

Unsupervised Change Detection using DRE-CUSUM

Sudarshan Adiga Ravi Tandon
Department of Electrical and Computer Engineering
University of Arizona, Tucson, AZ, USA
E-mail: {adiga, tandonr}@email.arizona.edu

Abstract—This paper presents DRE-CUSUM, an unsupervised density-ratio estimation (DRE) based approach to determine statistical changes in time-series data when no knowledge of the pre-and post-change distributions are available. The core idea behind the proposed approach is to split the time-series at an arbitrary point and estimate the ratio of densities of distribution (using a parametric model such as a neural network) before and after the split point. The DRE-CUSUM change detection statistic is then derived from the cumulative sum (CUSUM) of the logarithm of the estimated density ratio. We present a theoretical justification as well as accuracy guarantees which show that the proposed statistic can reliably detect statistical changes, irrespective of the split point. While there have been prior works on using density ratio based methods for change detection, to the best of our knowledge, this is the first unsupervised change detection approach with a theoretical justification and accuracy guarantees. The simplicity of the proposed framework makes it readily applicable in various practical settings (including high-dimensional time-series data); we also discuss generalizations for online change detection. We experimentally show the superiority of DRE-CUSUM using both synthetic and real-world datasets over existing state-of-the-art unsupervised algorithms (such as Bayesian online change detection, its variants as well as several other heuristic methods).

I. INTRODUCTION

Change detection is the process of identifying deviations in the statistical behavior of time series data, and finds numerous applications, such as detection of distributed denial of service (DDoS) attacks [4], real-time surveillance [15], video segmentation [5], [13], event prediction [31], [41] and healthcare monitoring [25], [36]. To describe the canonical problem of change detection, let us consider a time-series data, denoted by $X_{[1:n]} \triangleq (x_1, x_2, \dots, x_n)$ with a single change point at some unknown time T^* . Elements of the sub-sequence $X_{[1:T^*-1]}$ are i.i.d. and sampled from a distribution P_1 , whereas the elements of sub-sequence $X_{[T^*:n]}$ are sampled from a distribution P_2 . The goal of offline change detection is to efficiently determine T^* [3].

When the pre- and post- change distributions P_1 , and P_2 are known, one can obtain the maximum-likelihood (ML) estimate for the change point using cumulative-sum (CUSUM) of log-likelihood ratios based statistic [28], [39] (denoted as $S_k = \sum_{t=0}^k \log(P_2(x_t)/P_1(x_t))$). The main intuition behind

CUSUM statistic stems from the expected values of the log-likelihood ratio $P_2(\cdot)/P_1(\cdot)$, before and after T^* , which is,

$$\mathbb{E}_{x_t} \log \left(\frac{P_2(x_t)}{P_1(x_t)} \right) = \begin{cases} -KL(P_1||P_2), & t < T^* \\ KL(P_2||P_1), & t \geq T^* \end{cases} \quad (1)$$

Since Kullback-Leibler (KL) divergence is non-negative, the CUSUM statistic has a negative expected slope for any $t < T^*$, and conversely, positive expected slope for $t \geq T^*$. However, the limitation of the ML-and CUSUM approaches is that they can be applied only when $P_2(x)/P_1(x)$ can be accurately computed for any x . Moreover, in several real-world applications the distributions before and after the change point (denoted by P_1, P_2 , respectively) are unknown [12], and hence these approaches are impracticable.

Main Contributions: In this work, we focus on the challenging setting for change detection when we have no knowledge about pre-and post-change distributions. We do not make any assumptions on the underlying probability distributions, i.e., we consider a non-parametric setting. The core idea of our proposed methodology is as follows: suppose we observe a time series $X_{[1:n]}$ with an unknown change point at T^* . We split the time-series data at an arbitrarily chosen time T_{split} (say $n/2$) to obtain two sub-sequences as $X_{[1:T_{\text{split}}-1]} \sim P_{\text{left}}$, and $X_{[T_{\text{split}}:n]} \sim P_{\text{right}}$. We then propose DRE-CUSUM, an unsupervised change detection statistic which *mimics* the conventional CUSUM statistic, with the difference that $P_2(x)/P_1(x)$ is replaced by the estimate of the density ratio $P_{\text{left}}(x)/P_{\text{right}}(x)$. In Proposition 1, we prove the surprising result that the DRE-CUSUM statistic possesses theoretical properties analogous to the conventional CUSUM statistic, by showing that

$$\mathbb{E}_{x_t} \log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) = \begin{cases} > 0, & \text{for } t < T^* \\ < 0, & \text{for } t \geq T^* \end{cases} \quad (2)$$

The highlight of (2) is the fact that it always holds true irrespective of the choice of T_{split} . In addition, we also prove accuracy guarantees for DRE-CUSUM by determining the bounds on the probability of error of the estimated change point given that the estimator can correctly compute the density ratio with high probability.

Furthermore, the theoretical results supporting the use of DRE-CUSUM statistic for unsupervised change detection do not make any assumptions on the density ratio estimators. Therefore, in practice, one can leverage and choose from

This work has been supported in part by NSF Grants CAREER 1651492, CNS 1715947, CCF 2100013 and the 2018 Keysight Early Career Professor Award.

a wide variety of density ratio estimation techniques [18], [33] to estimate $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$. This allows a quite general and efficient framework for unsupervised change detection applicable for high-dimensional data.

We then present generalizations of the DRE-CUSUM approach for detecting multiple changes as well as for online-change detection. We also discuss the possible failure modes of the proposed approach and some possible methods to overcome them. Lastly, we present a comprehensive set of experimental results¹ to demonstrate the superiority in performance of DRE-CUSUM over other change detection methods using both synthetic and real-world datasets.

Related work: Several existing algorithms such as sequential probability ratio test (SPRT), generalized likelihood ratio test (GLRT), CUSUM and its variants such as weighted CUSUM are based on the assumption that the density ratios can be readily computed for devising test-statistics for change detection [3]. To overcome this, one can estimate the distributions using Kernel density estimation (KDE) and histogram-based approach [17], [29]. Density estimation techniques for change detection has several shortcomings. Firstly, density estimation based approaches are limited to applications with low-dimensional observations [22]. Secondly, the performance of change detection using density estimation techniques could deteriorate when division by the estimated density magnifies the estimation error of the log-likelihood ratio [32]. Lastly, the samples from pre- and post-change distributions must be labelled.

