Cohort-Level Protection and Individualized Inference in AI-Based Monitoring

Vishal Subedi, Snigdhansu Chatterjee

Department of Mathematics & Statistics, UMBC

Motivation

Artificial intelligence (AI) tools are increasingly used to monitor large groups of similar units (e.g., patients, credit cards), but critical decisions must often be made at the **individual level**.

We propose a framework that:

- Borrows strength across a cohort for better accuracy,
- Enables automated, individualized inference, and
- Supports early detection of system breaches (e.g., tumor growth, fraud).

Applications include:

- Cancer screening via image analysis
- Credit card fraud detection
- Cybersecurity and ecological monitoring

Our work offers a scalable, mathematically grounded approach that blends cohort-level learning with personalized monitoring, forming a key step toward digital twins and precision healthcare.

Model Structure

We observe **data for I individuals over T time points**:

$$\{(Y_{i,t}, X_{i,t}) \in \mathbb{R}^{d_1 \times d_2} \times \mathbb{R}^p \;:\; i = 1, \ldots, I;\; t = 1, \ldots, T\}$$

Main Model:

$$\mathbf{Y}_{i,t} = \mu(X_{i,t}) + U_i + E_{i,t}$$

$\mathbf{Y}_{i,t} \in \mathbb{R}^d$: Vectorized form of the $d_1 \times d_2$ **image**, obtained by stacking columns

$\mu : \mathbb{R}^p \to \mathbb{R}^d$ is the conditional mean function, where:

$$\mu(X_{i,t}) = \mathbb{E}[\mathbf{Y}_{i,t} \mid X_{i,t}]$$

$U_i \sim \mathcal{N}(0, \Sigma_u)$: **Individual-level** random effect

$E_{i,t} = \mathbf{\Phi}_i E_{i,t-1} + e_{i,t}$: VAR(1) **colored noise**, $\quad e_{i,t} \overset{\text{iid}}{\sim} \mathcal{N}_d(0, \sigma_e^2 I_d)$

Bayesian Framework

Simplified setup:

Assume a **single pixel** ($d_1 = d_2 = 1$), so,

$Y_{i,t} \in \mathbb{R}$

$\mu(X_{i,t}) = \beta_0 + X_{i,t}^\top \beta$

$U_i \sim \mathcal{N}(0, \sigma_u^2)$

$E_{i,t} = \phi E_{i,t-1} + e_{i,t}, \quad e_{i,t} \sim \mathcal{N}(0, \sigma_e^2)$

Model:

$$Y_{i,t} = \beta_0 + X_{i,t}^\top \beta + U_i + E_{i,t}$$

Target:

$$\theta_i = X_{i,t}^\top \beta + U_i$$

Likelihood:

$$Y_{i,t} \mid \mathcal{F}_{i,t-1} \sim \mathcal{N}\left(X_{i,t}^\top \beta + U_i + \phi(Y_{i,t-1} - X_{i,t-1}^\top \beta - U_i),\; \sigma_e^2\right)$$

Posterior of U_i:

$$U_i \mid \text{data} \sim \mathcal{N}(\mu_u, \sigma_v^2)$$

$$\mu_u = \frac{\sigma_v^2(1-\phi)}{\sigma_e^2} \sum_{t=2}^{T} R_{i,t}, \quad \sigma_v^2 = \left[\frac{1}{\sigma_u^2} + \frac{(1-\phi)^2(T-1)}{\sigma_e^2}\right]^{-1}$$

Adjusted Residuals:

$$R_{i,t} = Y_{i,t} - \phi Y_{i,t-1} - X_{i,t}^\top \beta + \phi X_{i,t-1}^\top \beta$$

