned}$$

one has

$$\phi_2(X) = \log_{\mathbf{G}}^{\vee}(\exp_{\mathbf{G}}(\mu_1)^{-1} \exp_{\mathbf{G}}(\phi_1(X) + \mu_1)). \quad (3)$$

Taking the Taylor series of (3) at $\phi_1(X) = 0$ up to first order yields:

$$\begin{aligned} \phi_2(X) &\approx D \log_{\mathbf{G}}^{\vee}(\text{id}) \circ D L_{\exp_{\mathbf{G}}(\mu_1)^{-1}} \circ D \exp_{\mathbf{G}}(\mu_1)[\phi_1(X)] \\ &= J_{\mu_1} \phi_1(X). \end{aligned}$$

Substitute the result into (2):

$$\begin{aligned} \Sigma_2 &= \mathbb{E}_p[\phi_2(X) \phi_2(X)^\top] \\ &\approx \mathbb{E}_p[J_{\mu_1} \phi_1(X) (J_{\mu_1} \phi_1(X))^\top] \\ &= J_{\mu_1} \mathbb{E}_p[\phi_1(X) \phi_1(X)^\top] J_{\mu_1}^\top \\ &= J_{\mu_1} \Sigma_1 J_{\mu_1}^\top, \end{aligned}$$

where the last equality follows from the definition of $\Sigma_1 = \mathbb{E}_p[\phi_1(X) \phi_1(X)^\top]$.

Corollary 2.2. Given an arbitrary zero-mean concentrated Gaussian distribution $p(X) = \mathbf{N}_{X_1}(0, \Sigma_1)$, choose and fix $X_2 \in \mathbf{SO}(3)$, then the concentrated Gaussian $q(X) = \mathbf{N}_{X_2}(\mu_2, \Sigma_2)$ for parameters

$$\begin{aligned} \mu_2 &= \log_{\mathbf{G}}^{\vee}(X_2^{-1}X_1) \\ \Sigma_2 &= J_{\mu_2}^{-1} \Sigma_1 J_{\mu_2}^{-\top} \end{aligned}$$

minimises the Kullback-Leibler divergence with $p(X)$ up to second-order linearisation error.

Corollary 2.2 follows directly from Lemma 2.1.

3. PROBLEM FORMULATION

The problem we target is to design a fully distributed algorithm to estimate the absolute attitude of individual agents collaboratively in a multi-agent system. Each agent is equipped with an onboard inertial measurement unit (IMU) that provides bias-free angular velocity $\omega \in \mathbb{R}^3$ in the body frame. With non-rotating, flat earth assumption, the deterministic system kinematics are given by

$$R_i(t+1) = R_i(t) \exp_{\mathbf{G}}(\Delta t \omega_i(t)^\wedge). \quad (4)$$

Note that the subscripted index i refers to i^{th} agent and the index t refers to the time step. We will drop the time step notation where it is clear in order to simplify notation.

The agent i also has extrinsic sensors such as a magnetometer that measures known directions (magnetic field) in the body-frame. The ℓ th direction measurement ${}^i z_\ell$ for agent i is given by

$${}^i z_\ell = R_i^\top d_\ell \quad (5)$$

where d_ℓ for $\ell = 1, \dots, n$ are a collection of known reference directions. Using these measurements, the agent can run a filter-based algorithm locally to estimate its own state and the associated covariance (Markley and Crassidis, 2014)(Fornasier et al., 2022).

In addition, each agent can communicate with agents in its neighbourhood and has a sensor capable of measuring relative attitude of its neighbouring agents. If two agents i and j have states $R_i, R_j \in \mathbf{SO}(3)$ then the relative state ${}^j R_i$ of j with respect to i is defined to be

$${}^j R_i := R_j^{-1} R_i \in \mathbf{SO}(3).$$

The relative state can be thought of as the attitude of agent i expressed as perceived by a sensor on agent j . Alternatively, consider left translation of the whole space by $L_{R_j^{-1}}$. Then $\text{id} = L_{R_j^{-1}} R_j = R_j^{-1} R_j = I_3$ and ${}^j R_i = L_{R_j^{-1}} R_i = R_j^{-1} R_i$. That is, the relative state is the coordinates of agent i with respect to new group parametrisation that place agent j at the identity attitude.

