

QUICKEST CHANGE DETECTION WITH LEAVE-ONE-OUT DENSITY ESTIMATION

Yuchen Liang and Venugopal V. Veeravalli

ECE Department and Coordinated Science Lab
University of Illinois at Urbana-Champaign
Email: yliang35, vvv@illinois.edu

ABSTRACT

The problem of quickest change detection in a sequence of independent observations is considered. The pre-change distribution is assumed to be known, while the post-change distribution is completely unknown. A window-limited leave-one-out (LOO) CuSum test is developed, which does not assume any knowledge of the post-change distribution, and does not require any post-change training samples. It is shown that, with certain convergence conditions on the density estimator, the LOO-CuSum test is first-order asymptotically optimal, as the false alarm rate goes to zero. The analysis is validated through numerical results, where the LOO-CuSum test is compared with baseline tests that have distributional knowledge.

Index Terms— Quickest change detection (QCD), non-parametric statistics, (kernel) density estimation.

1. INTRODUCTION

The problem of quickest change detection (QCD) is of fundamental importance in mathematical statistics (see, e.g., [1, 2] for an overview). Given a sequence of observations whose distribution changes at some unknown change-point, the goal is to detect the change in distribution as quickly as possible after it occurs, while controlling the false alarm rate. In classical formulations of the QCD problem, it is assumed that the pre- and post-change distributions are known, and that the observations are independent and identically distributed (i.i.d.) in either the pre-change or the post-change regime. However, in many practical situations, while it is reasonable to assume that we can accurately estimate the pre-change distribution, the post-change distribution is rarely completely known.

There have been extensive efforts to address pre- and/or post-change distributional uncertainty in QCD problems. In the case where both distributions are not fully known, one approach is to assume that they are indexed by a (low-dimensional) parameter that comes from a pre-defined pa-

rameter set, and employ a generalized likelihood ratio (GLR) approach to detection – this was first introduced in [3] and later analyzed in more detail in [4]. In particular, in [4], it is assumed that the pre-change distribution is known and that the post-change distribution comes from a parametric family, with the parameter being finite-dimensional. A window-limited GLR test is proposed, which is shown to be asymptotically optimal under certain smoothness conditions. This work has recently been extended to non-stationary post-change settings [5, 6].

Another approach to dealing with distributional uncertainty in QCD problems is the minimax robust approach [7], where it is assumed that the pre- and post-change distributions come from (known) mutually exclusive uncertainty classes, and the goal is to optimize the performance for the worst-case choice of distributions in the uncertainty classes. Under certain conditions, e.g., joint stochastic boundedness (see, e.g., [8] for a definition) and weak stochastic boundedness [9], robust solutions can be found [10, 9]. However, these robust tests can have suboptimal performance for the actual distributions encountered in practice.

In this paper, we will assume complete knowledge of the pre-change distribution, while not making any parametric assumptions about the post-change distribution. There have also been approaches to deal with non-parametric uncertainty in the distributions in QCD problems. One approach is to replace the log-likelihood ratio by some other statistic and formulate the test in the non-parametric setting. Examples of this approach include the use of kernel M-statistics [11, 12], one-class SVMs [13], nearest neighbors [14, 15], and Geometric Entropy Minimization [16]. In [11], a test is proposed that compares the kernel maximum mean discrepancy (MMD) within a window to a given threshold. A way to set the threshold is also proposed that meets the false alarm rate asymptotically [11]. Another approach is to estimate the log-likelihood ratio and thus the CuSum test statistic through a pre-collected training set. This includes direct kernel estimation [17] and, more recently, neural network estimation [18]. However, *the tests proposed in [11]–[18] lack explicit performance guarantees on the detection delay*. The closest work to ours is [19], where a binning approach is proposed to solve the QCD problem asymptotically without any pre-collected

This work was supported in part by the National Science Foundation under grant ECCS-2033900, and by the Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196, through the University of Illinois at Urbana-Champaign.

training set. In particular, in [19], an asymptotically optimal solution is established for the case where the pre-change distribution is known, the post-change distribution is distinguishable from the pre-change with binning, and both distributions have discrete support.

Contributions:

1. We propose a window-limited leave-one-out (LOO) CuSum test, which does not assume any knowledge of the post-change distribution, and does not require any post-change training samples.
2. We provide a way to set the test threshold that asymptotically meets the false alarm constraint.
3. We show that the proposed LOO-CuSum test is first-order asymptotically optimum, as the false alarm rate goes to zero.
4. We validate our analysis through numerical results, in which we compare the LOO-CuSum test with baseline tests that have distributional knowledge.

