DAKF: TRACKING NONSTATIONARY DYNAMICS WITH PRECISION-WEIGHTED EXCITATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Tracking changing dynamics requires deciding when to retain a model and when to adapt from noisy, dependent observations. We introduce the Dual-Arm Kalman Filter (DAKF), coupling retained-model filtering and adaptive identification through hard selection and majority synchronization. Precision-weighted persistent excitation (PW-PE) yields a finite-sample identification bound under calibrated conditional noise, with explicit filtering-bias and drift costs and a matching scalar precision lower bound. A filtered-response bias bound connects this guarantee to model mismatch, state uncertainty and selection. Experiments cover synthetic changes, variance diagnostics and recorded UAV braking, with covariance-matched classical baselines and an official KalmanNet implementation. Synthetic ablations support the architecture; the braking benchmark uses common ground-truth initialization, with cold-start sensitivity reported separately.

1 INTRODUCTION

Motion models support state estimation in autonomous vehicles, aerial robots and tracking systems. Braking or a maneuver can invalidate a previously accurate model, causing systematic prediction errors. Updating the model online is a natural response, but noisy observations can make an apparent change look real. Updating too quickly can replace a reliable model with an inaccurate one; updating too slowly prolongs the use of an outdated model.

We observe z_t from a changing linear system

$$x_t = A_t x_{t-1} + w_t, \qquad z_t = H x_t + v_t, \tag{1}$$

where x_t is hidden, A_t may change, and the process and observation noises have time-varying covariances Q_t, R_t. At time t, the estimator must produce $\widehat{x}_t$ using only $z_{0:t}$. It must also identify a usable transition from correlated, filtered states rather than direct noise-free state observations. Observation quality matters: a short accurate interval can carry more identification information than a long noisy interval, provided it excites the necessary directions (Narendra & Annaswamy, 1987).

The research question is how an online filter can retain, adapt and replace its motion model after a change, while using unequal sample precision to obtain an identification guarantee that accounts for temporal dependence and filtering-induced bias.

We introduce DAKF, a dual-arm estimator for nonstationary dynamics. One arm retains a piecewise-fixed transition; the other combines Kalman filtering with precision-weighted recursive identification. A discounted loss score selects their output, and sustained preference for the adaptive arm triggers synchronization. This separates retaining a usable model from learning its replacement. The arms are estimators, not actions in a bandit exploration problem.

Our contributions are threefold. First, DAKF combines online transition identification, hard selection, and majority synchronization for changing dynamics. Second, the PW-PE analysis separates precision, temporal dependence, filtering bias, and local drift, with a matching scalar precision lower bound and a conditional transfer to state error. Third, we test the mechanism with controlled variance and architecture ablations, an official learned-filter baseline, covariance-matched classical methods, and a recorded UAV braking case.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

2 Related Work

Model-based filtering and recursive adaptation. The Kalman filter (Kalman, 1960; Kailath et al., 2000), innovation-based covariance adaptation (Mohamed & Schwarz, 1999), strong tracking (Zhou & Frank, 1996), and interacting multiple models (IMM) (Blom & Bar-Shalom, 1988) provide the classical context. Optimized Kalman filtering fits noise covariances to state-estimation loss (Greenberg et al., 2023), while Lai & Bernstein (2024a) transfer RLS forgetting rules to adaptive Kalman covariance updates. Their generalized-forgetting analysis gives noiseless fixed-parameter stability conditions and robustness bounds for changing parameters with measurement and regressor noise (Lai & Bernstein, 2024b). DAKF differs in maintaining a retained transition model alongside an identified transition, then synchronizing after persistent preference for the adaptive arm. This architecture addresses when to retain an identified model and when to replace it.

Learned filtering under model mismatch. KalmanNet learns a gain within a filtering recursion (Revach et al., 2022); Split-KalmanNet separates state and innovation covariance learning (Choi et al., 2023). Adaptive KalmanNet uses context-dependent modulation for changing noise conditions (Ni et al., 2024), and MAML-KalmanNet uses model-agnostic meta-learning for adaptation (Chen et al., 2025). DANSE learns a state prior from measurement trajectories without a known process model or state labels (Ghosh et al., 2024). Cholesky-KalmanNet enforces positive-definite covariance structure (Ko & Shafieezadeh, 2025), while Recursive KalmanNet combines recursive covariance propagation with a likelihood-based training objective (Mortada et al., 2025). These approaches modify learned gains, state priors or covariance propagation; DAKF identifies an explicit transition within an online two-filter architecture.

Online adaptation and changing regimes. OTAKNet addresses noise-statistics drift through observation-driven optimal-transport adaptation (He et al., 2025). Two 2026 preprints are especially close in motivation: CASA-KalmanNet uses neural change indicators to regulate online adaptation (Zhang et al., 2026), whereas EM-KalmanNet unfolds parameter estimation and smoothing for blockwise changing state-space models (Cohen et al., 2026). DAKF instead uses explicit recursive transition identification and a two-filter selection rule, without an offline-trained neural module. Its PW-PE result characterizes the precision-weighted design and the bias in filtered-state regression. We train the original KalmanNet architecture as a direct learned-filter baseline. The context-adaptive and 2026 variants are related methods rather than evaluated baselines.

Nonstationarity, dependence, and guarantees. Classical identification and excitation theory (Ljung, 1999; Narendra & Annaswamy, 1987) is complemented by finite-sample identification without mixing (Simchowitz et al., 2018), long-context identification (Yüksel et al., 2025), and continuous-time finite-observation analysis (Zhou et al., 2025). More directly for changing dynamics, Xin et al. (2025) analyze online change detection with correlated system trajectories; PLOT learns a changing target model with forgotten RLS and gives dynamic-regret guarantees for tracking control (Tsiamis et al., 2024). These results concern identification, change detection or tracking control; our target is parameter identification within a latent-state filter. We use self-normalized martingale concentration (Abbasi-Yadkori et al., 2011) to allow predictable dependent regressors. The active-design assumption behind the harmonic-precision bandit result of Zhao et al. (2026) is compared with PW-PE in Section 4.2. Our claim is a conditional precision-dependent identification bound with explicit drift and filtering-bias terms, followed by a tracking-error decomposition under stated stability conditions.

3 Dual-Arm Kalman Filtering

3.1 Setting and the two estimation roles

In Eq. (1), $x_t \in \mathbb{R}^n$, $z_t \in \mathbb{R}^m$, and A_t is unknown. We allow piecewise-constant or locally drifting transitions; drift enters the analysis explicitly. Both arms output state estimates online. Arm 1 uses $A_t^{(1)} = A_{\text{base}}$ between synchronization events. Arm 2 uses an estimate $A_t^{(2)} = \widehat{A}_t^{\text{pred}}$ constructed from completed transitions before observing z_t. Each arm performs a standard Kalman prediction and correction. Write $\bar{x}_{t-1}^{(i)}$ for its input state: this can be its own posterior, the previous selected posterior, or a synchronized copy, as specified by the task implementation. The estimate is

$$\widehat{x}_t^{(i)} = A_t^{(i)} \bar{x}_{t-1}^{(i)} + K_t^{(i)} (z_t - H A_t^{(i)} \bar{x}_{t-1}^{(i)}). \tag{2}$$

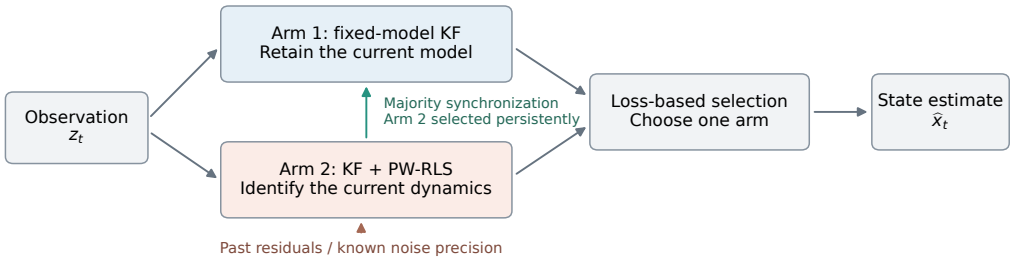

Figure 1: DAKF separates retaining a model from identifying its replacement. Both Kalman arms see the same observation. A hard score comparison selects the output; a sustained majority triggers synchronization. Precision weights for an identification sample are fixed using information available before that sample, and are retained if it is replayed.

The distinction matters: the synthetic implementation shares the selected posterior, while the UAV implementation maintains separate arm states. We preserve these benchmark choices, including their selection losses and synchronization schedules (Appendix A). Arm 1 is a retained-model reference; its name does not imply an unconditional stability guarantee after arbitrary switching.

3.2 Precision-weighted identification

For a completed filtered-state transition, let ϕ_{s-1} be the past regressor and y_s its response. In the scalar synthetic task these are consecutive selected posterior states. Let $\sigma^2_{\epsilon,s} > 0$ be a predictable effective regression-noise proxy, distinct from measurement covariance R_s. Precision-weighted RLS updates

$$g_s = \frac{P_{s-1}\phi_{s-1}}{\gamma_{\mathrm{RLS}}\sigma^2_{\epsilon,s} + \phi_{s-1}^\top P_{s-1}\phi_{s-1}}, \qquad \widehat{A}_s = \widehat{A}_{s-1} + (y_s - \widehat{A}_{s-1}\phi_{s-1})g_s^\top, \tag{3}$$

$$P_s = \gamma_{\mathrm{RLS}}^{-1}(P_{s-1} - g_s\phi_{s-1}^\top P_{s-1}), \qquad V_s = P_s^{-1} = \gamma_{\mathrm{RLS}}V_{s-1} + \frac{\phi_{s-1}\phi_{s-1}^\top}{\sigma^2_{\epsilon,s}}.$$

Here s indexes an identification update; replay is discussed below. We initialize the parameter at the baseline model and use a positive information prior. A common precision can be absorbed into the initial information matrix. Thus normalized constant weights give exactly the usual RLS recursion. The homoscedastic benchmarks use this constant-precision case; the variance experiment changes only identification weights.

The implementable proxy is the sample variance of at most 20 previously completed regression residuals, clipped to $[10^{-5}, 10^3]$, with initial reference 0.1. It is frozen before the current observation; causality does not prove calibration. The oracle ablation computes the selected posterior's conditional variance under the known generative model, accounting for hard selection. This conditional variance is not automatically a valid sub-Gaussian scale for the theorem.

3.3 Hard selection, majority vote, and synchronization

Each arm accumulates a discounted loss score $r_t^{(i)}$. The abrupt UAV version uses normalized squared innovations and a recent-loss variance penalty:

$$r_t^{(i)} = \gamma_{\mathrm{regret}} r_{t-1}^{(i)} + (1 - \gamma_{\mathrm{regret}})\big(\ell_t^{(i)} + \beta\,\mathrm{Var}_W(\ell^{(i)})\big). \tag{4}$$

The synthetic version uses posterior measurement residuals, with its existing task-specific risk adjustment and hysteresis. The lower-score arm supplies the output, subject to the task's switching threshold. A recent majority for Arm 2 triggers re-anchoring of the baseline and score reset. This "regret" terminology denotes an operational score: Eq. (4) is not a proved online-learning regret bound. Appendix A gives the exact four configurations, including a retained limitation of the gradual UAV synchronization code.

Algorithm 1 DAKF: causal processing order

Require: Baseline transition, Kalman covariances, RLS prior, task-specific score and synchronization settings

1: Initialize both arms and the identification estimate
2: **for** each observation time t **do**
3: Apply scheduled synchronization from completed votes, if the task synchronizes before prediction
4: Identify from completed transitions using Eq. (3) and their frozen precision weights
5: Form both predictions and gains; freeze the current regression variance proxy from past information
6: Receive z_t and correct both arms by Eq. (2)
7: Update loss scores; hard-select k_t and output $\widehat{x}_t = \widehat{x}_t^{(k_t)}$
8: Update majority history and perform or schedule the task's synchronization
9: Store the completed transition for subsequent identification
10: **end for**

4 Guarantees through Precision-Weighted Excitation

4.1 A local identification model

Let $\mathcal{F}_{s-1}$ contain observations and algorithmic decisions before physical sample s. On a local interval, write the filtered-state regression as

$$y_s = A_s\phi_{s-1} + b_s + \xi_s, \qquad \mathbb{E}[\xi_s \mid \mathcal{F}_{s-1}] = 0. \tag{5}$$

The transition A_s, regressor, scale $\sigma_{\epsilon,s}$, and bias b_s are predictable. Each coordinate of $\xi_s/\sigma_{\epsilon,s}$ is conditionally 1-sub-Gaussian. Regressors may depend arbitrarily on past observations subject to these conditions and $\|\phi_{s-1}\|_2 \le L$. No independence or mixing assumption is imposed on them. Persistent but directionally degenerate data can still fail excitation.

At a fixed horizon T, let $a_s \in [0,1]$ be deterministic discount coefficients, $q_s = a_s/\sigma_{\epsilon,s}^2$, and define

$$S_T = \sum_{s=1}^{T} q_s, \quad G_T = \sum_{s=1}^{T} q_s\phi_{s-1}\phi_{s-1}^\top, \quad V_T = \lambda I + G_T, \quad \lambda > 0. \tag{6}$$

The weighted ridge estimate minimizes $\sum_s q_s\|y_s - B\phi_{s-1}\|_2^2 + \lambda\|B - A_{\text{init}}\|_F^2$. Its recursive realization is Eq. (3) with the corresponding forgotten prior.

