Quickest Change Detection and its Application to Learning in Nonstationary Environments

Venu Veeravalli

ECE & Statistics & CSL

University of Illinois at Urbana-Champaign

http://vvv.ece.illlinois.edu

NASITS 2026

BYU, Provo, UT

Friday, June 26, 2026

Quickest Change Detection (QCD)

$$\underbrace{X_1, \ldots, X_{\gamma-1}}_{\text{dist } \boldsymbol{P}_0} \rightarrow \underbrace{X_\gamma, X_{\gamma+1} \ \ldots}_{\text{dist } \boldsymbol{P}_1}$$

- Sequential observations whose distribution changes due to anomaly at unknown time γ
  - Need to detect change online as observations are being received
- QCD algorithm chooses stopping time τ at which change is declared
- Optimize tradeoff between
  - Detection delay
  - Frequency of false alarms

Applications of QCD

- Manufacturing systems – quality control
- Surveillance systems – intruder detection
- Power systems – detecting line outages
- Infrastructure monitoring – fault detection
- Environmental monitoring
- Seismic event detection
- Neurological disorder detection
- Pandemic monitoring
- Learning in non-stationary environments

…

Change in Distribution

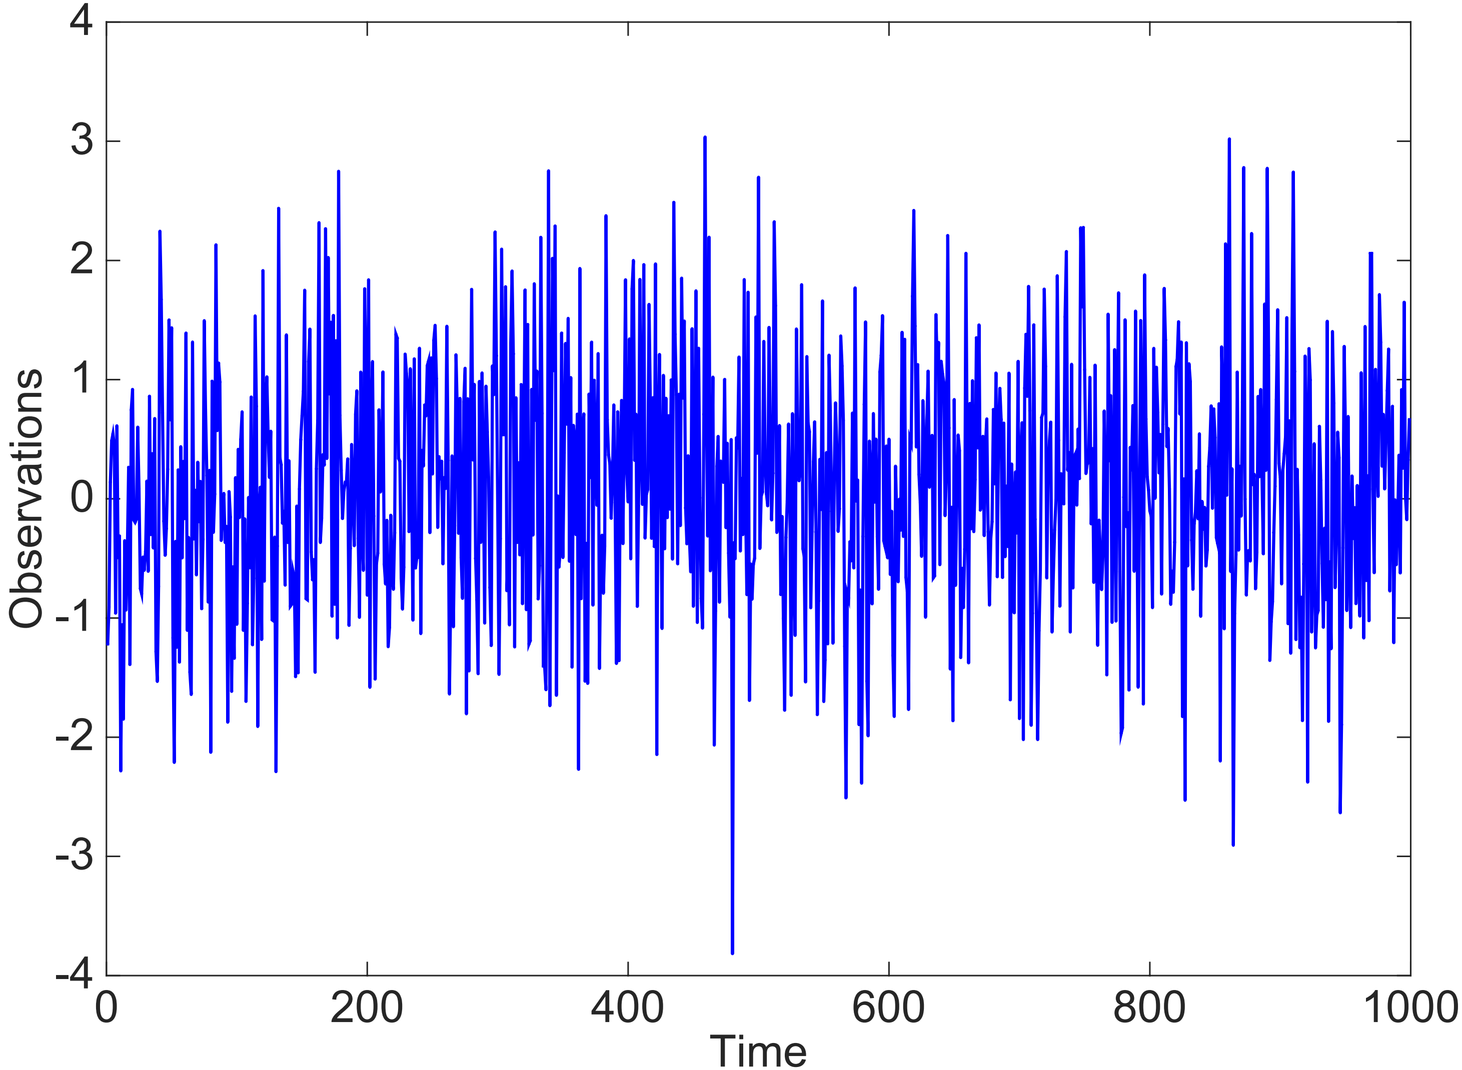

Change in Distribution

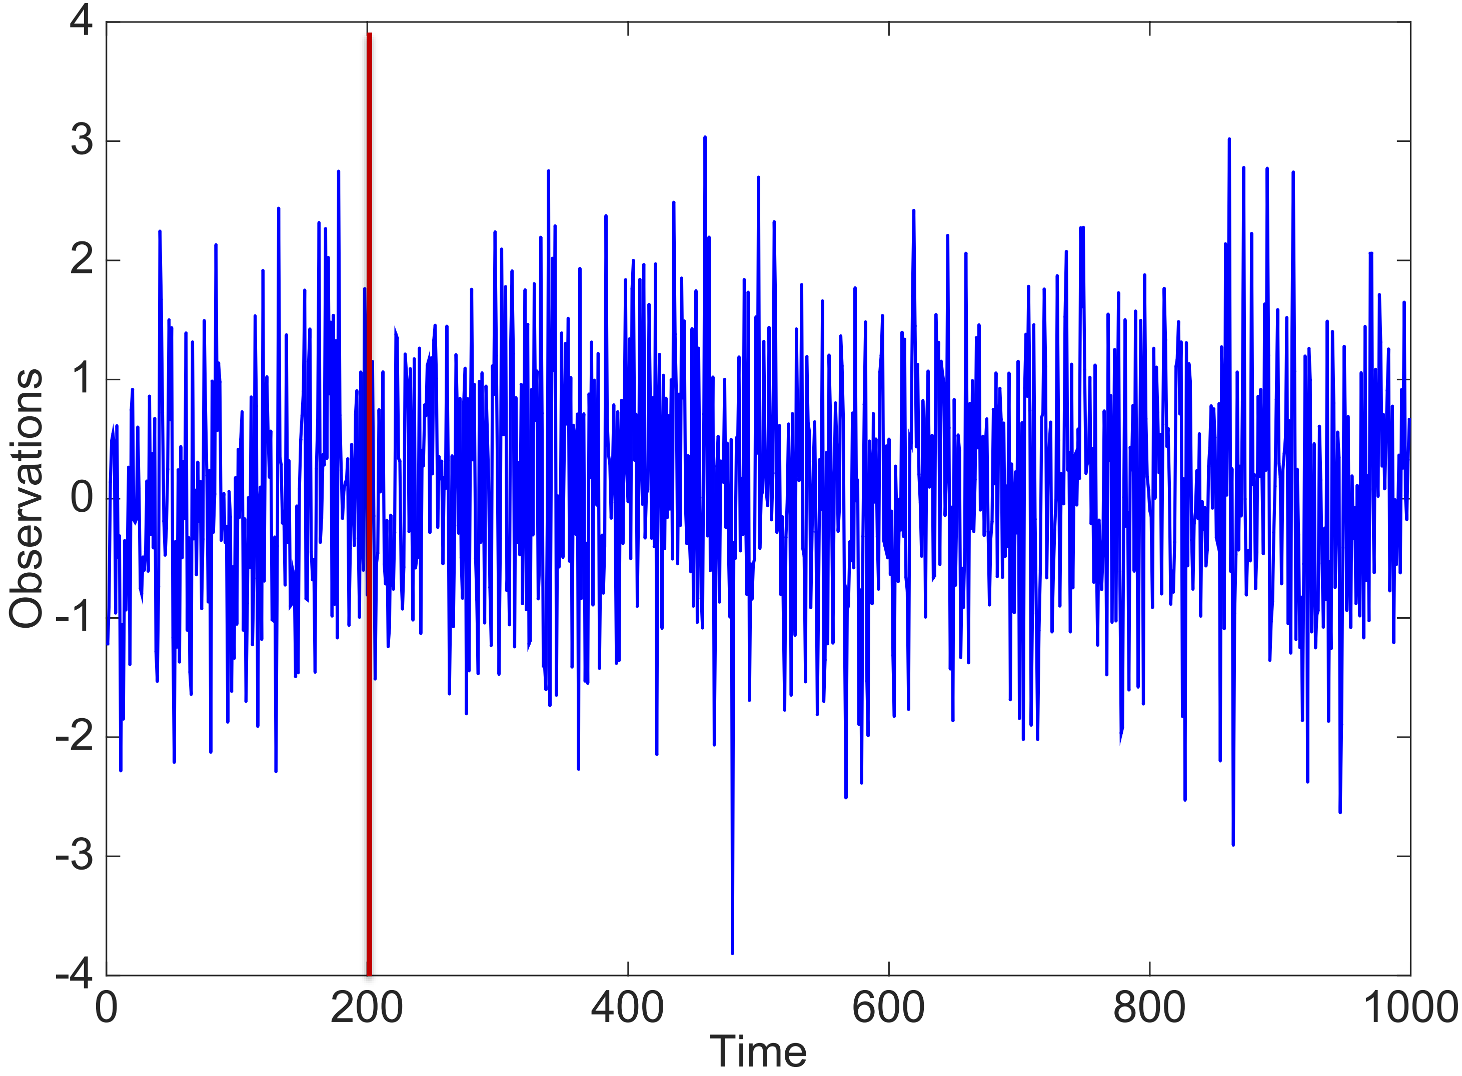

Change in Distribution

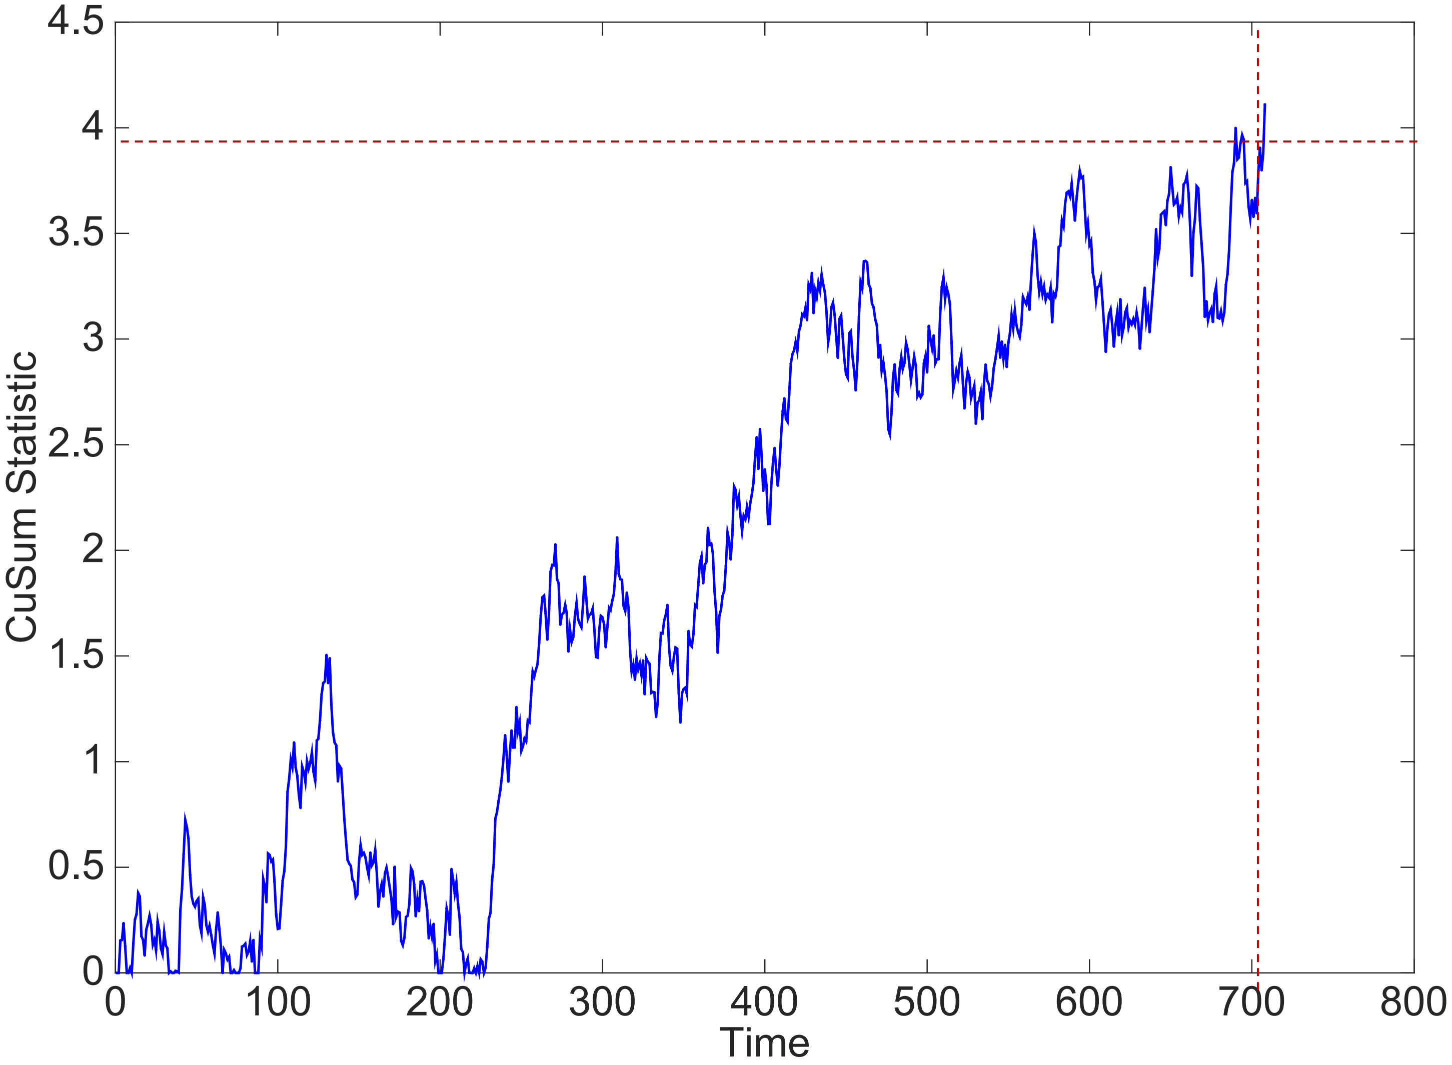

Mathematical Preliminaries

Stopping Time: Random variable τ is said to be stopping time on sequence of observations $Y_1, Y_2, \dots$ if

$$\{\tau = n\} \in \sigma(Y_1, Y_2, \dots, Y_n)$$

