

Sequential Change Detection by Optimal Weighted ℓ_2 Divergence

Liyan Xie[✉], *Graduate Student Member, IEEE*, and Yao Xie[✉], *Member, IEEE*

Abstract—We present a new non-parametric statistic, called the weighted ℓ_2 divergence, based on empirical distributions for sequential change detection. We start by constructing the weighted ℓ_2 divergence as a fundamental building block for two-sample tests and change detection. The proposed statistic is proved to attain the optimal sample complexity in the offline setting. We then study the sequential change detection using the weighted ℓ_2 divergence and characterize the fundamental performance metrics, including the average run length (ARL) and the expected detection delay (EDD). We also present practical algorithms to find the optimal projection to handle high-dimensional data and the optimal weights, which is critical to quick detection since, in such settings, there are not many post-change samples. Simulation results and real data examples are provided to validate the good performance of the proposed method.

Index Terms—Change-point detection, convex optimization, non-parametric methods, online algorithms, weighted Euclidean distance.

I. INTRODUCTION

SEQUENTIAL change detection is a classic problem in statistics and information theory. The goal is to detect the change in the underlying distribution as quickly as possible after it occurs. There is a wide range of applications, including sensor networks [1], seismology [2], social networks [3], power systems [4], and genomics [5]. Many classic results and detection procedures have been developed, see [6]–[8]. However, many widely used methods assume a parametric form of the distributions before and after the change. For high-dimensional data, such parametric methods can be difficult to implement since the post-change distribution is typically unknown. We cannot have a large number of samples to estimate the distribution. Recently, there have been many interests in developing a non-parametric change detection procedure for high-dimensional streaming data when we have limited post-change samples.

We focus on a type of *distribution-free* methods based on empirical distributions. Compared with parametric methods,

Manuscript received October 19, 2020; revised February 23, 2021; accepted April 5, 2021. Date of publication April 13, 2021; date of current version June 21, 2021. The work of Liyan Xie and Yao Xie was supported in part by an NSF CAREER Award under Grant CCF-1650913, Grant DMS-1938106, Grant DMS-1830210, Grant CCF-1442635, and Grant CMMI-1917624. (Corresponding author: Yao Xie.)

The authors are with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: lxie49@gatech.edu; yao.xie@isye.gatech.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/JSAIT.2021.3072960>, provided by the authors. The supplemental material includes the proofs to the main theorems.

Digital Object Identifier 10.1109/JSAIT.2021.3072960

such non-parametric tests are more flexible and can be more applicable for various real-world situations. They tend to perform better when (i) the data does not follow a parametric distribution or (ii) we do not have enough samples to estimate the underlying distribution reliably. However, one particular challenge is to establish performance guarantees and improve the sample efficiency of the non-parametric test statistic [9].

In this paper, we develop a new data-driven distribution-free sequential change detection procedure based on the *weighted ℓ_2 divergence* between empirical distributions as the test statistic, which is related to the idea of *testing closeness* between two distributions [10]. We introduce “weights” that are design parameters, which can be particularly important in achieving good performance in practice when we do not have a large number of samples. We show the optimality of the proposed procedure in achieving the theoretical lower bound of the sample complexity required for a low-risk test that meets the specifications. We also characterize the proposed detection procedure’s theoretical performance based on weighted ℓ_2 divergence for data in a sliding window. Moreover, we develop practical optimization procedures for selecting the optimal weights and the low-dimensional projections for high-dimensional data. The good performance of the proposed procedure is demonstrated using synthetic and real data.

The rest of the paper is organized as follows. Section II introduces preliminaries about the problem set-up and reviews related work. Section III proposes the weighted ℓ_2 test and shows its optimality in ℓ_2 sense. Section IV studies sequential change detection using the proposed statistic. Section V discusses different aspects to optimize the parameters involved in the proposed method. Section VI and Section VII demonstrate the performance of the proposed detection procedure using simulation and real-data study. All proofs are delegated to the supplementary material.

II. PRELIMINARIES

A. Weighted ℓ_2 Divergence for Two-Sample Test

We start by considering the problem of testing closeness between two discrete distributions from samples observed. The problem set-up is as follows. Let Ω be an n -element *observation space*. For notational convenience, we identify Ω with $\{1, \dots, n\}$ and a probability distribution on Ω with a vector $p \in \Delta_n = \{p \in \mathbb{R}^n : p \geq 0, \sum_i p_i = 1\}$, where p_i is the probability mass of the i -th element in Ω . Suppose we are given

two independent sample sets:

$$X^1 = \{x_1^1, \dots, x_{n_1}^1\} \text{ and } X^2 = \{x_1^2, \dots, x_{n_2}^2\},$$

where $x_1^1, \dots, x_{n_1}^1 \stackrel{\text{iid}}{\sim} p$, and $x_1^2, \dots, x_{n_2}^2 \stackrel{\text{iid}}{\sim} q$. Our goal is to design a *test* which, given observations X^1 and X^2 , claims one of the following hypotheses:

$$H_0 : p = q, \quad H_1 : \|p - q\|_2 \geq \epsilon \|p\|_2,$$

where $\|\cdot\|_2$ is the ℓ_2 norm on \mathbb{R}^n and $\epsilon > 0$ is a parameter that represents the relative difference of magnitude. Note that the alternative hypothesis H_1 considered here is slightly different but closely related to the traditional setting where H_1 is defined as $\|p - q\|_2 \geq \epsilon$.

Define the type-I risk of a test as the probability of rejecting hypothesis H_0 when it is true, i.e., the probability of claiming $\|p - q\|_2 \geq \epsilon \|p\|_2$ when $p = q$. The type-II risk is the probability of claiming $p = q$ when $\|p - q\|_2 \geq \epsilon \|p\|_2$. We aim at building a test which, given $0 < \alpha, \beta < 1/2$, has the type-I risk at most α (which we call *at level* α), and the type-II risk at most β (*of power* $1 - \beta$); and we aim to meet these specifications with sample sizes n_1 and n_2 as small as possible.

We propose a new type of test by considering a family of distance-based divergence between empirical distributions of the two sets of observations. More specifically, we consider tests that reject the null hypothesis H_0 (and accept the alternative H_1) when

$$D(X^1, X^2) > \ell,$$

where $D(\cdot, \cdot)$ is a proxy for the weighted ℓ_2 divergence between distributions p and q underlying X^1 and X^2 , and ℓ is a data-dependent (random) threshold.

Our motivation for considering the ℓ_2 divergence for the non-parametric test is twofold. First, the ℓ_2 divergence-based test has a certain (near) optimality that we will show in Section III. Second, the ℓ_2 divergence is more *robust* compared to other divergences such as χ^2 -divergence, which is commonly used when ℓ_1 separation between distributions is of interest. The χ^2 -divergence becomes numerically difficult to evaluate when there are “small” p_i (meaning some atoms have small probability), while the ℓ_2 distance remains bounded and in such cases. Similar argument holds for the α -and β -divergences [11]–[15] and detection statistic for robust change detection [16]. Moreover, here we focus on a new *weighted* ℓ_2 divergence, which emphasizes atoms that contribute most to $\|p - q\|_2$.

Below, we derive the optimality in the sense that the sample size needed by the ℓ_2 divergence test matches (up to a moderate constant factor) the minimax lower bound in [17]. More specifically, the optimal sample complexity of testing two distributions p and q with $\|p - q\|_2 \geq \epsilon / \sqrt{n}$ is $O(1) \sqrt{n} / \epsilon^2$ (i.e., it requires $O(1) \sqrt{n} / \epsilon^2$ samples from each distribution), whence any low-risk test requiring ϵ to be on the order of $O(1) / \sqrt{n}$ requires a sample size at least $O(1) \sqrt{n}$.

We also propose a general framework for selecting optimal test parameters utilizing convex optimization. Namely, when extra prior information about unknown distributions p and

q is available, such information can be used to improve the test’s quality by efficiently adjusting weights, which are adaptive to the “closeness” at different parts of the two distributions. Furthermore, we develop optimal projection for dimension reduction of high-dimensional data by maximizing the Wasserstein distance between two samples [18]. This sheds light on the potential extension of the proposed test statistic to high-dimensional continuous distributions.

B. Related Work

There is a long history of studying similar problems in both statistics and computer science. In statistics, a two-sample test is a fundamental problem in which one aims to decide if two sets of observations are drawn from the same distribution [19], with a wide range of applications [20]. Available approaches to the two-sample test can be largely divided into two categories: parametric and non-parametric. The parametric approach assumes that the data distribution belongs to certain parametric families, but the parameters can be unknown [21]. The non-parametric setting does not impose any assumption on the underlying distribution and therefore is widely applicable to real scenarios.

Classical approaches focus on the so-called “goodness-of-fit” test to decide whether the observations follow a pre-specified distribution. Non-parametric goodness-of-fit tests can be generalized for two-sample (and multi-sample) tests; in this case, the focus is the asymptotic analysis when the sample size goes to infinity. For instance, the Kolmogorov-Smirnov test [22], and the Anderson-Darling test [23] focus on univariate distributions and compute divergences between the empirical cumulative distributions of two (and multi) samples. The Wilcoxon-Mann-Whitney test [24], [25] is based on the data ranks and is also limited to univariate distributions. Van der Waerden tests are based on asymptotic approximation using quantiles of the standard Gaussian distribution [26], [27]. The nearest neighbors test for multivariate data is based on the proportion of neighbors belonging to the same sample [28].

There is much work aimed at extending univariate tests to the multivariate setting. A distribution-free generalization of the Smirnov two-sample test was proposed in [29] by conditioning on the empirical distribution functions. Wald-Wolfowitz run test and Smirnov two-sample test were generalized to multivariate setting using minimal spanning trees in [30]. A class of distribution-free multivariate tests based on nearest neighbors was studied in [28], [31], [32], and a multivariate k -sample test based on Euclidean distance between sample elements was proposed in [33]. Some recent work includes methods based on maximum mean discrepancy (MMD) [34] and the Wasserstein distance [35]. In particular, the ℓ_2 test enables us to draw a conclusion directly based on comparing empirical distributions. Compared with existing methods such as the MMD test, which requires a huge gram matrix when the sample size is large, the ℓ_2 test enables us to choose weights flexibly to better serve the testing task.

Another line of research in theoretical computer science deals with *closeness testing*. It was first studied in [10], [36], in which the testing algorithm with sub-linear sample complexity was presented; the lower bound to the sample complexity was

gave in [37]; a test that meets the optimal sample complexity was proposed in [17]; see [38] and [39] for recent surveys. The ℓ_2 case has also been studied in [10], [17], [40], and optimal algorithms are given. Many variants of closeness testing have also been studied recently. In [41], sublinear algorithms were provided for generalized closeness testing. In [42], the closeness testing was studied under the case where sample sizes are unequal for two distributions. In [43], a nearly optimal algorithm for closeness testing for discrete histograms was given. In [44], the problem was studied from a differentially private setting.

Outstanding early contributions of sequential change detection mainly focus on parametric methods [45]–[48] and is well-summarized in recent books [8], [49]. Recently, there have been growing interests in the non-parametric hypothesis test used in change detection problems. In [50], the “QuantTree” framework was proposed to define the bins in high-dimensional cases recursively, and the resulted histograms are used for change detection. In [51], a sequential change detection procedure using nearest neighbors was proposed. In the seminal work [52], a binning strategy was developed to discretize the sample space to construct the detection statistic to approximate the well-known generalized likelihood ratio test. The binned detection statistic’s asymptotic properties were studied, and it was shown to be asymptotically optimal when the pre-and post-change distributions are discrete. Note that here we do not rely on likelihood ratios and assume the pre- and post-change distributions are unknown, and all we have are some possible “training data.”

III. WEIGHTED ℓ_2 DIVERGENCE TEST

Our goal in this section is to develop a test statistic, the weighted ℓ_2 divergence, used as the basic building block of the change detection procedure. We aim to construct a test with the following properties. When applied to two independent sets of size N , i.i.d. samples $\{x_1^1, \dots, x_N^1\}$ and $\{x_1^2, \dots, x_N^2\}$ drawn from unknown distributions $p, q \in \Delta_n$, the test

- (i) rejects the null hypothesis with probability at most a given α under $H_0 : p = q$;
- (ii) accepts the null hypothesis with probability at most a given β when there is a relative difference “of magnitude at least a given $\epsilon > 0$,” i.e., under $H_1 : \|p - q\|_2 \geq \epsilon \|p\|_2$.

We want to meet these reliability specifications with as small sample size N .

