Improved robot pose estimation using relative bearing measurements to unknown landmarks

Behzad Zamani * **Jochen Trumpf** ** **Chris Manzie** *

* *Department of Electrical and Electronic Engineering, University of Melbourne, Melbourne, Australia. (e-mail: behzad.zamani@unimelb.edu.au, manziec@unimelb.edu.au)*
** *School of Engineering, The Australian National University, Canberra, Australia. (e-mail: Jochen.Trumpf@anu.edu.au)*

Abstract: In this paper, we propose a modular nonlinear least squares filtering approach for pose estimation of a moving robot that obtains relative bearing measurements to landmarks whose global position is unknown to the robot. Unlike in the well-studied simultaneous localization and mapping (SLAM) problem, we are not trying to localize these unknown landmarks as part of the problem specification; they merely serve to aid the robot pose estimation task. In particular, we are not tracking cross-covariance information between robot pose and landmarks in order to enable a modular filter design, where the landmark-aided module can be switched off and on at will. We integrate the Covariance Intersection (CI) algorithm as part of our solution in order to prevent double counting of information when filter modules share estimates with each other. An alternative derivation of the CI algorithm based on nonlinear least squares estimation makes this integration possible. In a randomized simulation study, we compare the proposed method against a standard INS-GPS error-state Extended Kalman Filter (EKF) for robot pose estimation, Sola (2017), and demonstrate that our proposed landmark-aided solution achieves the same positioning accuracy while drastically improving the 3D orientation estimate of the robot. Due to the modular design, this landmark-aided functionality can be retro-fitted to any existing robot pose filtering algorithm in the form of an additional measurement update step.

1. INTRODUCTION

Robot pose estimation using Inertial Navigation Systems (INS) that incorporate Global Positioning System (GPS) receivers is a very well studied problem, with error-state Extended Kalman Filter (EKF) designs based on quaternion kinematics, Sola (2017), and directional Inertial Measurement Unit (IMU) updates, Mahony et al. (2008), enabling a wide range of drone products through mature open source autopilot solutions such as ArduPilot (`https://ardupilot.org/`). In practice, the development of robotic systems proceeds through multiple design iterations and often arduous certification processes, and the addition of new sensor systems and algorithms to improve task performance can be prohibitively expensive unless a completely modular design is chosen. Modularity helps with managing the complexity of the algorithm (such as the size of state vector representations and error covariance matrices), separation of concerns, debugging, inter-operability and change management. Of particular interest to the robotics community is the incorporation of sensors that measure information about objects in the operating environment, such as cameras that measure relative bearings to recognizable landmarks; Radio Frequency (RF) receivers that measure the relative bearing of emitter beacons; or sonar, radar or lidar sensors that measure the relative position of objects. Simultaneous Localization and Mapping (SLAM), Durrant-Whyte and Bailey (2006); Ebadi et al. (2023), target tracking, Bar-Shalom et al. (2004), collaborative localization, Roumeliotis and Bekey (2002), self-driving cars, Patole et al. (2017), and augmented/virtual reality headsets, Park et al. (2021), are examples of well-studied robotics applications that all share this characteristic.

In this paper, we propose a modular nonlinear least squares filter design for 3D pose (position and orientation) of a mobile robot. In addition to the standard INS-GPS measurements (3-axis linear accelerometer, gyroscope and magnetometer, GPS position), the robot obtains noisy measurements of the relative bearing to multiple landmarks in the environment. This is in contrast to our work in Zamani and Trumpf (2025) that considered relative position measurements. We do not assume that the global position of these landmarks or their global identity is known to the robot, but merely that the robot can track and re-identify them over time, as is for example the case for visual features from an on-board camera or angle-of-arrival measurements from an antenna receiving transmitted ID codes from emitter beacons. We provide a modular design consisting of a robot filter module and a separate landmark filter module. We update the state and error covariance estimates of each module independently, even though the relative bearing measurements simultaneously depend on the states of both modules. In this setup, the information exchanged between the modules is not statistically independent and does not contain (tracked) cross correlations, Uhlmann (2003); Chen et al. (2002). Bayesian fusion algorithms such as the Kalman filter, Kalman (1960), or similar methods are not applicable because they require that statistically identically and independently distributed signals are exchanged or that cross-covariance information is available or tracked. We integrate the Covariance Intersection (CI) algorithm, Julier and Uhlmann (1997a), as part of our solution to prevent double counting of information, since in our situation the information flow graph describing how modules share estimates contains a loop. We use an alternative derivation of the CI algorithm based on nonlinear least squares filtering to make this integration possible. In a randomized simulation study, we compare the proposed method against a standard filter for robot pose estimation, Sola (2017), and demonstrate that our proposed landmark-aided solution achieves the same positioning accuracy while drastically improving the 3D orientation estimate of the robot.