equivalently if $\mathbb{I}\{\tau = n\}$ is function of only $Y_1, Y_2, \dots, Y_n$ (causality)

Wald's Lemma: Suppose $Y_1, Y_2, \dots$ are i.i.d. with mean μ, and τ is a stopping time on $\{Y_n\}$ with $\mathbb{E}[\tau] < \infty$, then

$$\mathbb{E}\left[\sum_{n=1}^{\tau} Y_n\right] = \mu\, \mathbb{E}[\tau]$$

Martingale: A sequence $Y_1, Y_2, \dots$ is said to be martingale w.r.t. another sequence $X_1, X_2, \dots$ (or w.r.t. filtration $\{\mathcal{F}_n = \sigma(X_1, X_2, \dots, X_n)\}$) if for all n, $\mathbb{E}[|Y_n|] < \infty$ and $\mathbb{E}[Y_{n+1} | X_1, X_2, \dots, X_n] = Y_n$

Optional Sampling Theorem (Doob): Expected value of martingale at stopping time is equal to its initial expected value, i.e.,

$$\mathbb{E}[Y_\tau] = \mathbb{E}[Y_1]$$

Shewhart Test [1925]

$$\underbrace{X_1, \dots, X_{\gamma-1}}_{\text{i.i.d. } p_0} \longrightarrow \underbrace{X_\gamma, X_{\gamma+1} \dots}_{\text{i.i.d. } p_1}$$

- Likelihood and Log-likelihood ratio:

$$L(X_k) = \frac{p_1(X_k)}{p_0(X_k)} \qquad \ln L(X_k)$$

- Shewhart test:

$$\tau_{\text{Shew}} = \inf\{n \geq 1 : \ln L(X_n) \geq b\}$$

where b chosen to meet FA constraint

- Motivation for τ_{Shew}:

$\mathrm{E}_{p_0}[\ln L(X_k)] = -D(p_0 || p_1) < 0$ and $\mathrm{E}_{p_1}[\ln L(X_k)] = D(p_1 || p_0) > 0$

- τ_{Shew} simple to implement, but ignores past observations

Cumulative Sum (CuSum) Test [Page 1954]

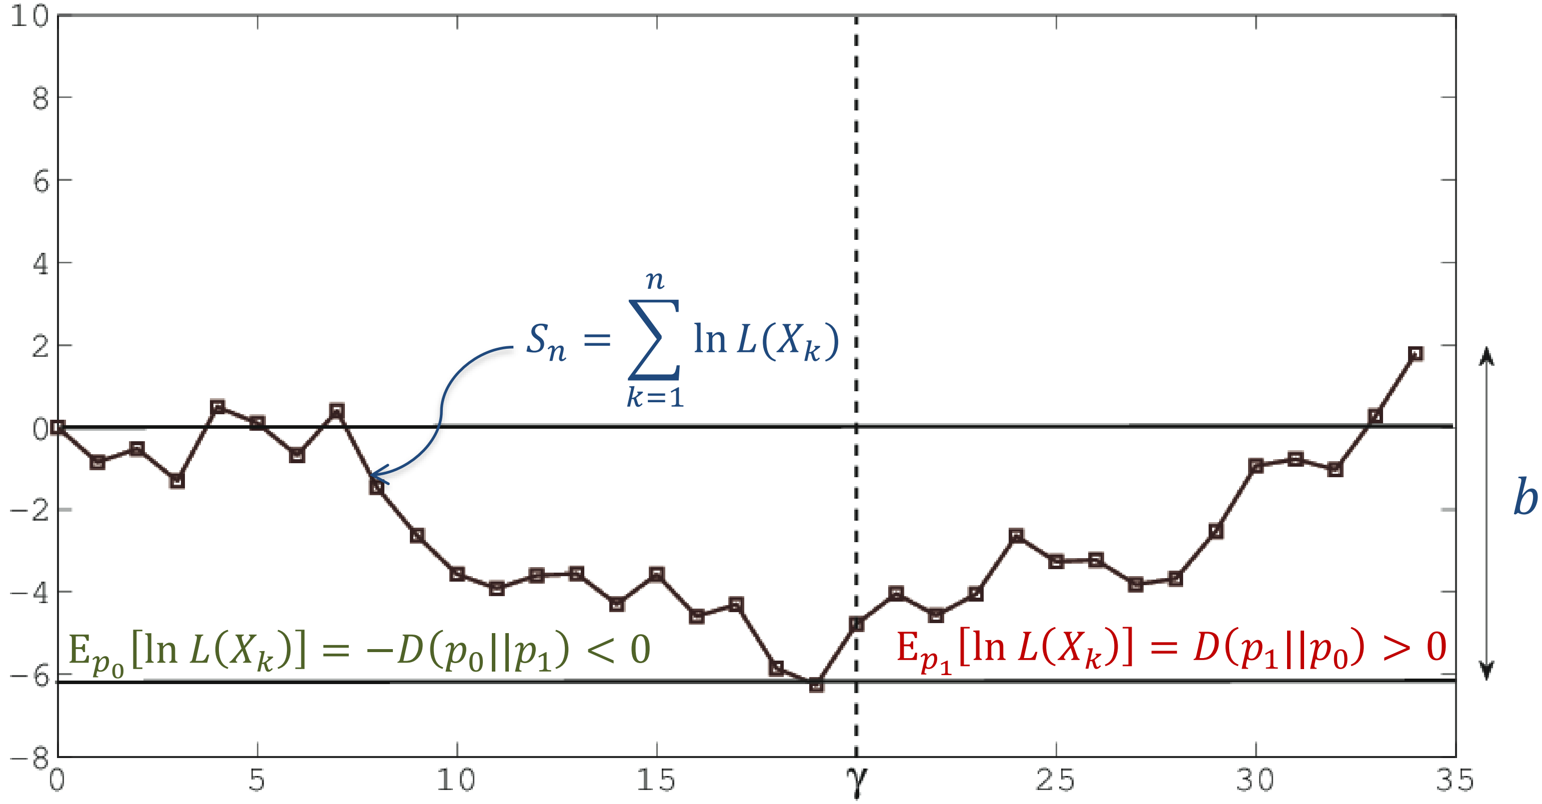

CuSum Test [Page 1954]