A. Test Statistic

The main ingredient of weighted ℓ_2 divergence test is the *individual test* built as follows. Let us fix “weights” $\sigma_i \geq 0$, $i = 1, \dots, n$, and let $\Sigma = \text{Diag}\{\sigma_1, \dots, \sigma_n\}$ be a diagonal matrix with diagonal entries being $\sigma_1, \dots, \sigma_n$. Given $\{x_1^1, \dots, x_N^1\}$ and $\{x_1^2, \dots, x_N^2\}$, we divide them into two *consecutive* (left) parts E, E' , of cardinality L each, and (right) parts F, F' , of cardinality R each, respectively. Note that the cardinality L and R are at most $N/2$ and can be less than $N/2$ if we do not use all N samples. Set

$$\gamma = \frac{R}{L+R}, \bar{\gamma} = 1 - \gamma = \frac{L}{L+R}, M = 2\gamma L = 2\bar{\gamma}R. \quad (1)$$

Let $\omega, \omega', \zeta, \zeta' \in \Delta_n$ be the empirical distributions of observations in sets E, E', F, F' , and χ be the weighted ℓ_2 test statistics defined as

$$\chi = (\omega - \zeta)^\top \Sigma (\omega' - \zeta') = \sum_{i=1}^n \sigma_i (\omega_i - \zeta_i)(\omega'_i - \zeta'_i). \quad (2)$$

The *weighted ℓ_2 divergence test* \mathcal{T} claims a change if and only if

$$|\chi| > \ell,$$

where ℓ is the threshold. The following lemma summarizes the properties of \mathcal{T} .

Proposition 1 (Test Properties): Let \mathcal{T} be the weighted ℓ_2 divergence test applied to a pair of samples drawn from distributions $p, q \in \Delta_n$, and let the threshold ℓ satisfy

$$\ell \geq 2\sqrt{2\theta} M^{-1} \sqrt{\sum_i \sigma_i^2 p_i^2}, \quad (3)$$

for some $\theta \geq 1$. Then

- 1) Risk: The type-I risk of \mathcal{T} is at most $1/\theta^2$.
- 2) Power: Under the assumption

$$\begin{aligned} & \sum_i \sigma_i (p_i - q_i)^2 \\ & > \ell + 2\sqrt{2\theta} \left[M^{-1/2} \sqrt{\sum_i \sigma_i^2 (p_i - q_i)^2 (\gamma p_i + \bar{\gamma} q_i)} \right. \\ & \quad \left. + M^{-1} \sqrt{\gamma \sum_i \sigma_i^2 p_i^2 + \bar{\gamma} \sum_i \sigma_i^2 q_i^2} \right], \end{aligned} \quad (4)$$

the power of \mathcal{T} is at least $1 - 3/\theta^2$.

For simplicity, in the rest of this section we assume that $\sigma_i = 1$, $1 \leq i \leq n$, so the left hand side of (4) reduces to $\|p - q\|_2^2$. In Section V-A, we will discuss how to utilize the non-uniform weights.

B. Special Case: ℓ_2 Test With Uniform Weights

The individual test \mathcal{T} in the previous section has two drawbacks: (i) to control the type-I risk, the threshold ℓ in (3) specifying \mathcal{T} must be chosen with respect to the magnitude $\|p\|_2$ which is typically *unknown*; (ii) to achieve a small type-I risk of \mathcal{T} we need to set a large θ , thus resulting in poor power of the test. This section will show that we can reduce these limitations by “moderately” increasing the sample sizes. To simplify the notation, from now on, we use the fixed value $\theta = 3$ (i.e., the type-I risk is at most $1/9$ and the power is at least $2/3$), and use $M = L = R$ as a special case of the definition in (1).

The testing procedure will be as follows. We first give the Algorithm 1 to specify the threshold ℓ that satisfies the condition (3) with high probability and then introduce the testing procedure.

Algorithm 1 Estimate a Tight Upper Bound on $\|p\|_2^2$

INPUT: Samples $X^1 := \{x_1, \dots, x_N\}$; Reliability $1 - \delta$; $m = \lceil \log_2(n) \rceil$; S in (6).

OUTPUT: A tight upper bound ϱ on $\|p\|_2^2$ satisfying the condition (5).

for all $i = 1$ to m **do**

$$\rho_i = 2^{-i/2};$$

Set $P_i \in \mathbb{R}_+$ as the solution to

$$3 \left[2^{7/4} P_i^{-1/2} \rho_i^{3/2} + 2 P_i^{-1} \rho_i \right] = \frac{1}{3} \rho_i^2;$$

Set $Q_i = \lceil P_i \rceil$;

Use $4S$ consecutive segments, of cardinality Q_i each, of the sample X^1 to build $2S$ pairs $\{(\xi_s, \xi'_s), s = 1, \dots, 2S\}$ of empirical distributions;

Set $\theta_s = \xi_s^T \xi'_s$ for $s = 1, \dots, 2S$;

Set Θ_i as the median of $\theta_1, \dots, \theta_{2S}$;

if $\Theta_i \geq 2\rho_i^2/3$ or $i = m$ or $N < 4SQ_{i+1}$ **then**
Terminate.

end if

end for

Return $\varrho = \Theta_i + \rho_i^2/3$.

1) Specifying Threshold: When the nominal distribution p is *unknown*, we perform a *training-step* – use part of the first set of observations to build, with desired reliability $1 - \delta$, a tight upper bound ϱ (the output of Algorithm 1) on the squared norm $\|p\|_2^2$ of the unknown distribution p such that

$$\mathbb{P} \left[\|p\|_2^2 \leq \varrho \leq 3\|p\|_2^2 \right] \geq 1 - \delta, \quad (5)$$

where the probability is taken with respect to the observations sampled from distribution p .

The training-step is organized in Algorithm 1, where the input parameter S is defined as

$$S = \min \left\{ S \in \mathbb{N} : \sum_{k=S}^{2S} \binom{2S}{k} \left(\frac{1}{3}\right)^k \left(\frac{2}{3}\right)^{2S-k} \leq \frac{\delta}{\lceil \log_2(n) \rceil} \right\}. \quad (6)$$

The definition in (6) has an intuitive explanation: S is the smallest number such that in $2S$ independent tosses of a coin, with probability of getting a head in each toss being $\leq 1/3$, the probability of getting at least S heads does not exceed δ/m , where $m = \lceil \log_2(n) \rceil$.

Properties of the training-step in Algorithm 1 can be summarized as follows.

Proposition 2 (Bounding $\|p\|_2^2$): Let $\rho_i = 2^{-i/2}$ and $i(p)$ be the smallest $i \leq m$ such that $\rho_i \leq \|p\|_2$ (note that $i(p)$ is well defined due to $\rho_m \leq n^{-1/2}$). Assume that the size of the first group of sample X^1 is at least $4SQ_{i(p)}$. Then the probability for the training-step to terminate in the first $i(p)$ stages and to output ϱ satisfying the condition (5) is at least $1 - \delta$, where δ is the reliability tolerance specifying the training-step. Besides this, the number of observations utilized in a successful training-step is at most

$$4SQ_{i(p)} = O(1) \ln(\ln(n)/\delta)/\|p\|_2. \quad (7)$$

2) Testing Procedure: After ϱ is built, we use the part of the first sample X^1 *not* used in the training-step and the entire second sample X^2 to run $K = K(\alpha, \beta)$ individual tests to make a decision. Here $\alpha < 1/2$ and $\beta < 1/2$ are pre-specified upper bounds on the type-I and type-II risks of the testing problem, and $K(\alpha, \beta)$ is the smallest integer such that the probability of getting at least $K/2$ heads in K independent tosses of a coin is

- (i) $\leq \alpha$, when the probability of getting head in a single toss is $\leq 1/9$,
- (ii) $\geq 1 - \beta$, when the probability of getting head in a single toss is $\geq 2/3$.

It is easy to check that $K \leq O(1)[\ln(1/\alpha) + \ln(1/\beta)]$.

The k -th individual test is applied to two $2M$ -long segments of observations taken first from the sample X^1 (and these are non-overlapping with the training-step observations), and second from X^2 , with non-overlapping segments of observations used in different individual tests. Here the positive integer M , same as the reliability tolerances δ, α, β , is a parameter of our construction, and the threshold ℓ for individual tests is chosen as

$$\ell = 6\sqrt{2}M^{-1}\sqrt{\varrho}. \quad (8)$$

After running K individual tests, we claim H_1 if and only if the number of tests where H_1 is claimed is at least $K/2$. The properties of the resulting ℓ_2 test are presented as follows.

Theorem 1 (Sample Complexity): Consider the ℓ_2 test above with design parameters $\delta, \alpha, \beta \in (0, 1/2)$ and M . Then for properly selected absolute constants $O(1)$, the following holds true. Let p, q be the true distributions from which X^1 and X^2 are sampled, and let the size N of X^1, X^2 satisfies

$$N \geq O(1)[\ln(\ln(n)/\delta)/\|p\|_2 + [\ln(1/\alpha) + \ln(1/\beta)]M]. \quad (9)$$

Then

- 1) The probability for the training-step in Algorithm 1 to be successful is at least $1 - \delta$, and when it happens there are enough observations to carry out K subsequent individual tests.
- 2) Under the condition that the training-step is successful:
 - a) The type-I risk (claiming H_1 when $p = q$) is at most α ;
 - b) For every $\epsilon > 0$, with positive integer M satisfying

$$M \geq O(1) \frac{1}{\epsilon^2 \|p\|_2}, \quad (10)$$

the type-II risk (claiming H_0 when $\|p - q\|_2 \geq \epsilon \|p\|_2$) is at most β .

C. Near-Optimality of Proposed ℓ_2 Divergence Test

From the above analysis, when testing a difference of magnitude $\|p - q\|_2 \geq \epsilon \|p\|_2$, reliable detection is guaranteed when the size N of samples X^1 and X^2 is at least $O(n^{1/2}\epsilon^{-2})$ (due to the fact that $\|p\|_2 \geq n^{-1/2}$), with just logarithmic in the reliability parameters factors hidden in $O(\cdot)$. We will show that the $O(n^{1/2})$ sample size is the best rate can achieve unless additional *a priori* information on p and q is available.

Proposition 3 (Optimality): Given cardinality n of the set Ω and sample size N . For i.i.d. N -observation samples X^1 and

X^2 , suppose there exists a low-risk test that can detect reliably any difference of magnitude $\|p - q\|_2 \geq \epsilon \|p\|_2$ for $0 < \epsilon < 1/2$ such that

- 1) for every distribution p , the type-I risk is at most a given $\alpha < 1/2$, and
- 2) for every distributions p, q satisfying $\|p - q\|_2 \geq \epsilon \|p\|_2$, the type-II risk is at most a given $\beta < 1/2$.

Then $N \geq O(1)\sqrt{n}$, with a positive absolute constant $O(1)$ that depends on α, β, ϵ .

D. Illustrating Example: Quasi-Uniform Distribution

Now we present an illustrative example using “quasi-uniform” distributions. Assume that the nominal distribution p and the alternative distribution q are *quasi-uniform*, i.e., there exists a known constant κ satisfying $2 \leq \kappa \leq n$ such that $\|p\|_\infty \leq \kappa/n$ and $\|q\|_\infty \leq \kappa/n$. Since $\|x\|_2 \leq \sqrt{\|x\|_1\|x\|_\infty}$, we have $\max[\|p\|_2, \|q\|_2] \leq \sqrt{\kappa/n}$, and hence the threshold

$$\ell = 6\sqrt{2}M^{-1}\sqrt{\kappa/n} \quad (11)$$

satisfies the condition (3) with $\theta = 3$ (recall that we are in the case of uniform weights $\sigma_i \equiv 1$). With this choice of ℓ , the right hand side of condition (4) is at most $6\sqrt{2}[2M^{-1}\sqrt{\kappa/n} + M^{-1/2}\sqrt{\kappa/n}\|p - q\|_2]$. To ensure the validity of condition (4) with $\theta = 3$, it suffices to have

$$\|p - q\|_2^2 \geq 6\sqrt{2}\left[2M^{-1}\sqrt{\kappa/n} + M^{-1/2}\sqrt{\kappa/n}\|p - q\|_2\right],$$

which holds when

$$\|p - q\|_2^2 \geq O(1)M^{-1}\sqrt{\kappa/n}, \quad (12)$$

with properly selected moderate absolute constant $O(1)$. For quasi-uniform distributions, $\|p - q\|_2$ is no larger than $2\sqrt{\kappa/n}$. Therefore, for $\|p - q\|_2 \geq \lambda n^{-1/2}$ with some $\lambda \in (0, 2\sqrt{\kappa}]$, the sample size M should satisfy

$$M \geq O(1) \frac{\sqrt{\kappa n}}{\lambda^2}$$

in order to ensure condition (12). We see that in the case of $L = R$, given $\alpha \ll 1, \beta \ll 1$, the sample size of

$$O(1)[\ln(1/\alpha) + \ln(1/\beta)] \frac{\sqrt{\kappa n}}{\lambda^2}$$

ensures that for the ℓ_2 test with the threshold (11), its type-I risk and type-II risk are upper bounded by α and β , respectively.

