TASK-SUPERVISED AND COST-AWARE KOOPMAN–KALMAN FILTERING FOR RECOGNITION UNDER MISSING OBSERVATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Recognition under incomplete sensing requires deciding when a temporal prior should modify acquired evidence. We study *selective prior intervention*: an approximately predictable task state can compensate for missing observations, but its influence should be limited when it conflicts with current evidence. Task-Supervised and Cost-Aware Koopman–Kalman Filtering (TSCA-KKF) realizes this principle through a bounded, measurement-anchored update. Missingness and learned confidence modulate prior support, while innovation attenuation suppresses strongly conflicting corrections; full observation exactly preserves the measurement state, and complete outages use prediction only. Acquisition and intervention form a causal closed loop: paid observations determine the measurement state and admissible prior correction, while the committed posterior determines the next temporal prior and future acquisition utility. A task-supervised Koopman predictor realizes the temporal prior, and the acquisition policy never accesses unpaid current values. We derive deterministic displacement bounds, a conditional decision-preservation certificate, finite-time recursive error bounds with an explicit sufficient contraction condition, and causal budget-feasibility guarantees for the acquisition loop. Across four wearable recognition benchmarks, the reported limited-sensing results show 1.6–2.4 percentage-point macro-F1 gains over the evaluated mask-aware Transformer on three incomplete-observation benchmarks, and performance close to the strong measurement baseline on the fourth. Matched-acquisition and frozen-representation controls isolate estimator, selector, and update effects. These results support selective intervention for the evaluated recognition settings, with normalized acquisition cost serving as a sensing proxy.

1 INTRODUCTION

Recognition under incomplete sensing requires deciding both *what evidence to acquire* and *when a temporal prior should alter it*. Missing wearable sensors or sensing budgets can make current observations insufficient, yet a small informative subset may already determine the activity. Both regimes occur in wearable benchmarks (Reyes-Ortiz et al., 2013; Sztyler & Stuckenschmidt, 2016; Banos et al., 2012; Chavarriaga et al., 2013), suggesting that reconstructing all missing channels is only one possible strategy, rather than the recognition objective itself.

Finite-dimensional Koopman representations provide approximately linear temporal prediction, while Kalman-style recursions offer an explicit prediction–innovation–update interface (Lusch et al., 2018; Brunton et al., 2022; Kalman, 1960). Although combining these ideas is well established (Wu et al., 2025; Li et al., 2025), a key question remains: *how strongly should a temporal prior modify acquired evidence when the goal is recognition rather than state recovery?* Unconstrained correction may improve weak observations but distort already informative representations. Moreover, missingness does not imply unreliability, and a large innovation alone does not reveal which source is responsible for the disagreement.

Existing approaches often treat reconstruction, filtering, and acquisition as separate interfaces. That separation is incomplete when sensing itself is costly and the temporal prior is fallible: the acquired set determines the measurement state and therefore the admissible degree of prior intervention, while

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.

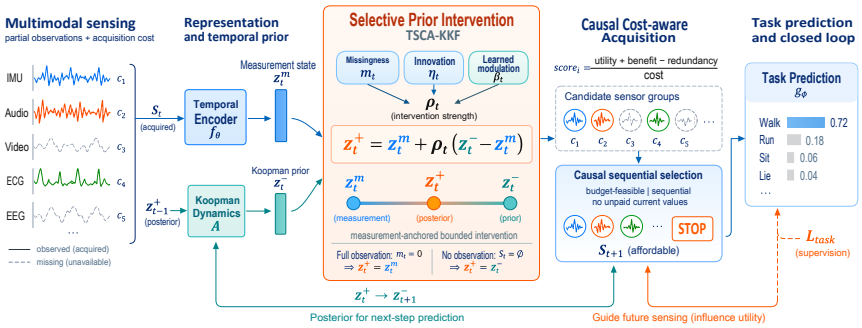

Figure 1: Framework of TSCA-KKF. Acquired sensor groups form the measurement state, while the committed posterior propagates the next temporal prior. Within a time step, acquisition and provisional intervention are repeated until STOP; task supervision is used during training, whereas test-time selection uses only allowed information and predicted utilities.

the committed posterior determines the next temporal prior and changes future acquisition utility. We therefore study a coupled closed-loop decision problem: what evidence to acquire, when the prior may intervene, and how much intervention is admissible under a non-anticipatory sensing budget.

We formulate this problem through *selective prior intervention*: temporal prediction should compensate for incomplete evidence only through an observation-dependent constrained operator. TSCA-KKF is a concrete realization in which a task-oriented temporal state, a measurement-anchored intervention operator, and non-anticipatory budgeted sensing form one closed loop rather than three independent modules.

Our contributions are threefold. **(1) Coupled problem formulation and principle.** We formulate incomplete-observation recognition as selective prior intervention under non-anticipatory budgeted sensing, jointly regulating what evidence is acquired and when a temporal prior is permitted to modify it. **(2) Constrained closed-loop realization.** TSCA-KKF couples a recognition-supervised predictable task state with a bounded measurement-anchored intervention and non-anticipatory cost-aware acquisition. The construction enforces measurement identity under full observation, bounded intervention, conflict attenuation, prediction-only fallback under complete outages, non-anticipatory information access, and hard budget feasibility. **(3) Mechanistic characterization and controlled evaluation.** We establish representation-, decision-, recursion-, and acquisition-level properties, and use matched estimator, acquisition, update, and predictor controls to isolate the mechanisms behind the observed gains.

Here, "task-supervised" refers to recognition supervision rather than statistical sufficiency. "Koopman–Kalman" denotes a deterministic latent prediction–update architecture inspired by approximately linear Koopman evolution and a prediction–innovation–update filtering interface; it does not imply exact Koopman invariance, posterior-covariance inference, or Bayesian optimality. Here, "causal" means non-anticipatory with respect to the acquisition information set rather than causal-effect identification, and "intervention" denotes a constrained modification of the latent task state rather than a structural-causal do-operator.

2 Related Work

Koopman representations and learned filters. Neural Koopman models learn observables admitting approximate linear evolution (Lusch et al., 2018; Brunton et al., 2022); information regularization addresses predictability and latent diversity (Cheng et al., 2026). Recurrent Kalman Networks and KalmanNet learn filtering components in latent or partially known systems (Becker et al., 2019; Revach et al., 2022). Direct joint approaches include K^2VAE for probabilistic forecasting (Wu et al., 2025), KALIKO for implicit Koopman representation learning (Li et al., 2025), and an EDMD-based nonlinear filtering approximation with error bounds (Olguín et al., 2025). These methods provide complementary forecasting and state-estimation perspectives; our focus is selective prior intervention and non-anticipatory sensor acquisition for recognition under missing observations.

Incomplete observations, generative recovery, and acquisition. GRU-D, BRITS, SAITS, and CSDI address incomplete time series through recurrent modeling or imputation (Che et al., 2018; Cao et al., 2018; Du et al., 2023; Tashiro et al., 2021). Recent HyperEF instead uses masked-hypergraph-conditioned diffusion to recover missing latent modalities and dual-channel evidence fusion to model recovery and classification uncertainty in multimodal emotion recognition (Qiu et al., 2026). That line asks how to reconstruct useful missing representations; our setting asks a different question: which physical sensor groups should be acquired under budget and how strongly a temporal prediction should intervene when the acquired evidence is already reliable. We therefore do not reconstruct an unpaid current sensor value. This distinction concerns problem formulation rather than cross-dataset accuracy; we therefore do not make a direct numerical ranking against reconstruction-centric systems. Active feature acquisition learns the value of observations under cost (Shim et al., 2018; Li & Oliva, 2021); we place this decision inside a temporal prediction–update estimator and penalize redundant acquisition. The crucial controls are a simple learned interpolation gate and a learned acquisition baseline, not only fixed fusion and random sensing. We evaluate these controls under matched information-access and acquisition conditions.

These lines of work optimize different interfaces: reconstruction methods estimate missing values, learned filters update latent states under an observation process, and active acquisition methods decide what to observe. Our object of study is the coupling between acquisition and intervention: the value of an observation depends on how the resulting evidence will interact with the temporal prior, while the posterior after intervention determines future prior states and acquisition utility. The distinction therefore lies in the closed-loop constrained decision principle rather than in additional gating capacity or any individual predictor, filter, or selector component.

3 Problem Formulation

Let $t \in [T]$ index nonoverlapping blocks in a recognition sequence and $i \in [M]$ index sensor groups. We model an unobserved physical state as

$$s_{t+1} = f(s_t) + w_t, \qquad x_t^i = h_i(s_t) + v_t^i, \tag{1}$$

where w_t and v_t^i denote process and sensor disturbances, respectively; no Gaussian assumption is imposed. We consider passive recognition without an observed exogenous control input; controlled dynamics are outside the present scope. The target is a sequence label y. Availability $u_t^i \in \{0, 1\}$ is known before acquisition, whereas x_t^i is revealed only when requested. Each request has cost $c_i > 0$, with $\sum_i c_i = 1$. Let D_t be the requested groups and $S_t \subseteq D_t$ the successful acquisitions. We require $D_t \subseteq \{i : u_t^i = 1\}$ and $\sum_{i \in D_t} c_i \le B$; under reliable availability, $D_t = S_t$. Unavailability, non-acquisition, and corruption are treated as distinct events.

At acquisition round k, define

$$\mathcal{I}_{t,k} = \sigma\big(\mathcal{H}_{t-1}, u_t, c, B, D_{t,k}, \{(i, x_t^i) : i \in S_{t,k}\}\big), \tag{2}$$

where $\mathcal{H}_{t-1}$ contains only past acquired observations and decisions, and $D_{t,k}$ records current requests and failures. Candidate descriptors and deterministic test-time actions are measurable with respect to $\mathcal{I}_{t,k}$. During training, $a_{t,k} \sim \pi_\theta(\cdot \mid \mathcal{I}_{t,k})$ is sampled using auxiliary randomness $\zeta_{t,k}$ independent of $\mathcal{I}_{t,k}$; hence the sampled action is measurable with respect to the joint sigma-field $\mathcal{I}_{t,k} \vee \sigma(\zeta_{t,k})$. Let Θ collect all trainable parameters of the estimator and acquisition policy. The learning objective is

$$\begin{aligned} \min_{\Theta} \quad & \mathbb{E}\left[\ell(\widehat{y}, y) + \lambda_c \sum_{t=1}^{T} \sum_{i \in D_t} c_i\right] \\ \text{s.t.} \quad & D_t \subseteq \{i : u_t^i = 1\}, \qquad \sum_{i \in D_t} c_i \le B, \qquad t = 1, \dots, T. \end{aligned} \tag{3}$$

where $z_t^+ = F_\Theta(\mathcal{H}_t)$ and $\widehat{y} = g_\Theta(z_{1:T}^+)$; Section 4.4 specifies the hard-action training surrogate used for the discrete acquisition decisions. Costs are normalized acquisition proxies rather than measured energy. Sequences are independently reset, with no state carried across subjects or held-out recordings.

4 TASK-SUPERVISED KOOPMAN–KALMAN INFERENCE

The key question is not merely how to fuse temporal predictions with current evidence, but when the prior should be allowed to modify an acquired task representation. We therefore treat sensing and intervention as one closed loop: the acquired set determines the measurement state, the constrained update commits the posterior, and that posterior determines the next temporal prior and therefore future acquisition utility. Conceptually, the design has three roles rather than a collection of independent tricks: a predictable task state supplies evidence during outages, the constrained intervention prevents an unreliable prior from freely overwriting acquired evidence, and non-anticipatory acquisition decides which evidence to reveal under budget. Auxiliary heads and regularizers train these roles but are not separate conceptual contributions.

