When Does Sequential Detection Collapse to a Scalar?

A necessary and sufficient characterisation

Prakul Sunil Hiremath
Peerahammad M. Bagawan
Sahil Bhekane

Department of Computer Science and Engineering
Visvesvaraya Technological University, Belagavi, Karnataka, India
{prakulhiremath, peerahamadmb, sahilbhekane}@vtu.ac.in

PROJECT PAGE

Outline

1. **Problem:** scalar scores for $K \geq 3$ regimes
2. **Setup:** HMM, belief filter, stopping value
3. **EDS:** three linear conditions
4. **Main theorem:** scalar collapse $\iff$ EDS
5. **Consequences:** impossibility, horizon, robustness
6. **Verification:** algorithm and deployment workflow
7. **Experiments:** synthetic, CICIDS2017, controls
8. **Scope, limits, takeaways**

Central message

Scalar optimality is a property of the *model*, not the algorithm.
The optimal stopping value collapses to a scalar if and only if the model satisfies EDS.

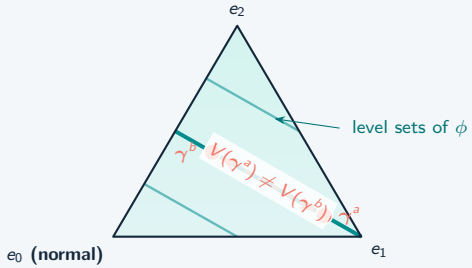
The problem: deployed detectors reduce a simplex to a scalar

What is known, and what is not

- $K = 2$ regimes: scalar thresholding of the posterior is Bayes-optimal.
- $K \geq 3$: the posterior γ_t lives on the $(K-1)$-simplex Δ^{K-1}.
- CUSUM variants and deep detectors still output one score.
- Open: *when* is this valid, *when* is it suboptimal, and *by how much*?

Structural question

Is there a scalar $T : \Delta^{K-1} \to \mathbb{R}$ with $V(\gamma) = \tilde{V}(T(\gamma))$? Equivalently, is V constant on the level sets of some one-dimensional statistic?



Simplex Δ^{K-1} ($K = 3$); $\phi(\gamma^a) = \phi(\gamma^b)$

Scalar sufficiency holds iff EDS: without EDS, equal-ϕ beliefs can have different optimal values.

Setup: HMM, Bayes filter, and the scalar statistic

Model

- Hidden regimes $s_t \in [K] = \{0\} \cup \mathcal{S}_{\mathrm{atk}}$; $0 =$ normal, $\mathcal{S}_{\mathrm{atk}} =$ attack regimes.
- Irreducible, aperiodic $\mathbf{\Pi}$; emissions p_k with $\ell \le p_k \le u$.
- Posterior $\gamma_t \in \Delta^{K-1}$, updated by the Bayes filter $\mathcal{B}$.

Scalar statistic

$$\phi(\gamma) = \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \, \gamma^{(k)} \in [0, C], \quad C = \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k.$$

Stopping problem

- $t^\star = \inf\{t \ge 1 : s_t \in \mathcal{S}_{\mathrm{atk}}\}$ (attack onset).
- Lorden-type constraint: $\mathbb{P}(\tau < t^\star - W) \le \epsilon$ (do not alarm too early).

$$V(\gamma) = \sup_{\tau \in \mathcal{T}_\epsilon} \mathbb{E}_\gamma \left[(t^\star - \tau)^+ \right]$$

Bellman operator

$$(\mathcal{T} g)(\gamma) = \max\big(r(\gamma),\, \mathbb{E}_{\mathbf{y}|\gamma}[g(\mathcal{B}(\gamma, \mathbf{y}))]\big)$$

ρ-contraction on $C(\Delta^{K-1})$, so V is its unique fixed point. The constraint enters through a penalty in r (Lagrangian form).

Extended Decision Sufficiency (EDS): three linear conditions

RO: Rank-One Emissions

$p_k = c_k\, h$ for all $k \in \mathcal{S}_{\mathrm{atk}}$: attack regimes share one shape h and differ only in scale c_k.

MF: Markov Factorisation

$\sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \Pi_{jk} = \alpha\, c_j$ for all $j \in [K]$: expected attack mass after one step is $\alpha\, \phi(\boldsymbol{\gamma})$.

NF: Normal Factorisation

$\Pi_{j0} = \beta\,(1 - c_j/C)$ for all $j \in [K]$: the normal-state probability is also a function of $\phi(\boldsymbol{\gamma})$.