In the following, we provide numerical examples to validate the optimality results in Proposition 3. Suppose the support size n is even and set $L = R = M$ for simplicity. The experiment set-up is described as the following two steps.

- (i) Draw two $n/2$ -element subsets independently, Ω_1 and Ω_2 , of Ω from the uniform distribution on the family of all subsets of Ω of cardinality $n/2$.
- (ii) The samples X^1 are i.i.d. drawn from the uniform distribution on Ω_1 , denoted as p ; and the second group of samples X^2 are i.i.d. drawn from the uniform distribution on Ω_2 , denoted by q .

Therefore we have $\max[\|p\|_2, \|q\|_2] \leq \sqrt{2/n}$, implying that we can set the threshold as

$$\ell = 12M^{-1}n^{-1/2}.$$

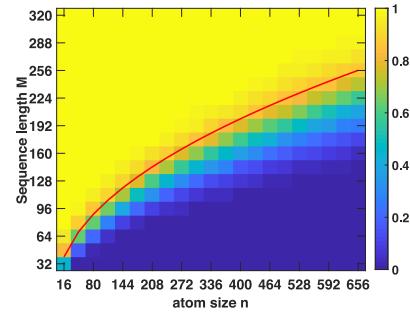


Fig. 1. Validation of the theoretical $O(\sqrt{n})$ bound by plotting the empirical test power of “quasi-uniform” in Section III-D, averaged over 1000 random trials. The Type-I risk is controlled to be less than 10^{-3} . The theoretical lower bound to sample complexity $O(\sqrt{n})$ is shown in red line, which match the empirical phase-transition “watershed.”

In all simulations, the individual test was applied. We perform the simulation for various n and M values. The power is shown in Fig. 1, averaged over 1000 trials. The results show that for magnitude $\|p - q\|_2 = O(1/\sqrt{n})$, at least $O(\sqrt{n})$ samples are required in order to detect the difference between p and q with high probability.

E. Comparison With Standard Identity Test

The most attractive feature of the ℓ_2 test is that to reliably test a distribution separation of relative “magnitude” ϵ , the sizes of samples should be of the order of $O(\sqrt{n}/\epsilon^2)$, just logarithmic in n and reliability tolerance (type-I and type-II risks) factors hidden in $O(\cdot)$. In contrast, the conceptually simplest approach to test – comparing norm $\|\hat{p} - \hat{q}\|_\pi$ of the difference of two empirical distributions with theoretical quantile of the recovery error as measured in the same norm – turns out to require sample sizes of the order of n . To the best of our knowledge, the only test proposed in the literature where the required sample size grows with n sublinearly is the “identity test” from [53] designed for the test *with known in advance nominal distribution p (the null distribution)*. It is shown in [53] that the number of observations allowing for the identity test to detect reliably the shift p to q is, up to logarithmic in the reliability parameters factors, $\mathcal{M}_p = \sqrt{n}/\|p - q\|_1^2$. When applying in the same situation of known p the ℓ_2 test, invoking our Theorem 1, the required number of observations is $\mathcal{M} = \max[1, \|p\|_2/\|p - q\|_2^2]$.

A natural way to compare the quality of the identity test and the ℓ_2 tests is to look at the *performance ratio* $\mathcal{C} = \mathcal{M}_p/\mathcal{M}$; $\mathcal{C} \gg 1$ means that the ℓ_2 test outperforms the identity test, while $\mathcal{C} \ll 1$ indicates the opposite situation. In the case of $\|p\|_2/\|p - q\|_2^2 \geq 1$, the performance ratio becomes

$$\mathcal{C} = \frac{\sqrt{n}}{\|p\|_2} \cdot \frac{\|p - q\|_2^2}{\|p - q\|_1^2},$$

which is the product of two factors in the range of $[O(\sqrt{n}), O(n)]$ and $[O(1/n), 1]$, respectively. Thus, \mathcal{C} is in the range $[O(1/\sqrt{n}), O(n)]$, implying that no one of the two tests always outperforms the other one. However, *with proper implementation*, ℓ_2 test always outperforms the identity test.

Specifically, with p known in advance, let us apply ℓ_2 test to an *equivalent* problem built as follows. Given p , let us define $k_i = \lfloor p_i n \rfloor$ for each point $i \in \Omega = \{1, \dots, n\}$ and define its p -children (i, k) , $1 \leq k \leq k_i$. Let $\bar{\Omega}$ be the set of children of all points from Ω . Observing a random variable ξ taking values in Ω , we can convert these observations into observations of random variable $\bar{\xi}$ taking values in $\bar{\Omega}$ as follows: observing value i of the random variable ξ , we throw a perfect dice with k_i faces and take (i, k) as the value of $\bar{\xi}$, where k being the observed face of the dice. Note that the distribution \bar{p} of $\bar{\xi}$ is obtained from the distribution p of ξ by spreading, for every $i \in \Omega$, the probability mass p_i of i in equal portions p_i/k_i between the children of i . In particular, \bar{p} has all entries at most $1/n$, while the cardinality \bar{n} of $\bar{\Omega}$ clearly is at most $2n$. Now, given samples X^1 and X^2 of cardinality N each, the distribution being (known) p and (unknown) q , let us use the above randomization to convert the samples into i.i.d. samples, of the same size, drawn from distributions \bar{p} and \bar{q} , and use these resulting samples to perform testing. With this approach, the sample size allowing to detect reliably the difference between p and q , up to factor logarithmic in reliability tolerances, will become $\bar{\mathcal{M}} = \max[1, \|\bar{p}\|_2/\|\bar{q} - \bar{p}\|_2^2]$. In this case of $\|\bar{p}\|_2/\|\bar{q} - \bar{p}\|_2^2 \geq 1$, we have

$$\frac{\mathcal{M}_p}{\bar{\mathcal{M}}} = \frac{\sqrt{n}}{\|\bar{p}\|_2} \cdot \frac{\|\bar{p} - \bar{q}\|_2^2}{\|p - q\|_1^2}. \quad (13)$$

Taking into account that $\|\bar{p}\|_2 \leq \sqrt{2\bar{n}}$, $\|p - q\|_1 = \|\bar{p} - \bar{q}\|_1$ and $\bar{n} \leq 2n$, the first factor in the right hand side of (13) is at least $O(n)$, and the second is at least $O(1/n)$, meaning that $\mathcal{M}_p/\bar{\mathcal{M}}$ is at least $O(1)$. Therefore, with the above implementation, ℓ_2 test is never much worse and in fact can be much better than the identity test.

F. Comparison With Generalized π -Tests

Here we compare ℓ_2 test with π -tests based on estimating $\|\cdot\|_\pi$ -distance between distribution p and q given samples, and will make an argument that ℓ_2 metric is preferred than other metrics. Specifically, given samples X^1 and X^2 of cardinality $2N$ each and $\pi \in [1, \infty]$, consider the test as follows:

- We split the first set of sample X^1 into two non-overlapping samples – *training subsample* x^{tr} and *testing subsample* x^{ts} , of cardinality N each, and compute the empirical distributions $\hat{f}[x^{\text{tr}}]$, $\hat{f}[x^{\text{ts}}]$ of these samples. In the same fashion, we divide the second set of sample X^2 , thus getting two N -observation non-overlapping training subsample \bar{x}^{tr} and testing subsample \bar{x}^{ts} and the corresponding empirical distributions $\hat{f}[\bar{x}^{\text{tr}}]$, $\hat{f}[\bar{x}^{\text{ts}}]$;
- We use the full training sample $[x^{\text{tr}}, \bar{x}^{\text{tr}}]$ to build a *separator* vector $e \in \mathbb{R}^n$ such that

$$e^\top [\hat{f}[x^{\text{tr}}] - \hat{f}[\bar{x}^{\text{tr}}]] = \|\hat{f}[x^{\text{tr}}] - \hat{f}[\bar{x}^{\text{tr}}]\|_\pi, \quad \|e\|_{\pi_*} = 1,$$

where $\pi_* = \pi/(\pi - 1)$ and set $h = e/\|e\|_\infty$;

- We compute the quantity $\chi = h^\top [\hat{f}[x^{\text{ts}}] - \hat{f}[\bar{x}^{\text{ts}}]]$ and claim H_1 if and only if

$$\chi > \ell,$$

where the threshold ℓ is the parameter of our test.

The properties of the resulting test can be summarized as follows. Let $\rho_{\beta, n, \pi}(N)$ be an upper bound on the $(1 - \beta)$ -quantile of the $\|\cdot\|_\pi$ -error of recovering a probability distribution p on observation space $\Omega = \{1, \dots, n\}$ by the empirical distribution of N -element i.i.d. sample drawn from p . In other words, for every probability distribution f on Ω , the empirical distribution $\hat{f}[x^N]$ associated with drawn from f N -element i.i.d. sample x^N satisfies the relation

$$\mathbb{P}\left\{\|\hat{f}[x^N] - f\|_\pi > \rho_{\beta, n, \pi}(N)\right\} \leq \beta.$$

It can be shown that a tight, within absolute constant factor, choice of $\rho_{\beta, n, \pi}(N)$ is given by

$$\rho_{\beta, n, \pi}(N) = 2\sqrt{3 \ln(2n/\beta)} n^\pi / \sqrt{N}, \quad \pi = \max\left[\frac{1}{\pi} - \frac{1}{2}, 0\right]. \quad (14)$$

We have the following result.

Theorem 2 (Properties of Generalized π -Test): Let N be a positive integer and $\pi \in [1, \infty]$. For $\alpha \in (0, 1)$, let $\tau_N(\alpha)$ be the positive root of the equation

$$\frac{\tau^2}{1 + \frac{2}{3}\tau} = \frac{2}{N} \ln(2/\alpha), \quad (15)$$

and define $\tau'_N(\alpha) = \min\{1, \tau_N(\alpha)\}$. For the π -test utilizing samples X^1 and X^2 of cardinality $2N$ each, if the threshold is set as

$$\ell = 2\tau'_N(\alpha) \quad (16)$$

(let us denote this test $\mathcal{T}_{\pi, \alpha}^N$), one has:

- The type-I risk of $\mathcal{T}_{\pi, \alpha}^N$ is $\leq \alpha$;
- Given $\beta \in (0, 1/4)$, under the assumption

$$\|p - q\|_\pi \geq 6\rho_{\beta, n, \pi}(N) + 2\tau'_N(\alpha), \quad (17)$$

where p, q are true distribution underlying samples X^1 and X^2 , the type-II risk of the test $\mathcal{T}_{\pi, \alpha}^N$ is at most 4β .

Note that the power of $\mathcal{T}_{\pi, \alpha}^N$ is the same as of the “straightforward” test, which uses samples X^1 and X^2 to build empirical approximations of the corresponding distributions and claims H_1 when the $\|\cdot\|_\pi$ -distance between these approximations is “essentially larger” than (theoretical upper bound on the) $(1 - \alpha)$ -quantile of the $\|\cdot\|_\pi$ -norm of the approximation error. The advantage of the π -test over the straightforward one is that the π -test is less conservative in terms of the type-I risk. Let us compare π -tests with the ℓ_2 -test. It is easy to see that there is no reason to use π -tests with $\pi > 2$; indeed, when $\pi > 2$, the right hand side in (17) is the same as when $\pi = 2$, see (14), while the left hand side in (17) is non-increasing in π . Thus, considering the power of π -tests, we can restrict $\pi \in [1, 2]$.