Similar ideas to our modular estimation approach can be found in Julier and Uhlmann (2007); Arambel et al. (2001); Carrillo-Arce et al. (2013); Chang et al. (2021) and the related papers they cite. In Julier and Uhlmann (2007), a CI based distributed robot landmark update method is proposed towards an EKF based formulation of the SLAM problem. In contrast to the SLAM problem, we do not aim at estimating the global position of the unknown landmarks (we are not trying to also build a map), and it turns out that not doing so still leads to better robot filter performance. Cooperative self-localization algorithms based on CI are explored in Arambel et al. (2001); Chang et al. (2021) where every subsystem estimates the joint state of the overall network. The approach in Carrillo-Arce et al. (2013) is closest to ours, albeit specialized to an EKF-CI based multi-robot cooperative localization algorithm. In contrast to these approaches, we concentrate on maximizing the performance of a single-robot pose estimation algorithm.

1.1 Notation

Let $\mathcal{A}$ denote the global reference frame and let $\mathcal{B}_r$ denote a body fixed frame of a robot moving in 3D space. Translation of the origin of $\mathcal{B}_r$ with respect to $\mathcal{A}$, expressed in $\mathcal{A}$, is denoted as $p_r \in \mathbb{R}^3$. The 3D attitude (orientation) of $\mathcal{B}_r$ relative to the reference frame $\mathcal{A}$, expressed in $\mathcal{A}$, is represented by a rotation matrix $R_r \in \mathrm{SO}(3)$. The rotation group is denoted by $\mathrm{SO}(3) = \{R \in \mathbb{R}^{3\times3} \,|\, R^\top R = I, \det(R) = 1\}$, where $(.)^\top$ is the transpose map, I is the 3 by 3 identity matrix and $\det(.)$ is the matrix determinant. The associated Lie algebra $\mathfrak{so}(3)$ is the set of skew-symmetric matrices, $\mathfrak{so}(3) = \{\omega_\times \in \mathbb{R}^{3\times3} \,|\, \omega_\times = -\omega_\times^\top\}$. For $\omega = [\omega_1, \omega_2, \omega_3]^\top \in \mathbb{R}^3$, the lower index operator $(.)_\times : \mathbb{R}^3 \longrightarrow \mathfrak{so}(3)$ yields the skew-symmetric matrix

$$\omega_\times := \begin{bmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{bmatrix}.$$

The skew-symmetric matrix is associated with the cross product by ω, *i.e.*, $\omega_\times v = \omega \times v$, for all $v \in \mathbb{R}^3$. Inversely, the operator $(.)^\vee : \mathfrak{so}(3) \longrightarrow \mathbb{R}^3$ extracts the skew coordinates, $\omega_\times^\vee := \omega$. The unit circle group is denoted by S^1. Let $\theta \in S^1 \subset \mathbb{R}$ and $a \in \mathbb{R}^3, \|a\| = 1$ denote the angle-axis representation of the attitude $R \in SO(3)$. It holds that $R = \exp(\theta a_\times)$ and $\log(R) = \theta a_\times$. We denote $e_3 := [0, 0, 1]^\top$ and $\|a\|_M := \sqrt{a^\top M a}$ where $M > 0 \in \mathbb{R}^{3\times3}$ is a symmetric positive definite matrix.

2. PROBLEM FORMULATION

In this section we introduce models for the robot, robot measurements and relative bearing measurements with respect to a group of stationary landmarks. Subsequently, we formalize the filtering problem using these models.

2.1 Robot State Kinematics

Let the robot state be denoted by $x_r := [p_r^\top,\ v_r^\top,\ (\log^\vee(R_r))^\top,\ b_a^\top,\ b_w^\top]^\top$. Here $p_r \in \mathbb{R}^3$ is the translation, $v_r \in \mathbb{R}^3$ is the linear velocity, and $R_r \in \mathrm{SO}(3)$ is the attitude of the robot body fixed frame $\mathcal{B}_r$ with respect to the global reference frame $\mathcal{A}$. The remaining states $b_a \in \mathbb{R}^3$ and $b_w \in \mathbb{R}^3$ denote the unknown bias states of the linear acceleration and angular velocity measurements introduced later. The kinematics model is given by

$$\dot{p}_r = v_r,\ \dot{v}_r = R_r^\top a,\ \dot{R}_r = R_r w_\times,\ \dot{b}_a = \epsilon_a,\ \dot{b}_w = \epsilon_w. \quad (1)$$

Here, $w \in \mathbb{R}^3$ is the angular velocity and $a \in \mathbb{R}^3$ is the linear acceleration applied to the robot in the body fixed frame $\mathcal{B}_r$. The signals $\epsilon_a,\ \epsilon_w \in \mathbb{R}^3$ are zero mean disturbance signals (with known error covariance matrices $Q_a > 0 \in \mathbb{R}^{3\times3}$ and $Q_w > 0 \in \mathbb{R}^{3\times3}$, respectively) causing the unknown bias states to slowly vary with time.