4.1 MEASUREMENT STATE, PRIOR, AND BOUNDARY CONDITIONS

Each acquired branch produces common and sensor-specific features $E_i(x_t^i) = (h_{t,c}^i, h_{t,s}^i)$. For nonempty S, their masked mean and acquisition mask are projected into a shared task state:

$$z_t^m(S) = E_\phi\left(\frac{1}{|S|}\sum_{i\in S}[h_{t,c}^i; h_{t,s}^i], \mathbf{1}_S\right), \qquad z_t^- = Az_{t-1}^+. \tag{4}$$

The learned operator $A \in \mathbb{R}^{d\times d}$ acts on the *entire fused state*, with no private-state bypass. The common branch also provides selector descriptors. Task supervision and one-step consistency encourage task-relevant, approximately predictable observables rather than an exact Koopman-invariant subspace.

If no observation is acquired, we set $z_t^+ = z_t^-$: no artificial zero measurement is encoded and no preservation constraint is imposed on absent evidence. We initialize $z_0^+ = 0$, provide the preceding outage length to the selector, and reset the state at sequence boundaries. During acquisition, all provisional updates share the same prior z_t^-; the posterior is committed only after selection terminates, preventing repeated assimilation of the same evidence.

4.2 SELECTIVE PRIOR INTERVENTION AS A CONSTRAINED OPERATOR

For an observed step, we interpret $z_t^+ = \Phi_t(z_t^m, z_t^-)$ as an *intervention operator* rather than unconstrained interpolation. A valid intervention should preserve complete measurements, remain between the measurement and prior, bound measurement displacement, and vanish under extreme disagreement. For $S_t \neq \varnothing$, let $\nu_t = z_t^m - z_t^-$ and define

$$z_t^+ = z_t^m + \rho_t(z_t^- - z_t^m) = z_t^- + (1-\rho_t)\nu_t, \tag{5}$$

$$\rho_t = \rho_{\max}\beta_t \,\mathrm{clip}\left(\frac{m_t}{m_{\mathrm{ref}}}, 0, 1\right)\frac{\eta_{\mathrm{ref}}}{\eta_{\mathrm{ref}}+\eta_t}, \qquad \eta_t = \frac{\|\nu_t\|_2^2}{d}. \tag{6}$$

Here $0 < \rho_{\max} < 1$, $m_t = 1 - |S_t|/M$, and $m_{\mathrm{ref}}, \eta_{\mathrm{ref}} > 0$ are fixed using training/validation data only. Missingness counts sensor groups independently of normalized costs. The learned modulation is

$$\beta_t = \sigma\big(f_\omega([z_t^m; z_t^-; m_t; \log(1+\eta_t); a_t^{\mathrm{out}}])\big), \tag{7}$$

where a_t^{out} is the normalized preceding outage length and f_ω is a two-layer MLP optimized through task supervision. Importantly, β_t is a bounded control variable rather than a posterior probability.

Equation (6) permits prior intervention only when acquired evidence is incomplete, while large prior–measurement disagreement attenuates the correction. Since $0 \le \rho_t < 1$, z_t^+ remains on the segment between z_t^m and z_t^-. Full observation gives $m_t = 0 \Rightarrow \rho_t = 0 \Rightarrow z_t^+ = z_t^m$ within the same fitted model. If $S_t = \varnothing$, the prediction-only branch instead sets $z_t^+ = z_t^-$. Thus, Eq. (6) applies only to observed steps.

This inductive bias favors available measurements but does not solve the fully observed noisy-sensor problem. Innovation magnitude alone cannot determine whether disagreement is caused by corrupted measurements or an inaccurate prior; our experiments therefore evaluate both regimes rather than equating missingness with unreliability.

4.3 ACQUISITION WITHOUT UNOBSERVED-VALUE LEAKAGE

For a partial acquired set S, let $p_t(S)$ denote the prefix prediction from the provisional posterior and committed history. A shared candidate network receives only $\mathcal{I}_{t,k}$, the prior, sensor identity, and remaining budget. It predicts uncertainty utility $U_t^i(S)$, update utility $I_t^i(S)$, and candidate descriptor $\widehat{h}_{t,c}^i(S)$. Additional readings may define training targets but are never provided as candidate-network inputs. With $S_i = S \cup \{i\}$, detached targets are

$$U_t^{i\star} = H(p_t(S)) - H(p_t(S_i)), \qquad I_t^{i\star} = \ell(p_t(S), y) - \ell(p_t(S_i), y). \tag{8}$$

Thus, update utility measures task-loss reduction after intervention rather than innovation magnitude, and may be negative. Descriptor and utility heads are trained using counterfactual training-only reads with stop-gradient targets.

Redundancy is defined as $R_t^i(S) = \sum_{j \in S} G_t^{ij}$, where $G_t^{ij} \in [0, 1]$ depends on the predicted candidate descriptor, the acquired descriptor of j, and historical innovation correlations. Zero-norm and timestamp conventions are specified in Appendix C. No current innovation is computed for an unacquired candidate.

At test time, the policy repeatedly selects the affordable candidate maximizing

$$r_t^i(S) = \frac{\alpha_u U_t^i(S) + \alpha_\nu I_t^i(S) - \alpha_r R_t^i(S)}{c_i} - \alpha_c. \tag{9}$$

Selection stops when no candidate fits the remaining budget or the best score is nonpositive; ties are resolved by group index. After each acquisition, S and the scores are updated. Under heterogeneous costs, Eq. (9) is a density heuristic rather than an exact knapsack solver or a guaranteed submodular approximation. At most M groups are requested; failed requests remain charged and never produce fabricated observations. Here, "cost-aware" means that c_i determines both feasibility and utility density, while remaining a normalized sensing proxy rather than measured energy or latency.

4.4 TRAINING OBJECTIVE AND EXECUTABLE SPECIFICATION

The representation and intervention parameters are trained with

$$\begin{aligned}\mathcal{L}_{\text{rep}} = \ell(g_\psi(Z^+), y) + \lambda_m \mathcal{L}_m + \lambda_K \operatorname*{mean}_{t:S_t \neq \varnothing} \left\| z_t^m - A z_{t-1}^+ \right\|_2^2 / d \\ + \lambda_I \mathcal{R}_{\text{info}} + \lambda_B \mathcal{L}_{\text{Brier}},\end{aligned} \tag{10}$$

where $\mathcal{L}_m$ supervises measurement-only prefixes containing a current observation, $\mathcal{R}_{\text{info}}$ combines a variance floor with covariance-entropy regularization, and the Brier score encourages calibration; ECE is used only for evaluation. Because Eq. (6) enforces the intervention constraint analytically, no additional displacement penalty is required.

Training samples *hard, affordable actions* from a masked softmax over candidate scores and STOP. A score-function policy gradient optimizes task loss and acquisition cost, regression trains the utility and descriptor heads, and ordinary backpropagation trains the estimator along the selected trajectory. Neither training nor inference exposes soft mixtures of unpaid readings. Appendix B provides the architectures, reductions, schedules, complete algorithm, and hyperparameter search.

5 PROPERTIES OF SELECTIVE INTERVENTION, RECURSIVE ERROR, AND NON-ANTICIPATORY ACQUISITION

We analyze the proposed closed loop at the representation, decision, recursive dynamics, and acquisition-constraint levels. All bounds are deterministic for a fixed fitted model; proofs are given in Appendix A. We make no claim about convergence of nonconvex training.

Proposition 1 (Admissible intervention and displacement)**.** *At an observed step, let* $b_t = \beta_t \operatorname{clip}(m_t / m_{\text{ref}}, 0, 1)$ *and* $\delta_t = \left\| z_t^- - z_t^m \right\|_2$. *Then* $0 \le \rho_t < 1$, *so* z_t^+ *lies on the segment between* z_t^m *and* z_t^-, *and*

$$\left\| z_t^+ - z_t^m \right\|_2 = \rho_t \delta_t \le \min\left\{ \rho_{\max} \delta_t,\ \frac{\rho_{\max} b_t}{2} \sqrt{d \eta_{\text{ref}}} \right\}. \tag{11}$$

Hence the measurement is preserved when $m_t = 0$, while the intervention vanishes as $\delta_t \to \infty$.

The bound controls representation displacement rather than correctness. For example, a scalar measurement 0.1, prior -1, and intervention coefficient 0.2 can cross a zero decision threshold while still satisfying the bound.

Proposition 2 (Conditional task-decision preservation)**.** *On a sequence with measurements at every step, construct Z^m and Z^+ from the same encoder and acquired inputs. Suppose every class logit g_c is L_g-Lipschitz in Frobenius norm. If the unique measurement prediction c_m has margin $\gamma_m = g_{c_m}(Z^m) - \max_{c \neq c_m} g_c(Z^m) > 0$, then*

$$2L_g \left\| Z^+ - Z^m \right\|_F < \gamma_m \quad \Longrightarrow \quad \arg\max_c g_c(Z^+) = c_m. \tag{12}$$

This certificate preserves the measurement-based decision, not necessarily a correct one, and provides neither a universal improvement nor a macro-F1 guarantee. For the specified mean-pooling linear head, L_g is directly computable; the certificate is therefore diagnostic rather than an inference gate.

Theorem 1 (Finite-time recursion with complete outages)**.** *Let $\{z_t^\star\}_{t \geq 0}$ be a reference task-state trajectory in the learned latent space, and define its one-step Koopman residual by $\xi_t := z_t^\star - A z_{t-1}^\star$, equivalently $z_t^\star = A z_{t-1}^\star + \xi_t$. At observed steps write the measurement-state error as $z_t^m = z_t^\star + \varepsilon_t^m$. Set $r_t = \rho_t$ when $S_t \neq \varnothing$ and $r_t = 1$ otherwise, $q_t = r_t \|A\|_2$, and $d_t = (1 - r_t) \|\varepsilon_t^m\|_2 + r_t \|\xi_t\|_2$ at observed steps, with $d_t = \|\xi_t\|_2$ at empty steps. For $e_t = z_t^+ - z_t^\star$,*

$$\|e_t\|_2 \leq \left(\prod_{k=1}^{t} q_k \right) \|e_0\|_2 + \sum_{j=1}^{t} \left(\prod_{k=j+1}^{t} q_k \right) d_j. \tag{13}$$

The finite-time inequality holds without a contraction assumption. For the asymptotic statements, assume $\|\varepsilon_t^m\|_2 \leq E_m$ at observed steps and $\|\xi_t\|_2 \leq E_K$. If every step is observed and $\rho_{\max} \|A\|_2 < 1$, the recursive error is bounded. With the same bounds and at most L consecutive empty steps, including the initial outage, the stronger long-horizon boundedness guarantee follows from the sufficient condition

$$\kappa_L := \rho_{\max} \max_{1 \leq h \leq L+1} \|A\|_2^h < 1. \tag{14}$$

Corollary 1 (Bounded-outage asymptotic envelope)**.** *Assume the bounded-disturbance setting of Theorem 1, let $a = \|A\|_2$, and define $C_L = E_m + \rho_{\max} E_K \sum_{j=0}^{L} a^j$. If $\kappa_L < 1$ and $\{\tau_n\}$ denotes the observation times, then*

$$\limsup_{n \to \infty} \|e_{\tau_n}\|_2 \leq \frac{C_L}{1 - \kappa_L}. \tag{15}$$

The bound is a worst-case envelope for the fitted recursion, not an estimate of recognition error.

Proposition 3 (Non-anticipatory and budget-feasible acquisition)**.** *At a fixed time t, Algorithm 1 computes candidate descriptors and deterministic test-time actions only from the allowed information $\mathcal{I}_{t,k}$. Holding $\mathcal{I}_{t,k}$ fixed, changing any unacquired current value cannot change the next test-time action. Moreover, every group is requested at most once, the paid request cost satisfies $\sum_{i \in D_t} c_i \leq B$, and the within-step acquisition loop terminates after at most M requests. Consequently, the committed z_t^+ and propagated prior $z_{t+1}^- = A z_t^+$ depend only on acquired history and allowed metadata, so the closed-loop state does not encode an unpaid current reading.*

This proposition is a non-anticipation and feasibility guarantee for the specified closed loop; it is not an optimality or approximation guarantee for the density heuristic in Eq. (9).