A relative state sensor on agent j may directly measure the physical relative attitude of agent i or may measure the relative angular difference between the states. Physical attitude measurements are associated with directly measuring the direction cosines that make up the entries of the matrix ${}^j R_i \in \mathbf{SO}(3)$. Such measurements are typically inner products like ${}^i z_\ell^\top d_\ell$ that correspond to cosines of angles between known or measured vectors. For a physical relative attitude measurement, an appropriate model is

$${}^j y_i = R_j^{-1} R_i \boxplus \kappa_j, \quad \kappa_j \sim \mathbf{N}(0, Q_j). \quad (6)$$

That is, the generative measurement noise model is a zero-mean concentrated Gaussian process ${}^j y_i \sim \mathbf{N}_{j R_i}(0, Q_j)$. Conversely, a relative angular sensor measures the underlying angle from one attitude to another. For example, if two attitudes are connected through a physical gimbal system then the sensor will measure Euler angles between the two states. Another example is when a vision system or similar system estimates the relative axis of rotation and relative angle from itself to another agent rather than the direction cosines (Chatterji et al., 1998). The measurement in this case, which we denote by z to distinguish it from (6), is appropriately modelled by

$${}^j z_i = \exp_{\mathbf{G}}\left({}^j \mu_i + \kappa_j^\wedge\right), \quad \kappa_j \sim \mathbf{N}(0, Q_j) \quad (7)$$

where ${}^j \mu_i = \log_{\mathbf{G}}(R_j^{-1} R_i) \in \mathfrak{so}(3)$ is the angle-axis ${}^j \mu_i = \theta a^\wedge$ for a rotation of θ rad around an axis $a \in \mathbb{S}^2$. In this case, the generative noise measurement model is a concentrated Gaussian ${}^j z_i \sim \mathbf{N}_{\text{id}}({}^j \mu_i, Q_j)$ with non-zero mean.

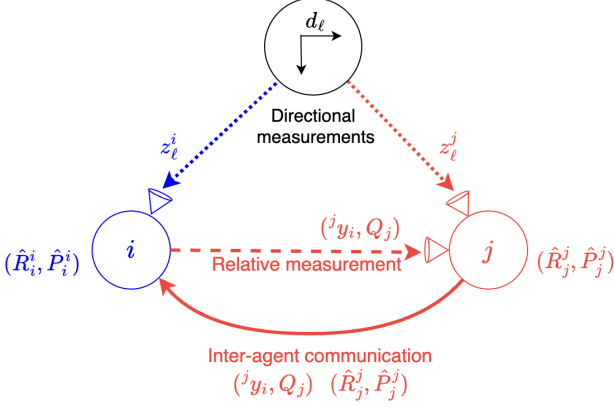


Fig. 1. Illustration of the experimental setup to improve the estimate of agent i . The dotted lines represent each agent’s measurements of known directions which is used to locally estimate their own states. The dashed line refers to the inter-agent measurement ${}^j y_i$ taken by agent j . The solid line represents the communication from agent j to agent i .

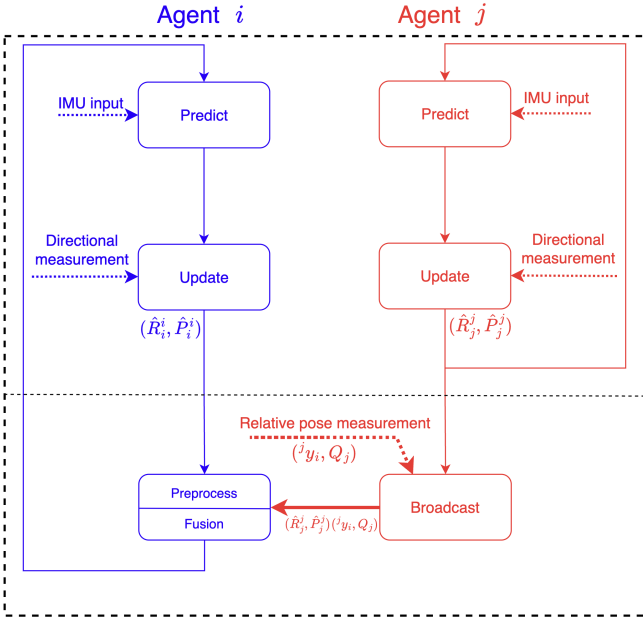


Fig. 2. Implementation structure of the estimators and the information flow on each agent.