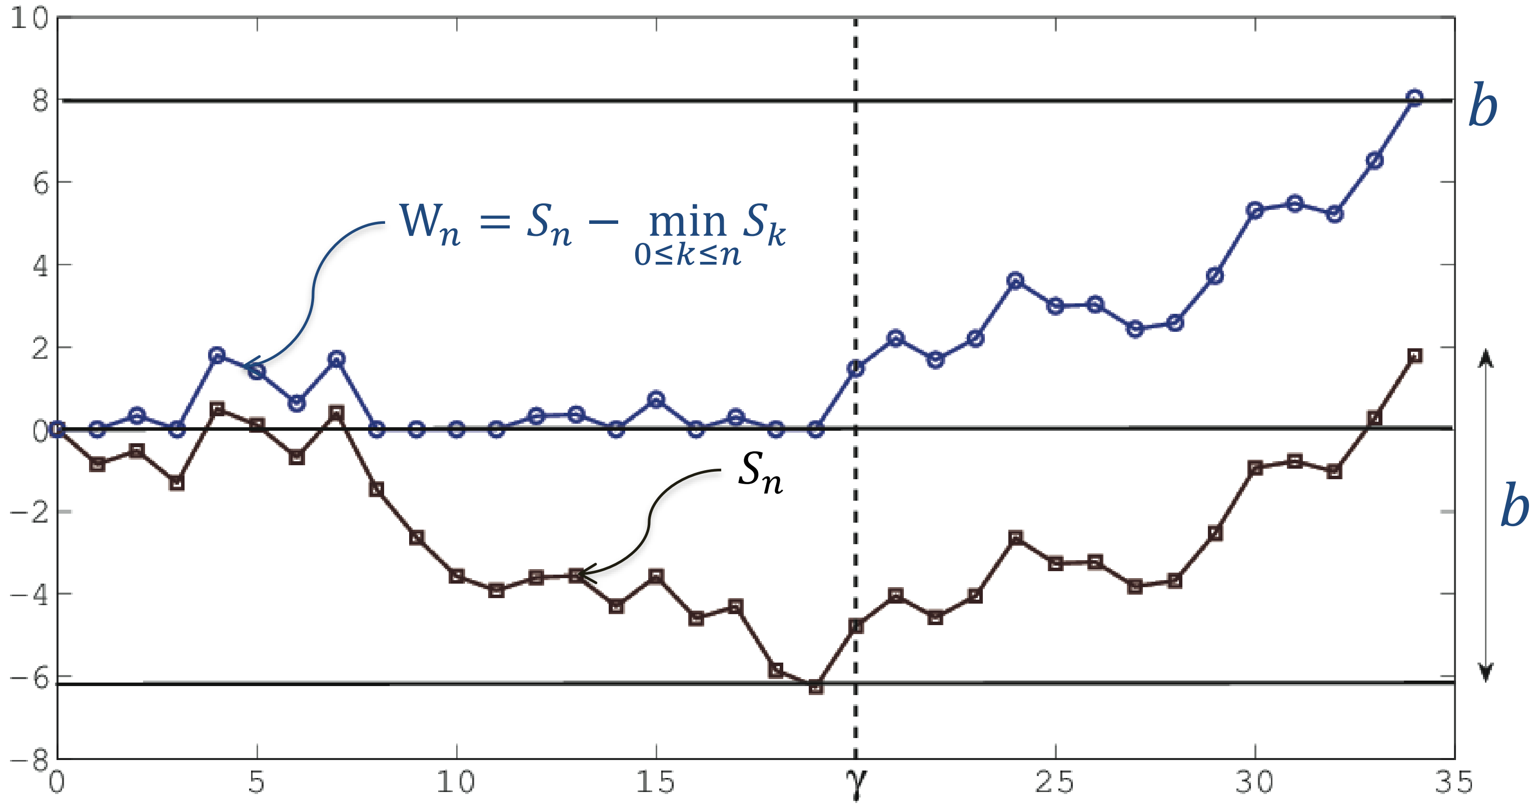

CuSum Test

CuSum statistic:

$$W_n = S_n - \min_{0 \le k \le n} S_k = \max_{0 \le k \le n} S_n - S_k$$

$$= \max_{0 \le k \le n} \sum_{\ell = k+1}^{n} \ln L(X_\ell)$$

Recursion:

$$W_n = (W_{n-1} + \ln L(X_n))^+, \qquad W_0 = 0$$

CuSum Test:

$$\tau_C = \inf\{n \ge 1 : W_n \ge b\}$$

CuSum Test – GLR Interpretation

- At time n:
  - Likelihood of no change before time n ($\mathcal{H}_0^{(n)}$): $\prod_{\ell=1}^{n} p_0(X_\ell)$
  - Likelihood of change at $k \le n$: $\prod_{\ell=1}^{k-1} p_0(X_\ell) \prod_{\ell=k}^{n} p_1(X_\ell)$
  - Generalized likelihood of change at some time before time n

    ($\mathcal{H}_1^{(n)}$): $\max_{1 \le k \le n} \prod_{\ell=1}^{k-1} p_0(X_\ell) \prod_{\ell=k}^{n} p_1(X_\ell)$

- Generalized log-likelihood ratio:

$$
\begin{aligned}
C_n &= \ln \frac{\max_{1 \le k \le n} \prod_{\ell=1}^{k-1} p_0(X_\ell) \prod_{\ell=k}^{n} p_1(X_\ell)}{\prod_{\ell=1}^{n} p_0(X_\ell)} \\
&= \max_{1 \le k \le n} \sum_{\ell=k}^{n} \ln L(X_\ell)
\end{aligned}
$$

- Recursion:

$$
C_n = (C_{n-1})^+ + \ln L(X_n)\,, \qquad C_0 = 0
$$

CuSum Test - C_n versus W_n

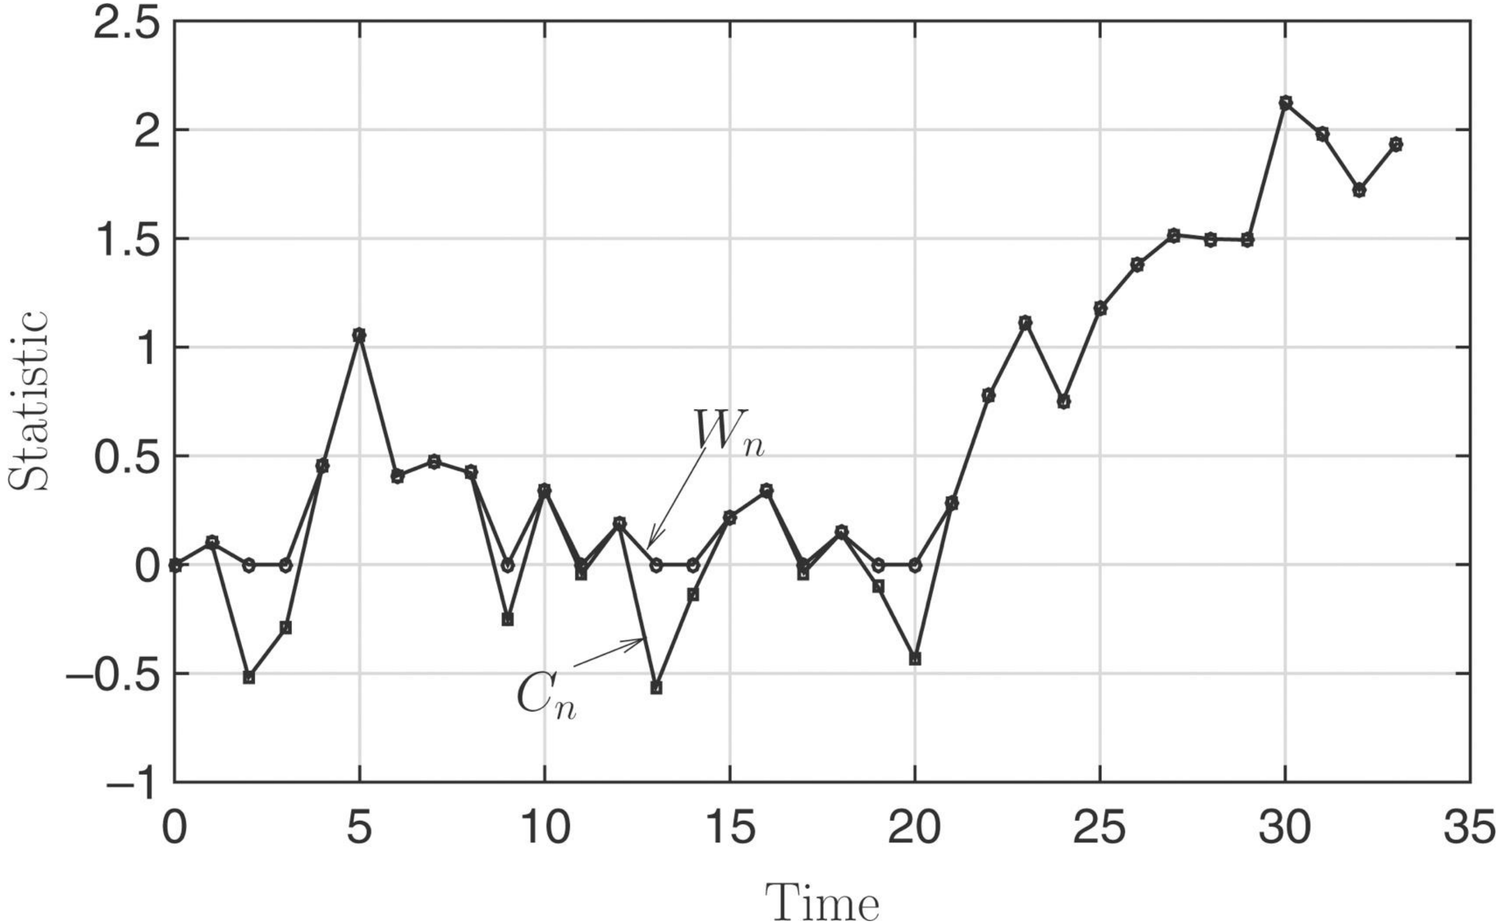

$$\tau_C = \inf\{n \geq 1: W_n \geq b\} = \inf\{n \geq 1: C_n \geq b\}$$

Shiryaev-Roberts (SR) Test

Generalized likelihood ratio:

$$L_G(X_1^n) = \frac{\max_{1 \le k \le n} \prod_{\ell=1}^{k-1} p_0(X_\ell) \prod_{\ell=k}^{n} p_1(X_\ell)}{\prod_{\ell=1}^{n} p_0(X_\ell)}$$

$$= \max_{1 \le k \le n} \prod_{\ell=k}^{n} L(X_\ell)$$

Replace max by sum:

$$R_n = \sum_{k=1}^{n} \prod_{\ell=k}^{n} L(X_\ell)$$

Recursion:

$$R_n = (1 + R_{n-1}) L(X_n), \qquad R_0 = 0$$

SR Test:

$$\tau_{\mathrm{SR}} = \inf\{n \ge 1 : R_n \ge e^b\}$$

Martingale Property of SR Statistic

- Observation sequence: $X_1, X_2, X_3, \ldots$
- Filtration: $\mathcal{F}_n = \sigma(X_1, X_2, \ldots, X_n)$
- Measure and Expectation with change at γ: $\mathbb{P}_\gamma,\ \mathbb{E}_\gamma$
- Pre-Change Measure and Expectation: $\mathbb{P}_\infty,\ \mathbb{E}_\infty$

- Result: $R_n - n$ is martingale w.r.t. filtration of observations $\{\mathcal{F}_n\}$ under $\mathbb{P}_\infty$

  Proof:

$$\begin{aligned}\mathbb{E}_\infty[R_{n+1} - (n+1)|\mathcal{F}_n] &= (R_n + 1)\mathbb{E}_\infty[L(X_{n+1})] - (n+1)\\ &= (R_n + 1)(1) - (n+1) = R_n - n\end{aligned}$$

QCD Minimax Problem Formulations

$$\underbrace{X_1, \ldots, X_{\gamma-1}}_{\text{dist } \boldsymbol{P}_0} \longrightarrow \underbrace{X_\gamma, X_{\gamma+1} \ldots}_{\text{dist } \boldsymbol{P}_1}$$

False Alarm Rate: $\mathrm{FAR}(\tau) = \dfrac{1}{\mathbb{E}_\infty[\tau]}$

Worst-Case Average Detection Delay (Lorden):

$$\mathrm{WADD}\ (\tau) = \sup_{\gamma \geq 1} \operatorname{ess\,sup} \ \mathbb{E}_\gamma\left[(\tau - \gamma + 1)^+ \mid X_1, \cdots X_{\gamma-1}\right]$$

Conditional Average Detection Delay (Pollak):

$$\mathrm{CADD}\ (\tau) = \sup_{\gamma \geq 1} \mathbb{E}_\gamma[(\tau - \gamma + 1)^+ \mid \tau \geq \gamma]$$

Easy to see that $\mathrm{CADD}\ (\tau) \leq \mathrm{WADD}\ (\tau)$

Goal: choose τ to minimize $\mathrm{WADD}\ (\tau)$ (or $\mathrm{CADD}\ (\tau)$) subject to constraint $\mathrm{FAR}(\tau) \leq \alpha$

QCD Minimax Problem Formulations

$$\underbrace{X_1, \ldots, X_{\gamma-1}}_{\text{i.i.d. } p_0} \longrightarrow \underbrace{X_\gamma, X_{\gamma+1} \ldots}_{\text{i.i.d. } p_1}$$

False Alarm Rate: $\mathrm{FAR}(\tau) = \frac{1}{\mathbb{E}_\infty[\tau]}$

Worst-Case Average Detection Delay (Lorden):

$$\mathrm{WADD}\,(\tau) = \sup_{\gamma \geq 1} \operatorname{ess\,sup}\ \mathbb{E}_\gamma\left[(\tau - \gamma + 1)^+ \mid X_1, \cdots X_{\gamma-1}\right]$$

Conditional Average Detection Delay (Pollak):

$$\mathrm{CADD}\,(\tau) = \sup_{\gamma \geq 1} \mathbb{E}_\gamma[(\tau - \gamma + 1)^+ \mid \tau \geq \gamma]$$

Easy to see that $\mathrm{CADD}\,(\tau) \leq \mathrm{WADD}\,(\tau)$

Goal: choose τ to minimize $\mathrm{WADD}\,(\tau)$ (or $\mathrm{CADD}\,(\tau)$) subject to constraint $\mathrm{FAR}(\tau) \leq \alpha$

CuSum and SR Tests

- CuSum:

$$\tau_{\mathrm{C}} = \inf\{n \geq 1 : C_n \geq b\}$$

$$C_n = \max\{C_{n-1}, 0\} + \ln L(X_n), \qquad C_0 = 0$$

$$e^{C_n} = \max\{e^{C_{n-1}}, 1\}\, L(X_n), \qquad e^{C_0} = 1$$

- SR:

$$\tau_{\mathrm{SR}} = \inf\{n \geq 1 : R_n \geq e^b\}$$

$$R_n = (1 + R_{n-1})L(X_n), \qquad R_0 = 0$$

- Comparison:

$$e^{C_n} \leq R_n, \qquad \text{for all } n \geq 1$$

$$\Rightarrow \mathbb{E}_\infty[\tau_{\mathrm{C}}] \geq \mathbb{E}_\infty[\tau_{\mathrm{SR}}]$$

$$\mathbb{E}_1[\tau_{\mathrm{C}}] \geq \mathbb{E}_1[\tau_{\mathrm{SR}}]$$

FAR for CuSum and SR Tests

Optional Sampling Theorem: Expected value of martingale at stopping time is equal to its initial expected value

Recall: $R_n - n$ is martingale w.r.t. filtration of observations $\{\mathcal{F}_n\}$ under $\mathbb{P}_\infty$

Therefore

$$\mathbb{E}_\infty\left[R_{\tau_{\mathrm{SR}}} - \tau_{\mathrm{SR}}\right] = \mathbb{E}_\infty[R_0 - 0] = 0 \Rightarrow \mathbb{E}_\infty[\tau_{\mathrm{SR}}] = \mathbb{E}_\infty\left[R_{\tau_{\mathrm{SR}}}\right] \geq e^b$$

$$\Rightarrow \mathbb{E}_\infty\ [\tau_{\mathrm{C}}] \geq\ \mathbb{E}_\infty[\tau_{\mathrm{SR}}] \geq e^b$$

$$\Rightarrow \mathrm{FAR}(\tau_{\mathrm{C}}) \leq \mathrm{FAR}(\tau_{\mathrm{SR}}) \leq e^{-b}$$

Choosing $b = |\ln \alpha| \implies \mathrm{FAR}(\tau_{\mathrm{C}}) \leq \mathrm{FAR}(\tau_{\mathrm{SR}}) \leq \alpha$

QCD Minimax Problem Formulations

$$\underbrace{X_1, \ldots, X_{\gamma-1}}_{\text{i.i.d. } p_0} \longrightarrow \underbrace{X_\gamma, X_{\gamma+1} \ldots}_{\text{i.i.d. } p_1}$$

False Alarm Rate: $\mathrm{FAR}(\tau) = \frac{1}{\mathbb{E}_\infty[\tau]}$

Worst-Case Average Detection Delay (Lorden):

$$\mathrm{WADD}\,(\tau) = \sup_{\gamma \geq 1} \operatorname{ess\,sup}\, \mathbb{E}_\gamma\left[(\tau - \gamma + 1)^+ \mid X_1, \cdots X_{\gamma-1}\right]$$

Conditional Average Detection Delay (Pollak):

$$\mathrm{CADD}\,(\tau) = \sup_{\gamma \geq 1} \mathbb{E}_\gamma[(\tau - \gamma + 1)^+ \mid \tau \geq \gamma]$$

Easy to see that $\mathrm{CADD}\,(\tau) \leq \mathrm{WADD}\,(\tau)$

Goal: choose τ to minimize $\mathrm{WADD}\,(\tau)$ (or $\mathrm{CADD}\,(\tau)$) subject to constraint $\mathrm{FAR}(\tau) \leq \alpha$

CuSum Test [Page 1954]