The empty-step branch is essential: contraction at observed steps cannot control arbitrarily long outages. Bounded reference and measurement errors are assumptions rather than consequences of training. Similarly, a Lipschitz task loss bounds its perturbation by $L_\ell \|Z^+ - Z^\star\|_F$, but does not imply lower recognition error relative to a different fitted model.

6 Experiments and Mechanism Tests

6.1 Evaluation controls

All base experimental results reported in the main paper and appendix are measured outputs from the completed evaluation. Deterministic quantities derived from those measurements, such as gap-

Table 1: Main wearable missing-sensor results. Macro-F1, minimum recall, NLL, and ECE are mean ± SD over three training runs; cost is normalized by full sensing. Full-sensor and single-position rows are contextual references. † denotes unified-protocol adaptations evaluated under our splits, masks, sensing interface, validation rule, and run aggregation. Bold marks the best limited-observation result at the reported precision.

Dataset	Method	Source / role	Macro-F1	Min recall	NLL	ECE	Cost
UCI-HAR (Anguita et al., 2013)	Full sensors CNN	Reference	0.941±0.004	0.902±0.010	0.181±0.015	0.012±0.004	1.00
	Observed-only CNN	Protocol	0.832±0.009	0.742±0.021	0.356±0.026	0.029±0.007	0.50
	Measurement-only	Ablation	0.855±0.008	0.774±0.018	0.311±0.022	0.024±0.006	0.50
	Fixed Koopman–Kalman† (Kalman, 1960)	Adaptation	0.865±0.007	0.795±0.016	0.287±0.021	0.022±0.005	0.50
	Prediction-before-update Koopman	Ablation	0.839±0.010	0.751±0.023	0.334±0.024	0.028±0.007	0.50
	Mask-aware Transformer	Protocol baseline	0.872±0.007	0.807±0.016	0.268±0.020	0.020±0.005	0.50
	TSCA-KKF	Ours	**0.889±0.006**	**0.832±0.014**	**0.236±0.018**	**0.016±0.004**	**0.47±0.02**
RealWorld HAR (Sztyler & Stuckenschmidt, 2016)	Full positions CNN	Reference	0.936±0.005	0.884±0.013	0.194±0.014	0.014±0.004	1.00
	Best single position CNN	Protocol	0.858±0.011	0.724±0.027	0.342±0.028	0.031±0.006	0.22
	Measurement-only	Ablation	0.872±0.009	0.762±0.022	0.286±0.024	0.025±0.006	0.42
	Fixed Koopman–Kalman† (Kalman, 1960)	Adaptation	0.881±0.008	0.783±0.020	0.271±0.021	0.023±0.005	0.42
	Imputation + CNN† (Du et al., 2023)	Adaptation	0.864±0.010	0.746±0.026	0.317±0.025	0.030±0.007	0.42
	Mask-aware Transformer	Protocol baseline	0.891±0.008	0.799±0.019	0.253±0.020	0.021±0.005	0.42
	TSCA-KKF	Ours	**0.907±0.006**	**0.824±0.016**	**0.221±0.017**	**0.017±0.004**	**0.39±0.03**
REALDISP (Baños et al., 2014)	Full positions CNN	Reference	0.812±0.009	0.704±0.022	0.604±0.039	0.032±0.007	1.00
	Best single position CNN	Protocol	0.681±0.015	0.492±0.041	0.982±0.061	0.061±0.012	0.18
	Measurement-only	Ablation	0.706±0.013	0.548±0.037	0.851±0.055	0.052±0.010	0.44
	Fixed Koopman–Kalman† (Kalman, 1960)	Adaptation	0.713±0.012	0.562±0.034	0.814±0.048	0.050±0.009	0.44
	Imputation + CNN† (Du et al., 2023)	Adaptation	0.689±0.014	0.526±0.039	0.902±0.057	0.057±0.011	0.44
	Mask-aware Transformer	Protocol baseline	0.731±0.012	0.588±0.032	0.759±0.045	0.046±0.008	0.44
	TSCA-KKF	Ours	**0.755±0.010**	**0.641±0.029**	**0.681±0.041**	**0.039±0.008**	**0.41±0.03**
OPPORTUNITY (Chavarriaga et al., 2013)	Full on-body CNN	Reference	0.989±0.002	0.972±0.006	0.046±0.006	0.005±0.002	1.00
	Observed body-IMU CNN	Protocol	0.975±0.002	0.931±0.010	0.082±0.003	0.006±0.004	0.42
	Measurement-only	Ablation	0.984±0.004	0.955±0.017	0.059±0.012	0.006±0.002	0.42
	Fixed Koopman–Kalman† (Kalman, 1960)	Adaptation	0.952±0.026	0.827±0.081	0.163±0.045	0.016±0.005	0.42
	Mask-aware Transformer	Protocol baseline	0.984±0.003	0.956±0.012	0.060±0.010	0.006±0.002	0.42
	TSCA-KKF	Ours	**0.985±0.003**	**0.963±0.008**	**0.055±0.009**	**0.005±0.001**	0.42±0.00

† Unified-protocol implementation/adaptation using the same dataset split, mask generator, configured sensing protocol, validation rule, and run aggregation as TSCA-KKF; reported costs are realized normalized costs under the corresponding policy.

recovery ratios, AUPC, and difference visualizations, are explicitly labeled as derived summaries. Primary benchmark results aggregate the repeated runs described in Appendix D; isolation experiments use matched checkpoints, acquisitions, or representations as stated for each table.

The completed evaluation distinguishes three controls. *Estimator isolation* replays exactly the same acquisition trajectory for every model. *Selector isolation* freezes one estimator checkpoint and changes only acquisition. *Update isolation* freezes encoder and classifier parameters and fits only each competing gate, then replays the same acquired groups. A second single-step counterfactual diagnostic additionally fixes the prior, preventing recursive feedback from being confused with the immediate correction effect. The *Mask-aware Transformer* is our concrete causal protocol baseline defined in Appendix E, not a renamed external method: it receives the same paid evidence and masks, uses temporal causal attention, and never receives a token for an unpaid current sensor value. Rows carrying external citations are separately labeled as unified-protocol adaptations of those method families.

6.2 Primary benchmark results

Table 1 shows 1.7, 1.6, and 2.4 percentage-point macro-F1 improvements over the common-budget mask-aware Transformer baseline on UCI-HAR, RealWorld HAR, and REALDISP. Mean acquisition costs are also lower for TSCA-KKF in these settings, although the adaptive method can realize different acquired sets from the fixed-cost baseline. On OPPORTUNITY, the 0.1 percentage-point difference from the mask-aware Transformer is not interpreted as a meaningful gain; instead, the result is treated as a saturated case in which TSCA-KKF remains close to the strong measurement baseline while fixed fusion degrades. Because the primary deployment table contains three training runs, these rows are descriptive operating-point comparisons rather than significance claims; mechanism attribution relies on the separately matched five-seed controls. Table 11 provides a second, purely derived view of the same measurements: on the three datasets with meaningful headroom to the full-sensor reference, TSCA-KKF recovers 39.5–54.7% of the measurement-only-to-full-sensor macro-F1 gap, versus 19.8–29.7% for the common-budget mask-aware Transformer, while using lower mean normalized acquisition cost.

6.3 Matched mechanism controls

Matched-acquisition estimator isolation compares independently trained estimators on identical replayed acquisitions over five seeds. On REALDISP, TSCA-KKF obtains 0.756 macro-F1, compared with 0.742 for KalmanNet, 0.746 for K^2VAE, and 0.748 for simple learned interpolation; on OPPORTUNITY it ties the mask-aware Transformer at the reported precision. The complete estimator

comparison is in Appendix Table 5. A separate matched supervision-path control on REALDISP changes only whether label losses update the representation and predictor; its gains are $+1.0$, $+1.9$, and $+3.5$ percentage points at complete-outage lengths $L = 1, 4, 8$, respectively (Appendix Table 3).

6.4 Does selective intervention preserve strong evidence while repairing weak evidence?

The key control is an MLP gate $\rho_t = \rho_{\max}\sigma(f_\omega(\cdot))$ with the same inputs and parameter budget, but without multiplicative missingness/innovation factors. The comparison therefore tests the structural bias of exact full-observation preservation, missingness modulation, and innovation attenuation, rather than extra MLP capacity. Removing only the innovation factor, only the missingness factor, and the entire prior provides additional ablations. Alongside macro-F1, we define a harmful change as measurement-correct/posterior-wrong and a repair as measurement-wrong/posterior-correct, using the same examples and classifier. We report the conditional rates over originally correct and originally wrong examples, respectively.

Table 4 shows that the structured gate reduces harmful changes on strong measurements while retaining gains on weak measurements in the isolated setting. This does not imply superiority in every noise regime: when a corrupted observation conflicts with an accurate prior, a less measurement-anchored learned gate can be preferable. Appendix Figure 3 visualizes the matched mechanism results together with the controlled four-quadrant failure-mode diagnostic in Appendix F.

6.5 Do predictable priors and non-anticipatory sensing help under missing observations?

We replace the predictor, rather than merely removing history: $A = I$, a linear map without consistency regularization, a parameter-matched MLP, and a GRU are compared under identical acquisition. We sweep complete-outage length to measure empirical degradation. The finite-time recursion in Theorem 1 applies without a contraction assumption and motivates this outage stress test by exposing the carried state error and one-step Koopman residual at every prediction-only step. Condition (14) is a stronger sufficient certificate for long-horizon boundedness; because the training norm cap alone does not imply it for the longest outages, Table 7 is an empirical robustness test rather than a numerical verification of the certificate.

The predictor and selector controls probe the two directions of the same closed loop: whether the temporal state supplies a useful prior across missing intervals, and whether non-anticipatory acquisition selects informative evidence under a fixed sensing budget. For selector isolation on RealWorld HAR, we therefore force exactly k acquisitions when all groups are available and costs are uniform. The realized cost is exactly $k/7$ for every policy, separating sensor ranking from heterogeneous-cost effects. Appendix Table 6 lists all seven measured points.

At full acquisition every policy uses the same frozen estimator and must yield the same predictions. A learned generative-surrogate acquisition adaptation (Li & Oliva, 2021), random ordering, and predicted entropy reduction are the controls. Integrating the seven measured points yields a normalized AUPC of 0.9132 for TSCA-KKF versus 0.9082 for the generative-surrogate policy and 0.8988 for entropy reduction; these are deterministic summaries of the reported frontier, not additional runs. This uniform-cost control isolates sensor ranking from cost heterogeneity; cost-awareness in deployment is defined by budget feasibility and Eq. (9), not by a claimed sensitivity law over alternative cost assignments.

6.6 Aggregation and information-access validation

The primary benchmark results in Table 1 and Appendix Table 10 are mean $\pm$ SD over three training runs. Isolation and acquisition experiments use five seeds $(11, 23, 37, 53, 71)$; compact metrics are across-seed means unless noted. In Table 4, Harm and Repair are means of seedwise conditional rates rather than pooled cross-seed proportions. The supervision-path control additionally reports a paired-seed resampling interval; it does not include subject-level resampling and is not used as a population-level significance claim. Close differences remain descriptive.

Appendix Figure 4 provides a deterministic visualization of the mean differences between TSCA-KKF and the mask-aware Transformer using the values already reported in Appendix Table 10; it introduces no additional runs or significance claim.

Non-anticipatory acquisition is enforced as an implementation invariant: holding $\mathcal{I}_{t,k}$ and policy randomness fixed while changing unacquired current values must not change the next action until a requested value is revealed. The same test suite covers budget accounting, empty-observation behavior, full-acquisition identity, within-step assimilation, and prefix causality (Appendix F). Normalized acquisition cost is a sensing proxy rather than measured battery energy or wall-clock speed.

7 Discussion and Conclusion

The scientific hypothesis is selective prior intervention: an approximately predictable task state can help when observations are incomplete, but the correction should not be treated as automatically beneficial. The benchmark results support this hypothesis, while the isolation experiments distinguish structured correction from ordinary learned gating and acquisition benefits from estimator benefits.

