

Guardrails for Learning-Based Navigation Filters

Preprint - Under Review

Barak Or

Abstract—Learning-augmented navigation filters introduce data-driven process-noise adaptation into recursive state estimation, yet erroneous learned parameters can compromise measurement acceptance and uncertainty reporting.

This work investigates runtime guardrails that constrain learned process-noise proposals in a hybrid inertial navigation system with ultra-wideband aiding (INS/UWB). Local analysis characterizes how variations in process-noise parameters influence innovation-based measurement gating, while experiments examine magnitude, rate, and projection restrictions within an error-state extended Kalman filter. Across eight held-out quadrotor recordings, magnitude restrictions suppress two severe Ridge-based failures, reducing recording-level position root-mean-square error (RMSE) from over 650 m to below 0.5 m. Under a substantial upward perturbation of the learned process-noise proposals, magnitude restriction reduces mean position RMSE by 11.9%, with improvement across all eight recordings. Channel-specific tests identify effective angular-rate process-noise bounds, while discrete projection reveals that lower error may result from suppressed adaptation. However, improved accuracy does not necessarily imply calibrated uncertainty or recovery following estimator divergence. These findings characterize the benefits and limitations of runtime guardrails for learning-based navigation estimation.

Index Terms—Error-state extended Kalman filter, hybrid INS/UWB navigation, learned process noise, physical AI, estimation-level guardrails, guardrail, uncertainty quantification.

I. INTRODUCTION

LEARNED process-noise adaptation inserts a data-driven decision into a model-based measurement chain. In a hybrid inertial navigation system with ultra-wideband aiding (INS/UWB), an error-state extended Kalman filter (es-EKF) combines inertial propagation and radio updates with a learned noise updater. Process covariance affects predicted uncertainty, Kalman gain, innovation normalization, and measurement acceptance [1], [2]. An erroneous noise proposal can therefore corrupt the reported position even when numerical checks pass. Learning suitable noise parameters and controlling their influence on the estimator are distinct engineering problems. We study estimation-level runtime guardrails that constrain learned process-noise proposals before they affect the navigation filter. The underlying Kalman filter remains model-based.

Fig. 1 summarizes the estimator-level intervention; the detailed filter architecture is shown in Fig. 2.

Learning-augmented Kalman filtering combines state-space models with data-driven components to address incomplete dynamics and uncertain sensing. KalmanNet [3] integrates a recurrent neural network into the filtering recursion to learn the




Fig. 1. Estimation-level runtime guardrail: learned process noise is constrained before filter propagation.

correction gain under partial model knowledge; differentiable filters [4]–[6] permit learned representations and end-to-end estimation. A recent review [7] systematizes such AI-aided Kalman architectures. In navigation, learned inertial odometry [8], [9] and hybrid model-learning filters [10] illustrate complementary uses of data-driven estimation. Adaptive noise models [11]–[14] address uncertain process statistics, while learned trajectory uncertainty [15] addresses predictive uncertainty.

These advances motivate a distinct reliability question. The learned mapping is data-dependent, and its behavior under erroneous predictions or operating conditions not represented during training need not be constrained by the surrounding model-based recursion. For example, an inappropriate learned gain can alter a state correction, whereas an inappropriate process-noise proposal can change covariance propagation, subsequent gains, and innovation gating. Either may influence later estimates through feedback. Thus, preserving Kalman-filter structure and reporting improved accuracy do not alone establish bounded influence or reliable uncertainty reporting.

Classical innovation-, correlation-, likelihood-, and Bayesian-based covariance adaptation [16]–[23] addresses noise estimation. Other navigation interfaces include learned integration schedules [24], [25], inertial speed and alignment estimation [26]–[28], vibration-based position inference [29], and hybrid INS/DVL fusion [30]. Proper scoring [31], distribution-shift evaluation [32], conformal prediction [33], and covariance-consistency studies [34], [35] address complementary aspects of uncertainty assessment; metrological terminology and uncertainty-evaluation guidance are provided in [36], [37]. Recent measurement-oriented studies examine degraded UWB/inertial aiding [38], [39], adaptive covariance estimation [40]–[43], robust state estimation under outliers [44], [45], and uncertainty credibility [46]. These studies motivate evaluating accuracy and covariance calibration separately, but their methods are not equivalent to explicitly restricting learned process-noise outputs at runtime.

Earlier hybrid-navigation work [10], [14] demonstrates learned process-noise adaptation, but improved navigation accuracy alone does not establish how reliably a learned

component interacts with measurement acceptance and reported uncertainty. The measurement-oriented contribution of this study is a controlled characterization of *estimation-level runtime guardrails*: restrictions applied to learned process-noise proposals before covariance propagation. We examine their effects on innovation-based measurement gating, position-error behavior, covariance coverage, and continued adaptation, and identify limits under aiding faults and short-horizon recovery tests. The nominal dynamics and observation model remain unchanged; the analysis does not claim calibrated uncertainty or physical-safety guarantees. Conventional measurement gates serve different purposes, as do robot safety filters [47]–[51]; action-level guardrails [52]–[56] do not confer physical-safety guarantees on the estimation restrictions studied here.

Recent Physical AI studies expose related reliability limitations in learned robotic policies. VLA-Arena [57] identifies generalization and safety-constraint failures across systematically varied vision-language-action (VLA) tasks, while SafeLab [58] distinguishes successful task completion from safe execution. World-model-based VLA training [59] further addresses error accumulation in imagined rollouts. These studies concern policy evaluation, action-level safety, or predictive rollouts rather than recursive navigation estimation. They motivate complementary scrutiny of learned components at the estimator interface, without implying that the guardrails studied here certify physical safety.

The main contributions of this work are as follows:

- (i) We compare guarded and unguarded navigation filters with fixed-noise and innovation-adaptive baselines under nominal and perturbed conditions, documenting severe-error mitigation in Ridge and reduced perturbation-induced degradation in ExtraTrees alongside trade-offs in adaptation and covariance consistency.
- (ii) We characterize how learned process-noise adaptation influences innovation-based measurement gating through a local analytical relationship between applied noise parameters and measurement acceptance.
- (iii) We systematically examine runtime estimation guardrails through channel-specific restrictions, adaptation loss under discrete projection, and recovery tests.

The rest of this paper is organized as follows. Sections II and III describe the hybrid INS/UWB es-EKF and its process-noise restrictions. Sections IV and V present the evaluation protocol and results. Sections VI and VII discuss the findings and conclude the work.

II. HYBRID INS/UWB ES-EKF

A. Experimental Platform and Reference Measurements

Fig. 2 details the learned updater and guardrail in the hybrid INS/UWB es-EKF. The updater changes process-noise parameters before covariance propagation; the nominal INS dynamics and UWB observation equations remain model-based.

The measurand is the vehicle position at its body reference point, expressed in the dataset-aligned optical world frame. The coordinates use the dataset reference frame, rather than an assumed geographic navigation frame. We evaluate the method using publicly available indoor quadrotor recordings containing

inertial measurements, UWB observations, and an optical position reference. The public Ultra-wideband Time-difference-of-arrival Indoor Localization (UTIL) dataset [60] uses a customized Crazyflie Bolt-based quadrotor, a vertical DWM1000-based Loco Positioning tag, and eight anchors per constellation. Ten Vicon Vantage+ cameras cover a $7.0 \times 8.0 \times 3.5$ m arena. The Leica-surveyed anchor positions are aligned to the optical frame using six markers. The reported alignment root-mean-square error (RMSE) is 1.12 mm, and communication latency is approximately 10 ms. Events and received reference poses share the onboard microsecond clock [60], [61]. We replay the inertial measurement unit (IMU) and UWB TDOA2 streams offline. Fig. 3 shows a recorded reference trajectory and its direction.

With p^{ref} denoting the recorded optical position, position errors $\hat{p} - p^{\text{ref}}$, RMSE, and covariance coverage are evaluated relative to this reference.

B. Filter Framework

The nominal state is $x = (p, v, R, b_a, b_g)$: $p, v \in \mathbb{R}^3$ are position (m) and velocity (m s^{-1}), $R \in \text{SO}(3)$ is the body-to-world rotation, and $b_a, b_g \in \mathbb{R}^3$ are accelerometer (m s^{-2}) and gyroscope (rad s^{-1}) biases. Hats denote estimates; superscripts $-$ and $+$ denote prior and corrected quantities. The filter index is k ; i indexes learned-parameter updates. $\mathbf{I}_d$ is the d -dimensional identity, and $[u]_{\times} w = u \times w$. Unsubscripted norms are Euclidean; $\|\cdot\|_{\infty}$ is the maximum absolute component for vectors and maximum absolute row sum for matrices. The right-multiplicative error state is [62]

$$\delta x = [\delta p^{\top}, \delta v^{\top}, \delta \phi^{\top}, \delta b_a^{\top}, \delta b_g^{\top}]^{\top} \in \mathbb{R}^{15}. \quad (1)$$

Here $\delta \phi$ is the local rotation error, defined by $R_{\text{true}} = \hat{R} \exp([\delta \phi]_{\times})$. With measured specific force f_m and angular rate ω_m , define $f = f_m - b_a$ and $\omega = \omega_m - b_g$. The nominal dynamics are $\dot{p} = v$, $\dot{v} = Rf + g$, and $\dot{R} = R[\omega]_{\times}$, where $g = (0, 0, -9.81)^{\top} \text{ m s}^{-2}$. Nominal biases are assumed to be constant during propagation; their random walks enter the covariance.