Definition 1 (PW-PE). *The realized regressors satisfy precision-weighted persistent excitation at T with constant $\alpha > 0$ when $G_T \succeq \alpha S_T I$.*

PW-PE asks whether informative directions are present among the high-precision samples, rather than only among all samples. A small noise variance increases S_T, but cannot compensate for missing directions in G_T. The condition is stated on the actual weighted design; we do not infer it from ordinary unweighted PE. In the scalar fixed-design diagnostic, $\phi_{s-1} = 1$ and $\alpha = 1$ exactly.

4.2 Identification accuracy and harmonic precision

Theorem 1 (Local identification with bias and drift). *Under the preceding conditional-noise and regressor conditions, suppose PW-PE holds. Set $S_A = \max_j \|(A_T - A_{\text{init}})_{j,:}\|_2$. With probability at least $1-\delta$, at the fixed horizon T,*

$$\|\widehat{A}_T - A_T\|_F \le \frac{\sqrt{n}\,\beta_T + \left(\sum_s q_s\|b_s\|_2^2\right)^{1/2}}{\sqrt{\lambda + \alpha S_T}} + \frac{L^2\sum_s q_s\|A_s - A_T\|_F}{\lambda + \alpha S_T}, \tag{7}$$

$$\beta_T = \sqrt{\lambda}S_A + \sqrt{n\log\left(1 + \frac{L^2 S_T}{n\lambda}\right) + 2\log\frac{n}{\delta}}.$$

If excitation holds only on an event, the probability statement is the intersection of the martingale confidence event with that excitation event.

The proof normalizes each distinct sample by its precision and applies the self-normalized inequality of Abbasi-Yadkori et al. (2011). Weighted Cauchy–Schwarz controls filtering bias, and the normal equations isolate drift relative to A_T; Appendix F gives the details. For constant A, zero bias, fixed dimension, and sufficient excitation, the leading stochastic term is $\widetilde{O}((\alpha S_T)^{-1/2})$. This is a conditional identification guarantee, rather than a proof that the full selector is minimax optimal.

For $a_s = 1$, write $\sigma^2_{\mathrm{harm}} = T/(\sum_s \sigma^{-2}_{\epsilon,s})$. Then

$$S_T^{-1/2} = \frac{\sigma_{\mathrm{harm}}}{\sqrt{T}} \le \frac{\sqrt{\sum_s \sigma^2_{\epsilon,s}}}{T}. \tag{8}$$

The inequality follows from Cauchy–Schwarz. In a scalar Gaussian constant-design regression, the left side is the exact standard deviation of unregularized precision-weighted least squares, whereas the right side is that of the unweighted mean. This illustrates the value of a few low-variance samples without asserting the same exact comparison for endogenous filtered-state regression.

Proposition 1 (Scalar precision lower bound). *For $y_s = A + \epsilon_s$, independent $\epsilon_s \sim \mathcal{N}(0, \sigma^2_{\epsilon,s})$ with known deterministic variances, there exists a universal $c > 0$ such that*

$$\inf_{\widehat{A}} \sup_{A\in[-1,1]} \mathbb{E}_A|\widehat{A} - A| \ge c \min\{1, S_T^{-1/2}\}, \qquad S_T = \sum_s \sigma^{-2}_{\epsilon,s}. \tag{9}$$

Proposition 1 identifies the precision scale that no estimator can improve on uniformly, even for a scalar constant-design problem. Inverse-variance estimation is classical (Kailath et al., 2000); the issue here is how that information enters identification from a changing, partially observed system. Simchowitz et al. (2018) analyze a directly observed state trajectory, and PLOT (Tsiamis et al., 2024) learns a changing target from sequentially revealed target states for tracking control. In Eq. (5), both the regressor and response are filtered estimates. Theorem 1 therefore retains their conditional bias and local parameter drift, while PW-PE specifies whether accurate samples cover the required directions.

Zhao et al. (2026, Section 3) study rewards $r_t = \langle a_t, \theta^* \rangle + \eta_t$ with a fixed parameter: the learner chooses a_t from a fixed action set and observes the noise scale after the reward. Their VAEE algorithm actively probes uncertain directions; its near-harmonic bound includes a correction removing a dimension-dependent number of the largest precision contributions. DAKF cannot choose its trajectory's regressors, and filtered responses introduce the bias b_s absent from that centered reward model. Our guarantee consequently requires $G_T \succeq \alpha S_T I$ on the realized trajectory and includes explicit bias and drift terms. The distinction is how directional information is obtained and what is estimated, rather than whether past-dependent data are allowed. The harmonic scale alone does not turn their best-action guarantee into a transition-identification or state-tracking guarantee.

Remark 1 (Low-variance intervals accelerate information accumulation). With m accurate samples of variance σ_L^2 and $T - m$ noisy samples of variance σ_H^2,

$$S_T = \frac{m}{\sigma_L^2} + \frac{T-m}{\sigma_H^2}.$$

When the first term dominates, the stochastic scale is approximately $\sigma_L/\sqrt{m}$: a short accurate interval can supply most of the identification information. For $T = 100$, $m = 10$, $\sigma_L = 0.1$ and $\sigma_H = 1$, the scalar Gaussian weighted standard error is 0.0303, versus 0.0949 for the unweighted mean. Each accurate sample contributes as much precision as 100 noisy samples. Under PW-PE, this information also reduces the stochastic term in Theorem 1; bias and drift remain separate costs.

For example, consider UAV identification across gusty and calm operating intervals. If calmer conditions reduce unpredictable aerodynamic disturbances and hence the calibrated effective regression variance, while the motion still excites the relevant directions, the calm interval can substantially accelerate identification. This is an application insight, not a wind-data experiment: wind speed alone does not determine sensor noise, remove filtering bias or guarantee excitation.

Proposition 2 (Bias of an actual filtered response). *Condition on $\mathcal{F}_{s-1}$, with predictable A_s, H and shared input ϕ. Let $\Pi = \mathbb{E}[(x_{s-1} - \phi)(x_{s-1} - \phi)^\top \mid \mathcal{F}_{s-1}]$, with conditionally centered process/observation noises of covariances Q, R, uncorrelated with each other and with the prior centered state. Arm i uses predictable transition B_i, gain K_i, and posterior $(I - K_iH)B_i\phi + K_iz_s$.*

If $p = \Pr(k_s = 2 \mid \mathcal{F}_{s-1})$, $p_1 = 1 - p$, $p_2 = p$, and $\bar{K} = \sum_i p_i K_i$, then the selected regression bias satisfies

$$\begin{aligned} \|b_s\|_2 \le B_s := &\sum_{i=1}^{2} p_i \|(I - K_i H)(B_i - A_s)\phi\|_2 + \|\bar{K} H A_s\|_2 \sqrt{\operatorname{tr} \Pi} \\ &+ \sqrt{p(1-p)\operatorname{tr}[(K_2 - K_1)\bar{\Omega}(K_2 - K_1)^\top]}, \\ \bar{\Omega} := &H(A_s \Pi A_s^\top + Q)H^\top + R. \end{aligned} \tag{10}$$

For a predictable single arm, take its selection probability to be one; the last term vanishes.

This bounds b_s by model mismatch, actual conditional state error and selection between unequal gains. Substituting B_s into Theorem 1 makes its bias cost explicit. Importantly, Π is a conditional error second moment, not the algorithm's nominal Riccati matrix; replacing it by P requires a proved envelope $\Pi \preceq c_P P$. Appendix F.5 proves the bound and audits the executed scalar recursion.

4.3 Forgetting, replay, and transfer to state tracking

For one pass with forgetting $\gamma = \gamma_{\mathrm{RLS}}$,

$$S_{T,\gamma} = \sum_{s=1}^{T} \frac{\gamma^{T-s}}{\sigma_{\epsilon,s}^2}, \qquad \operatorname{Var}(\widehat{A}_T) = \frac{\sum_s \gamma^{2(T-s)}/\sigma_{\epsilon,s}^2}{S_{T,\gamma}^2} \le S_{T,\gamma}^{-1}, \tag{11}$$

The equality holds for scalar Gaussian constant-design regression; at $\gamma = 1$ it becomes $1/S_T$. Forgetting trades retained precision against stale-model drift. Synthetic window replay reuses noise rather than creating independent samples: Appendix F.3 aggregates sample coefficients and rescales the prior, with numerical agreement between recursive and batch estimates.

For $y_s = \widehat{x}_s$, $\phi_{s-1} = \widehat{x}_{s-1}$, the residual is $w_s + A_s e_{s-1} - e_s$ and need not be centered. Proposition 2 controls its mean. If the true centered sub-Gaussian scale is at most $\sqrt{\kappa}$ times the estimated scale, the theorem's stochastic radius acquires the factor $\sqrt{\kappa}$; a causal sample variance alone does not certify this calibration.

Finally, let $\bar{e}_{t-1}^{(i)} = x_{t-1} - \bar{x}_{t-1}^{(i)}$ denote the actual input error in Eq. (2). Direct substitution gives

$$e_t^{(i)} = (I - K_t^{(i)} H) A_t^{(i)} \bar{e}_{t-1}^{(i)} + (I - K_t^{(i)} H)(A_t - A_t^{(i)}) x_{t-1} + \eta_t^{(i)}, \tag{12}$$

where $\eta_t^{(i)} = (I - K_t^{(i)} H) w_t - K_t^{(i)} v_t$. If products along the actual shared/synchronized state ancestry obey $\|M_{t:k+1}\|_2 \le C\mu^{t-k}$, $\mu < 1$, with bounded gains/states, unrolling bounds tracking error by decaying initialization error, accumulated model mismatch and filtered noise (Appendix F.6). Identification reduces the mismatch contribution; this separate stability condition does not guarantee an optimal selector.

5 Experiments

5.1 Protocol and baselines

Our primary evaluation targets abrupt changes: the original synthetic jump and an author-selected UAV braking segment. Gradual synthetic drift and the extended UAV task remain in the appendix. Synthetic benchmarks use 200 paired trials; the braking case uses 20 observation-noise trials. Observations are shared within a trial. Tables report mean $\pm$ SD; bold marks the lowest displayed mean, including ties. Curves show pointwise approximate 95% intervals for the Monte Carlo mean. Paired intervals quantify effect uncertainty separately from bold formatting. The UAV intervals concern repeated observations of one fixed maneuver, not independent flights.

Baselines are Fixed KF, STF, IMM, Optimized KF, RVFF-KF, and KalmanNet. STF/IMM now use the true UAV observation standard deviation, 0.05 m. A strengthened Optimized KF fits a full positive-definite 2×2 process covariance and scalar observation variance on separate flights; the original scalar-multiplier fit is retained as a control. KalmanNet uses the authors' GRU/FC architecture, three training restarts and validation-only checkpoint selection. Training, data splits, initialization,

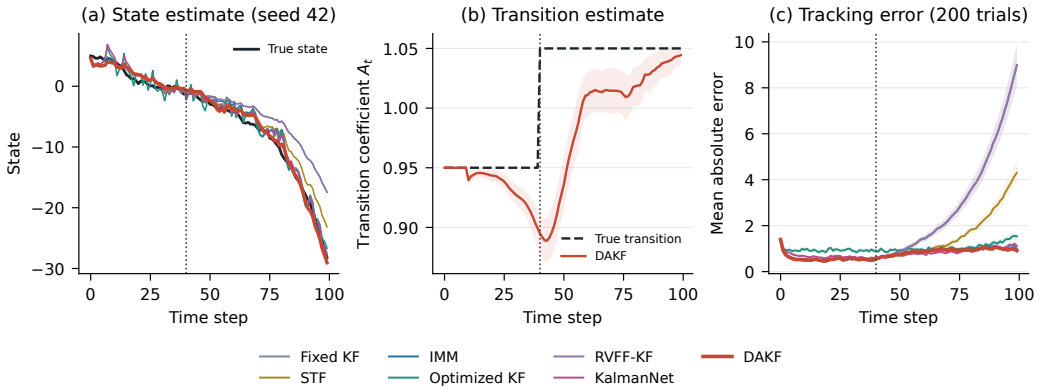

Figure 2: Synthetic abrupt change, preserving the original $0.95 \to 1.05$ transition at step 40. Left: a fixed test trajectory (seed 42). Middle: DAKF's prediction-time transition estimate across 200 trials. Right: mean state error for all methods. Shading denotes pointwise 95% Monte Carlo confidence intervals. The coefficient curve is an empirical estimate, not a confidence bound supplied by the theorem.

Table 1: Abrupt-change benchmarks, mean $\pm$ SD. Synthetic: 200 trials; selected UAV braking segment: 20 observation-noise trials. UAV units are cm. All methods share observations and initial states. KalmanNet averages errors over three training restarts.

Method	Synthetic MAE	Synthetic RMSE	UAV MAE	UAV RMSE
Fixed KF	2.042 ± 1.266	3.071 ± 2.051	6.369 ± 0.153	8.612 ± 0.141
STF	1.184 ± 0.565	1.653 ± 0.886	1.141 ± 0.066	1.435 ± 0.075
IMM	0.788 ± 0.144	1.005 ± 0.181	1.981 ± 0.120	2.498 ± 0.131
Optimized KF (scale)	1.009 ± 0.120	1.265 ± 0.155	1.182 ± 0.076	1.484 ± 0.092
Optimized KF (full)	–	–	1.033 ± 0.069	1.299 ± 0.078
RVFF-KF	2.035 ± 1.264	3.067 ± 2.051	6.258 ± 0.149	8.504 ± 0.148
KalmanNet	0.785 ± 0.125	0.993 ± 0.157	1.385 ± 0.062	1.737 ± 0.076
DAKF	$\mathbf{0.755 \pm 0.186}$	$\mathbf{0.948 \pm 0.236}$	$\mathbf{0.851 \pm 0.095}$	$\mathbf{1.070 \pm 0.123}$

covariance fitting and numerical safeguards are specified in Appendix A.5. No test flight or test score selects a model.