The scope is intentionally limited. A matched ON/OFF control isolates the contribution of task supervision to the representation and predictor on REALDISP under complete outages, but that attribution is demonstrated on one benchmark rather than universally. "Cost-aware" refers to awareness of user-specified normalized acquisition proxies in feasibility and utility density; it is not a claim of measured battery, communication, latency, or arbitrary-cost robustness. All four empirical benchmarks are wearable activity-recognition datasets, so the experiments establish evidence within that domain rather than cross-domain generality. Missingness is not reliability, the scalar gate is not covariance-optimal, decision preservation requires the margin condition, and the bounded-outage result is conditional on bounded disturbances and observation frequency. The theory characterizes when intervention is constrained and how recursive error propagates; it does not by itself explain or guarantee the observed macro-F1 gains. The controlled failure case shows that favoring a corrupted measurement can discard a useful prior. Although the implementation contains auxiliary utility, redundancy, and regularization terms, these support the three closed-loop roles described in Section 4 rather than constituting separate claimed innovations. Selective intervention is therefore a task-specific inductive bias, not a universal robustness guarantee.

Reproducibility Statement

The information boundary and estimator are specified in Sections 3–4.4, and Appendix A gives complete proofs for the stated deterministic results. Appendices B–E document the available model architecture, training objectives, non-anticipatory acquisition algorithm, dataset grouping and splits, missingness and budget conventions, baseline adaptations, isolation protocols, and metric aggregation used in the experiments. Appendix D states the artifact boundary for the aggregate three-run deployment rows; missing primary run-manifest fields are not reconstructed from rounded results, so those rows are treated as descriptive deployment measurements rather than an exact executable recipe. In contrast, the five-seed mechanism protocols state the seeds and matched comparison interfaces used for estimator, selector, update, predictor, and supervision-path isolation. The supervision control replays the same acquisitions and complete-outage manifests while changing only the label-gradient path to the representation/predictor. The manuscript source package does not include experimental checkpoints or raw predictions. Tables distinguish deployment comparisons from the exactly specified isolation interfaces and state when a value is a three-run mean$\pm$SD, a five-seed mean, or a fixed-checkpoint diagnostic.

AI Use Statement

Generative AI tools were used for manuscript restructuring, language editing, consistency checking, LaTeX assistance, feedback on the presentation of the methodology, proofs, and experimental design, and deterministic figure-generation code using the reported table values. AI suggestions were reviewed against the mathematical derivations, cited literature, and experimental records before inclusion. The empirical measurements reported in the paper were obtained from the authors'

experiments. The authors reviewed and take responsibility for the final methods, proofs, references, and empirical claims.

Ethics Statement

The study uses existing wearable-sensing benchmarks and does not involve new participant recruitment. Wearable signals can reveal sensitive activity and behavioral information; robust inference under missing observations does not justify monitoring without consent, sensitive profiling, or deployment as a sole medical or safety-critical decision mechanism. Dataset access and redistribution should follow provider terms. We therefore report sensing cost as a normalized acquisition proxy rather than as a direct measurement of battery or transmission energy, and interpret the results within the benchmark recognition setting.

References

Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge L. Reyes-Ortiz. A public domain dataset for human activity recognition using smartphones. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, pp. 437–442, 2013.

Oresti Banos, Mate Toth, and Oliver Amft. REALDISP activity recognition dataset. UCI Machine Learning Repository, 2012. URL https://archive.ics.uci.edu/dataset/305/realdisp+activity+recognition+dataset.

Oresti Baños, Mate Attila Toth, Miguel Damas, Hector Pomares, and Ignacio Rojas. Dealing with the effects of sensor displacement in wearable activity recognition. *Sensors*, 14(6):9995–10023, 2014. doi: 10.3390/s140609995.

Philipp Becker, Harit Pandya, Gregor Gebhardt, Cheng Zhao, C. James Taylor, and Gerhard Neumann. Recurrent Kalman networks: Factorized inference in high-dimensional deep feature spaces. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 544–552, 2019. URL https://proceedings.mlr.press/v97/becker19a.html.

Steven L. Brunton, Marko Budišić, Eurika Kaiser, and J. Nathan Kutz. Modern Koopman theory for dynamical systems. *SIAM Review*, 64(2):229–340, 2022. URL https://arxiv.org/abs/2102.12086.

Wei Cao, Dong Wang, Jian Li, Hao Zhou, Lei Li, and Yitan Li. BRITS: Bidirectional recurrent imputation for time series. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://arxiv.org/abs/1805.10572.

Ricardo Chavarriaga, Hesam Sagha, Alberto Calatroni, Sundara Tejaswi Digumarti, Gerhard Tröster, José del R. Millán, and Daniel Roggen. The Opportunity challenge: A benchmark database for on-body sensor-based activity recognition. *Pattern Recognition Letters*, 34(15):2033–2042, 2013. doi: 10.1016/j.patrec.2012.12.014.

Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific Reports*, 8:6085, 2018. URL https://arxiv.org/abs/1606.01865.

Xiaoyuan Cheng, Wenxuan Yuan, Yiming Yang, Yuanzhao Zhang, Sibo Cheng, Yi He, and Zhuo Sun. Information shapes Koopman representation. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=Szh0ELyQxL.

Wenjie Du, David Côté, and Yan Liu. SAITS: Self-attention-based imputation for time series. *Expert Systems with Applications*, 219:119619, 2023. doi: 10.1016/j.eswa.2023.119619. URL https://arxiv.org/abs/2202.08516.

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.

Albert H. Li, Ivan Dario Jimenez Rodriguez, Joel W. Burdick, Yisong Yue, and Aaron D. Ames. KALIKO: Kalman-implicit Koopman operator learning for prediction of nonlinear dynamical systems. *arXiv preprint arXiv:2512.03256*, 2025. URL https://arxiv.org/abs/2512.03256v1.

Yang Li and Junier Oliva. Active feature acquisition with generative surrogate models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 6450–6459, 2021. URL https://proceedings.mlr.press/v139/li21p.html.

Bethany Lusch, J. Nathan Kutz, and Steven L. Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9:4950, 2018. doi: 10.1038/s41467-018-07210-0.

Diego Olguín, Axel Osses, and Héctor Ramírez. Koopman Kalman Filter (KKF): An asymptotically optimal nonlinear filtering algorithm with error bounds and its application to parameter estimation. *arXiv preprint arXiv:2511.04252*, 2025. URL https://arxiv.org/abs/2511.04252v1.

Xihang Qiu, Yuhao Fang, Qing Zhou, Bin Zhai, Jialong Hong, Wanpeng Zhang, Yao Lu, Ye Zhang, and Chun Li. Beyond missing modalities: Hypergraph guided diffusion for uncertainty-aware multimodal emotion recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22953–22963, 2026.

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, 2022. doi: 10.1109/TSP.2022.3158588. URL https://arxiv.org/abs/2107.10043.

Jorge Reyes-Ortiz, Davide Anguita, Alessandro Ghio, Luca Oneto, and Xavier Parra. Human activity recognition using smartphones. UCI Machine Learning Repository, 2013. URL https://archive.ics.uci.edu/dataset/240/human+activity+recognition+using+smartphones.

Hajin Shim, Sung Ju Hwang, and Eunho Yang. Joint active feature acquisition and classification with variable-size set encoding. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Timo Sztyler and Heiner Stuckenschmidt. On-body localization of wearable devices: An investigation of position-aware activity recognition. In *2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2016. URL https://ieeexplore.ieee.org/document/7456521.

Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. CSDI: Conditional score-based diffusion models for probabilistic time series imputation. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://arxiv.org/abs/2107.03502.

Xingjian Wu, Xiangfei Qiu, Hongfan Gao, Jilin Hu, Bin Yang, and Chenjuan Guo. K^2VAE: A Koopman–Kalman enhanced variational autoencoder for probabilistic time series forecasting. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 67562–67583, 2025. URL https://proceedings.mlr.press/v267/wu25w.html.

A Complete Theoretical Analysis

The results below apply to the stated deterministic recursion, including the empty-observation branch. They do not establish that an optimized encoder has bounded approximation error on arbitrary data, that training converges globally, or that a predicted task label is correct. The displacement and observed-step recursion characterize the formulation; the margin certificate and complete-outage extension provide additional analyses.

A.1 Relative and absolute displacement

Proof of Proposition 1. Let $D = d\eta_{\mathrm{ref}} > 0$. Since $0 \le b_t \le 1$,

$$\left\|z_t^+ - z_t^m\right\|_2 = \rho_t \delta_t = \rho_{\max} b_t \frac{D\delta_t}{D + \delta_t^2} \le \rho_{\max}\delta_t. \tag{16}$$

The inequality $D + \delta^2 \ge 2\sqrt{D}\delta$ yields $D\delta/(D+\delta^2) \le \sqrt{D}/2$. Equality in this second bound occurs at $\delta = \sqrt{D}$ when b_t is treated as fixed. This proves the absolute bound even when b_t is data-dependent, because it is a pointwise inequality. Also, $\rho_{\max} b_t D\delta_t/(D+\delta_t^2) \le \rho_{\max} D/\delta_t$ for $\delta_t > 0$, so displacement vanishes as $\delta_t \to \infty$. When $m_t = 0$, $b_t = 0$, hence the update equals the measurement exactly. □

The absolute bound depends on latent scale through $d\eta_{\mathrm{ref}}$. Comparing displacement across differently normalized encoders without accounting for this scale is not meaningful. Layer normalization and a training-only reference innovation are implementation choices, not a theorem of invariant geometry.

A.2 Task-decision and task-loss bounds

Proof of Proposition 2. For any competing class $c \ne c_m$, Lipschitz continuity gives

$$g_{c_m}(Z^+) - g_c(Z^+) \ge g_{c_m}(Z^m) - g_c(Z^m) - 2L_g\left\|Z^+ - Z^m\right\|_F \tag{17}$$

$$\ge \gamma_m - 2L_g\left\|Z^+ - Z^m\right\|_F > 0. \tag{18}$$

Thus the measurement winner remains the unique posterior winner. No true label enters this certificate. □

For the specified mean-pooling linear head,

$$g_c(Z) = w_c^\top\left(\frac{1}{T}\sum_{t=1}^{T} z_t\right) + b_c, \qquad L_g = \frac{\max_c \|w_c\|_2}{\sqrt{T}} \tag{19}$$

is a valid constant by Cauchy–Schwarz. Combining Proposition 1 across observed steps gives a sufficient, possibly conservative, certificate with

$$\left\|Z^+ - Z^m\right\|_F \le \left(\sum_{t=1}^{T} \min\left\{\rho_{\max}\delta_t, \frac{\rho_{\max} b_t}{2}\sqrt{d\eta_{\mathrm{ref}}}\right\}^2\right)^{1/2}. \tag{20}$$

One can compute a tighter pairwise certificate using $\|w_{c_m} - w_c\|_2 \left\|T^{-1}\sum_t (z_t^+ - z_t^m)\right\|_2$ for every competitor c. The bound is evaluated on sequences with a measurement comparator at every step; it is not applied to fabricated measurements in complete outages.

If a scalar task loss $\mathcal{J}(Z, y)$ is $L_{\mathcal{J}}$-Lipschitz on the relevant states, then

$$|\mathcal{J}(Z^+, y) - \mathcal{J}(Z^\star, y)| \le L_{\mathcal{J}}\left(\sum_t \|e_t\|_2^2\right)^{1/2}. \tag{21}$$

For cross-entropy composed with the specified linear head, its logit gradient has norm at most $\sqrt{2}$, so one may take $L_{\mathcal{J}} = \sqrt{2}\,\|W\|_2/\sqrt{T}$ as a valid Lipschitz constant. This controls loss perturbation relative to the chosen reference; it does not bound a difference between independently trained models or prove lower classification error.

