

Article

Online Bias Estimation for Single-Platform Airborne Radar Using Bias-Subspace Information-Guided MAP-EKF

Junwu Luo ¹, Xujun Guan ^{2,*}, Chuang Song ² and Hai Zhang ¹

¹ School of Automation Science and Electrical Engineering, Beihang University, Beijing 100191, China; luojunwu@buaa.edu.cn (J.L.); zhanghai@buaa.edu.cn (H.Z.)

² National Key Laboratory of Complex System Control and Intelligent Agent Cooperation, Beijing 100074, China; 13381055619@163.com

* Correspondence: guanxujun_buaa@163.com

Abstract

Systematic measurement biases are a persistent source of degradation in airborne radar target tracking. Online bias estimation is especially difficult for single-platform operations because only one measurement stream is available, and the target state and radar biases are coupled in the same nonlinear observation model. Under weakly observable geometries, directly augmenting bias states into a recursive filter may lead to slow convergence, while fixed or overly frequent batch optimization may inject updates from weakly informative windows. To address this problem, this paper proposes a Bias Subspace Information-Guided MAP-EKF (maximum a posteriori–extended Kalman filter) (BI-MAP-EKF) for online bias estimation in single-platform airborne radar. A nine-dimensional augmented state jointly describes target motion and range, azimuth, and elevation biases. A posterior Cramér–Rao lower bound (PCRLB) is constructed for this state, where the EKF prior, windowed radar measurements, and process noise propagation are jointly considered. The bias subspace is then extracted by Schur complement, and the resulting azimuth bias PCRLB is used to decide whether the MAP update is reliable enough for EKF injection. The scheduler also includes a cooldown interval and a maximum-window safeguard, which respectively limit excessive updates under informative maneuvers and prevent indefinite waiting under weak geometries. Fifty-run Monte Carlo simulations under different maneuvering conditions show that the proposed scheduler is particularly effective in weakly informative geometries, where it improves azimuth bias and horizontal position estimation while reducing the number of accepted MAP refinements compared with the tested fixed-period MAP and moving-horizon estimation (MHE) baselines. In stronger maneuvering scenarios, it provides a competitive accuracy–cost trade-off rather than uniformly outperforming aggressive MHE in every channel.

Keywords: single-platform radar; systematic bias estimation; weak observability; MAP-EKF



Academic Editor: David J. Daniels

Received: 26 June 2026

Revised: 20 July 2026

Accepted: 1 August 2026

Published: 6 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Airborne radar is a remote sensing payload for unmanned aerial vehicle (UAV) platforms, providing range, angular, and Doppler measurements when optical sensing is unreliable [1,2]. In such radar sensing systems, the measurements are affected not only by random noise and target maneuver uncertainty, but also by systematic errors caused by antenna installation, platform attitude errors, timing offsets, and signal processing delays. These biases are sensor-level errors, but once they enter a recursive tracking estimator,

they are propagated together with the target state and can lead to persistent localization errors [3,4]. The problem studied in this paper is therefore an online radar measurement calibration and sensor data fusion problem rather than a separate offline calibration task.

Most existing bias estimation and sensor registration methods are developed for multi-sensor or cooperative settings. Representative works have addressed coordinate alignment, asynchronous bias estimation, least-squares registration, and performance bounds for cooperative sensing systems [5–7]. In these settings, redundant observations, common targets, and distributed geometric constraints can be used to separate sensor biases from target states [8,9]. Related calibration problems have also been studied in multi-modal sensing systems, including radar–camera and camera–LiDAR–IMU calibration [10–13]. Motion geometry-based decoupling also appears in moving-camera perception, where Ma et al. [14] extracted principal motion from optical flow to separate camera ego motion from foreground motion.

This redundancy is not available in the single-platform airborne radar case considered here. The target is observed only by one onboard radar, and no auxiliary radar, cooperative transponder, or external calibration sources are assumed. As a result, the target state and the radar biases must be inferred from the same nonlinear measurement stream.

The main difficulty appears in the angular channels. A small azimuth bias can be absorbed by a cross-track displacement of the target estimate, so different state–bias combinations may explain nearly the same measurements over a limited interval. Similar geometry-dependent identifiability issues have been reported in bearings-only tracking and target motion analysis, where observer maneuvering is required to improve state distinguishability [15–18]. Although the extended Kalman filter (EKF), unscented Kalman filter (UKF), and cubature Kalman filter (CKF), as well as particle filters, are widely used in nonlinear target tracking [19–22], augmenting the bias terms into the filter state does not by itself remove this ambiguity. In weakly observable geometries, the filter may remain numerically stable while converging slowly or settling near a coupled state–bias solution.

This difficulty is fundamentally related to time-varying observability. Classical nonlinear observability theory shows that state distinguishability depends on the system dynamics and measurement functions [23]. From a statistical viewpoint, the Fisher Information Matrix (FIM) and Cramér–Rao-type bounds quantify how sensing geometry limits achievable estimation accuracy [24–26]. Sliding-window methods, including moving-horizon estimation (MHE) and factor graph smoothing, exploit temporal constraints beyond a single recursive update [27–33]. However, more data do not necessarily imply more usable information. An informative window can help separate the bias from the target state, whereas poor geometry may only reinforce the existing ambiguity. Thus, the issue is whether maximum a posteriori (MAP) refinement is reliable enough for injection into the recursive EKF.

Motivated by this observation, this paper proposes a Bias Subspace Information-Guided MAP-EKF (BI-MAP-EKF) framework for online bias calibration in a single-platform airborne radar sensing system. The recursive augmented-state EKF remains responsible for real-time estimation and covariance propagation, while MAP refinement is used only as an intermittent sensor data fusion update. The trigger is not determined by elapsed time alone. Instead, a posterior Cramér–Rao lower bound (PCRLB) is constructed for the candidate window from the EKF prior, windowed radar measurements, and process noise propagation. The bias subspace bound is extracted by Schur complement, and MAP is injected only when the azimuth bias bound indicates that the candidate window is suitable for a reliable bias update. A cooldown interval and a maximum-window safeguard regularize the injection frequency.

The main contributions of this paper are summarized as follows:

1. The online radar measurement bias calibration problem is formulated for a single-platform airborne radar, where angular biases and target position are strongly coupled under weak geometric excitation.
2. A complete posterior window PCRLB is constructed for MAP window evaluation. Compared with the deterministic backward propagation FIM, the revised formulation incorporates process noise propagation and the EKF prior, so that the bias subspace bound is consistent with the actual MAP-EKF injection problem.
3. A PCRLB-guided, resource-aware MAP-EKF scheduler is proposed. The scheduler combines an azimuth bias PCRLB threshold, a cooldown interval, and a maximum-window safeguard, allowing MAP refinement to be scheduled according to posterior bias subspace information rather than a fixed period alone.
4. Monte Carlo simulations with MHE and periodic MAP baselines show that the proposed scheduler is especially effective in weakly informative geometries, improving azimuth bias and horizontal position estimation with fewer MAP refinements. In stronger maneuvering geometries, it maintains competitive accuracy with reduced MAP calls, rather than claiming uniform dominance in every channel.

The remainder of this paper is organized as follows. Section 2 formulates the augmented state-space model and the nonlinear measurement equations. Section 3 presents the proposed BI-MAP-EKF framework, including the coupling analysis, the recursive PCRLB for maneuver-level diagnosis, and the complete posterior window PCRLB for online MAP triggering. Section 4 evaluates observability characteristics under different maneuvering modes and reports the proposed method's performance with parameter sensitivity analysis. Section 5 provides comparative simulation results with MHE and periodic MAP baselines. Section 6 discusses the limitations of the current approach. Section 7 concludes the paper.

2. Problem Formulation

We consider an airborne radar mounted on a UAV platform tracking a ground object. The radar operates in a spherical coordinate system and provides four types of measurements at each sampling instant: slant range r_m , azimuth angle θ_m , elevation angle φ_m , and radial velocity v_{rm} (derived from Doppler shift). The following assumptions are adopted throughout this work:

- The UAV platform's own position $\mathbf{p}_{u,k} = [x_u, y_u, z_u]^T$ and velocity $\mathbf{v}_{u,k} = [v_{xu}, v_{yu}, v_{zu}]^T$ are known from onboard inertial navigation system (INS)/global navigation satellite system (GNSS) data at time k .
- The target is a low-dynamic ground object; its motion is adequately described by the Constant Velocity (CV) model.
- The radar measurements are corrupted by both zero-mean Gaussian random noise and unknown systematic biases Δr (range), $\Delta\theta$ (azimuth), and $\Delta\varphi$ (elevation). These biases are assumed constant or slowly varying over the tracking interval.
- Only a single airborne platform is available; no cooperative transponder, auxiliary radar, or multi-sensor redundancy can be exploited.