Scalar errors are $|\widehat{x}_t - x_t|$; extended 3D errors are position-error norms. MAE, RMSE and MaxAE exclude only the initial sample. Ground-truth transition errors are reported only for synthetic data. The original homoscedastic identification uses constant normalized precision, preserving ordinary RLS exactly.

5.2 Abrupt Changes: Synthetic Drift and UAV Braking

In the synthetic experiment, we simulate $T = 100$ scalar steps, $x_0 = 5$, $H = 1$, $Q = 0.1$, $R = 3$, and $A_t = 0.95$ before $t = 40$, then 1.05. DAKF uses $W = 10$, $\gamma_{\mathrm{RLS}} = 0.95$, and $\gamma_{\mathrm{regret}} = 0.9$. The near-unit dynamics generate persistent, dependent states; the post-change system is locally expansive, making retained-model mismatch visible. Figure 2 shows the first test seed, the mean identified coefficient used for prediction, and error across all 200 trials.

DAKF's synthetic abrupt MAE is 0.755, versus 0.788 for IMM; the paired difference is -0.033 with approximate 95% interval $[-0.053, -0.013]$. After the change, their MAEs are 0.888 and 0.936. Table 1 includes KalmanNet and the braking-case comparison.

On this fixed braking segment, DAKF has MAE 0.851 cm, versus 1.033 cm for full-covariance Optimized KF and 1.385 cm for KalmanNet. The comparison keeps the recorded motion, observation noise and initial state identical across methods. It tests recovery from the selected rapid change

For UAV braking, we retain the author-selected braking segment of the UZH-FPV `outdoor_forward_5` ground-truth trajectory (Delmerico et al., 2019), its Y coordinate from 10 to 14 seconds (2,001 samples), and independent Gaussian position-observation noise with standard deviation 5 cm, exactly as in the original benchmark. The segment captures rapid deceleration after forward motion. This uses real recorded motion with simulated observations; it is not an evaluation of raw visual-inertial odometry. All methods receive the same initial position and velocity computed from the first two ground-truth samples. This privileged initialization is retained for comparability and separately stress-tested below. Optimized KF fitting uses `outdoor_forward_3`, not the test flight.

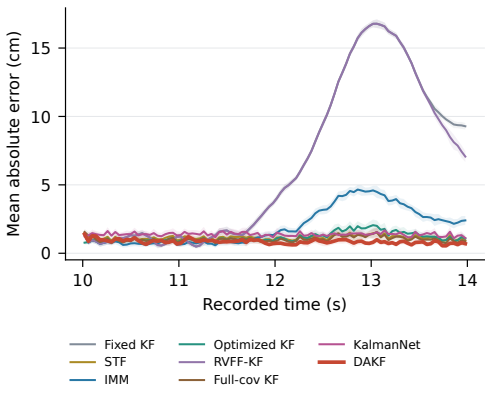

Figure 3: UAV braking case: mean absolute error and 95% confidence bands over 20 noise trials. Display uses 20-sample bins; metrics use full resolution.

Table 2: Architecture ablation, mean $\pm$ SD of state MAE. Synthetic/UAV: 200/20 paired trials; low-noise: 200 trials, $R = 0.03$ on $[40, 75)$. First two columns use the full trajectory; the last uses $t \geq 40$. All variants retain their parent arm settings; UAV uses the shared ground-truth initialization.

Variant	Synthetic	UAV (cm)	Low-noise, post-change
Fixed arm only	2.042 ± 1.266	6.429 ± 0.155	2.829 ± 2.071
Adaptive arm only	0.786 ± 0.174	0.856 ± 0.095	0.732 ± 0.249
Two arms, no sync	0.784 ± 0.202	0.887 ± 0.086	0.701 ± 0.252
DAKF	$\mathbf{0.755 \pm 0.186}$	$\mathbf{0.851 \pm 0.095}$	$\mathbf{0.641 \pm 0.186}$

rather than average performance over arbitrary flight regimes. The three KalmanNet restarts and all MAE/RMSE/MaxAE results are retained in Appendix A.6.

5.3 Does the dual-arm architecture help?

We remove the fixed arm, the adaptive arm, or synchronization while retaining the parent update rules and hyperparameters. Each single arm feeds back its own posterior; the no-sync variant retains both arms and selection. Table 2 uses the same paired observations as the parent experiments, including the low-noise setting below. This tests architectural components, not a different RLS tuning.

DAKF reduces synthetic abrupt MAE from 0.786 to 0.755 (3.9%) relative to the adaptive arm; the paired difference has approximate 95% interval $[-0.045, -0.017]$. Low-noise post-change MAE falls from 0.732 to 0.641 (12.4%). Removing synchronization increases error in both synthetic settings. On the UAV maneuver, MAEs are 0.851 versus 0.856 cm, with paired interval $[-0.0105, 0.00015]$ cm containing zero. The synthetic improvements therefore do not establish a practically meaningful architecture gain on this braking case. Appendix A.4 provides full traces.

5.4 Does low variance accelerate identification?

We retain the same synthetic plant, initial state, horizon, and DAKF selector. Only measurement variance changes: it is $R_t = 3$ outside the window $40 \leq t < 75$, and $R_{\text{low}} \in \{3, 0.3, 0.03\}$ inside. Process variance remains 0.1, and the Kalman arms retain nominal $R = 3$; the added adaptive quantity is the lagged regression-variance proxy. We use 200 new paired seeds (12000–12199), and also retain a delayed window $55 \leq t < 90$ in the appendix.

Reducing R_{low} from 3 to 0.03 lowers post-change parameter MAE from 0.0910 to 0.0725 (20.3%), state MAE from 0.846 to 0.641 (24.1%), and censored adaptation lag from 23.35 to 18.99 steps (18.7%). Recovery requires parameter error ≤ 0.05 for five consecutive predictions; failures are

censored at 60 steps. Lower noise also directly benefits filtering, so this comparison alone does not establish mediation through identification.

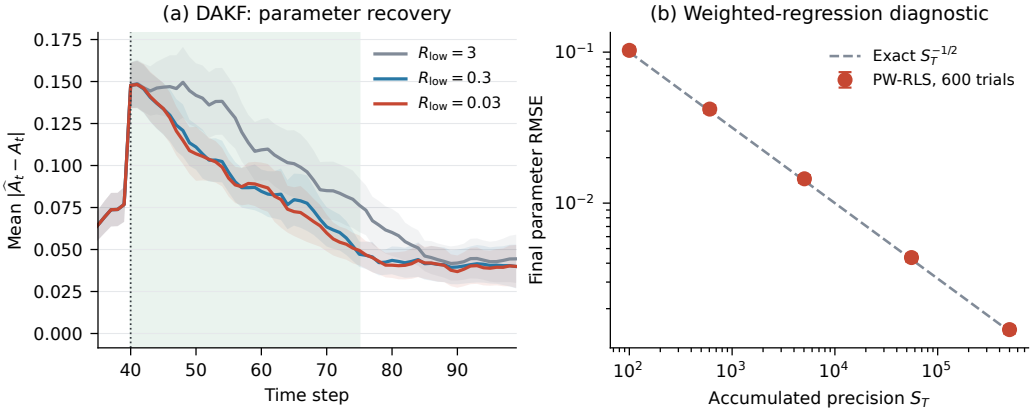

Figure 4: Low variance and identification information. Left: full DAKF with lagged variance estimates; mean parameter error and pointwise 95% confidence bands over 200 paired trials. Shading marks the low-noise window. Right: independent scalar regression (600 trials, $\gamma = 1$), comparing PW-RLS RMSE with exact $S_T^{-1/2}$ scaling; bars show Monte Carlo uncertainty.

Table 3: Low-variance availability in full DAKF (200 trials; mean $\pm$ SD). Errors use $t = 40, \ldots, 99$. Lag is censored at 60 steps; recovery requires five consecutive accurate predictions. Bold compares noise conditions, not methods.

R_{low}	Parameter MAE	State MAE	State RMSE	Censored lag	Recovery
3	0.091 ± 0.073	0.846 ± 0.265	1.029 ± 0.314	23.35 ± 16.93	92.0%
0.3	0.074 ± 0.052	0.672 ± 0.210	0.833 ± 0.252	19.57 ± 15.05	95.0%
0.03	$\mathbf{0.072 \pm 0.049}$	$\mathbf{0.641 \pm 0.186}$	$\mathbf{0.808 \pm 0.233}$	$\mathbf{18.99 \pm 13.81}$	**98.5%**

The scalar diagnostic uses $y_t = A + \epsilon_t$, a constant regressor and 600 realizations; low-noise standard deviations $1, 0.3, 0.1, 0.03, 0.01$ apply at steps 20–69. PW-RLS follows harmonic-precision scaling (Figure 4, right). Appendix E compares Eq. (11)'s exact variance, its $1/S_{T,\gamma}$ bound and uniform weights for $\gamma \in \{0.98, 0.95\}$.

Uniform weights give lower immediate low-noise parameter error, but estimated weights slightly improve state MAE; two-window estimated/uniform/oracle state MAEs are $0.676/0.696/0.698$. Appendix A.7 reports paired differences, the signed error decomposition and an AR(0.98)-regressor diagnostic. Calibrated precision improves diagnostic identification, while full-filter weighting gains remain mixed.

6 Discussion and Limitations

PW-PE and Proposition 2 distinguish precision, filtering bias and drift. The sliding variance proxy still requires calibration: its conditional-variance/proxy ratio reaches 10.65 at the audited 99th percentile. Finite ancestry-product measurements do not prove future stability, and recursive clipping in the extended UAV task lies outside the weighted-normal-equation theorem.

The retained UAV benchmark uses ground-truth initial position and velocity on one selected braking segment; noise repeats do not establish cross-flight generalization. Zero initial velocity causes the separate sensitivity failure in Appendix C. Synthetic architecture gains and the small, unresolved UAV gain should be distinguished. Gradual results and weighting reversals remain disclosed; calibrated scales, robust initialization and stability-preserving adaptation remain open.

Reproducibility Statement

The supplementary package includes the four preserved source programs, selected UAV ground-truth files with attribution, executable configurations, raw per-trial CSV/JSON metrics, state and parameter traces, and scripts for every figure and table. Training, validation, and test seeds are disjoint and recorded. The numerical precision audit checks causality, constant-weight equivalence, recursive-versus-batch estimates with replay, and selected-posterior moments. The accompanying README gives the execution sequence. All proofs and task-specific implementation details are included below.

Ethics Statement

This work uses synthetic processes and public UAV trajectories. It does not collect personal data or deploy a controller. The recorded data retain their original attribution and noncommercial share-alike license. The reported sensitivity to initialization and unmodeled noise should be considered before using these estimators in a safety-critical system.

Statement of AI Usage

AI assistance was used for drafting, code implementation, experimental organization, and mathematical checking. All reported measurements were generated by the supplied programs. The authors remain responsible for verifying the final manuscript, proofs, citations, and submission.

References

Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 24, 2011. URL https://proceedings.neurips.cc/paper/2011/hash/e1d5be1c7f2f456670de3d53c7b54f4a-Abstract.html.

Henk A. P. Blom and Yaakov Bar-Shalom. The interacting multiple model algorithm for systems with Markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, 1988. doi: 10.1109/9.1299.

Shanli Chen, Yunfei Zheng, Dongyuan Lin, Peng Cai, Yingying Xiao, and Shiyuan Wang. MAML-KalmanNet: A neural network-assisted Kalman filter based on model-agnostic meta-learning. *IEEE Transactions on Signal Processing*, 73:988–1003, 2025. doi: 10.1109/TSP.2025.3540018.

Geon Choi, Jeonghun Park, Nir Shlezinger, Yonina C. Eldar, and Namyoon Lee. Split-KalmanNet: A robust model-based deep learning approach for state estimation. *IEEE Transactions on Vehicular Technology*, 72(9):12326–12331, 2023. doi: 10.1109/TVT.2023.3270353.

Ori Cohen, Nir Shlezinger, and Tirza Routtenberg. EM-KalmanNet: Learned expectation-maximization for adaptive tracking in partially known, block-wise time-varying state-space models. *arXiv preprint arXiv:2608.24404*, 2026. URL https://arxiv.org/abs/2608.24404.

Jeffrey Delmerico, Titus Cieslewski, Henri Rebecq, Matthias Faessler, and Davide Scaramuzza. Are we ready for autonomous drone racing? the UZH-FPV drone racing dataset. In *IEEE International Conference on Robotics and Automation*, 2019. URL https://fpv.ifi.uzh.ch/.

Anubhab Ghosh, Antoine Honoré, and Saikat Chatterjee. DANSE: Data-driven non-linear state estimation of model-free process in unsupervised learning setup. *IEEE Transactions on Signal Processing*, 72:1824–1838, 2024. doi: 10.1109/TSP.2024.3383277.

Ido Greenberg, Netanel Yannay, and Shie Mannor. Optimization or architecture: How to hack Kalman filtering. In *Advances in Neural Information Processing Systems*, volume 36, pp. 50482–50505, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/9dfcc83c01e94d02c751c47517855c9f-Abstract.html.