A.3 FINITE-TIME ERROR, INCLUDING EMPTY STEPS

Proof of Theorem 1. At an observed step,

$$e_t = (1-\rho_t)(z_t^\star + \varepsilon_t^m) + \rho_t A z_{t-1}^+ - z_t^\star \tag{22}$$

$$= \rho_t A e_{t-1} + (1-\rho_t)\varepsilon_t^m - \rho_t \xi_t. \tag{23}$$

At an empty step the definition $z_t^+ = Az_{t-1}^+$ gives $e_t = Ae_{t-1} - \xi_t$. Both cases imply $\|e_t\|_2 \le q_t \|e_{t-1}\|_2 + d_t$. Repeated substitution proves Eq. (13), taking every empty product as one. State dependence of ρ_t does not invalidate this pathwise bound; no independence assumption is used.

For the all-observed case let $a = \|A\|_2$, assume $\|\varepsilon_t^m\|_2 \le E_m$ and $\|\xi_t\|_2 \le E_K$, and write $\gamma = \rho_{\max} a < 1$. Since $d_t \le E_m + \rho_{\max} E_K$,

$$\|e_t\|_2 \le \gamma^t \|e_0\|_2 + (E_m + \rho_{\max} E_K)\frac{1-\gamma^t}{1-\gamma}. \tag{24}$$

This is a sufficient worst-case bound, not an estimate of actual recognition error.

For bounded outages, let two successive observed times be τ and $\tau + h$, where $1 \le h \le L+1$. Expanding through $h-1$ prediction-only steps and the final observed update gives

$$\|e_{\tau+h}\|_2 \le \rho_{\max} a^h \|e_\tau\|_2 + E_m + \rho_{\max} E_K \sum_{j=0}^{h-1} a^j \le \kappa_L \|e_\tau\|_2 + C_L, \tag{25}$$

where $C_L = E_m + \rho_{\max} E_K \sum_{j=0}^{L} a^j$. If $\kappa_L < 1$, the error at observation times is bounded by the corresponding geometric series, with limiting upper bound $C_L/(1-\kappa_L)$. Between observations, for $0 \le k \le L$,

$$\|e_{\tau+k}\|_2 \le a^k \|e_\tau\|_2 + E_K \sum_{j=0}^{k-1} a^j. \tag{26}$$

There are only finitely many such k, so all errors are bounded. The initial outage must also have length at most L, or its finite initial error must be treated separately. This proves the sufficient block-contraction condition. □

The condition is conservative: actual coefficients may be much smaller than $\rho_{\max}$. Conversely, $\rho_{\max} a < 1$ alone is inadequate for unrestricted gaps when $a > 1$. Even a finite-time bound is useful when the sufficient asymptotic condition fails, but failure of the condition proves neither stability nor instability. For example, the default cap $\|A\|_2 \le 1.1$ and $\rho_{\max} = 0.6$ imply only the loose worst-case bound $\kappa_8 \le 0.6\,1.1^9 \approx 1.415 > 1$. Thus the cap alone cannot certify eight-step outages. Table 7 should be read as an empirical robustness result; Eq. (14) is claimed only conditionally, not asserted to hold for every trained checkpoint.

Proof of Corollary 1. Equation (25) gives $\left\|e_{\tau_{n+1}}\right\|_2 \le \kappa_L \|e_{\tau_n}\|_2 + C_L$ for successive observation times. Iterating this affine recursion yields $\|e_{\tau_n}\|_2 \le \kappa_L^n \|e_{\tau_0}\|_2 + C_L(1-\kappa_L^n)/(1-\kappa_L)$. Taking $n \to \infty$ proves Eq. (15). □

A.4 NON-ANTICIPATORY ACQUISITION AND BUDGET FEASIBILITY

Proof of Proposition 3. At test time every candidate feature and score is computed from $\mathcal{I}_{t,k}$, so the deterministic tie-broken maximizer is $\mathcal{I}_{t,k}$-measurable and cannot change when only unacquired current values are perturbed. Algorithm 1 selects a candidate only when $c_i \le b$, then updates $b \leftarrow b - c_i$, so b remains nonnegative and the cumulative paid request cost cannot exceed the initial budget B. Requested indices are removed from future candidate sets, hence each group is requested at most once and at most M requests can occur within a step.

It remains to verify the closed-loop information claim. The initial state z_0^+ is fixed. Suppose the propagated prior z_t^- is a function only of previously acquired observations and allowed metadata. Within step t, each successful request adds only its paid reading to the current information set, and every provisional or committed update is a deterministic function of z_t^- and those acquired

readings. Hence z_t^+ is measurable with respect to the acquired history at the end of the step. Because $z_{t+1}^- = Az_t^+$ is deterministic given z_t^+, induction over t proves that no deployment state or future action can depend on an unpaid current reading before that reading is requested. □

B Implementation and Training Specification

The following specification defines the implementation used for the reported experiments. Dimensions, optimization settings, and hyperparameter searches are fixed as described below.

B.1 Representation, prefix head, and gate

One latent step contains 8 raw samples per group. Each branch flattens its standardized sample block and per-channel validity mask, then applies a linear layer to 128 units, GELU, and a linear layer to 64 units. Split this into a 32-dimensional common feature and a 32-dimensional sensor-specific feature. Group identity is encoded by separate branch parameters. The fusion MLP maps the 64-dimensional masked mean concatenated with the M-bit acquired mask to 128 hidden units and a $d = 64$ output, followed by LayerNorm. No unacquired branch is evaluated at deployment. Accumulate acquired features in ascending group-index order for deterministic, order-independent full-set evaluation.

The task logits are $W\bar{z} + b$, where $\bar{z}$ is the mean of committed states, including prediction-only steps, up to the current prefix. Final sequence prediction uses all T states. The affine head is shared between measurement and posterior supervision. Initialize $A = 0.95I$. After each optimization step, optionally enforce the specified default norm cap by

$$A \leftarrow A \min\{1, 1.1/\|A\|_2\}. \tag{27}$$

The norm cap is a training choice and does not by itself guarantee the bounded-outage condition. The gate MLP has input size $2d + 3$, hidden width 64, GELU, and one sigmoid output. Set $a_t^{\text{out}} = \min\{\text{consecutive preceding empty steps}, T\}/T$.

At time t, $p_t(S)$ is computed by replacing the current provisional state in the prefix mean, while keeping all earlier committed states fixed. For $S = \varnothing$, use z_t^-. Prefix logits never use future values. Measurement supervision uses the mean of available measurement states within the prefix, and includes a term only at prefixes whose current acquired set is nonempty. If no supervised measurement prefix exists, its loss contribution is zero; the posterior task loss still applies.

B.2 Loss definitions and reductions

With $\mathcal{O} = \{t : S_t \neq \varnothing\}$ and $\bar{z}_{1:t}^m = |\mathcal{O} \cap [1,t]|^{-1} \sum_{s \in \mathcal{O}, s \le t} z_s^m$,

$$\mathcal{L}_m = \frac{1}{\max\{1, |\mathcal{O}|\}} \sum_{t \in \mathcal{O}} \text{CE}(\text{softmax}(W\bar{z}_{1:t}^m + b), y), \tag{28}$$

$$\mathcal{L}_{\text{Brier}} = \sum_{c=1}^{C} (p_T(c) - \mathbf{1}\{y = c\})^2. \tag{29}$$

Cross-entropy is $-\log p(y)$. Task, Brier, and prefix losses are averaged over sequences; the dynamics term averages only observed steps per sequence and is zero if none exist.

For the information regularizer, form the covariance C_z of observed measurement states in the training minibatch. For fewer than two states, set this regularizer to zero. Let λ_j be the eigenvalues of $C_z + \epsilon I$, $p_j = \lambda_j / \sum_k \lambda_k$, and $\epsilon = 10^{-5}$. Define

$$\mathcal{R}_{\text{info}} = \frac{1}{d} \sum_j [\max\{0, 1 - \sqrt{(C_z)_{jj} + \epsilon}\}]^2 + \log d + \sum_j p_j \log p_j. \tag{30}$$

The variance term addresses the zero-variance loophole of normalized spectral entropy. This regularizer encourages diversity; it does not prove statistical sufficiency or identify a physical state.

Train the candidate network with squared error for $U_t^{i\star}$ and $I_t^{i\star}$ and mean squared descriptor error $\left\| \widehat{h}_{t,c}^i - \text{sg}(h_{t,c}^i) \right\|_2^2 / 32$. Average over training prefixes and candidates. Sample up to four affordable unacquired candidates per prefix uniformly; evaluating all candidates is a diagnostic variant.

Counterfactual teacher evaluations use the same provisional state interface and detached targets. The complete training sample may supply these labels, but its unacquired values never enter the policy input. Naturally absent training channels supply no target, and are not imputed using held-out data.

B.3 HARD-ACTION POLICY TRAINING

During joint training, logits consist of r_t^i/τ for affordable, available, unrequested groups, plus a zero STOP logit. Sample one hard action. Only its reading is added to the estimator. This is not a straight-through weighted mixture of all current sensors.

For a minibatch of N sampled sequences, use batch-averaged representation and auxiliary losses and set

$$J = \mathcal{L}_{\text{rep}} + \lambda_{\text{aux}}\mathcal{L}_{\text{aux}} + \frac{\lambda_c}{N}\sum_{n=1}^{N}\sum_{t,i\in D_{n,t}} c_i, \tag{31}$$

where $D_{n,t}$ includes any failed requests. Let $\mathcal{A}$ collect all acquisition decisions in that minibatch (with the sequence index suppressed). A Monte Carlo surrogate is

$$\widehat{\mathcal{L}} = J + \sum_{(t,k)\in\mathcal{A}} \operatorname{sg}\big(J - b_v(\mathcal{I}_{t,k})\big)\log\pi(a_{t,k}\mid\mathcal{I}_{t,k}) + \lambda_v\sum_{(t,k)\in\mathcal{A}}\big(b_v(\mathcal{I}_{t,k}) - \operatorname{sg}(J)\big)^2. \tag{32}$$

The positive sign is appropriate for minimizing a cost. Ordinary backpropagation through J and the score-function term account for the differentiable and discrete paths, respectively. The baseline is detached in the policy term and is independent of the current action conditional on the prefix. Its regression shares no estimator parameters; use a separate width-64 MLP. The common minibatch return accounts for the batch-covariance term and is valid but can be high-variance; this is a practical limitation rather than a convergence claim.

Table 2: Defaults and finite validation-search choices used in the experiments.

Item	Default	Search / convention
Latent dimension	64	$\{32, 64, 128\}$
$\rho_{\max}$	0.6	$\{0.3, 0.6, 0.8\}$
m_{ref}	0.5	$\{0.25, 0.5, 1\}$
η_{ref}	training median	warm-up innovations; floor 10^{-6}
λ_K	0.1	$\{0.01, 0.1, 1\}$
$\lambda_m, \lambda_I, \lambda_B$	$1, 0.01, 0.1$	fixed in primary protocol
$\lambda_{\text{aux}}, \lambda_v$	$1, 0.1$	fixed in primary protocol
λ_c	0.05	$\{0, 0.01, 0.05, 0.1\}$
α_u, α_ν	$1, 1$	fixed in primary protocol
α_r, α_c	$0.1, 0.05$	each $\{0, 0.05, 0.1\}$
Optimizer	AdamW	learning rate 3×10^{-4}
Batch / weight decay	128 / 10^{-4}	same feasible batch for all models
Epochs / warm-up	100 / 10	patience 15 after warm-up
Policy temperature	$1 \to 0.2$	linear anneal over 50 joint epochs
Isolation/acquisition seeds	$11, 23, 37, 53, 71$	same data splits and masks

Warm up the estimator using uniformly sampled nonempty available subsets and dropout masks; if none are available, use the prediction-only branch. Fix η_{ref} from training-only innovations after warm-up, and initialize the candidate heads from training-only counterfactual targets before joint optimization. Use 10 candidate-pretraining epochs, with the estimator frozen, before the joint stage. At most 100 estimator epochs are used including warm-up; head pretraining is additional and must be charged to training compute. Sample 20 configurations from the stated finite search with seed 0, include the default, and select by validation macro-F1, then NLL, then smaller model size. Retain the chosen configuration for all reported seeds; do not select a separate best test seed.