An alternate technique to determine the change points is by directly estimating the density ratios [18], [26], [33], [37], an overview which has been presented in Appendix C. An unsupervised change detection algorithm using a DRE framework has been proposed in [14]. However, no theoretical justification and accuracy guarantees have been provided for this approach. Clustering, and Bayesian change detection (BCD) are some of the other techniques that do not require explicit labelling of pre- and post- change samples [1], [9], [43]. However, such approaches often lack a statistical justification and are dependent on ordering of data as well as the initial parameters [24]. Furthermore, BCD based methods [1], [9] make assumptions on the underlying data generating distributions and on the change points (also referred to as data, and change-point priors, respectively). Other limitations of the BCD based methods are that: (a) these approaches are suitable for low-dimensional data and are difficult to implement for high-dimensional data, and (b) are prone to high false alarm rates.

There are several other heuristic approaches for unsupervised change detection. For instance, dynamic programming techniques can be used for change detection by performing an ordered search to determine the best possible set of partitions, given the total number of change points apriori [34]. Linearized penalty segmentation (Pelt) overcomes the need for the number of change points by minimizing a penalty term that

is itself a function of the estimated change points [19], [40]. These methods however, have a higher computational cost. Binary segmentation (BinSeg) finds an approximate solution by first finding the best instance (first change point estimate) to split the time-series into two sub-sequences [10]. This operation is then executed in the derived sub-sequences in an iterative manner to determine all possible change points. The main drawback of this approach is that errors in the initial stages of the algorithm can propagate for subsequent steps. An adaptive search method to find the change point by maximizing a CUSUM based statistic when the pre-and post-change distributions are from Gaussian distribution has been proposed in [21].

In Section III, we compare the performance of the proposed DRE-CUSUM technique with its unsupervised change detection counterparts such as BCD and its variants, as well as several other heuristic approaches as discussed above.

II. UNSUPERVISED CHANGE DETECTION

Problem formulation- Consider the time-series data $X_{[1:n]}$ which undergoes a change at an unknown time T^* . The samples $X_{[1:T^*-1]}$ are i.i.d., drawn from a distribution P_1 , and samples $X_{[T^*:n]}$ are i.i.d., drawn from a distribution P_2 . The goal of unsupervised offline change detection is to estimate T^* , when P_1, P_2 are unknown. We discuss the online setting later in this section.

When pre-and post-change distributions are known, one can obtain the change point estimate (\hat{T}_{ML}) using maximum likelihood (ML) [3]:

$$\hat{T}_{ML} = \arg \max_t \sum_{i=t}^n \log \left(\frac{P_2(x_i)}{P_1(x_i)} \right) \quad (3)$$


However, the ML approach can be applied if either the distributions P_1 and P_2 are known, or the density ratio P_2/P_1 can be accurately computed. The need for the information on the distributions and their corresponding order in the time series makes the ML approach infeasible for most change detection applications.

To alleviate the disadvantages of the ML approach, we consider a setting in which we do not know the pre-and post-change distributions. Given the time-series $X_{[1:n]}$, we split the time-series at a point T_{split} , to obtain two sub-sequences: $X_{[1:T_{\text{split}}-1]} \sim P_{\text{left}}$, $X_{[T_{\text{split}}:n]} \sim P_{\text{right}}$, where P_{left} and P_{right} denote the corresponding distributions of the samples (see Fig. 1a). Based on the relative position of T_{split} with respect to T^* , either P_{left} or P_{right} is a mixture distribution, and conversely, the other is a pure distribution (either P_1 or P_2). We next define the density-ratio (DR) based cumulative-sum (CUSUM) of likelihood ratio based statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$, $\forall t \in [1, n]$.


$$S_{\text{DR}}^{T_{\text{split}}}(t) \triangleq \sum_{j=1}^t \log \left(\frac{P_{\text{left}}(x_j)}{P_{\text{right}}(x_j)} \right). \quad (4)$$

Fig. 1b depicts $S_{\text{DR}}^{T_{\text{split}}}(t)$, $\forall t \in [1, n]$ for different values of T_{split} (i.e. both $T_{\text{split}} \geq T^*$ and $T_{\text{split}} < T^*$) for a 10-dimensional multivariate Gaussian time-series undergoing a mean change

¹Source code for experiments available at: <https://www.dropbox.com/sh/xkchf2iajge57jq/AACztXuP-W16VSkUuoTcNY0ya?dl=0>



(a) Time-series data with a single change point at T^* , when split at $T_{\text{split}} \geq T^*$ yields two distributions: $P_{\text{left}} = \alpha^{T_{\text{split}}} P_1 + (1 - \alpha^{T_{\text{split}}}) P_2$, and $P_{\text{right}} = P_2$. Alternatively, when $T_{\text{split}} \leq T^*$, we have $P_{\text{left}} = P_1$, and $P_{\text{right}} = \alpha^{T_{\text{split}}} P_1 + (1 - \alpha^{T_{\text{split}}}) P_2$.



(b) Plot of density-ratio based CUSUM statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$ vs t for 10-dimensional time-series with 500 samples, and an unknown change point $T^* = 150$, such that $X_{[1:149]} \sim P_1 = \mathcal{N}(\bar{\mu}_1, I)$, and $X_{[150:500]} \sim P_2 = \mathcal{N}(\bar{\mu}_2, I)$, where entries of the mean vectors $\bar{\mu}_1, \bar{\mu}_2$ are sampled from $\text{Unif.}[0.0, 0.4]$, $\text{Unif.}[0.6, 1.0]$, respectively.

Fig. 1: Unsupervised change detection statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$ for different T_{split} values. From Fig. 1b, we observe that the slope of $S_{\text{DR}}^{T_{\text{split}}}(t)$ changes at T^* irrespective of the value of T_{split} .

at $T^* = 150$. As seen from this example, the change point T^* manifests itself in $S_{\text{DR}}^{T_{\text{split}}}(t)$ through a slope change at T^* , *irrespective* of the choice of T_{split} . Furthermore, we note that $S_{\text{DR}}^{T_{\text{split}}}(t)$ for $T_{\text{split}} = T^*$ corresponds to the maximum likelihood-estimate in (3). We can make the following observations about the nature of $S_{\text{DR}}^{T_{\text{split}}}(t)$: (a) the slope before the change point T^* is always positive, (b) the slope after T^* is always negative. This intuition is formalized in our first result which is stated in Proposition 1. Subsequently, we use Proposition 1 to show that the expected value of the DR-CUSUM statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$ takes its maximum value at the true change point T^* . In order to present our theoretical results, we define a mixture distribution as well as a non-negative function composed of weighted KL divergence.