Yangguang He, Wenhao Li, Minzhe Li, Juan Zhang, Xiangfeng Wang, and Bo Jin. Differentiable adaptive Kalman filtering via optimal transport. *arXiv preprint arXiv:2508.07037*, 2025. URL https://arxiv.org/abs/2508.07037.

Thomas Kailath, Ali H. Sayed, and Babak Hassibi. *Linear Estimation*. Prentice Hall, 2000. URL https://web.stanford.edu/~tkailath/cgi-bin/PubsList.php.

R. E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960. doi: 10.1115/1.3662552.

Minhyeok Ko and Abdollah Shafieezadeh. Cholesky-KalmanNet: Model-based deep learning with positive definite error covariance structure. *IEEE Signal Processing Letters*, 32:326–330, 2025. doi: 10.1109/LSP.2024.3519265.

Brian Lai and Dennis S. Bernstein. Adaptive Kalman filtering developed from recursive least squares forgetting algorithms. In *American Control Conference*, 2024a. doi: 10.23919/ACC60939.2024.10644929.

Brian Lai and Dennis S. Bernstein. Generalized forgetting recursive least squares: Stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 69(11):7646–7661, 2024b. doi: 10.1109/TAC.2024.3394351.

Lennart Ljung. *System Identification: Theory for the User*. Prentice Hall, second edition, 1999. URL https://rt.isy.liu.se/books/sysid/.

A. H. Mohamed and K. P. Schwarz. Adaptive Kalman filtering for INS/GPS. *Journal of Geodesy*, 73 (4):193–203, 1999. doi: 10.1007/s001900050236.

Hassan Mortada, Cyril Falcon, Yanis Kahil, Mathéo Clavaud, and Jean-Philippe Michel. Recursive KalmanNet: Deep learning-augmented Kalman filtering for state estimation with consistent uncertainty quantification. In *33rd European Signal Processing Conference*, pp. 885–889, 2025. URL https://eusipco2025.org/wp-content/uploads/pdfs/0000885.pdf.

Kumpati S. Narendra and Anuradha M. Annaswamy. Persistent excitation in adaptive systems. *International Journal of Control*, 45(1):127–160, 1987. doi: 10.1080/00207178708933715.

Xiaoyong Ni, Guy Revach, and Nir Shlezinger. Adaptive KalmanNet: Data-driven Kalman filter with fast adaptation. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 5970–5974, 2024. doi: 10.1109/ICASSP48485.2024.10447012.

Guy Revach, Nir Shlezinger, Xiaoyong Ni, Adria Lopez Escoriza, Ruud J. G. van Sloun, and Yonina C. Eldar. KalmanNet: Neural network aided Kalman filtering for partially known dynamics. *IEEE Transactions on Signal Processing*, 70:1532–1547, 2022. doi: 10.1109/TSP.2022.3158588.

Max Simchowitz, Horia Mania, Stephen Tu, Michael I. Jordan, and Benjamin Recht. Learning without mixing: Towards a sharp analysis of linear system identification. In *Proceedings of the 31st Conference on Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pp. 439–473, 2018. URL https://proceedings.mlr.press/v75/simchowitz18a.html.

Anastasios Tsiamis, Aren Karapetyan, Yueshan Li, Efe C. Balta, and John Lygeros. Predictive linear online tracking for unknown targets. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 48657–48694, 2024. URL https://proceedings.mlr.press/v235/tsiamis24a.html.

Lei Xin, George T.-C. Chiu, and Shreyas Sundaram. Online change points detection for linear dynamical systems with finite sample guarantees. *Automatica*, 171:111894, 2025. doi: 10.1016/j.automatica.2024.111894.

Oğuz Kaan Yüksel, Mathieu Even, and Nicolas Flammarion. Long-context linear system identification. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/9062b7d6e522dadf4f7d85d49b60d81e-Abstract-Conference.html.

Wenyi Zhang, Xiaoyong Ni, Nir Shlezinger, and Zengfu Wang. Change-aware self-adaptive AI-aided Kalman filters with neural change point detection. *arXiv preprint arXiv:2607.13387*, 2026. URL https://arxiv.org/abs/2607.13387.

Heyang Zhao, Tianyuan Jin, Weixin Wang, Vincent Tan, Pan Xu, and Quanquan Gu. Breaking the total variance barrier: Sharp sample complexity for linear heteroscedastic bandits with fixed action set. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/166cdd595aea8bf6cd8de4998d49f205-Abstract-Conference.html.

D. H. Zhou and P. M. Frank. Strong tracking filtering of nonlinear time-varying stochastic systems with coloured noise: Application to parameter estimation and empirical robustness analysis. *International Journal of Control*, 65(2):295–307, 1996. doi: 10.1080/00207179608921698.

Hongyi Zhou, Jingwei Li, and Jingzhao Zhang. Finite sample analyses for continuous-time linear systems: System identification and online control. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1072. URL https://proceedings.nips.cc/paper_files/paper/2025/hash/2df20ebc2bdaca5b8b48966cbe0ab3cf-Abstract-Conference.html.

A Experimental Details and Data Provenance

A.1 Preserved task implementations

The supplied `originals/` directory preserves the four existing programs. An AST-based adapter loads their definitions without running top-level plotting, platform-specific paths, or synthetic data fallbacks. The scalar tasks call `dual_arm_tiger`; the UAV tasks call `run_1d_dakf` and `dakf`. We use floating-point covariance arguments, avoiding integer-array truncation inside the synthetic IMM routine. No selector, majority-vote, or synchronization rule is replaced in the primary benchmarks.

Table 4: DAKF settings retained in the four task implementations. Losses differ across the existing benchmarks; these distinctions are not hidden behind a single canonical loss formula.

Setting	Synthetic abrupt	Synthetic gradual	UAV abrupt	UAV extended
Vote/replay window W	10	5	15	10
γ_{RLS}	0.95	0.98	0.98	0.96
γ_{regret}	0.90	0.92	0.90	0.90
Risk coefficient β	0	0.05	0.20	0.20
RLS initial P	100	100	100	10
Synchronization vote	Unanimous	Unanimous	> 0.8	> 0.9
Arm input states	Shared	Shared	Separate	Separate

Synthetic abrupt. The two filters predict from the previous selected state and covariance. Before step $t \geq W$, RLS replays the $W - 1$ completed transitions among selected states $t - W, \ldots, t - 1$. Synchronization checks whether the previous W choices are all Arm 2 and re-anchors the baseline before prediction; the next score comparison resets scores to current losses. Selection uses squared posterior measurement residuals, zero hysteresis, and Arm 1 during warm-up. Initial state is z_0, covariance is $R = 3$, and the initial transition is 0.95. The process uses $x_0 = 5$, $Q = 0.1$, and the noise generator first draws process noise, then measurement noise.

Synthetic gradual. The horizon is 200 steps, $x_0 = 0$, and A_t increases linearly from 0.95 to 1.05. The generator interleaves process and observation draws, retaining the source's random-number order. Both arms share the selected posterior. The selection loss is posterior measurement residual squared multiplied by $1 + \beta P_t^{(i)}$, with $\beta = 0.05$, hysteresis 0.01, and unanimous synchronization. The risk adjustment is a posterior-covariance multiplier, not a sample-variance penalty. Both Kalman arms retain $Q = 0.1, R = 3$.

UAV abrupt. The state is position and velocity, with $F_t = \left[\begin{smallmatrix} 1 & \Delta t \\ 0 & a_t \end{smallmatrix}\right]$ and position observation. Arm 1 and Arm 2 process covariances are respectively $\mathrm{diag}(0.005^2, 0.1^2)\overline{\Delta t}$ and $\mathrm{diag}(0.01^2, 0.5^2)\overline{\Delta t}$; both use $R = 0.05^2$. The adaptive arm identifies the velocity coefficient from completed estimated-velocity pairs. Loss is normalized innovation squared, with population variance over the recent loss window. The first ten steps use Arm 1. Synchronization copies Arm 2's state, covariance, and transition coefficient into Arm 1 and resets the scores and loss histories. The same initial position and velocity are supplied to all compared methods.

Extended UAV. The state contains three positions and three velocities. Initial position is the first observation and velocity is zero. The nominal process-noise scales are position 0.01, velocity 0.1; the adaptive arm uses position 0.05, velocity 2.5 in the source's integrated covariance construction. Each identified velocity coefficient is clipped to $[0.8, 1.2]$. Loss is the unnormalized squared innovation norm plus its recent sample-variance penalty. Early output alternates in two-step blocks over the first 20 steps. Synchronization copies states and covariances and assigns the identified coefficients to a baseline variable; the existing code does *not* insert that variable into Arm 1's transition matrix. We retain and disclose this limitation rather than silently fixing it while claiming the same benchmark. Thus the full parameter re-anchoring interpretation applies to the abrupt UAV implementation, whereas the extended-flight implementation transfers state and covariance only.

A.2 Real-data scope and splits

The real trajectories are the existing ground-truth files from the UZH-FPV dataset (Delmerico et al., 2019). The package records SHA-256 hashes of the copied files. Their columns are timestamp, three position coordinates, and four orientation coordinates; the current position-tracking task does not use orientation. Time is measured relative to the first timestamp before slicing. There is no smoothing, resampling, or trajectory modification. Position observations are constructed by adding independent Gaussian noise with standard deviation 0.05 m to each retained position coordinate. Raw camera, event, and inertial observations are not used to claim a sensor-fusion result.

Table 5: Recorded trajectories and their roles. Test runs repeat observation noise on the stated fixed trajectory.

Role	Sequence	Interval	Coordinates
Abrupt test	`outdoor_forward_5`	10–14 s; 2,001 samples	Y
Abrupt training/validation	`outdoor_forward_3`	10–14 s	Y
Extended test	`indoor_forward_6`	Full 30.198 s; 15,100 samples	X, Y, Z
Extended training/validation	`indoor_forward_7`	First 30.198 s	X, Y, Z

The source's "gradual" label for the extended flight describes its role in the benchmark; it does not certify a slowly drifting ground-truth transition matrix. Real motion can include rapid local maneuvers. We therefore refer to it as an extended trajectory when interpreting the results. Training and validation use independent observation-noise realizations on the same designated fitting flight; test uses another flight. This is a trajectory-level split for fitting versus testing, not a multi-flight population study.

A.3 Baseline implementations and fitting

The KF, STF, and IMM functions are loaded directly from the corresponding task sources. Their hyperparameters are not retuned on test realizations. In the synthetic IMM, the transition bank is $(0.95, 1.05)$ and model process covariances are $(0.05, 3)$ for abrupt and $(0.05, 0.3)$ for gradual. Switching matrices are $\left[\begin{smallmatrix} .98 & .02 \\ .10 & .90 \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} .95 & .05 \\ .08 & .92 \end{smallmatrix}\right]$, respectively. This endpoint information is unavailable to DAKF's adaptive identification. The archived UAV source assumes observation standard deviation 0.08 m for STF/IMM. We correct both to 0.05 m, equal to the data generator and KF/DAKF, and rerun the nominal and initialization-sensitivity comparisons. IMM likelihood weights use log normalization to avoid floating-point underflow under large initialization mismatch; its models and transition probabilities are unchanged. The superseded results remain only in the versioned artifact archive.

Optimized KF uses the supervised state-MSE objective motivated by Greenberg et al. (2023), restricted to positive multipliers of the source KF's nominal Q and R. A grid Q multiplier in $\{0.1, 1, 10, 100\}$ and R multiplier in $\{0.25, 1, 4\}$ supplies the three best training starts. Each start is refined by L-BFGS-B in log coordinates, with log-Q bounds $[-8, 10]$, log-R bounds $[-8, 8]$, 50 iterations, function tolerance 10^{-10} and gradient tolerance 10^{-9}. The candidate with smallest validation MSE is selected. Training seeds are 30000–30015 and validation seeds 40000–40015. Optimizer convergence flags and all candidates are retained in `results/tuning.json`. This two-parameter specialization is reproducible but less expressive than general matrix-valued covariance optimization.

RVFF-KF is our descriptive label for the robust variable-forgetting-factor covariance-inflation construction of Lai & Bernstein (2024a), Eqs. (30)–(31) and (36)–(40) of the conference paper and arXiv v1, with forgetting bounds $[0.5, 1]$, $K_\alpha = 2$, $K_\beta = 10$, and denominator safeguard 10^{-6}. The state dimension is one in the synthetic task and two per position/velocity axis in the UAV task. The three residual statistics start at one. Statistics updated with the completed posterior set the inflation for the next prediction, avoiding future access. Because UAV axes are uncoupled in the nominal model, this is an axis-wise scalar-observation implementation. It does not change the transition model and need not overcome persistent transition mismatch by covariance inflation alone. The generic implementation with forgetting fixed at one agrees with every original fixed-KF output to better than 10^{-8}; observed numerical differences are near floating-point precision.

The expanded comparison trains the official KalmanNet architecture and fits a full process covariance on disjoint flights (Appendix A.5). Scores from other protocols are not transferred into our

Table 6: Optimized KF: selected covariance multipliers and training/validation MSE. The same two-parameter family is optimized for every task; no test state enters fitting or selection. UAV fitting uses separate flights.