2.2 INS-GPS Measurements

The considered Inertial Navigation System (INS-GPS) comprises a GPS, an accelerometer, a gyroscope and a magnetometer. The INS-GPS measurement model is

$$\begin{aligned} u_p &= p_r + \delta_p,\ u_a = a - R_r^\top g + b_a + \delta_a, \\ u_w &= w + b_w + \delta_w,\ u_m = R_r^\top m + \delta_m. \end{aligned} \quad (2)$$

See Section 2.3 for an explanation of u_p, u_w, u_a and u_m. The gravitational acceleration $g \approx -9.8 e_3$ and $m \in \mathbb{R}^3$ denotes the known directional vector of the magnetic field in the local area of operations, both represented in $\mathcal{A}$. The additive noise signals are denoted by $\delta_p, \delta_a, \delta_w, \delta_m \in \mathbb{R}^3$ which are unknown (zero mean) noises with known positive definite covariance matrices $Q_p, Q_a, Q_w, Q_m \in \mathbb{R}^{3\times3}$.

2.3 INS-GPS Processing

We will utilise a nonlinear error-state Kalman filter described in Sola (2017) to estimate the robot states, denoted by $\hat{x}_r \in \mathbb{R}^{15}$, together with an associated error covariance $P_r \in \mathbb{R}^{15\times15}$.

The inertial measurements u_a and u_w are available at high frequencies and will be used in the prediction part of the filter. The lower frequency GPS measurement u_p will be used in the update part of the filter for error correction on the predicted states. The accelerometer and magnetometer measurements u_a and u_m will also be used in the update part of the filter in order to provide error correction on the predicted attitude state (as well as the other states). This follows the approach taken in Mahony et al. (2008) and Sola (2017).

2.4 Relative Landmark Measurements

Let $\{z_i \in \mathbb{R}^3\}_{i=1}^m$ denote a set of m bearing measurements in $\mathcal{B}_r$ relative to associated landmarks located at $\{p_{l_i} \in \mathbb{R}^3\}_{i=1}^m$ (in the global reference frame $\mathcal{A}$). We assume that the values of p_{l_i}, $i = 1, \dots, m$ are unknown to the robot. In practice, the robot measures the azimuth and elevation angles relative to the landmark positions. Denote $\theta_i, \psi_i \in \mathcal{S}^1$ the true azimuth and elevation bearing angles with respect to landmark i; $\theta_i^m, \psi_i^m \in \mathcal{S}^1$ their measured versions; and, $\delta_{\theta,i}, \delta_{\psi,i} \in \mathcal{S}^1$ their additive error signals. The errors are drawn from normal distributions $\delta_{\theta,i} \sim \mathcal{N}(0, \sigma_\theta^2)$ and $\delta_{\psi,i} \sim \mathcal{N}(0, \sigma_\psi^2)$ with known (in practice, computed from camera and lens characteristics) standard deviations $\sigma_\theta > 0$ and $\sigma_\psi > 0$, respectively, i.e.

$$\theta_i^m = \theta_i + \delta_{\theta,i}, \; \psi_i^m = \psi_i + \delta_{\psi,i}. \tag{3}$$

The spherical coordinate models for the true and measured bearing directions are respectively

$$\begin{aligned} \phi_i &:= \begin{bmatrix} \cos(\theta_i)\cos(\psi_i) \\ \sin(\theta_i)\cos(\psi_i) \\ \sin(\psi_i) \end{bmatrix} = \frac{R_r^\top (p_{l_i} - p_r)}{\|p_{l_i} - p_r\|} \quad \text{and} \\ z_i &:= \begin{bmatrix} \cos(\theta_i^m)\cos(\psi_i^m) \\ \sin(\theta_i^m)\cos(\psi_i^m) \\ \sin(\psi_i^m) \end{bmatrix}. \end{aligned} \tag{4}$$

We use this relative measurement model to update the state and error covariance estimates of the robot and the stationary landmarks. Joint landmark estimates and their joint error covariance are denoted $\hat{p}_l =: [\cdots \hat{p}_{l_i}^\top \cdots]^\top \in \mathbb{R}^{3m}$, $P_l \in \mathbb{R}^{3m \times 3m}$.

2.5 Bearing Cost

Relative bearing measurements can provide information for error correction updates to the state estimates of both the robot and the landmark filter modules. The following cost captures the uncertainty of the relative bearing measurement z_i with respect to both robot state x_r and landmark state p_{l_i}.

$$\begin{aligned} J(x_r, p_{l_i}) &= \frac{1}{2}\|D_{z_i} R_r^\top (p_{l_i} - p_r)\|^2, \\ D_{z_i} &:= \frac{z_{\theta_i}^\perp (z_{\theta_i}^\perp)^\top}{\sigma_\theta} + \frac{z_{\psi_i}^\perp (z_{\psi_i}^\perp)^\top}{\sigma_\psi}, \\ z_{\theta_i}^\perp &:= \begin{bmatrix} -\sin(\theta_i^m) \\ \cos(\theta_i^m) \\ 0 \end{bmatrix}, \; z_{\psi_i}^\perp := \begin{bmatrix} -\cos(\theta_i^m)\sin(\psi_i^m) \\ -\sin(\theta_i^m)\sin(\psi_i^m) \\ \cos(\psi_i^m) \end{bmatrix}. \end{aligned} \tag{5}$$