The learner adjusts two dimensionless base-10 log-density multipliers $\lambda = (\lambda_a, \lambda_{\omega})^{\top}$:

$$Q_c(\lambda) = \text{blkdiag}(0.04 \cdot 10^{\lambda_a} \mathbf{I}_3, 10^{-4} \cdot 10^{\lambda_{\omega}} \mathbf{I}_3, 10^{-6} \mathbf{I}_3, 10^{-8} \mathbf{I}_3). \quad (2)$$

Here blkdiag assembles 3×3 blocks. They describe acceleration noise, angular-rate noise, and their bias random walks, with units $\text{m}^2 \text{ s}^{-3}$, $\text{rad}^2 \text{ s}^{-1}$, $\text{m}^2 \text{ s}^{-5}$, and $\text{rad}^2 \text{ s}^{-3}$, respectively. Only the first two vary. Let $P_k \in \mathbb{R}^{15 \times 15}$ be the error-state covariance and $P_{pp,k} \in \mathbb{R}^{3 \times 3}$ its position block. For integration-substep duration $\delta \tau_k$ (s), $F_k \in \mathbb{R}^{15 \times 15}$ is the frozen error-dynamics matrix, $G_k \in \mathbb{R}^{15 \times 12}$ maps driving noise to state error, and $\Phi_k = \exp(F_k \delta \tau_k)$ is the state transition. Then

$$P_{k+1}^{-} = \Phi_k P_k^{+} \Phi_k^{\top} + Q_{d,k}, \quad (3)$$

where [63]

$$Q_{d,k} = \int_0^{\delta \tau_k} e^{F_k \tau} G_k Q_c(\lambda) G_k^{\top} e^{F_k^{\top} \tau} d\tau. \quad (4)$$

Here $Q_{d,k}$ is the discrete process covariance, τ is the integration variable, and $Q_c \in \mathbb{R}^{12 \times 12}$ is the continuous driving-noise




Fig. 2. Hybrid INS/UWB es-EKF with a learned process-noise updater (taupe) and runtime guardrail (mint). The guardrail constrains the family-limited proposal before covariance propagation; nominal INS dynamics and NIS-gated UWB corrections remain model-based. Applied noise parameters are updated every 0.2 s and held between updates. Dashed arrows indicate causal feedback.




Fig. 3. Recorded optical reference trajectory (C3-T4).

spectral-density matrix. Appendix A specifies the matrices and discretization.

For anchor positions $a_a, a_b \in \mathbb{R}^3$, the tag position is $r = p + R\ell$, where $\ell \in \mathbb{R}^3$ is the fixed body-to-tag lever arm. The TDOA timestamps are represented as a scalar range-difference observation in meters. The scalar observation is $z_k = h_{ab}(x_k) + \epsilon_k$ (m), with zero-mean noise ϵ_k and

$$h_{ab}(x) = \|r - a_b\| - \|r - a_a\|. \quad (5)$$

Observation variance $\sigma_z^2 = 0.05 \text{ m}^2$ is a fixed estimator setting. Defining $u = (r - a_b)/\|r - a_b\| - (r - a_a)/\|r - a_a\|$ gives the 1×15 observation Jacobian $H_k = \partial h_{ab}/\partial \delta x$, distinguishing the scalar observation function h_{ab} from its linearization

$$H_k = [u^\top \quad 0 \quad -u^\top \hat{R}_k^{-1}[\ell]_\times \quad 0 \quad 0], \quad (6)$$

where zero blocks have three columns. Before gating, the innovation, its predicted variance, and normalized innovation squared (NIS) are respectively

$$\nu_k = z_k - h_{ab}(\hat{x}_k^-), \quad (7)$$

$$S_k = H_k P_k^- H_k^\top + \sigma_z^2, \quad (8)$$

$$\text{NIS}_k = \nu_k^2 / S_k. \quad (9)$$

Let $\chi_d^2(\beta)$ denote the β -quantile of the chi-square distribution with d degrees of freedom, $0 < \beta < 1$. The observation is accepted when $\text{NIS}_k \leq \chi_1^2(0.999) \simeq 10.828$, the nominal threshold for the scalar UWB residual. For an accepted observation, the 15×1 Kalman gain and error correction

are

$$K_k = P_k^- H_k^\top S_k^{-1}, \quad (10)$$

$$\hat{\delta x}_k = K_k \nu_k. \quad (11)$$

Position, velocity, and biases are corrected additively using the Kalman update employed in hybrid navigation [10]. Attitude correction and covariance reset follow the right-error formulation [62], as detailed in Appendix A.

C. Analytical Characterization of Learned-Noise Influence on Measurement Gating

Consider one covariance-propagation substep of duration $\delta\tau = \delta\tau_k$, holding the nominal prediction and $F = F_k$, $G = G_k$, $\Phi = \exp(F\delta\tau)$ fixed. Linearity of (4) in the noise density gives

$$Q_d(\lambda) = 10^{\lambda_a} B_a + 10^{\lambda_\omega} B_\omega + B_b, \quad (12)$$

Here $B_a, B_\omega \in \mathbb{R}^{15 \times 15}$ are obtained from (4) using only the corresponding baseline block of (2), at zero log multiplier. B_b contains fixed bias noise. These contributions are positive semidefinite, so $B_c \succeq 0$ and $H B_c H^\top \geq 0$ for $c \in \{a, \omega\}$.

Also hold P^+ , innovation ν , observation Jacobian H , measurement variance $\sigma_z^2 > 0$, and the other log-density coordinate fixed. Substituting (3) into (8) gives the local innovation variance

$$S = H(\Phi P^+ \Phi^\top + Q_d(\lambda)) H^\top + \sigma_z^2 > 0. \quad (13)$$

Since $\partial 10^{\lambda_c} / \partial \lambda_c = \ln(10) 10^{\lambda_c}$,

$$\frac{\partial S}{\partial \lambda_c} = \ln(10) 10^{\lambda_c} H B_c H^\top \geq 0. \quad (14)$$

With ν^2 fixed, differentiating $\text{NIS} = \nu^2 / S$ contributes the factor $-\nu^2 / S^2$, giving

$$\frac{\partial (\nu^2 / S)}{\partial \lambda_c} = -\frac{\nu^2 \ln(10) 10^{\lambda_c}}{S^2} H B_c H^\top \leq 0. \quad (15)$$

The derivative is strictly negative when $\nu \neq 0$ and $H B_c H^\top > 0$, and zero otherwise. Increasing either applied noise intensity therefore increases or preserves the denominator S , decreasing or preserving the NIS for the unchanged residual. A rejected UWB observation may then pass the fixed gate solely because

TABLE I
ORDERED CAUSAL INPUT FEATURES ($\psi_i \in \mathbb{R}^{54}$)

Indices	Feature block, in implementation order	Count
1–24	Four statistics per measured IMU channel	24
25–39	$\log \max(P_{\ell\ell}, 10^{-30})$, $\ell = 1, \dots, 15$	15
40–42	Body-frame velocity $\hat{R}^\top \hat{v}$	3
43–48	Bias estimates $\hat{b}_a$, then $\hat{b}_g$	6
49–50	Previous applied $\lambda_{a,i-1}, \lambda_{\omega,i-1}$	2
51–54	History duration W_i , requested hold T_i , feature interval h_f , preceding event interval δt_{evt} (s)	4

Vector blocks use x, y, z order; covariance follows (1) in fixed SI/radian coordinates, before the update. $T_i = 0.2$ s; $h_f = 0.01$ s.

predicted uncertainty increased, not because its residual decreased. Passing the gate assumes valid inputs and does not establish that the measurement is correct.

Equation (15) is local to the fixed innovation and linearization. In the complete recursion, covariance and state updates alter subsequent innovations, gains, and learned proposals; the relationship does not guarantee improved accuracy, monotonic NIS evolution, or filter stability.

III. LEARNED PROPOSALS AND GUARDRAILS

A. Supervision and Controls

At each of 482 training checkpoints per objective, a candidate λ is evaluated over a 0.5 s horizon. Let n be the number of supported scored samples in that horizon and $e_k(\lambda) = \hat{p}_k(\lambda) - p_k^{\text{ref}}$, expressed in meters. The mean-squared-error (MSE) objective is

$$J_{\text{MSE}}(\lambda) = \frac{1}{n} \sum_{k=1}^n e_k(\lambda)^\top e_k(\lambda), \quad (16)$$

For a zero-mean Gaussian position-error model with symmetric positive-definite $P_{pp,k}$, the mean negative log-likelihood (NLL) is

$$J_{\text{NLL}}(\lambda) = \frac{1}{2n} \sum_{k=1}^n [e_k^\top P_{pp,k}^{-1} e_k + \log \det P_{pp,k} + 3 \log(2\pi)]. \quad (17)$$

Both e_k and $P_{pp,k}$ depend on λ ; $\log$ is natural, with densities expressed in fixed meter coordinates. Invalid Cholesky pivots give infinite NLL. This scoring model does not assert Gaussian or calibrated replay errors. Appendix B gives the search settings and recording roles.