4. ALGORITHM

We propose a filter-based algorithm to solve the attitude estimation problem. Each agent estimates its own state $\hat{R} \in \mathbf{SO}(3)$ and the associated covariance $\hat{P} \in \mathbb{S}_+(3)$. We use subscript and superscript for the agent that is being estimated and the agent that is making the estimation, respectively. Fig 2 demonstrates the overview of different steps in the proposed filter algorithm for an ego-agent (shown in blue) and an altruistic agent (shown in red). In this section, we focus on the details of the filter running on the ego-agent i , which can be separated into three stages: predict (using IMU input), update (using directional measurement) and fusion (using relative measurements).

4.1 Predict and update

The filter follows the conventional EKF design methodology, including the predict and update step. The information state of the filter, which approximates the true distribution of the system state on $\mathbf{SO}(3)$, is a concentrated Gaussian distribution, given by

$$R_i(t) \sim \mathbf{N}_{\hat{R}_i^i(t)}(0, \hat{P}_i^i(t)).$$

When an IMU input is available, it will be used to propagate the state estimate \hat{R}_i^i , which also acts as the reference point in the information state of the filter, using the full nonlinear model (4). The covariance estimate \hat{P}_i^i will be propagated by linearising the system model. When the agent receives the directional measurement, under the assumption that every extrinsic measurement is independent, we can perform the standard Kalman fusion in logarithmic coordinates, followed by a covariance reset step which transforms the posterior into a zero-mean concentrated Gaussian (Mueller et al. (2017); Ge et al. (2023)). For detailed implementation of the filter, see Fornasier et al. (2022); Ge et al. (2023).

4.2 Fusion using relative measurements:

The core contribution of this work lies in the preprocessing and fusion steps shown in Figure 2. When agent j makes a relative attitude measurement ${}^j y_i$ or ${}^j z_i$ of agent i , it broadcasts the measurement and the associated noise covariance Q_j of the measurement model, as well as agent j ’s own state estimate $(\hat{R}_j^j, \hat{P}_j^j)$ to agent i . As demonstrated in Fig 2, it takes two steps to fuse the inter-agent information with agent i ’s own estimate. Firstly in the preprocessing stage, the relative measurement is combined with both agent i and agent j ’s state estimate, which generates a new estimate of agent i , denoted by $(\hat{R}_i^j, \hat{P}_i^j)$. We use the superscript j to distinguish it from agent i ’s estimate. Then we implement the CCE fusion method to fuse $(\hat{R}_i^j, \hat{P}_i^j)$ with $(\hat{R}_i^i, \hat{P}_i^i)$.

Geometric transformation of the measurement model: We consider both measurement models in (6) and (7). The second measurement model requires an extra step to transform it into a zero-mean concentrated Gaussian. As given in (7), the relative measurement can be expressed as a concentrated Gaussian with non-zero mean,

$${}^j z_i \sim \mathbf{N}_{\text{id}}(\log_{\mathbf{G}}(R_j^{-1} R_i), Q_j).$$

By applying Lemma 2.1, one gets the following approximation:

$${}^j y_i \sim \mathbf{N}_{R_j^{-1} R_i}(0, J_{\log_{\mathbf{G}}}(R_j^{-1} R_i) Q_j J_{\log_{\mathbf{G}}}(R_j^{-1} R_i)^\top).$$