This cost was previously proposed in the target tracking literature, Miles et al. (2024); Lingren and Gong (1978), and captures the measurement error in a subspace perpendicular to the measured bearing z_i. Note that when there is no noise, the measured bearing direction z_i coincides with the true bearing direction ϕ_i. Moreover, in the same ideal scenario, D_{z_i} is a projection onto a subspace (plane) perpendicular to the true bearing direction ϕ_i. Hence, $D_{z_i} R_r^\top (p_{l_i} - p_r) = D_{z_i} \phi_i \|p_{l_i} - p_r\| = 0$ in that scenario. Thus, minimising this cost picks a robot-landmark estimate with the least expected bearing error with respect to the measured bearing directions. Note that $p_{l_i} = p_r$ also minimises this cost which is an undesired effect that is avoided when there is *persistency of excitation* in the measurements, Deghat et al. (2014). Nevertheless, this is an attractive formulation for the target tracking problem due to the fact that robot position and landmark position states feature linearly in the error $D_{z_i} R_r^\top (p_{l_i} - p_r)$.

3. MODULAR ROBOT AND LANDMARK ESTIMATION

In this paper, our focus is on a modular estimation algorithm for updating the estimates of the landmarks $(\hat{p}_l, P_l)$ and the robot states $(\hat{x}_r, P_r)$ based on their relative bearing measurements z_i. The modular approach updates each module independently, in contrast to a joint filter that couples all the estimates, i.e. with a joint stacked estimate of dimension $\mathbb{R}^{(3m+15)}$ and joint error covariance matrix of dimension $\mathbb{R}^{(3m+15)\times(3m+15)}$. The modular approach is desirable as it does not need to directly update the landmark estimates when processing the INS-GPS measurements (which are most relevant to the robot states). We reiterate that we are not interested in the performance of the estimate $(\hat{p}_l, P_l)$ itself. Figure 1 depicts the information flow between the two modules.

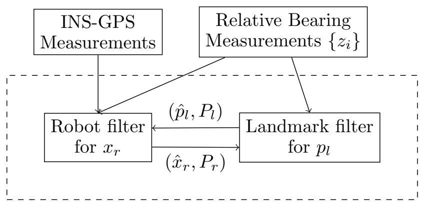

Fig. 1. Modular Filter Architecture

3.1 Modular Landmark Estimation

In this section, we show how the relative bearing measurements can be used to update the landmark estimation module.

First, we combine the uncertainty of the prior robot estimate with that of the relative bearing measurements, using the following cost

$$J_l(x_r, p_l) = \frac{1}{2}\|e(x_r, p_l, \{z_i\})\|^2 + \frac{1}{2}\|x_r - \hat{x}_r\|_{P_r^{-1}}^2, \tag{6}$$

where the error function $e(x_r, p_l, \{z_i\})$ and the combined landmark state $p_l := [\cdots p_{l_i}^\top \cdots]^\top \in \mathbb{R}^{3m}$ are defined as

$$e(x_r, p_l, \{z_i\}) := [\cdots (p_{l_i} - p_r)^\top R_r D_{z_i} \cdots]^\top.$$

Recall that $D_{z_i} \in \mathbb{R}^{3\times 3}$ was defined in (5). Taylor expansion of the error function and its first derivative with respect to x_r yields

$$\begin{aligned} e(x_r, p_l, \{z_i\}) &= e(\hat{x}_r, p_l, \{z_i\}) \\ &\quad + \frac{\partial}{\partial x_r} e(\hat{x}_r, p_l, \{z_i\})(x_r - \hat{x}_r) + r_r \end{aligned} \tag{7}$$

$$\frac{\partial}{\partial x_r} e(x_r, p_l, \{z_i\}) = \frac{\partial}{\partial x_r} e(\hat{x}_r, \hat{p}_l, \{z_i\}) + r_l \tag{8}$$

where $r_r, r_l \in \mathbb{R}^3$ are the remainder terms. Note that both of these terms are second order with respect to the error function. We now introduce $M := I_{3m\times 3m} + H_r P_r H_r^\top$, where

$$H_r := \begin{bmatrix} \vdots & \vdots & \vdots & \vdots \\ -D_{z_i}\hat{R}_r^\top & 0_{3\times 3} & D_{z_i}\left(\hat{R}_r^\top(\hat{p}_l - \hat{p}_r)\right)_\times & 0_{3\times 6} \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} \in \mathbb{R}^{3m\times 15}.$$

We can use the following lemma to obtain a bearing measurement cost only dependent on landmark states.

Lemma 1. Assume that the remainders r_r in (7) and r_l in (8) are negligible. Then the optimal cost

$$J_l(p_l) \coloneqq \min_{x_r} J_l(x_r, p_l) = \frac{1}{2}\|e(\hat{x}_r, p_l, \{z_i\})\|^2_{M^{-1}} + r_{J_l}, \quad (9)$$

where r_{J_l} is negligible.