Definition 1. Given two distributions P_1, P_2 , we define a parametric mixture distribution (for any $0 \leq \lambda \leq 1$) as

$$P(\lambda) = \lambda P_1 + (1 - \lambda) P_2 \quad (5)$$

Furthermore, for any $\lambda \in (0, 1), \gamma \in [0, 1]$, we define the

following non-negative function $f_i(\gamma, \lambda)$ for $i = \{1, 2\}$ as follows:

$$f_i(\gamma, \lambda) \triangleq \frac{1}{\gamma} KL(P(\lambda) \| P_i) + \frac{1 - \gamma}{\gamma} KL(P_i \| P(\lambda))$$

Our first result is stated in the following Proposition. Proof of Proposition 1 is presented in Appendix A.

Proposition 1. If $T_{\text{split}} \leq T^*$, then

$$\mathbb{E}_{x_t} \left[\log \frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right] = \begin{cases} KL(P_1 \| P(1 - \alpha_1)), & t \in [1, T^*] \\ -f_1(\alpha_1, 1 - \alpha_1), & t \in [T^*, n] \end{cases}$$

If $T_{\text{split}} \geq T^*$, then

$$\mathbb{E}_{x_t} \left[\log \frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right] = \begin{cases} f_2(\alpha_2, \alpha_2), & t \in [1, T^*] \\ -KL(P_2 \| P(\alpha_2)), & t \in [T^*, n] \end{cases}$$

where, $\alpha_1 = \frac{n - T^*}{n - T_{\text{split}}}$ and $\alpha_2 = \frac{T^*}{T_{\text{split}}}$.

We next present Corollary 1, which is a direct consequence of the result in Proposition 1.

Corollary 1. The DR-CUSUM statistic satisfies

$$T^* = \arg \max_t \mathbb{E} \left[S_{\text{DR}}^{T_{\text{split}}}(t) \right] \quad \forall T_{\text{split}} \quad (6)$$

The key significance of the above result is that the expected value of the DR-CUSUM statistic takes its maximum value at the unknown change point T^* , irrespective of the split point T_{split} . This result provides the motivation for the density ratio based estimator, denoted by $\hat{T}_{\text{DR-CUSUM}}$ for unsupervised change detection :

$$\hat{T}_{\text{DR-CUSUM}} = \arg \max_t S_{\text{DR}}^{T_{\text{split}}}(t) \quad (7)$$

In order to quantify the performance of the change point estimate $\hat{T}_{\text{DR-CUSUM}}$, we adopt the notion of (α, β) -accuracy as introduced in [6], and defined next.

Definition 2. A change point estimate \hat{T} is (α, β) -accurate for a change point T^* if for $\alpha, \beta \in [0, 1]$ -

$$P[|\hat{T} - T^*| < \alpha] \geq 1 - \beta. \quad (8)$$

We next present our second main result, which provides (α, β) -accuracy guarantees for the density ratio based estimator $\hat{T}_{\text{DR-CUSUM}}$, under the assumption that the log-likelihood ratio $\log(P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot))$ is bounded by a constant A .

Theorem 1. For any $0 < \beta < 1$, $\hat{T}_{\text{DR-CUSUM}}$ satisfies (α, β) -accuracy for

$$\alpha = \frac{2A^2}{C^2} \log \left(\frac{32}{3\beta} \right), \quad (9)$$

where the constant C represents the minimum value of the expected log-likelihood ratio $\frac{P_{\text{left}}(x)}{P_{\text{right}}(x)}$ before and after the

change point T^* , and is given as follows

$$C = \begin{cases} \min(KL(P_1||P(1-\alpha_1)), f_1(\alpha_1, 1-\alpha_1)), T_{split} \leq T^* \\ \min(f_2(\alpha_2, \alpha_2), KL(P_2||P(\alpha_2))), T_{split} \geq T^*, \end{cases}$$

and $\alpha_1 = \frac{n-T^*}{n-T_{split}}$ and $\alpha_2 = \frac{T^*}{T_{split}}$.

Proof of Theorem 1 is presented in Appendix B. The main idea behind the proof of Theorem 1 follows from considering the upper bound of the probability with which DR-CUSUM statistic assumes its maximum value at time instances at least α away from the true change point T^* . This can then be expressed in terms of the probability with which sum of the log-likelihood ratios of data points in the time-series between T^* and all other time instances at distance α from T^* is greater than 0. By subtracting the expected value of the log-likelihood ratio (i.e., KL divergence), this simplifies to the sum of zero mean i.i.d random variables to which we apply Ottaviani's inequality [35] to obtain the (α, β) accuracy guarantees for the density ratio based estimator $\hat{T}_{DR-CUSUM}$.

Remark 1. *Dependence on KL divergence:* For fixed A, β , we note that α is inversely proportional to C^2 which is related to the KL-divergence between pure and mixture distributions (subject to T_{split}), and this implies that the larger the change in distributions before and after T^* , higher is the accuracy of $\hat{T}_{DR-CUSUM}$.

A. The DRE-CUSUM Estimator

We now present the DRE-CUSUM estimator $\hat{T}_{DRE-CUSUM}$. Specifically, we split the time series at T_{split} and compute the statistic $S_{DRE}^{T_{split}}$ as follows:

$$S_{DRE}^{T_{split}}(t) = \sum_{i=1}^t \log(\hat{w}(x)), \quad (10)$$

where $\hat{w}(x)$ is an estimate of the density ratio $P_{left}(x)/P_{right}(x)$ which is obtained by density ratio estimation (DRE) models using samples from distributions P_{left} and P_{right} . The DRE-CUSUM estimator $\hat{T}_{DRE-CUSUM}$ is then obtained as follows:

$$\hat{T}_{DRE-CUSUM} = \arg \max_t S_{DRE}^{T_{split}}(t) \quad (11)$$

We label the algorithm using the DRE for unsupervised change detection as DRE-CUSUM, whose steps are described in Algorithm 1. Assuming that the DRE models correctly estimate the density ratio with high probability [30], [38], then the Proposition 1 and Theorem 1 extend to $S_{DRE}^{T_{split}}(t)$ and $\hat{T}_{DRE-CUSUM}$, respectively.

Remark 2. *Role and Impact of DRE estimation models:* we highlight that the proposed approach can in principle rely upon any of the existing methods (such as [18], [33]) for density ratio estimation given the samples from P_{left} and P_{right} . We study the impact of the choice of DRE estimation on the change detection performance in Section III.