Such modification requires the knowledge of the noiseless configuration output $R_j^{-1} R_i$ which is not available in practice. In this paper, we will assume the measurement noise is small and use the measurement ${}^j z_i$ as a proxy for the relative state. Alternative algorithms that exploit the two state estimates \hat{R}_j^j and \hat{R}_i^i directly are also possible. The measurement model can now be transformed as follow:

$${}^j y_i \sim \mathbf{N}_{R_j^{-1} R_i}(0, J_{\log_{\mathbf{G}}}({}^j z_i) Q_j J_{\log_{\mathbf{G}}}({}^j z_i)^\top),$$

or equivalently,

$${}^j y_i \approx R_j^{-1} R_i \boxplus \kappa_j, \quad \kappa_j \sim \mathbf{N}(0, Q_j^*), \quad (8)$$

with $Q_j^* = J_{\log_{\mathbf{G}}^{\vee}({}^j y_i)} Q_j J_{\log_{\mathbf{G}}^{\vee}({}^j y_i)}^{\top}$.

Preprocessing relative measurements: Given the measurement model

$${}^j y_i = R_j^{-1} R_i \boxplus \kappa_j, \quad \kappa_j \sim \mathbf{N}(0, Q_j) \quad (9)$$

and the error state models

$$R_i = \hat{R}_i \boxplus \varepsilon_i, \quad \varepsilon_i \sim \mathbf{N}(0, \hat{P}_i) \quad (10)$$

$$R_j = \hat{R}_j \boxplus \varepsilon_j, \quad \varepsilon_j \sim \mathbf{N}(0, \hat{P}_j) \quad (11)$$

substituting (11) into (9) yields

$$\begin{aligned} {}^j y_i &= (\hat{R}_j \boxplus \varepsilon_j)^{-1} R_i \boxplus \kappa_j \\ &= \exp_{\mathbf{G}}(-\varepsilon_j) \hat{R}_j^{-1} R_i \boxplus \kappa_j \\ &= \hat{R}_j^{-1} R_i \boxplus \text{Ad}_{(\hat{R}_j^{-1} R_i)^{-1}}(-\varepsilon_j) \boxplus \kappa_j. \end{aligned}$$

Replace R_i using (10),

$$\begin{aligned} {}^j y_i &= \hat{R}_j^{-1} R_i \boxplus \text{Ad}_{(\hat{R}_j^{-1} (\hat{R}_i \boxplus \varepsilon_i))^{-1}}(-\varepsilon_j) \boxplus \kappa_j \\ &= \hat{R}_j^{-1} R_i \boxplus \text{Ad}_{\exp_{\mathbf{G}}(-\varepsilon_i)} \text{Ad}_{(\hat{R}_j^{-1} \hat{R}_i)^{-1}}(-\varepsilon_j) \boxplus \kappa_j. \end{aligned}$$

Take the Taylor expansion of $\text{Ad}_{\exp_{\mathbf{G}}(-\varepsilon_i)}^{\vee}$, one has

$$\text{Ad}_{\exp_{\mathbf{G}}(-\varepsilon_i)}^{\vee} = I_3 - \text{ad}_{\varepsilon_i}^{\vee} + \mathcal{O}(|\varepsilon_i|^2).$$

Assume that both ε_i and ε_j are small, then $\text{ad}_{\varepsilon_i}^{\vee}(\varepsilon_j)$ and the higher-order terms can be approximated to be zero:

$$\begin{aligned} {}^j y_i &= \hat{R}_j^{-1} R_i \boxplus \text{Ad}_{(\hat{R}_j^{-1} \hat{R}_i)^{-1}}(-\varepsilon_j) \boxplus \kappa_j \\ &\approx \hat{R}_j^{-1} R_i \boxplus \kappa_j^+, \end{aligned} \quad (12)$$

where the new measurement noise κ_j^+ is given by

$$\kappa_j^+ \sim \mathbf{N}(0, \text{Ad}_{(\hat{R}_j^{-1} \hat{R}_i)^{-1}}^{\vee} \hat{P}_j^+ \text{Ad}_{(\hat{R}_j^{-1} \hat{R}_i)^{-1}}^{\top} + Q_j).$$