Proof. Note that the bearing measurement error is due to the sensor used and therefore is independent of the prior estimation error of the robot. The first order condition for minimising (6) over x_r yields

$$\left(\frac{\partial}{\partial x_r} e(x_r^*, p_l, \{z_i\})\right)^\top e(x_r^*, p_l, \{z_i\}) + P_r^{-1}(x_r^* - \hat{x}_r) = 0,$$

where $x_r^* \coloneqq \operatorname{argmin}_{x_r} J_l(x_r, p_l)$. The error function can be replaced from (7), and we have (using left parameterisation of TSO(3))

$$\frac{\partial}{\partial x_r} e(x_r^*, p_l, \{z_i\}) \approx \frac{\partial}{\partial x_r} e(\hat{x}_r, \hat{p}_l, \{z_i\}) = H_r \quad (10)$$

from (8) and our assumption that $r_l \approx 0$. By the assumption that $r_r \approx 0$ it follows that

$$H_r^\top \left(e(\hat{x}_r, p_l, \{z_i\}) + H_r(x_r^* - \hat{x}_r)\right) + P_r^{-1}(x_r^* - \hat{x}_r) \approx 0$$

and therefore $x_r^* \approx \hat{x}_r - (H_r^\top H_r + P_r^{-1})^{-1} H_r^\top e(\hat{x}_r, p_l, \{z_i\})$. Replacing x_r^* into $J_l(x_r, p_l)$ and applying the Woodbury matrix identity then yields the stated formula for $J_l(p_l)$.

Next, we combine this cost with a cost on the prior landmark estimates. Note that the uncertainty captured in the measurement cost (9) is not independent from that of the prior landmark estimates. This is because in obtaining (9), we utilised the robot estimates which in turn are computed incorporating the landmark estimates in the previous iteration. This information loop (see Figure 1) can result in double counting of information and inconsistencies. Therefore, we employ Covariance Intersection (CI), Julier and Uhlmann (1997b), to prevent the potential double counting problem. A least squares derivation of CI yields the following modified combined cost, Zamani and Trumpf (2025), for $0 \leq \alpha_l \leq 1$,

$$J_l^+(p_l) \coloneqq \frac{\alpha_l}{2}\|p_l - \hat{p}_l\|^2_{P_l^{-1}} + \frac{(1-\alpha_l)}{2}\|e(\hat{x}_r, p_l, \{z_i\})\|^2_{M^{-1}}.$$

Minimizing this cost with respect to p_l yields the following nonlinear filter equations for updating the landmark estimates.

$$\begin{aligned}
\hat{p}_l^+ &= \hat{p}_l - (1-\alpha_l^*)P_l^+ H_l^\top M^{-1} e(\hat{X}_r, \hat{p}_l, \{z_i\}),\\
P_l^+ &= \left(\alpha_l^* P_l^{-1} + (1-\alpha_l^*)H_l^\top M^{-1} H_l\right)^{-1},\\
H_l &\coloneqq \operatorname{Diag}(D_{z_i}\hat{R}_r^\top) \in \mathbb{R}^{3m\times 3m},\\
\alpha_l^* &= \operatorname*{argmin}_{0\leq\alpha_l\leq 1} \det\left\{\left(\alpha_l P_l^{-1} + (1-\alpha_l)H_l^\top M^{-1} H_l\right)^{-1}\right\}.
\end{aligned} \quad (11)$$

3.2 Modular Robot Estimation

This section parallels the previous one and shows how the relative bearing measurements can be used to update the robot estimation module.

This time, we first combine the uncertainty of the prior landmark estimates with that of the relative bearing measurements.

$$J_r(x_r, p_l) = \frac{1}{2}\|e(x_r, p_l, \{z_i\})\|^2 + \frac{1}{2}\|p_l - \hat{p}_l\|^2_{P_l^{-1}}. \quad (12)$$

Recall $H_l = \operatorname{Diag}(D_{z_i}\hat{R}_r^\top) \in \mathbb{R}^{3m\times 3m}$ and denote

$$W \coloneqq I_{3m\times 3m} + H_l P_l H_l^\top. \quad (13)$$

Similar to the previous section, Taylor expansion of the error function and its first derivative with respect to p_l yields

$$\begin{aligned}
e(x_r, p_l, \{z_i\}) &= e(x_r, \hat{p}_l, \{z_i\})\\
&\quad + \frac{\partial}{\partial p_l} e(x_r, \hat{p}_l, \{z_i\})(p_l - \hat{p}_l) + \bar{r}_l
\end{aligned} \quad (14)$$

$$\frac{\partial}{\partial p_l} e(x_r, p_l, \{z_i\}) = \frac{\partial}{\partial p_l} e(\hat{x}_r, \hat{p}_l, \{z_i\}) + \bar{r}_r, \quad (15)$$

where $\bar{r}_l, \bar{r}_r \in \mathbb{R}^3$ are second order remainder terms with respect to the error function.