Algorithm 1 Unsupervised Single Change Point Detection using DRE-CUSUM.

Input time-series data: $(x_1, x_2, \dots, x_{T^*}, \dots, x_n)$

1. Density Ratio Estimator (DRE) Training

Divide the time-series data at T_{split} (say $T_{split} = \frac{n}{2}$) to obtain-

(i) $X_{[1:T_{split}-1]} \sim P_{left}$, (ii) $X_{[T_{split}:n]} \sim P_{right}$.

for number of epochs **do**

- a. Sample N_1, N_2 samples from P_{left}, P_{right} , respectively.
- b. Train DRE to determine $\hat{w}(x)$, an estimate of the density ratio $P_{left}(x)/P_{right}(x)$. (see Appendix C.)

end for

2. DRE-CUSUM based Change Detection

- a. Compute $S_{DRE}^{T_{split}}(t) = \sum_{j=1}^t \log(\hat{w}(x_j))$
- b. List the time instance \hat{T} (estimated change point) at which there is a change in slope.

3. Verification Step

Repeat steps 1,2 setting $T'_{split} = \hat{T}$ (but not equal to $\frac{n}{2}$), and find $\hat{T}_{DRE-CUSUM} = \arg \max_t S_{DRE}^{T'_{split}}(t)$. Verify that $\hat{T} = \hat{T}_{DRE-CUSUM}$ is the only slope change in $S_{DRE}^{T'_{split}}(t)$.


Remark 3. *Justification of the verification Step in Algorithm 1:* When implementing DRE-CUSUM on real-world data, it is possible to observe two slope changes: one at T_{split} (the initial split point) and another at the potential change point. Therefore, we first list all possible slope changes in $S_{DRE}^{T_{split}}(t)$ as mentioned in step 2 of Algorithm 1. In order to correctly declare the change point, we propose a verification step: repeat steps 1,2 in Algorithm 1 by setting T'_{split} equal to the all instances in this list; and find $\hat{T}_{DRE-CUSUM} = \arg \max_t S_{DRE}^{T'_{split}}(t)$. If $\hat{T}_{DRE-CUSUM} = T^*$, we expect to observe only one slope change in $S_{DRE}^{T'_{split}}(t)$ at T^* , and thereby allowing us to rule out T_{split} as a change point.

Remark 4. *Robustness of DRE-CUSUM to T_{split} and the choice of T_{split} :* a natural question to ask is the following: how sensitive is the DRE-CUSUM estimator to the choice of T_{split} ? We present experimental results in Section III which show the impact of choice of T_{split} and the distance $|T_{split} - T^*|$ between the split point and the unknown change point. The second question to ask is: what is the natural choice of T_{split} ? For sufficiently large number of samples from P_{left} and P_{right} , [37] shows that $\hat{w}(x)$ converges in probability to the true density ratio $P_{left}(x)/P_{right}(x)$. This provides an intuitive explanation behind the choice of $T_{split} = \frac{n}{2}$ in Algorithm 1.


B. Generalizations of DRE-CUSUM

In this section, we discuss generalizations of the DRE-CUSUM algorithm for detecting multiple changes, adaptation for online change detection, as well as some approaches for overcoming potential failure modes.

(a) Multiple change detection



(a) Multiple change point time-series data, where $X_{[T_{j-1}^*, T_j^*]} \sim P_j$.



(b) $S_{DR}^{T_{split}}(t)$ vs t for 10-dimensional time-series of length 600 with two change points $T_1^* = 150$, $T_2^* = 450$. $X_{[1:149]}$, $X_{[150:449]}$ and $X_{[450:599]}$ follow multivariate gaussian distributions with mean vectors are sampled from $\text{Unif.}[0, 0.4]$, $\text{Unif.}[0.6, 1.0]$, and $\text{Unif.}[1.6, 2.0]$, respectively, and identity covariance matrix.

Fig. 2: Unsupervised multiple change detection statistic $S_{DR}^{T_{split}}(t)$ for different T_{split} values.

Consider the time series $X_{[1:n]}$ with $K \geq 1$ change points, denoted as $T_1^* \leq T_2^* \leq \dots \leq T_K^*$. The sub-sequence $X_{[T_{j-1}^*, T_j^*]}$ in the j th segment is i.i.d. with samples drawn from an unknown distribution P_j for $j = 1, 2, \dots, K$ (see Fig. 2a). We show that a similar approach of splitting the time-series followed by computing the DRE-CUSUM statistic can be leveraged for detecting more than one change points. To provide the intuition behind this, consider any split point T_{split} , and as before, suppose that we can compute the ratio $P_{left}(x)/P_{right}(x)$. Analogous to Proposition 1, it can be readily shown that for every $t \in [T_{j-1}^*, T_j^*]$, the expected value of the $\log(\cdot)$ of the density ratio is given as:

$$\mathbb{E}_{x_t} \left[\log \frac{P_{left}(x_t)}{P_{right}(x_t)} \right] = \underbrace{KL(P_j || P_{right}) - KL(P_j || P_{left})}_{=\Delta_j}$$

As discussed in the previous section, the slope of the DRE-CUSUM statistic will be proportional to the quantity Δ_j . Thus, as long as $\Delta_j \neq \Delta_{j-1}$ and $\Delta_j \neq \Delta_{j+1}$ for all $j = 1, 2, \dots, K$, we can expect distinct slopes in the DRE-CUSUM statistic for each segment in the time-series. In Fig. 2b, we show this behaviour for a synthetic 10-dim multivariate Gaussian time-series with two change points. The instances of the slope change are potential candidates for the estimated change points. In Section III, we also provide a comprehensive set of experiments on the use of DRE-CUSUM for detecting multiple change points for a variety of real-world datasets.






Fig. 3: Online adaptation of DRE-CUSUM algorithm. The window size can be either of a fixed size or it can be selected in an adaptive manner based on changes detected in the past.



(a) Failure mode: DRE-CUSUM in Algorithm 1 fails to detect the changes T_1^* , T_2^* when $P_{left} \approx P_{right}$.



(b) Proposed approach: Compute DRE-CUSUM statistic for multiple T_{split} values followed by a combined decision (e.g., majority vote).

Fig. 4: Overcoming failure mode of DRE-CUSUM.