The rest of the paper is structured as follows. In Section 2, we describe several properties required of the density estimators for asymptotically optimal detection. In Section 3, we propose the LOO-CuSum test, and analyze its theoretical performance. In Section 4, we present numerical results that validate the theoretical analysis. In Section 5, we provide some conclusions.

2. LEAVE-ONE-OUT (LOO) DENSITY ESTIMATOR

Let $X_1, X_2, \dots \in \mathbb{R}^d$ be i.i.d. samples drawn from an unknown distribution p . Denote by $\text{supp}(p)$ the support of p . Denote by E_p and V_p , the expectation and variance operator, respectively, under p . Denote the LOO estimated density as $\hat{p}_{-i}^{n,k}$, where the subscript $-i$ represents that X_i , with $k \leq i \leq n$, is the sample that is left out. Note that the density is a function of $X_{-i}^{n,k} := X_k, \dots, X_{i-1}, X_{i+1}, \dots, X_n$, and thus $\hat{p}_{-i}^{n,k}$ and X_i are independent. The estimation procedure is assumed to be sample-homogeneous, i.e., $\hat{p}_{-i}^{n,k} \stackrel{d.}{=} \hat{p}_{-j}^{n,k}, \forall k \leq i < j \leq n$. The Kullback-Leibler (KL) divergence between distributions p and q is $D(p||q) := \int_{\text{supp}(p)} \log(p(x)/q(x))dx$.

Suppose that, for large enough n , there exist constants $0 < \beta_1, C_1, C_2 < \infty$ and $1 < \beta_2 < 2$ (that depend only on the distribution p and the estimation procedure) such that, for each $1 \leq i \leq n+1$, the KL loss [20] of the leave-one-out (LOO) estimator satisfies

$$\text{KL-loss}(\hat{p}_{-i}^{n,k}) := E_p \left[D(p||\hat{p}_{-i}^{n,k}) \right] \leq \frac{C_1}{(n-k)\beta_1} \quad (1)$$

where the expectation E_p is over the randomness of $\hat{p}_{-i}^{n,k}$. Also, the total variance of the estimator satisfies

$$V_p \left(\sum_{i=k}^n \log \frac{p(X_i)}{\hat{p}_{-i}^{n,k}(X_i)} \right) \leq C_2(n-k+1)^{\beta_2}. \quad (2)$$

One typical loss measure for a density estimator is the mean-integrated squared error (MISE), defined as (see, e.g., [21, Chap. 2])

$$\begin{aligned} \text{MISE}(p, \hat{p}_{-i}^{n,k}) &= E_p \left[\int (\hat{p}_{-i}^{n,k}(x_i) - p(x_i))^2 dx_i \right] \\ &= E_p \left[\left\| \hat{p}_{-i}^{n,k} - p \right\|_2^2 \right], \quad \forall k \leq i \leq n. \end{aligned} \quad (3)$$

The following lemma connects the MISE with the bounds in (1) and (2).

Lemma 2.1. *Suppose that there exist $\bar{\zeta}, \underline{\zeta}$ such that*

$$0 < \underline{\zeta} \leq p(x), \hat{p}_{-i}^{n,k}(x) \leq \bar{\zeta} < \infty, \quad \forall x \in \text{supp}(p). \quad (4)$$

for any $k \leq i \leq n$. If the estimator achieves

$$\text{MISE}(p, \hat{p}_{-i}^{n,k}) \leq \frac{C_3}{(n-k)\beta_1} \quad (5)$$

for some constant $C_3 < \infty$, then (1) and (2) are satisfied with

$$C_1 = \frac{C_3}{\underline{\zeta}}, \quad C_2 = \frac{\bar{\zeta}^2 C_3}{\underline{\zeta}^2}, \quad \beta_2 = 2 - \beta_1$$

where $r := \left(\frac{\log(\bar{\zeta}/\underline{\zeta})}{(\bar{\zeta}/\underline{\zeta})-1} \right)^2$ is a constant.