One can now reconstruct a new estimate of R_i from (12) by left multiplying by \hat{R}_j . The new estimate is a zero mean Gaussian $\mathbf{N}_{\hat{R}_i^j}(0, \hat{P}_i^j)$ where the parameters are given by

$$\begin{aligned} \hat{R}_i^j &= \hat{R}_j {}^j y_i \\ \hat{P}_i^j &= \text{Ad}_{(\hat{R}_j^{-1} \hat{R}_i)^{-1}}^{\vee} \hat{P}_j^+ \text{Ad}_{(\hat{R}_j^{-1} \hat{R}_i)^{-1}}^{\top} + Q_j. \end{aligned}$$

Geometric correction of distributions: The concentrated Gaussian distributions that are being fused can be written as $\mathbf{N}_{\hat{R}_i^i}(0, \hat{P}_i^i)$ and $\mathbf{N}_{\hat{R}_i^j}(0, \hat{P}_i^j)$. Although the covariance \hat{P}_i^i and \hat{P}_i^j are both symmetric matrices, they are still associated with distributions expressed in different coordinates. Fusing these covariance matrices directly, without correcting for the associated change of coordinates, will introduce artifacts and errors in the information state of the resulting filter, decreasing consistency and compromising performance of the algorithm.

In order for the ego-agent to compensate for the change of coordinates in the measurement received from the altruistic agent, it must transform the measurement concentrated Gaussian into a concentrated Gaussian in its local coordinates. That is, it must solve for μ_i^{j*} and \hat{P}_i^{j*} such that

$$\mathbf{N}_{\hat{R}_i^j}(0, \hat{P}_i^j) \approx \mathbf{N}_{\hat{R}_i^i}(\mu_i^{j*}, \hat{P}_i^{j*}).$$

Applying Corollary 2.2 then one has

$$\begin{aligned} \mu_i^{j*} &:= \log_{\mathbf{G}}^{\vee}(\hat{R}_i^{i-1} \hat{R}_j {}^j y_i) \\ \hat{P}_i^{j*} &= J_{\mu_i^{j*}}^{-1} \hat{P}_j^+ J_{\mu_i^{j*}}^{\top}. \end{aligned}$$

Data Fusion Now the targeting distributions are transported into the same coordinate, and the next step is to perform data fusion to the two distributions. Rewrite the distributions into ellipsoidal sets on $\mathfrak{so}(3)$, defined as $\mathcal{E}(0, \hat{P}_i^i) = \{u^\wedge \in \mathfrak{so}(3) : \|u\|_{\hat{P}_i^i}^2 \leq 1\}$ and $\mathcal{E}(\mu_i^{j*}, \hat{P}_i^{j*}) = \{u^\wedge \in \mathfrak{so}(3) : \|u - \mu_i^{j*}\|_{\hat{P}_i^{j*}}^2 \leq 1\}$. Given these two prior ellipsoids have nonempty intersection, the convex combination $\mathcal{E}(\hat{u}^+, \hat{P}_i^+)$ of them is given by

$$\begin{aligned} \hat{P}_i^+ &= kX, \quad X = \left(\alpha \hat{P}_i^i + (1 - \alpha) \hat{P}_i^{j*} \right)^{-1}, \\ k &= 1 - d^2, \quad d^2 = \left\| \mu_i^{j*} \right\|_{\left(\frac{\hat{P}_i^i}{\alpha} + \frac{\hat{P}_i^{j*}}{1-\alpha} \right)^{-1}}^2, \\ \hat{u}^+ &= X \left((1 - \alpha) \hat{P}_i^{j*} \mu_i^{j*} \right), \end{aligned}$$

where $\alpha \in [0, 1]$ is a freely chosen gain in this paper. Alternatively, one can choose an optimal α^* such that $\alpha^* = \text{argmin}_{\alpha} \det(\hat{P}_i^+)$.

Note that given the outputs of CCE fusion method, the posterior is a concentrated Gaussian with non-zero mean, that is, $R_i \sim \mathbf{N}_{\hat{R}_i}(\hat{u}^+, \hat{P}_i^+)$. However, the next fusion iteration requires the distribution to have a zero mean, hence the goal of the reset step is to identify \hat{P}_i° such that

$$R_i \sim \mathbf{N}_{\hat{R}_i}(\hat{u}^+, \hat{P}_i^+) \approx \mathbf{N}_{\hat{R}_i \exp_{\mathbf{G}} \hat{u}^+}(0, \hat{P}_i^\circ). \quad (13)$$

Similar to the coordinate change in the previous steps, this may be solved by using Lemma 2.1. The reset covariance \hat{P}_i° is found to be

$$\hat{P}_i^\circ = J_{\hat{u}^+} \hat{P}_i^+ J_{\hat{u}^+}^{\top}.$$

Note that the reset step only modifies the covariance estimate and does not change the attitude estimate of agent i .