Lemma 2. Assume that the remainders $\bar{r}_l$ in (14) and $\bar{r}_r$ in (15) are negligible. Then the optimal cost

$$J_r(x_r) \coloneqq \min_{p_l} J_r(x_r, p_l) = \frac{1}{2}\|e(x_r, \hat{p}_l, \{z_i\})\|^2_{W^{-1}} + r_{J_r}, \quad (16)$$

where r_{J_r} is negligible.

Proof. Follows a similar approach to Lemma 1.

Next, we combine this cost with a cost on the uncertainty of the prior robot state estimate. Again, we note that the uncertainty captured in the measurement cost (16) is not independent from the uncertainty of this prior estimate due to the information loop existing between the two modules.

To prevent the potential double counting problem we define the following safe fusion cost for $0 \leq \alpha_r \leq 1$,

$$J_r^+(x_r) = \frac{\alpha_r}{2}\|x_r - \hat{x}_r\|^2_{P_r^{-1}} + \frac{(1-\alpha_r)}{2}\|e(x_r, \hat{p}_l, \{z_i\})\|^2_{W^{-1}}.$$

Minimizing this cost with respect to x_r yields the following nonlinear filter equations for updating the robot estimate.

$$\begin{aligned}
\hat{x}_r^+ &= \hat{x}_r - (1-\alpha_r^*)P_r^+ H_r^\top W^{-1} e(\hat{x}_r, \hat{p}_l, \{z_i\}),\\
P_r^+ &= (\alpha_r^* P_r^{-1} + (1-\alpha_r^*)H_r^\top W^{-1} H_r)^{-1},\\
\alpha_r^* &= \min_{0\leq\alpha_r\leq 1} \det\left\{\left(\alpha_r P_r^{-1} + (1-\alpha_r)H_r^\top W^{-1} H_r\right)^{-1}\right\}.
\end{aligned} \quad (17)$$

4. SIMULATION STUDY

In this section we provide a randomized simulation study to assess the performance of the modular landmark-aided robot localisation algorithm proposed in Section 3. We will compare the robot localization performance of our *landmark-aided* approach against a baseline algorithm (denoted as *unaided*) that only uses INS-GPS measurements. Those measurements are also used by the proposed landmark-aided approach. From the perspective of the robot filter module, the landmark-aided approach merely adds an additional filter update step. We also consider another version of the proposed landmark-aided algorithm that has reduced computational requirements. The reduced computation variant (naively) neglects the fact that estimates shared between subsystems are correlated and performs a standard least squares [1] instead of the CI least squares fusion step. This amounts to not requiring the computation of the optimal weights α_l^* and α_r^*, and uses $\alpha_l^*, \alpha_r^* \longleftarrow 1$ and $(1-\alpha_l^*), (1-\alpha_r^*) \longleftarrow 1$ instead. We

[1] A standard least squares fusion step corresponds to Kalman Fusion, Julier and Uhlmann (1997a).

found, however, that this reduced-computation version is unstable (numerically blows up) in almost every run of our experiments. This further highlights the importance of the CI least squares fusion step in avoiding an unstable information loop. The baseline 'unaided' algorithm does not encounter this problem as it avoids the information loop all together.

4.1 Scenario

Consider a quadrotor aerial robot with a flight controller interface that accepts angular velocity and linear acceleration input commands. At the start of the scenario, the robot is stationary, located at the origin, and with its body-fixed frame $\mathcal{B}_r$ aligned with the global reference frame $\mathcal{A}$ (pinned at the origin). A cascade feedback controller generates angular velocity and linear acceleration commands that guide the robot through a sequence of 3D coordinates (in $\mathcal{A}$) arranged as columns of

$$P = \begin{bmatrix} 0 & 40 & 50 & 80 & 110 & 80 & 50 \\ 0 & 20 & 80 & 110 & 80 & 50 & 80 \\ 10 & 10 & 20 & 20 & 20 & 20 & 20 \end{bmatrix} (m). \quad (18)$$

The controller also simultaneously aligns the robot heading (x-axis of the body-fixed frame) with the following associated heading angles (while maintaining zero elevation and roll angles).

$$h = [60,\ 20,\ 11.3,\ 101.3,\ 171.9,\ 98.1,\ 11.3]\ (\text{deg}). \quad (19)$$

The controller restricts the control inputs to respect physical constraints in terms of a maximum angular rate $\|w\| \le 5\frac{\text{deg}}{s}$, maximum linear velocity rate $\|v_r\| \le 2\frac{m}{s}$ and maximum linear acceleration $\|a\| \le 80\frac{m}{s^2}$.