Table I defines the 54 causal features. Each of the six IMU channels contributes detrended root-mean-square (RMS), first-difference RMS, median absolute first difference, and mean; remaining features summarize the filter state and update context. Detailed sampling conventions appear in Appendix B.

Ridge uses training-fitted standardization and regularization parameter 1; ExtraTrees uses unstandardized features. Invalid or insufficient causal history suppresses learned updates, and no optical reference enters runtime inputs. Implementation details are given in Appendix B and the reproduction package.

Twelve model–history–objective configurations were considered using Ridge [64], ExtraTrees [65], and a two-hidden-layer perceptron. Validation selected MSE-trained Ridge (0.5 s history) and NLL-trained ExtraTrees (1 s history; 16 trees,

depth six, minimum leaf size four). Models remain frozen during evaluation and predict noise parameters, not states.

At model update i , the frozen predictor $f_\theta : \mathbb{R}^{54} \rightarrow \mathbb{R}^2$ returns the raw proposal $\lambda_i^{\text{raw}} = f_\theta(\psi_i)$. Every learned branch enforces the common family $\mathcal{D} = [-5, 5]^2$:

$$\bar{\lambda}_i = \Pi_{\mathcal{D}} f_\theta(\psi_i). \quad (18)$$

For a closed set $\mathcal{S}$, $\Pi_{\mathcal{S}}(x) \in \arg \min_{\xi \in \mathcal{S}} \|\xi - x\|_2$ denotes nearest-point projection; fixed enumeration resolves nonunique minima. $\bar{\lambda}_i$ is the family-limited proposal; λ_i is the value actually applied. Updates occur every 0.2 s once the history is available; $\lambda = (0, 0)$ before eligibility and is held between updates. *OFF* means no restriction beyond (18); the observation gate remains active. Each ON branch replays the complete filter, so altered states can change subsequent proposals.

Three nonlearned controls distinguish offline tuning from adaptation. The original pair $(0, -3)$ is selected from nine validation candidates. The development-tuned pair $(-3, -1)$ minimizes equal-recording mean MSE on four recording grids. A matched-startup control switches from zero to $(-2, -2)$ at the first eligible MSE update.

B. Magnitude, Rate, and Projection Guardrails

Here, a runtime process-noise guardrail is a deterministic constraint applied to learned noise-parameter proposals before covariance propagation. It limits the learner’s influence on the estimator but does not enforce physical-state safety constraints or guarantee estimation accuracy. All learned configurations, including OFF, enforce $\mathcal{D}$. The *magnitude guardrail* (MAG) adds the narrower box $\mathcal{M} = [-3, 3]^2$, enabling matched comparisons with otherwise identical learned-filter configurations:

$$\lambda_i^{\text{mag}} = \Pi_{\mathcal{M}}(\bar{\lambda}_i). \quad (19)$$

Let r_λ denote the maximum permitted change per second in either base-10 log-density coordinate, and Δt_i the elapsed time since the previous update. The fixed reference setting is $r_\lambda = 4$ log units s^{-1} . The rate-feasible box is

$$\mathcal{B}_i = \{\lambda : \|\lambda - \lambda_{i-1}\|_\infty \leq r_\lambda \Delta t_i\}. \quad (20)$$

The *rate guardrail* (RATE) applies

$$\lambda_i^{\text{rate}} = \Pi_{\mathcal{B}_i}(\bar{\lambda}_i), \quad (21)$$

and the *combined magnitude-rate guardrail* (BOTH) applies magnitude projection before rate projection:

$$\lambda_i^{\text{both}} = \Pi_{\mathcal{B}_i} \Pi_{\mathcal{M}}(\bar{\lambda}_i). \quad (22)$$

Each branch uses its own previous applied value. The magnitude box bounds acceleration and angular-rate densities to $[4 \times 10^{-5}, 40] \text{ m}^2 \text{ s}^{-3}$ and $[10^{-7}, 0.1] \text{ rad}^2 \text{ s}^{-1}$. At the 0.2 s update interval, the evaluated rate limit permits at most 0.8 log units per coordinate, corresponding to a multiplicative density change of up to $10^{0.8} \simeq 6.31$.

One-sided guardrails apply only the lower -3 or upper 3 bound, to both coordinates or either coordinate alone. *Channel-specific guardrails* bound both sides of one coordinate; others retain $\mathcal{D}$. These guardrails restrict the process-noise intensities used by the Kalman filter to propagate estimation uncertainty. They modify the estimator’s covariance model, not sensor measurements, nominal vehicle-dynamics equations, or control commands. Development-only sensitivity tests use the boxes

$[-2, 2]^2$ and $[-4, 4]^2$, and rates $r_\lambda \in \{1, 8\}$ log units s^{-1} . The confirmation settings are not retuned.

The *discrete-map guardrail* uses a set $\mathcal{G} \subset \mathcal{D}$ of 106 development-supported unit-grid points. For the rate-limited proposal $s_i = \Pi_{\mathcal{B}_i}(\bar{\lambda}_i)$ and grid candidate ξ , it applies

$$\lambda_i^{\text{disc}} = \begin{cases} \arg \min_{\xi \in \mathcal{G} \cap \mathcal{B}_i} \|\xi - s_i\|^2, & \mathcal{G} \cap \mathcal{B}_i \neq \emptyset, \\ \lambda_{i-1}, & \text{otherwise.} \end{cases} \quad (23)$$

Fixed enumeration resolves ties, using a feasibility tolerance of 10^{-12} . The *continuous-map guardrail* includes every closed unit rectangle whose four corner nodes belong to $\mathcal{G}$. Enumerating adjacent grid nodes gives 85 such rectangles, denoted by $\mathcal{C}_q$; this count follows from the 106 retained nodes rather than a separately selected parameter (Fig. 4). The feasible union is

$$\mathcal{U}_i = \bigcup_{q: \lambda_{i-1} \in \mathcal{C}_q} (\mathcal{C}_q \cap \mathcal{B}_i \cap \mathcal{D}). \quad (24)$$

It projects s_i onto $\mathcal{U}_i$ when nonempty; otherwise it retains λ_{i-1} and marks the update unsupported. Restricting rectangles to those containing the previous value prevents jumps between disconnected components. Neither map guardrail adds the narrower magnitude box.

Let $d_{\min}$ be the minimum nonzero infinity-norm distance between grid points. Once $\lambda_{i-1} \in \mathcal{G}$,

$$r_\lambda \Delta t_i < d_{\min} \implies \mathcal{G} \cap \mathcal{B}_i = \{\lambda_{i-1}\}. \quad (25)$$

All other grid nodes lie outside the permitted step. Here $d_{\min} = 1$ and the regular step is 0.8; a longer initial interval permits entry before locking. Reaching the spacing removes this obstruction but does not guarantee motion.

Algorithm 1 summarizes the event-driven integration. The guardrail policy $\mathcal{R}_i$ denotes OFF, MAG, RATE, BOTH, or the map projection defined in (18)–(24).

Algorithm 1. Runtime Guardrail Integration

Input: time-ordered IMU/UWB events, f_θ , $\mathcal{R}_i$, $(\hat{x}_0, P_0)$; $\lambda_0 = (0, 0)$.

```

1 for each event block do
2   if update due and  $\psi_i$  valid:  $\bar{\lambda}_i \leftarrow \Pi_{\mathcal{D}} f_\theta(\psi_i)$  (18);
3    $\lambda_i \leftarrow \mathcal{R}_i(\bar{\lambda}_i, \lambda_{i-1})$  (19)–(24);
4   else:  $\lambda_i \leftarrow \lambda_{i-1}$ ;
5   Hold  $\lambda_i$ ; propagate  $(\hat{x}^-, P^-)$  via  $Q_c(\lambda_i)$  (2)–(4);
6   At UWB events, compute  $(\nu_k, S_k, \text{NIS}_k)$  (7)–(9);
7   Apply gated correction (10)–(11), or retain the prior;
8   Retain  $(\hat{x}^+, P^+, \lambda_i)$  for causal feedback.
```

Output: $\hat{x}_k^+$, P_k^+ , λ_i and intervention/gate records.