Monitoring & Automated Inference

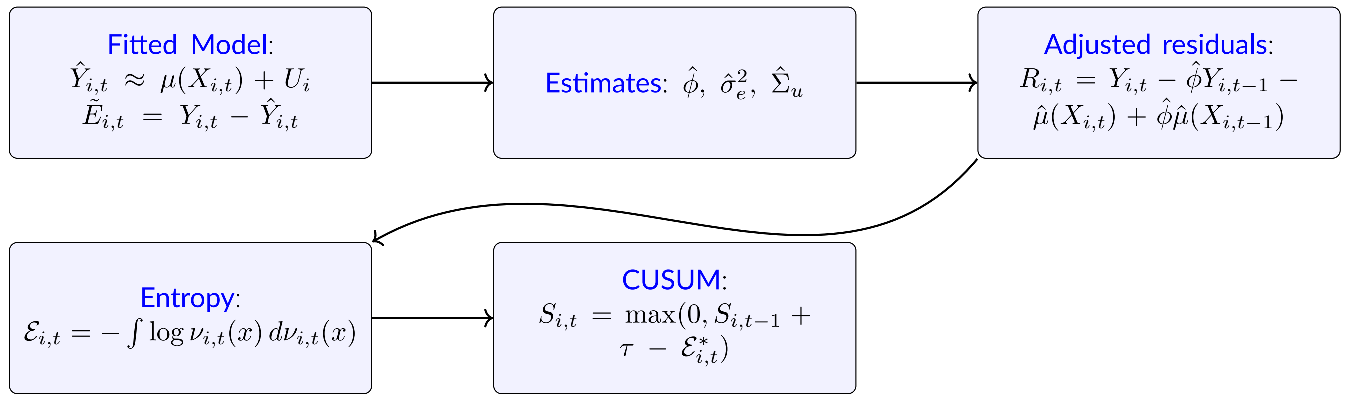

Simulation Setup

Cohort: $I = 400$ individuals, each with $T = 100$ baseline frames and $n = 100$ post-change frames. Images $Y_{i,t} \in \mathbb{R}^{d_1 \times d_2}$ evolve over time as

$$Y_{i,t} = Y_{i,t-1} + E_{i,t}$$

Noise Model: For $t \le T$, all $E_{i,t} \sim \mathcal{N}(0, \sigma^2)$. For $t > T$, a fixed patch has $E_{i,t} \sim \mathcal{N}(\delta_i, \sigma^2)$ and the rest remain zero-mean.

Group Assignment:

$$\delta_i = \begin{cases} 0.0 & i \in [1, 100] \\ 0.5 & i \in [101, 200] \\ 1.0 & i \in [201, 300] \\ 2.0 & i \in [301, 400] \end{cases}$$

Entropy: Each image is quantized into 8 bins and a PMF $p = (p_1, \ldots, p_8)$ is computed. The entropy score is:

$$U_{i,t} = \log\left(-\frac{1}{8}\sum_{j=1}^{8} p_j \log p_j\right)$$

CUSUM Monitoring: After standardizing $U_{i,t}$ using $t \le 30$, define:

$$C_{i,t} = \max\{C_{i,t-1} - U_{i,t}^* + \gamma,\ 0\}, \quad \gamma = -0.5$$

Detection Rule: Change is flagged at

$$\hat{T}_i = \min\left\{t > T \mid C_{i,t} > \max_{t \le T} C_{i,t}\right\}$$

or $\hat{T}_i = \text{NA}$ if no detection occurs.

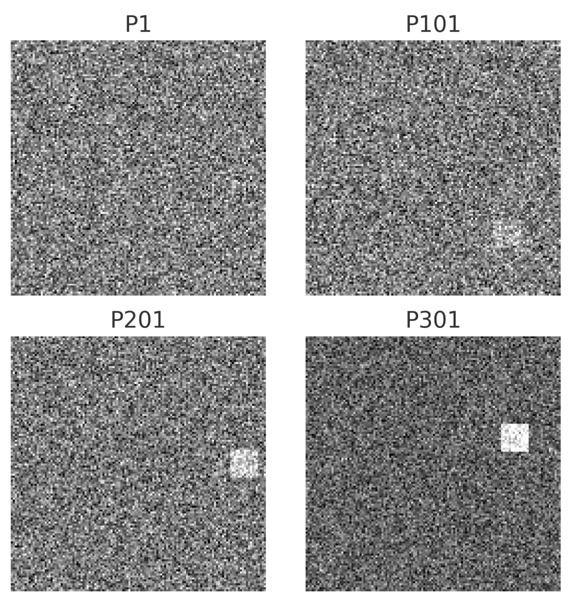

Black-and-white synthetic image samples

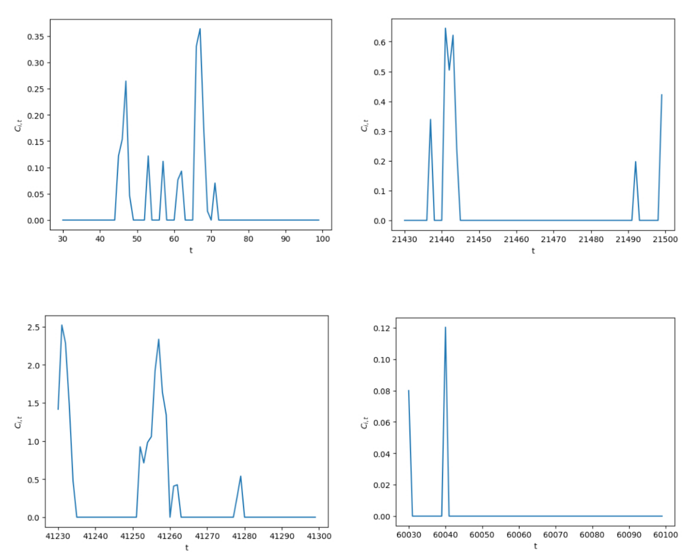

CUSUM plots

Results

Detection Time Summary: Mean and variance of post change detection times across δ-groups (based on non-NA detections):