5. NUMERICAL EXPERIMENT

In this section, we provide a numerical evaluation of the algorithm proposed in Section 4. A Monte-Carlo simulation is undertaken to validate the performance of both the proposed geometric modifications and the fusion algorithm.

5.1 System implementation

In the Monte-Carlo setup we use the following randomisation and run 1000 experiments. We simulate two independent oscillatory trajectories for agent i and j , with the noise-less angular velocity generated by

$$\begin{aligned} \omega_i(\tau) &:= (10|\sin(\tau)|, |\cos(\tau)|, 0.1|\sin(\tau)|) \text{ rad/s} \\ \omega_j(\tau) &:= (|\sin(\tau)|, 0.5|\cos(\tau)|, 5|\sin(\tau)|) \text{ rad/s}, \end{aligned}$$

and subsequently corrupted by additive Gaussian noise with covariance $\text{diag}(0.3^2, 0.2^2, 0.1^2)$. The trajectory is realized using Euler integration (4) and a time step $\Delta t = 0.02s$. The initial states estimates of the agents are offset from each other by 180 degrees with initial errors sampled from $\mathbf{N}_{R(0)}(0, I_3)$.

The extrinsic sensor on agent j measures two known directions

$$d_1 = (0, 1, 0) \quad \text{and} \quad d_2 = (1, 0, 0).$$

with the output function (5), while the sensor on agent i only measures the first direction. The measurements for each direction are corrupted with additive Gaussian noise sampled from $\mathbf{N}(0, \text{diag}(0.2^2, 0.1^2, 0.3^2))$. In this experiment, we design the agent i to use only the directional measurement of d_1 , while agent j has access to measurements of both directions. In consequence, without the relative measurements from agent j , the state of the ego agent i is unobservable — a single directional measurement is insufficient to determine the full attitude of the vehicle. In this way, the experiment emphasises the role of the shared information in the filter and exacerbates errors due to information incest. Both agents receive directional measurements at 20Hz.

Agent j makes relative attitude measurement of agent i at 1Hz, which is corrupted by Gaussian noise $\mathbf{N}(0, Q_j)$. The non-homogeneous noise covariance is given by $\text{diag}(0.5^2, 0.3^2, 0.2^2)$. We run separate simulations with both of the measurement noise models considered in (6) and (7).

For comparison, we implement two filters on agent i using different measurements, aside from the proposed algorithm. One filter only uses the extrinsic directional measurements and disregards relative inter-agent measurements. The second filter uses both extrinsic and relative measurements, however, it only does a naive fusion in the logarithmic coordinates, as in (Wolfe and Mashner, 2011).

5.2 Result

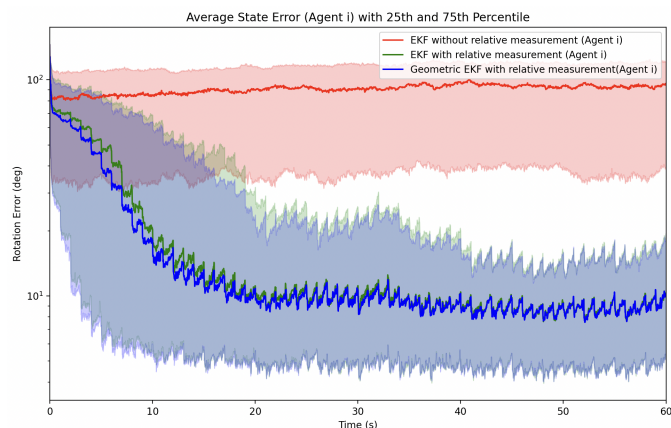


Fig. 3. Direct physical measurement of relative state (6). Mean rotation error ($e := \arccos((\text{tr}(R^{-1}\hat{R}) - 1)/2)$) with a shaded area representing the 25th and 75th percentiles.