The INS-GPS measurements are simulated according to (2) with zero mean Gaussian additive noise signals with standard deviations $\sigma_p = 0.3m$, $\sigma_a = 0.1\frac{m}{s^2}$, $\sigma_w = 0.01\frac{rad}{s}$, and $\sigma_m = 5$deg, respectively. The initial angular velocity and linear acceleration measurement bias terms are $b_w(0) = 0.5\frac{deg}{s}$ and $b_a(0) = 0.05\frac{m}{s^2}$, respectively, while their driving noise terms ϵ_w, ϵ_a are drawn from a zero mean Gaussian distribution with standard deviations $0.01\frac{deg}{s^2}$ and $0.001\frac{m}{s^3}$, respectively. The relative robot-landmark azimuth and elevation bearing angle measurements are simulated according to (3) with zero mean Gaussian additive noise signals with standard deviations $\sigma_\theta = \sigma_\psi = 0.3$ deg. We consider seven landmarks, always detectable by the robot, with their true 3D coordinates (unknown to the robot) arranged as the columns of

$$L = \begin{bmatrix} 30 & 50 & 30 & 50 & 80 & 90 & 70.5 \\ 30 & 120 & 90 & 30 & 120 & 80 & 80.5 \\ 50 & 4 & 15 & 70 & 60 & 10 & 42 \end{bmatrix} (m). \quad (20)$$

The robot does have an initial estimate for these landmarks but with a very large error (in the order of 100 meters in each coordinate), $\hat{L}(0) = L + \epsilon_L,\ \epsilon_L \sim \mathcal{N}(0, 100^2)$. Therefore, we pick $P_L(0) = 100^2 I_{21\times21}$. The robot's state estimate is initialized as

$$\begin{aligned} \hat{p}(0) &= p(0) + \epsilon_p,\ \epsilon_p \sim \mathcal{N}(0, 0.5^2)(m), \\ \hat{v}(0) &= v(0) + \epsilon_v,\ \epsilon_v \sim \mathcal{N}(0, 0.1^2)(m/s), \\ \hat{h}(0) &= h(0) + \epsilon_h,\ \epsilon_h \sim \mathcal{N}(0, 5^2)(deg), \\ \hat{b}_a(0) &= 0_3(m/s^2),\ \hat{b}_w(0) = 0_3(deg/s). \end{aligned} \quad (21)$$

Each Monte Carlo run lasts 140 seconds (coinciding with the robot reaching its last waypoint (18)). The robot's control and movement is simulated at every simulation step $dt = 0.01s$ (100 Hz) based on exact state feedback [2]. INS measurements for the prediction part of the robot state EKF are incorporated at $100Hz$ while the INS-GPS and relative landmark measurements used for the update equations are incorporated at $1Hz$.

4.2 Results

We completed 10000 Monte Carlo simulations of the randomised scenario in Section 4.1 and compared the performance of the proposed landmark-aided method with the baseline unaided approach. As already mentioned, the reduced-computation approach was unstable in almost every run and therefore is omitted from this comparison. We evaluate the performance of these methods based on estimation error signals of the main robot states achieved at the end of each repeat. Figures 2, 3 and 4 depict the distribution of the final robot heading, position and velocity errors over all runs. Table 1 summarizes the mean and standard deviation of these distributions.

As is evident from these results, the proposed landmark-aided approach achieves a significant (almost 34%) error reduction in the mean heading error while maintaining comparable performance in position and velocity error metrics. The latter was expected as the GPS sensor already provides sufficient information for estimating position and velocity of the robot. This study further emphasizes that our proposed landmark-aided approach can be considered as a modular addition to the existing localization algorithms of many practical robotic systems, requiring minimal additional computation or architectural change.

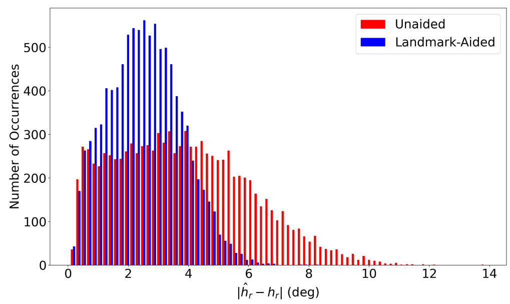

Fig. 2. Histogram of final robot heading errors of Monte Carlo runs.

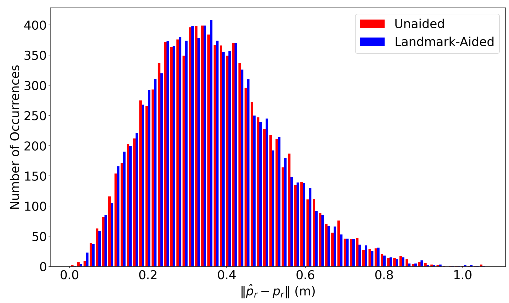

Fig. 3. Histogram of final robot position errors of Monte Carlo runs.

[2] Our focus is on estimation and not closed-loop control based on estimated states. Our setup simulates realistic control values for this purpose.