δ	$\mathbb{E}[\hat{T}_i]$	$\text{Var}(\hat{T}_i)$
0.0	19.57	651.90
0.5	25.50	483.57
1.0	14.59	77.92
2.0	7.76	25.01

Ribbon Plot (Mean $\pm$ IQR): The plot below shows the CUSUM trajectory $C_{i,t}$ (log-scaled) over time for each δ group.

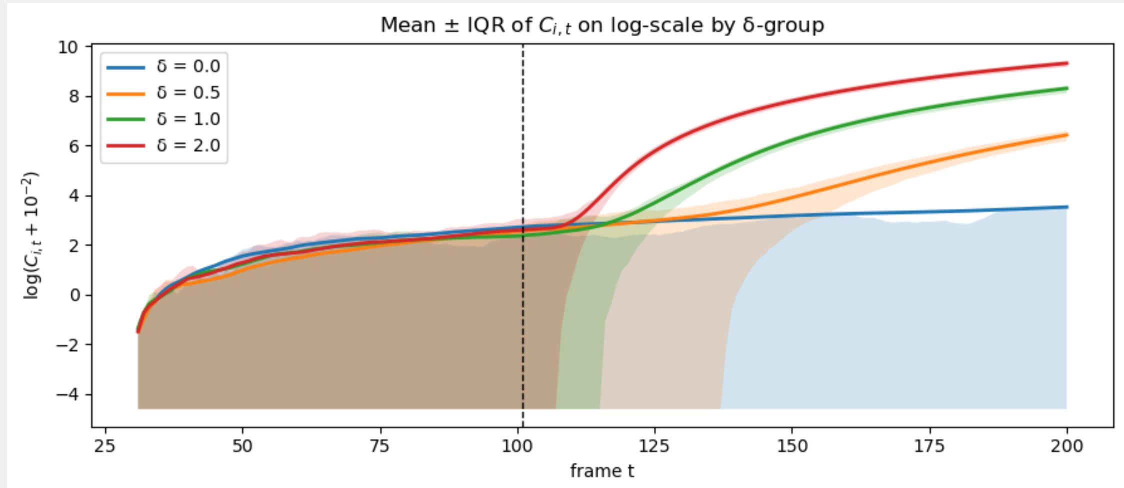

Each curve is the average $\log(C_{i,t} + 10^{-2})$ across 100 individuals in a group. Ribbons show interquartile range (IQR). Higher δ causes sharper, earlier rises.

Conclusion & Future Work

Summary. We develop a **Bayesian monitoring framework** that combines **cohort-level learning** with **individualized inference**, using an **entropy-based CUSUM** approach for timely detection in image sequences.

Simulation Highlights:
- Detects early-stage changes even at low SNR ($\delta = 0.5$)
- Shows faster detection as signal increases ($\delta = 2.0$)
- Control group ($\delta = 0$) maintains low false positives

Future Directions:
- Extend $\mu(X_{i,t})$ via deep learning architectures
- Develop theoretical guarantees for entropy-based sequential tests
- Explore other metrics: Wasserstein, graph curvature
- Apply to clinical imaging and real-time surveillance

References

H. Farooq, Y. Chen, T. T. Georgiou, A. Tannenbaum, and C. Lenglet, Network curvature as a hallmark of brain structural connectivity, Nature communications, 10 (2019), p. 4937.

G. V. Moustakides, Optimal stopping times for detecting changes in distributions, the Annals of Statistics, 14 (1986), pp. 1379–1387.

R. Sandhu, T. Georgiou, E. Reznik, L. Zhu, I. Kolesov, Y. Senbabaoglu, and A. Tannenbaum, Graph curvature for differentiating cancer networks, Scientific reports, 5 (2015), p. 12323.

R. S. Sandhu, T. T. Georgiou, and A. R. Tannenbaum, Ricci curvature: An economic indicator for market fragility and systemic risk, Science advances, 2 (2016), p. e1501495.

C. Wiedeman, A. Sarmakeeva, E. Sizikova, D. Filienko, M. Lago, J. G. Delfino, and A. Badano, T-synth: A knowledge-based dataset of synthetic breast images, arXiv preprint arXiv:2507.04038, (2025).

vsubedi1@umbc.edu