Intuition

To the sensor, all attack "flavours" look the same up to loudness (RO). They switch among themselves in a way that depends only on total attack mass (MF, NF). So one never needs to know *which* attack is active, only *how much* attack mass there is.

Cheap to check

EDS is verifiable in $O(K^2)$ operations from $(\boldsymbol{\Pi}, \mathbf{c}, h)$ alone. It is a condition on the *model parameters*, not on a trained detector.

A concrete EDS model: $K = 3$ with $c_1 \neq c_2$

Take $c_1 = 1$, $c_2 = 2$ (regime 2 is twice as "loud"), $C = 3$, $\alpha = \frac{3}{4}$, $\beta = \frac{1}{2}$.

$$\mathbf{\Pi} = \begin{pmatrix} \frac{1}{3} & \frac{7}{12} & \frac{1}{12} \\ \frac{1}{3} & \frac{7}{12} & \frac{1}{12} \\ \frac{1}{6} & \frac{1}{6} & \frac{2}{3} \end{pmatrix}, \qquad \phi(\gamma) = \gamma^{(1)} + 2\gamma^{(2)}.$$

- MF: $\Pi_{j1} + 2\Pi_{j2} = \frac{3}{4}\, c_j$ for each j.
- NF: $\Pi_{j0} = \frac{1}{2}(1 - c_j/3)$ for each j.
- All entries positive: irreducible and aperiodic.

Why it matters

Scalar Bayes-optimal stopping is achievable for a strictly broader class than the identically-distributed-attack case. The set of EDS models with $c_1 \neq c_2$ has non-empty interior in parameter space.

Strict hierarchy

observation sufficiency $\subsetneq$ EDS $\subsetneq$ unconstrained

Gaussian HMMs with distinct attack means violate RO, so EDS is not generic.

Main result: the EDS theorem

Theorem (Extended Decision Sufficiency)

Under the HMM, boundedness, and non-degeneracy assumptions, the following are equivalent:

(i) V is constant on every level set of ϕ (i.e. $V \in \mathcal{G}$).

(ii) There exist *some* measurable scalar $T : \Delta^{K-1} \to \mathbb{R}$ and $\tilde{V}$ with $V = \tilde{V} \circ T$.

(iii) EDS holds (RO, MF, and NF).

When EDS holds

- $Z_t = \phi(\gamma_t)$ is a scalar Markov process.
- ϕ is minimally decision-sufficient: any scalar T through which V factors is ϕ-measurable.

Simplex Collapse Principle

Scalar optimality is a property of the model, not the algorithm. When EDS fails, no scalar reduction, whatever the model class, capacity, or training, can be Bayes-optimal.

Proof idea 1 (EDS $\Rightarrow$ scalar): Bellman closure

Four algebraic steps

1. MF: expected next attack mass $= \alpha\,\phi(\gamma)$.
2. NF: next normal mass $= \beta\,(1-\phi(\gamma)/C)$.
3. RO: the filter's normaliser depends on γ only via $\phi(\gamma)$.
4. So $\mathcal{T}$ maps functions of ϕ into functions of ϕ; the closed set $\mathcal{G}$ contains the unique fixed point V.

Scalar filter recursion

$$Z_{t+1} = \frac{h(\mathbf{y}_{t+1})\,\alpha Z_t}{p_0(\mathbf{y}_{t+1})\,\beta\big(1 - Z_t/C\big) + h(\mathbf{y}_{t+1})\,\alpha Z_t}$$

Z_{t+1} depends on $(Z_t, \mathbf{y}_{t+1})$ alone, so $\{Z_t\}$ is Markov and the Bellman recursion closes on one dimension.

The only thing the filter must remember is Z_t.

Proof idea 2 (scalar $\Rightarrow$ EDS): coupling contradiction

Contrapositive

Suppose some EDS condition fails but $V = \tilde{V} \circ \phi$. For each failure mode (RO, MF, NF) construct beliefs γ^a, γ^b with

$$\phi(\gamma^a) = \phi(\gamma^b), \qquad \mathrm{TV}(m_a, m_b) \geq 2\delta_0 > 0,$$

where m_i are the one-step predictive densities.

- A coupling lemma converts the TV gap into a gap $|\mathbb{E}_{m_a} V(\mathcal{B}) - \mathbb{E}_{m_b} V(\mathcal{B})| \geq c_0 > 0$, using only continuity of V.
- A case analysis of the Bellman `max` then gives $(\mathcal{T}V)(\gamma^a) \neq (\mathcal{T}V)(\gamma^b)$.