Under these assumptions, the unknown systematic biases are augmented into the target kinematic state to form a joint estimation vector. The target state in Cartesian coordinates is $\mathbf{x}_t = [x \ y \ z \ v_x \ v_y \ v_z]^T$ and the bias vector is $\mathbf{b} = [\Delta r \ \Delta\theta \ \Delta\varphi]^T$. The 9-dimensional augmented state is defined as

$$\mathbf{X} = [\mathbf{x}_t^T \ \mathbf{b}^T]^T = [x \ y \ z \ v_x \ v_y \ v_z \ \Delta r \ \Delta\theta \ \Delta\varphi]^T \quad (1)$$

The discrete-time state dynamics follow

$$\mathbf{X}_k = \mathbf{A}\mathbf{X}_{k-1} + \mathbf{w}_{k-1} \quad (2)$$

where $\mathbf{w}_{k-1} \sim \mathcal{N}(0, \mathbf{Q})$ is the process noise. The state transition matrix is block-diagonal in the kinematic and bias subspaces:

$$\mathbf{A} = \begin{bmatrix} \mathbf{I}_{3 \times 3} & \Delta t \cdot \mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{I}_{3 \times 3} \end{bmatrix} \quad (3)$$

where Δt is the sampling interval. The biases are modeled as a random walk with small process noise intensity to accommodate slow environmental drifts. Since the kinematics are expressed in Cartesian coordinates and the bias subspace evolves as a random constant, the prediction step is strictly linear and does not introduce linearization errors [19].

Let $\Delta \mathbf{p} = \mathbf{p}_{t,k} - \mathbf{p}_{u,k} = [\Delta x, \Delta y, \Delta z]^T$ and $\Delta \mathbf{v} = \mathbf{v}_{t,k} - \mathbf{v}_{u,k} = [\Delta v_x, \Delta v_y, \Delta v_z]^T$ denote the relative position and velocity vectors. Define the slant range $R = \sqrt{\Delta x^2 + \Delta y^2 + \Delta z^2}$, the horizontal range $R_{xy} = \sqrt{\Delta x^2 + \Delta y^2}$, and the unit line-of-sight vector $\mathbf{u} = [u_x, u_y, u_z]^T = \Delta \mathbf{p} / R$. The nonlinear measurement model is

$$\mathbf{z}_k = h(\mathbf{X}_k) + \mathbf{v}_k = \begin{bmatrix} R + \Delta r \\ \arctan(\Delta y / \Delta x) + \Delta \theta \\ \arctan(\Delta z / R_{xy}) + \Delta \varphi \\ (\Delta \mathbf{p}^T \Delta \mathbf{v}) / R \end{bmatrix} + \begin{bmatrix} v_r \\ v_\theta \\ v_\varphi \\ v_{vr} \end{bmatrix} \quad (4)$$

where $\mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R})$ with $\mathbf{R} = \text{diag}(\sigma_r^2, \sigma_\theta^2, \sigma_\varphi^2, \sigma_{vr}^2)$. All four measurement channels are jointly affected by the kinematic states and the biases; the range measurement directly shifts by Δr , the two angular measurements shift by $\Delta \theta$ and $\Delta \varphi$ respectively, while the radial velocity measurement is bias-free.

The Jacobian of the measurement function with respect to the augmented state plays a central role in both the recursive filtering and the FIM-based observability analysis. It takes the form $\mathbf{H}_k = \partial h / \partial \mathbf{X} \in \mathbb{R}^{4 \times 9}$, with the following row-wise structure:

$$\mathbf{H}_k = \begin{bmatrix} u_x & u_y & u_z & 0 & 0 & 0 & 1 & 0 & 0 \\ -\frac{\Delta y}{R_{xy}^2} & \frac{\Delta x}{R_{xy}^2} & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ -\frac{\Delta z \Delta x}{R^2 R_{xy}} & -\frac{\Delta z \Delta y}{R^2 R_{xy}} & \frac{R_{xy}}{R^2} & 0 & 0 & 0 & 0 & 0 & 1 \\ \frac{\partial \dot{R}}{\partial x} & \frac{\partial \dot{R}}{\partial y} & \frac{\partial \dot{R}}{\partial z} & u_x & u_y & u_z & 0 & 0 & 0 \end{bmatrix} \quad (5)$$

where the radial velocity Jacobian block for the position components is $[\partial \dot{R} / \partial x, \partial \dot{R} / \partial y, \partial \dot{R} / \partial z] = \frac{1}{R}(\mathbf{I}_{3 \times 3} - \mathbf{u}\mathbf{u}^T)\Delta \mathbf{v}$. The structure of $\mathbf{H}_k$ directly reflects the geometry of the problem: The first three columns couple position states to all four measurements; columns 4–6 couple velocity only to the radial velocity measurement; and columns 7–9 encode a one-to-one additive bias model for range, azimuth, and elevation, respectively. The fact that the angular bias columns (8 and 9) share nonzero entries only with the position columns (1–3) is the root of the identifiability difficulty analyzed in Section 3.

The estimation objective is to recover, at each time step k , both the target kinematic state $\mathbf{x}_{t,k}$ and the systematic biases $\mathbf{b}$ from the incoming measurement stream $\{\mathbf{z}_1, \dots, \mathbf{z}_k\}$. The central difficulty is that under straight-line flight the position and angular bias columns of $\mathbf{H}_k$ are nearly linearly dependent: a shift in the azimuth bias can be perfectly compen-

sated by a cross-track position adjustment, making the state-bias pair unidentifiable from measurements alone. The platform must therefore maneuver to generate the geometric diversity needed for successful bias decoupling.

3. Methodology

Building on the problem formulated in Section 2, this section first analyzes the state-bias coupling mechanism that limits recursive filter performance. A recursive PCRLB is then introduced as a maneuver-level diagnosis tool for evaluating theoretical estimation bounds under different platform trajectories. Finally, the core contribution—the Bias Subspace Information-Guided MAP-EKF (BI-MAP-EKF)—is presented, where a posterior window PCRLB serves as the online trigger for MAP refinement and evaluates whether the candidate window, together with the EKF prior, is reliable enough for MAP injection.

3.1. State Coupling Mechanism and Observability Analysis

The primary challenge of single-platform augmented bias estimation lies in the geometric coupling between the target's physical position and the radar's angular biases. This coupling renders the system weakly observable under certain trajectories (e.g., straight-line flight) [23].

To illustrate this coupling effect, consider the relative 2D position vector $\mathbf{p}_{xy} = [\Delta x, \Delta y]^T$ and the true azimuth angle $\theta_{true} = \arctan(\Delta y / \Delta x)$. Suppose the radar measures an azimuth $\theta_m = \theta_{true} + \Delta\theta_{true}$, where $\Delta\theta_{true}$ is the true bias. If the filter incorrectly estimates the azimuth bias with a small error $\delta\theta$, such that $\Delta\hat{\theta} = \Delta\theta_{true} + \delta\theta$, the filter compensates by rotating the estimated position vector to match the observation θ_m . The resulting position estimate $\hat{\mathbf{p}}_{xy}$ can be expressed using a 2D rotation matrix:

$$\hat{\mathbf{p}}_{xy} = \mathbf{R}(-\delta\theta)\mathbf{p}_{xy} = \begin{bmatrix} \cos(\delta\theta) & \sin(\delta\theta) \\ -\sin(\delta\theta) & \cos(\delta\theta) \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} \quad (6)$$

Applying the small-angle approximations $\cos(\delta\theta) \approx 1$ and $\sin(\delta\theta) \approx \delta\theta$, the induced position error $\delta\mathbf{p}_{xy}$ is

$$\delta\mathbf{p}_{xy} = \hat{\mathbf{p}}_{xy} - \mathbf{p}_{xy} \approx \begin{bmatrix} 1 & \delta\theta \\ -\delta\theta & 1 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} - \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} \Delta y \\ -\Delta x \end{bmatrix} \delta\theta \quad (7)$$

It can be seen from (7) that an azimuth bias error $\delta\theta$ is equivalent to a cross-track position displacement of magnitude $\|\Delta\mathbf{p}_{xy}\| \cdot \delta\theta$ perpendicular to the line of sight. Under the idealized uniform straight-line geometry, such state bias pairs can become locally indistinguishable over a limited observation interval. The augmented system is therefore rank-deficient along this coupled direction, and no estimator can separate the true state from the coupled alternative without a change in the sensing geometry.