C. Innovation-Adaptive Comparator

The completion study includes *innovation-based adaptive estimation (IAE)-projected-Q*, an explicit adaptation of the innovation covariance-matching construction discussed in [10]. Let $\mathcal{V}_i$ contain finite pre-gate innovation indices in the trailing window, including rejected observations, and K_i^{last} be the most recent pre-gate gain. The target is

$$\hat{Q}_{d,i} = K_i^{\text{last}} \left(\frac{1}{|\mathcal{V}_i|} \sum_{k \in \mathcal{V}_i} \nu_k^2 \right) (K_i^{\text{last}})^\top. \quad (26)$$

The most recently available pre-gate Kalman gain K_i^{last} is held fixed across the innovation window $\mathcal{V}_i$, rather than using a separate historical gain for each residual. Using the current state and inertial input, the bases in (12) are integrated over the median positive radio interval in $\mathcal{V}_i$. With P_0 denoting the

initial error-state covariance and $D_0 = \text{diag}(\sqrt{\text{diag } P_0})$, define $\|X\|_{D_0} = \|D_0^{-1} X D_0^{-1}\|_F$, where $\|\cdot\|_F$ is the Frobenius norm. Two positive density-scale factors $\kappa = (\kappa_a, \kappa_\omega)^\top$ are fitted by $\hat{\kappa}_i = \arg \min_{\kappa \in [10^{-5}, 10^5]^2} \|\kappa_a B_a + \kappa_\omega B_\omega + B_b - \hat{Q}_{d,i}\|_{D_0}^2$. (27)

The applied update is $\Pi_{\mathcal{D}}[0.8\lambda_{i-1} + 0.2 \log_{10} \hat{\kappa}_i]$, initialized at $(-3, -1)$. An update is attempted every 0.2 s when the window contains at least five finite innovations, a positive radio interval, and a gain recorded within that window; otherwise the previous parameters are held. Development selection chooses a 0.5 s window from $\{0.5, 1, 2\}$ s.

IV. EXPERIMENTAL METHODOLOGY

A. Data Partition and Evidence Units

The dataset comprises 26 distinct constellation/trial recording groups (18 non-holdout and eight held-out), with one TDOA2 variant per group. The holdout evaluation spans 19,735 evaluated position samples (984.850 s) and 454,012 UWB arrivals. The original completion campaign comprises 552 unique full replays and 816 matched-state continuations; the complete computational inventory appears in Table VI of Appendix B. For the original holdout campaign, predictors, guardrail settings, perturbation scenarios, and data partitions were fixed before evaluation; only the IAE window was selected on four development groups. All recordings remain within UTIL and the same sensor platform.

A 2 s sensor-only prefix initializes the vehicle reference-point position by robust radio multilateration with the tag lever arm and initializes tilt from gravity; velocity and biases start at zero. Appendix B specifies the initialization parameters. Original sensor events are retained with causal inertial holds. Scoring uses a 0.05 s grid, excludes reference-interpolation gaps over 0.05 s, and integrates only consecutive valid endpoints at most 0.075 s apart. All methods share the same scoring support; the confirmation data contain 4.950 s of excluded reference intervals.

B. Experimental Conditions and Perturbation Scenarios

The eight primary methods are Fixed- Q (dev-tuned), IAE-projected- Q , Ridge-OFF, Ridge-MAG, and ExtraTrees (ET) with OFF, MAG, RATE, or BOTH. Each branch recomputes its own features and complete recursion, not a clipped OFF trace. One-sided interventions use four non-holdout diagnostic groups; sensitivity uses four development groups. Both compare against the same campaign's baselines.

Each confirmation group is replayed under nominal conditions and five perturbations, each active from 40% to 70% of the recording duration. Learned-low and learned-high add -2 and $+2$ log units, respectively, to both raw coordinates before family clipping. Learned-osc adds a sinusoid with amplitude 2 log units and period 0.4 s. UWB-bias adds 1.5 m to each radio observation, while UWB-gap removes the radio observations and retains inertial propagation. Scores cover the full recording, including the intervals before and after the perturbation. Nonlearned baselines are reused for learned-output perturbations because they are unaffected. Confirmation results are not used to retune thresholds.




Fig. 4. Both panels use the same local map, previous value $\lambda_{i-1} = (-2, -2)$, example proposal $\bar{\lambda}_i = (2, 2)$, and rate-limited proposal $s_i = (-1.2, -1.2)$. (a) The permitted step $r_\lambda \Delta t_i = 0.8$ is smaller than the unit grid spacing, leaving only the previous grid node feasible. (b) Replacing isolated nodes with the same map's supported cells yields a continuous feasible intersection containing s_i , so this update can move. Gray shows map support, taupe the proposal, and mint the applied update.

A follow-up experiment extends the learned-output perturbations to signed shifts of $-4, -2, 0, +2$, and $+4$ base-10 log units, applied equally to both raw process-noise outputs during 40%–70% of each recording, before family projection and guardrailing. The frozen ExtraTrees predictor is evaluated with OFF, MAG, RATE, and BOTH using unchanged sensor data, guardrail settings, and full-recording scoring. The resulting 160 replays reuse the same eight recordings after the original holdout analysis; they assess perturbation sensitivity rather than independent generalization.

For each matched-state probe, the candidate continuations share the state and covariance saved immediately before an eligible update near 30%, 50%, or 70% of one of four non-holdout diagnostic recordings, separately for Ridge and ET. Each of the 24 states has 34 two-second continuations with fixed noise: 25 grid pairs from $\{-5, -2.5, 0, 2.5, 5\}^2$, OFF, Fixed- Q (dev-tuned), and seven guarded proposals. MSE and NLL are evaluated on common scoring support without further learned updates.

C. Error, Coverage, and Parameter Motion

Let j index recordings, m methods, and $\mathcal{I}_j$ the set of supported scoring-time indices k , with $N_j = |\mathcal{I}_j|$ and $e_{jk}^{(m)} = \hat{p}_{jk}^{(m)} - p_{jk}^{\text{ref}}$. The indicator $\mathbf{1}\{E\}$ equals one if E holds. Recording-level position error is

$$\text{RMSE}_j^{(m)} = \sqrt{\frac{1}{N_j} \sum_{k \in \mathcal{I}_j} \|e_{jk}^{(m)}\|^2}. \quad (28)$$

The paired change is $\Delta_j^{(m)} = \text{RMSE}_j^{(m)} - \text{RMSE}_j^{(\text{off})}$, with negative values favoring the restriction.

Let $\eta > 0$ be an evaluation error threshold (m), $\mathcal{K}_j$ the supported consecutive interval indices, and $\delta t_{jk} = t_{j,k+1} - t_{jk}$ their duration (s). Exceedance time is

$$T_{\eta,j}^{(m)} = \sum_{k \in \mathcal{K}_j} \delta t_{jk} \mathbf{1}\{\|e_{jk}^{(m)}\| > \eta\}. \quad (29)$$

This left-endpoint rule excludes unsupported gaps; $\eta = 0.5$ m is the main reporting limit. These errors are evaluated offline.

Unlike the innovation-based NIS used for online measurement gating, normalized estimation error squared (NEES) evaluates the consistency of the reported position covariance

against the optical reference and is computed only offline. For the position block,

$$\text{NEES}_{jk}^{(m)} = (e_{jk}^{(m)})^\top (P_{pp,jk}^{(m)})^{-1} e_{jk}^{(m)}, \quad (30)$$

computed by the Cholesky factorization $P_{pp} = LL^\top$: solve the lower-triangular system $L\zeta = e$ and evaluate $\|\zeta\|^2$ [66]. The three degrees of freedom correspond to the three position-error coordinates. For a zero-mean Gaussian error with matching covariance, its whitened components are independent standard normals and their squared norm follows χ_3^2 . Empirical inclusion in the corresponding nominal ellipsoid is

$$C_j^{(m)}(\beta) = \frac{1}{N_j} \sum_{k \in \mathcal{I}_j} \mathbf{1}\{\text{NEES}_{jk}^{(m)} \leq \chi_3^2(\beta)\}. \quad (31)$$

The nominal levels are 0.50, 0.68, 0.80, 0.90, 0.95, 0.975, 0.99, and 0.999. The primary tables and coverage plot use $\bar{C}(\beta) = N_{\text{rec}}^{-1} \sum_{j=1}^{N_{\text{rec}}} C_j(\beta)$, where N_{rec} is the number of recordings. Pooled coverage from archived diagnostics is not mixed into these macro-averages. With $\text{eig}_{\max}(P)$ denoting the largest eigenvalue of P , the largest semi-axis quantifies region size:

$$a_{\max,jk}^{(m)}(\beta) = \sqrt{\chi_3^2(\beta) \text{eig}_{\max}(P_{pp,jk}^{(m)})}. \quad (32)$$

The table statistic $\tilde{a}_{95}$ is the median over recordings of each recording's median $a_{\max}(0.95)$. Ellipsoid volume is $(4\pi/3)[\chi_3^2(\beta)]^{3/2} \sqrt{\det P_{pp,jk}^{(m)}}$, computed through log determinants. These are model-based regions evaluated relative to the recorded optical reference. Fig. 5 illustrates why smaller position error after guardrailing need not imply reference inclusion.

For $M_j^{(m)}$ model updates in method m on recording j , actual parameter changes after the first update are counted as

$$N_{\Delta,j}^{(m)} = \sum_{i=2}^{M_j^{(m)}} \mathbf{1}\{\|\lambda_i^{(m)} - \lambda_{i-1}^{(m)}\|_\infty > 10^{-9}\}. \quad (33)$$

Numerical status is separate from error and parameter motion. Rejection fractions divide rejected by attempted observations in one-second bins; empty bins remain missing. Pre-gate NIS is retained before selection. A parameter guardrail intervenes when $\|\lambda_i - \bar{\lambda}_i\|_\infty > 10^{-12}$ at an eligible update. Reported intervention fractions and N_Δ are averaged over recordings. Here N_Δ counts changes in applied parameters (either coordi-




Fig. 5. Schematic position-uncertainty regions with identical axes and a common reference. MAG reduces position error while excluding the reference from its nominal 95% region.

nate), not raw proposals or interventions. Parameter motion is not itself evidence of useful adaptation.