Task	Q multiplier	R multiplier	Training MSE	Validation MSE
synthetic abrupt	57.901	1.7271	1.7172	1.5641
uav abrupt	22026	252.5	0.00019871	0.00018908
synthetic gradual	98.392	4.0654	1.5734	1.4828
uav gradual	8582.5	21.855	0.00019103	0.00018855

Table 7: Additional architecture-ablation metrics (mean $\pm$ SD). Parameter error is reported only for synthetic data with known transition. Full/post-change regions match Table 2; all regions and adaptation-lag outcomes are retained in the raw metrics.

Variant	MAE	RMSE	MaxAE	Parameter MAE
Synthetic abrupt (all)				
Fixed arm only	2.042 ± 1.266	3.071 ± 2.051	9.163 ± 6.278	$\mathbf{0.061 \pm 0.000}$
Adaptive arm only	0.786 ± 0.174	0.989 ± 0.219	2.627 ± 0.686	0.094 ± 0.086
Two arms, no sync	0.784 ± 0.202	0.993 ± 0.268	2.614 ± 0.792	0.066 ± 0.040
DAKF	$\mathbf{0.755 \pm 0.186}$	$\mathbf{0.948 \pm 0.236}$	$\mathbf{2.452 \pm 0.666}$	0.066 ± 0.041
UAV maneuver (cm, all)				
Fixed arm only	6.429 ± 0.155	8.674 ± 0.142	17.606 ± 0.527	–
Adaptive arm only	0.856 ± 0.095	1.075 ± 0.124	$\mathbf{4.895 \pm 1.676}$	–
Two arms, no sync	0.887 ± 0.086	1.111 ± 0.103	5.246 ± 1.530	–
DAKF	$\mathbf{0.851 \pm 0.095}$	$\mathbf{1.070 \pm 0.123}$	$\mathbf{4.895 \pm 1.677}$	–
Low-noise synthetic (post-change)				
Fixed arm only	2.829 ± 2.071	3.643 ± 2.664	8.555 ± 6.252	0.100 ± 0.000
Adaptive arm only	0.732 ± 0.249	0.899 ± 0.281	2.096 ± 0.623	0.101 ± 0.090
Two arms, no sync	0.701 ± 0.252	0.910 ± 0.359	2.277 ± 0.912	0.073 ± 0.048
DAKF	$\mathbf{0.641 \pm 0.186}$	$\mathbf{0.808 \pm 0.233}$	$\mathbf{1.997 \pm 0.597}$	$\mathbf{0.072 \pm 0.049}$

tables. Theoretical work on identification, change detection, and control supplies context rather than additional filtering baselines.

A.4 Architecture Ablation

The four predeclared variants are the fixed arm alone, the adaptive arm alone, both arms with synchronization disabled, and full DAKF. We use 200 seeds 42–241 for the original synthetic abrupt task, 20 seeds 42–61 for the UAV maneuver, and 200 seeds 12000–12199 for the low-noise synthetic task. The latter retains $R = 3$ except on $40 \le t < 75$, where $R = 0.03$, and uses the same lagged residual-variance rule for every adaptive variant. The synthetic model, replay window, forgetting factor and prior are unchanged. Single-arm removal also removes selection and synchronization; the no-sync variant changes only synchronization.

In synthetic runs the adaptive arm alone regresses its own completed posterior transitions, using the parent's replay schedule. Before identification begins, its state update is identical to the fixed arm's because both use the same initial transition and covariances. The UAV adaptive arm alone retains its own velocity-history RLS and covariance settings from the first prediction; a single-arm filter has no two-arm warm-up selection. Each UAV arm keeps its native process covariance, $Q_1 = \mathrm{diag}(0.005^2, 0.1^2)\overline{\Delta t}$ or $Q_2 = \mathrm{diag}(0.01^2, 0.5^2)\overline{\Delta t}$. Consequently, the fixed-arm-only result is not identical to the standalone Fixed KF baseline, whose position process-noise setting differs. No ablation parameter is retuned.

The full variants exactly reproduce the previously saved parent state/error arrays; future-observation perturbations leave earlier outputs unchanged for all variants. Generated source programs, configuration, per-trial metrics and traces are included. In the low-noise post-change interval, DAKF's parameter MAE is 0.0725 versus 0.1015 for the adaptive arm alone. These results support the architecture in the evaluated synthetic settings. The UAV paired MAE interval contains zero for

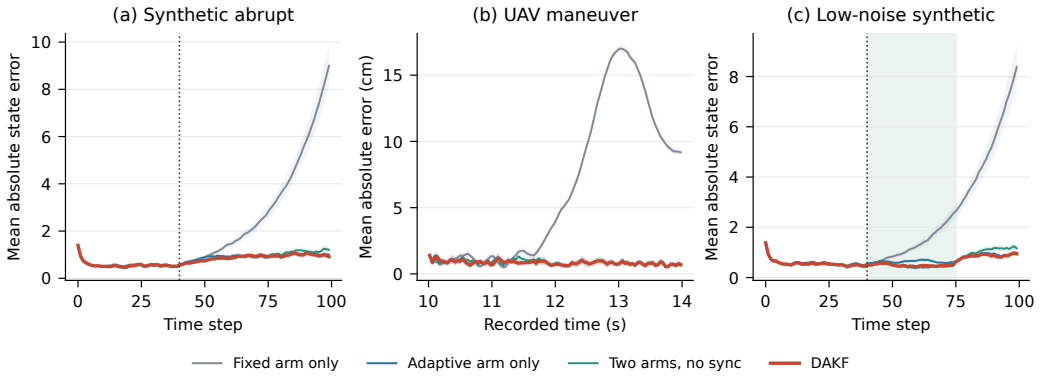

Figure 5: Architecture ablation on paired observations. All four variants are shown, including the near-overlap of DAKF and the adaptive arm on the UAV maneuver. Bands are pointwise approximate 95% Monte Carlo confidence intervals. UAV errors are displayed in 20-sample bins; metrics use all samples. The green interval marks the low-noise window.

DAKF versus the adaptive arm; it would be incorrect to infer a clear gain there from the smallest displayed mean alone.

A.5 Expanded comparison: fair covariances, splits, and learned filtering

Covariance correction and numerical normalization. In the original maneuver and observed-position/zero-velocity sensitivity experiment, we rerun STF and IMM with $R = 0.05^2$, replacing the archived 0.08^2. KF and DAKF are unchanged. The IMM likelihood update is evaluated in log space and normalized by log-sum-exp: direct exponentiation underflowed for some zero-velocity trials. This numerical repair preserves the two transition models and their mixing probabilities. No failed trial is silently dropped. The original four-task data and superseded scores remain versioned in the artifact; all current tables and plots use the corrected results.

Data splits and braking-case scope. Training uses `outdoor_forward_3` and `indoor_forward_7`; validation uses `indoor_forward_3`. The test case is the original author-selected 10–14 s Y-axis segment from `outdoor_forward_5`, chosen to study abrupt braking after forward motion. This fixed segment precedes the present baseline expansion and is not reselected for each method. Its twenty observation-noise realizations (seeds 42–61, standard deviation 0.05 m) are generated after slicing, preserving the original task. Test states do not enter covariance fitting, gradient updates or checkpoint selection. The real trajectory is combined with simulated position observations; no raw IMU/camera fusion is claimed.

Initialization. Every tested method receives the same initial ground-truth position and velocity in the primary braking comparison. The observed-position/zero-velocity sensitivity test remains in Appendix C. Training data for the added learned and full-covariance baselines retain their original mixture of true-state and causal initializations: the latter uses a least-squares line fitted to noisy positions strictly within the preceding 0.2 s. This training convention is unchanged; it does not provide future test information. Test metrics exclude the initial sample.

Full-covariance Optimized KF. In addition to the original scalar-multiplier implementation, we fit the entire symmetric positive-definite 2×2 process covariance and scalar observation variance. The diagonal entries use positive exponentials and the correlation is $0.99 \tanh(c)$, guaranteeing positive definiteness. Four unconstrained coordinates control the two diagonal scales, correlation and R. Three fixed starts use L-BFGS-B with at most 70 iterations each; log-process scales lie in $[-10, 14]$, correlation coordinate in $[-3, 3]$, and log-observation scale in $[-8, 8]$. Training uses eight noise seeds per training flight and both initializers on the 10–14 s interval. Four independent noise seeds on the validation flight select the candidate. The objective is position MSE, excluding the first sample. Full parameter values, objectives, convergence flags and iteration counts are saved. The three starts converge to similar held-out objectives; test data do not select the result.

Table 8: Original abrupt tasks with corrected covariances and added baselines. Mean ± SD across 200 synthetic or 20 UAV noise trials. KalmanNet errors are averaged over three training restarts within each noise trial.

Method	MAE	RMSE	MaxAE
Synthetic abrupt			
Fixed KF	2.042 ± 1.266	3.071 ± 2.051	9.163 ± 6.278
STF	1.184 ± 0.565	1.653 ± 0.886	4.856 ± 2.769
IMM	0.788 ± 0.144	1.005 ± 0.181	2.877 ± 0.654
Optimized KF (scale)	1.009 ± 0.120	1.265 ± 0.155	3.497 ± 0.703
RVFF-KF	2.035 ± 1.264	3.067 ± 2.051	9.170 ± 6.269
KalmanNet	0.785 ± 0.125	0.993 ± 0.157	2.833 ± 0.650
DAKF	$\mathbf{0.755 \pm 0.186}$	$\mathbf{0.948 \pm 0.236}$	$\mathbf{2.452 \pm 0.666}$
Original UAV maneuver (cm)			
Fixed KF	6.369 ± 0.153	8.612 ± 0.141	17.433 ± 0.532
STF	1.141 ± 0.066	1.435 ± 0.075	5.547 ± 1.262
IMM	1.981 ± 0.120	2.498 ± 0.131	6.359 ± 1.034
Optimized KF (scale)	1.182 ± 0.076	1.484 ± 0.092	4.939 ± 0.647
Optimized KF (full)	1.033 ± 0.069	1.299 ± 0.078	5.238 ± 1.549
RVFF-KF	6.258 ± 0.149	8.504 ± 0.148	17.433 ± 0.532
KalmanNet	1.385 ± 0.062	1.737 ± 0.076	5.951 ± 0.529
DAKF	$\mathbf{0.851 \pm 0.095}$	$\mathbf{1.070 \pm 0.123}$	$\mathbf{4.895 \pm 1.677}$

KalmanNet implementation and training. We use the authors' batched architecture #2 from https://github.com/KalmanNet/KalmanNet_TSP, commit `828a2cf529bc84f43b37d543d916fe5858054457`. Its three GRUs and seven fully connected blocks are unchanged. Task-specific data preparation and training/evaluation drivers are supplied. This is an official-architecture baseline under our protocol, not a reproduction of the original paper's benchmark scores. The input/output multipliers are 5/40, batch size 20, Adam learning rate 10^{-3}, weight decay 10^{-4} and gradient-norm clipping threshold 10. Three independent restarts use seeds 71001–71003; each runs 1,000 optimizer steps and validates every 50 steps. The lowest validation-MSE checkpoint within each restart is retained; no best test restart is chosen.

The recursive gain-output layer starts with zero weights and a nominal steady-state KF-gain bias. This task-specific initialization avoids numerical divergence observed in pre-test random-initialization pilots; the architecture and subsequent learned gain remain unconstrained. Pilot logs are retained. The scalar nominal transition is 0.95. For UAV, the nominal model is constant velocity with 0.002 s sampling; subtracting the initial position and expressing velocity as displacement per nominal sample improves coordinate scaling. This invertible transformation uses the shared initial estimate only. Training loss is position MSE, matching the reported target.

Synthetic training/validation sets have 256/20 independent sequences (seeds starting at 31000/41000), disjoint from the 200 test seeds. UAV training/validation use 256/20 four-second windows sampled on the designated flights, with fresh observation noise; half use true-state initialization and half the causal initializer. The same pool is used across restarts. UAV training uses truncated backpropagation in 100-sample blocks, carrying and detaching recurrent states between blocks; it resets at the end of a sequence or after validation. Validation and testing unroll the complete sequence. Checkpoints, every recorded training/validation loss, exact data arrays, window provenance and software versions are included. Table entries average per-trial errors over the three trained networks, not predictions from a test-time ensemble. Noise SD and training-restart variability are reported separately.

A.6 Expanded results and statistical interpretation

Uncertainty unit. The original synthetic and braking comparisons use paired Monte Carlo realizations. UAV confidence intervals quantify variability from observation noise on this one fixed motion sequence. Training restarts are averaged within each noise trial and separately listed, rather than counted as independent flights. Bold formatting indicates the displayed minimum; it is distinct from a significance test.

Table 9: All KalmanNet training restarts. Each checkpoint minimizes validation MSE within its restart. Test MAE uses the original synthetic task or fixed UAV braking segment (cm), not an ensemble prediction.

Task	Train seed	Best step	Validation MSE	Test MAE
Synthetic	71001	950	1.01254	0.797
Synthetic	71002	800	1.03119	0.803
Synthetic	71003	1000	0.909686	0.756
UAV braking	71001	150	0.000244057	1.241
UAV braking	71002	150	0.000352321	1.504
UAV braking	71003	450	0.000301119	1.411

The paired DAKF-minus-adaptive-arm interval on the braking segment contains zero, despite DAKF's lower mean. Thus the favorable comparison against other baselines does not independently establish a dual-arm gain on this segment. The observed-position/zero-velocity failure remains reported. These limits apply to the selected-case evidence, even though the main task is intentionally focused on abrupt changes.