To quantify the theoretical lower bound under different maneuvering modes, the posterior Cramér–Rao Lower Bound (PCRLB) for discrete-time systems follows the Riccati recursion [24,25]:

$$\mathbf{P}_{k|k-1} = \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^T + \mathbf{Q} \quad (8)$$

$$\mathbf{P}_{k|k}^{-1} = (\mathbf{P}_{k|k-1})^{-1} + \mathbf{H}_k^T \mathbf{R}^{-1} \mathbf{H}_k \quad (9)$$

where $\mathbf{F} = \mathbf{A}$ is from (3) and $\mathbf{H}_k$ is from (5). This recursion mirrors the EKF covariance propagation and provides a cumulative information bound over the full trajectory. In Section 4, it is used to compare how different platform trajectories affect the theoretical accuracy limit for bias estimation.

Critically, this bound is **global**; it aggregates past information without isolating whether the *current* window alone is informative. If recent data were collected under a weakly observable geometry, the recursive PCRLB could still appear favorable due to earlier informative measurements. The online trigger therefore requires a window-level metric, developed in the following subsection.

3.2. Bias Subspace Information-Guided MAP-EKF (BI-MAP-EKF)

To make the trigger consistent with the finite-window MAP refinement, we construct a posterior window PCRLB on the same candidate window used by the optimizer. This metric should not be interpreted as the pure geometric information of the window alone. It evaluates whether the selected window, together with the current EKF prior, is reliable enough to be fed back into the EKF.

Recursive EKF Backbone: The 9-dimensional augmented-state EKF backbone operates continuously to provide real-time tracking and propagate the statistical covariance. The complete predict and update iterations are as follows:

Time Update (Prediction):

$$\hat{\mathbf{X}}_{k|k-1} = \mathbf{A}\hat{\mathbf{X}}_{k-1|k-1} \quad (10)$$

$$\mathbf{P}_{k|k-1} = \mathbf{A}\mathbf{P}_{k-1|k-1}\mathbf{A}^T + \mathbf{Q} \quad (11)$$

Measurement Update (Correction):

$$\tilde{\mathbf{z}}_k = \mathbf{z}_k - h(\hat{\mathbf{X}}_{k|k-1}) \quad (12)$$

$$\mathbf{S}_k = \mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^T + \mathbf{R} \quad (13)$$

$$\mathbf{K}_k = \mathbf{P}_{k|k-1}\mathbf{H}_k^T\mathbf{S}_k^{-1} \quad (14)$$

$$\hat{\mathbf{X}}_{k|k} = \hat{\mathbf{X}}_{k|k-1} + \mathbf{K}_k\tilde{\mathbf{z}}_k \quad (15)$$

$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1} \quad (16)$$

where $\tilde{\mathbf{z}}_k$ is the innovation, $\mathbf{S}_k$ is the innovation covariance, and $\mathbf{K}_k$ is the optimal Kalman gain.

Complete Posterior Window PCRLB for MAP Injection: For the candidate window $W_k = [k - N_w + 1, k]$, the window-end state $\mathbf{X}_k$ is selected as the MAP optimization variable. A historical measurement $\mathbf{z}_j$ at time $j \leq k$ is originally a function of the historical state $\mathbf{X}_j$. To express this measurement as a constraint on $\mathbf{X}_k$, the deterministic part of the state transition is written in the backward form

$$\mathbf{X}_j \simeq \Phi_{j,k}\mathbf{X}_k + \boldsymbol{\eta}_{j,k}, \quad j \leq k, \quad (17)$$

where $\Phi_{j,k}$ maps the window-end state back to the historical time j , and $\boldsymbol{\eta}_{j,k}$ represents the equivalent uncertainty induced by process noise over the interval $[j, k]$. Although $\Phi_{j,k}$ is a backward transition, the resulting Jacobian derived below is taken with respect to $\mathbf{X}_k$. Therefore, the information contribution obtained in this section is still an information matrix for the window-end state.

For the constant-velocity model used in this paper, $\Phi_{j,k}$ has the same block form as the state transition matrix in (3), with the time interval replaced by $(t_j - t_k)$. The covariance of the backward equivalent error $\boldsymbol{\eta}_{j,k}$ is denoted by $\Gamma_{j,k}$. In the linear time-invariant case, it can be written as

$$\Gamma_{j,k} = \sum_{\ell=j}^{k-1} \Phi_{j,\ell+1} \mathbf{Q} \Phi_{j,\ell+1}^T, \quad \Gamma_{k,k} = 0. \quad (18)$$

This term accounts for the process uncertainty accumulated when an older measurement is interpreted as a constraint on the current window-end state.

Linearizing the historical measurement model around the state implied by $\mathbf{X}_k$ gives

$$\mathbf{z}_j \simeq h_j(\Phi_{j,k}\mathbf{X}_k) + \mathbf{H}_j\boldsymbol{\eta}_{j,k} + \mathbf{v}_j, \quad (19)$$

where

$$\mathbf{H}_j = \left. \frac{\partial h_j}{\partial \mathbf{X}_j} \right|_{\hat{\mathbf{X}}_{j|k}}, \quad \bar{\mathbf{H}}_{j,k} = \frac{\partial h_j}{\partial \mathbf{X}_k} = \mathbf{H}_j\Phi_{j,k}. \quad (20)$$

Here, $\mathbf{H}_j$ is the local measurement Jacobian with respect to the historical state $\mathbf{X}_j$, while $\bar{\mathbf{H}}_{j,k}$ is the equivalent Jacobian with respect to the window-end state $\mathbf{X}_k$. The effective covariance of the historical measurement constraint becomes

$$\bar{\mathbf{R}}_{j,k} = \mathbf{R} + \mathbf{H}_j\boldsymbol{\Gamma}_{j,k}\mathbf{H}_j^\top. \quad (21)$$

Therefore, the measurement information contributed by the candidate window is

$$\mathbf{J}_{\text{meas},W}(k) = \sum_{j \in W_k} \bar{\mathbf{H}}_{j,k}^\top \bar{\mathbf{R}}_{j,k}^{-1} \bar{\mathbf{H}}_{j,k}. \quad (22)$$

Combining the prior information from the EKF with the window measurement information, the complete posterior window information matrix is written as

$$\mathbf{J}_{\text{post},W}(k) = \mathbf{J}_{\text{prior}}(k) + \mathbf{J}_{\text{meas},W}(k), \quad (23)$$

where $\mathbf{J}_{\text{prior}}(k) = \mathbf{P}_{k|k}^{-1}$ is the EKF posterior information at the window end. In practice, the EKF covariance in this problem converges to a slowly varying steady state under quasi-static geometries, so using $\mathbf{P}_{k|k}$ rather than a strictly measurement-free prior has a negligible effect on the resulting PCRLB trends. The metric is therefore treated as a practical scheduling indicator for the tested operating regime.

The bias subspace is extracted from $\mathbf{J}_{\text{post},W}$ via the Schur complement. Partition $\mathbf{J}_{\text{post},W}$ into kinematic ($\mathbf{x}_t$) and bias ($\mathbf{b}$) blocks:

$$\mathbf{J}_{\text{post},W} = \begin{bmatrix} \mathbf{J}_{xx} & \mathbf{J}_{xb} \\ \mathbf{J}_{bx} & \mathbf{J}_{bb} \end{bmatrix}. \quad (24)$$

The posterior bias subspace information matrix is

$$\mathbf{J}_{b|x} = \mathbf{J}_{bb} - \mathbf{J}_{bx}\mathbf{J}_{xx}^{-1}\mathbf{J}_{xb}, \quad (25)$$

and the corresponding posterior bias PCRLB is $\mathbf{C}_{b|x} = \mathbf{J}_{b|x}^{-1}$. The azimuth bias bound

$$\sigma_{\Delta\theta}^{\text{post}} = \sqrt{[\mathbf{C}_{b|x}]_{2,2}} \quad (26)$$

is used to judge whether the candidate window provides a sufficiently reliable MAP update for injection into the EKF.

The complete MAP injection rule at time step k is

$$\text{Inject MAP if } \begin{cases} \text{first MAP: } T_W(k) \geq T_{\text{init}}, \\ \text{subsequent MAP: } (\sigma_{\Delta\theta}^{\text{post}} \leq \tau_\theta \vee T_W(k) \geq W_{\text{max}}) \wedge (t_k - t_{\text{last}} \geq \Delta t_{\text{min}}), \end{cases} \quad (27)$$

where $T_W(k)$ is the duration of the current candidate MAP window, $W_{\max}$ is the maximum allowed window time, T_{init} is the minimum initial window time, and $\Delta t_{\min}$ is the cooldown interval.

The three terms in (27) serve distinct roles:

- *PCRLB threshold* τ_θ evaluates whether the candidate window, together with the EKF prior, provides sufficient azimuth-bias information. When $\sigma_{\Delta\theta}^{\text{post}}$ falls below τ_θ , the candidate window is considered sufficiently informative for MAP injection.
- *Maximum-window trigger* $W_{\max}$ prevents indefinite postponement when the geometry remains persistently weak. If the candidate window reaches $W_{\max}$ before the PCRLB threshold is satisfied, MAP is activated by the maximum-window safeguard.
- *Cooldown interval* $\Delta t_{\min}$ limits MAP frequency when the PCRLB condition is repeatedly satisfied under highly informative maneuvers.