Next, assume that α and N are such that $\tau = \tau_N(\alpha) < 1$ (this is the only nontrivial case, since otherwise the test $\mathcal{T}_{\pi, \alpha}^N$ never claims H_1). In this case, (15) implies that $\sqrt{\Theta} \leq \tau \leq \frac{4}{3}\sqrt{\Theta}$, with $\Theta = 2 \ln(2/\alpha)/N$. Finally, assume that $n/\beta \geq 2/\alpha$; under these assumptions (14) says that the first term in the right hand side of (17) dominates the second one, so that (17) implies that

in order for H_1 to be detected with reliability $1-4\beta$, we should have

$$\|p - q\|_\pi \geq O(1)n^{\frac{1}{\pi} - \frac{1}{2}}\sqrt{\ln(2n/\beta)/N}, \quad (18)$$

with moderate absolute constant $O(1)$. Condition (18) implies that the best of π -tests is the one with $\pi = 2$. Indeed, when $\pi \in [1, 2]$, by Hölder inequality, we have $\|p - q\|_\pi \leq n^{\frac{1}{\pi} - \frac{1}{2}}\|p - q\|_2$, implying that whenever condition (18) is satisfied by some p, q and some $\pi \in [1, 2]$, it is satisfied by the same p, q with $\pi = 2$.

Another implication is that the ℓ_2 test significantly outperforms the best of π -tests. Indeed, by Theorem 1, for all $\delta, \alpha, \beta \in (0, 1/2)$, samples of cardinality N , with

$$N \geq N_{\ell_2} = O(1) \frac{\epsilon^2 \ln(\ln(n)/\delta) + \ln(1/\alpha) + \ln(1/\beta)}{\epsilon^2 \|p\|_2}$$

allow to ensure conditional, the probability of the condition being at least $1 - \delta$, type-I and type-II risks are at most α and β , respectively. By (18), to ensure similar reliability properties for π -test with $\pi = 2$, we should have a much larger, for small $\|p\|_2$, the sample size

$$N \geq N_2 = O(1) \frac{\ln(2n/\beta)}{\epsilon^2 \|p\|_2^2}.$$

In the “extreme case” of $\|p\|_2 = O(1)/\sqrt{n}$, N_2/N_{ℓ_2} is as large as $O(1)\sqrt{n}$. The discussion above supports the near-optimality of the ℓ_2 test.

IV. SEQUENTIAL CHANGE DETECTION PROCEDURES

In this section, we construct the change detection procedure based on the proposed weighted ℓ_2 divergence test. Change detection is an important instance of the sequential hypothesis test, but it has unique characteristics that require a separate study due to different performance metrics considered. Since we do not know the change location, we have to perform scanning when forming the detection statistic. We discuss two settings: the offline scenario where we have fixed samples and the online setting where the data come sequentially.

A. Offline Change Detection by “Scan” Statistic

In the offline setting, we observe samples $X^T = \{x_1, \dots, x_T\}$ on a time horizon $t = 1, \dots, T$, with x_t ’s taking values in an n -element set $\Omega = \{1, \dots, n\}$. Assume there exists time $K \in \{1, \dots, T\}$ such that for $t \leq K$, x_t are i.i.d. drawn from some *pre-change* distribution p , and for $t \geq K+1$, x_t are i.i.d. drawn from the *post-change* distribution q . Our goal is to design a test which, based on the samples X^T , decides on the null hypothesis $K = T$ (“no change”) versus the alternative $K < T$ (“change”). Meanwhile, we want to control the probability of false alarm to be at most a given $\alpha > 0$, and under this restriction to make the probability of successfully detecting the change as large as possible, at least when K and $T - K$ both are moderately large and q “significantly differs” from p .

We use the proposed test in Section III to construct a scan statistic for change detection. Given T , we select a collection of *bases* B_j , $1 \leq j \leq J$. A base B is a segment of $\{1, \dots, T\}$ partitioned into three consecutive parts: *pre-change*

part B_{lf} , middle part B_{md} , and *post-change* part B_{rg} ; the last instant in B_{lf} is the first instant in B_{md} , and the first instant in B_{rg} is by 1 larger than the last instant in B_{md} . For example: $B_{\text{lf}} = \{1, \dots, 10\}$, $B_{\text{md}} = \{10, 11\}$, $B_{\text{rg}} = \{12, \dots, 20\}$. We associate with base B an *individual test* \mathcal{T}_B which operates with observations $\{x_t, t \in B_{\text{lf}} \cup B_{\text{rg}}\}$ only. This test aims at deciding on two hypotheses: (1) “No change:” there is no change on B , that is, either K is less than the first, or larger than or equal to the last time instant from B ; (2) “Change:” the change point K belongs to the middle set B_{md} .

Given $\alpha > 0$ and a base B , we call individual test \mathcal{T}_B associated with this base α -*feasible*, if the probability of false alarm for \mathcal{T}_B is at most α , meaning that whenever there is no change on the base B of the test, the probability for the test to claim change is at most α . Our “overall” test \mathcal{T} works as follows: we equip bases B_j , $1 \leq j \leq J$, with tolerances $\alpha_j > 0$ such that $\sum_j \alpha_j = \alpha$, and then associate with each base B_j with a α_j -feasible individual test \mathcal{T}_{B_j} (as given by the ℓ_2 test in Section III-A). Given observations X^T , we perform one by one the individual tests in some fixed order, until either (i) the current individual test claims change; when it happens, the overall test claims change and terminates, or (ii) all J individual tests are performed and no one of them claimed change; in this case, the overall test claims no change and terminates.

Proposition 4 (False Alarm Rate for Offline Change Detection): With the outlined structure of the overall test and under condition $\sum_j \alpha_j = \alpha$, the probability of false alarms for \mathcal{T} (of claiming change when $K = T$) is at most α .

B. Online Change Detection

Instead of giving a fixed duration of samples in the offline setting, the observations arrive sequentially for online detection tasks. The goal is to detect the change as quickly as possible, under the constraint that the false alarm rate is under control.

The proposed detection procedure based on ℓ_2 test is illustrated in Fig. 2. Given a sequence $\{x_t, t = 1, 2, \dots\}$, as each time t , we search over all possible change-points $k < t$. In particular, we form two sequences before k and two sequences between $[k, t]$ with the length all equal to $M_{t,k} = \lceil (t-k)/2 \rceil$; their corresponding empirical distributions are denoted as $\xi_{t,k}$, $\xi'_{t,k}$, and $\eta_{t,k}$, $\eta'_{t,k}$. The detection statistic $\chi_{t,k}$ is formed as:

$$\chi_{t,k} = M_{t,k}(\xi_{t,k} - \eta_{t,k})^\top \Sigma(\xi'_{t,k} - \eta'_{t,k}). \quad (19)$$

We note that the multiplicative term $M_{t,k}$ can be viewed as a *scaling* parameter (which is proportional to the standard deviation of the test statistic) such that the variance of $\chi_{t,k}$ is of a constant order as $t - k$ increases.

The online change-point detection procedure is given by a stopping time

$$\mathcal{T} := \inf \left\{ t : \max_{0 \leq k \leq t} \chi_{t,k} \geq b \right\}, \quad (20)$$

where b is a pre-specified threshold that needs to be determined by controlling the false alarm rate. An intuitive interpretation of \mathcal{T} is that at each time t , we search over all possible change-points $k < t$, and raise alarm if the maximum statistic exceeds the threshold.

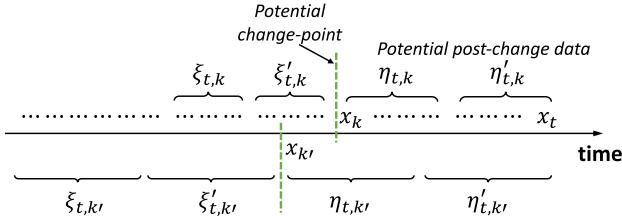


Fig. 2. Illustration of the sequential change detection procedure.

Remark 1 (Window-Limited Procedure): In practice, we can also adopt a *window-limited* version of \mathcal{T} as

$$\mathcal{T}' = \inf \left\{ t : \max_{m_0 \leq t-k \leq m_1} \chi_{t,k} \geq b \right\}, \quad (21)$$

where m_0 and m_1 are the lower and upper bounds of the window size that we would like to scan for the possible changes. Usually m_0 can be set such that the resulted sequences are long enough to have meaningful empirical distributions for constructing the detection statistic in (19). For practical considerations, we usually require the window size m_1 to be at least the expected detection delay, as discussed in [54] where the original theoretical study of the window-limited test was proposed.

Remark 2 (Comparison With the Binning Approach [52]): We note that the binning approach in [52] also considers discretized space and scans over all possible change-points when approximating log-likelihood ratio statistics. Compared with [52] which assumes the pre-change distribution is known, the detection procedure (20) and its window-limited version (21) do not need the prior knowledge of the pre-change distribution. We did not use the log-likelihood ratio statistic here but scan over all neighboring time windows directly to detect any significant difference in empirical distributions.

C. Theoretical Analysis for Online Change Detection

Now we characterize the two fundamental performance metrics for sequential change detection, namely the average run length (ARL) and the expected detection delay (EDD). We cannot use the previous method in Proposition 1 to determine the threshold because the bound is too conservative and will be intractable when the ARL is large. Here we present an asymptotic theoretical approximation.

1) *Theoretical Approximation to ARL:* To compute the ARL, we need to quantify the distribution of \mathcal{T} when data are sampled from the same distribution p . Intuitively, the detection statistic $\chi_{t,k}$ is small when the samples are from the same distribution. A relatively standard result is that when the threshold b tends to infinity, the stopping time's asymptotic distribution is approximately exponential when there is no change. This is proven true in various scenarios [55]–[57]. The main idea is to show that the number of boundary cross events for detection statistics over disjoint intervals converges to Poisson random variable in the total variation norm; the result can be established by invoking the Poisson limit theorem for dependent samples. Detailed proofs by adapting those techniques into the specific ℓ_2 test setting are left for future work. Under such

approximation, we have

$$\mathbb{P}_\infty(\mathcal{T}' \geq m) = \mathbb{P}_\infty \left(\max_{1 \leq t \leq m} \max_{m_0 \leq t-k \leq m_1} \chi_{t,k} \geq b \right) \approx e^{-\lambda m},$$

where \mathbb{P}_∞ is the probability measure when the change-point equals to ∞ , i.e., the change never happens; and \mathbb{E}_∞ denotes the corresponding expectation under this probability measure. Therefore, we only need to compute the probability $\mathbb{P}_\infty(\mathcal{T}' \geq m)$ and find the parameter λ , then the expectation of \mathcal{T}' equals $1/\lambda$. We adopt the change-of-measure transformation [58]–[60] and characterize the local properties of a random field. We first quantify the correlation between $\chi_{t,k}$ and $\chi_{\tau,s}$ in order to find the probability $\mathbb{P}_\infty(\mathcal{T}' \geq m)$ theoretically.

Proposition 5 (Temporal Correlation of Sequential Detection Statistics): Suppose all samples are i.i.d. drawn from the same distribution p , denote $M = t - k = \tau - s$, then the correlation between $\chi_{t,k}$ and $\chi_{\tau,s}$ is

$$\text{Corr}(\chi_{t,k}, \chi_{\tau,s}) = 1 - \frac{2}{M}(|t - \tau| + |k - s|) + \frac{(t - \tau)^2 + (s - k)^2}{M^2}.$$

Based on the correlation result, we have the following Theorem characterizing the ARL of the proposed ℓ_2 sequential detection procedure. The main idea is to use a linear approximation for the correlation between detection statistics $\chi_{t,k}$ and $\chi_{\tau,s}$. Then the behavior of the detection procedure can be related to a random field. By leveraging the localization theorem [60], we can obtain an asymptotic approximation for ARL when the threshold b is large enough (in the asymptotic sense). Define a special function $v(\cdot)$ which is closely related to the Laplace transform of the overshoot over the boundary of a random walk [61]:

$$v(x) \approx \frac{\frac{x}{2}[\Phi(\frac{x}{2}) - 0.5]}{\frac{x}{2}\Phi(\frac{x}{2}) + \phi(\frac{x}{2})}, \quad (22)$$

where $\phi(x)$ and $\Phi(x)$ are the probability density function and cumulative density function of the standard Gaussian distribution. For simplicity, we denote the variance of $\chi_{t,k}$ as

$$\sigma_p^2 := \text{Var}[\chi_{t,k}] = 4 \left[\sum_{i=1}^n \sigma_i^2 p_i^2 (1 - p_i)^2 + \sum_{i \neq j} \sigma_i \sigma_j p_i^2 p_j^2 \right]. \quad (23)$$