(b) Online change point detection

DRE-CUSUM can be readily applied for online change detection by recursively performing Steps 1-3 in Algorithm 1 on real-time data. As shown in Fig. 3, a simple approach is to consider a window of length L (with L most recent samples collected). Steps 1-3 in Algorithm 1 can be performed on this window of L samples to determine all change points within this time interval. We slide this window across the time series to consider new observations. A generalization of this approach is to use *adaptive* window sizes depending on past detected changes. Specifically, if we have reliably detected changes in the previous window, then one only needs to keep the most recent samples from the past after the latest detected change point. In Section III, we provide experimental results for DRE-CUSUM for detecting multiple changes in both offline as well as online setting.

(c) Reducing Errors in DRE-CUSUM

We now discuss how to overcome one of the most common failure modes of the DRE-CUSUM approach using an example as shown in Fig. 4a, in which $X_{[1:T_{split}-1]} \sim P_{left}$, and $X_{[T_{split}:n]} \sim P_{right}$. If for a T_{split} , it happens that $P_{left}(x) \approx P_{right}(x)$, $\forall x$, then as a consequence, the KL divergence $KL(P_{left} || P_{right}) \approx 0$. In such a scenario, the DRE-CUSUM statistic $S_{DRE}(t)$ can fail to exhibit a slope change at the unknown change points. To alleviate this phenomenon, we

propose a simple and efficient modification of the Algorithm 1 by considering multiple distinct T_{split} as shown in the Fig. 4b, i.e., we run the DRE-CUSUM algorithm for multiple distinct split points (say $T_{\text{split}}^{(1)}, T_{\text{split}}^{(2)}, \dots, T_{\text{split}}^{(r)}$). The change points in the time-series can then be determined by applying a combined decision across the slope changes exhibited by the multiple DRE-CUSUM statistic(s). Some examples of the combined decision techniques that can be applied here are: (i) majority voting, (ii) weighted sum technique, wherein the weight corresponds to the probability that the slope change at a time instance corresponds to the true change point and is determined by the extent of the slope change. Furthermore, by using multiple values of T_{split} , we enhance the change detection framework in Algorithm 1 through reduction in the detection errors (i.e. false alarms and mis-detections). Another refinement to Algorithm 1 to minimize the errors is by searching for the best T_{split} according to the proposed adaptive methods in [21]. The subsequent T_{split} can be selected to maximize the value of the statistic $S_{\text{DRE}}^{T_{\text{split}}}$ at time instances with a slope change.

III. EXPERIMENTS

In this section, we present a comprehensive set of experiments to show: (i) the robustness of the DRE-CUSUM algorithm, (ii) the superiority of the DRE-CUSUM approach with other unsupervised techniques on both synthetic and real-world datasets, (iii) capability of detecting changes in high-dimensional video datasets. Particularly, the experiments on the event detection in video frames highlight the key aspect that *DRE-CUSUM is capable of demarcating the change points in very high-dimensional time-series data.*

Performance metrics: For evaluating DRE-CUSUM with other approaches, we use *false alarm rate* (FAR) and *missed detection rate* (MDR) [2] which is computed as,

$$FAR = \frac{FP}{FP + TN} \quad MDR = \frac{FN}{FN + TP} \quad (12)$$


where, TP, TN, FP, and FN denote the true positives, true negatives, false positives and false negatives, respectively.

Architecture of DRE models ($\hat{w}(x)$, estimate of $P_{\text{left}}/P_{\text{right}}$): For the scope of the experiments, we consider DRE modeled using kernels and deep neural networks (DNN's). We use the package provided in [8] for the kernel-based DRE. For the synthetic datasets, 4-layered feed-forward neural network based DRE is used with sigmoid, and softplus activations in the hidden, and final layers, respectively. For the change detection on video datasets, we use 4-layered convolutional neural network, with sigmoid, and softplus activations used in the hidden layers, and final layer, respectively. To train a DRE, a wide variety of training objectives such as KLIEP and LSIF have been widely accepted and used [18], [33], which we adopt to train the DRE's in the experimental section. Details on objective functions to train DRE, input to the DRE and more are presented in Appendix C.


A. Experiments on synthetic datasets

We first demonstrate the robustness of DRE-CUSUM to $|T^* - T_{\text{split}}|$, and distance between pre-change (P_1) and post-change distributions (P_2).

Robustness of DRE-CUSUM To demonstrate the robustness of the DRE-CUSUM to distance $|T^* - T_{\text{split}}|$, we consider a 10-dimensional time-series data with 1000 samples whose pre- and post-change distributions are sampled from multivariate Gaussian distributions with mean shift at time T^* as described in Fig. 5a. We set $T_{\text{split}} = 500$, the change point in the time-series data T^* is varied (say 20, 50, 100), thereby, varying the number of points in the time-series sampled from distributions P_1 and P_2 . From Fig. 5a, we infer that *DRE-CUSUM statistic $S_{\text{DRE}}^{T_{\text{split}}}(t)$ changes slope at T^* irrespective of $|T^* - T_{\text{split}}|$.* For checking the robustness of DRE-CUSUM to distance between P_1 and P_2 , we consider 10-dimensional time-series data, with a mean shift at time-instance $T^* = 350$. P_1 and P_2 are multivariate Gaussian distributions with same covariance matrix. We set mean $\bar{\mu}_1$ corresponding to P_1 as shown in Fig. 5b, and vary the difference $\Delta\bar{\mu} = |\bar{\mu}_2 - \bar{\mu}_1|$. From Fig. 5b, it is clear that *slope of DRE-CUSUM statistic $S_{\text{DRE}}^{T_{\text{split}}}(t)$ changes at $T^* = 350$ for relatively small $\Delta\bar{\mu}$.* We next compare the DRE-CUSUM approach with other unsupervised change detection approaches, particularly Bayesian change detection and its variants [9], [20]



(a) Robustness to $|T^* - T_{\text{split}}|$ when $T_{\text{split}} = 500$: P_1 and P_2 are multivariate Gaussian distributions whose mean vectors $\bar{\mu}_1$ and $\bar{\mu}_2$ are sampled from $\text{Unif.}[-1, 1]$ and $\text{Unif.}[-2, 2]$, respectively, with the covariance matrices equal to the identity matrix.