Based on empirical tuning (Section 4), the parameters are set to $\tau_\theta = 0.06^\circ$, $\Delta t_{\min} = 15$ s, $W_{\max} = 80$ s (including 20 s overlap), and the first MAP is forced at $T_{\text{init}} = 30$ s with checks performed every 5 s.

MAP Optimization with EKF Prior: When the trigger conditions are satisfied, the MAP estimation maximizes the posterior probability of the window-end state $\mathbf{X}_k$ given the windowed measurements $\mathbf{Z}_{k-N_w:k}$,

$$\hat{\mathbf{X}}_{MAP} = \arg \max_{\mathbf{X}_k} p(\mathbf{X}_k | \mathbf{Z}_{k-N_w:k}) \propto \arg \max_{\mathbf{X}_k} (p(\mathbf{Z}_{k-N_w:k} | \mathbf{X}_k) p(\mathbf{X}_k)), \quad (28)$$

where $\hat{\mathbf{X}}_{k|k}$ and $\mathbf{P}_{k|k}$ from the EKF serve as the Gaussian prior. Converting to a nonlinear least-squares cost yields

$$\hat{\mathbf{X}}_{MAP} = \arg \min_{\mathbf{X}_k} \left\{ \frac{1}{2} \|\mathbf{X}_k - \hat{\mathbf{X}}_{k|k}\|_{\mathbf{P}_{k|k}^{-1}}^2 + \frac{1}{2} \sum_{i=k-N_w}^k \|\mathbf{z}_i - h_i(\mathbf{X}_k)\|_{\mathbf{R}^{-1}}^2 \right\}. \quad (29)$$

The cost is minimized via Gauss–Newton iteration. Upon convergence, the posterior covariance is

$$\mathbf{P}_{MAP} = \left(\mathbf{H}_{aug}^T \mathbf{R}_{aug}^{-1} \mathbf{H}_{aug} + \mathbf{P}_{k|k}^{-1} \right)^{-1}, \quad (30)$$

where $\mathbf{H}_{aug}$ stacks the window Jacobians and $\mathbf{R}_{aug}$ is the block-diagonal measurement noise covariance. Including $\mathbf{P}_{k|k}^{-1}$ is essential. Using only the measurement FIM may underestimate the uncertainty of bias components weakly constrained within the window, which can lead to over-confident covariance injection into the subsequent EKF recursion.

The complete pipeline of the proposed BI-MAP-EKF is summarized in Figure 1. The EKF runs at every filtering step, while the posterior window PCRLB is evaluated only at admissible check instants. When either the PCRLB condition or the maximum-window condition is satisfied after the cooldown interval, MAP refinement is performed, and the optimized state-covariance pair is injected back into the EKF.

In the implementation, $\tilde{\mathbf{H}}_{j,k}$ is evaluated directly by using the relative time offset $t_j - t_k$, which is equivalent to applying the chain rule $\mathbf{H}_j \Phi_{j,k}$.




Figure 1. Workflow of the Bias Subspace Information-Guided maximum a posteriori (MAP)–extended Kalman filter (EKF), abbreviated BI-MAP-EKF. The EKF runs continuously while the posterior Cramér–Rao lower bound (PCRLB) is checked at admissible instants. MAP refinement is performed under the initialization, PCRLB, or maximum-window condition, subject to cooldown, and the optimized state and covariance are then injected into the EKF.

3.3. Computational Complexity

Let n be the dimension of the augmented state, m be the number of measurements per time step, N_w be the number of measurements in the MAP window, and I be the number of Gauss–Newton iterations. The EKF recursion has a per-step cost dominated by the covariance propagation and update. Since $n = 9$ and $m = 4$ in the considered problem, these operations are nearly constant in practice.

Let N_c denote the number of admissible PCRLB-check instants, and let N_M denote the number of accepted MAP refinements. Beyond the EKF recursion, BI-MAP-EKF has two additional costs. First, at each check instant, the posterior window PCRLB must be evaluated by accumulating the equivalent Jacobians and effective covariance terms over the candidate window. This information-evaluation layer has a cost of approximately $O(N_w m n^2)$ per check and is incurred even when MAP is not finally accepted. Second, once a MAP update is accepted, the Gauss–Newton solver requires residual and Jacobian accumulation over the selected window and solution of the $n \times n$ normal equation, with cost approximately $O(I N_w m n^2 + I n^3)$.

Thus, the additional cost can be summarized as

$$O(N_c N_w m n^2) + O(N_M (I N_w m n^2 + I n^3)). \quad (31)$$

The first term corresponds to window information evaluation, and the second term corresponds to accepted MAP optimization. Therefore, under the same number of MAP refinements, BI-MAP-EKF may still have extra PCRLB-checking overhead. Its runtime benefit mainly comes from reducing unnecessary MAP optimizations through scheduling, rather than from a cheaper MAP solver.

4. Simulation Part I: Observability Evaluation and Method Validation

Simulation experiments are conducted in two stages. First, maneuver-dependent observability characteristics are evaluated using a recursive PCRLB analysis, which provides a theoretical diagnosis of how different platform trajectories influence the estimability of radar biases. Second, the proposed Bias Subspace Information-Guided MAP-EKF (BI-MAP-EKF) is examined through a representative single-run experiment and a Monte Carlo comparison of different MAP scheduling strategies. The former illustrates the correction mechanism in detail, whereas the latter evaluates how information-guided scheduling changes estimation accuracy, MAP call frequency, and runtime relative to fixed-period alternatives. The core simulation parameters are summarized in Table 1. These settings are reported explicitly to make the radar sensing noise model, MAP-window configuration, and online fusion procedure reproducible. All simulations are performed in MATLAB R2022a on an Intel Core i7-13790F processor.

Table 1. Radar/process-noise and Bias Subspace Information-Guided maximum a posteriori-extended Kalman filter (BI-MAP-EKF) scheduling parameters.

Parameter	Value
Range Noise (σ_r)	5.0 m
Azimuth Noise (σ_θ)	0.12 deg
Elevation Noise (σ_ϕ)	0.12 deg
Radial Velocity Noise (σ_{vr})	0.6 m/s
Sampling Interval (Δt)	0.1 s
Maximum MAP window time ($W_{\max}$)	80 s (including 20 s overlap)
MAP Cooldown Interval ($\Delta t_{\min}$)	15 s
MAP Trigger Threshold (τ_θ)	0.06°
Initialization MAP trigger (T_{init})	30 s
PCRLB check interval	5 s
Maximum MAP Iterations	8
Process noise (discrete-time, $\Delta t = 0.1$ s)	
Target accel. std (σ_a)	0.5 m/s ² (Scenario A)/0.4 m/s ² (Scenario B)
Range bias RW std (σ_{b_r})	0.05 m
Azimuth bias RW std (σ_{b_θ})	0.005°
Elevation bias RW std (σ_{b_ϕ})	0.005°

4.1. Observability Comparison Under Different Maneuvering Modes

To clarify how platform maneuvering affects bias-state observability, three representative trajectories are compared: circling, constant-velocity straight-line, and accelerated straight-line, as shown in Figure 2. These trajectories provide different levels of geometric diversity and are used as typical diagnostic cases rather than an exhaustive set of possible flight profiles.




Figure 2. Platform and target trajectories used for maneuver-level PCRLB analysis: (a) circling; (b) constant-velocity straight-line motion; and (c) accelerated straight-line motion. Filled and open markers denote the start and end positions, respectively.

The recursive PCRLB results in Figure 3 show that the range bias is reduced rapidly in all three modes, which is consistent with its direct additive coupling to the range measurement. The azimuth bias, however, is much more sensitive to the platform trajectory. Under constant-velocity flight, the azimuth bound remains high because the azimuth bias and the cross-track target position are difficult to decouple. Accelerated straight-line motion provides partial improvement by changing the relative geometry over time, whereas circling leads to the fastest decrease due to continuous line-of-sight rotation. The elevation bias shows an intermediate trend. These results indicate that the azimuth bias is the most geometry-sensitive component, which motivates its use as the primary indicator in the MAP-window scheduling rule.

This recursive PCRLB is used only as a maneuver-level diagnostic tool. It explains why geometric excitation is important for bias identifiability, but it is not directly used as the online MAP trigger. The trigger itself is based on the posterior window PCRLB introduced in Section 3.2.




Figure 3. Recursive PCRLB evolution of (a) range bias, (b) azimuth bias, and (c) elevation bias under the three platform maneuvers. Lower bounds indicate stronger trajectory-induced identifiability.