Proof Sketch. The key is to use the fact that $\log s \leq s-1$ and that $(\log s)^2 \leq r(s-1)^2$ on $s \geq \bar{\zeta}/\underline{\zeta}$ if $r = \left(\frac{\log(\bar{\zeta}/\underline{\zeta})}{(\bar{\zeta}/\underline{\zeta})-1} \right)^2$. Thus,

$$\begin{aligned} E_p \left[\log \frac{p(X_i)}{\hat{p}_{-i}^{n,k}(X_i)} \right] &\leq E_p \left[\frac{p(X_i)}{\hat{p}_{-i}^{n,k}(X_i)} - 1 \right] \\ V_p \left(\log \frac{p(X_i)}{\hat{p}_{-i}^{n,k}(X_i)} \right) &\leq E_p \left[r \left(\frac{p(X_i)}{\hat{p}_{-i}^{n,k}(X_i)} - 1 \right)^2 \right] \end{aligned}$$

Note the definition of MISE in (3). Reordering the terms gives the desired result. $\square$

An example of a LOO estimator that satisfies (1) and (2) (under condition (4)) is the LOO kernel density estimator (LOO-KDE), defined as

$$\hat{p}_{-i}^{n,k}(x_i) = \frac{1}{(n-k)h} \sum_{\substack{j=k \\ j \neq i}}^n K \left(\frac{x_i - x_j}{h} \right) \quad (6)$$

where $K(\cdot) \geq 0$ is a kernel function and $h > 0$ is a smoothing parameter. The KL loss for kernel density estimators is analyzed carefully in [20], where it is shown that the rate of convergence in KL loss is slower than that of MISE for most well-behaved densities. Nevertheless, this loss indeed converges to zero with a polynomial decay rate with the use of appropriate kernel functions, and thus (1) is satisfied. Furthermore, using (4), it can be shown that (2) is also satisfied. We note that the actual choices of β_1 and β_2 do not affect the first-order asymptotic optimality result given in Thm 3.3.

3. QCD WITH LOO DENSITY ESTIMATION

Let $X_1, X_2, \dots, X_n, \dots \in \mathbb{R}^d$ be a sequence of independent random variables (or vectors), and let ν be a change-point. Assume that $X_1, \dots, X_{\nu-1}$ all have density p_0 with respect to some measure μ . Furthermore, assume that $X_\nu, X_{\nu+1}, \dots$ have densities p_1 also with respect to μ . Here p_0 is assumed to be completely known. While p_1 is completely unknown, we assume that (1) and (2) are satisfied for LOO estimators of p_1 .

Let $\mathbb{P}_\nu$ denote the probability measure on the entire sequence of observations when the change-point is ν , and let $\mathbb{E}_\nu[\cdot]$ denote the corresponding expectation. The change-time ν is assumed to be unknown but deterministic. The problem is to detect the change quickly, while controlling the false alarm rate. Let τ be a stopping time [8] defined on the observation sequence associated with the detection rule, i.e. τ is the time at which we stop taking observations and declare that the change has occurred.

3.1. Classical Results

When p_1 is known, Lorden [3] proposed solving the following optimization problem to find the best stopping time τ :

$$\inf_{\tau \in \mathcal{C}_\alpha} \text{WADD}(\tau) \quad (7)$$

where

$$\text{WADD}(\tau) := \sup_{\nu \geq 1} \text{ess sup } \mathbb{E}_\nu \left[(\tau - \nu + 1)^+ | \mathcal{F}_{\nu-1} \right] \quad (8)$$

characterizes the worst-case delay, and $\mathcal{F}_n$ denotes the sigma algebra generated by $X_1, \dots, X_n$, i.e., $\mathcal{F}_n = \sigma(X_1, \dots, X_n)$. The constraint set in (7) is

$$\mathcal{C}_\alpha := \{\tau : \text{FAR}(\tau) \leq \alpha\} \quad (9)$$

with $\text{FAR}(\tau) := \frac{1}{\mathbb{E}_\infty[\tau]}$ which guarantees that the false alarm rate of the algorithm does not exceed α . Here, $\mathbb{E}_\infty[\cdot]$ is the expectation operator when the change never happens, and $(\cdot)^+ := \max\{0, \cdot\}$.