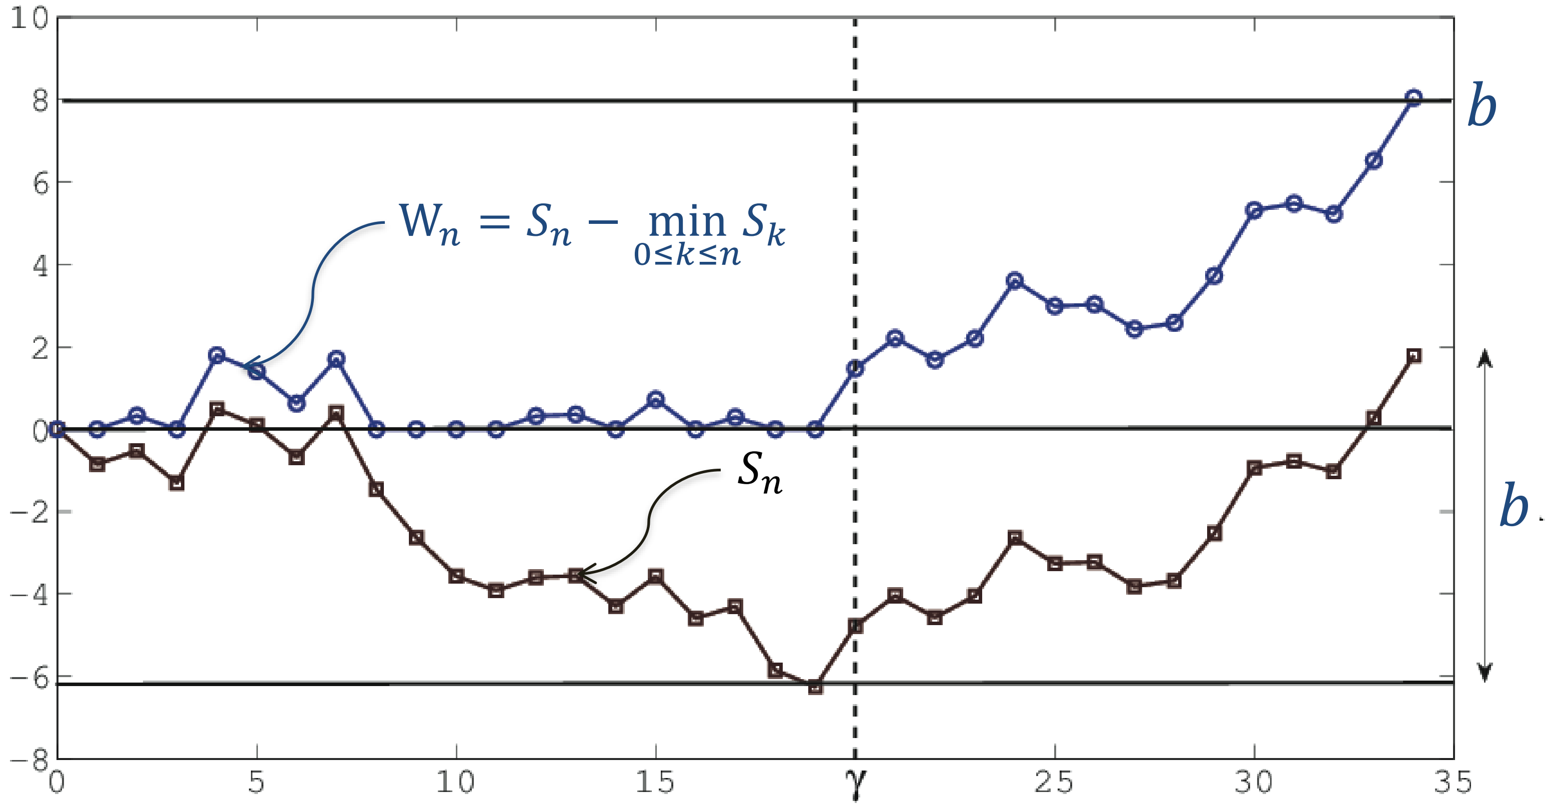

Delay for CuSum and SR Tests

Result: Both CuSum and SR tests attain worst-case delay when change occurs from start ($\gamma = 1$), i.e.,

$$\text{WADD}\,(\tau_{\text{C}}) = \text{CADD}\,(\tau_{\text{C}}) = \mathbb{E}_1[\tau_{\text{C}}]$$
$$\text{WADD}\,(\tau_{\text{SR}}) = \text{CADD}\,(\tau_{\text{SR}}) = \mathbb{E}_1[\tau_{\text{SR}}]$$

Recall: $\mathbb{E}_1[\tau_{\text{SR}}] \leq \mathbb{E}_1[\tau_{\text{C}}]$

Also,

$$\tau_{\text{C}} = \inf\{n \geq 1 : W_n \geq b\} \leq \inf\{n \geq 1 : S_n \geq b\} =: \tau^*$$
$$\Longrightarrow \mathbb{E}_1[\tau_{\text{C}}] \leq \mathbb{E}_1[\tau^*]$$

$\mathbb{E}_1[\tau^*]$ is mean time it takes for random walk $S_n = \sum_{k=1}^{n} \ln L(X_k)$ with positive drift $\mathrm{E}_{p_1}[\ln L(X_k)] = D(p_1||p_0)$ to cross threshold b

Result:

$$\mathbb{E}_1[\tau^*] = \frac{b}{D(p_1||p_0)}\big(1 + o(1)\big) \text{ as } b \to \infty$$

Asymptotic Optimality

Asymptotic Upper Bounds: For $\tau = \tau_{\mathrm{C}}$ and $\tau = \tau_{\mathrm{SR}}$, as $b \to \infty$

$$\mathrm{WADD}\,(\tau) = \mathrm{CADD}\,(\tau) \leq \frac{b}{D(p_1||p_0)}\big(1 + o(1)\big)$$

Universal Lower Bound: [Lorden 1971, Lai 1988] Consider any stopping time that satisfies $\mathrm{FAR}(\tau) \leq \alpha$, then

$$\mathrm{WADD}\,(\tau) \geq \mathrm{CADD}\,(\tau) \geq \frac{|\ln \alpha|}{D(p_1||p_0)}\big(1 + o(1)\big)$$

Asymptotic Optimality: Recall that for τ_{C} and τ_{SR}, setting $b = |\ln \alpha|$ implies that $\mathrm{FAR}(\tau_{\mathrm{C}}) \leq \mathrm{FAR}(\tau_{\mathrm{SR}}) \leq \alpha$. Therefore, τ_{C} and τ_{SR} are both asymptotically optimal as $b \to \infty$ with

$$\mathrm{WADD}\,(\tau) \sim \mathrm{CADD}\,(\tau) \sim \frac{|\ln \alpha|}{D(p_1||p_0)}$$

CuSum Performance Tradeoff

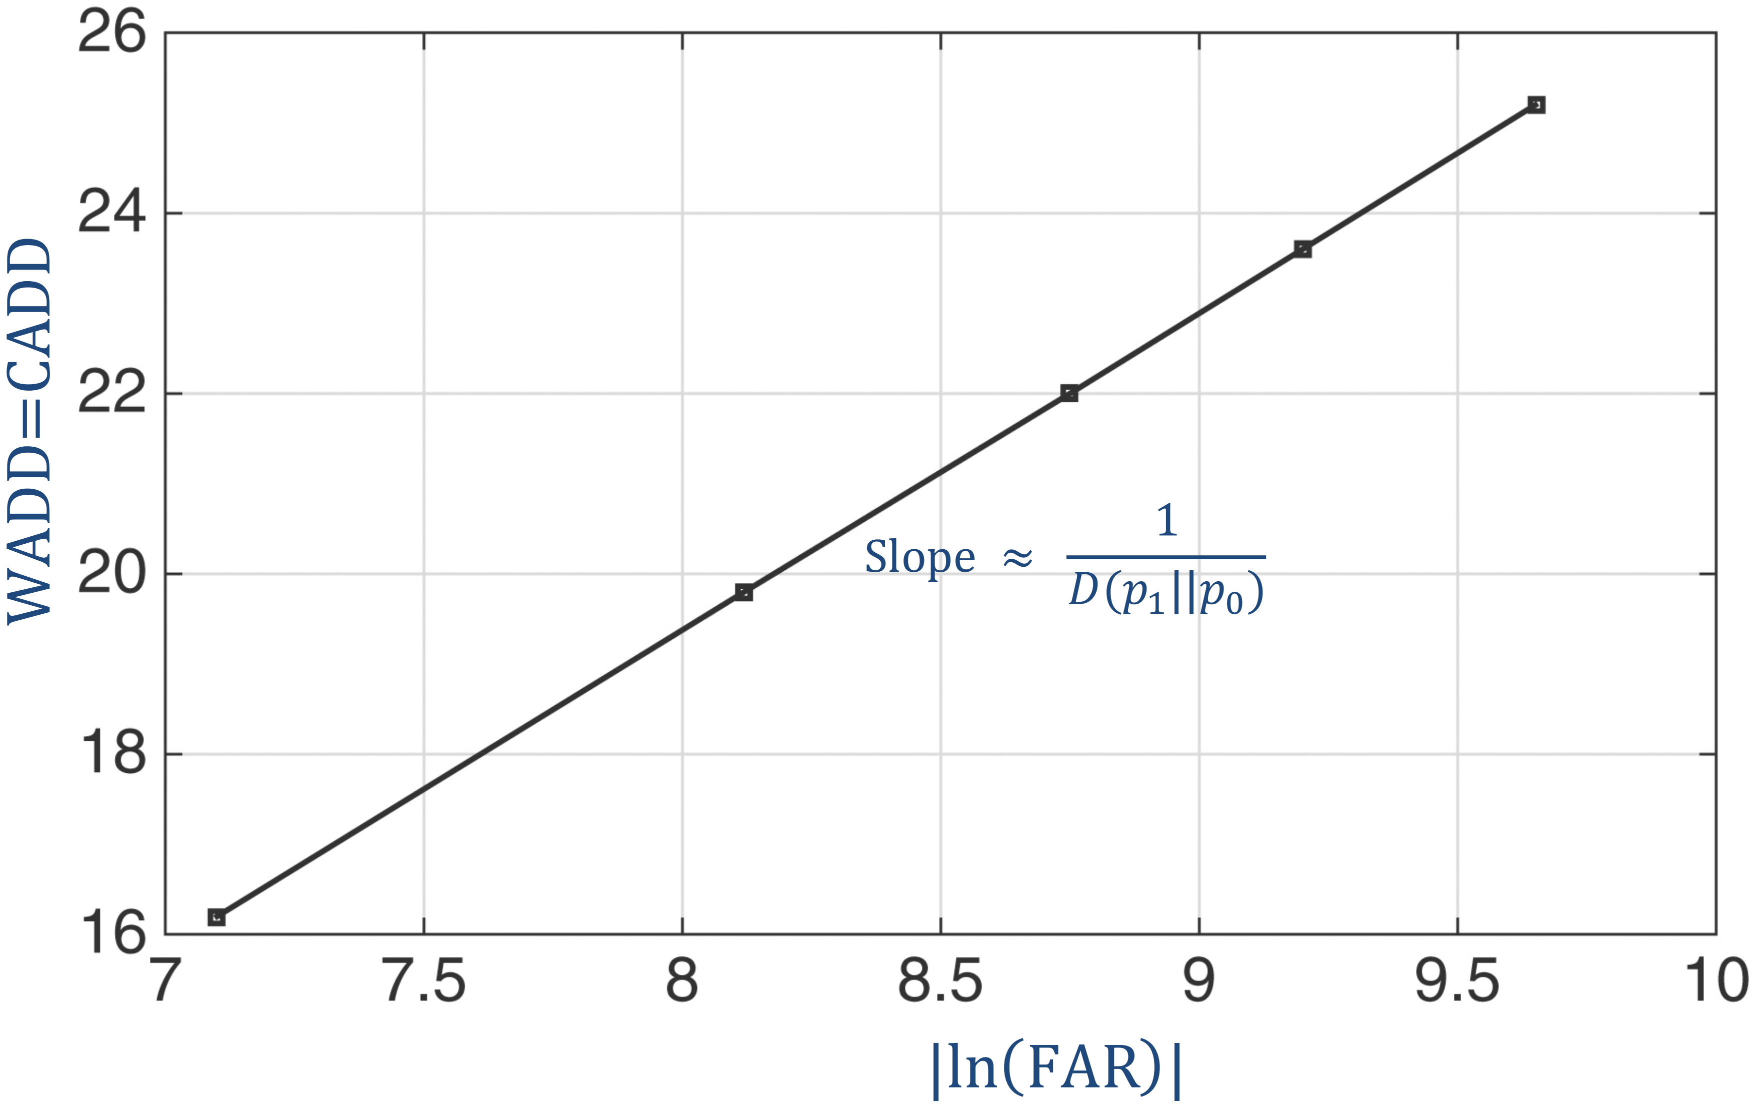

Bayesian QCD

$$\underbrace{X_1, \dots, X_{\Gamma-1}}_{\text{i.i.d. } p_0} \longrightarrow \underbrace{X_\Gamma, X_{\Gamma+1} \dots}_{\text{i.i.d. } p_1}$$