4.2. Single-Run Demonstration Under Two Maneuver Scenarios

To illustrate how the proposed method behaves under qualitatively different observability conditions, two representative single-run experiments are presented: Scenario A

employs a mild heading-change maneuver, while Scenario B follows an aggressive accelerating straight-line profile with stronger observability. The proposed method uses the balanced parameter set determined in Section 4.3: $\tau_\theta = 0.06^\circ$, $\Delta t_{\min} = 15$ s, and $W_{\max} = 80$ s (including 20 s overlap). The scheduling metric is the complete posterior window PCRLB described in Section 3.2, which combines the EKF prior with process noise-consistent window measurement information and extracts the azimuth bias bound through the Schur complement. In the scheduling figures, three trigger types are distinguished by color and marker: initialization trigger (blue diamond), PCRLB trigger (teal circle), and maximum-window trigger (amber square).

In the scheduling subfigures, the upper panel shows the posterior azimuth bias PCRLB evaluated at the admissible check instants. The dashed horizontal line denotes the threshold τ_θ . The lower panel shows the MAP windows actually used for refinement; each horizontal bar represents the time span of one accepted MAP window, and the marker at the right end indicates the injection time and trigger type. A PCRLB trigger means that $\sigma_{\Delta\theta}^{\text{post}} \leq \tau_\theta$, whereas a maximum-window trigger means that the candidate window has reached the maximum allowed time $W_{\max}$ before the PCRLB condition is satisfied.

4.2.1. Scenario A: Mild-Maneuver Experiment

Scenario A is used as a weak and uneven information case. The platform performs only a mild heading change, so the angular bias information is accumulated gradually and unevenly. This scenario examines the conservative side of the scheduler; when the angular information grows slowly, the method should avoid early low-information injection, but it should also avoid waiting indefinitely.

As shown in Figure 4, the posterior azimuth bias bound decreases only gradually after the initial maneuver. The first MAP update is forced after the minimum initial window. Among the six subsequent updates, four refinements are triggered by the PCRLB condition, while two updates are activated by the maximum-window trigger. This behavior indicates that the maximum-window safeguard remains necessary when the posterior bound does not fall below the threshold within the allowed window length.




Figure 4. Scenario A (mild maneuver): (a) platform target geometry and speed profile; (b) posterior azimuth bias PCRLB at the admissible check instants (upper) and accepted MAP windows (lower). Horizontal bars show the window spans, and right-end markers identify the injection times and trigger types.

Figure 5 shows that the range and elevation biases settle relatively quickly, whereas the azimuth bias requires a longer transient. This is consistent with the weaker angular observability in Scenario A and explains why the azimuth bias PCRLB is used as the scheduling indicator.




Figure 5. Single-run estimates of the range, azimuth, and elevation biases in Scenario A. Magenta markers denote accepted MAP injection instants.

4.2.2. Scenario B: Strong-Maneuver Experiment

Scenario B is used as a more informative case. The stronger acceleration profile alters the relative geometry more clearly, so the posterior azimuth bias bound is expected to satisfy the PCRLB threshold more regularly. This case highlights the role of the cooldown interval when the information condition is frequently met.

Figure 6 shows that the posterior azimuth bias bound falls below the threshold at nearly every admissible check. Therefore, no maximum-window trigger activation is needed in this scenario. The accepted MAP updates are mainly limited by the cooldown interval, which prevents excessive optimization when the information condition is repeatedly satisfied.




Figure 6. Scenario B (strong maneuver): (a) platform target geometry and speed profile; (b) posterior azimuth bias PCRLB at the admissible check instants (upper) and accepted MAP windows (lower). Horizontal bars show the window spans, and right-end markers identify the injection times and trigger types.

The bias estimates in Figure 7 converge faster than in Scenario A because the maneuver provides stronger geometric excitation. Together, the two scenarios show the two roles of the proposed scheduler: the maximum-window trigger prevents indefinite waiting under weak information, while the cooldown interval limits unnecessary MAP calls under informative motion.




Figure 7. Single-run estimates of the range, azimuth, and elevation biases in Scenario B. Magenta markers denote accepted MAP injection instants.

Table 2 summarizes the single-run bias accuracy and trigger composition of the two demonstration scenarios. The bias RMSEs are computed over the final 50 s, so they reflect the post-convergence behavior rather than the initial transient. Scenario A requires maximum-window activations, indicating that the mild maneuver does not always satisfy the PCRLB threshold before the window reaches $W_{\max}$. In contrast, Scenario B is dominated by PCRLB-triggered updates, and the shorter average window length shows that informative measurements arrive more regularly. The difference between the two rows explains why the same scheduler behaves differently in the two cases: in Scenario A it sometimes waits until the window reaches $W_{\max}$, whereas in Scenario B the PCRLB condition is usually satisfied before the maximum-window limit is reached.

Table 2. Final-50-s single-run bias root-mean-square error (RMSE) and MAP-trigger statistics for Scenarios A and B.

Scenario	Range RMSE (m)	Azimuth RMSE (deg)	Elevation RMSE (deg)	Init. Trigger	PCRLB Trigger	Max-Win. Trigger	Total MAP	Avg. Win. Len. (s)
Scenario A	0.56	0.036	0.027	1	4	2	7	49.9
Scenario B	0.98	0.023	0.015	1	13	0	14	35.6

4.3. Parameter Analysis and Strategy Comparison

4.3.1. Parameter Sensitivity Analysis

The scheduler parameters are examined through 50-run Monte Carlo sweeps. In all sweeps, the same complete posterior window PCRLB as in Section 3.2 is used. The sensitivity analysis provides a trend analysis for each parameter under the tested ranges rather than a global optimum search. In each sweep, one parameter is varied while the other two are held at their baseline values ($\Delta t_{\min} = 15$ s, $W_{\max} = 80$ s, $\tau_{\theta} = 0.06^{\circ}$). The tested ranges are chosen according to online implementation constraints and the physical meaning of each parameter. In particular, an overly large maximum window increases latency, computational load, and the risk of model mismatch within the constant-velocity target model; an overly small cooldown interval may cause excessive MAP refinements under informative maneuvers; and an overly loose PCRLB threshold may accept short

windows with weaker bias state separation. Therefore, when the best value within a scanned range appears at a boundary, it should be interpreted as a trend under the tested constraints rather than as evidence of a global optimum.

The accuracy–runtime trade-off for each sweep is shown in Figure 8. In the maximum-window sweep (Scenario A), $W_{\max} = 80$ s achieves the best azimuth accuracy within the tested delay limit, suggesting that the weak-information case benefits from longer information accumulation within the tested range. In the cooldown sweep (Scenario B), $\Delta t_{\min} = 5$ s gives the lowest azimuth RMSE but at 38.0 MAP calls per run, whereas $\Delta t_{\min} = 15$ s reduces this to 14.0 calls with only a 0.0024° accuracy loss, showing how the update frequency can be limited under strong maneuvering. The selected setting is therefore treated as a practical operating point under the online latency and model validity constraints, rather than as a global optimum. The numerical results of all tested parameter values are summarized in Table 3.




Figure 8. Accuracy–runtime trade-off across the three parameter sweeps: (a) cooldown $\Delta t_{\min}$ sweep; (b) max window $W_{\max}$ sweep; (c) threshold τ_θ sweep. Each marker corresponds to one tested parameter value.

Table 3. Sensitivity of estimation accuracy, MAP calls, and runtime to the three scheduling parameters over 50 Monte Carlo runs. Boldface marks the best value within each sweep, and the selected operating point is listed in the last column.

Sweep	Scn.	Value	Az. RMSE (deg)	XY RMSE (m)	MAP /Run	Runtime (s)	Selected
$\Delta t_{\min}$	B	5 s	0.023 ± 0.011	8.64 ± 6.01	38.0	0.55	$\Delta t_{\min} = 15$ s
		10 s	0.028 ± 0.011	10.06 ± 6.08	20.0	0.37	
		15 s	0.026 ± 0.009	8.47 ± 5.27	14.0	0.30	
		20 s	0.029 ± 0.012	10.04 ± 6.22	11.0	0.26	
$W_{\max}$	A	40 s	0.072 ± 0.052	17.80 ± 14.03	11.0	0.26	$W_{\max} = 80$ s
		50 s	0.071 ± 0.048	17.32 ± 13.01	8.0	0.27	
		60 s	0.059 ± 0.045	14.44 ± 11.96	7.0	0.29	
		70 s	0.052 ± 0.033	12.17 ± 8.89	7.0	0.30	
		80 s	0.047 ± 0.033	11.15 ± 8.83	7.0	0.29	
τ_θ	A	0.06°	0.047 ± 0.033	11.15 ± 8.83	7.0	0.31	$\tau_\theta = 0.06^\circ$
		0.08°	0.049 ± 0.033	11.72 ± 8.75	11.0	0.36	
		0.10°	0.049 ± 0.033	11.72 ± 8.75	11.0	0.34	
		0.12°	0.049 ± 0.033	11.72 ± 8.75	11.0	0.32	
		0.15°	0.078 ± 0.059	19.49 ± 15.74	13.8	0.31	

4.3.2. Comparison with Periodic MAP Baselines

Based on the sensitivity analysis, a balanced parameter set is selected: $\Delta t_{\min} = 15$ s, $W_{\max} = 80$ s, and $\tau_\theta = 0.06^\circ$. This configuration is evaluated against three periodic MAP

baselines (10 s/30 s, 20 s/30 s, and 10 s/50 s) and EKF only over 50 MC runs on the mild maneuver trajectory.