Contradiction

$V = \mathcal{T}V$ and V constant on ϕ-level sets force $(\mathcal{T}V)(\gamma^a) = (\mathcal{T}V)(\gamma^b)$. The constructed gap says otherwise.

Ingredients

- Lipschitz bound on the filter $\mathcal{B}$: $L_{\mathcal{B}} \leq u^2/\ell^2$.
- ρ-contraction of $\mathcal{T}$ (Doeblin).
- No structure on V assumed beyond continuity.

Consequence 1: an irreducible cost when EDS fails

Theorem (Value-function separation)

If EDS fails with structural TV gap $\Delta p = \mathrm{TV}(m_a, m_b) > 0$ for the witnessing pair, and $\kappa = \ell/(Ku^2)$, then for *every* measurable $T : \Delta^{K-1} \to \mathbb{R}$ and $\tilde{V} : \mathbb{R} \to \mathbb{R}$,

$$\sup_{\gamma \in \Delta^{K-1}} \left| V(\gamma) - \tilde{V}(T(\gamma)) \right| \;\geq\; \frac{\kappa\, W\, \Delta p}{2} \;>\; 0.$$

Algorithm-independent

- Applies to every scalar statistic, learned or hand-crafted.
- The constant is computable from model parameters.

No fix from scale

More data, larger models, or a better optimiser cannot remove this loss. It is attributable to scalar reduction alone.

Consequence 2: closed-form Anticipation Horizon

$$\mathrm{ELD}^\star = \mathbf{v}_0^\top \left(I - P_Z^{\leq}\right)^{-1} \mathbf{1} + O(WC/N)$$

Ingredients

- $Z_t = \phi(\gamma_t)$ discretised to N levels on $[0, C]$.
- $P_Z^{\leq}$: sub-stochastic transition matrix on states below the optimal threshold $\delta^\star$.
- $\mathbf{v}_0$: initial distribution of Z_0.
- $I - P_Z^{\leq}$ is invertible since $\rho(P_Z^{\leq}) < 1$.

What is new

The fundamental-matrix hitting-time formula is classical. The contributions are: (i) EDS is exactly the condition under which $\{Z_t\}$ is Markov, so the formula applies; (ii) an explicit discretisation error $O(WC/N)$.

Without EDS

$\{Z_t\}$ is not Markov in general, and no closed form for $\mathrm{ELD}^\star$ exists.

Consequence 3: robustness under approximate EDS

$$\mathrm{ELD}^\star - \mathbb{E}\big[(t^\star - \tau_\eta^\star)^+\big] \;\le\; \frac{W\eta}{1-\rho} + \frac{W\eta}{\Delta_{\min}}$$

Reading the bound

- $\eta = \eta_{\mathrm{RO}} + \eta_{\mathrm{MF}} + \eta_{\mathrm{NF}}$ is the total EDS deviation.
- First term: cumulative path-measure error (Mitrophanov-type accumulation).
- Second term: sensitivity of the Lorden threshold; $\Delta_{\min}$ is the initial distance to threshold.

A practical diagnostic

η is computable from parameters and estimable from held-out data. Small η implies modest loss from scalarisation.

Where it is loose

Both terms blow up as $\rho \to 1$ (slow mixing) or $\Delta_{\min} \to 0$ (near indifference). In experiments the bound is conservative.

Verification algorithm and deployment workflow

EDS verification (structural check, not a detector)

1. **RO:** $c_k = \|p_k\|_1 / \|p_1\|_1$;
   $\eta_{\mathrm{RO}} = \max_k \|p_k / c_k - p_1\|_1$.
2. **MF:** $r_j = \sum_{k \in \mathcal{S}_{\mathrm{atk}}} c_k \Pi_{jk}$; $\hat{\alpha} = r_0 / c_0$;
   $\eta_{\mathrm{MF}} = \max_j |r_j - \hat{\alpha} c_j|$.
3. **NF:** $\hat{\beta} = \Pi_{00} / (1 - c_0 / C)$;
   $\eta_{\mathrm{NF}} = \max_j |\Pi_{j0} - \hat{\beta}(1 - c_j / C)|$.
4. **Return** $\eta = \eta_{\mathrm{RO}} + \eta_{\mathrm{MF}} + \eta_{\mathrm{NF}}$; declare EDS if $\eta \le \epsilon_{\mathrm{EDS}}$.