Algorithm 1 One non-anticipatory recognition sequence under a sensing budget

Require: Prior A, branch/fusion encoders, gate, task and candidate heads; costs c; budget B
$z_0^+ \leftarrow 0$; clear history, innovation records, and last-read ages
for $t = 1, \dots, T$ **do**
  $z_t^- \leftarrow A z_{t-1}^+$; $D \leftarrow \varnothing$; $S \leftarrow \varnothing$; $b \leftarrow B$
  Read availability metadata u_t, not current sensor values
  loop
    $z^{\mathrm{prov}} \leftarrow z_t^-$ if $S = \varnothing$; otherwise apply Eqs. (4)–(7) using S and fixed z_t^-
    $C \leftarrow \{i : u_t^i = 1,\ i \notin D,\ c_i \le b\}$
    if $C = \varnothing$ **then**
      break
    end if
    Evaluate candidate heads and scores using only $\mathcal{I}_{t,k}$
    At test: pick smallest-index maximizer in C; STOP if its score ≤ 0
    At training: sample one action from the masked softmax including STOP
    if action is STOP **then**
      break
    end if
    $D \leftarrow D \cup \{i\}$; $b \leftarrow b - c_i$; request x_t^i
    if request succeeds **then**
      Encode x_t^i; record its innovation against the pre-request prediction
      $S \leftarrow S \cup \{i\}$; add only the paid reading to current information
    end if
  end loop
  Commit $z_t^+ \leftarrow z_t^-$ if $S = \varnothing$, else recompute one update from S
  Save paid cost $\sum_{i \in D} c_i$, successful mask, states, and allowed history
end for
return $\mathrm{softmax}(W(T^{-1} \sum_t z_t^+) + b_{\mathrm{head}})$

C Non-Anticipatory Acquisition Algorithm and Relation Interface

Every policy input is produced without evaluating an unrequested current branch. Let D denote requests and $S \subseteq D$ successful acquisitions. Under reliable availability $D = S$; otherwise failed requests consume budget and are excluded from further consideration.

C.1 Candidate descriptors and redundancy

The candidate MLP input consists of the prior (d), provisional posterior (d), mean previous committed state (d, zero initially), availability and successful-acquisition masks ($2M$), learned candidate identity embedding (16), candidate cost (1), remaining budget (1), normalized time since its last successful reading (1), and a sequence-start indicator (1). Its total input dimension is $3d + 2M + 20$. The prefix state is computed from the currently acquired set using the fixed prior. Remaining budget is based on all requests, not only successful ones. Use two width-128 GELU hidden layers, with unconstrained scalar outputs U, I and a 32-dimensional predicted common descriptor. Selection feasibility separately excludes all requested indices, including failures.

For $j \in S$, define

$$G_t^{ij} = \sigma\Big(\cos_0(\widehat{h}_{t,c}^{\,i}, h_{t,c}^{j}) + \mathrm{corr}_0(V^i, V^j)\Big). \tag{33}$$

Here $\cos_0$ is cosine similarity, set to zero if either vector norm is below 10^{-8}. The stored branch innovation at a successful acquisition is the acquired common feature minus its prediction computed *before that acquisition*. For each pair, V^i, V^j contain innovations at common timestamps in the preceding 8 latent steps, never the current step. Flatten matching vectors; compute their centered correlation, or zero if fewer than two common timestamps or either centered norm is below 10^{-8}. This avoids aligning irregular observations by array position. The two relation weights are fixed to one in the primary configuration. No innovation for the current hidden reading is stored or queried.

For exactly-k uniform-cost selector isolation, disable voluntary STOP until k successful acquisitions; this control assumes reliable availability and all groups present. The normal deployment policy uses STOP. No selector claims a nontrivial advantage at full acquisition under a common frozen estimator.

Cost and complexity. Candidate descriptors depend on the current provisional posterior, acquired mask, and remaining budget, so they are recomputed after each acquisition. At round k, evaluating redundancy for the $M - k$ remaining candidates against k acquired groups costs $O((M - k)kd_c)$ with $d_c = 32$; summing over rounds gives a conservative $O(M^3 d_c)$ worst-case relation cost for the algorithm as specified. Candidate-head evaluation contributes $O(M^2)$ network evaluations over a complete step, one prior matrix–vector product costs $O(d^2)$ per time step, and only acquired branches are encoded. If candidate descriptors are frozen within a step, caching reduces the relation term to $O(M^2 d_c)$, but that approximation is not assumed in the reported algorithm. With at most nine groups in the evaluated wearable settings, the selector remains a small-group procedure.

D Evaluation Protocol and Statistical Reporting

Public dataset facts below follow the cited providers. We use the stated grouping, splits, windowing, and optimization choices consistently across the reported experiments.

D.1 Data grouping and splits

UCI-HAR. The smartphone dataset contains six activities from 30 subjects, with 50 Hz inertial signals and provided 128-sample windows (Anguita et al., 2013; Reyes-Ortiz et al., 2013). Use two groups: body acceleration and body gyroscope, each with three axes. Exclude total acceleration and engineered 561-dimensional features. Retain the official train/test subject partition. For validation, sort training subject IDs by the hexadecimal SHA-256 digest of the ASCII string `TSCA-val-ID` and reserve the first $\lceil 0.2n \rceil$ IDs. Each provided window is an independent sequence of 16 latent steps; never concatenate windows as if their physical ordering were established.

RealWorld HAR. The dataset covers eight activities, 15 participants, and seven body positions (Sztyler & Stuckenschmidt, 2016). Group accelerometer and gyroscope axes by chest, forearm, head, shin, thigh, upper arm, and waist, giving seven six-channel groups. Use 50 Hz alignment, 128-sample windows, stride 64, and 16 latent steps. Subject IDs sorted numerically are assigned to five folds by zero-based rank modulo five. In outer fold k, use fold k for test, $(k+1) \bmod 5$ for validation, and the other three for training. Report every subject in a held-out test fold once per seed. The seven-point selector-isolation table uses this grouping and uniform costs.

REALDISP. The dataset provides nine sensor positions, 33 activities, and ideal, self-placement, and induced-displacement scenarios (Baños et al., 2014; Banos et al., 2012). Use three accelerometer and three gyroscope axes per position, excluding magnetometer and quaternion channels. Use 50 Hz signals, 128-sample windows, stride 64, and the same subject-rank five-fold rule as above. The primary missing-observation protocol pools eligible recordings from the available scenarios within each subject split. All REALDISP benchmark numbers reported here refer to this pooled protocol.

OPPORTUNITY. The source data support multiple recognition tasks and contain recordings from four participants (Chavarriaga et al., 2013). We use four-class locomotion (standing, walking, sitting, lying), *not* the gesture task. Use the five upper-body IMUs, six accelerometer/gyroscope channels per group, at the native 30 Hz rate. Exclude shoe and ambient sensors. Use 80-sample windows, stride 40, and 10 latent steps. In four outer folds, test subject k, validate subject $(k + 1) \bmod 4$, and train on the other two; all recordings of a subject remain in that split. Determine the window target by majority label including the null label; break ties by lowest label ID and discard windows whose majority is null. Store the exact label-name mapping in the manifest rather than guessing label codes.

Preprocessing boundary. Split subjects and recordings before forming windows. Compute per-channel normalization statistics from training data only, with standard-deviation floor 10^{-6}. Align timestamps using preceding samples only, with a maximum carry-forward age of 100 ms. No interpolation may consult a future timestamp. Remaining invalid values are represented by zero

after normalization plus the per-channel validity mask. A group with no valid value in a block is unavailable; known channel-validity metadata are assumed accessible before requests, but neither signal amplitudes nor signal-quality estimates are free. At least one valid value permits a paid, partially corrupted reading. This availability assumption must be revisited for devices that cannot report status without a paid query.

D.2 Missingness, budgets, and sequences

Apply masks at the group/block level after preprocessing. Training draws a missing probability uniformly from $\{0, 0.25, 0.5, 0.75\}$ per sequence and uses independent Bernoulli availability for each group and block. Do not condition on at least one available group; complete outages exercise the explicit prediction-only branch. Random-missingness tests use each rate separately. For a block-outage test, choose one group uniformly and one valid start uniformly, then remove that group for $L \in \{1, 2, 4, 8\}$ latent steps; run a separate *all-group* outage test with the same lengths. Label these two tests differently.

The normal policy is evaluated at budgets $B \in \{0.25, 0.5, 0.75, 1\}$, with realized costs reported separately. In the primary wearable benchmark, each dataset uses one fixed positive proxy vector c normalized so that $\sum_i c_i = 1$; all methods within that dataset use the same vector, and the costs in Table 10 are computed from those run configurations. The selector-isolation frontier is a distinct control that uses uniform costs $c_i = 1/M$, so its one-group cost is not intended to equal the single-position reference costs in Table 10. These quantities are acquisition proxies rather than measured transmission energy. Acquisition below the minimum configured group cost legitimately produces an empty observation.

Save a mask manifest before training and reuse it for all methods and seeds at test time. Derive each mask generator seed from the experiment ID and recording/window ID, independently of method identity. Reset hidden state and acquisition history at the start of every sequence, even for adjacent or overlapping windows. Non-anticipatory inference is required within each sequence; offline smoothing variants are separately labeled and excluded from the primary online comparison.

Artifact boundary for the aggregate primary rows. Table 1 and Appendix Table 10 report completed three-run deployment measurements, but the retained submission artifact does not contain the corresponding run manifests identifying the three primary seeds, numerical dataset-specific proxy vectors, or the per-dataset operating-point selection. We therefore neither infer those fields from rounded aggregates nor present the primary rows as an exact executable reproduction recipe. Their role is descriptive deployment comparison. Mechanism attribution is evaluated separately under five-seed matched protocols whose seeds and comparison interfaces are stated explicitly.

D.3 Metrics and uncertainty

Compute macro-F1 over the predefined label set, minimum per-class recall, NLL, and ECE with 15 equal-width confidence bins. Clip probabilities only for numerical log evaluation at 10^{-8} and renormalize. Empty calibration bins contribute zero. A class absent from the entire evaluation partition makes its recall undefined for a full-class minimum rather than being silently set to zero or removed. Pool held-out predictions across outer folds within each training run before computing the primary metric. Table 1 and Appendix Table 10 report mean and SD over three primary benchmark runs. The estimator-, selector-, update-, and predictor-isolation experiments use the five seeds $11, 23, 37, 53, 71$; their compact tables report the across-seed mean unless an SD is shown explicitly. Fold variability is not substituted for seed variability.

Any population-level paired uncertainty analysis would require resampling held-out subjects or recordings as intact clusters rather than treating overlapping windows as independent. No cluster-bootstrap confidence intervals are reported for the primary or estimator/update/predictor tables. The task-supervision control instead reports a paired-seed resampling interval over its five matched training runs and is labeled accordingly. Small differences elsewhere are interpreted descriptively rather than assigned significance labels; this is especially important for OPPORTUNITY, which has only four subjects.

Let $\widehat{y}^m$ and $\widehat{y}^+$ be the same-checkpoint predictions. Define

$$\text{Harm} = \frac{\sum_n \mathbf{1}\{\widehat{y}_n^m = y_n,\ \widehat{y}_n^+ \neq y_n\}}{\sum_n \mathbf{1}\{\widehat{y}_n^m = y_n\}}, \tag{34}$$

$$\text{Repair} = \frac{\sum_n \mathbf{1}\{\widehat{y}_n^m \neq y_n,\ \widehat{y}_n^+ = y_n\}}{\sum_n \mathbf{1}\{\widehat{y}_n^m \neq y_n\}}. \tag{35}$$

Use "not defined" for an empty denominator. These conditional rates are not comparable by simple subtraction. Table 4 computes each rate within each training seed and then averages the five seedwise rates; it is not a pooled ratio across seeds. Certificate coverage is the fraction satisfying Eq. (12); a certified prediction change is a correctness bug in the certificate or implementation, not ordinary stochastic model error.