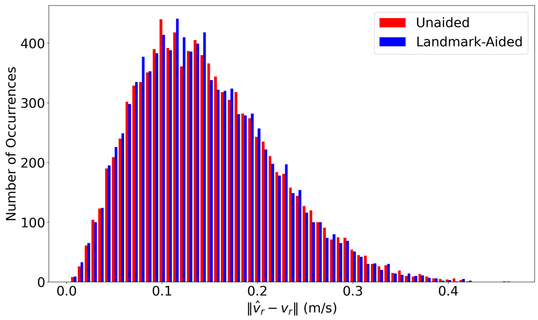

Fig. 4. Histogram of final robot velocity errors of Monte Carlo runs.

Method	Measure	Mean	Std
Unaided	Heading Error (deg)	3.84	2.22
Landmark-Aided	Heading Error (deg)	**2.54**	**1.21**
Unaided	Position Error (m)	0.37	0.16
Landmark-Aided	Position Error (m)	**0.36**	0.16
Unaided	Velocity Error (m/s)	0.15	0.07
Landmark-Aided	Velocity Error (m/s)	**0.14**	0.07

Table 1. Robot estimation error statistics after 10000 random trials.

5. CONCLUSIONS

In this paper, we provided a modular non-linear least squares filter for landmark-aided INS-GPS 3D pose estimation of a mobile robot using relative bearing measurements to landmarks whose position in the global reference frame is unknown. We demonstrate in a large simulation study that the additional landmark filter module significantly improves the 3D orientation estimate of the robot while not degrading the 3D position and linear velocity estimates. The proposed filter module can be used with other existing 3D pose estimation filters in the form of an additional filter update step with minimal requirements in terms of architectural change or additional computation.

REFERENCES

Arambel, P.O., Rago, C., and Mehra, R.K. (2001). Covariance intersection algorithm for distributed spacecraft state estimation. In *Proceedings of the 2001 American Control Conference.(Cat. No. 01CH37148)*, volume 6, 4398–4403. IEEE.

Bar-Shalom, Y., Li, X.R., and Kirubarajan, T. (2004). *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons.

Carrillo-Arce, L.C., Nerurkar, E.D., Gordillo, J.L., and Roumeliotis, S.I. (2013). Decentralized multi-robot cooperative localization using covariance intersection. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, 1412–1417. IEEE.

Chang, T.K., Chen, K., and Mehta, A. (2021). Resilient and consistent multirobot cooperative localization with covariance intersection. *IEEE Transactions on Robotics*, 38(1), 197–208.

Chen, L., Arambel, P.O., and Mehra, R.K. (2002). Estimation under unknown correlation: Covariance intersection revisited. *IEEE Transactions on Automatic Control*, 47(11), 1879–1882.

Deghat, M., Shames, I., Anderson, B.D., and Yu, C. (2014). Localization and circumnavigation of a slowly moving target using bearing measurements. *IEEE Transactions on Automatic Control*, 59(8), 2182–2188.

Durrant-Whyte, H. and Bailey, T. (2006). Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2), 99–110.

Ebadi, K., Bernreiter, L., Biggie, H., Catt, G., Chang, Y., Chatterjee, A., Denniston, C.E., Deschênes, S.P., Harlow, K., Khattak, S., et al. (2023). Present and future of slam in extreme environments: The darpa subt challenge. *IEEE Transactions on Robotics*, 40, 936–959.

Julier, S.J. and Uhlmann, J.K. (1997a). 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. (1997b). A non-divergent estimation algorithm in the presence of unknown correlations. In *American Control Conference, 1997. Proceedings of the 1997*, 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.

Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, Transactions of the ASME*, 82(1), 35–45.

Lingren, A.G. and Gong, K.F. (1978). Position and velocity estimation via bearing observations. *IEEE Transactions on Aerospace and electronic systems*, AES-14(4), 564–577.

Mahony, R., Hamel, T., and Pflimlin, J.M. (2008). Nonlinear complementary filters on the special orthogonal group. *IEEE Transactions on Automatic Control*, 53(5), 1203–1218.

Miles, A., Zamani, B., Shames, I., and Arulampalam, S. (2024). Recursive least squares bearing only tracking. *Submitted to the IEEE Transactions on Aerospace and Electronic Systems*.

Park, S., Bokijonov, S., and Choi, Y. (2021). Review of microsoft hololens applications over the past five years. *Applied sciences*, 11(16), 7259.

Patole, S.M., Torlak, M., Wang, D., and Ali, M. (2017). Automotive radars: A review of signal processing techniques. *IEEE Signal Processing Magazine*, 34(2), 22–35.

Roumeliotis, S.I. and Bekey, G.A. (2002). Distributed multirobot localization. *IEEE Transactions on Robotics and Automation*, 18(5), 781–795.

Sola, J. (2017). Quaternion kinematics for the error-state Kalman filter. *arXiv preprint, arXiv:1711.02508*.

Uhlmann, J.K. (2003). Covariance consistency methods for fault-tolerant distributed data fusion. *Information Fusion*, 4(3), 201–215.

Zamani, B. and Trumpf, J. (2025). A modular nonlinear continuous-discrete filter for concurrent 3d robot pose and landmark position estimation. In *2025 29th International Conference on Methods and Models in Automation and Robotics (MMAR)*, 467–472. IEEE.