Theorem 3 (ARL Approximation): For large values of threshold $b \rightarrow \infty$, the ARL of the test \mathcal{T}' can be approximated as

$$\mathbb{E}_\infty[\mathcal{T}'] = \frac{1}{2} \frac{b^{-1} e^{b^2/(2\sigma_p^2)} [2\pi\sigma_p^2]^{1/2}}{\int_{[4b^2/(m_0\sigma_p^2)]^{1/2}}^{[4b^2/(m_1\sigma_p^2)]^{1/2}} y v^2(y) dy} (1 + o(1)). \quad (24)$$

The main contribution of Theorem 3 is to provide a theoretical method to set the threshold that can avoid the Monte Carlo simulation, which could be time-consuming, especially when ARL is large. Although there is no close-form analytical solution for b , when we let the right-hand side of

TABLE I
COMPARISON OF THE THRESHOLD b OBTAINED FROM SIMULATIONS AND
THE APPROXIMATION (24). SCANNING WINDOW $m_0 = 10, m_1 = 50$,
SUPPORT SIZE $n = 20$, NOMINAL DISTRIBUTION p IS UNIFORM

ARL	5k	10k	20k	30k	40k	50k
Simulation	2.0000	2.1127	2.2141	2.2857	2.3333	2.3750
Theoretical	1.8002	1.8762	1.9487	1.9897	2.0183	2.0398

Equation (24) equals to a specific ARL value (lower bound), we can numerically compute the right-hand side of (24) for any given threshold value b . Then we search over a grid to find the corresponding threshold values. Table I validates the approximation's good accuracy by comparing the threshold obtained from Equation (24) and compares it with that obtained by the Monte Carlo simulation. In detail, we generate 2000 independent trials of data from nominal distribution p and perform the detection procedure \mathcal{T}' for each trial; the ARL for each threshold b is estimated by the average stopping time over 2000 trials. In Table I, we report the threshold obtained through Monte Carlo simulation (as a proxy for the ground-truth) and on the approximation (24), for a range of ARL values. The ARL values in Table I correspond to the lower bound of an ARL; since ARL will increase when increasing the threshold. So if we have a good approximation, this can help us to calibrate the threshold and control the false alarm rate. The results in Table I indicate that the approximation is reasonably accurate since the relative error is around 10% for all specified ARL values. It is worth mentioning that ARL is very sensitive to the choice of threshold, making it challenging to estimate the threshold with high precision. However, the EDD is not that sensitive to the choice of the threshold, which means that a small difference in the threshold will not significantly change EDD.

2) *Theoretical Characterization of EDD*: After the change occurs, we are interested in the expected detection delay, i.e., the expected number of additional samples to detect the change. There are a variety of definitions for the detection delay [8], [48], [62], [63]. To simplify the study of EDD, it is customary to consider a specific definition $\mathbb{E}_0[\mathcal{T}']$, which is the expected stopping time when the change happens at time 0 and only depends on the underlying distributions p, q . It is not always true that $\mathbb{E}_0[\mathcal{T}']$ is equivalent to the standard worst-case EDD in literature [48], [62]. However, since $\mathbb{E}_0[\mathcal{T}']$ is certainly of interest and is reasonably easy to approximate, we consider it as a surrogate here. We adopt the convention that there are certain pre-change samples $\{x_{-1}, x_{-2}, \dots\}$ available before time 0, which can be regarded as reference samples.

Note that for any $t > 0$ and $k = 0$, the sequences $\xi_{t,0}$ and $\xi'_{t,0}$ come from the pre-change distribution p since they belong to the reference sequence $\{x_{-1}, x_{-2}, \dots\}$, and the sequences $\eta_{t,0}$ and $\eta'_{t,0}$ are from the post-change distribution q . Therefore, the expectation of the detection statistic $\chi_{t,k}$ is $\mathbb{E}[\chi_{t,k}] = \lceil(t - k)/2\rceil(p - q)^\top \Sigma(p - q)$, which determines the asymptotic growth rate of the detection statistic after the change. Using Wald's identity [64], we are able to obtain a first-order approximation for the detection delay, provided that the maximum window size m_1 is large enough compared to the EDD.

Theorem 4 (EDD Approximation): Suppose $b \rightarrow \infty$, with other parameters held fixed. If the window size m_1 is sufficiently large and greater than $2b/[(p - q)^\top \Sigma(p - q)]$, then the expected detection delay

$$\mathbb{E}_0[\mathcal{T}] = \frac{b(1 + o(1))}{(p - q)^\top \Sigma(p - q)/2}. \quad (25)$$

From the approximation result in Theorem 3, we note that the ARL is of order $O(e^{b^2/(2\sigma_p^2)})$ with respect to the threshold value b , which means that the threshold b chosen for a fixed ARL value γ should be on the order of $O(\log \sqrt{\gamma})$. Moreover, by taylor expansion and Csiszár-Kullback-Pinsker inequality [65], we have the Kullback–Leibler (KL) divergence is equivalent to squared ℓ_2 norm up to certain constants, i.e., $C_0 \|p - q\|_2^2 \leq \text{KL}(p, q) \leq C_1 \|p - q\|_2^2$, therefore the EDD is also of order $b/\text{KL}(p, q)$, which matches the theoretical lower bound of EDD in first-order.

Remark 3 (Optimize Weights to Minimize EDD): From the EDD approximation in (25), it is obvious that we can minimize EDD by optimizing over the weights matrix Σ . In particular, the EDD can be minimized when we can find the weights Σ such that the weighted ℓ_2 divergence between p and q is maximized. This is consistent with the subsequent discussion in Section V-A. In particular, when we have certain prior information about the distributions p and q , we could apply the optimization-based method in Section V-A to find the optimal weights to reduce the detection delay.

V. OPTIMIZED WEIGHTS AND PROJECTION OF HIGH-DIMENSIONAL DATA

This section discusses setting optimal weights that adapt to the closeness at different elements in Ω , given some *a prior* information on p and q . In addition, we tackle the data high-dimensionality by adopting the Wasserstein-based principal differences analysis [18] to find the optimal projection.

A. Optimize Weights for ℓ_2 Test

So far, we primarily focused on the case with uniform weights $\sigma_i \equiv 1$. In this section, we will discuss how to further improve performance by choosing the optimal weights. In the simplest case, when we know in advance (or can infer from additional “training” samples) that the distribution shift $p \rightarrow q$ (nearly) does not affect probabilities with indexes from some *known* set I , we can set $\sigma_i = 0$ for $i \in I$ and $\sigma_i = 1$ for $i \notin I$. This will keep the magnitude $\sum_i \sigma_i (p_i - q_i)^2$ on the left hand side of (4), as compared to uniform weights, intact, but will reduce the right hand side of (4).

A framework to optimize over σ_i 's is as follows. Assume that we know distributions p, q belong to a set $\mathcal{X} \subset \Delta_n$, which is defined by a set of quadratic constraints:

$$\mathcal{X} = \{p \in \Delta_n : p^\top Q_k p \leq 1, k = 1, \dots, K\}, \quad (26)$$

where $Q_k \in \mathbb{R}^{n \times n}$ are positive semi-definite ($Q_k \succeq 0$).

A natural way to measure “magnitude of difference” is to use $\|p - q\|_2$ (the case using $\|p - q\|_1$ can be similarly defined and solved). Assume we want to select $\sigma = [\sigma_1, \dots, \sigma_n] \geq 0$ to make reliable detection of difference $\|p - q\|_2 \geq \rho$, for

some given $\rho > 0$. To achieve this, we can impose a fixed upper bound on the right hand side in (4) when $p = q \in \mathcal{X}$, i.e., to require σ to satisfy

$$g_*(\sigma) := \max_{p \in \mathcal{X}} \sum_i \sigma_i^2 p_i^2 \leq a$$

with some given a , and to maximize under this constraint the quantity

$$f_*(\sigma) := \min_{p, q} \left\{ \sum_i \sigma_i (p_i - q_i)^2 : p, q \in \mathcal{X}, \|p - q\|_2 \geq \rho \right\}.$$

For any σ that satisfies $g_*(\sigma) \leq a$, the associated test which claims H_1 when the statistics (defined in (2)) $|\chi| > 2\sqrt{2\theta}M^{-1}\sqrt{a}$ is with type-I risk at most $1/\theta^2$. At the same time, large $f_*(\sigma)$ is in favor of good detection of distribution shift of magnitude $\|p - q\|_2 \geq \rho$. By the homogeneity in σ , we can set $a = 1$ without loss of generality.

In general, both g_* and f_* are difficult to compute. Therefore, we replace the problem

$$\max_{\sigma \geq 0} \{f_*(\sigma) : g_*(\sigma) \leq 1\}$$

with its *safe tractable approximation*:

$$\max_{\sigma \geq 0} \{f(\sigma) : g(\sigma) \leq 1\}, \quad (27)$$

where f is a concave efficiently computable *lower bound* on f_* , and g is a convex efficiently computable *upper bound* on g_* .

To build g , note that when $p \in \mathcal{X}$, the matrix $P = pp^\top \in \mathbb{R}^{n \times n}$ is positive semi-definite ($P \succeq 0$), non-negative in each entry ($P \geq 0$), $\sum_{i,j=1}^n P_{ij} = 1$, and $\text{Tr}(PQ_k) \leq 1$, $k \leq K$, by (26). Consequently, the function

$$g(\sigma) := \max \left\{ \text{Tr}(\Sigma^2 P) : P \succeq 0, P \geq 0, \sum_{i,j=1}^n P_{ij} = 1, \right. \\ \left. \text{Tr}(PQ_k) \leq 1, 1 \leq k \leq K \right\}$$

with $\Sigma := \text{Diag}\{\sigma_1, \dots, \sigma_n\}$ is an efficiently computable convex upper bound on g_* . Similarly, to build f , observe that the matrix $S = (p - q)(p - q)^\top$ stemming from $p, q \in \mathcal{X}$ with $\|p - q\|_2 \geq \rho$ belongs to the convex set

$$\mathcal{S} = \left\{ S : S \succeq 0, \sum_{i,j=1}^n |S_{ij}| \leq 4, \sum_{i,j=1}^n S_{ij} = 0, \right. \\ \left. \text{Tr}(S) \geq \rho^2, \text{Tr}(SQ_k) \leq 4, 1 \leq k \leq K \right\}.$$

Therefore,

$$f_*(\sigma) \geq f(\sigma) := \min_{S \in \mathcal{S}} \text{Tr}(\Sigma S)$$

and the function $f(\sigma)$ is concave and efficiently computable.

To implement the problem (27) efficiently, we derive the tractable dual formulation in the following. Note that these constraints can be greatly simplified if Q_k are *diagonal* matrices, especially for the high dimensional case.

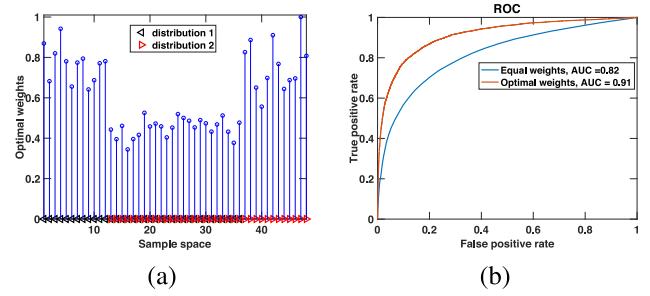


Fig. 3. Illustration of optimal weights on a simulated example. (a): Optimal weights; (b): The ROC curves under optimal weights and equal weights.

Proposition 6 (Dual Reformulation): The dual formulation of the optimization problem (27) is

$$\begin{aligned} & \max \lambda \rho^2 - 4 \sum_k x_k - 4\xi \\ \text{s.t. } & \lambda \geq 0, P \succcurlyeq 0, \xi \geq 0, x_k \geq 0, U \geq 0, W \geq 0, \\ & \Lambda \succcurlyeq 0, V \geq 0, \mu_k \geq 0, 1 \leq k \leq K, \\ & \sum_k x_k Q_k + U - W - P - rJ - \lambda I_n \succcurlyeq -\Sigma, \\ & U_{ij} + W_{ij} \leq \xi, 1 \leq i \leq n, 1 \leq j \leq n, \\ & \sum_k \mu_k - v \leq 1, -\Lambda - V + \sum_k \mu_k Q_k - vJ \succcurlyeq \Sigma^2, \end{aligned}$$

where $\Sigma = \text{Diag}\{\sigma_1, \dots, \sigma_n\}$ and $J \in \mathbb{R}^{n \times n}$ is a matrix with all elements equal to 1.