A.7 Connecting the component theory to the executed filter

Exact decomposition on actual filtered regressors. We audit all 200 saved two-window full-DAKF trajectories for each weight choice. The saved conditional selected-posterior means and actual filtered states define b_s and ξ_s in Eq. (5). Aggregating replay coefficients exactly as in Appendix F.3, we compute the signed centered-noise, conditional-bias, drift and prior contributions in Eq. (14), using the actual V_T. Their sum reproduces the prediction-time coefficient error to below 10^{-9} throughout every audited recursion. These are oracle diagnostics of existing full-filter runs, not quantities assumed available online.

At the final horizon, the average absolute noise/bias contributions are approximately $0.0166/0.00946$ (uniform), $0.0146/0.00920$ (oracle) and $0.0179/0.0105$ (estimated). Signed terms can cancel, so separate absolute contributions must not be summed and called the observed error. Weighting changes the filtered regressors, the selected state and the conditional bias, in addition to changing precision. This explains why a centered-regression variance improvement does not imply a deterministic ranking of coupled-filter errors.

Excitation and calibration. We save G_T/S_T for each actual scalar design. At the final horizon its minimum across trials is approximately 0.0595–0.0653, while its median is much larger; realized excitation is highly variable. An observed positive Gram value is evidence about the sampled design, not a lower-bound guarantee for future runs. For estimated scales, conditional variance divided by the proxy has median 0.458, 90th percentile 3.48, 99th percentile 10.65 and maximum 193.24. These ratios already rule out treating the uninflated proxy as uniformly conservative on the tested paths. They do not estimate a distribution-wide sub-Gaussian constant: conditional variance alone is insufficient after hard selection. Theorem 1 consequently remains conditional on a valid noise-scale calibration.

Ancestry products on the actual scalar recursion. We reconstruct the selected Kalman gain and $M_t = (1 - K_t)A_t^{\text{selected}}$ from the saved choices and identified coefficients. Reconstructed states agree exactly with the saved full-DAKF states. For every trial and horizon, we enumerate all products along the shared selected-state ancestry. Individual absolute factors can exceed one (maximum 1.184 after rounding), whereas the largest ten-step product across all modes/trials is 0.260 after rounding and the largest forty-step product is below 0.001. Thus the measured scalar paths contract over blocks despite occasional single-step expansion. This finite-horizon audit does not certify the $C\mu^{t-k}$ condition on arbitrary future paths, nor does it transfer that condition to the separate-state UAV implementation.

Dependent-regressor diagnostic. To test the theorem's non-independent design setting beyond the earlier constant-regressor example, let $\phi_t = 0.98\phi_{t-1} + \sqrt{1 - 0.98^2}u_t$ with stationary Gaussian initialization, and $y_t = A_t\phi_t + \epsilon_t$. Innovations u_t and observation noises are independent. We use 300 samples and 500 paired seeds 92000–92499. The coefficient is either constant 0.7 or changes

Table 10: Dependent-regressor diagnostic: late parameter MAE on the fixed interval $[200, 240)$, mean $\pm$ SD over 500 paired trials. All dynamics, forgetting factors and noise schedules are reported. This is regression identification, not a full-DAKF benchmark.

Dynamics	γ	Noise window	Uniform	Oracle	Estimated
Constant	1	Homoscedastic	$\mathbf{0.063 \pm 0.048}$	$\mathbf{0.063 \pm 0.048}$	0.066 ± 0.049
Constant	1	Immediate	0.051 ± 0.041	$\mathbf{0.011 \pm 0.009}$	0.013 ± 0.011
Constant	1	Delayed	0.053 ± 0.045	$\mathbf{0.013 \pm 0.011}$	0.018 ± 0.016
Constant	0.98	Homoscedastic	$\mathbf{0.088 \pm 0.060}$	$\mathbf{0.088 \pm 0.060}$	0.093 ± 0.064
Constant	0.98	Immediate	0.024 ± 0.023	$\mathbf{0.012 \pm 0.009}$	0.014 ± 0.011
Constant	0.98	Delayed	0.031 ± 0.031	$\mathbf{0.014 \pm 0.010}$	0.019 ± 0.014
Abrupt	1	Homoscedastic	0.316 ± 0.127	0.316 ± 0.127	$\mathbf{0.315 \pm 0.128}$
Abrupt	1	Immediate	0.316 ± 0.118	$\mathbf{0.020 \pm 0.020}$	0.042 ± 0.048
Abrupt	1	Delayed	0.317 ± 0.120	$\mathbf{0.027 \pm 0.029}$	0.067 ± 0.069
Abrupt	0.98	Homoscedastic	$\mathbf{0.151 \pm 0.106}$	$\mathbf{0.151 \pm 0.106}$	0.153 ± 0.106
Abrupt	0.98	Immediate	0.123 ± 0.089	$\mathbf{0.012 \pm 0.009}$	0.017 ± 0.013
Abrupt	0.98	Delayed	0.125 ± 0.094	$\mathbf{0.014 \pm 0.011}$	0.022 ± 0.017

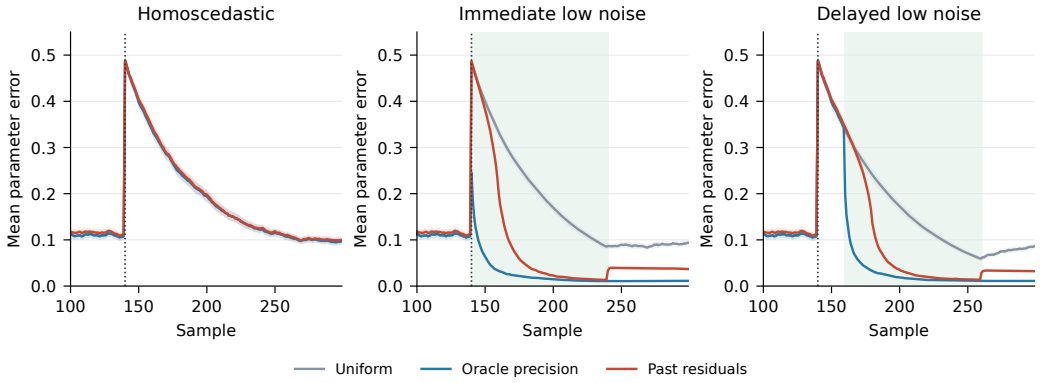

Figure 6: Separate identification diagnostic with AR(0.98) regressors and an abrupt coefficient change. The three panels show homoscedastic, immediate low-noise and delayed low-noise schedules, all with $\gamma = 0.98$, 500 paired trials and pointwise 95% Monte Carlo intervals. All schedules were fixed before this diagnostic was run; complete constant/abrupt and $\gamma = 1/0.98$ results are retained. This experiment removes endogenous filtered-state bias and does not operate the DAKF selector.

from 0.7 to 1.2 at sample 140. Observation standard deviation is one, with either no change, a 0.1 interval $[140, 240)$, or a delayed interval $[160, 260)$. Ordinary and precision-weighted RLS use $\gamma \in \{1, 0.98\}$ and ridge precision 0.001; estimated weights use only the preceding 20 completed prediction residuals, with initial variance one and clipping to $[0.001, 100]$. The common late interval $[200, 240)$ lies within both low-noise schedules and was fixed in the configuration before evaluation. Final-time error is also retained, including deterioration after low-noise data are forgotten.

For abrupt dynamics, immediate low noise and $\gamma = 0.98$, late parameter MAE is approximately 0.123 with uniform weights, 0.0122 with oracle precision and 0.0167 with estimated precision. Under homoscedastic noise, oracle and uniform are identical; estimated weights need not help. All factorial outcomes are reported. The effective precision is $S_{T,\gamma} = \sum_t \gamma^{T-t}/\sigma_t^2$ for oracle weights; the actual Gram also contains ϕ_t^2, as required by PW-PE. This isolates centered regression noise and predictable correlated inputs. Gaussian regressors have no deterministic uniform amplitude bound, so the diagnostic examines realized information trends rather than certifying every theorem condition. It does not show that the current full filter has solved its bias/calibration problem.

Table 11: Gradual synthetic drift and extended UAV tracking, mean ± SD. Synthetic: 200 realizations; UAV: 20 noise realizations on one extended recorded trajectory. UAV errors are Euclidean 3D position errors in centimeters.

Method	Synthetic gradual			Extended UAV (cm)		
	MAE	RMSE	MaxAE	MAE	RMSE	MaxAE
Fixed KF	1.492 ± 0.821	2.475 ± 1.583	9.955 ± 6.845	17.69 ± 0.04	20.19 ± 0.03	43.39 ± 0.52
STF	0.922 ± 0.105	1.202 ± 0.122	4.232 ± 0.802	2.81 ± 0.03	3.07 ± 0.03	10.72 ± 1.87
IMM	$\mathbf{0.631 \pm 0.082}$	$\mathbf{0.801 \pm 0.103}$	2.600 ± 0.609	$\mathbf{2.19 \pm 0.02}$	2.39 ± 0.03	9.66 ± 2.73
Optimized KF	0.938 ± 0.105	1.187 ± 0.148	3.713 ± 0.859	$\mathbf{2.19 \pm 0.03}$	$\mathbf{2.38 \pm 0.03}$	$\mathbf{7.85 \pm 1.58}$
RVFF-KF	1.489 ± 0.820	2.473 ± 1.584	9.953 ± 6.847	6.70 ± 0.10	9.41 ± 0.19	38.77 ± 1.96
DAKF	0.645 ± 0.113	0.815 ± 0.143	$\mathbf{2.428 \pm 0.593}$	2.39 ± 0.04	2.63 ± 0.05	10.90 ± 2.04

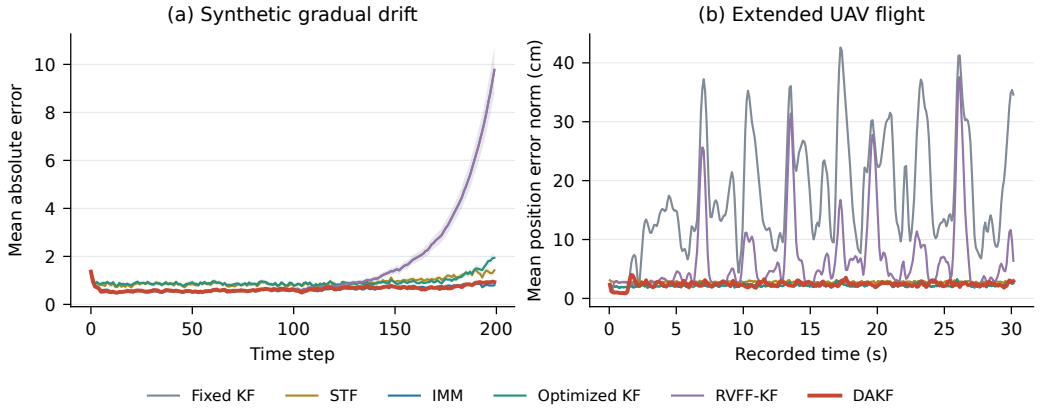

Figure 7: Complete gradual-setting evaluation. Left: linearly drifting synthetic transition, 200 realizations. Right: full 3D position error on `indoor_forward_6`, 20 measurement-noise realizations. Bands are pointwise 95% confidence intervals for the mean. The UAV plot averages consecutive 50-sample bins for display; tables use all samples. DAKF does not attain the lowest mean MAE in either setting.

B Gradual and Extended-Trajectory Results

DAKF improves substantially over the nominal KF on both gradual settings but does not dominate tuned or multiple-model alternatives. Its synthetic MAE exceeds IMM by 0.0140, with paired approximate 95% interval $[0.0019, 0.0260]$. On the extended UAV trajectory its MAE exceeds Optimized KF by 0.00196 m, with interval $[0.00181, 0.00212]$ m. These differences are computed on paired noise realizations; the intervals quantify Monte Carlo noise, not uncertainty across unknown future flights. The retention/adaptation split is most compelling when a previous regime becomes abruptly obsolete. More aggressive forgetting is not guaranteed to help gradual drift, because it can also increase identification variance and trigger unnecessary switching.

C Initialization Sensitivity

The original maneuver benchmark initializes all methods from the first two ground-truth positions. Table 12 instead sets $\widehat{p}_0 = z_0$ and $\widehat{v}_0 = 0$ for all methods, retaining the same 20 noise seeds and previously fitted covariances. Its DAKF failure is reproduced exactly from the original implementation; it is not a transcription or unit-conversion error.

To distinguish initial position error from velocity mismatch, we cross ground-truth versus observed initial position with ground-truth versus zero initial velocity, keeping DAKF unchanged. With the correct initial velocity, either position source gives approximately 0.00851 m MAE. With zero initial velocity, either position source gives approximately 1.96 m MAE. Excluding the first 0.5 seconds does

Table 12: Observation-only initialization sensitivity on the same UAV maneuver: $\widehat{p}_0 = z_0$, $\widehat{v}_0 = 0$. All values in meters, mean $\pm$ SD over 20 paired noise realizations. Existing hyperparameters remain unchanged.

Method	MAE	RMSE	MaxAE
Fixed KF	0.074 ± 0.001	0.099 ± 0.001	0.318 ± 0.009
STF	$\mathbf{0.013 \pm 0.001}$	$\mathbf{0.018 \pm 0.001}$	$\mathbf{0.132 \pm 0.013}$
IMM	0.029 ± 0.001	0.052 ± 0.001	0.310 ± 0.009
Optimized KF	0.024 ± 0.001	0.052 ± 0.001	0.292 ± 0.009
RVFF-KF	0.073 ± 0.001	0.098 ± 0.002	0.318 ± 0.009
DAKF	1.957 ± 0.049	2.134 ± 0.044	3.014 ± 0.047

Table 13: DAKF initialization audit on the same maneuver, 20 paired noise trials per condition. Values are in meters. Late MAE excludes the first 0.5 seconds. No filtering or identification hyperparameter is changed.