Random Change-Point: Γ with prior $\pi_\gamma = \mathbb{P}\{\Gamma = \gamma\}$

Probability of False Alarm :

$$\mathrm{PFA}(\tau) = \mathbb{P}\{\tau < \Gamma\} = \sum_{\gamma=1}^{\infty} \mathbb{P}\{\tau < \gamma\}\, \pi_\gamma$$

Average Detection Delay:

$$\mathrm{ADD}(\tau) = \mathbb{E}[(\tau - \Gamma + 1)^+] = \sum_{\gamma=1}^{\infty} \mathbb{E}[(\tau - \Gamma + 1)^+]\, \pi_\gamma$$

Goal: choose τ to minimize $\mathrm{ADD}\,(\tau)$ subject to $\mathrm{PFA}(\tau) \leq \alpha$

Shiryaev Test

Geometric Prior: $\pi_\gamma = \mathbb{P}\{\Gamma = \gamma\} = \rho(1-\rho)^{\gamma-1}, \quad \gamma \geq 1$

A posteriori Probability of Change :

$$G_n = \mathbb{P}\{\Gamma \leq n | X_1, X_2, \dots, X_n\}$$

Recursion:

$$G_n = \frac{\tilde{G}_{n-1} L(X_n)}{\tilde{G}_{n-1} L(X_n) + (1 - \tilde{G}_{n-1})}$$

$$\tilde{G}_{n-1} = G_n + \rho\,(1 - G_n), \qquad G_0 = 0$$

Result [Shiryaev 1961]: Optimal solution to Bayes QCD problem with geometric prior is

$$\tau_S = \inf\{n \geq 1 \colon G_n \geq a\}$$

if a is chosen so that $\mathrm{PFA}(\tau_S) = \alpha$

Connection to SR Test

$$G_n = \mathbb{P}\{\Gamma \leq n | X_1, X_2, \dots, X_n\}$$

Define:

$$R_{n,\rho} = \frac{G_n}{\rho\,(1 - G_n)}$$

$$= \frac{1}{(1-\rho)^n} \sum_{k=1}^{n} (1-\rho)^{k-1} \prod_{\ell=k}^{n} L(X_\ell)$$

Recursion:

$$R_{n,\rho} = \frac{1 + R_{n-1,\rho}}{(1-\rho)}\, L(X_n), \qquad R_{0,\rho} = 0$$

Taking limit as $\rho \to 0$

$$R_{n,0} = \left(1 + R_{n-1,0}\right) L(X_n), \qquad R_{0,0} = 0$$

Shiryaev-Roberts (SR) Test

Generalized likelihood ratio:

$$L_G(X_1^n) = \frac{\max_{1\le k\le n} \prod_{\ell=1}^{k-1} p_0(X_\ell) \prod_{\ell=k}^{n} p_1(X_\ell)}{\prod_{\ell=1}^{n} p_0(X_\ell)}$$

$$= \max_{1\le k\le n} \prod_{\ell=k}^{n} L(X_\ell)$$

Replace max by sum:

$$R_n = \sum_{k=1}^{n} \prod_{\ell=k}^{n} L(X_\ell)$$

Recursion:

$$R_n = (1 + R_{n-1}) L(X_n), \qquad R_0 = 0$$

SR Test:

$$\tau_{\mathrm{SR}} = \inf\{n \ge 1 : R_n \ge e^b\}$$

Performance Analysis of τ_S

$$\tau_S = \inf\{n \geq 1 : G_n \geq a\} \qquad\qquad G_n = \mathbb{P}\{\Gamma \leq n | X_1, \dots, X_n\}$$

False Alarm Control:

$$\begin{aligned}
\mathrm{PFA}(\tau_S) &= \mathbb{P}\{\tau_S < \Gamma\} \\
&= \mathbb{E}\left[\mathbb{P}\left\{\tau_S < \Gamma \,| X_1, \dots, X_{\tau_S}\right\}\right] \\
&= 1 - \mathbb{E}\left[G_{\tau_S}\right] \\
&\leq 1 - a
\end{aligned}$$

thus, setting $a = 1 - \alpha$ controls $\mathrm{PFA}(\tau_S) \leq \alpha$

Delay Analysis: with $a = 1 - \alpha,$ as $\alpha \to 0$

$$\mathrm{ADD}\,(\tau_S) \sim \frac{|\ln \alpha|}{D(p_1||p_0) + |\ln(1 - \rho)|}$$

[Tartakovsky&VVV, 2005]

Variants of QCD Problem

- Extensions to allow for parametric distributional uncertainty
- Extensions to non-i.i.d. pre- and post-change
- Non-parametric settings
- Transient changes and persistent changes with transient dynamics
- Multichannel/multi-stream settings
- Distributed and decentralized (multi-sensor) settings
- Data-efficient QCD
- …

CuSum Test – i.i.d. setting

$$\underbrace{X_1, \dots, X_{\gamma-1}}_{\text{i.i.d. } p_0} \longrightarrow \underbrace{X_\gamma, X_{\gamma+1} \dots}_{\text{i.i.d. } p_1}$$

- CuSum statistic:

$$C_n = \max_{1 \le k \le n} \sum_{\ell=k}^{n} \ln \frac{p_1(X_\ell)}{p_0(X_\ell)}$$

- Recursion:

$$C_n = (C_{n-1})^+ + \ln \frac{p_1(X_n)}{p_0(X_n)}, \qquad C_0 = 0$$

- CuSum Test :

$$\tau_C = \inf\{n \ge 1 : C_n \ge b\}$$

GLR-CuSum Test – i.i.d. setting

$$\underbrace{X_1, \ldots, X_{\nu-1}}_{\text{i.i.d. } p_{\theta_0}} \longrightarrow \underbrace{X_\nu, X_{\nu+1} \ldots}_{\text{i.i.d. } p_\theta} \qquad \theta \in \Theta, \theta_0 \notin \Theta$$

- GLR-CuSum statistic:

$$C_n^{\mathrm{G}} = \max_{1 \le k \le n} \sup_{\theta \in \Theta} \sum_{\ell=k}^{n} \ln \frac{p_\theta(X_\ell)}{p_{\theta_0}(X_\ell)}$$

- GLR-CuSum Test:

$$\tau_{\mathrm{G}} = \inf\left\{n \ge 1 : C_n^{\mathrm{G}} \ge b\right\}$$

- Asymptotically optimum, but no recursion (exception: $|\Theta| < \infty$)!

- Window-limited GLR-CuSum Statistic:

$$C_n^{\mathrm{G}} = \max_{n-m(\alpha) \le j \le n} \sup_{\theta \in \Theta} \sum_{\ell=k}^{n} \ln \frac{p_1^\theta(X_\ell)}{p_0(X_\ell)}$$

Reinforcement Learning

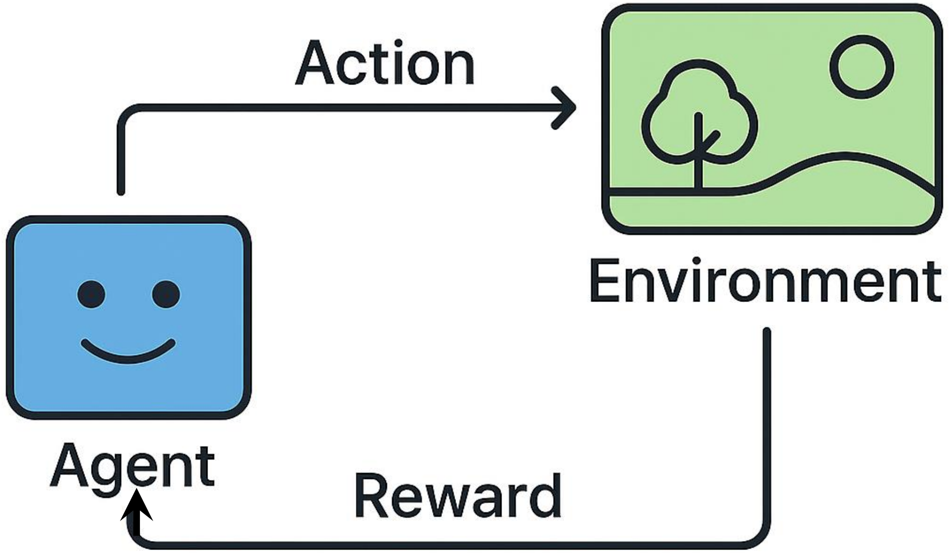

- RL Agent learns to make decisions by trial and error through repeated interaction with an environment
  - maximize some notion of cumulative reward
  - equivalently minimize regret with respect to optimal decision maker
- Most existing RL algorithms assume that environment is stationary – resulting policies may no longer be useful if environment changes

RL in Non-Stationary Environments

- Non-stationary (NS) environments of two types
  - piecewise stationary (PS)
  - drifting NS

- Approaches to dealing with NS environments
  - adaptive – run continuously without restarting
  - restarting based -- restart learning process in response to changes

- Existing adaptive methods are prior-based, i.e., they need to have prior knowledge about non-stationarity in order to ensure adaptation. Also, computational complexity can be prohibitive [Peng & Papadimitriou, COLT 2024]

- Restarting methods can be prior-based or prior-free
  - prior-based require knowledge of total non-stationarity budget and choose restarting rate based on budget
  - prior-free methods are more practical, but require careful design of non-stationarity/change detectors to trigger restarts