D.4 Task-supervision attribution control

We isolate the supervision path on REALDISP with five matched seeds $(11, 23, 37, 53, 71)$. ON and OFF branches start from the same 10-epoch label-free warm-up state and optimizer state, use the same frozen gate, costs, budget, acquisition replay, masks, and complete-outage manifests, and train the same classifier head. The only intended difference is whether posterior/measurement label losses backpropagate through the representation and predictor; in OFF those states are detached for label losses while the non-label dynamics/information objectives remain active.

Table 3: Matched task-supervision attribution on REALDISP. Complete-outage length L is measured in latent steps. Values are mean $\pm$ sample SD over five paired seeds. The interval is the percentile interval obtained by exhaustive resampling of the five paired seed differences (5^5 index tuples); it captures seed-pair variation only and is not a subject-level confidence interval.

L	OFF F1	ON F1	Paired ΔF1	Paired-seed interval
1	0.7460 ± 0.0081	0.7560 ± 0.0072	0.0100 ± 0.0028	$[0.007583, 0.011950]$
4	0.7280 ± 0.0096	0.7470 ± 0.0096	0.0190 ± 0.0049	$[0.015249, 0.022751]$
8	0.6890 ± 0.0131	0.7240 ± 0.0120	0.0350 ± 0.0068	$[0.030127, 0.040723]$

E Baseline Implementations and Fair Controls

For each external baseline, the citation identifies the source method and the text below specifies our task/interface adaptation. All numerical rows use these stated adaptations and the same information-access constraints. The uncited mask-aware Transformer is an in-paper non-anticipatory protocol baseline rather than an adaptation attributed to another publication.

Measurement and simple interpolation. Use the same group encoders, fusion projection, and mean-pooling task head. Measurement-only sets $\rho = 0$ on observed steps and uses a declared missing-state baseline (zero state) on empty steps. The *prediction-before-update* control uses z_t^- for current prefix logits, then assimilates the paid current measurement to form the committed state used to predict $t + 1$. Thus $z_0^+ = 0$ does not force a constant sequence prediction, and the acquired reading is counted in cost. In the isolated update experiment, include only sequences with a current measurement at every compared step; the separate outage experiment tests the empty-state choice. Fixed interpolation selects one scalar $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$ by validation only. The simple learned gate uses exactly the gate MLP input and width, with $\rho = \rho_{\max}\sigma(f_\omega)$ but no multiplicative factors. "Without innovation" and "without missingness" remove only the named factor, retaining all other components and the same $\rho_{\max}$.

Recurrent and imputation controls. GRU-D receives acquired values, masks, and times since acquisition; a linear task head replaces its original dataset-specific output (Che et al., 2018). SAITS is rerun at each prefix using only samples available through the current time and the paid-group mask, followed by the same branch/classifier interface (Du et al., 2023). This prefix-only use is an adaptation evaluated under the same non-anticipatory information boundary. BRITS and CSDI are related reconstruction approaches rather than numerical baselines in the reported tables; using future

samples would make such a comparison offline (Cao et al., 2018; Tashiro et al., 2021). The mask-aware Transformer uses two layers, four heads, width 64, temporal causal masking, and no token for an unpaid sensor value. This is the concrete baseline labeled "Mask-aware Transformer" in Table 1 and Appendix Table 10; it is selected by validation under the same dataset budget and information boundary, never by test performance. The similarly named row in Table 5 is its estimator-isolation rerun under replayed identical acquisitions, so its score need not equal the jointly trained deployment score exactly.

Learned filters and direct Koopman–Kalman work. RKN retains its latent prediction and uncertainty update, with the branch observation embedding and recognition head adapted to this task (Becker et al., 2019). KalmanNet uses the same learned linear latent transition and an identity latent observation interface; its learned gain network supplies the update (Revach et al., 2022). Both receive the same acquired evidence. The K^2VAE adaptation retains its native Koopman/Kalman and variational objectives, attaches the recognition head, and tunes the auxiliary-loss weight on validation data (Wu et al., 2025). This tests a substantive direct alternative but is not a claim to reproduce that paper's forecasting benchmark. These adaptations preserve the defining update mechanism of each baseline while matching the recognition interface and acquired evidence.

Predictor replacement. Identity continuation changes only A to I. Unregularized linear prediction removes the dynamics consistency term while retaining task supervision. A width-adjusted two-layer GELU MLP replaces Az for the nonlinear predictor. A GRU receives the previous posterior, with its own hidden state reset per sequence; this extra memory is disclosed. We match total predictor parameters within 10% where possible; equal latent width alone is not treated as evidence of equal model size. Each replacement is retrained; the gate, cost policy, masks, and validation rule are held fixed.

Acquisition baselines. For random acquisition, generate one seeded random ranking without inspecting hidden values. Predicted-entropy acquisition uses only the U head with U/c scores. For the generative-surrogate adaptation (Li & Oliva, 2021), condition a surrogate on exactly the allowed history/acquired set, draw 32 counterfactual candidate samples with fixed test randomness, and average the same prefix entropy-reduction criterion before dividing by cost. Its conditional model must be fit on training data only. This is a task/interface adaptation under the common non-anticipatory information boundary. Full acquisition is a common-estimator reference, not a mathematical upper bound on accuracy.

Three isolation protocols. For estimator isolation, we precompute a common acquisition manifest with seed-fixed uniform random feasible sets and replay it for every estimator. For update isolation, we first fit and freeze one encoder/head/prior checkpoint, use fixed observed subsets, and fit only each gate. We recurse separately for the multi-step effect and also evaluate each update against the *same fixed prior* for the one-step effect. For selector isolation, we freeze the whole estimator including the gate and change only the selection rule. Separately trained, jointly optimized deployment models are reported separately from the identical-checkpoint isolation tests.

F Additional Diagnostic Matrices

All populated values in this section are measured outputs from the corresponding completed diagnostic or isolation experiment.

Table 4: Frozen-representation update isolation. Weak-observation macro-F1 uses REALDISP; strong-observation macro-F1 and conditional change rates use OPPORTUNITY. All entries are means over five seeds; Harm and Repair are computed within each seed on their respective conditional denominators before averaging. Rows within a column share inputs, encoder, classifier, and acquired evidence.

Update rule	Weak F1	Strong F1	Harm (%)	Repair (%)
Measurement only	0.706	0.984	0.000	0.000
Fixed interpolation	0.714	0.953	3.500	14.900
Simple learned gate	0.748	0.981	0.700	18.500
Without innovation factor	0.753	0.972	1.500	16.100
Without missingness factor	0.755	0.980	0.800	21.000
Structured gate	0.756	0.984	0.200	14.800

Gate parameters alone are fitted. A second diagnostic fixes the prior as well. Harm and repair use different denominators and are not subtracted.

Table 5: Estimator-isolation macro-F1 at identical replayed acquisitions within each dataset. Entries are means over five training seeds. Each model is trained independently; this differs from the frozen-representation update isolation in Table 4.

Model / adaptation	UCI-HAR	RealWorld	REALDISP	OPPORTUNITY
Measurement-only network	0.853	0.869	0.704	0.983
GRU-D adaptation	0.866	0.883	0.720	0.982
SAITS prefix-only adaptation	0.870	0.886	0.726	0.981
Mask-aware Transformer	0.874	0.891	0.733	0.984
RKN adaptation	0.878	0.895	0.738	0.979
KalmanNet adaptation	0.881	0.897	0.742	0.981
K^2VAE adaptation	0.882	0.899	0.746	0.981
Simple learned interpolation	0.883	0.900	0.748	0.981
TSCA-KKF	0.889	0.907	0.756	0.984

Neural-filter and forecasting baselines use the task/interface adaptations specified in Appendix E; values shown here are from our matched evaluation rather than scores copied from the source papers.

Table 6: Complete seven-point macro-F1 frontier for the same five-seed frozen-estimator protocol as Table 12. The budget and realized cost coincide by construction.

Policy	$1/7$	$2/7$	$3/7$	$4/7$	$5/7$	$6/7$	$7/7$
Random acquisition	0.803	0.839	0.862	0.887	0.908	0.930	0.938
Predicted entropy reduction	0.827	0.861	0.885	0.907	0.923	0.934	0.938
Generative-surrogate acquisition	0.844	0.879	0.899	0.916	0.928	0.936	0.938
TSCA-KKF acquisition	0.852	0.887	0.907	0.922	0.932	0.936	0.938

Table 7: Predictor replacement on REALDISP with identical replayed acquisition trajectories and the structured gate. Entries are mean $\pm$ sample SD over five training seeds; columns vary complete-outage length in latent steps. This empirical outage test is not a certificate that Eq. (14) holds.

Predictor	$L = 1$	$L = 4$	$L = 8$
Identity continuation $A = I$	0.739 ± 0.0086	0.702 ± 0.0114	0.652 ± 0.0143
Linear, no consistency loss	0.746 ± 0.0082	0.718 ± 0.0110	0.671 ± 0.0137
Parameter-matched MLP	0.753 ± 0.0079	0.737 ± 0.0105	0.702 ± 0.0131
Parameter-matched GRU	0.754 ± 0.0075	0.740 ± 0.0100	0.710 ± 0.0125
Task-supervised Koopman	0.756 ± 0.0072	0.747 ± 0.0096	0.724 ± 0.0120

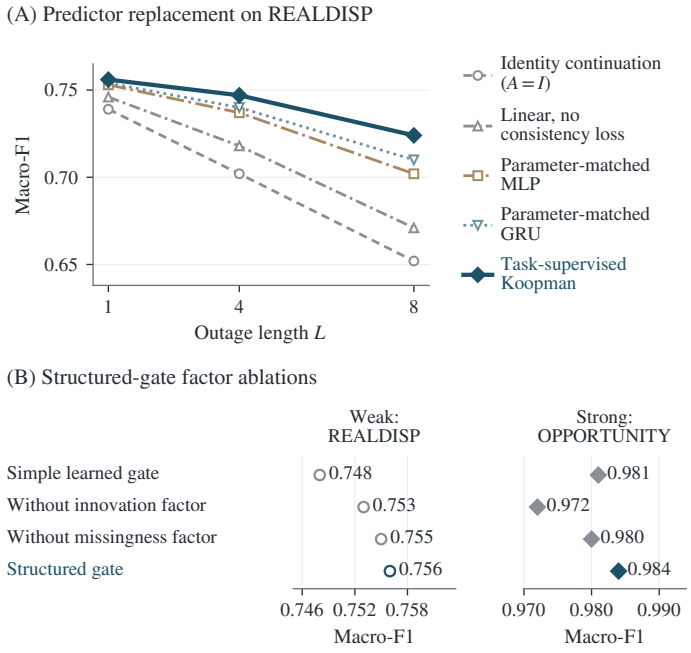

Figure 2: Matched component-isolation summary. Panel A visualizes predictor robustness across complete-outage lengths from Table 7; Panel B visualizes the structured-gate factor ablations from Table 4. No additional aggregation across datasets or seeds is introduced.

Table 8: Controlled fixed-checkpoint gate diagnostic: pooled classification accuracy under good/bad measurement evidence and good/bad temporal priors. Strong-observation columns use $m = 0$; weak-observation columns use $m = 0.5$. The same measurement input is used within each column.

Update	Good / good	Good / bad	Bad / good	Bad / bad
Measurement only	0.960	0.960	0.700	0.700
Prior only	0.955	0.640	0.955	0.640
Fixed interpolation	0.949	0.752	0.846	0.652
Simple learned gate	0.959	0.952	0.855	0.701
Structured gate	0.960	0.960	0.828	0.694

Column order is measurement / prior quality. This controlled latent perturbation test isolates update behavior rather than inferring latent reliability from wearable data. When $m = 0$, the structured update equals measurement-only exactly; under bad measurement/good prior it remains measurement-biased by design.