Initial position	Initial velocity	MAE	Late MAE
Observed	Ground truth	$\mathbf{0.00851 \pm 0.00095}$	$\mathbf{0.00823 \pm 0.00090}$
Observed	Zero	1.95669 ± 0.04936	2.05106 ± 0.05469
Ground truth	Ground truth	$\mathbf{0.00851 \pm 0.00095}$	$\mathbf{0.00823 \pm 0.00090}$
Ground truth	Zero	1.95547 ± 0.05227	2.04970 ± 0.05796

not remove the failure (late MAE is approximately 2.05 m). Initial velocity mismatch is therefore the dominant factor among these tested interventions, rather than simply the initial position observation noise or a short startup transient. The audit does not prove failure under every imperfect velocity initialization.

The runs also record the identified coefficient and synchronization history for diagnosis; early velocity transients can interact with recursive identification and subsequent synchronization. Designing and independently validating a velocity warm start or robust identification constraint remains future work. We retain the sensitivity result because it materially limits interpretation of the favorable maneuver benchmark.

D Variance Experiments and Ablations

D.1 Low-variance windows on the unchanged synthetic plant

All variants retain $T = 100$, $A_t = 0.95$ before 40 and 1.05 afterward, $x_0 = 5$, $Q = 0.1$, the original Kalman gains computed with nominal $R = 3$, and the same selector and replay logic. Measurement variance alone is changed inside either $[40, 75)$ or $[55, 90)$, using $R_{\text{low}} \in \{3, 0.3, 0.03\}$. Seeds 12000–12199 share identical process and standardized observation-noise draws across schedules and weight choices. The identical $R_{\text{low}} = 3$ control appears under both timing labels and must not be counted twice as independent evidence.

The table labels Known variance, Past residuals, and Uniform denote the following identification-weight choices: (i) conditional oracle variance, (ii) lagged estimated variance, and (iii) uniform precision. With $\phi_{s-1} = \widehat{x}_{s-1}$, the estimated proxy is the sample variance of completed residuals $\widehat{x}_j - \widehat{A}_j^{\text{pred}} \widehat{x}_{j-1}$ for $j = \max(1, s-20), \ldots, s-1$. At least five residuals are required; otherwise the proxy is 0.1. The sample variance uses denominator $m - 1$, is clipped to $[10^{-5}, 10^3]$, and its value is frozen at the original physical sample time. Identification uses normalized precision $0.1/\widehat{\sigma}^2_{\epsilon,s}$, making uniform precision exactly one.

Every weight/timing combination is reported in Table 14. Uniform precision remains competitive in parameter error, including the smallest parameter MAE for the immediate $R_{\text{low}} = 0.03$ case. Estimated precision has the smallest state MAE in that case, but not for every noise level. This differs from clean exogenous regression because the weighted estimates alter the filter and selector, which in turn alter subsequent regression inputs and bias. Lower measurement variance does not imply that every oracle-variance weighting choice must improve every full-filter metric.

Table 14: All low-variance-window outcomes, including all three weight choices and both timings. Each cell is mean ± SD over 200 paired trials. Post-change parameter and state MAE and right-censored lag are reported. The repeated $R_{\text{low}} = 3$ rows are the same control, not independent experiments.

Timing	R_{low}	Weights	Parameter MAE	State MAE	Lag
Immediate	3	**Known variance**	$\mathbf{0.089 \pm 0.073}$	0.860 ± 0.272	23.40 ± 18.04
Immediate	3	**Past residuals**	0.091 ± 0.073	$\mathbf{0.846 \pm 0.265}$	$\mathbf{23.35 \pm 16.93}$
Immediate	3	**Uniform**	$\mathbf{0.089 \pm 0.071}$	0.865 ± 0.274	24.70 ± 17.71
Immediate	0.3	**Known variance**	$\mathbf{0.071 \pm 0.048}$	0.683 ± 0.217	19.80 ± 16.41
Immediate	0.3	**Past residuals**	0.074 ± 0.052	$\mathbf{0.672 \pm 0.210}$	19.57 ± 15.05
Immediate	0.3	**Uniform**	$\mathbf{0.071 \pm 0.049}$	0.680 ± 0.210	$\mathbf{18.77 \pm 14.98}$
Immediate	0.03	Known variance	0.070 ± 0.044	0.646 ± 0.202	20.00 ± 16.56
Immediate	0.03	**Past residuals**	0.072 ± 0.049	$\mathbf{0.641 \pm 0.186}$	18.99 ± 13.81
Immediate	0.03	**Uniform**	$\mathbf{0.068 \pm 0.046}$	0.646 ± 0.190	$\mathbf{18.50 \pm 14.91}$
Delayed	3	**Known variance**	$\mathbf{0.089 \pm 0.073}$	0.860 ± 0.272	23.40 ± 18.04
Delayed	3	**Past residuals**	0.091 ± 0.073	$\mathbf{0.846 \pm 0.265}$	$\mathbf{23.35 \pm 16.93}$
Delayed	3	**Uniform**	$\mathbf{0.089 \pm 0.071}$	0.865 ± 0.274	24.70 ± 17.71
Delayed	0.3	**Known variance**	$\mathbf{0.077 \pm 0.053}$	0.646 ± 0.223	21.68 ± 16.58
Delayed	0.3	**Past residuals**	0.080 ± 0.057	$\mathbf{0.634 \pm 0.196}$	$\mathbf{21.59 \pm 15.01}$
Delayed	0.3	**Uniform**	$\mathbf{0.077 \pm 0.053}$	0.651 ± 0.212	22.36 ± 15.75
Delayed	0.03	**Known variance**	0.076 ± 0.050	0.616 ± 0.224	$\mathbf{20.61 \pm 14.91}$
Delayed	0.03	**Past residuals**	0.079 ± 0.056	$\mathbf{0.605 \pm 0.189}$	21.02 ± 14.26
Delayed	0.03	**Uniform**	$\mathbf{0.075 \pm 0.051}$	0.620 ± 0.203	21.23 ± 14.40

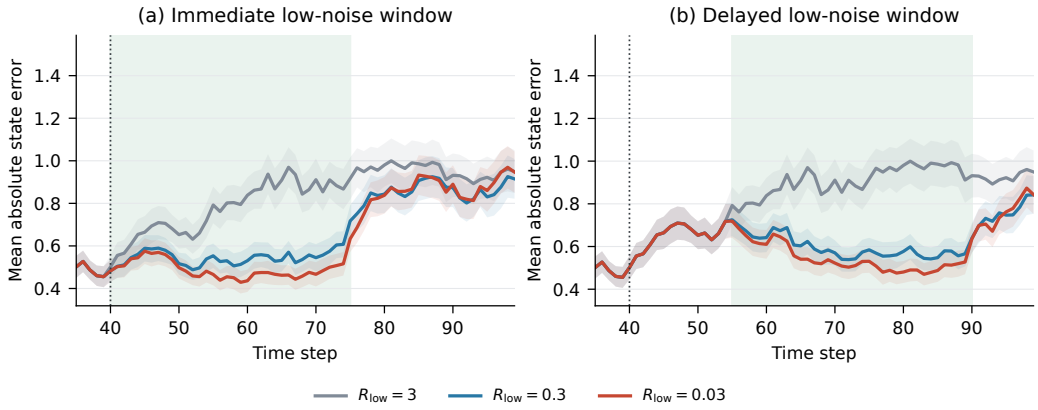

Figure 8: Full DAKF state tracking with the estimated-variance proxy, for both low-noise-window timings. Each curve is the mean over 200 paired realizations with pointwise 95% confidence bands. Shaded intervals mark the reduced measurement variance. Lower-noise observations can benefit both the direct filter update and subsequent identification, so this figure does not establish a purely parameter-mediated effect.

D.2 A SEPARATE TWO-WINDOW ABLATION

We also run 200 independent seeds (10000–10199) with low measurement variance 0.03 during $20 \leq t < 35$ and $55 \leq t < 85$, and variance 3 elsewhere. This includes a low-noise interval before and after the transition. Table 15 reports all weight choices. Compared with uniform weighting, oracle weighting slightly improves parameter MAE but not state MAE; the estimated proxy improves state MAE but worsens parameter MAE. These mixed outcomes are useful evidence about the gap between the centered-regression guarantee and the coupled full filter.

Table 15: Two-window variance ablation: $R = 0.03$ during $20 \le t < 35$ and $55 \le t < 85$, otherwise $R = 3$. Mean $\pm$ SD over 200 trials; all errors evaluated after $t = 40$. The entire DAKF architecture is active.

Weights	Parameter MAE	State MAE	State RMSE	State MaxAE
Known variance	$\mathbf{0.071 \pm 0.049}$	0.698 ± 0.247	0.858 ± 0.281	2.033 ± 0.631
Past residuals	0.079 ± 0.061	$\mathbf{0.676 \pm 0.200}$	$\mathbf{0.833 \pm 0.234}$	$\mathbf{1.990 \pm 0.573}$
Uniform	0.073 ± 0.053	0.696 ± 0.216	0.849 ± 0.247	2.008 ± 0.581

D.3 Conditional oracle variance and causality

At time s, both predicted states p_i, Kalman gains k_i, and past scores are known. The scalar corrected state in arm i is $h_i(z) = k_i z + (1 - k_i)p_i$. Comparing its squared posterior residual losses makes the hard-selection region a quadratic inequality in z. The selected state is therefore piecewise affine. We compute its first two conditional moments by integrating affine and quadratic polynomials over the Gaussian intervals defined by the score-comparison roots. The observation's predictive distribution is obtained from the known generative model conditioned on past observations, with deterministic true $x_0 = 5$. No realized current state or future observation is used in the weights.

This calculation provides the exact conditional variance of the selected regression response; subtracting $A_s\phi_{s-1}$ does not change conditional variance because the regressor is already known. It does not remove the response's conditional bias. Nor does conditional variance certify its sub-Gaussian moment-generating-function scale after hard selection. The main theorem therefore applies with this proxy only under the stated additional calibration condition, not merely because the diagnostic variance is exact.

Perturbing observations from time 50 onward leaves all variance weights through time 50 unchanged, for both oracle and estimated modes. Uniform weighting reproduces the preserved source to numerical precision. Across the three modes, the empirical mean of squared centered selected-state error divided by its oracle variance is between 0.989 and 0.993. Independent numerical quadrature checks three different selection configurations, with maximum moment discrepancy below 10^{-7}. These checks validate the implemented conditional moments and causal order; they do not prove a distribution-wide calibration constant for estimated variance.

E Independent Harmonic-Precision Diagnostic

The diagnostic is deliberately simpler than DAKF: $\phi_{t-1} = 1$, $y_t = 1.05 + \epsilon_t$, and independent Gaussian ϵ_t with known variance. There are 100 samples per realization, 600 seeds (20000–20599), and standard deviation one outside $20 \le t < 70$. Inside, standard deviation is one of $1, 0.3, 0.1, 0.03, 0.01$. For each of $\gamma = 1, 0.98, 0.95$, we compute ordinary exponentially weighted least squares and precision-weighted least squares recursively without a prior. Full per-time errors, variances, and accumulated precisions are saved.

Writing $d_t = \gamma^{T-t}$, the ordinary estimator is $\sum_t d_t y_t / \sum_t d_t$ and its variance is $\sum_t d_t^2 \sigma_t^2 / (\sum_t d_t)^2$. The precision-weighted estimator is $\sum_t d_t y_t / \sigma_t^2 / S_{T,\gamma}$, with the exact variance in Eq. (11). For $\gamma < 1$, the numerator contains squared discount factors; treating its variance as exactly $1/S_{T,\gamma}$ would be incorrect. The theoretical mean absolute error is $\sqrt{2/\pi}$ times the exact standard deviation. No multiplicative constant is fitted to the measured curves.