Lorden also showed that Page's Cumulative Sum (CuSum) algorithm [22] whose test statistic is given by:

$$W(n) = \max_{1 \leq k \leq n+1} \sum_{i=k}^n \log \frac{p_1(X_i)}{p_0(X_i)}$$

solves the problem in (7) asymptotically as $\alpha \rightarrow 0$. The CuSum stopping rule is given by:

$$\tau_{\text{Page}}(b) := \inf\{n : W(n) \geq b\} \quad (10)$$

where the threshold is set as $b = |\log \alpha|$. It was shown by Moustakides [23] that the CuSum algorithm is exactly optimal for the problem in (7). The asymptotic performance is

$$\inf_{\tau \in \mathcal{C}_\alpha} \text{WADD}(\tau) \sim \text{WADD}(\tau_{\text{Page}}(|\log \alpha|)) \sim \frac{|\log \alpha|}{D(p_1 || p_0)} \quad (11)$$

as $\alpha \rightarrow 0$. Here $Y_\alpha \sim G_\alpha$ is equivalent to $Y_\alpha = G_\alpha(1+o(1))$. Also, we use the notation $o(1)$ to denote a quantity that goes to 0, as $\alpha \rightarrow 0$ or $b \rightarrow \infty$.

When the post-change distribution has parametric uncertainties, Lai [4] generalized this performance guarantee with the following assumptions. Suppose that p_0 and p_1 satisfy

$$\sup_{\nu \geq 1} \mathbb{P}_\nu \left\{ \max_{t \leq n} \sum_{i=\nu}^{\nu+t} Z_i \geq (1+\delta)nI \right\} \xrightarrow{n \rightarrow \infty} 0 \quad (12)$$

for any $\delta > 0$, and

$$\sup_{t \geq \nu \geq 1} \mathbb{P}_\nu \left\{ \sum_{i=t}^{t+n} Z_i \leq (1-\delta)nI \right\} \xrightarrow{n \rightarrow \infty} 0 \quad (13)$$

for any $\delta \in (0, 1)$, with some constant $I > 0$. Also, suppose that the window size m_α satisfies

$$\liminf m_\alpha / |\log \alpha| > I^{-1} \quad \text{and} \quad \log m_\alpha = o(|\log \alpha|).$$

Then, the window-limited (WL) GLR CuSum test:

$$\tilde{\tau}_{GLR}(b) := \inf \left\{ n : \max_{n-m_\alpha \leq k \leq n+1} \sup_{\theta \in \Theta} \sum_{i=k}^n Z_{i,k}^\theta \geq b \right\} \quad (14)$$

with some test threshold $b_\alpha \sim |\log \alpha|$ solves the problem in (7) asymptotically as $\alpha \rightarrow 0$. The asymptotic performance is

$$\inf_{\tau \in \mathcal{C}_\alpha} \text{WADD}_\theta(\tau) \sim \text{WADD}_\theta(\tilde{\tau}_{\text{Page}}(b_\alpha)) \sim \frac{|\log \alpha|}{I}. \quad (15)$$

Note that $I = D(p_1 || p_0)$ when p_0 and p_1 are independent.

3.2. Leave-one-out (LOO) CuSum Test

For the case when p_1 is unknown, we define the LOO log-likelihood ratio as

$$\hat{Z}_i^{n,k} = \log \frac{\hat{p}_{-i}^{n,k}(X_i)}{p_0(X_i)}, \quad \forall k \leq i \leq n. \quad (16)$$

and the LOO CuSum stopping rule as

$$\hat{\tau}(b) := \inf \left\{ n : \max_{(n-m_\alpha)^+ \leq k \leq n-1} \sum_{i=k}^n \hat{Z}_i^{n,k} \geq b \right\}. \quad (17)$$

Here the window size m_α is designed to satisfy

$$\liminf m_\alpha / |\log \alpha| > fI^{-1} \quad \text{with} \quad \log m_\alpha = o(|\log \alpha|) \quad (18)$$

where $f > 1$ is a constant.

In Lemma 3.1, we show that $\hat{\tau}$ with a properly chosen threshold b_α satisfies the false alarm constraint in (9) asymptotically. In Lemma 3.2, we establish an asymptotic upper bound on $\text{WADD}(\hat{\tau}(b_\alpha))$. Finally, in Theorem 3.3, we combine the two lemmas and establish the first-order asymptotic optimality of $\hat{\tau}(b_\alpha)$.

Lemma 3.1. Suppose that b_α satisfies

$$b_\alpha = |\log \alpha| + \log(8m_\alpha). \quad (19)$$

Then,

$$\mathbb{E}_\infty [\hat{\tau}(b_\alpha)] \geq \alpha^{-1}.$$

Remark. If m_α satisfies (18), then $b_\alpha = |\log \alpha| (1 + o(1))$ as $\alpha \rightarrow 0$.