(b) Robustness to change in distributions: $T_{\text{split}} = 500$ and $T^* = 350$. P_1 and P_2 are multivariate Gaussian distributions with identity covariance matrix, and mean vectors $\bar{\mu}_1$ is sampled from $\text{Unif.}[-1, 1]$, and we plot $S_{\text{DR}}^{T_{\text{split}}}(t)$ for different values of $\Delta\bar{\mu} = |\bar{\mu}_1 - \bar{\mu}_2|$

Fig. 5: Robustness of the DRE-CUSUM algorithm
Comparison with other approaches We consider a 50-dimensional time-series data with 2000 samples generated from multivariate Gaussian distribution with same covariance matrix, such that it undergoes mean changes at 10 intervals. The mean vectors of the Gaussian distribution in different

Methodology	FAR	MDR
DRE-CUSUM (DNN, KLIEP)	0%	0%
DRE-CUSUM (DNN, LSIF)	0%	14.3%
DRE-CUSUM (Kernel, LSIF)	0.0005%	14.3%
Online BCD	~ 30%	~ 0%
Robust Online BCD	0.04%	42%

TABLE I: Comparison of online DRE-CUSUM with Online BCD [1], and Robust Online BCD [20]).

segments are sampled from uniform distributions (more details on the dataset generation have been included in the appendix). The results of DRE-CUSUM (online variant) along with other approaches have been tabulated in Table I, from which we infer that *DRE-CUSUM (for KLIEP objective) outperforms Bayesian approach*.

B. Experiments on real-world datasets


In this section, we compare performance of DRE-CUSUM with other unsupervised approaches on two real-world datasets. Furthermore, we present the results of the DRE-CUSUM for event detection tasks on video dataset.

Performance comparison on real-world datasets In this section, we perform evaluation against dynamic programming [34], Linearized penalty segmentation (Pelt) [40], and Binary segmentation (BinSeg) [10] techniques on the following *labeled* datasets: (i) HASC dataset [16], (ii) USC-HAD dataset [44], which we described next.


For change detection using HASC dataset, we consider the time-series data with 11857 samples with 11 change points (for example, activity shift from walking to jogging). Each datapoint in time-series is 3-dimensional corresponding to the recordings of the accelerometer along x, y, z axis. For change detection using USC dataset, we consider a time-series of 93635 samples with 35 change points [7], [44]. Each sample in the time-series is the accelerometer reading along x -axis that corresponds to the state at that instance (for example, sitting). *As seen from Table II, DRE-CUSUM outperforms its counterparts, particularly Pelt and dynamic programming methods, both of which are known to be accurate in a low-dimensional setting [34]. Furthermore, Pelt always has smaller FAR compared to other approaches.*

Video event detection using DRE-CUSUM To highlight the main advantage offered by the DRE-CUSUM approach, we consider 2012-Dataset, an *unlabeled* high dimensional video dataset [11]. We applied the DRE-CUSUM on two video/image sequences: canoe and overpass, with the objective of finding instances that demarcate the start or end of an event (for example, entry/exit of a boat). Fig. 6a and Fig. 6b correspond to DRE-CUSUM statistics for video frames from canoe and overpass dataset, respectively, and the details of the experimental setup are as described next.

(a) Canoe dataset: The time-series in Fig. 6a has 1189 video frames. We set $T_{\text{split}} = 580$. Frames 908, and 1056 marks the entry, and the exit of the boat, respectively. At the



(a) DRE-CUSUM algorithm on Canoe dataset to detect entry and exit of a boat.



(b) DRE-CUSUM algorithm on Overpass dataset for detecting entry and exit of both person/boat.

Fig. 6: Video event detection using DRE-CUSUM.

corresponding instances, we observe slope changes in DRE-CUSUM statistic. On visual inspection, we note that there are no significant changes at frame 336 (the slope change at $t = 336$ in DRE-CUSUM is observed for different values of T_{split}). We declare the slope change at 336 as a false alarm.

(b) Overpass dataset: In Fig. 6b, the time-series has 1500 samples, wherein we set $T_{\text{split}} = 700$. Slope changes present in DRE-CUSUM statistic around frames 553 and 684 corresponds to the object entry and exit frames, respectively. However, the slope change around the frame 332 is a false alarm.

Additional experimental results on video sequences in 2012-Dataset are provided in Appendix D. *We can extrapolate from these experimental results that: one can in principle use DRE-CUSUM to determine the change point estimates across any time-series data (both high-and low-dimensional). The slope changes can then be used for interpretation.*

IV. DISCUSSION AND FUTURE WORK

In this paper, we proposed DRE-CUSUM, a novel approach for unsupervised change detection, and showed its broad applicability on a wide range of applications backed by theoretical guarantees and experimental results. The salient aspect of DRE-CUSUM is that it does not require any knowledge/specification of the underlying distributions, nor an estimate of the number of underlying change points, and is universally applicable for high-dimensional data. To the best of our knowledge, our work is the first to provide theoretical justification and accuracy guarantees for the use of density ratio based unsupervised change detection. There are several possible directions for future work and we list some of them below:

Methodology	HASC		USC	
	FAR	MDR	FAR	MDR
Pelt [40]	0%	21.42%	0.0010%	74.28%
Dynamic Programming [34]	0.033%	21.42%	0.018%	37.14%
Binary segmentation [10]	0.033%	26.67%	0.024%	60%
DRE-CUSUM (Kernel,LSIF)	0%	0%	0.0075%	14.2%

TABLE II: Performance comparison of online DRE-CUSUM with Pelt, dynamic programming, and binary segmentation methods on HASC- [16] and USC- datasets [7], [44].

- a) Obtaining accuracy guarantees for DRE-CUSUM when there are multiple change points is an immediate interesting research direction.
- b) The accuracy guarantees for online adaptation of the DRE-CUSUM can be derived to understand the theoretical trade-off between between fixed vs. adaptive window size for a given time-series data.
- c) In real-world applications, we use density ratio estimators (DREs) to implement DRE-CUSUM. Studying the impact of sample complexity (i.e. number of samples required for a good estimation of the density ratio) of different DRE models on change point detection accuracy is also an important direction.