Intervals for paired RMSE changes use 2,000 bootstrap resamples of whole recording pairs and the 2.5th and 97.5th percentiles. Intervals for the follow-up comparisons are descriptive.

D. Reference-Perturbation Sensitivity

To separate large reductions from small reference-sensitive differences, keep the computed estimates fixed and perturb the scored reference by vectors d_{jk} . Let ε_{ref} (m) bound the RMS magnitude of this perturbation over $\mathcal{I}_j$. If $\tilde{\Delta}_j^{(m)}$ is the paired RMSE difference evaluated against $p_{jk}^{\text{ref}} + d_{jk}$, the reverse triangle inequality gives

$$|\tilde{\Delta}_j^{(m)} - \Delta_j^{(m)}| \leq 2\varepsilon_{\text{ref}}. \quad (34)$$

Each individual RMSE changes by at most ε_{ref} ; subtracting two such quantities proves (34). A reduction larger than $2\varepsilon_{\text{ref}}$ therefore retains its sign under this assumed perturbation.

V. RESULTS

A. Nominal Navigation Performance

Table II summarizes the eight held-out recordings. Ridge-MAG suppresses two Ridge-OFF failures on C4-T5 and C4-T6, reducing their RMSEs from 1290.297 and 653.553 m to 0.413 and 0.449 m (Fig. 6). The mean reduction from 243.110 to 0.227 m is dominated by these cases (paired change -242.883 m; descriptive 95% bootstrap interval $[-567.392, -0.006]$ m). Four recordings improve and four deteriorate, the latter by at most 4.336 mm.

MAG reduces Ridge parameter changes from 272.250 to 8.125 per recording, accompanied by lower nominal 95% coverage (62.932% to 18.885%). Fixed- Q (dev-tuned) and IAE-projected- Q attain similarly low nominal RMSEs (0.217 and 0.227 m) but substantially lower coverage than ExtraTrees (Table II). For the already accurate ET-OFF, MAG intervenes on only 0.267% of updates and changes mean RMSE negligibly (-0.007 mm). Thus, nominal guardrailing principally mitigates Ridge's extreme cases rather than yielding a general accuracy gain.

B. Channel-Specific Intervention Effects

Table III isolates one-sided and channel-specific Ridge restrictions on four non-holdout diagnostic recordings. The angular-rate upper bound $\lambda_{\omega} \leq 3$ (a process-noise density ceiling of $0.1 \text{ rad}^2 \text{ s}^{-1}$), alone or with the acceleration upper bound, suppresses all three baseline extreme cases. The acceleration upper bound alone suppresses none, while lower-only restrictions leave two. The tested bound is sufficient for these recordings, not proven optimal or universally necessary.

On development recordings, Ridge boxes $[-2, 2]^2$, $[-3, 3]^2$, and $[-4, 4]^2$ yield mean RMSEs of 0.129, 0.132, and 39.011 m, with 0, 0, and 1 extreme cases. The tightest box allows only 2.250 mean parameter changes. Rates of 1 and 8 log units s^{-1} yield 3725.669 and 602.114 m, respectively, with one extreme case each. ExtraTrees remains near 0.138 m without extreme cases across these four variants. Confirmation settings are unchanged.

C. Adaptation Under Projection Constraints

In eight archived discrete-map replays, parameters stop changing after the first update. Six select $(-2, -2)$ and exactly reproduce matched-startup fixed state and covariance; two select $(0, -2)$. Discrete-map mean RMSE is 0.148 m versus 0.158 m for the original constant. Continuous projection permits 299–620 changes but yields 0.171 m versus 0.145 m for development-tuned fixed- Q . Fig. 7 shows the parameter trace; Fig. 4 illustrates the mechanism.

Sixteen operator tests combine rates $\{2, 4, 5, 8\}$ and intervals $\{0.1, 0.2, 0.25, 0.5\}$ s. All seven sub-spacing combinations lock, while the other nine permit motion, consistent with (25). These tests do not evaluate navigation accuracy after clock changes.

D. Robustness Under Perturbations

Table IV covers the original five perturbations. Under a 1.5 m UWB bias, ET-BOTH reduces mean RMSE from 5.665 to 5.135 m, but its paired interval $[-2.109, 0.532]$ m includes zero. With UWB removal, every primary method exceeds 10 m RMSE on all eight recordings. ET-RATE reduces mean RMSE from 286.460 to 86.592 m without preventing these failures (paired interval $[-522.186, 4.027]$ m). The large IAE-projected- Q errors characterize this specific comparator, not adaptive filtering generally.

TABLE II
HOLDOUT-SET NAVIGATION PERFORMANCE AND UNCERTAINTY REPORTING

Method	Mean RMSE (m)	$T_{0.5}$ (s)	$\bar{C}(0.95)$ (%)	$\tilde{a}_{95}$ (m)	$\bar{N}_{\Delta}$	Intervention (%)
Fixed-Q (dev-tuned)	0.217	59.300	11.603	0.066	0.000	0.000
IAE-projected-Q	0.227	51.400	15.031	0.074	615.125	0.000
Ridge-OFF	243.110	166.900	62.932	0.365	272.250	0.000
Ridge-MAG	0.227	58.600	18.885	0.106	8.125	100.000
ET-OFF	0.226	51.200	60.736	0.207	579.625	0.000
ET-MAG	0.226	51.150	60.825	0.208	579.625	0.267
ET-RATE	0.226	51.350	60.797	0.207	584.375	16.897
ET-BOTH	0.226	51.350	60.822	0.207	584.250	16.759

Eight held-out recordings; 19,735 common scored position samples. Values are equal-recording means except summed $T_{0.5}$ and median $\tilde{a}_{95}$. Bold indicates the lowest RMSE and $T_{0.5}$, respectively.




Fig. 6. The two nominal extreme-error cases among eight confirmation groups. Magnitude restriction suppresses the Ridge-OFF excursion on both recordings; all four branches pass numerical checks. These illustrative cases are identified from the completed results; Table II reports all eight groups.




Fig. 7. Applied angular-rate log multiplier on archived C3-T4. Discrete projection coincides with the matched-startup fixed control at -2 .

TABLE III
ONE-SIDED AND CHANNEL-SPECIFIC INTERVENTIONS

Additional Ridge restriction	Mean RMSE (m)	Extreme	$\bar{N}_{\Delta}$
None beyond family	1005.395	3/4	173.000
Both channels, both bounds	0.206	0/4	37.750
Both lower bounds	264.373	2/4	219.250
Both upper bounds	0.206	0/4	39.000
Acceleration, both bounds	239.360	3/4	181.500
Angular rate, both bounds	0.217	0/4	53.250
Acceleration lower only	245.806	2/4	224.000
Acceleration upper only	133.313	3/4	204.750
Angular-rate lower only	992.492	3/4	179.250
Angular-rate upper only	0.217	0/4	53.250

Table V extends ExtraTrees output shifts to $\{-4, -2, 0, +2, +4\}$ log units on the same recordings.

At +4, MAG reduces mean RMSE from 0.257916 to 0.227186 m (11.9%), improving all eight recordings; the paired change is -0.030730 m (descriptive 95% bootstrap interval $[-0.037811, -0.024303]$ m). Time above 0.5 m decreases from 67.15 to 51.30 s. BOTH performs similarly (0.227230 m), whereas RATE alone does not mitigate the shift (0.257932 m). At other shifts, mean MAG–OFF differences remain below 1 mm.

At +4, MAG reduces mean parameter changes from 582.375 to 444.375, while 95% ellipsoid coverage decreases from 72.892% to 67.885%. All 160 replays complete numerically, with no recording RMSE above 10 m. The benefit is mitigation of accuracy degradation, not prevention of extreme failures.

E. Uncertainty Consistency and Recovery

Fig. 8 shows undercoverage relative to nominal confidence. For Ridge, 95% ellipsoid coverage is 62.932% without MAG and 18.885% with MAG (Table II), despite lower RMSE. At six of twelve Ridge checkpoints, OFF has two-second RMSE above 10 m, and the lowest RMSE among 34 tested continuations remains above 300 m. All 816 continuations complete numerically; at all twelve ExtraTrees checkpoints, MAG leaves the proposal unchanged. These findings delimit covariance reliability and recovery under the tested restrictions.