Proof Sketch. For all thresholds b , we upper bound $\mathbb{P}_\infty \{\nu \leq \hat{\tau}(b) < \nu + m_\alpha\}$ by upper bounding $\mathbb{P}_\infty \{\tau_k(b) \leq k + m_\alpha\}$ for each k , where τ_k is an auxiliary stopping time as seen in [4, Proof of Lemma 2]. We use the fact that the density estimator produces a density that is independent of X_i for each i . Then, we use [24, Lemma 2.2(ii)] to finish the proof. $\square$

Lemma 3.2. Let $b_\alpha = |\log \alpha| (1 + o(1))$ and m_α satisfy (18). Suppose that (1) and (2) hold. Further, suppose (13) holds. Then,

$$\text{WADD}(\hat{\tau}(b_\alpha)) \leq \frac{|\log \alpha|}{D(p_1 \| p_0)} (1 + o(1))$$

as $\alpha \rightarrow 0$.

Proof Sketch. The key part in the proof is to show that

$$\sup_{t \geq \nu \geq 1} \text{ess sup } \mathbb{P}_\nu \left\{ \sum_{i=t}^{t+n_b-1} \hat{Z}_i^{t+n_b-1, t} < b \mid \mathcal{F}_{t-1} \right\} < 2\delta_b^2 \quad (20)$$

with $n_b := \left\lfloor \frac{b}{I(1-\delta_b)} \right\rfloor$ and some δ_b satisfying $\delta_b \searrow 0$ as $b \nearrow \infty$. The left-hand side in (20) can be upper bounded by

$$\mathbb{P}_1 \left\{ \sum_{i=1}^{n_b} Z_i \leq b + \epsilon \right\} + \mathbb{P}_1 \left\{ \frac{1}{n_b} \sum_{i=1}^{n_b} (Z_i - \hat{Z}_i^{n_b, 1}) \geq \frac{\epsilon}{n_b} \right\}$$

and it remains to choose a proper $\epsilon = \epsilon_b$ in order to keep both terms small. The idea is to choose ϵ_b by controlling the second term via concentration inequalities. Then, it can be verified that ϵ_b is small enough for the first term to vanish when b is large. Indeed, we have $\epsilon_b = o(n_b)$ as $b \rightarrow \infty$. Then the rest of the proof is similar to [4, Proof of Theorem 4]. $\square$

Theorem 3.3. Suppose that b_α is chosen as in (19), with a window size m_α large enough to satisfy (18), and suppose that (12), (13) hold for the true log-likelihood ratio. Then $\hat{\tau}(b_\alpha)$ with $\hat{\tau}$ defined in (17) solves the problem in (7) asymptotically as $\alpha \rightarrow 0$. The worst case delay is

$$\inf_{\tau \in \mathcal{C}_\alpha} \text{WADD}(\tau) \sim \text{WADD}(\hat{\tau}(b_\alpha)) \sim \frac{|\log \alpha|}{D(p_1 \| p_0)}$$

as $\alpha \rightarrow 0$.

Proof. The asymptotic lower bound on the delay follows from (13) by using [4, Thm. 1]. The asymptotic optimality of $\hat{\tau}(b_\alpha)$ follows from Lemma 3.1 and Lemma 3.2. $\square$




Fig. 1. Comparison of operating characteristics of LOO-CuSum (solid lines) with the CuSum test (in red) and the WL-GLR-CuSum test (dotted lines) in detecting a shift in mean of a Gaussian. The pre- and post-change distributions are $\mathcal{N}(0, 1)$ and $\mathcal{N}(0.5, 1)$. The change-point $\nu = 1$. The kernel width parameter $h = (\min\{n, m\} - 1)^{-1/5}$ where m is the window size.

4. NUMERICAL RESULTS

In Fig. 1, we study the performance of the proposed LOO-CuSum test defined in (17) through Monte Carlo (MC) simulations when the pre-change distribution is $\mathcal{N}(0, 1)$. The LOO-KDE (defined in (6)) is used to estimate the density. The actual post-change distribution is $\mathcal{N}(0.5, 1)$, but the LOO-CuSum test has no knowledge of it. The performance of the LOO-CuSum test is compared with that of the following tests:

1. the CuSum test (in (10)), which assumes full knowledge of the post-change distribution;
2. the WL-GLR-CuSum test, which assumes that the post-change distribution belongs to $\{\mathcal{N}(\theta, 1)\}_{\theta > 0}$.

Different window sizes are considered, among which window sizes of 100 and 200 are sufficiently large to cover the full range of delay. It is seen that the expected delay of the LOO-CuSum test is close to that of the WL-GLR-CuSum test for all window sizes considered.