Fig 3 and 4 demonstrate the performance of the proposed algorithm (in blue) compared with the EKF using only directional measurement (in red) and the EKF using a naive fusion scheme (in green). Note that the noise parameters were chosen to demonstrate the relative advantages of the geometric modifications to be proposed. That is, while we found that the proposed method outperformed the alternatives regardless of the noise model, the particular parameters used to generate the plots shown here

were chosen to emphasise the performance advantage. In particular, the measurement errors are larger than would be desirable in a real application, leading to relatively large attitude error in the plots.

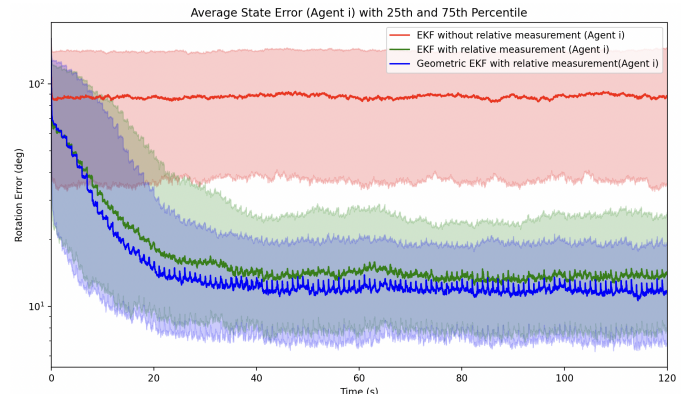


Fig. 4. Relative angle measurement (7). Mean rotation error ($e := \arccos((\text{tr}(R^{-1}\hat{R}) - 1)/2)$) with a shaded area representing the 25th and 75th percentiles.

As expected, the local filter solution with only a single direction measurement available does not converge as the system is unobservable, as shown in both Fig 3 and 4. Fig 3 shows the algorithm performance in the case where the direct physical measurement model (6) is used. It demonstrates an advantage in the proposed fusion algorithm during the transient of the algorithm but has similar asymptotic performance to the naive filter. This is to be expected, since the geometric modification in the data fusion step corrects for the difference between the filter estimate and relative state measurement coordinates. As the filter converges this difference becomes negligible and the benefit of the geometric correction is lost.

This is not the case for measurement model (7), as shown in Fig 4, since the relative state measurement ${}^j z_i$ is taken in coordinates around agent j . The relative state between the two agents i and j does not converge to identity in general, and the geometric correction to compensate for the coordinate representation remains critical.

6. CONCLUSION

In this paper, we propose a collaborative attitude estimation problem where agents run local filter-based algorithms which fuse the estimates and relative measurements communicated by neighboring agents. We utilize the concept of concentrated Gaussians on $\mathbf{SO}(3)$ and exploit the geometric properties of the underlying space. The proposed algorithm combines an EKF running locally on the agent with the CCE fusion of relative state measurements. We provide geometric corrections in the algorithm to incorporate Lie group geometric insights that improve filter performance. The simulation results demonstrate the convergence of the proposed fusion method, and shows the improvements gained from applying the correct geometric modifications with different measurement noise models.

REFERENCES

- Barfoot, T.D. (2017). *State estimation for robotics*. Cambridge University Press.