In general, increasing the MAP update frequency or enlarging the optimization window tends to improve estimation accuracy because the former applies corrections more frequently and the latter accumulates more measurement information. Both changes, however, increase computational cost or update latency. The three periodic baselines are therefore arranged as controlled comparisons. Per 10 s/30 s and Per 20 s/30 s use the same 30 s window and isolate the effect of update frequency, whereas Per 10 s/30 s and Per 10 s/50 s use the same 10 s interval and isolate the effect of window length. Per 10 s/50 s represents a relatively aggressive high-frequency, long-window configuration. This design allows us to examine whether the gain of BI-MAP-EKF comes simply from more frequent or longer optimization or from selecting the MAP injection times according to the posterior window PCRLB.

Table 4 reports the quantitative results, and Figure 9 shows the corresponding azimuth-RMSE evolution. The periodic baselines follow the expected trend. With the window fixed at 30 s, shortening the update interval from 20 s to 10 s reduces the azimuth RMSE from 0.105° to 0.076° but increases the MAP calls from 11 to 22 and the runtime from 0.17 s to 0.31 s. With the interval fixed at 10 s, enlarging the window from 30 s to 50 s further reduces the azimuth RMSE to 0.056° , while increasing the runtime to 0.45 s. These results confirm that higher update frequency and longer windows can improve accuracy, but at additional computational cost.

BI-MAP-EKF achieves an azimuth RMSE of 0.047° with 7 accepted MAP updates. Compared with the aggressive Per 10 s/50 s baseline, it provides lower azimuth error with fewer MAP calls and lower runtime. Compared with Per 10 s/30 s, it achieves better accuracy and fewer MAP calls at a similar runtime. In contrast, Per 20 s/30 s remains faster because it uses short, infrequent optimization, but its azimuth error is substantially larger. Therefore, the proposed method should not be interpreted as universally faster than periodic MAP. Its time advantage appears mainly relative to high-frequency and/or long-window optimization, while its broader benefit is the selective use of MAP when the available window is informative. The RMSE evolution in Figure 9 also shows that the largest difference between the scheduling strategies appears in the azimuth channel, which is the dominant coupled direction in the mild-maneuver case.




Figure 9. Ensemble azimuth-bias RMSE over 50 Monte Carlo runs for BI-MAP-EKF, three periodic MAP configurations, and EKF-only in Scenario A. The inset enlarges the final 50 s.

Table 4. Estimation RMSE and computational statistics for BI-MAP-EKF, periodic MAP configurations, and EKF-only in Scenario A over 50 Monte Carlo runs. RMSE values are reported as mean $\pm$ standard deviation.

Method	Proposed	Per 10 s/30 s	Per 20 s/30 s	Per 10 s/50 s	EKF
Metric					
Range RMSE (m)	0.64 $\pm$ 0.42	0.76 $\pm$ 0.50	0.78 $\pm$ 0.52	0.64 $\pm$ 0.41	0.93 $\pm$ 0.68
Azimuth RMSE (deg)	0.047 $\pm$ 0.033	0.076 $\pm$ 0.050	0.105 $\pm$ 0.084	0.056 $\pm$ 0.039	0.591 $\pm$ 0.438
Elevation RMSE (deg)	0.018 $\pm$ 0.005	0.021 $\pm$ 0.008	0.028 $\pm$ 0.015	0.019 $\pm$ 0.006	0.048 $\pm$ 0.034
XY Position RMSE (m)	11.15 $\pm$ 8.83	19.30 $\pm$ 13.20	26.80 $\pm$ 22.00	13.70 $\pm$ 10.36	155.10 $\pm$ 115.40
Time/run (s)	0.31	0.31	0.17	0.45	0.03
MAP calls/run	7.0	22.0	11.0	20.0	0.0

Per 10 s/30 s: periodic MAP with 10 s interval and 30 s window; Per 20 s/30 s: 20 s interval and 30 s window; Per 10 s/50 s: 10 s interval and 50 s window.

5. Simulation Part II: MHE and MAP Strategy Comparison

This section compares BI-MAP-EKF with three optimization-based baselines under the two maneuver scenarios introduced in Section 4.2. The two MHE variants represent fixed-lag refinement with different update frequencies and window lengths: MHE 5 s/50 s is an aggressive setting with frequent long-window optimization, whereas MHE 15 s/30 s is a lighter setting with fewer updates and a shorter window. The periodic MAP baseline uses the same augmented state model but injects MAP refinements at a fixed interval, without checking whether the accepted window is informative for azimuth bias separation. A full-batch factor graph smoother is not included as a direct online baseline because it solves a different smoothing problem and may use measurements beyond the current filtering instant. The two fixed-lag MHE configurations are therefore used as causal smoothing baselines under the same measurement history. Incremental factor graph implementations will be considered in future work.

All methods use the same radar measurement sequences and the same target bias state definition. The comparison is therefore not intended to show that one optimizer is universally better in every channel. Instead, it examines whether the proposed scheduler can use fewer accepted MAP refinements while preserving the accuracy gains that matter most for this problem, namely azimuth bias and horizontal position estimation.

5.1. Scenario A: Mild-Maneuver Comparison

Scenario A uses the mild heading-change trajectory. This is the more difficult case for online bias calibration because the angular information is not distributed uniformly over time. The range bias is directly observed through the range channel, but the azimuth bias is mainly separated from horizontal target position when the platform geometry changes. As a result, a fixed-period optimizer may refine the estimate even when the current window adds little new azimuth bias information.

This scenario is therefore used as the main test for the proposed scheduling idea. If the PCRLB-based window check is meaningful, the method should not simply increase the number of MAP refinements. It should instead reject or postpone low-information windows and keep the useful correction effect of long-window optimization with fewer accepted MAP calls.

Figure 10 gives a compact view of the azimuth error distribution and the runtime–accuracy trade-off, while Table 5 lists the full bias and position metrics. The main gain appears in the azimuth and horizontal-position channels. This is consistent with the coupling analysis in Section 3.1: The dominant ambiguity in this trajectory is not the range bias, but the trade-off between azimuth bias and cross-track target position.

The two MHE settings illustrate the accuracy–cost tension. MHE 5 s/50 s performs frequent long-window refinements and therefore keeps a good azimuth estimate, but the number of optimizations is high. MHE 15 s/30 s reduces runtime, but the shorter and less frequent windows provide weaker bias state separation in this trajectory. Periodic MAP 10 s/50 s benefits from a long window, but it still refines according to elapsed time rather than window informativeness. BI-MAP-EKF differs from these baselines by accepting fewer but more selective MAP updates. Its advantage in Scenario A should therefore be interpreted as a scheduling gain under weak and uneven observability rather than as a cheaper MAP solver. Its runtime is lower than that of the aggressive MHE 5 s/50 s and Periodic MAP 10 s/50 s baselines but higher than that of the lightweight MHE 15 s/30 s configuration. Therefore, no universal runtime advantage is claimed; the gain lies in obtaining better azimuth/XY accuracy with fewer accepted MAP refinements under weak and uneven information.




Figure 10. Scenario A (mild maneuver): (a) distribution of the per-run azimuth-bias RMSE; (b) mean azimuth bias RMSE versus runtime for each method.

Table 5. Estimation RMSE and computational statistics for moving-horizon estimation (MHE), periodic MAP, and BI-MAP-EKF in Scenario A from 50 Monte Carlo runs. RMSE values are reported as mean $\pm$ standard deviation.