5. CONCLUSION

We studied a window-limited LOO-CuSum test for QCD that does not assume any knowledge of the post-change distribution, and does not require post-change training samples. We established the first-order asymptotic optimality of the test, and validated our analysis through numerical results.

6. REFERENCES

- [1] V. V. Veeravalli and T. Banerjee, "Quickest change detection," in *Academic press library in signal processing: Array and statistical signal processing*. Cambridge, MA: Academic Press, 2013.
- [2] 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*, vol. 2, no. 2, pp. 494–514, 2021.
- [3] G. Lorden, "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [4] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2917–2929, November 1998.
- [5] Y. Liang and V. V. Veeravalli, "Quickest detection of composite and non-stationary changes with application to pandemic monitoring," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 5807–5811.
- [6] Y. Liang, A. G. Tartakovsky, and V. V. Veeravalli, "Quickest change detection with non-stationary post-change observations," *IEEE Transactions on Information Theory*, 2022.
- [7] P. J. Huber, "A robust version of the probability ratio test," *The Annals of Mathematical Statistics*, vol. 36, no. 6, pp. 1753–1758, Dec. 1965.
- [8] P. Moulin and V. V. Veeravalli, *Statistical Inference for Engineers and Data Scientists*. Cambridge, UK: Cambridge University Press, 2018.
- [9] T. L. Molloy and J. J. Ford, "Misspecified and asymptotically minimax robust quickest change detection," *IEEE Transactions on Signal Processing*, vol. 65, no. 21, pp. 5730–5742, 2017.
- [10] J. Unnikrishnan, V. V. Veeravalli, and S. P. Meyn, "Minimax robust quickest change detection," *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1604–1614, 2011.
- [11] S. Li, Y. Xie, H. Dai, and L. Song, "M-statistic for kernel change-point detection," in *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds., vol. 28. Curran Associates, Inc., 2015, pp. 3366–3374.
- [12] T. Flynn and S. Yoo, "Change detection with the kernel cumulative sum algorithm," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 6092–6099.
- [13] F. Desobry, M. Davy, and C. Doncarli, "An online kernel change detection algorithm," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 2961–2974, 2005.
- [14] L. Chu and H. Chen, "Sequential change-point detection for high-dimensional and non-euclidean data," *IEEE Transactions on Signal Processing*, vol. 70, pp. 4498–4511, 2022.
- [15] H. Chen, "Sequential change-point detection based on nearest neighbors," *The Annals of Statistics*, vol. 47, no. 3, pp. 1381–1407, 2019. [Online]. Available: <https://doi.org/10.1214/18-AOS1718>
- [16] M. N. Kurt, Y. Yilmaz, and X. Wang, "Real-time non-parametric anomaly detection in high-dimensional settings," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 7, pp. 2463–2479, 2021.
- [17] Y. Kawahara and M. Sugiyama, "Sequential change-point detection based on direct density-ratio estimation," *Statistical Analysis and Data Mining: The ASA Data Science Journal*, vol. 5, no. 2, pp. 114–127, 2012.
- [18] G. V. Moustakides and K. Basioti, "Training neural networks for likelihood/density ratio estimation," *arXiv preprint arXiv:1911.00405*, Nov. 2019.
- [19] T. Lau, W. P. Tay, and V. Veeravalli, "A binning approach to quickest change detection with unknown post-change distribution," *IEEE Transactions on Signal Processing*, vol. 67, no. 3, pp. 609–621, Nov. 2018.
- [20] P. Hall, "On Kullback-Leibler loss and density estimation," *The Annals of Statistics*, vol. 15, no. 4, pp. 1491–1519, 1987.
- [21] D. W. Scott, *Multivariate Density Estimation: Theory, Practice, and Visualization*, 2nd ed. Hoboken, NJ: John Wiley & Sons, Inc., 2015.
- [22] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115, Jun. 1954.
- [23] G. V. Moustakides, "Optimal stopping times for detecting changes in distributions," *Annals of Statistics*, vol. 14, no. 4, pp. 1379–1387, Dec. 1986.
- [24] A. G. Tartakovsky, *Sequential Change Detection and Hypothesis Testing: General Non-i.i.d. Stochastic Models and Asymptotically Optimal Rules*, ser. Monographs on Statistics and Applied Probability 165. Boca Raton, London, New York: Chapman & Hall/CRC Press, Taylor & Francis Group, 2020.