TABLE IV
ALL PRESPECIFIED PERTURBATIONS ON THE EIGHT CONFIRMATION GROUPS

Method	Learned low	Learned high	Learned oscillation	UWB bias	UWB gap
Fixed-Q (dev-tuned)	0.217 (0)	0.217 (0)	0.217 (0)	3.521 (0)	44.701 (8)
IAE-projected-Q	0.227 (0)	0.227 (0)	0.227 (0)	1914.356 (7)	1032.246 (8)
Ridge-OFF	211.165 (2)	159.811 (1)	224.857 (2)	927.742 (6)	3533.351 (8)
Ridge-MAG	0.227 (0)	0.227 (0)	0.227 (0)	21.869 (3)	570.631 (8)
ET-OFF	0.225 (0)	0.228 (0)	0.227 (0)	5.665 (2)	286.460 (8)
ET-MAG	0.225 (0)	0.228 (0)	0.227 (0)	5.665 (2)	286.473 (8)
ET-RATE	0.225 (0)	0.229 (0)	0.226 (0)	5.134 (3)	86.592 (8)
ET-BOTH	0.225 (0)	0.228 (0)	0.226 (0)	5.135 (3)	160.809 (8)

Each cell is mean recording RMSE in meters, followed by the number of recordings with RMSE > 10 m out of eight.




Fig. 8. Position-ellipsoid calibration on eight held-out recordings. Horizontal: nominal confidence β ; vertical: mean fraction of reference positions inside the ellipsoids. The diagonal denotes ideal calibration, $\bar{C}(\beta) = \beta$.

TABLE V
EXTRATREES PERFORMANCE UNDER GRADED LEARNED-OUTPUT
PERTURBATIONS (MEAN POSITION RMSE, M)

Shift (log units)	ET-OFF	ET-MAG	ET-RATE	ET-BOTH
-4	0.2226	0.2226	0.2221	0.2222
-2	0.2251	0.2250	0.2250	0.2250
0	0.2258	0.2258	0.2257	0.2257
+2	0.2284	0.2276	0.2286	0.2275
+4	0.2579	0.2272	0.2579	0.2272

VI. DISCUSSION

A. Interpretation

Equation (15) identifies how applied learned noise locally affects measurement gating. Its recursive consequences cannot be attributed to gating alone, because covariance propagation, Kalman gains, and subsequent learned inputs also change. The different responses of MAG and RATE to a sustained positive ExtraTrees perturbation suggest that bounding proposal magnitude and limiting its rate address distinct failure mechanisms; neither guarantees robustness to aiding faults.

The innovation-adaptive comparator achieves nominal accuracy comparable to the learned filters but markedly lower empirical position-ellipsoid coverage in these recordings. Its outcome depends on innovation-window selection, initialization, and covariance-matching settings, while learned models require training and validation. Earlier studies report navigation gains from learned process-noise adaptation [10], [14]; the present study instead evaluates restrictions on its runtime influence, not the general superiority of either adaptation paradigm.

The evidence separates three criteria that should not be conflated: position accuracy, continued parameter adaptation, and uncertainty consistency. The discrete-projection lock-in condition (25) demonstrates that reduced error can coexist with suppressed adaptation; the coverage and recovery tests likewise delimit what process-noise restrictions can establish without additional fault handling.

B. Limitations

The evidence is limited to the indoor dataset and the hybrid INS/UWB es-EKF with two noise coordinates. Selective-bound tests and recovery probes use non-holdout diagnostic recordings; IAE is an adapted comparator, and Ridge frequently requires clipping.

VII. CONCLUSION

Runtime restrictions on learned process-noise proposals mitigate severe Ridge errors and reduce ExtraTrees degradation under the evaluated perturbation, but neither calibrated uncertainty nor recovery after divergence follows from parameter constraints alone. The findings support evaluating accuracy, adaptation, and uncertainty consistency separately when deploying learning-based navigation filters. The interface-level approach may extend to other aided Kalman estimators, subject to evaluation with their sensing models.

APPENDIX A

LINEARIZATION AND NUMERICAL IMPLEMENTATION

For the state ordering in (1), the local linear error model is $\delta\dot{x} = F\delta x + Gw$, where $w \in \mathbb{R}^{12}$ is the zero-mean driving-noise vector with continuous density Q_c . The filter models

driving and measurement noises as white, zero-mean, mutually uncorrelated, and uncorrelated with the initial error; these are modeling assumptions, not verified properties of the recordings. Each zero block is 3×3 , and I below denotes $\mathbf{I}_3$:

$$F = \begin{bmatrix} 0 & I & 0 & 0 & 0 \\ 0 & 0 & -R[f]_{\times} & -R & 0 \\ 0 & 0 & -[\omega]_{\times} & 0 & -I \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (35)$$

The driving-noise order is acceleration, angular rate, and both bias random walks:

$$G = \begin{bmatrix} 0 & 0 & 0 & 0 \\ -R & 0 & 0 & 0 \\ 0 & -I & 0 & 0 \\ 0 & 0 & I & 0 \\ 0 & 0 & 0 & I \end{bmatrix}. \quad (36)$$

For a rotation increment $\phi \in \mathbb{R}^3$ (rad), let $\alpha = \|\phi\|$. The 3×3 right Jacobian [62] is

$$J_r(\phi) = \mathbf{I}_3 - \frac{1 - \cos \alpha}{\alpha^2} [\phi]_{\times} + \frac{\alpha - \sin \alpha}{\alpha^3} [\phi]_{\times}^2, \quad (37)$$

with its limit at zero evaluated by series. Let $\widehat{\delta\phi}_k$ be the attitude component of $\widehat{\delta x}_k$ in (11). The attitude correction is $\hat{R}_k^+ = \hat{R}_k^- \exp([\delta\phi_k]_{\times})$. Define $A_k = \mathbf{I}_{15} - K_k H_k$ and $J_k = \text{diag}(\mathbf{I}_3, \mathbf{I}_3, J_r(\delta\phi_k), \mathbf{I}_3, \mathbf{I}_3)$. The implementation uses

$$\tilde{P}_k^+ = A_k P_k^- A_k^T + \sigma_z^2 K_k K_k^T, \quad (38)$$

followed by the error-coordinate reset

$$P_k^+ = \text{sym}(J_k \tilde{P}_k^+ J_k^T), \quad (39)$$

where $\text{sym}(M) = (M + M^T)/2$ removes floating-point asymmetry.

The integrator holds body-frame inputs constant per substep. Each substep is at most 0.02 s and limits the product of its duration and the start-of-substep error-dynamics matrix infinity norm to 0.25. Covariance propagation then uses midpoint linearization and four-node positive-weight Gauss–Legendre quadrature for (4). Matrix exponentials use a degree-12 Taylor approximation, with scaling and squaring when required. The wrapper stops intervals needing over 5000 substeps.

APPENDIX B DATA PARTITION AND REPRODUCTION SETTINGS

TABLE VI
EXPERIMENTAL DATASET AND EVALUATION SUMMARY

Evidence unit	Extent
Distinct recording groups	26: 18 non-holdout + 8 held-out
New nominal scoring support	19,735 evaluated position samples; 984.850 s
New nominal UWB arrivals	454,012; shared across methods
Frozen training targets	482 checkpoints per objective
Initial constant- Q sweep	484 replays on 4 recordings
Initial prescribed schedules	216 replays on 2 recordings
Initial paired comparisons	64 method/recording cases; 56 follow-up cases
Original completion replays	552 unique full replays, including selection
Original completion entries	592 rows; 544 unique nonselection cases
Matched-state probes	24 states $\times$ 34 candidates = 816 branches
Graded ExtraTrees follow-up	160 replays; 64 additional recording–method–perturbation combinations

Feature preprocessing. IMU channels are ordered acceleration x, y, z , then angular rate x, y, z . First differences are not divided

by time; causal 100 Hz zero-order-hold windows include both endpoints (51/101 samples for 0.5/1 s). Ridge standardization uses training-weighted means and standard deviations (scales below 10^{-10} become one); scalars remain frozen. ExtraTrees does not standardize. Missing startup history or nonfinite features skip learned updates; invalid sensor measurements are rejected, not imputed.

Graded follow-up. No recordings are added: 96 cases at shifts $-2, 0$, and $+2$ repeat evaluated conditions; 64 at -4 and $+4$ extend the perturbation grid.

Recording identifiers. $Cx\text{-}Ty$ denotes constellation x , trial y in the UTIL TDOA2 dataset. C4-T1–T6 use the `-traj1` variant; C3-T7 and C4-T7 use `-manual1`. Table VII lists the experimental roles.

TABLE VII
RECORDING GROUPS BY EXPERIMENTAL ROLE

Role	Recording identifiers
Original training	C1-T4, C2-T5, C3-T3, C3-T6
Original selection	C2-T2
Original calibration	C1-T1, C2-T3, C3-T1, C3-T5
Original evaluation	C1-T5, C2-T6, C3-T4, C4-T3, C4-T4
Development	C1-T1, C1-T5, C2-T2, C3-T4
Extended evaluation	C1-T2, C2-T1, C3-T7, C4-T7
Held-out evaluation	C1-T3, C1-T6, C2-T4, C3-T2, C4-T1, C4-T2, C4-T5, C4-T6
Mechanism probes	C1-T5, C3-T4, C3-T7, C4-T7

Non-holdout roles overlap. The eight held-out groups are disjoint from all fitting and development groups; the union contains 26 recording groups.