Method	MHE 5 s/50 s	MHE 15 s/30 s	Periodic MAP 10 s/50 s	Proposed
Metric				
Range RMSE (m)	0.68 $\pm$ 0.35	0.41 $\pm$ 0.32	0.64 $\pm$ 0.41	0.64 $\pm$ 0.42
Azimuth RMSE (deg)	0.061 $\pm$ 0.049	0.074 $\pm$ 0.071	0.056 $\pm$ 0.039	0.047 $\pm$ 0.033
Elevation RMSE (deg)	0.107 $\pm$ 0.065	0.020 $\pm$ 0.015	0.019 $\pm$ 0.006	0.018 $\pm$ 0.005
XY Position RMSE (m)	15.32 $\pm$ 12.54	19.22 $\pm$ 18.88	13.70 $\pm$ 10.36	11.15 $\pm$ 8.83
Runtime/run (s)	0.81	0.19	0.44	0.31
MAP calls/run	39.0	15.0	20.0	7.0

5.2. Scenario B: Strong-Maneuver Comparison

Scenario B uses the accelerating straight-line platform. Compared with Scenario A, the stronger maneuver produces informative windows more regularly. This changes the role of the scheduler. In Scenario A, the main risk is accepting weak windows too early. In Scenario B, the main risk is the opposite: The PCRLB condition can be satisfied frequently, so the cooldown interval becomes important for limiting unnecessary refinements.

For this reason, Scenario B is not expected to produce the same relative advantage as Scenario A. When useful information is available in most windows, a fixed-period MAP strategy is less likely to choose a severely poor window, and aggressive MHE can also benefit from frequent refinement. This scenario is included to check whether BI-MAP-EKF remains competitive when the observability condition is already favorable.

The results in Figure 11 and Table 6 follow this expectation. MHE 5 s/50 s obtains the smallest azimuth RMSE, which is reasonable because it applies the most aggressive long-window refinement. The periodic MAP baseline is also strong in this scenario, since fixed intervals are less harmful when informative windows occur regularly. BI-MAP-EKF does not achieve the minimum azimuth RMSE here. Its role is better described as a lower-call alternative; it keeps the azimuth error in a competitive range, gives the lowest XY position error among the tested methods, and uses fewer accepted MAP refinements than the high-rate MHE and periodic MAP settings. Its runtime is lower than those of MHE 5 s/50 s and Periodic MAP 10 s/50 s but not lower than that of the lightweight MHE 15 s/30 s setting.




Figure 11. Scenario B (strong maneuver): (a) distribution of the per-run azimuth bias RMSE; (b) mean azimuth bias RMSE versus runtime for each method.

Table 6. Estimation RMSE and computational statistics for MHE, periodic MAP, and BI-MAP-EKF in Scenario B over 50 Monte Carlo runs. RMSE values are reported as mean $\pm$ standard deviation.

Method	MHE 5 s/50 s	MHE 15 s/30 s	Periodic MAP 10 s/50 s	Proposed
Metric				
Range RMSE (m)	0.59 $\pm$ 0.39	0.27 $\pm$ 0.21	0.76 $\pm$ 0.50	0.77 $\pm$ 0.56
Azimuth RMSE (deg)	0.019 $\pm$ 0.013	0.029 $\pm$ 0.022	0.024 $\pm$ 0.011	0.026 $\pm$ 0.009
Elevation RMSE (deg)	0.096 $\pm$ 0.041	0.024 $\pm$ 0.015	0.017 $\pm$ 0.004	0.019 $\pm$ 0.005
XY Position RMSE (m)	8.90 $\pm$ 6.40	14.00 $\pm$ 10.00	8.70 $\pm$ 6.10	8.50 $\pm$ 5.30
Runtime/run (s)	0.82	0.19	0.45	0.29
MAP calls/run	39.0	15.0	20.0	14.0

This result also clarifies the boundary of the proposed scheduler. When computation is unconstrained and the geometry is consistently informative, aggressive smoothing can obtain a slightly better azimuth estimate. The proposed method is designed for online operation, where every MAP injection has to be justified by its information contribution and computational cost. Therefore, the value of BI-MAP-EKF in Scenario B lies less in absolute azimuth dominance and more in maintaining competitive azimuth/XY accuracy with fewer refinements than the high-frequency, long-window baselines.

Taken together, the two scenarios show that the clearest benefit of BI-MAP-EKF appears when the information content of the candidate window varies over time. In the mild-maneuver case, the posterior PCRLB check postpones or rejects low-value refinements and retains the accuracy benefit of long-window MAP with fewer accepted updates. In the strong-maneuver case, informative windows occur regularly, so fixed-period MAP and aggressive MHE are already competitive; the cooldown interval then mainly prevents the scheduler from becoming an unnecessarily frequent optimizer.

BI-MAP-EKF should therefore be interpreted as an information-aware MAP scheduler rather than a universally more accurate or faster optimizer. Its runtime benefit appears only when the reduction in MAP optimizations is large enough to offset the additional posterior-PCRLB evaluation, particularly relative to high-frequency and/or long-window baselines. Its primary purpose is to determine when MAP refinement is worth injecting into a recursive online estimator.

6. Discussion

The simulation results support the proposed scheduling idea, but the current validation is mechanism-level rather than field-level. Public airborne radar datasets rarely provide synchronized range, azimuth, elevation, and Doppler measurements together with independently calibrated systematic-bias ground truth, which limits direct evaluation of online bias registration. Future work will therefore include higher-fidelity measurement models, hardware-in-the-loop tests, and real-data validation when suitable calibration references are available. One technical limitation is that the posterior window PCRLB computation uses the current EKF covariance as a prior. When the EKF has not yet converged, the prior may be overly diffuse, making the PCRLB metric too conservative. The initialization trigger at $T_{\text{init}} = 30$ s partially compensates for this, but a more adaptive prior-scaling strategy could be considered. Under non-Gaussian noise or intermittent outliers, the quadratic MAP cost may produce over-weighted residuals; robust loss functions or measurement quality gating would be required.

The present PCRLB trigger is derived under a CV target model, Gaussian measurement noise, a slowly varying bias model, and accurate platform state assumptions. If the target performs abrupt maneuvers, the MAP window may contain model mismatch and the PCRLB may become overconfident. If radar biases change abruptly, the random-walk model may delay adaptation, and a change detector or IMM-type bias model would be needed. If the platform INS/GNSS solution contains non-negligible errors, the platform state uncertainty should be augmented into the state or folded into the measurement covariance.

Another limitation is that the parameter tuning relies on empirical sweeps. The scheduling parameters are not claimed to be globally optimal; they are selected from bounded ranges that reflect the intended online operating regime. Extending these ranges may further improve one metric, such as steady-state RMSE or runtime, but it would also change the latency, responsiveness, or model validity assumptions of the estimator. The sensitivity analysis in Section 4.3 clarifies the main accuracy, runtime, and update frequency trade-offs around the selected values, but the preferred operating point may change for substantially different platform dynamics or sensor characteristics. A systematic, context-aware parameter selection guideline would improve practical applicability.

Finally, the current evaluation is limited to simulated trajectories with a single non-maneuvering target. For other flight geometries, the expected scheduling behavior follows the same principle: Slowly varying line-of-sight geometry tends to produce longer accepted windows or maximum-window activations, whereas persistently informative geometry shifts the main scheduling role to the cooldown interval. In multi-target or cluttered environments, data association errors could corrupt the MAP window and degrade the trigger

decisions. Extending the framework to incorporate measurement-to-target assignment uncertainty is a direction for future work.

7. Conclusions

The results suggest that the posterior window PCRLB is useful for deciding when MAP refinement should be injected into a recursive EKF. The benefit is most visible when the available geometry is weak or uneven, where fixed-period optimization may spend updates on windows with limited azimuth bias information. Under stronger maneuvering, aggressive MHE can still obtain the smallest azimuth RMSE, but the proposed scheduler provides a lower-call alternative with competitive accuracy. Given that posterior window PCRLB evaluation introduces additional computation, a runtime advantage is expected only when the resulting reduction in MAP refinements offsets this overhead, as observed relative to the high-frequency or long-window baselines. Future work will focus on real-data validation, abrupt bias changes, and extensions to platform-state uncertainty.