- Bertsekas, D. and Rhodes, I. (1971). Recursive state estimation for a set-membership description of uncertainty. *IEEE Transactions on Automatic Control*, 16(2), 117–128.
- Bourmaud, G., Mégret, R., Giremus, A., and Berthoumieu, Y. (2016). From intrinsic optimization to iterated extended kalman filtering on lie groups. *Journal of Mathematical Imaging and Vision*, 55, 284–303.
- Carrillo-Arce, L.C., Nerurkar, E.D., Gordillo, J.L., and Roumeliotis, S.I. (2013). Decentralized multi-robot cooperative localization using covariance intersection. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1412–1417. IEEE.
- Chatterji, G.B., Menon, P.K., and Sridhar, B. (1998). Vision-Based Position and Attitude Determination for Aircraft Night Landing. *Journal of Guidance, Control, and Dynamics*, 21(1), 84–92. doi:10.2514/2.4201.
- Chirikjian, G. and Kobilarov, M. (2014). Gaussian approximation of non-linear measurement models on lie groups. In *53rd IEEE Conference on Decision and Control*, 6401–6406. IEEE.
- Chirikjian, G.S. (2011). *Stochastic Models, Information Theory, and Lie Groups, Volume 2: Analytic Methods and Modern Applications*. Springer.
- Fornasier, A., Ng, Y., Brommer, C., Böhm, C., Mahony, R., and Weiss, S. (2022). Overcoming bias: Equivariant filter design for biased attitude estimation with online calibration. *IEEE Robotics and Automation Letters*, 7(4), 12118–12125.
- Ge, Y., Van Goor, P., and Mahony, R. (2023). A note on the extended kalman filter on a manifold. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, 7687–7694. IEEE.
- Julier, S. and Uhlmann, J.K. (2017). General decentralized data fusion with covariance intersection. In *Handbook of multisensor data fusion*, 339–364. CRC Press.
- Julier, S.J. and Uhlmann, J.K. (1997). A non-divergent estimation algorithm in the presence of unknown correlations. In *Proceedings of the 1997 American Control Conference (Cat. No. 97CH36041)*, volume 4, 2369–2373. IEEE.
- Julier, S.J. and Uhlmann, J.K. (2007). Using covariance intersection for slam. *Robotics and Autonomous Systems*, 55(1), 3–20.
- Kahan, W. (1968). Circumscribing an ellipsoid about the intersection of two ellipsoids. *Canadian Mathematical Bulletin*, 11(3), 437–441.
- Luft, L., Schubert, T., Roumeliotis, S.I., and Burgard, W. (2018). Recursive decentralized localization for multi-robot systems with asynchronous pairwise communication. *The International Journal of Robotics Research*, 37(10), 1152–1167.
- Mahony, R., Hamel, T., and Trumpf, J. (2020). Equivariant systems theory and observer design. *arXiv preprint arXiv:2006.08276*.
- Markley, F.L. and Crassidis, J.L. (2014). *Fundamentals of spacecraft attitude determination and control*, volume 1286. Springer.
- Mueller, M.W., Hehn, M., and D’Andrea, R. (2017). Covariance correction step for kalman filtering with an attitude. *Journal of Guidance, Control, and Dynamics*, 40(9), 2301–2306.
- Noack, B., Sijs, J., Reinhardt, M., and Hanebeck, U.D. (2017). Decentralized data fusion with inverse covariance intersection. *Automatica*, 79, 35–41.
- Olfati-Saber, R. (2005). Distributed kalman filter with embedded consensus filters. In *Proceedings of the 44th IEEE Conference on Decision and Control*, 8179–8184. IEEE.
- Olfati-Saber, R. (2006). Distributed kalman filtering and sensor fusion in sensor networks. In *Networked Embedded Sensing and Control: Workshop NESC’05: University of Notre Dame, USA October 2005 Proceedings*, 157–167. Springer.
- Ros, L., Sabater, A., and Thomas, F. (2002). An ellipsoidal calculus based on propagation and fusion. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 32(4), 430–442.
- Roumeliotis, S.I. and Bekey, G.A. (2002). Distributed multirobot localization. *IEEE transactions on robotics and automation*, 18(5), 781–795.
- Schweppe, F. (1968). Recursive state estimation: Unknown but bounded errors and system inputs. *IEEE Transactions on Automatic Control*, 13(1), 22–28.
- Wolfe, K.C. and Mashner, M. (2011). Bayesian fusion on lie groups. *Journal of Algebraic Statistics*, 2(1).
- Zamani, M., Trumpf, J., and Manzie, C. (2023). Collaborative bearing estimation using set membership methods. *arXiv preprint arXiv:2304.07996*.
- Zhu, P. and Ren, W. (2020). Fully distributed joint localization and target tracking with mobile robot networks. *IEEE Transactions on Control Systems Technology*, 29(4), 1519–1532.