REFERENCES

- [1] Ryan Prescott Adams and David JC MacKay. Bayesian online change-point detection. *arXiv preprint arXiv:0710.3742*, 2007. **2, 7**
- [2] Samaneh Aminikhanghahi and Diane J Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017. **6**
- [3] Michele Basseville, Igor V Nikiforov, et al. *Detection of abrupt changes: theory and application*, volume 104. Prentice Hall Englewood Cliffs, 1993. **1, 2**
- [4] Alvaro A Cardenas, John S Baras, and Vahid Ramezani. Distributed change detection for worms, ddos and other network attacks. In *Proceedings of the 2004 American control conference*, volume 2, pages 1008–1013. IEEE, 2004. **1**
- [5] Costas Cotsaces, Nikos Nikolaidis, and Ioannis Pitas. Video shot detection and condensed representation. a review. *IEEE signal processing magazine*, 23(2):28–37, 2006. **1**
- [6] Rachel Cummings, Sara Krehbiel, Yajun Mei, Rui Tuo, and Wanrong Zhang. Differentially private change-point detection. *arXiv preprint arXiv:1808.10056*, 2018. **3, 10, 11**
- [7] Shohreh Deldari, Daniel V Smith, Hao Xue, and Flora D Salim. Time series change point detection with self-supervised contrastive predictive coding. In *Proceedings of the Web Conference 2021*, pages 3124–3135, 2021. **7, 8**
- [8] densratio. <https://pyipi.org/project/densratio/>, 2016. **6, 13, 14**
- [9] Paul Fearnhead. Exact and efficient bayesian inference for multiple changepoint problems. *Statistics and computing*, 16(2):203–213, 2006. **2, 6**
- [10] Piotr Fryzlewicz. Wild binary segmentation for multiple change-point detection. *The Annals of Statistics*, 42(6):2243–2281, 2014. **2, 7, 8**
- [11] Nil Goyette, Pierre-Marc Jodoin, Fatih Porikli, Janusz Konrad, and Prakash Ishwar. Changedetection. net: A new change detection benchmark dataset. In *2012 IEEE computer society conference on computer vision and pattern recognition workshops*, pages 1–8. IEEE, 2012. **7, 13**
- [12] Douglas M Hawkins and Qiqi Deng. A nonparametric change-point control chart. *Journal of Quality Technology*, 42(2):165–173, 2010. **1**
- [13] Chung-Lin Huang and Bing-Yao Liao. A robust scene-change detection method for video segmentation. *IEEE transactions on circuits and systems for video technology*, 11(12):1281–1288, 2001. **1**
- [14] Mikhail Hushchyn and Andrey Ustyuzhanin. Generalization of change-point detection in time series data based on direct density ratio estimation. *Journal of Computational Science*, 53:101385, 2021. **2**
- [15] Stefan Huwer and Heinrich Niemann. Adaptive change detection for real-time surveillance applications. In *Proceedings Third IEEE International Workshop on Visual Surveillance*, pages 37–46. IEEE, 2000. **1**
- [16] Haruyuki Ichino, Katsuhiko Kaji, Ken Sakurada, Kei Hiroi, and Nobuo Kawaguchi. Hasc-pac2016: Large scale human pedestrian activity corpus and its baseline recognition. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*, pages 705–714, 2016. **7, 8**
- [17] Alan Julian Izenman. Review papers: Recent developments in non-parametric density estimation. *Journal of the American Statistical Association*, 86(413):205–224, 1991. **2**
- [18] Takafumi Kanamori, Shohei Hido, and Masashi Sugiyama. A least-squares approach to direct importance estimation. *Journal of Machine Learning Research*, 10(Jul):1391–1445, 2009. **2, 4, 6, 12**
- [19] Rebecca Killick, Paul Fearnhead, and Idris A Eckley. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500):1590–1598, 2012. **2**
- [20] Jeremias Knoblauch, Jack E Jewson, and Theodoros Damoulas. Doubly robust bayesian inference for non-stationary streaming data with β -divergences. In *Advances in Neural Information Processing Systems*, pages 64–75, 2018. **6, 7**
- [21] Solt Kovács, Housen Li, Lorenz Haubner, Axel Munk, and Peter Bühlmann. Optimistic search strategy: Change point detection for large-scale data via adaptive logarithmic queries. *arXiv preprint arXiv:2010.10194*, 2020. **2, 6**
- [22] Danga Li, Kun Yang, and Wing Hung Wong. Density estimation via discrepancy based adaptive sequential partition. *arXiv preprint arXiv:1404.1425*, 2014. **2**
- [23] Song Liu, Akiko Takeda, Taiji Suzuki, and Kenji Fukumizu. Trimmed density ratio estimation. In *Advances in Neural Information Processing Systems*, pages 4518–4528, 2017. **12**
- [24] Luigi Tommaso Luppino, Stian Normann Anfinsen, Gabriele Moser, Robert Jenssen, Filippo Maria Bianchi, Sebastiano Serpico, and Gregoire Mercier. A clustering approach to heterogeneous change detection. In *Scandinavian Conference on Image Analysis*, pages 181–192. Springer, 2017. **2**
- [25] Martin McKeown, Colin Humphries, Peter Achermann, Alexander Borbély, and Terrence Sejnowsk. A new method for detecting state changes in the eeg: exploratory application to sleep data. *Journal of sleep research*, 7(S1):48–56, 1998. **1**
- [26] Hyunha Nam and Masashi Sugiyama. Direct density ratio estimation with convolutional neural networks with application in outlier detection. *IEICE TRANSACTIONS on Information and Systems*, 98(5):1073–1079, 2015. **2, 12**
- [27] XuanLong Nguyen, Martin J Wainwright, and Michael I Jordan. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010. **12**
- [28] Ewan S Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, 1954. **1**
- [29] Raquel Sebastião, João Gama, Pedro Pereira Rodrigues, and João Bernardes. Monitoring incremental histogram distribution for change detection in data streams. In *International Workshop on Knowledge Discovery from Sensor Data*, pages 25–42. Springer, 2008. **2**
- [30] A Stefanyuk. Estimation of the likelihood ratio function in the “disorder” problem of random processes. *Automation and Remote Control*, 9:53–59, 1986. **4**
- [31] Mark Steyvers and Scott Brown. Prediction and change detection. In *Advances in neural information processing systems*, pages 1281–1288, 2006. **1**