Frozen training. Models use seed 1605 in scikit-learn [67]. Target searches allow 120 evaluations, starting in $[-3, 3]^2$ and expanding to $[-5, 5]^2$. Near-optimal tolerances are 1% plus 10^{-6} m² (MSE) and 0.01 nats (NLL). Quality/group weights are retained; group-exclusion fits do not replace the replay models.

Initialization. The body-to-tag lever arm ℓ has components $-0.01245, 0.00127$, and 0.0908 m in body-axis order. Initial orientation uses the minimum rotation aligning mean specific force with world vertical; heading is not observed. P_0 is diagonal. Its standard deviations, in state order (1), are 0.5 m for position, 0.3 m s^{-1} for velocity, $(0.15, 0.15, \pi/3)$ rad for rotation error, 0.1 m s^{-2} for accelerometer bias, and 0.02 rad s^{-1} for gyroscope bias. Scalar values apply to all three components; diagonal entries are their squares.

The reproduction package documents features, target-resolution diagnostics, and map screening. Supported corners do not certify interiors.

REFERENCES

- [1] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *J. Basic Eng.*, vol. 82, pp. 35–45, 1960, doi: 10.1115/1.3662552.
- [2] S. J. Julier and J. K. Uhlmann, “Unscented filtering and nonlinear estimation,” *Proc. IEEE*, vol. 92, pp. 401–422, 2004, doi: 10.1109/JPROC.2003.823141.
- [3] G. Revach, N. Shlezinger, X. Ni, A. López Escoriza, R. J. G. van Sloun, and Y. C. Eldar, “KalmanNet: neural network aided Kalman filtering for partially known dynamics,” *IEEE Trans. Signal Process.*, vol. 70, pp. 1532–1547, 2022, doi: 10.1109/TSP.2022.3158588.
- [4] T. Haamoja, A. Ajay, S. Levine, and P. Abbeel, “Backprop KF: learning discriminative deterministic state estimators,” in *Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 4376–4384.

- [5] P. Becker, H. Pandya, G. Gebhardt, C. Zhao, C. J. Taylor, and G. Neumann, "Recurrent Kalman networks: factorized inference in high-dimensional deep feature spaces," in *Proc. Mach. Learn. Res.*, vol. 97, 2019, pp. 544–552.
- [6] A. Kloss, G. Martius, and J. Bohg, "How to train your differentiable filter," *Auton. Robots*, vol. 45, pp. 561–578, 2021, doi: 10.1007/s10514-021-09990-9.
- [7] N. Shlezinger, G. Revach, A. Ghosh, S. Chatterjee, S. Tang, T. Imbiriba, J. Dunik, O. Straka, P. Closas, and Y. C. Eldar, "Artificial intelligence-aided Kalman filters: AI-augmented designs for Kalman-type algorithms," *IEEE Signal Process. Mag.*, vol. 42, no. 3, pp. 52–76, 2025, doi: 10.1109/MSP.2025.3569395.
- [8] M. Brossard, A. Barrau, and S. Bonnabel, "AI-IMU dead-reckoning," *IEEE Trans. Intell. Veh.*, vol. 5, pp. 585–595, 2020, doi: 10.1109/TIV.2020.2980758.
- [9] W. Liu *et al.*, "TLIO: tight learned inertial odometry," *IEEE Robot. Autom. Lett.*, vol. 5, pp. 5653–5660, 2020, doi: 10.1109/LRA.2020.3007421.
- [10] B. Or and I. Klein, "A hybrid model and learning-based adaptive navigation filter," *IEEE Trans. Instrum. Meas.*, vol. 71, 2022, art. no. 2516311, doi: 10.1109/TIM.2022.3197775.
- [11] N. Cohen and I. Klein, "Adaptive Kalman-informed transformer," *Eng. Appl. Artif. Intell.*, vol. 146, 2025, art. no. 110221, doi: 10.1016/j.engappai.2025.110221.
- [12] A. Levy and I. Klein, "Adaptive neural unscented Kalman filter," *IEEE Trans. Intell. Veh.*, vol. 11, pp. 428–437, 2026, doi: 10.1109/TIV.2026.3657014.
- [13] B. Diker and I. Klein, "Neural aided adaptive innovation-based invariant Kalman filter," 2026, arXiv:2603.26709. [Online]. Available: <https://arxiv.org/abs/2603.26709>
- [14] B. Or and I. Klein, "ProNet: adaptive process noise estimation for INS/DVL fusion," in *2023 IEEE Underwater Technology (UT)*, 2023, doi: 10.1109/UT49729.2023.10103449.
- [15] B. Or and I. Klein, "Learning vehicle trajectory uncertainty," *Eng. Appl. Artif. Intell.*, vol. 122, 2023, art. no. 106101, doi: 10.1016/j.engappai.2023.106101.
- [16] R. K. Mehra, "On the identification of variances and adaptive Kalman filtering," *IEEE Trans. Autom. Control*, vol. 15, pp. 175–184, 1970, doi: 10.1109/TAC.1970.1099422.
- [17] R. K. Mehra, "Approaches to adaptive filtering," *IEEE Trans. Autom. Control*, vol. 17, pp. 693–698, 1972, doi: 10.1109/TAC.1972.1100100.
- [18] K. A. Myers and B. D. Tapley, "Adaptive sequential estimation with unknown noise statistics," *IEEE Trans. Autom. Control*, vol. 21, pp. 520–523, 1976, doi: 10.1109/TAC.1976.1101260.
- [19] A. H. Mohamed and K. P. Schwarz, "Adaptive Kalman filtering for INS/GPS," *J. Geod.*, vol. 73, pp. 193–203, 1999, doi: 10.1007/s001900050236.
- [20] B. J. Odelson, M. R. Rajamani, and J. B. Rawlings, "A new autocovariance least-squares method for estimating noise covariances," *Automatica*, vol. 42, pp. 303–308, 2006, doi: 10.1016/j.automatica.2005.09.006.
- [21] J. Dunik, O. Straka, O. Kost, and J. Havlík, "Noise covariance matrices in state-space models: a survey and comparison of estimation methods-Part I," *Int. J. Adapt. Control Signal Process.*, vol. 31, pp. 1505–1543, 2017, doi: 10.1002/acs.2783.
- [22] S. Särkkä and A. Nummenmaa, "Recursive noise adaptive Kalman filtering by variational Bayesian approximations," *IEEE Trans. Autom. Control*, vol. 54, pp. 596–600, 2009, doi: 10.1109/TAC.2008.2008348.
- [23] Y. Huang, Y. Zhang, Z. Wu, N. Li, and J. A. Chambers, "A novel adaptive Kalman filter with inaccurate process and measurement noise covariance matrices," *IEEE Trans. Autom. Control*, vol. 63, pp. 594–601, 2018, doi: 10.1109/TAC.2017.2730480.
- [24] B. Or, B. Z. Bobrovsky, and I. Klein, "Kalman filtering with adaptive step size using a covariance-based criterion," *IEEE Trans. Instrum. Meas.*, vol. 70, pp. 1–10, 2021, doi: 10.1109/TIM.2021.3063191.
- [25] B. Or and I. Klein, "Adaptive step size learning with applications to velocity aided inertial navigation system," *IEEE Access*, vol. 10, pp. 85 818–85 830, 2022, doi: 10.1109/ACCESS.2022.3198672.
- [26] M. Freydin and B. Or, "Learning car speed using inertial sensors for dead reckoning navigation," *IEEE Sensors Lett.*, vol. 6, 2022, art. no. 6003104, doi: 10.1109/LSSENS.2022.3201731.
- [27] B. Or, "Learning-based speed estimation from accelerometer-only inertial sensing," 2026, arXiv:2401.07468v4. [Online]. Available: <https://arxiv.org/abs/2401.07468v4>
- [28] M. Freydin, N. Segol, N. Sfaradi, A. Eweida, and B. Or, "Deep learning for inertial sensor alignment," *IEEE Sensors J.*, vol. 24, pp. 17 282–17 290, 2024, doi: 10.1109/JSEN.2024.3384302.
- [29] B. Or, N. Segol, A. Eweida, and M. Freydin, "Learning position from vehicle vibration using an inertial measurement unit," *IEEE Trans. Intell. Transp. Syst.*, vol. 25, pp. 10 766–10 776, 2024, doi: 10.1109/TITS.2024.3422037.
- [30] B. Or and I. Klein, "A hybrid adaptive velocity aided navigation filter with application to INS/DVL fusion," in *OCEANS 2022 Hampton Roads*, 2022, doi: 10.1109/OCEANS47191.2022.9977082.
- [31] T. Gneiting and A. E. Raftery, "Strictly proper scoring rules, prediction, and estimation," *J. Am. Stat. Assoc.*, vol. 102, pp. 359–378, 2007, doi: 10.1198/016214506000001437.
- [32] Y. Ovadia *et al.*, "Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift," in *Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 13 991–14 002.
- [33] A. N. Angelopoulos and S. Bates, "A gentle introduction to conformal prediction and distribution-free uncertainty quantification," 2022, arXiv:2107.07511v6. [Online]. Available: <https://arxiv.org/abs/2107.07511v6>
- [34] I. Klein, G. Revach, N. Shlezinger, J. E. Mehr, R. J. G. van Sloun, and Y. C. Eldar, "Uncertainty in data-driven Kalman filtering for partially known state-space models," in *2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 3194–3198, doi: 10.1109/ICASSP43922.2022.9746732.
- [35] B. Or, "Selection-induced contraction of innovation statistics in gated Kalman filters," 2026, arXiv:2512.18508v3. [Online]. Available: <https://arxiv.org/abs/2512.18508v3>
- [36] Joint Committee for Guides in Metrology, *International Vocabulary of Metrology: Basic and General Concepts and Associated Terms (VIM)*, 3rd ed., BIPM, 2012, jCGM 200:2012, doi: 10.59161/JCGM200-2012.
- [37] Joint Committee for Guides in Metrology, *Evaluation of Measurement Data: Guide to the Expression of Uncertainty in Measurement*, BIPM, 2008, jCGM 100:2008(E), doi: 10.59161/JCGM100-2008E.
- [38] J. Qu, Y. Yuan, F. Shen, W. Li *et al.*, "A generalized maximum correntropy extended kalman filter for uwb/imu localization," *IEEE Sensors J.*, vol. 25, no. 19, pp. 36 267–36 280, 2025.
- [39] F. Chen, L. Zhao, L. Hao, Z. Tang *et al.*, "Adaptive fusion localization optimization framework for ranging signal loss mitigation in wireless multisensor systems," *IEEE Trans. Instrum. Meas.*, vol. 75, 2026.
- [40] L. Qian, L. Chang, K. Li, F. Qin, T. Zhu, and P. Gao, "The optimized SINS/DVL integration based on $SE_3(3)$ quaternion error model and adaptive estimator," *IEEE Trans. Instrum. Meas.*, vol. 74, 2025, doi: 10.1109/TIM.2025.3545191.
- [41] F. Zhu, S. Zhang, X. Li, Y. Huang, and Y. Zhang, "Adaptive kalman filters with small-magnitude and inaccurate process noise covariance matrix—part i: Theoretical results," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 61, no. 3, pp. 6964–6986, 2025.
- [42] F. Zhu, S. Zhang, X. Li, Y. Huang, and Y. Zhang, "Adaptive kalman filters with small-magnitude and inaccurate process noise covariance matrix—part ii: Application to inertial-based integrated navigation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 61, no. 3, pp. 7205–7235, 2025.
- [43] G. Huang, J. Liu, Y. Zhao *et al.*, "Vehicle inertial localization with adaptive noise estimation and pseudo-measurement constraints in gnss-denied environment," *IEEE Internet Things J.*, vol. 12, no. 21, pp. 44 763–44 779, 2025.
- [44] Q. Wang, D. Zhou, Y. Li, and S. Li, "A stability guaranteed distributionally robust adaptive kalman filter for dual uncertainties in model and noise," *IEEE Trans. Instrum. Meas.*, vol. 75, 2026.
- [45] Z. Jiang, G. Battistelli, L. Chisci, N. Forti, and W. Zhou, "Outlier-robust centralized and distributed variational bayesian moving horizon estimation," *IEEE Trans. Signal Process.*, vol. 73, pp. 2527–2541, 2025.
- [46] P. Yan, X. Zhan, R. Sun, and L.-T. Hsu, "Credible uncertainty quantification under noise and system model mismatch," *IEEE Trans. Instrum. Meas.*, vol. 75, 2026, art. no. 8506712, doi: 10.1109/TIM.2026.3682809.
- [47] L. Brunke *et al.*, "Safe learning in robotics: from learning-based control to safe reinforcement learning," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 5, pp. 411–444, 2022, doi: 10.1146/annurev-control-042920-020211.
- [48] M. Alshiekh *et al.*, "Safe reinforcement learning via shielding," in *Proc. AAAI Conf. Artif. Intell.*, vol. 32, 2018, pp. 2669–2678, doi: 10.1609/aaai.v32i1.11797.
- [49] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, pp. 3861–3876, 2017, doi: 10.1109/TAC.2016.2638961.
- [50] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," *Automatica*, vol. 129, 2021, art. no. 109597, doi: 10.1016/j.automatica.2021.109597.