F.1 QUALITY-CROSSING TEST AND FAILURE CASES

For the quality-crossing diagnostic, we use a fixed held-out checkpoint and distinguish the two information sources explicitly. A good measurement is the unperturbed current measurement latent; a bad measurement adds independently seeded Gaussian noise scaled coordinate-wise by multiplier 2 of the training-state coordinate standard deviation. A good prior is the unperturbed one-step temporal prior z_t^- from the same held-out trajectory, whereas a bad prior applies the same multiplier-2 perturbation independently to that temporal prior. Table 8 reports these clean-versus-multiplier-2 endpoints. The measurement and prior are therefore distinct vectors even in the Good/Good column, which permits their standalone accuracies to differ.

Table 8 is a latent controlled test: strong-observation columns set $m = 0$ directly and weak-observation columns set $m = 0.5$ directly, rather than mapping these values to a physical subset when M is odd. This isolates gate behavior from sensor-group combinatorics. True transitions can invalidate the prior, whereas sensor corruption can invalidate the measurement; neither condition is identified by innovation magnitude alone.

The intended interpretation is deliberately limited: a structured gate that preserves a wrong confident measurement is behaving according to its measurement-anchored bias, not satisfying a general safety guarantee. Table 8 therefore serves as a controlled failure-mode diagnostic rather than evidence of universal robustness.

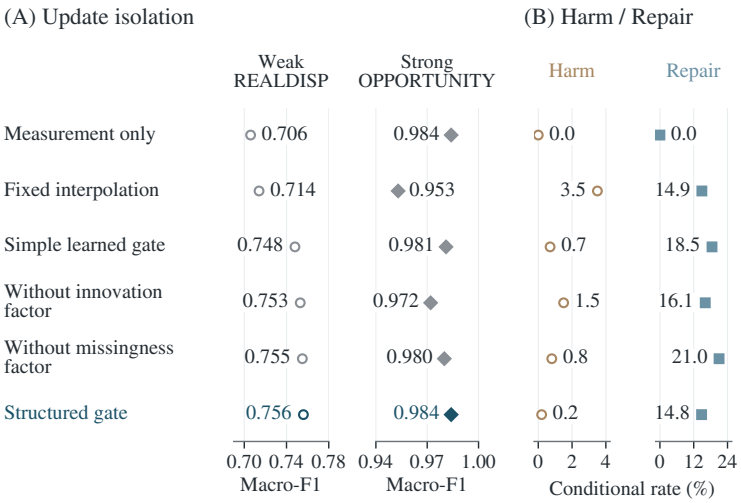

(C) Controlled measurement / prior quality

	Pooled accuracy			
	Good/Good	Good/Bad	Bad/Good	Bad/Bad
Measurement only	0.960	0.960	0.700	0.700
Prior only	0.955	0.640	0.955	0.640
Fixed interpolation	0.949	0.752	0.846	0.652
Simple learned gate	0.959	0.952	0.855	0.701
Structured gate	0.960	0.960	0.828	0.694

Figure 3: Mechanism-isolation visualizations derived from Tables 4 and 8. The panels summarize weak/strong-observation performance, harmful/repair conditional rates, and the controlled measurement/prior quality diagnostic. No additional mechanism experiment is introduced by this figure.

F.2 IMPLEMENTATION INVARIANTS

Independent NumPy reference tests check the algebraic update, finite-time bound, budget logic, and margin certificate on controlled arrays. These are pass/fail implementation assertions rather than empirical rates. They complement, rather than replace, the benchmark evaluations.

Table 9: Implementation invariants enforced by the evaluation code and reference tests.

Check	Required property
Unpaid-value intervention	Same next action with fixed allowed information and policy randomness.
Budget accounting	Total paid cost, including failures, never exceeds B; no repeated request.
Empty observation	Committed state equals the single propagated prior, not a zero-measurement mixture.
All groups acquired	$m = 0$ and posterior equals measurement exactly.
Within-step acquisition	Provisional updates always share the same prior; commit only once.
Full-budget isolation	Same checkpoint and complete acquired set produce identical predictions for every selector.
Prefix causality	Replacing future readings leaves the present action and provisional state unchanged.
Margin certificate	No class change on any sequence satisfying the strict certificate.

G Complete Wearable Benchmark Results

Table 10 reports the complete wearable benchmark metrics used in the primary evaluation, including recognition quality, calibration, and normalized acquisition cost. Baseline names denote the concrete implementations evaluated under the corresponding dataset protocol.

Table 10: Complete wearable benchmark results. Recognition and calibration metrics are mean ± SD across three training runs. Costs are realized normalized acquisition proxies from the fixed dataset-specific benchmark configuration; adaptive costs retain their across-run SD. Full- and single-position references are context rather than equal-budget competitors.

Method	Macro-F1	Min. recall	NLL	ECE	Cost
UCI-HAR					
Full sensors CNN	0.941±0.004	0.902±0.010	0.181±0.015	0.012±0.004	1.00
Observed-only CNN	0.832±0.009	0.742±0.021	0.356±0.026	0.029±0.007	0.50
Measurement-only	0.855±0.008	0.774±0.018	0.311±0.022	0.024±0.006	0.50
Fixed Koopman–Kalman	0.865±0.007	0.795±0.016	0.287±0.021	0.022±0.005	0.50
Prediction-before-update Koopman	0.839±0.010	0.751±0.023	0.334±0.024	0.028±0.007	0.50
Mask-aware Transformer	0.872±0.007	0.807±0.016	0.268±0.020	0.020±0.005	0.50
TSCA-KKF	0.889±0.006	0.832±0.014	0.236±0.018	0.016±0.004	0.47±0.02
RealWorld HAR					
Full positions CNN	0.936±0.005	0.884±0.013	0.194±0.014	0.014±0.004	1.00
Best single position CNN	0.858±0.011	0.724±0.027	0.342±0.028	0.031±0.006	0.22
Measurement-only	0.872±0.009	0.762±0.022	0.286±0.024	0.025±0.006	0.42
Fixed Koopman–Kalman	0.881±0.008	0.783±0.020	0.271±0.021	0.023±0.005	0.42
Imputation + CNN	0.864±0.010	0.746±0.026	0.317±0.025	0.030±0.007	0.42
Mask-aware Transformer	0.891±0.008	0.799±0.019	0.253±0.020	0.021±0.005	0.42
TSCA-KKF	0.907±0.006	0.824±0.016	0.221±0.017	0.017±0.004	0.39±0.03
REALDISP					
Full positions CNN	0.812±0.009	0.704±0.022	0.604±0.039	0.032±0.007	1.00
Best single position CNN	0.681±0.015	0.492±0.041	0.982±0.061	0.061±0.012	0.18
Measurement-only	0.706±0.013	0.548±0.037	0.851±0.055	0.052±0.010	0.44
Fixed Koopman–Kalman	0.713±0.012	0.562±0.034	0.814±0.048	0.050±0.009	0.44
Imputation + CNN	0.689±0.014	0.526±0.039	0.902±0.057	0.057±0.011	0.44
Mask-aware Transformer	0.731±0.012	0.588±0.032	0.759±0.045	0.046±0.008	0.44
TSCA-KKF	0.755±0.010	0.641±0.029	0.681±0.041	0.039±0.008	0.41±0.03
OPPORTUNITY					
Full on-body CNN	0.989±0.002	0.972±0.006	0.046±0.006	0.005±0.002	1.00
Observed body-IMU CNN	0.975±0.002	0.931±0.010	0.082±0.003	0.006±0.004	0.42
Measurement-only	0.984±0.004	0.955±0.017	0.059±0.012	0.006±0.002	0.42
Fixed Koopman–Kalman	0.952±0.026	0.827±0.081	0.163±0.045	0.016±0.005	0.42

Continued on next page

Table 10 continued

Method	Macro-F1	Min. recall	NLL	ECE	Cost
Mask-aware Transformer	0.984 ± 0.003	0.956 ± 0.012	0.060 ± 0.010	0.006 ± 0.002	0.42
TSCA-KKF	0.985 ± 0.003	0.963 ± 0.008	0.055 ± 0.009	0.005 ± 0.001	0.42 ± 0.00

G.1 Derived Frontier and Problem-Setting Summaries

Table 11: Derived gap-recovery and cost-efficiency summary from Table 1. Gap recovery is $(F_1 - F_1^{\text{meas}})/(F_1^{\text{full}} - F_1^{\text{meas}})$ and is reported only where the full-sensor reference leaves a nontrivial headroom. No additional experiments are introduced by this table.

Dataset	Transformer gap recovered	TSCA-KKF gap recovered	Transformer cost	TSCA-KKF cost
UCI-HAR	19.8%	**39.5%**	0.50	**0.47**
RealWorld HAR	29.7%	**54.7%**	0.42	**0.39**
REALDISP	23.6%	**46.2%**	0.44	**0.41**

Table 12: Selector isolation on RealWorld HAR with one frozen estimator, seven available groups, uniform costs, and exactly k acquisitions. Entries are means over five training seeds; full acquisition yields the same prediction for every selector by construction.

Policy	$k = 1$	$k = 3$	$k = 5$	$k = 7$
Random acquisition	0.803	0.862	0.908	0.938
Predicted entropy reduction	0.827	0.885	0.923	0.938
Generative-surrogate acquisition	0.844	0.899	0.928	0.938
TSCA-KKF acquisition	0.852	0.907	0.932	0.938

The corresponding realized costs are exactly $1/7$, $3/7$, $5/7$, and 1. All seven budget points appear in Table 6.

Table 13: Area under the seven-point macro-F1–cost frontier (AUPC) on RealWorld HAR, computed by trapezoidal integration of the measured points in Table 6 and normalized by the covered cost interval. Higher is better.

Policy	Normalized AUPC
Random acquisition	0.8828
Predicted entropy reduction	0.8988
Generative-surrogate acquisition	0.9082
TSCA-KKF acquisition	**0.9132**

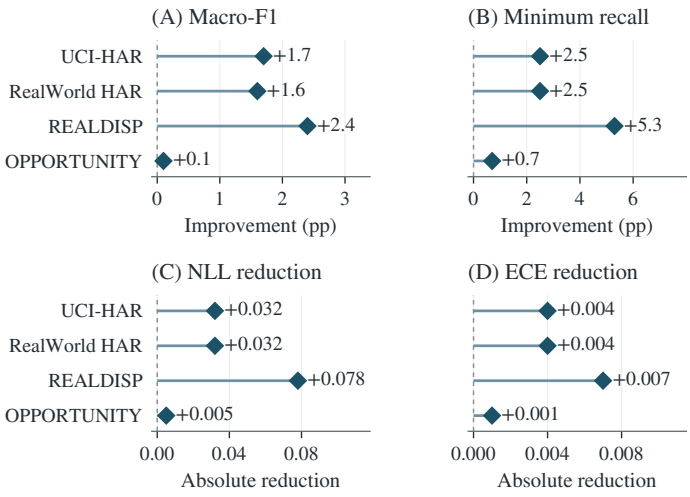

Figure 4: Derived mean differences between TSCA-KKF and the mask-aware Transformer in the primary limited-observation results of Appendix Table 10. Panels A and B report percentage-point differences in macro-F1 and minimum recall, while Panels C and D report absolute reductions in NLL and ECE. Positive values favor TSCA-KKF. All values are deterministic functions of the reported means and introduce no additional experimental runs or significance claims.

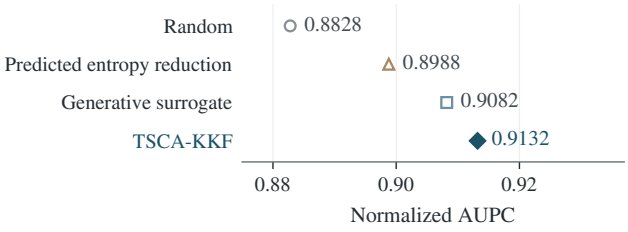

Figure 5: Selector-isolation summary derived from Tables 6 and 13. Panel A shows the seven-point uniform-cost macro-F1 frontier for the four reported policies; Panel B shows the corresponding normalized AUPC values. This is the same selector-isolation protocol summarized in Table 6.