Detection-Augmented Learning (DAL)

Stationary learning ($\mathscr{L}$) …*used as is*

+

Change detector ($\mathscr{D}$) …*with guaranteed probabilities of false alarm and delayed detection*

+

Forced exploration …*required to bypass focusing on optimal actions*

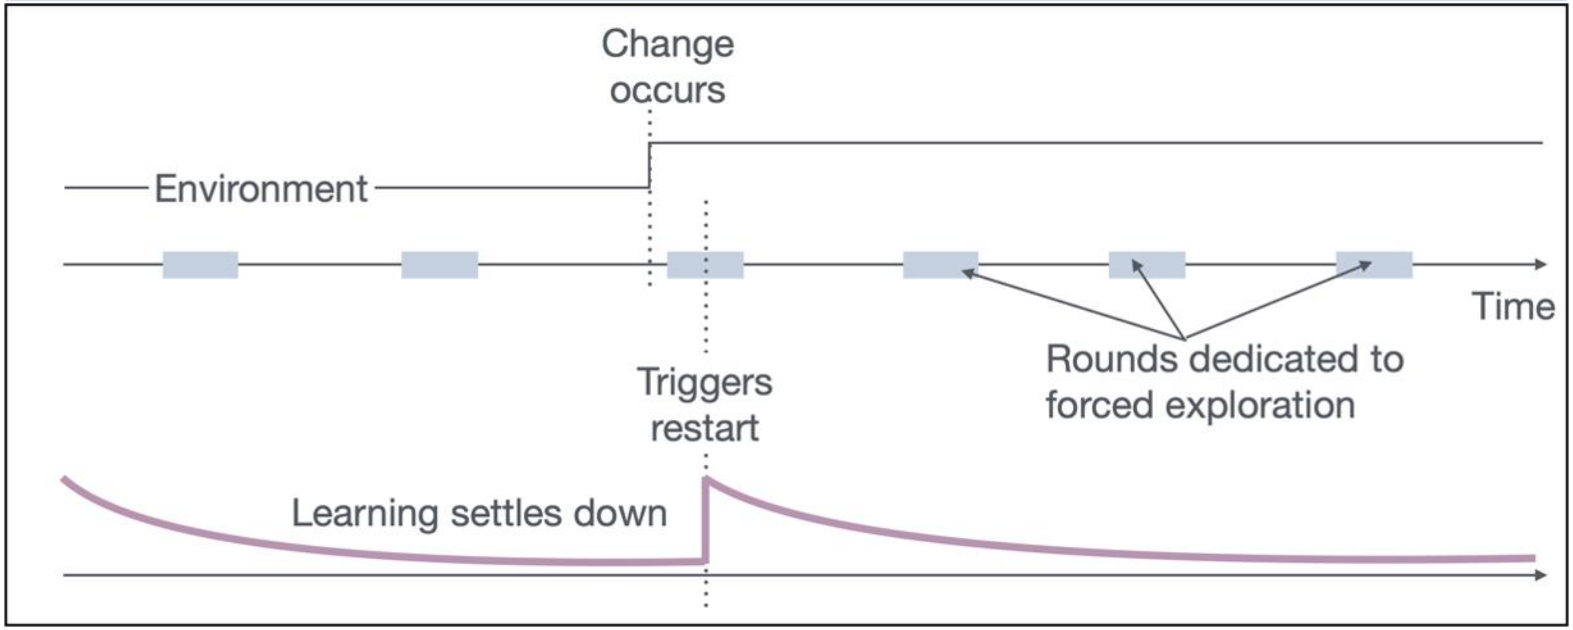

QCD for Learning in PS Environments

- Standard minimax QCD problem formulation: minimize (worst-case) expected detection delay, subject to constraint on expected time to false alarm, over possibly infinite horizon
- Learning in piecewise stationary (PS) environments: underlying distribution of environment undergoes changes at multiple change-points, and remains stationary in between
- Effective way to learn in PS environments: detect changes in underlying distribution, and restart learning after change detection [Auer et al, 2019], [Besson et al, 2022], [Huang et al, 2026], …
- To achieve good (order-optimal) performance, these algorithms need to:
  - control probability of false alarm
  - control latency, which is smallest value such that probability of detection delay exceeding this value is upper bounded by predetermined level
  - operate over finite horizon T

QCD Problem Formulation

$$\underbrace{X_1, \dots, X_{\gamma-1}}_{\text{i.i.d. } p_0} \longrightarrow \underbrace{X_\gamma, X_{\gamma+1}, \dots, X_T}_{\text{i.i.d. } p_1}$$

- p_0, μ_0: pre-change density/mean, p_1, μ_1: post-change density/mean
- in PS learning, p_0 and p_1 are unknown, but typically assumed to be sub-Gaussian or sub-Bernoulli and $\Delta = |\mu_1 - \mu_0| > 0$ (Δ unknown)
- when p_0 unknown, assume change occurs after time m, i.e., $\gamma > m$
- change is declared at stopping time τ

- Latency (at level $\delta_\mathrm{D} \in (0,1)$):

$$\ell = \min\left\{ d \in [T] : \mathbb{P}_\gamma(\tau \geq \gamma + d) \leq \delta_\mathrm{D}, \forall \gamma \in \{m+1, \dots, T-d\}\right\}$$

- Optimization problem:

$$\underset{\tau}{\text{minimize}} \quad \ell \qquad \text{s.t.} \qquad \mathbb{P}_\infty(\tau \leq T) \leq \delta_\mathrm{F}$$

- Horizon T may be unknown to change detector

Universal Lower Bound on Latency

- Theorem 1: If $\delta_{\mathrm{F}} + \delta_{\mathrm{D}} < 1$,

$$\ell \geq \left(\frac{1}{D(p_1 || p_0)} + o(1) \right) [\log T + \log(1/\delta_{\mathrm{F}}) + \log(1 - \delta_{\mathrm{F}} - \delta_{\mathrm{D}}) + o(1)]$$

as $T \to \infty$

- Property 1: In regret analysis for piecewise stationary bandits and MDPs, good change detectors need to satisfy

$$\ell + m = \mathcal{O}(\log T + \log(1/\delta_{\mathrm{F}}) + \log(1/\delta_{\mathrm{D}}))$$

- Goal: develop change detectors that satisfy Property 1 and approach lower bound

TVT-CuSum Test

- First assume p_0 and p_1 known, with pre-change window $m = 0$
- CuSum statistic:

$$C_n = \max_{1 \le k \le n} \sum_{\ell=k}^{n} \ln \frac{p_1(X_\ell)}{p_0(X_\ell)}$$

which has recursion:

$$C_n = (C_{n-1})^+ + \ln \frac{p_1(X_n)}{p_0(X_n)}, \qquad C_0 = 0$$

- Stopping time for standard CuSum test:

$$\tau_{\mathrm{C}} = \inf\{n \ge 1 : C_n \ge b\}$$

  - do not know $T \Rightarrow$ Cannot tune threshold b according to T
  - $\mathbb{E}_\infty[\tau] < \infty$ for all constant $b > 0 \Rightarrow \mathbb{P}_\infty(\tau \le T) \to 1$ as $T \to \infty$
  - $\Rightarrow b$ should increase with time to ensure $\mathbb{P}_\infty(\tau \le T) \le \delta_{\mathrm{F}} \in (0,1)$

TVT-CuSum Test

- Time varying threshold (TVT)-CuSum test:

$$\tau := \inf\{n \in \mathbb{N}: C_n > \beta(n, \delta_{\mathrm{F}})\}$$

where $\beta(n, \delta_{\mathrm{F}}) := 2\log n + \log(1/\delta_{\mathrm{F}}) + \log\left(\pi^2/6\right)$

- Theorem 2: With $m = 0$, TVT-CuSum satisfies false alarm constraint $\mathbb{P}_\infty(\tau \leq T) \leq \delta_{\mathrm{F}}$. In addition, the latency satisfies

$$\ell \leq \inf_{\theta \in (0,1)} \left\{ \frac{1}{|\Lambda(\theta)|} \left[\theta \cdot \beta(T, \delta_{\mathrm{F}}) + \log(1/\delta_{\mathrm{D}})\right] \right\}$$

  - where $\Lambda(\theta)$ is cgf of $\log \frac{f_0(X)}{f_1(X)}$ with $X \sim f_1$
  - $\beta(T, \delta_{\mathrm{F}}) = \mathcal{O}\left(\log T + \log(1/\delta_{\mathrm{F}})\right)$
  - $\lim_{\theta \to 0^+} \frac{|\Lambda(\theta)|}{\theta} = D(f_1 || f_0)$

TVT-CuSum Test

- TVT-CuSum latency satisfies

$$\ell \leq \inf_{\theta \in (0,1)} \left\{ \frac{1}{|\Lambda(\theta)|} [\theta \cdot \beta(T, \delta_{\mathrm{F}}) + \log(1/\delta_{\mathrm{D}})] \right\}$$

$$= \mathcal{O}\big(\log T + \log(1/\delta_{\mathrm{F}}) + \log(1/\delta_{\mathrm{D}})\big)$$

- TVT-CuSum attains $\ell = \mathcal{O}\big(\log T + \log(1/\delta_{\mathrm{F}}) + \log(1/\delta_{\mathrm{D}})\big)$ with $m = 0$ i.e., it satisfies Property 1, but needs to know p_0 and p_1

GLR test with unknown p_0 and p_1

- p_0, p_1 unknown σ^2 –sub-Gaussian $\Delta = |\mu_1 - \mu_0| > 0$ (Δ unknown)
- GLR statistic:

$$G_n \coloneqq \sup_{1 \le k \le n} \log \frac{\sup_{\mu_0', \mu_1'} \prod_{i=1}^{k} f_{\mu_0'}(X_i) \prod_{i=k+1}^{n} f_{\mu_1'}(X_i)}{\sup_{\mu} \prod_{i=1}^{n} f_{\mu}(X_i)}$$

f_μ is Gaussian pdf with mean μ and variance σ^2

- GLR test:

$$\tau_{\mathrm{GLR}} \coloneqq \inf\{n \in \mathbb{N} : G_n > \beta_{\mathrm{GLR}}(n, \delta_{\mathrm{F}})\}$$

where $\beta_{\mathrm{GLR}}(n, \delta_{\mathrm{F}}) \coloneqq 6 \log(1 + \log n) + \frac{5}{2} \log\left(\frac{3n^{3/2}}{\delta_{\mathrm{F}}}\right) + 11$