We present an illustrative simulation example to show the benefits of optimizing weights σ . The experimental set-up is as follows. Consider the sample space $\Omega = \{1, \dots, n\}$ with $n = 48$. The distributions p and q are set as uniform distributions on the subset $\Omega_1 = \{1, \dots, 3n/4\}$ and $\Omega_2 = \{n/4 + 1, \dots, n\}$, respectively. The common support of p and q consists of $n/2$ elements. We first use training data to estimate the matrix Q in our formulation. Specifically, we sample 32 observations from each distribution and compute the empirical distribution of all observations. This process is repeated for $m = 50$ trials, and the resulting Q is solved from the following optimization problem

$$\begin{aligned} & \min \log \det A^{-1} \\ \text{s.t. } & \|Ap_i\|_2 \leq 1, i = 1, \dots, m, \end{aligned}$$

where $A = Q^{1/2}$ and p_i is the empirical distribution in the i -th trial. The volume of the ellipsoid defined with Q is proportional to $\det A^{-1}$. Thus the solution to the above optimization problem is the *minimum ellipsoid* that contains the m empirical distributions [66]. The optimal weights are shown in Fig. 3(a). Moreover, we compare the ROC curve of the test under equal weights $\sigma_i = 1$ and optimal weights in Fig. 3(b), averaged over 10,000 trials. The result shows the benefits of using optimal weights.

B. Optimal Projection for High-Dimensional Data

Assume that the two distributions p and q , rather than discrete distributions on a given finite set, are continuous distributions on \mathbb{R}^d . In this situation, we may try to convert

observations $x \in \mathbb{R}^d$ into observations $f(x)$ taking values in a *finite* set and apply the proposed test to the transformed observations.

One way to build f is to project observations x onto one-dimensional subspace and then split the range of the projection into bins. We propose to select this subspace using, when available, “training sample” x_1, \dots, x_{2T} , with the first T observations drawn, independently of each other, from the nominal distribution p , and the last T observations drawn independently of each other and of x_1, \dots, x_T , from the distribution q . A natural selection of the one-dimensional subspace can be as follows. Denote by e the unit vector spanning the subspace. Let us look at the sample empirical distributions of the projections of the observations x_1, \dots, x_{2T} on e , and try to find unit vector e for which the Wasserstein distance between the distributions of the first half and the second half of the projections is as large as possible [18]. The distance above is, up to factor $1/T$, the quantity

$$\begin{aligned} \phi(e) &= \min_{\omega_{ij}, 1 \leq i, j \leq 2T} \left\{ \sum_{i,j} |e^\top (x_i - x_j)| \omega_{ij} : \right. \\ &\quad \sum_j \omega_{ij} = \begin{cases} 1, & i \leq T \\ 0, & i > T \end{cases}; \sum_i \omega_{ij} = \begin{cases} 0, & j \leq T \\ 1, & j > T \end{cases} \\ &\quad \omega_{ij} \geq 0, 1 \leq i, j \leq 2T \\ &= \max_{\lambda} \left\{ \sum_{i=1}^T \lambda_i - \sum_{i=T+1}^{2T} \lambda_i : \right. \\ &\quad \lambda_i - \lambda_j \leq |e^\top (x_i - x_j)|, 1 \leq i, j \leq 2T \left. \right\} \\ &= \Phi(E[e]), \end{aligned}$$

where

$$\begin{aligned} E[e] &= ee^\top, \\ \Phi(E) &= \max_{\lambda} \left\{ \sum_{i=1}^T \lambda_i - \sum_{i=T+1}^{2T} \lambda_i : \lambda_i - \lambda_j \right. \\ &\quad \left. \leq \sqrt{[x_i - x_j]^\top E[x_i - x_j]}, 1 \leq i, j \leq 2T \right\}. \end{aligned}$$

Note that function $\Phi(E)$ is concave and the goal is to maximize $\Phi(E)$ over positive semi-definite rank-one matrices $E = ee^\top$ with trace 1. An efficiently solvable convex relaxation after relaxing the rank-one constraints is:

$$\begin{aligned} \max_{E, \lambda} \left\{ \sum_{i=1}^T \lambda_i - \sum_{i=T+1}^{2T} \lambda_i : E \succeq 0, \text{Tr}(E) = 1, \right. \\ \left. \lambda_i - \lambda_j \leq \sqrt{[x_i - x_j]^\top E[x_i - x_j]}, 1 \leq i, j \leq 2T \right\}. \end{aligned}$$

After the optimal solution E_* to the problem is found, we can use standard methods to obtain a reasonably good e , e.g., take e as the leading eigenvector of E_* .

Here we present a simple numerical illustration for the optimal project. Consider the two-dimensional Gaussian distributions with same mean value and different covariance

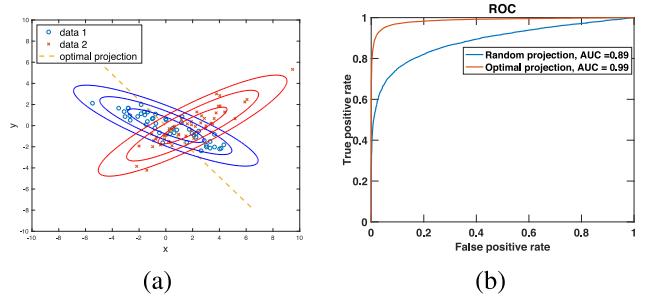


Fig. 4. Illustration of optimal projection on simulated data. (a): Optimal projection for two training sets; (b): The ROC curves for optimal projection and random projection.

structures. More specifically, let the data X_1 to be sampled i.i.d. from $\mathcal{N}(\mu_1, \Sigma_1)$ and data X_2 to be sampled from $\mathcal{N}(\mu_2, \Sigma_2)$, where $\mu_1 = [0 \ 0]^\top$, $\mu_2 = [2 \ 0]^\top$, and

$$\Sigma_1 = \begin{bmatrix} 5.03 & -2.41 \\ -2.41 & 1.55 \end{bmatrix}, \quad \Sigma_2 = \begin{bmatrix} 5.50 & 3.30 \\ 3.30 & 2.53 \end{bmatrix}.$$

Fig. 4 shows the optimal projection obtained from 50 training samples from each distribution (which can be seen to optimally “separate” the two distributions), and the ROC curve averaged over 10,000 trials that demonstrates the performance gain of the optimal projection.

VI. NUMERICAL EXAMPLES

In this section, we perform some simulations to validate the performance of the ℓ_2 test and compare with two benchmarks: (i) the classical parametric Hotelling’s T^2 test [67]; and (ii) the non-parametric maximum mean discrepancy (MMD) test [34]. More specifically, we study the test power of the two-sample test for Gaussian distributions under various dimensions. Moreover, we show the performance in change detection by studying the detection power in the offline case and the expected detection delay in the online case, respectively.

We first introduce briefly the two benchmark procedures.

a) *Hotelling’s T^2 statistic*: The Hotelling’s T^2 statistic is a classical parametric test designed utilizing the mean and covariance structures of data, and thus it can detect both the mean and covariance shifts [67]. Given two set of samples $\{x_1, \dots, x_{n_1}\}$ and $\{y_1, \dots, y_{n_2}\}$, the Hotelling’s T^2 statistic is defined as

$$t^2 = \frac{n_1 n_2}{(n_1 + n_2)} (\bar{x} - \bar{y})^\top \widehat{\Sigma}^{-1} (\bar{x} - \bar{y}), \quad (28)$$

where \bar{x} and \bar{y} are the sample mean and $\widehat{\Sigma}$ is the pooled covariance matrix estimate.

b) *MMD statistic*: The MMD test is a non-parametric benchmark for two-sample test and change detection [34], [68]. Given a class of functions \mathcal{F} and two distributions p and q , the MMD distance between p and q is defined as $\text{MMD}_{\mathcal{F}}(p, q) = \sup_{f \in \mathcal{F}} (\mathbb{E}_{x \sim p}[f(x)] - \mathbb{E}_{y \sim q}[f(y)])$. For MMD in reproducing kernel Hilbert spaces (RKHS), given samples $\{x_1, \dots, x_{n_1}\}$ and $\{y_1, \dots, y_{n_2}\}$, an unbiased estimate of squared MMD distance is given by

$$\text{MMD}_u^2 = \frac{1}{n_1(n_1 - 1)} \sum_{i=1}^{n_1} \sum_{j \neq i} k(x_i, x_j)$$

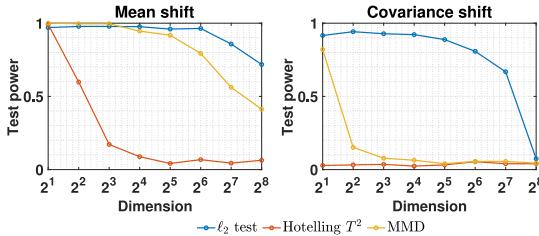


Fig. 5. Comparison of test power of the proposed test versus classic Hotelling's T^2 statistic and the MMD statistic, when performing a two-sample test on two Gaussian distributions, with significance level $\alpha = 0.05$. (Left) Gaussian distributions having the same variance and different means; (Right) Gaussian distributions having same mean and different variances.

$$+ \frac{1}{n_2(n_2 - 1)} \sum_{i=1}^{n_2} \sum_{j \neq i} k(y_i, y_j) \\ - \frac{2}{n_1 n_2} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} k(x_i, y_j), \quad (29)$$

where $k(\cdot, \cdot)$ is the kernel function associated with RKHS.

A. Two-Sample Test

Following a similar setup as in [34], we investigate the performance of various tests as a function of the dimension d of the sample space \mathbb{R}^d , when both p and q are Gaussian distributions. We consider values of d up to 256. The type-I risk for all tests is set as $\alpha = 0.05$. The sample size is chosen as $n_1 = n_2 = 100$, and results are averaged over 500 independent trials. In the first case, the distributions p, q have different means and the same variance. More specifically, $p = \mathcal{N}(0, I_d)$ and $q = \mathcal{N}(\mu \mathbf{1}/\sqrt{d}, I_d)$ with $\mu = 0.8$. Note that the division of each element of the mean vector by \sqrt{d} makes the difficulty of the hypothesis testing similar across all d values. In the second case, the distributions p, q have the same means but different variance. More specifically, $p = \mathcal{N}(0, I_d)$ and $q = \mathcal{N}(0, \Sigma)$ with $\Sigma = \text{Diag}\{0.25, 1, \dots, 1\}$, i.e., we only scale the first diagonal entry in the covariance matrix to make the hypothesis testing problem challenging to perform.

The test power for different methods is shown in Fig. 5. The test power drops when the dimension increases, which is consistent with the results in [9]. Hotelling's T^2 test performs good in low dimensions, but its performance degrades quickly when we consider higher dimensional problems. The MMD test is comparable to ℓ_2 test in low dimensions, but the ℓ_2 test tends to outperform the MMD test in high dimensions. The reason can be that by projecting to one-dimensional spaces using a good projection, the power of ℓ_2 test tends to decrease slower compared to Hotelling's T^2 and MMD tests.

B. Offline Change Detection

As an extension and application of the proposed ℓ_2 test, we investigate the performance for the *offline* change detection and compare the detection power, i.e., the probability of successfully detecting the change when there is a change.

Assume we have sample x_1, \dots, x_T with a fixed time horizon $T = 200$, when there is a change, we set the change-point $K = 100$. The ℓ_2 detection statistic at each time t is $M\chi$ with

TABLE II
DETECTION POWER IN OFFLINE CHANGE DETECTION. THE SEQUENCE OF LENGTH IS 200. THRESHOLDS FOR ALL METHODS ARE CALIBRATED SO THAT THE SIGNIFICANCE LEVEL IS $\alpha = 0.10$ AND $\alpha = 0.25$. AVERAGED OVER 500 TRIALS

	$\alpha = 0.10$				$\alpha = 0.25$			
	Case 1	Case 2	Case 3	Case 4	Case 1	Case 2	Case 3	Case 4
ℓ_2 test	0.52	0.85	0.18	0.56	0.70	0.90	0.35	0.71
MMD	0.32	0.90	0.16	0.43	0.60	0.95	0.34	0.69
Hotelling's T^2	0.07	0.23	0.09	0.06	0.20	0.23	0.23	0.23

χ defined in (2) (here $M = 2LR/(L + R)$ is the normalizing constant). To avoid the segment being too short, we compute the detection statistics for time instances $t \in [w, T - w]$ with $w = 20$, and then take the maximum. Similarly, the Hotelling's T^2 statistic at each time t is computed using (28) by treating data before t as one sample and after t as another sample; the MMD statistic is computed in a similar way from (29). We claim there is a change-point when the maximum of the detection statistics within window $t \in [w, T - w]$ exceeds the threshold. The thresholds for different methods will be chosen by Monte Carlo simulation to control the false alarm rate.