Deployment pipeline

1. Fit the HMM.
2. Compute η.
3. If η is small, form $Z_t = \phi(\gamma_t)$.
4. Discretise $[0, C]$; estimate P_Z on held-out data.
5. Set $\delta^\star$ to meet the Lorden-type constraint.
6. Predict $\mathrm{ELD}^\star$; deploy $\tau = \inf\{t : Z_t \ge \delta^\star\}$.

If η is large

Keep the full belief state instead of any scalar reduction.

Experiments: four structural predictions

Predictions tested

- **P1** The horizon formula is exact under EDS.
- **P2** Graceful degradation as $\hat{\eta}$ grows.
- **P3** Complete scalar collapse when EDS fails.
- **P4** The lead-time advantage is structural, not regularisation.

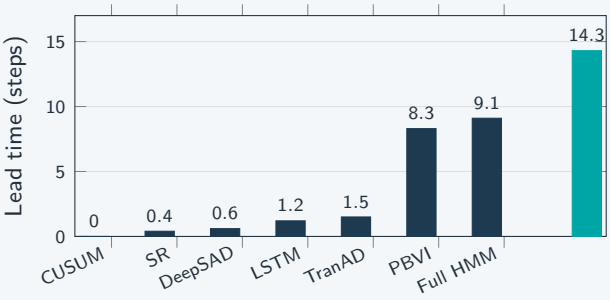
Mean ± 1.96 SE over 5 seeds; paired t-test at $\alpha = 0.05$.

P1: 72 synthetic configurations

Setting	MAE (steps)	R^2	$\hat{\eta}$
$K = 3$	0.28 ± 0.05	.994	$< .001$
$K = 4$	0.35 ± 0.06	.990	$< .001$
$K = 5$	0.40 ± 0.07	.981	$< .001$
All	$\mathbf{0.34 \pm 0.04}$	**.989**	—

Predicted and observed lead times agree within discretisation error (≤ 0.4 steps for $K = 5$, $N = 200$).

P2/P3: CICIDS2017 benchmark (matched FPR ≤ 0.04)



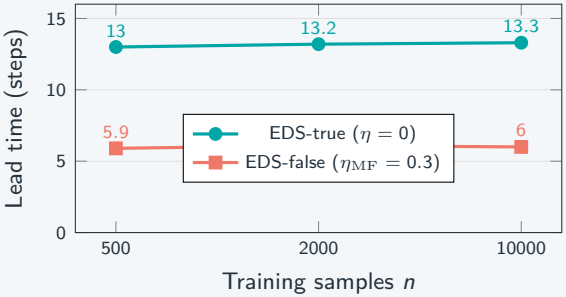
Five seeds; training on days 1–3, evaluation on days 4–5.

Method	F1	Lead	Pred.
PBVI	.85	8.3	7.9
Full HMM	.84	9.1	8.7
EDS	**.92**	**14.3**	**13.8**

Reading

EDS leads PBVI by $+6.0$ steps and full HMM by $+5.2$ steps ($p < 10^{-3}$). Deep pointwise detectors reach ≤ 1.5 steps. Estimated $\hat{\eta}^{\mathrm{raw}} = 0.058 \pm 0.009$.

P4: structure, not regularisation



Controlled experiment

- $K = 3$ synthetic HMM, matched FPR ≤ 0.01.
- Both conditions have the *same* three free parameters (α, β, c_1).
- EDS-false perturbs the MF direction only.

Finding

The gap (≈ 7.2 steps) does not shrink as n grows. With parameter counts matched, it is attributable to the structural property.

$\hat{\eta}$ is estimated non-parametrically on held-out data, independently of the detection model, so the evaluation is not circular.

Controlled violations: MF breakdown and RO failure

MF violation ($\hat{\eta} = 0.31$)

Condition	EDS	Full HMM	Bound
EDS intact (.041)	14.3	9.1	≤ 30.3
MF broken (.31)	6.1	8.7	≤ 56.1

- Observed degradation: 8.2 steps, inside the bound.
- The full posterior is unaffected, so the loss is specific to scalar reduction.
- The bound is worst-case over belief pairs, hence conservative.

RO violation ($\hat{\eta} = 0.41$)

0.3
scalar lead (steps)

11.7
full posterior lead

Detection without anticipation

F1 is unchanged: the model still identifies attacks. But all anticipatory capability is lost under scalarisation.

When does EDS hold in practice?

Plausible

Attack regimes share a common behavioural structure and differ mainly in intensity or scale: related exploit variants, repeated attack patterns with similar traffic signatures. Expect small $\hat{\eta}$ and near-scalar behaviour.