- [51] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: a unified view of safety-critical control in autonomous systems," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 7, pp. 47–72, 2024, doi: 10.1146/annurev-control-071723-102940.
- [52] Z. Ravichandran, A. Robey, V. Kumar, G. J. Pappas, and H. Hassani, "Safety guardrails for LLM-enabled robots," *IEEE Robot. Autom. Lett.*, vol. 11, no. 4, pp. 4649–4656, 2026, doi: 10.1109/LRA.2026.3667488.
- [53] J. Kim, W. Chen, D. Soleymanzadeh, Y. Ding *et al.*, "Modular safety guardrails are necessary for foundation-model-enabled robots in the real world," *arXiv preprint arXiv:2602.04056*, 2026. [Online]. Available: <https://arxiv.org/abs/2602.04056>
- [54] A. Lin, S. Peng, and S. Bansal, "One filter to deploy them all: Robust safety for quadrupedal navigation in unknown environments," *IEEE Trans. Robot.*, vol. 42, pp. 545–560, 2026, doi: 10.1109/TRO.2025.3644957.
- [55] Z. Mei, A. Dixit, M. Booker, E. Zhou *et al.*, "Perceive with confidence: Statistical safety assurances for navigation with learning-based perception," *Int. J. Robot. Res.*, vol. 45, no. 6, pp. 938–967, 2026.
- [56] J. Seo, K. Nakamura, and A. Bajcsy, "Uncertainty-aware latent safety filters for avoiding out-of-distribution failures," in *Proc. Conf. Robot Learning*, 2025, pp. 4442–4472. [Online]. Available: <https://proceedings.mlr.press/v305/seo25a.html>
- [57] B. Zhang, J. Li, J. Shen, Y. Zhang, Y. Cai, Y. Chen, J. Dai, J. Ji, and Y. Yang, "VLA-arena: An open-source framework for benchmarking vision-language-action models," in *Proceedings of the 43rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 306, 2026, pp. 156 567–156 614. [Online]. Available: <https://proceedings.mlr.press/v306/zhang26cp.html>
- [58] F. Bai, Y. Li, R. Wu, P. Wang, Y. Wang, B. H. Zhu, Y. Wang, T. Chou, J. Gao, R. Zhu, Y. Wen, Y. Yang, and Y. Chen, "SafeLab: An interactive high-fidelity benchmark for embodied safety in scientific robotics," in *Proceedings of the 43rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 306, 2026, pp. 5324–5347. [Online]. Available: <https://proceedings.mlr.press/v306/bai26e.html>
- [59] Z. Zhang, H. Ren, Y. Sun, Y. Sheng, H. Wang, Z. Wu, H. Lin, P.-L. Bacon, and Y. Yu, "Towards practical world model-based reinforcement learning for vision-language-action models," in *Proceedings of the 43rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, vol. 306, 2026, pp. 158 191–158 214. [Online]. Available: <https://proceedings.mlr.press/v306/zhang26ez.html>
- [60] W. Zhao, A. Goudar, X. Qiao, and A. P. Schoellig, "UTIL: an ultra-wideband time-difference-of-arrival indoor localization dataset," *Int. J. Robot. Res.*, vol. 43, pp. 1443–1456, 2024, doi: 10.1177/02783649241230640.
- [61] Dynamic Systems Lab, University of Toronto, "UTIL: Ultra-wideband Dataset, Experimental Setup and Data Documentation," Dataset documentation, accessed: Oct. 9, 2026. [Online]. Available: <https://learnsyslab.github.io/util-uwbd-dataset/>
- [62] J. Solà, "Quaternion kinematics for the error-state Kalman filter," 2017, arXiv:1711.02508. [Online]. Available: <https://arxiv.org/abs/1711.02508>
- [63] C. F. Van Loan, "Computing integrals involving the matrix exponential," *IEEE Trans. Autom. Control*, vol. 23, pp. 395–404, 1978, doi: 10.1109/TAC.1978.1101743.
- [64] A. E. Hoerl and R. W. Kennard, "Ridge regression: biased estimation for nonorthogonal problems," *Technometrics*, vol. 12, pp. 55–67, 1970, doi: 10.1080/00401706.1970.10488634.
- [65] P. Geurts, D. Ernst, and L. Wehenkel, "Extremely randomized trees," *Mach. Learn.*, vol. 63, pp. 3–42, 2006, doi: 10.1007/s10994-006-6226-1.
- [66] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users' Guide*, 3rd ed. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 1999.
- [67] F. Pedregosa *et al.*, "Scikit-learn: machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.



Barak Or received the B.Sc. degree in aerospace engineering and the B.A. degree (cum laude) in economics and management from the Technion–Israel Institute of Technology in 2016, the M.Sc. degree in aerospace engineering from the Technion in 2018, and the Ph.D. degree from the University of Haifa in 2022.

He is the Founder and Chief Executive Officer of STATE16 Technologies Ltd., Haifa, Israel, and Chairman of ArtificialGate Ltd. He is also an Adjunct Lecturer at the Technion and Reichman University, and Academic Director of the Google–Reichman Tech School. His research interests include navigation, sensor fusion, estimation theory, deep learning, and physical AI.