We consider the following cases (distribution changes) in the numerical experiments.

Case 1 (Discrete distributions): The support size is $n = 10$, distribution shifts from $p = 1/10$ (uniform) to $q = [1/30, 2/30, \dots, 5/30, 5/30, \dots, 2/30, 1/30]$ (non-uniform).

Case 2 (Gaussian distributions): The distribution shifts from two-dimensional standard Gaussian $\mathcal{N}(0, I_2)$ to $\mathcal{N}([0.5 \quad 0]^T, [1 \quad 0.7]^T [1 \quad 0.7] + [-1 \quad 0.4]^T [-1, \quad 0.4])$.

Case 3 (Gaussian to Gaussian mixture): The distribution shifts from $\mathcal{N}(0, I_{20})$ to the Gaussian mixture $0.8\mathcal{N}(0, I_{20}) + 0.2\mathcal{N}(0, 0.1I_{20})$.

Case 4 (Gaussian to Laplace): The distribution shifts from standard Gaussian $\mathcal{N}(0, 1)$ to Laplace distribution with zero mean and standard deviation 0.8.

The detection power is averaged over 500 repetitions and is reported in Table II. It shows that the proposed ℓ_2 test outperforms the classic Hotelling's T^2 and MMD tests, especially when the distribution change is difficult to detect (such as Case 3 and Case 4, where pre- and post-change distributions are close). For Case 2 detecting mean and covariance shifts, the MMD test performs slightly better. A possible explanation is that the MMD metric can capture the difference between pre- and post-change Gaussian distributions well in a fairly low-dimensional setting.

C. Online Change Detection

We further investigate the performance for *online* change detection and compare the average detection delay, i.e., the number of samples it takes to detect the change after the change happens. More specifically, the detection delay is the difference between the stopping time and the true change-point.

Assume we have samples $\{x_t, t = 1, 2, \dots\}$ that are available sequentially. We adopt the convention that there are pre-change samples available as $\{x_0, x_{-1}, \dots\}$, which are referred

as historical data and can be used during the detection procedure. Consider the window-limited ℓ_2 detection procedure defined in (21) with parameter $m_0 = 20$ and $m_1 = 100$. The Hotelling's T^2 detection statistic at each time t is constructed as $(\bar{x}_t - \hat{\mu})^\top \hat{\Sigma}^{-1} (\bar{x}_t - \hat{\mu})$ where \bar{x}_t is the average of samples within window $[t - m_0 + 1, t]$, and $\hat{\mu}$, $\hat{\Sigma}$ are estimated from historical data. The MMD statistic is constructed in the same way as in [68] with block size $B_0 = 20$ and number of blocks $N = 5$. We will claim change and stop the detection procedures when the detection statistic exceeds the threshold; the thresholds for different methods are chosen by Monte Carlo simulation to control the average run length.

We consider the following four cases, which are modified slightly from the offline case. We have increased the signal-to-noise ratio in certain cases to increase the detectability in the online setting.

Case 1 (Discrete distributions): The support size is $n = 10$, distribution shifts from $p = 1/10$ (uniform) to $q = [0.04, 0.14, 0.32, 0, 0, 0, 0, 0.32, 0.14, 0.04]$ (non-uniform).

Case 2 (Gaussian distributions): The distribution shifts from two-dimensional standard Gaussian $\mathcal{N}(0, I_2)$ to $\mathcal{N}([0.5 \quad 0]^\top, [1 \quad 0.7]^\top [1 \quad 0.7] + [-1 \quad 0.4]^\top [-1 \quad 0.4])$.

Case 3 (Gaussian to Gaussian mixture): The distribution shifts from $\mathcal{N}(0, I_{20})$ to the Gaussian mixture $0.4\mathcal{N}(0, I_{20}) + 0.6\mathcal{N}(0, 0.1I_{20})$.

Case 4 (Gaussian to Laplace): The distribution shifts from standard Gaussian $\mathcal{N}(0, 1)$ to Laplace distribution with zero mean and standard deviation 0.7.

The evolution paths of detection statistics for all cases are given in Fig. 6. To simulate EDD, we let the change occur at the first time instant of the testing data. The detection delay is averaged over 500 repetitions and reported in Table III.

VII. REAL-DATA STUDY: ONLINE GESTURE CHANGE DETECTION

In this section, we apply our method to the sequential gesture detection problem using a real dataset: the Microsoft Research Cambridge-12 (MSRC-12) Kinect gesture dataset [69]. This dataset consists of sequences of human skeletal body part movements (represented as body part locations) collected from 30 people performing 12 gestures. There are 18 sensors in total, and each sensor records the coordinates in the three-dimensional Cartesian coordinate system at each time. Therefore there are 54 attributes, denoted by $y_t \in \mathbb{R}^{54}$, $t = 1, 2, \dots, T$. The goal is to detect the transition of gestures from the sequences of sensor observations.

We apply the proposed online change detection procedure defined in (21) to the MSRC-12 dataset, and the detailed scheme is outlined as follows. We first preprocess the data by removing the frames that the person is standing still or with little movements. Then we select a *unit-norm* vector $u \in \mathbb{R}^{54}$ and project data into this direction to obtain a univariate sequence: $x_t = u^\top y_t$. The projection vector u is found by finding the optimal projection to maximize the Wasserstein distance described in Section V-B. Then we discretize the

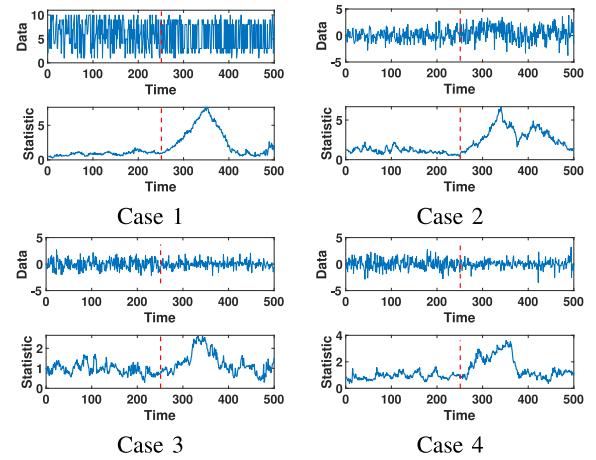


Fig. 6. Illustration of online change detection using the ℓ_2 divergence under four simulated cases explained in Section VI-C. For each case, the upper plot shows the raw data and the bottom plot shows the evolution path of the ℓ_2 detection statistic, with true change-point indicated in red dash lines.

TABLE III
COMPARISON OF EDD FOR ONLINE CHANGE DETECTION USING THE PROPOSED STATISTIC, THE MMD, AND THE HOTELLING'S T^2 STATISTIC. THE PARAMETER IS $n = 10$, $m_0 = 20$, $m_1 = 100$ AND THRESHOLDS FOR ALL METHODS ARE CALIBRATED SO THAT ARL = 500. THE DASHED LINE INDICATES THE METHOD FAILS TO DETECT THE CHANGE (I.E., THE DELAY IS LARGER THAN THE TIME HORIZON)

	Case 1	Case 2	Case 3	Case 4
ℓ_2 test	20.34	89.66	69.23	92.49
MMD	258.02	47.72	—	394.91
Hotelling's T^2	406.42	36.79	—	370.61

univariate sequence into n bins. At each time t , we construct the detection statistic $\max_{m_0 \leq t-k \leq m_1} \chi_{t,k}$ as illustrated in Fig. 2.

The parameters are set as $m_0 = 20$, $m_1 = 300$ for the detection procedure. The detection statistics are shown in Fig. 7, with the true change indicated by red dash lines. We also compare the ℓ_2 test with tests based on Hotelling's T^2 statistic, ℓ_1 distance, and KL divergence. Using the ℓ_1 approach detailed in Section III-F, we build the test statistic $\delta_{t,k}$ for time t and potential change-point $k < t$, then the detection statistic is computed by maximizing over all potential change-points: $\max_{m_0 \leq t-k \leq m_1} \delta_{t,k}$. For the KL divergence test, at time t and for $k < t$, set $N_{t,k} = \lceil t - k \rceil$ and denote the empirical distribution of two segments $[k - N_{t,k} + 1, k]$, $[k + 1 : t]$ as \hat{f}_1 and \hat{f}_2 (with zero value adjusted to a small constant $\delta = 0.01$), then compute their KL divergence as detection statistic.

From the results in Fig. 7, we can observe that our sequential detection procedure based on the ℓ_2 divergence can detect the change right after it happens. This is because the detection statistic before the change has a smaller variance, which indicates that we can set the threshold to be reasonably low for quicker detection. Moreover, there is a clear linear increasing trend after the change, enabling quick and reliable detection. In contrast, Hotelling's T^2 statistic does not have the desired online change detection behavior. The detection statistic is noisy before the change and does not have a consistent positive shift after the change; the KL test is even worse in this regard. The ℓ_1 divergence-based test has a similar behavior as the ℓ_2

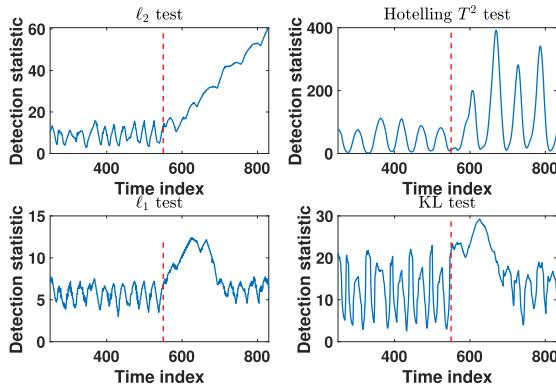


Fig. 7. Real-data example using online gesture change detection. Comparison of detection statistics (under uniform weights) for “bow” to “throw,” for the proposed procedure, the Hotelling’s T^2 test, ℓ_1 test, and the KL test. Red dash lines indicate the true change-point (hand-labeled).

divergence. However, the ℓ_1 divergence has smaller “signal-to-noise” ratio in that the variance between the change is larger, and post-change distribution drift seems to be smaller.

VIII. CONCLUSION

We have presented a new non-parametric change detection procedure based on the optimal weighted ℓ_2 divergence. We studied the optimality and various theoretical properties of the weighted ℓ_2 divergence for the offline and online change-point detection. We also studied the practical aspects, including calibration threshold using training data, optimizing weights, and finding an optimal projection for high-dimensional data. We demonstrate the good performance of the proposed method using simulated and real-data for human gesture detection.

ACKNOWLEDGMENT

The authors are grateful to Professor Arkadi Nemirovski, Professor Anatoli Juditsky, and Professor Yihong Wu for their helpful discussions regarding optimality proof. They also thank the Guest Editor and the anonymous reviewers for their helpful comments.