Author Contributions: Conceptualization, J.L. and X.G.; methodology, J.L.; software, J.L. and C.S.; validation, X.G. and H.Z.; formal analysis, J.L.; investigation, X.G. and C.S.; data curation, X.G.; writing—original draft preparation, J.L.; writing—review and editing, H.Z., X.G. and C.S.; visualization, J.L.; supervision, X.G.; project administration, X.G. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported by the National Natural Science Foundation of China (No. 62373031).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The simulation data and source code that support the findings of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Corradi, F.; Fioranelli, F. Radar perception for autonomous unmanned aerial vehicles: A survey. In *Proceedings of System Engineering for Constrained Embedded Systems (DroneSE and RAPIDO '22)*, Budapest, Hungary, 21 June 2022; Association for Computing Machinery: New York, NY, USA, 2022; pp. 14–20. <https://doi.org/10.1145/3522784.3522787>.
- Zhou, X.; Jia, W.; He, R.; Sun, W. High-precision localization tracking and motion state estimation of ground-based moving target utilizing unmanned aerial vehicle high-altitude reconnaissance. *Remote Sens.* **2025**, *17*, 735. <https://doi.org/10.3390/rs17050735>.
- Bar-Shalom, Y.; Li, X.R.; Kirubarajan, T. *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*; John Wiley & Sons: New York, NY, USA, 2002. <https://doi.org/10.1002/0471221279>.
- Ren, B.; Jia, T.; Liu, H.; Wang, Y.; Yan, J. Efficient estimation for sensor biases and target states in the presence of sensor position errors. *IEEE Sens. J.* **2024**, *24*, 16551–16562. <https://doi.org/10.1109/JSEN.2024.3384810>.
- Bu, S.; Zhou, G. Joint data association, spatiotemporal bias compensation and fusion for multisensor multitarget tracking. *IEEE Trans. Signal Process.* **2023**, *71*, 1509–1523. <https://doi.org/10.1109/TSP.2023.3268472>.
- Pu, W.; Liu, Y.-F.; Luo, Z.-Q. Efficient estimation of sensor biases for the 3-D asynchronous multi-sensor system. *IEEE Trans. Signal Process.* **2023**, *71*, 2420–2433. <https://doi.org/10.1109/TSP.2023.3289706>.
- Fortunati, S.; Farina, A.; Gini, F.; Graziano, A.; Greco, M.S.; Giompapa, S. Least squares estimation and Cramér–Rao type lower bounds for relative sensor registration process. *IEEE Trans. Signal Process.* **2011**, *59*, 1075–1087. <https://doi.org/10.1109/TSP.2010.2097258>.
- Gao, L.; Battistelli, G.; Chisci, L.; Wei, P. Distributed joint sensor registration and multitarget tracking via sensor network. *IEEE Trans. Aerosp. Electron. Syst.* **2020**, *56*, 1301–1317. <https://doi.org/10.1109/TAES.2019.2929972>.
- Pu, W.; Liu, Y.-F.; Yan, J.; Liu, H.; Luo, Z.-Q. Optimal estimation of sensor biases for asynchronous multi-sensor data fusion. *Math. Program.* **2018**, *170*, 357–386. <https://doi.org/10.1007/s10107-018-1304-2>.
- Qiu, Z.; Martinez-Sanchez, J.; Arias-Sanchez, P.; Rashdi, R. External multi-modal imaging sensor calibration for sensor fusion: A review. *Inf. Fusion* **2023**, *97*, 101806. <https://doi.org/10.1016/j.inffus.2023.101806>.

11. Liu, X.; Deng, Z.; Zhang, G. Targetless radar-camera extrinsic parameter calibration using track-to-track association. *Sensors* **2025**, *25*, 949. <https://doi.org/10.3390/s25030949>.
12. Durmaz, O.; Cevikalp, H. Targetless radar-camera calibration via trajectory alignment. *Sensors* **2025**, *25*, 7574. <https://doi.org/10.3390/s25247574>.
13. Lyu, X.; Liu, S.; Qiao, R.; Jiang, S.; Wang, Y. Camera, LiDAR, and IMU spatiotemporal calibration: Methodological review and research perspectives. *Sensors* **2025**, *25*, 5409. <https://doi.org/10.3390/s25175409>.
14. Ma, Y.; Dong, G.; Zhao, C.; Basu, A.; Wu, Z. Background subtraction based on principal motion for a freely moving camera. In *Smart Multimedia*; McDaniel, T., Berretti, S., Curcio, I., Basu, A., Eds.; Lecture Notes in Computer Science; Springer: Cham, Switzerland, 2020; Volume 12015, pp. 67–78. https://doi.org/10.1007/978-3-030-54407-2_6.
15. Wang, Q.; Li, Z.; Peng, J.; Lu, K. Bio-inspired observability enhancement method for UAV target localization and sensor bias estimation with bearing-only measurement. *Biomimetics* **2025**, *10*, 336. <https://doi.org/10.3390/biomimetics10050336>.
16. Nardone, S.C.; Aidala, V.J. Observability criteria for bearings-only target motion analysis. *IEEE Trans. Aerosp. Electron. Syst.* **1981**, *AES-17*, 162–166. <https://doi.org/10.1109/TAES.1981.309141>.
17. Peng, J.; Wang, Q.; Jin, B.; Zhang, Y.; Lu, K. Trajectory optimization to enhance observability for bearing-only target localization and sensor bias calibration. *Biomimetics* **2024**, *9*, 510. <https://doi.org/10.3390/biomimetics9090510>.
18. Jiang, H.; Zhang, Y.; Wang, X.; Cai, Y. Observability-based Gaussian sum cubature Kalman filter for three-dimensional target tracking using a single two-dimensional radar. *Remote Sens.* **2024**, *16*, 4172. <https://doi.org/10.3390/rs16224172>.
19. Sun, T.; Xin, M. Bearings-only tracking using augmented ensemble Kalman filter. *IEEE Trans. Control Syst. Technol.* **2020**, *28*, 1009–1016. <https://doi.org/10.1109/TCST.2018.2890370>.
20. Qian, M.; Chen, W.; Sun, R. A maneuvering target tracking algorithm based on cooperative localization of multi-UAVs with bearing-only measurements. *IEEE Trans. Instrum. Meas.* **2024**, *73*, 1–11. <https://doi.org/10.1109/TIM.2024.3382741>.
21. Huang, J.; Hu, H.; Kang, L. Time convolutional network-based maneuvering target tracking with azimuth–Doppler measurement. *Sensors* **2024**, *24*, 263. <https://doi.org/10.3390/s24010263>.
22. Zhou, Z.; Hu, J.; Chen, B.; Shen, X.; Meng, B. Target tracking and circumnavigation control for multi-unmanned aerial vehicle systems using bearing measurements. *Actuators* **2024**, *13*, 323. <https://doi.org/10.3390/act13090323>.
23. Hermann, R.; Krener, A.J. Nonlinear controllability and observability. *IEEE Trans. Autom. Control* **1977**, *22*, 728–740. <https://doi.org/10.1109/TAC.1977.1101601>.
24. Massiani, P.-F.; Buisson-Fenet, M.; Solowjow, F.; Di Meglio, F.; Trimpe, S. Data-driven observability analysis for nonlinear stochastic systems. *IEEE Trans. Autom. Control* **2024**, *69*, 4042–4049. <https://doi.org/10.1109/TAC.2023.3346812>.
25. Tichavsky, P.; Muravchik, C.H.; Nehorai, A. Posterior Cramer–Rao bounds for discrete-time nonlinear filtering. *IEEE Trans. Signal Process.* **1998**, *46*, 1386–1396. <https://doi.org/10.1109/78.668800>.
26. Jia, T.; Liu, H.; Wang, P.; Wang, R.; Gao, C. Sensor error calibration and optimal geometry analysis of calibrators. *Signal Process.* **2024**, *214*, 109249. <https://doi.org/10.1016/j.sigpro.2023.109249>.
27. Schiller, J.D.; Muntwiler, S.; Köhler, J.; Zeilinger, M.N.; Müller, M.A. A Lyapunov function for robust stability of moving horizon estimation. *IEEE Trans. Autom. Control* **2023**, *68*, 7466–7481. <https://doi.org/10.1109/TAC.2023.3280344>.
28. Schiller, J.D.; Müller, M.A. Nonlinear moving horizon estimation for robust state and parameter estimation. *Automatica* **2026**, *185*, 112790. <https://doi.org/10.1016/j.automatica.2025.112790>.
29. Dellaert, F.; Kaess, M. Factor graphs for robot perception. *Found. Trends Robot.* **2017**, *6*, 1–139. <https://doi.org/10.1561/23000000043>.
30. Cao, Y.; Beltrame, G. VIR-SLAM: Visual, inertial, and ranging SLAM for single and multi-robot systems. *Auton. Robots* **2021**, *45*, 905–917. <https://doi.org/10.1007/s10514-021-09992-7>.
31. Wen, W.; Pfeifer, T.; Bai, X.; Hsu, L.-T. Factor graph optimization for GNSS/INS integration: A comparison with the extended Kalman filter. *Navigation* **2021**, *68*, 315–331. <https://doi.org/10.1002/navi.421>.
32. Chen, Z.; Fang, Y.; Zhang, R.; Fu, W. Factor graph method for target state estimation in bearing-only sensor network. *J. Syst. Eng. Electron.* **2025**, *36*, 380–396. <https://doi.org/10.23919/JSEE.2024.000122>.
33. Nguyen, T.H.; Xie, L. Relative transformation estimation based on fusion of odometry and UWB ranging data. *IEEE Trans. Robot.* **2023**, *39*, 2861–2877. <https://doi.org/10.1109/TRO.2023.3264946>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.