GLR test with unknown p_0 and p_1

- Theorem 3: With,

$$m > \frac{8\sigma^2}{\Delta^2}\beta_{\mathrm{GLR}}(T, \delta_{\mathrm{F}})$$

τ_{GLR} satisfies:

$$\mathbb{P}_\infty(\tau_{\mathrm{GLR}} \leq T) \leq \delta_{\mathrm{F}}$$

and

$$\ell \leq \max\left\{\frac{8m\sigma^2\beta_{\mathrm{GLR}}(T,\delta_{\mathrm{F}})}{\Delta^2 m - 8\sigma^2\beta_{\mathrm{GLR}}(T,\delta_{\mathrm{F}})}, \frac{{\delta_{\mathrm{F}}}^{2/3}}{2^{16/15}{\delta_{\mathrm{D}}}^{4/15}} - m\right\}$$

- Since $\beta_{\mathrm{GLR}}(T, \delta_{\mathrm{F}}) = \mathcal{O}\big(\log T + \log(1/\delta_{\mathrm{F}})\big)$, it is easily seen that τ_{GLR} satisfies Property 1, i.e., $\ell + m = \mathcal{O}\big(\log T + \log(1/\delta_{\mathrm{F}}) + \log(1/\delta_{\mathrm{D}})\big)$

Experiments

- Windowed GLR statistics: Since supremum imposes high computation complexity, we substitute Gaussian GLR statistic with corresponding windowed version

$$G'_n := \sup_{n-700 \le k \le n} \log \frac{\sup_{\mu'_0,\mu'_1} \prod_{i=1}^{k-1} f_{\mu'_0}(X_i) \prod_{i=k}^{n} f_{\mu'_1}(X_i)}{\sup_{\mu_0} \prod_{i=1}^{n} f_{\mu}(X_i)}$$

- $p_0 \sim \mathcal{N}(0,1), p_1 \sim \mathcal{N}(1,1), \quad \delta_\mathrm{D} = \delta_\mathrm{F} = 0.01, \quad T$ varied

- m and β chosen as function of δ_F, T

- γ varied between $m + 1$ and T

- Choose $100(1 - \delta_\mathrm{D})$-th percentile of $\tau - \gamma$ over 200000 trials as ℓ

Experiments

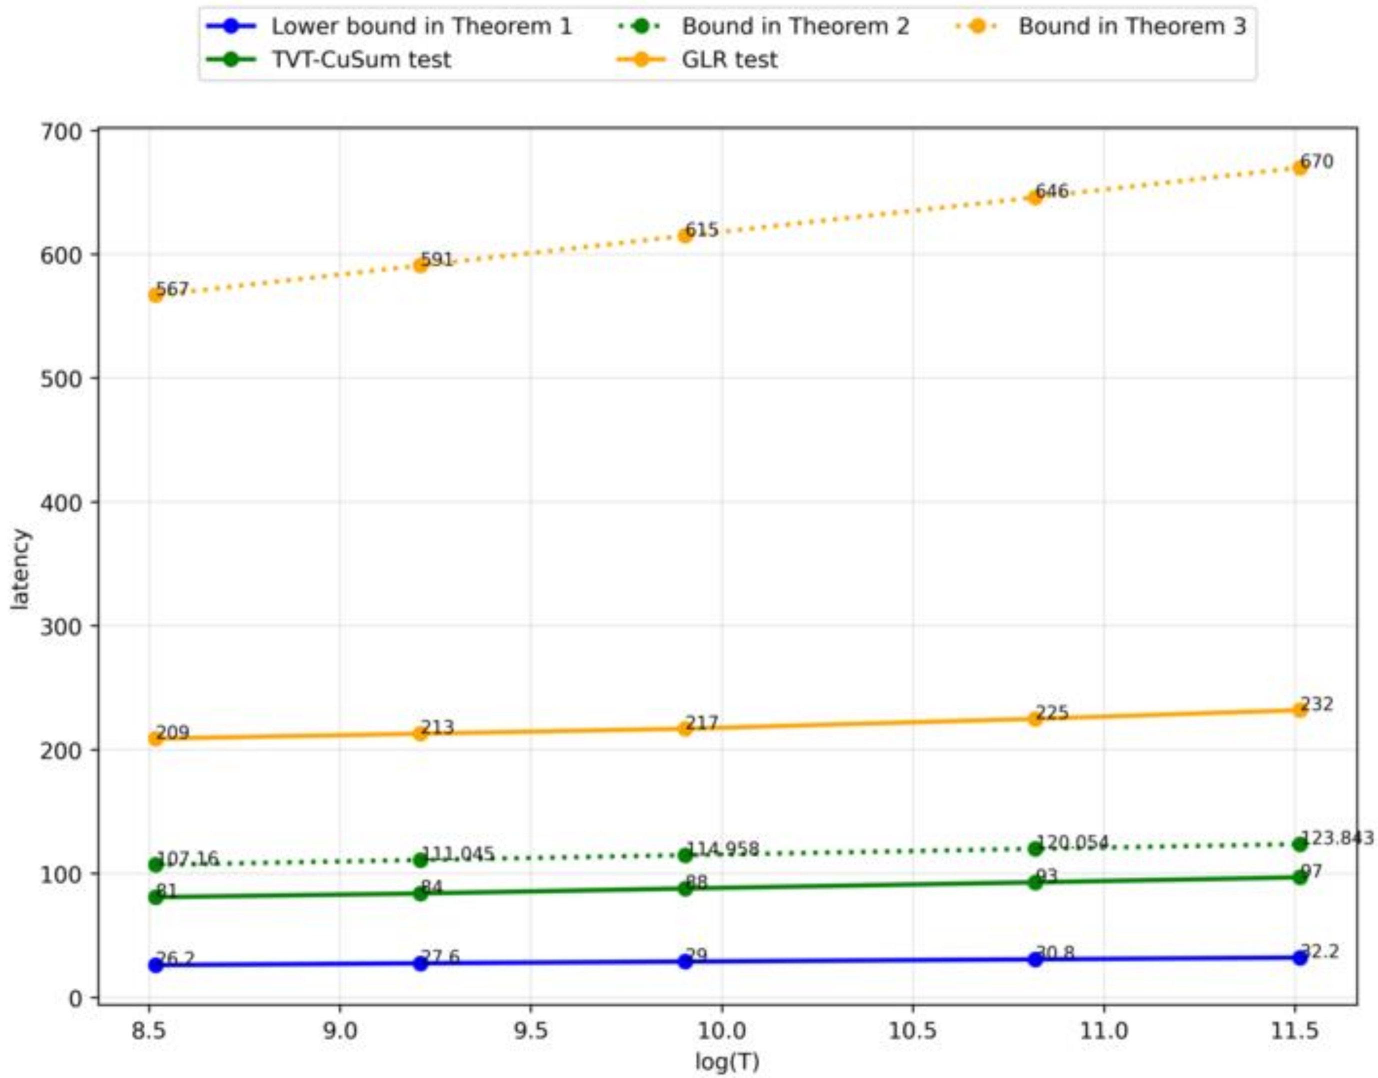

Recent Papers on Detection-Augmented Learning for PS Environments

- Y.-H. Huang and V.V. Veeravalli. "Finite-Horizon Quickest Change Detection Balancing Latency with False Alarm Probability." To appear in Sequential Analysis in 2026. Pre-print on ArXiv.

- Y-H. Huang, A. Gerogiannis, S. Bose and V.V. Veeravalli. "Detection Augmented Bandit Procedures for Piecewise Stationary MABs: A Modular Approach." To appear in IEEE Transactions on Information Theory in 2026. Pre-print on ArXiv.

- A.Gerogiannis, Y.-H. Huang, S. Bose and V.V. Veeravalli. "DAL: A Practical Prior-Free Black-Box Framework for Piecewise Stationary Bandits." In Proc. International Conference on Machine Learning (ICML), Seoul, South Korea, July 2026.

- A. Gerogiannis, Y.-H. Huang and V.V. Veeravalli. "Is Prior-Free Black-Box Non-Stationary Reinforcement Learning Feasible?" Proc. Artificial Intelligence and Statistics (AISTATS), Mai Khao, Thailand, May 2025.

- A. Gerogiannis, Y.-H. Huang and V.V. Veeravalli. "DARLING: Detection Augmented Reinforcement Learning with Non-Stationary Guarantees." arXiv preprint arXiv:2604.16684 (2026).

References

- W. A. Shewhart, “The application of statistics as an aid in maintaining quality of a manufactured product,” J. Amer. Statist. Assoc., vol. 20, pp. 546–548, Dec. 1925.
- E. S. Page, “Continuous inspection schemes,” Biometrika, vol. 41, pp. 100–115, June 1954.
- A. N. Shiryaev, “On optimum methods in quickest detection problems,” Theory of Prob and App., vol. 8, pp. 22–46, 1963.
- G. Lorden, “Procedures for reacting to a change in distribution,” Ann. Math. Statist., vol. 42, pp. 1897–1908, Dec. 1971.
- M. Pollak, “Optimal detection of a change in distribution,” Ann. Statist., vol. 13, pp. 206–227, Mar. 1985.
- G. V. Moustakides, “Optimal stopping times for detecting changes in distributions,” Ann. Statist., vol. 14, pp. 1379–1387, Dec. 1986.
- T. L. Lai, “Information bounds and quick detection of parameter changes in stochastic systems,” IEEE Trans. Inf. Theory, vol. 44, pp. 2917 –2929, Nov. 1998.
- A. G. Tartakovsky and V. V. Veeravalli, “General asymptotic Bayesian theory of quickest change detection,” SIAM Theory of Prob. and App., vol. 49, pp. 458–497, Sept. 2005.
- B. Peng and C. Papadimitriou. The complexity of non-stationary reinforcement learning. In Proc. Int. Conference on Algorithmic Learning Theory (COLT), PMLR, pp. 972–996, Feb 2024.
- D. Siegmund, Sequential Analysis: Tests and Confidence Intervals. Springer series in statistics, Springer-Verlag, 1985.
- H. V. Poor and O. Hadjiliadis, Quickest detection. Cambridge University Press, 2009.
- A. G. Tartakovsky, I. V. Nikiforov, and M. Basseville, Sequential Analysis: Hypothesis Testing and Change-Point Detection. Statistics, CRC Press, 2013.
- V.V. Veeravalli and T. Banerjee. “Quickest Change Detection.” In E-Reference Signal Processing. Elsevier, 2013. Also available on ArXiv.
- L. Xie, S. Zou, Y. Xie, and V.V. Veeravalli. “Sequential (Quickest) Change Detection: Classical Results and New Directions.” IEEE Journal on Selected Areas in Information Theory, Special Issue on Sequential Active and Reinforcement Learning, 2(2): 494-514, June 2021.