F Proofs and Scope of the Guarantees

F.1 Weighted identification with predictable dependent regressors

Fix T and deterministic coefficients $a_s \in [0, 1]$. Define $u_s = \sqrt{a_s}\phi_{s-1}/\sigma_{\epsilon,s}$ and $\zeta_{s,j} = \sqrt{a_s}\xi_{s,j}/\sigma_{\epsilon,s}$. The features are predictable and the normalized noises are conditionally 1-sub-Gaussian, since $a_s \le 1$. They need not be independent over time. The self-normalized martingale

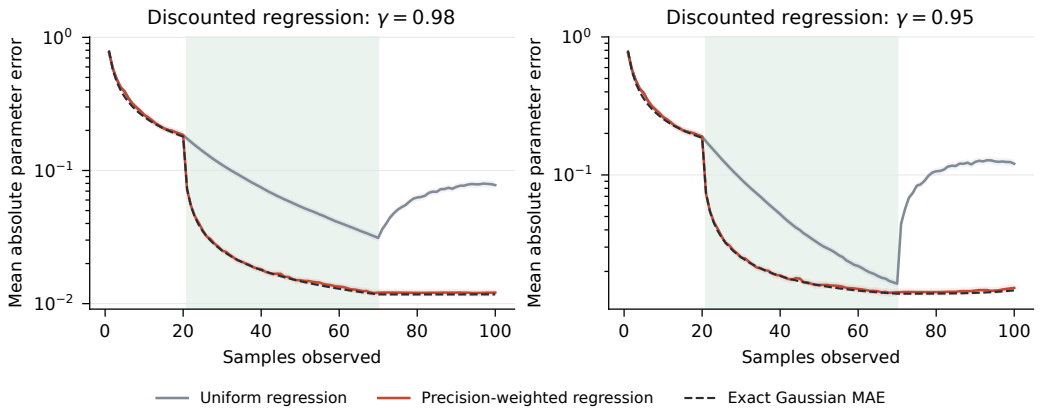

Figure 9: Independent scalar regression with forgetting, 600 paired realizations and low-noise standard deviation 0.1 in the shaded window. Curves compare uniform and precision-weighted mean absolute parameter error with the exact Gaussian prediction using squared discount factors. Bands are pointwise 95% confidence intervals for the Monte Carlo mean. These are diagnostic regression results, not full-DAKF tracking curves.

inequality, followed by a union bound over the n output coordinates, yields

$$\left\| \sum_s u_s \zeta_{s,j} \right\|_{V_T^{-1}} \le \sqrt{\log \frac{\det V_T}{\det(\lambda I)} + 2 \log \frac{n}{\delta}} \quad \text{for every } j \tag{13}$$

with probability at least $1 - \delta$. The normal equations give

$$\begin{aligned} \widehat{A}_T - A_T = & \left[\sum_s q_s \xi_s \phi_{s-1}^\top + \lambda (A_{\text{init}} - A_T) \right] V_T^{-1} \\ & + \left[\sum_s q_s b_s \phi_{s-1}^\top \right] V_T^{-1} + \left[\sum_s q_s (A_s - A_T) \phi_{s-1} \phi_{s-1}^\top \right] V_T^{-1}. \end{aligned} \tag{14}$$

For each row, the first line in the V_T norm is bounded by the self-normalized radius plus $\sqrt{\lambda} S_A$. Moreover,

$$\log \frac{\det V_T}{\det(\lambda I)} \le n \log \left(1 + \frac{\operatorname{tr} G_T}{n\lambda} \right) \le n \log \left(1 + \frac{L^2 S_T}{n\lambda} \right). \tag{15}$$

Since $\lambda_{\min}(V_T) \ge \lambda + \alpha S_T$, summing squared row bounds controls the first line of Eq. (14) by $\sqrt{n}\beta_T / \sqrt{\lambda + \alpha S_T}$.

For the bias term let $U = [\sqrt{q_s}\phi_{s-1}]_s$ and $B = [\sqrt{q_s} b_s]_s$. Then $UU^\top = G_T \preceq V_T$, so $\|U^\top V_T^{-1/2}\|_2 \le 1$ and

$$\|BU^\top V_T^{-1}\|_F \le \|B\|_F \|U^\top V_T^{-1/2}\|_2 \|V_T^{-1/2}\|_2 \le \frac{(\sum_s q_s \|b_s\|_2^2)^{1/2}}{\sqrt{\lambda + \alpha S_T}}. \tag{16}$$

For the drift term, apply the triangle inequality, $\|\phi_{s-1}\|_2^2 \le L^2$, and $\|V_T^{-1}\|_2 \le (\lambda + \alpha S_T)^{-1}$. Combining these bounds proves Theorem 1. For simultaneous statements at several horizons one must additionally union-bound the confidence level or use a suitable time-uniform construction. The displayed theorem is a fixed-horizon statement.

F.2 Harmonic inequality and scalar lower bound

Equation (8) follows from

$$\left(\sum_s \sigma_{\epsilon,s}^2 \right) \left(\sum_s \sigma_{\epsilon,s}^{-2} \right) \ge T^2. \tag{17}$$

For Proposition 1, take $A_0 = -\Delta/2$ and $A_1 = \Delta/2$. Their Gaussian observation laws have Kullback–Leibler divergence $\Delta^2 S_T/2$. Choose $\Delta = c_0 \min\{1, S_T^{-1/2}\}$ with a sufficiently small universal $c_0 > 0$. Both parameters lie in $[-1, 1]$, and their total variation distance is bounded away from one by Pinsker's inequality. Any estimator induces a nearest-parameter test; an estimation error smaller than $\Delta/2$ would imply a correct test decision. The two-point testing bound therefore gives a maximum expected absolute estimation error at least $c_1\Delta$, proving the claim. This scalar minimax construction establishes the necessity of the precision dependence in the centered stationary subproblem only.

F.3 Forgetting and information accounting under replay

For a fresh-sample recursion with $\gamma \in (0, 1]$, the normal equations have $a_s = \gamma^{T-s}$ and $\lambda = \gamma^T \lambda_0$. The normalized-noise proof above remains valid because $a_s \le 1$. In the independent scalar Gaussian diagnostic, taking the variance of the linear estimator directly gives Eq. (11). Since $\gamma^{2(T-s)} \le \gamma^{T-s}$, the exact variance is no larger than $1/S_{T,\gamma}$, with equality at $\gamma = 1$.

The synthetic DAKF implementation instead performs several RLS updates per physical time. At a fixed physical horizon, let J be the number of updates and let update j reuse distinct physical transition s_j. The original transition's frozen response, regressor, and scale are reused; its noise is not resampled. Its aggregate coefficient is

$$c_s = \sum_{j:s_j=s} \gamma^{J-j}. \tag{18}$$

The batch objective is $\sum_s c_s \|y_s - B\phi_{s-1}\|^2/\sigma_{\epsilon,s}^2 + \gamma^J \lambda_0 \|B - A_{\text{init}}\|_F^2$. Divide the entire objective by $c_{\max} = \max(1, \max_s c_s)$, setting $a_s = c_s/c_{\max}$ and $\lambda = \gamma^J \lambda_0/c_{\max}$. The minimizer is unchanged, all $a_s \le 1$, and the theorem applies to the distinct physical samples with these coefficients. In particular, replaying a sample nine times cannot be counted as nine independent noise realizations in the concentration argument.

The package saves both raw and normalized accumulated precision and Gram values over physical time, along with the exact update count and normalized prior. In the abrupt experiment $\gamma = 0.95$ is per internal RLS update, not per physical observation. Batch estimates with this accounting agree with the recursive prediction-time parameter estimates within 4×10^{-14} for all 200 trials of each variance mode. This verifies the weighting recursion and discount bookkeeping, not the excitation or calibration assumptions themselves.

F.4 Filtered-state bias and estimated variance

Using $e_s = x_s - \widehat{x}_s$ in Eq. (1) gives

$$\widehat{x}_s = A_s \widehat{x}_{s-1} + w_s + A_s e_{s-1} - e_s. \tag{19}$$

Conditional centering of the last three terms yields Eq. (5). The identification theorem concerns the unconstrained weighted-ridge recursion. The repeatedly clipped coefficient update in the extended UAV implementation is an empirical variant; its exact recursion is not covered by the weighted normal-equation proof. The filtered state is an endogenous response; even with zero-mean process and observation noise, the effective residual need not have zero conditional mean. The bias term in Eq. (7) explicitly retains this discrepancy. It need not vanish as precision increases.

For an estimated predictable proxy s_s^2, suppose the true centered residual is conditionally sub-Gaussian with scale at most $\sqrt{\kappa} s_s$ for a deterministic $\kappa \ge 1$. The normalized noise in the proof has scale at most $\sqrt{\kappa}$, multiplying only the stochastic self-normalized radius by $\sqrt{\kappa}$. The design and PW-PE condition must use the actual estimated weights a_s/s_s^2. A sliding sample variance is causal but may underestimate a noise increase, overreact to drift, or absorb filtering bias; we do not assert this calibration condition is automatic.

Gaussian states also do not satisfy a deterministic bounded-state assumption over all time. In finite-horizon simulations, the bounded-regressor condition should be interpreted on an explicitly controlled envelope event (or by a stopped-process argument), with its failure probability added to the confidence bound. Empirical bounded trajectories alone do not prove a distribution-wide guarantee.

F.5 Proof and interpretation of the filtered-response bias bound

All expectations below condition on $\mathcal{F}_{s-1}$. Write $m = \mathbb{E}[x_{s-1} \mid \mathcal{F}_{s-1}]$, $\mu = m - \phi$, and $C = \mathrm{Cov}(x_{s-1} \mid \mathcal{F}_{s-1})$. Thus $\Pi = C + \mu\mu^\top$, $C \preceq \Pi$, and $\|\mu\|_2 \le \sqrt{\mathrm{tr}\,\Pi}$. We assume the noises are conditionally uncorrelated with each other and with $x_{s-1} - m$; centeredness alone would not justify adding their covariances. For $u = z_s - HA_s m$, we have $\mathbb{E}[u] = 0$ and

$$\Omega = \mathbb{E}[uu^\top] = H(A_s C A_s^\top + Q)H^\top + R \preceq \bar{\Omega}.$$

The ith candidate response minus $A_s\phi$ is $d_i + K_i u$, where

$$d_i = (I - K_i H)(B_i - A_s)\phi + K_i H A_s \mu.$$

Let $I_2 = 1$ if $k_s = 2$ and zero otherwise. The selector may use the current observation; it need not be predictable. Taking conditional expectations yields the exact identity

$$b_s = \sum_i p_i d_i + \mathbb{E}[(I_2 - p)(K_2 - K_1)u \mid \mathcal{F}_{s-1}]. \tag{20}$$

For the last term, vector Cauchy–Schwarz and $\mathrm{Var}(I_2) = p(1-p)$ give

$$\left\|\mathbb{E}[(I_2 - p)(K_2 - K_1)u]\right\|_2 \le \sqrt{p(1-p)\,\mathrm{tr}[(K_2 - K_1)\Omega(K_2 - K_1)^\top]}.$$

Apply the triangle inequality to the mismatch terms, use $\|\bar{K}HA_s\mu\|_2 \le \|\bar{K}HA_s\|_2\sqrt{\mathrm{tr}\,\Pi}$, and replace Ω by its positive-semidefinite upper bound to obtain Eq. (10). No Gaussian assumption or independence between the selector and observation is used. With identical gains, the selection term is exactly zero; with a predictable single arm, $p(1-p) = 0$.

Connection to the implemented regressions. The synthetic implementation uses the previous selected posterior as both arms' input and the regression covariate, precisely the shared-ϕ case above. The abrupt UAV adaptive arm regresses its own consecutive velocity estimates. Apply the single-arm version to the full position/velocity response with its actual transition matrix, and then project onto the velocity coordinate. Since the first row of the constant-velocity transition is shared, the mismatch in the second row is the learned velocity coefficient mismatch. In a more general separate-input selector, an additional known term $(I - K_i H)B_i(\phi_i - \phi)$ must be included for each candidate; omitting it would incorrectly equate shared-state and separate-state implementations.

What this adds to the theorem. For the conditions of Theorem 1, its bias numerator is at most $(\sum_s q_s B_s^2)^{1/2}$. This replaces an unrestricted symbol by state-error, transition-mismatch and gain-disagreement terms. It does not establish their decay automatically. If a calibrated covariance envelope $\Pi \preceq c_P P$ is available, substitution provides an explicit bound in P, Q, R and model mismatch. Under adaptation or misspecification the reported Riccati P need not satisfy that envelope, so it cannot silently replace actual error moments. The bound motivates reducing initialization error before enabling identification, and controlling model/gain disagreement; it does not imply inverse-variance weights always reduce the coupled system's error.

Numerical bound audit. For three weights, 200 saved trajectories and 99 transitions (59,400 cases), a true-model oracle filter and integration over the actual selector's decision regions give conditional moments. These are diagnostic quantities, not online inputs. The maximum bound violation is below 10^{-12}; median $|b_s|/B_s$ is 0.376. Selection bias exceeds 10^{-8} in 49.2% of cases. Source and raw terms are supplied.

F.6 Conditional tracking along the actual state ancestry

Equation (12) follows by subtracting Eq. (2) from the true state. The relevant previous error is the input error $\bar{e}_{t-1}^{(i)}$, so the identity covers both separate-arm and shared-posterior implementations. A synchronization copy redirects an input to the copied arm; it does not create a new independent state estimate. A hard-selected state has one actual ancestry through these input choices.

Along such an ancestry π, define $M_s^\pi = (I - K_s^\pi H)A_s^\pi$, mismatch $\Delta_s^\pi = A_s - A_s^\pi$, and noise η_s^π. Suppose every realized ancestry product obeys $\|M_{t:k+1}^\pi\|_2 \le C\mu^{t-k}$ for $0 < \mu < 1$, and

$\|(I - K_s^\pi H)x\|_2 \le C_K\|x\|_2$ with $\|x_s\|_2 \le X$ on the considered event. Unrolling gives the pathwise bound

$$\|x_t - \widehat{x}_t\|_2 \le C\mu^t E_0 + CC_K X \sum_{s=1}^{t} \mu^{t-s}\|\Delta_s^\pi\|_2 + C\sum_{s=1}^{t} \mu^{t-s}\|\eta_s^\pi\|_2, \tag{21}$$

where E_0 bounds the initial input errors. The selected ancestry may depend on current observations, but the inequality is pathwise under the assumed product bound. Synchronization alters which mismatch terms occur later. A smaller adaptive identification bound can reduce the associated terms of Eq. (21); no inequality against the best arm is implied.

Fixed-filter Riccati stability does not establish this product bound under adaptive switching; pointwise spectral radii below one are also insufficient. Stability therefore remains an explicit condition, consistent with the native cold-start failure and gradual-drift results. These conditions do not establish an unconditional full-algorithm guarantee.