Unlikely

Heterogeneous, qualitatively distinct failure modes, such as industrial systems with unrelated fault types. Expect large $\hat{\eta}$ and substantial loss from scalar reduction, consistent with the separation theorem.

Compute $\hat{\eta}$ before deployment to choose between scalar and full-belief detection.

Practical workflow

Fit the HMM $\rightarrow$ compute $\eta \rightarrow$ small: scalar chain Z_t with predicted $\mathrm{ELD}^\star$ | large: retain the full belief state.

Scope, limitations, and open questions

What this work does not establish

- **Scalar only:** the characterisation covers one-dimensional statistics, not block-scalar ones.
- **Exact boundary:** RO is not a modelling choice but an exact condition; large η means keep the full belief.
- **Assumptions:** bounded emissions (truncation error is absorbed into η), stationary regimes, discrete time.
- **Bound looseness:** the robustness bound worsens as $\Delta_{\min} \to 0$ or $\rho \to 1$.
- **Optimiser scope:** the unconstrained-HMM comparison uses Baum–Welch with 50 restarts; no claim is made for every estimator.

Open questions

- **Beyond one dimension:** exact low-dimensional (block-scalar) statistics.
- **Beyond stationarity:** non-stationary regimes.
- **Beyond discrete time:** continuous-time models.
- **Beyond passive detection:** active detection and sensing.

Responsible use

A small $\hat{\eta}$ on one dataset does not certify scalar detection elsewhere. Re-estimate η under distribution shift, keep false-alarm constraints conservative, and retain human review of alarms.

Takeaways

1. **Scalar collapse has an exact characterisation.**
   The stopping value is constant on level sets of a scalar statistic if and only if RO, MF, and NF hold.
2. **When EDS fails, the loss is irreducible.**
   Every scalar rule pays a strictly positive, computable, algorithm-independent lead-time cost.
3. **When EDS holds, performance is predictable.**
   Optimal lead time is a closed-form hitting-time functional of a scalar Markov chain; predicted and observed lead times agree within 0.4 steps (MAE, 72 configurations).
4. η **is a practical diagnostic.**
   Degradation scales as $O(\eta/(1-\rho))$; on CICIDS2017, EDS reaches 14.3 steps versus 9.1 for a full HMM, and a matched-parameter control indicates the gain is structural.

Whether one number is enough is a property of the model.

Thank you

Backup: key equations

Bayesian belief update

$$[\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y})]^{(k)} = \frac{p_k(\mathbf{y}) \sum_j \Pi_{jk} \gamma^{(j)}}{\sum_\ell p_\ell(\mathbf{y}) \sum_j \Pi_{j\ell} \gamma^{(j)}}.$$

$$\phi(\boldsymbol{\gamma}) = \sum_{k \in \mathcal{S}_k} c_k \gamma^{(k)}.$$

Value function and Bellman operator

$$V(\boldsymbol{\gamma}) = \sup_{\tau \in \mathcal{T}_\epsilon} \mathbb{E}_{\boldsymbol{\gamma}}[(t^\star - \tau)^+].$$

$$(\mathcal{T} g)(\boldsymbol{\gamma}) = \max \Big\{ r(\boldsymbol{\gamma}), \mathbb{E}_{\mathbf{y}|\boldsymbol{\gamma}}[g(\mathcal{B}(\boldsymbol{\gamma}, \mathbf{y}))] \Big\}.$$

Structural assumptions

$$\text{RO:} \quad p_k = c_k h,$$

$$\text{MF:} \quad \sum_{k \in \mathcal{S}_k} c_k \Pi_{jk} = \alpha c_j,$$

$$\text{NF:} \quad \Pi_{j0} = \beta \left(1 - \frac{c_j}{C}\right).$$

Main theoretical bounds

$$\sup_{\boldsymbol{\gamma}} \left| V(\boldsymbol{\gamma}) - \widetilde{V}(T(\boldsymbol{\gamma})) \right| \geq \frac{\kappa W \Delta p}{2}.$$

$$\mathrm{ELD}^\star = \mathbf{v}_0^\top (I - P_Z^<)^{-1} \mathbf{1} + O(WC/N).$$

$$\mathrm{gap} \leq \frac{W\eta}{1-\rho} + \frac{W\eta}{\Delta_{\min}}.$$

Backup: assumptions and filter regularity

Assumptions

- **HMM:** irreducible, aperiodic $\mathbf{\Pi}$; conditionally independent observations.
- **Boundedness:** $\ell \le p_k \le u$ a.e. For Gaussian emissions, apply to a truncated or clipped model on a compact set.
- **Non-degeneracy:** $\Pi_{j0} > 0$ and some $\Pi_{jk} > 0$ with $k \in \mathcal{S}_{\mathrm{atk}}$, for each j.