REFERENCES

- V. Raghavan and V. V. Veeravalli, “Quickest change detection of a Markov process across a sensor array,” *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1961–1981, Apr. 2010.
- D. Amorése, “Applying a change-point detection method on frequency-magnitude distributions,” *Bull. Seismol. Soc. America*, vol. 97, no. 5, pp. 1742–1749, 2007.
- M. Raginsky, R. M. Willett, C. Horn, J. Silva, and R. F. Marcia, “Sequential anomaly detection in the presence of noise and limited feedback,” *IEEE Trans. Inf. Theory*, vol. 58, no. 8, pp. 5544–5562, Aug. 2012.
- Y. C. Chen, T. Banerjee, A. D. Dominguez-Garcia, and V. V. Veeravalli, “Quickest line outage detection and identification,” *IEEE Trans. Power Syst.*, vol. 31, no. 1, pp. 749–758, Jan. 2016.
- D. Siegmund, B. Yakir, and N. R. Zhang, “Detecting simultaneous variant intervals in aligned sequences,” *Ann. Appl. Stat.*, vol. 5, no. 2A, pp. 645–668, 2011.
- H. V. Poor and O. Hadjiliadis, *Quickest Detection*. Cambridge, U.K.: Cambridge Univ. Press, 2008.
- M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Englewood Cliff, NJ, USA: Prentice Hall, 1993.
- A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. (Monographs on Statistics and Applied Probability), vol. 136. Boca Raton, FL, USA, London, U.K., New York, NY, USA: Chapman & Hall/CRC Press, Taylor & Francis Group, 2015.
- A. Ramdas, S. J. Reddi, B. Póczos, A. Singh, and L. Wasserman, “On the decreasing power of kernel and distance based nonparametric hypothesis tests in high dimensions,” in *Proc. 29th AAAI Conf. Artif. Intell.*, 2015, pp. 3571–3577.
- T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White, “Testing that distributions are close,” in *Proc. 41st Annu. Symp. Found. Comput. Sci.*, 2000, pp. 259–269.
- H. Chernoff, “A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations,” *Ann. Math. Stat.*, vol. 23, no. 4, pp. 493–507, 1952.
- A. Rényi, “On measures of entropy and information,” in *Proc. 4th Berkeley Symp. Math. Stat. Probab.*, 1961, pp. 547–561.
- A. Basu, I. R. Harris, N. L. Hjort, and M. Jones, “Robust and efficient estimation by minimising a density power divergence,” *Biometrika*, vol. 85, no. 3, pp. 549–559, 1998.
- A. Cichocki and S.-I. Amari, “Families of alpha-beta-and gamma-divergences: Flexible and robust measures of similarities,” *Entropy*, vol. 12, no. 6, pp. 1532–1568, 2010.
- S.-I. Amari, *Information Geometry and Its Applications*, vol. 194. Tokyo, Japan: Springer, 2016.
- R. Zhang, Y. Mei, and J. Shi, “Robust real-time monitoring of high-dimensional data streams,” 2019. [Online]. Available: arXiv:1906.02265.
- S.-O. Chan, I. Diakonikolas, P. Valiant, and G. Valiant, “Optimal algorithms for testing closeness of discrete distributions,” in *Proc. 25th Annu. ACM-SIAM Symp. Discr. Algorithms*, 2014, pp. 1193–1203.
- J. W. Mueller and T. Jaakkola, “Principal differences analysis: Interpretable characterization of differences between distributions,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 1702–1710.
- E. L. Lehmann and J. P. Romano, *Testing Statistical Hypotheses*. New York, NY, USA: Springer, 2006.
- R. J. Larsen, *Statistics in the Real World: A Book of Examples*. New York, NY, USA: Macmillan, 1976.
- H. Cramér, *Mathematical Methods of Statistics*, vol. 43. Princeton, NJ, USA: Princeton Univ. Press, 1946.
- N. V. Smirnov, “Estimate of deviation between empirical distribution functions in two independent samples,” *Bull. Moscow Univ.*, vol. 2, no. 2, pp. 3–16, 1939.
- J. Hajek, Z. Sidak, and P. Sen, *Theory of Rank Tests*. San Diego, CA, USA: Academic, 1999.
- H. B. Mann and D. R. Whitney, “On a test of whether one of two random variables is stochastically larger than the other,” *Ann. Math. Stat.*, vol. 18, no. 1, pp. 50–60, 1947.
- W. H. Kruskal and W. A. Wallis, “Use of ranks in one-criterion variance analysis,” *J. Amer. Stat. Assoc.*, vol. 47, no. 260, pp. 583–621, 1952.
- W. J. Conover, *Practical Nonparametric Statistics*, vol. 350. New York, NY, USA: Wiley, 1998.
- F. Scholz and A. Zhu. (2019). *kSamples: K-Sample Rank Tests and Their Combinations*. [Online]. Available: <https://CRAN.R-project.org/package=kSamples>
- M. F. Schilling, “Multivariate two-sample tests based on nearest neighbors,” *J. Amer. Stat. Assoc.*, vol. 81, no. 395, pp. 799–806, 1986.
- P. J. Bickel, “A distribution free version of the Smirnov two sample test in the p-variate case,” *Ann. Math. Stat.*, vol. 40, no. 1, pp. 1–23, 1969.
- J. H. Friedman and L. C. Rafsky, “Multivariate generalizations of the Wald-Wolfowitz and Smirnov two-sample tests,” *Ann. Stat.*, vol. 7, no. 4, pp. 697–717, 1979.
- P. J. Bickel and L. Breiman, “Sums of functions of nearest neighbor distances, moment bounds, limit theorems and a goodness of fit test,” *Ann. Probab.*, vol. 11, no. 1, pp. 185–214, 1983.
- N. Henze, “A multivariate two-sample test based on the number of nearest neighbor type coincidences,” *Ann. Stat.*, vol. 16, no. 2, pp. 772–783, 1988.
- G. J. Székely and M. L. Rizzo, “Testing for equal distributions in high dimension,” *InterStat*, vol. 5, no. 16, pp. 1249–1272, 2004.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, “A kernel two-sample test,” *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 723–773, 2012.
- A. Ramdas, N. G. Trillos, and M. Cuturi, “On Wasserstein two-sample testing and related families of nonparametric tests,” *Entropy*, vol. 19, no. 2, p. 47, 2017.

- [36] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White, "Testing closeness of discrete distributions," *J. ACM*, vol. 60, no. 1, pp. 1–25, 2013.
- [37] P. Valiant, "Testing symmetric properties of distributions," *SIAM J. Comput.*, vol. 40, no. 6, pp. 1927–1968, 2011.
- [38] R. Rubinfeld, "Taming big probability distributions," *XRDS Crossroads ACM Mag. Students*, vol. 19, no. 1, pp. 24–28, 2012.
- [39] C. L. Canonne, "A survey on distribution testing: Your data is big. But is it blue?" *Theory Comput.*, vol. 9, pp. 1–100, Aug. 2020.
- [40] O. Goldreich and D. Ron, "On testing expansion in bounded-degree graphs," in *Studies in Complexity and Cryptography. Miscellanea on the Interplay Between Randomness and Computation*. Heidelberg, Germany: Springer, 2011, pp. 68–75.
- [41] J. Acharya, A. Jafarpour, A. Orlitsky, and A. T. Suresh, "Sublinear algorithms for outlier detection and generalized closeness testing," in *Proc. Int. Symp. Inf. Theory*, 2014, pp. 3200–3204.
- [42] B. Bhattacharya and G. Valiant, "Testing closeness with unequal sized samples," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2611–2619.
- [43] I. Diakonikolas, D. M. Kane, and V. Nikishkin, "Near-optimal closeness testing of discrete histogram distributions," 2017. [Online]. Available: [arXiv:1703.01913](https://arxiv.org/abs/1703.01913).
- [44] J. Acharya, Z. Sun, and H. Zhang, "Differentially private testing of identity and closeness of discrete distributions," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 6878–6891.
- [45] E. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, nos. 1–2, pp. 100–115, 1954.
- [46] E. Page, "A test for a change in a parameter occurring at an unknown point," *Biometrika*, vol. 42, nos. 3–4, pp. 523–527, 1955.
- [47] A. N. Shiryaev, "On optimum methods in quickest detection problems," *Theory Probab. Appl.*, vol. 8, no. 1, pp. 22–46, 1963.
- [48] G. Lorden, "Procedures for reacting to a change in distribution," *Ann. Math. Stat.*, vol. 42, no. 6, pp. 1897–1908, 1971.
- [49] T. L. Lai, "Sequential analysis: Some classical problems and new challenges," *Statistica Sinica*, vol. 11, no. 2, pp. 303–350, 2001.
- [50] G. Boracchi, D. Carrera, C. Cervellera, and D. Maccio, "QuantTree: Histograms for change detection in multivariate data streams," in *Proc. Int. Conf. Mach. Learn.*, vol. 80, 2018, pp. 638–647.
- [51] H. Chen, "Sequential change-point detection based on nearest neighbors," *Ann. Stat.*, vol. 47, no. 3, pp. 1381–1407, 2019.
- [52] T. S. Lau, W. P. Tay, and V. V. Veeravalli, "A binning approach to quickest change detection with unknown post-change distribution," *IEEE Trans. Signal Process.*, vol. 67, no. 3, pp. 609–621, Feb. 2019.
- [53] G. Valiant and P. Valiant, "An automatic inequality prover and instance optimal identity testing," *SIAM J. Comput.*, vol. 46, no. 1, pp. 429–455, 2017.
- [54] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Trans. Inf. Theory*, vol. 44, no. 7, pp. 2917–2929, Nov. 1998.
- [55] D. Siegmund and E. Venkatraman, "Using the generalized likelihood ratio statistic for sequential detection of a change-point," *Ann. Stat.*, vol. 23, no. 1, pp. 255–271, 1995.
- [56] D. Siegmund and B. Yakir, "Detecting the emergence of a signal in a noisy image," *Stat. Interface*, vol. 1, no. 1, pp. 3–12, 2008.
- [57] B. Yakir, "Multi-channel change-point detection statistic with applications in DNA copy-number variation and sequential monitoring," in *Proc. 2nd Int. Workshop Sequential Methodol.*, 2009, pp. 15–17.
- [58] B. Yakir, *Extremes in Random Fields: A Theory and Its Applications*. Chichester, U.K.: Wiley, 2013.
- [59] Y. Xie and D. Siegmund, "Sequential multi-sensor change-point detection," *Ann. Stat.*, vol. 41, no. 2, pp. 670–692, 2013.
- [60] D. Siegmund, B. Yakir, and N. Zhang, "Tail approximations for maxima of random fields by likelihood ratio transformations," *Sequential Anal.*, vol. 29, no. 3, pp. 245–262, 2010.
- [61] D. Siegmund and B. Yakir, *The Statistics of Gene Mapping*. New York, NY, USA: Springer, 2007.
- [62] M. Pollak, "Optimal detection of a change in distribution," *Ann. Stat.*, vol. 13, no. 1, pp. 206–227, Mar. 1985.
- [63] L. Pelkowitz and S. C. Schwarts, "Asymptotically optimum sample size for quickest detection," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-23, no. 2, pp. 263–272, Mar. 1987.
- [64] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*. New York, NY, USA: Springer, 1985.
- [65] A. Jüngel, *Entropy Methods for Diffusive Partial Differential Equations*. Cham, Switzerland: Springer, 2016.
- [66] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [67] H. Hotelling, "Multivariate quality control," in *Techniques of Statistical Analysis*. New York, NY, USA: McGraw-Hill, 1947.
- [68] S. Li, Y. Xie, H. Dai, and L. Song, "M-statistic for kernel change-point detection," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 3366–3374.
- [69] S. Fothergill, H. Mentis, P. Kohli, and S. Nowozin, "Instructing people for training gestural interactive systems," in *Proc. SIGCHI Conf. Human Factors Comput. Syst.*, 2012, pp. 1737–1746.
- [70] B. Yakir and M. Pollak, "A new representation for a renewal-theoretic constant appearing in asymptotic approximations of large deviations," *Ann. Appl. Probab.*, vol. 8, no. 3, pp. 749–774, 1998.
- [71] D. Siegmund and B. Yakir, "Tail probabilities for the null distribution of scanning statistics," *Bernoulli*, vol. 6, no. 2, pp. 191–213, 2000.
- [72] Y. Cao, Y. Xie, and N. Gebraeel, "Multi-sensor slope change detection," *Ann. Oper. Res.*, vol. 263, nos. 1–2, pp. 163–189, 2018.



Liyan Xie (Graduate Student Member, IEEE) received the B.Sc. degree in statistics from the University of Science and Technology, China, in 2016. She is currently pursuing the Ph.D. degree with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology. Her research interests are the intersection between statistics and optimization, with a primary focus on sequential change detection and their applications in sensor networks and healthcare. She is also a recipient of the 2020 ARC Student Fellowship and the IDEaS-TRIAD Research Scholarships from the Transdisciplinary Research Institute for Advancing Data Science, Georgia Tech. She was awarded the finalist of the INFORMS QSR Student Paper Competition and Runner-up for the INFORMS Computing Society Student Paper Prize in 2019.



Yao Xie (Member, IEEE) received the B.Sc. degree in electrical engineering and computer science from the University of Science and Technology of China, the M.Sc. degree in electrical and computer engineering from the University of Florida, and the Ph.D. degree in electrical engineering (minor in Mathematics) from Stanford University. She is an Associate Professor and Harold R. and Mary Anne Nash Early Career Professor with the Georgia Institute of Technology, H. Milton Stewart School of Industrial and Systems Engineering, and an Associate Director of the Machine Learning Center. She was a Research Scientist with Duke University. Her research areas are statistics (particularly change-point detection and spatio-temporal data modeling), machine learning, and signal processing, providing the theoretical foundation and developing computationally efficient and statistically powerful algorithms for engineering problems. She received the National Science Foundation CAREER Award in 2017. She is currently an Associate Editor for *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, *Sequential Analysis: Design Methods and Applications*, and *INFORMS Journal on Data Science*, and serves on the editorial board of *Journal of Machine Learning Research*.