Regularity (appendix of the paper)

- Filter denominator $\ge K\ell$; $\mathcal{B}$ is Lipschitz with $L_{\mathcal{B}} \le u^2/\ell^2$.
- Doeblin minorisation gives $\mathcal{T}$ a contraction rate $\rho = 1 - \varepsilon_0^{n_0} < 1$.
- $V \in C(\Delta^{K-1})$ is the unique fixed point and is non-constant.

Backup: witnessing belief pairs for each failure mode

Failure	Witnessing pair with $\phi(\gamma^a) = \phi(\gamma^b)$	Source of the gap
RO	$\gamma^a = e_{k_1}$, $\gamma^b = \mu e_{k_2} + (1-\mu)e_0$ with $\mu = c_{k_1}/c_{k_2}$	Emission shapes differ: $f = p_{k_1} - \mu p_{k_2}$ has positive mass, so $\mathrm{TV}(m_a, m_b) > 0$.
MF (given RO)	$\gamma^a = e_{j_0}$, $\gamma^b = (1-t)e_{j_0} + t\gamma^{\mathrm{ref}}$	Expected next attack mass differs: reward gap and predictive gap.
NF (given RO, MF)	$\gamma^a = e_{j_1}$, $\gamma^b = s\, e_{j_2} + (1-s)e_0$ with $s = c_{j_1}/c_{j_2}$	h-terms cancel; the gap is $p_0(\mathbf{y})D_{\mathrm{NF}}$ with $D_{\mathrm{NF}} \neq 0$.

Routing in the Bellman case analysis: RO and NF failures differ only in the continuation value (rewards equal); MF failure also differs in the reward.

Backup: estimating η and the discretisation error

Gaussian emissions, shared Σ

$$\eta_{\mathrm{RO}} = 2 \max_{k \in \mathcal{S}_{\mathrm{atk}}} \Phi\Big(-\tfrac{1}{2}\|\boldsymbol{\mu}_k - \boldsymbol{\mu}_1\|_{\boldsymbol{\Sigma}^{-1}}\Big)$$

RO holds exactly iff all attack means coincide.

Non-parametric estimate

Train pairwise regime classifiers on held-out labelled data and use a variational TV estimate. CICIDS2017 (78 features, 5-fold CV, 5 seeds):
$\hat{\eta}_{\mathrm{RO}}^{\mathrm{raw}} = 0.058 \pm 0.009$.

Discretisation

$|\mathrm{ELD}_N^\star - \mathrm{ELD}^\star| \le WC/N$.

Setting	N	Error bound
$K = 3,\ C = 2,\ W = 20$	200	≤ 0.2 steps
$K = 5,\ C = 4,\ W = 20$	200	≤ 0.4 steps

Backup: lead time vs. training size on CICIDS2017

Train n	EDS lead	Full HMM lead	PBVI lead
500	13.1	4.2	3.8
1000	13.8	7.3	6.9
2000	14.2	9.8	9.1
Full ($\sim$3400)	14.3	9.1	8.3

EDS lead time is stable from $n = 500$ (variance < 0.1 steps). Full HMM and PBVI improve with data but remain below EDS at all sample sizes, consistent with the irreducible-loss bound.

Backup: experimental details

- **HMM fitting:** Baum–Welch, 50 random restarts, convergence 10^{-6}; EDS regulariser $\ell_{\text{reg}} = \ell_{\text{HMM}} - \lambda_\eta \eta$ with $\lambda_\eta = 0.5$ (chosen by held-out F1).
- **PBVI:** 500 belief points from forward simulation; 50 backup iterations; no pruning.
- **TranAD, DeepSAD:** public implementations; hyperparameters chosen by held-out F1.
- **Scalar chain:** $N = 200$ levels on $[0, C]$; transitions via kernel-smoothed empirical estimates.
- **CICIDS2017:** 78 flow features; 10 PCA components ($> 95\%$ variance) for the HMM; raw 78 features for $\hat{\eta}^{\text{raw}}$; train days 1–3, test days 4–5; FPR ≤ 0.04.
- **Hardware and seeds:** Google Colab (NVIDIA T4), < 24 h total; seeds $\{0, 1, 2, 3, 4\}$.