- [32] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. Density-ratio matching under the bregman divergence: a unified framework of density-ratio estimation. *Annals of the Institute of Statistical Mathematics*, 64(5):1009–1044, 2012. [2](#)
- [33] Masashi Sugiyama, Taiji Suzuki, Shinichi Nakajima, Hisashi Kashima, Paul von Büna, and Motoaki Kawanabe. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics*, 60(4):699–746, 2008. [2](#), [4](#), [6](#), [12](#)
- [34] Charles Truong, Laurent Oudre, and Nicolas Vayatis. Selective review of offline change point detection methods. *Signal Processing*, 167:107299, 2020. [2](#), [7](#), [8](#)
- [35] Aad W Van Der Vaart, Adrianus Willem van der Vaart, Aad van der Vaart, and Jon Wellner. *Weak convergence and empirical processes: with applications to statistics*. Springer Science & Business Media, 1996. [4](#)
- [36] Wouter JC Van Elmpt, Tamara ME Nijsen, Paul AM Griep, and Johan BAM Arends. A model of heart rate changes to detect seizures in severe epilepsy. *Seizure*, 15(6):366–375, 2006. [1](#)
- [37] Vladimir Vapnik, Igor Braga, and Rauf Izmailov. Constructive setting of the density ratio estimation problem and its rigorous solution. *arXiv preprint arXiv:1306.0407*, 2013. [2](#), [4](#)
- [38] Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999. [4](#)
- [39] Abraham Wald, Jacob Wolfowitz, et al. Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 19(3):326–339, 1948. [1](#)
- [40] G Dorcas Wambui, Gichuhi Anthony Waititu, and Anthony Wanjoya. The power of the pruned exact linear time (pelt) test in multiple changepoint detection. *American Journal of Theoretical and Applied Statistics*, 4(6):581, 2015. [2](#), [7](#), [8](#)
- [41] Weng-Keen Wong, Andrew Moore, Gregory Cooper, and Michael Wagner. What’s strange about recent events (wsare): An algorithm for the early detection of disease outbreaks. *Journal of Machine Learning Research*, 6(Dec):1961–1998, 2005. [1](#)
- [42] Makoto Yamada, Taiji Suzuki, Takafumi Kanamori, Hirotaoka Hachiya, and Masashi Sugiyama. Relative density-ratio estimation for robust distribution comparison. In *Advances in neural information processing systems*, pages 594–602, 2011. [12](#)
- [43] Jesin Zakaria, Abdullah Mueen, and Eamonn Keogh. Clustering time series using unsupervised-shapelets. In *2012 IEEE 12th International Conference on Data Mining*, pages 785–794. IEEE, 2012. [2](#)
- [44] Mi Zhang and Alexander A Sawchuk. Usc-had: a daily activity dataset for ubiquitous activity recognition using wearable sensors. In *Proceedings of the 2012 ACM conference on ubiquitous computing*, pages 1036–1043, 2012. [7](#), [8](#)

APPENDIX A PROOF OF PROPOSITION 1

Proof. For the case when $T_{\text{split}} \leq T^*$, we have the following two distributions before and after T_{split} , respectively:

$$P_{\text{left}} = P_1, \quad P_{\text{right}} = \alpha^{T_{\text{split}}} P_1 + (1 - \alpha^{T_{\text{split}}}) P_2, \quad (13)$$

where, $\alpha^{T_{\text{split}}} = \frac{T^* - T_{\text{split}}}{n - T_{\text{split}}}$. The expected values of the log-likelihood ratio $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$ before T^* (i.e., $\forall t < T^*$) is,

$$\begin{aligned} & \mathbb{E}_{x_t} \left[\log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) \right] \\ &= \mathbb{E}_{x_t \sim P_1} \left[\log \left(\frac{P_1(x_t)}{\alpha^{T_{\text{split}}} P_1(x_t) + (1 - \alpha^{T_{\text{split}}}) P_2(x_t)} \right) \right] \\ &= \int_x P_1(x) \log \left(\frac{P_1(x)}{\alpha^{T_{\text{split}}} P_1(x) + (1 - \alpha^{T_{\text{split}}}) P_2(x)} \right) dx \\ &\stackrel{(a)}{=} KL(P_1 || P_{\text{right}}) \end{aligned} \quad (14)$$

where, (a) follows from the definition of KL-divergence. For simplicity of notation, let us define $\alpha_1 \triangleq \frac{n - T^*}{n - T_{\text{split}}}$, therefore, $\alpha^{T_{\text{split}}} = 1 - \alpha_1$, and from definition of the parametric mixture

distribution, we can write $P_{\text{right}} = P(1 - \alpha_1)$. Substituting in (14) we get,

$$\mathbb{E}_{x_t} \left[\log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) \right] = KL(P_1 || P(1 - \alpha_1)) \geq 0 \quad (15)$$

Similarly, we obtain the expected values of the log-likelihood ratio $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$ for any $t \geq T^*$ as follows:

$$\begin{aligned} & \mathbb{E}_{x_t} \left[\log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) \right] \\ &= \mathbb{E}_{x_t \sim P_2} \left[\log \left(\frac{P_1(x_t)}{\alpha^{T_{\text{split}}} P_1(x_t) + (1 - \alpha^{T_{\text{split}}}) P_2(x_t)} \right) \right] \\ &= - \mathbb{E}_{x_t \sim P_2} \left[\log \left(\frac{\alpha^{T_{\text{split}}} P_1(x_t) + (1 - \alpha^{T_{\text{split}}}) P_2(x_t)}{P_1(x_t)} \right) \right] \\ &= - \int_x P_2(x) \log \left(\frac{P_{\text{right}}(x)}{P_1(x)} \right) dx \\ &\stackrel{(a)}{=} - \frac{1}{1 - \alpha^{T_{\text{split}}}} \left[\int_x P_{\text{right}}(x) \log \left(\frac{P_{\text{right}}(x)}{P_1(x)} \right) dx \right. \\ &\quad \left. + \int_x \alpha^{T_{\text{split}}} P_1(x) \log \left(\frac{P_1(x)}{P_{\text{right}}(x)} \right) dx \right] \\ &\stackrel{(b)}{=} - \left[\frac{1}{1 - \alpha^{T_{\text{split}}}} KL(P_{\text{right}} || P_1) + \frac{\alpha^{T_{\text{split}}}}{1 - \alpha^{T_{\text{split}}}} KL(P_1 || P_{\text{right}}) \right] \\ &\stackrel{(c)}{=} - \left[\frac{1}{\alpha_1} KL((1 - \alpha_1)P_1 + \alpha_1 P_2 || P_1) \right. \\ &\quad \left. + \frac{1 - \alpha_1}{\alpha_1} KL(P_1 || (1 - \alpha_1)P_1 + \alpha_1 P_2) \right] \\ &\stackrel{(d)}{=} - f_1(\alpha_1, (1 - \alpha_1)) \leq 0 \end{aligned} \quad (16)$$

where, (a) and (c) follows from substituting $P_2(x_t) = (P_{\text{right}}(x_t) - \alpha^{T_{\text{split}}} P_1(x_t)) / (1 - \alpha^{T_{\text{split}}})$ and $\alpha_1 = 1 - \alpha^{T_{\text{split}}}$, respectively. (b) and (d) follow from definition of KL-divergence and definition 1 in the paper, respectively.

For the case when $T_{\text{split}} \geq T^*$, we have the following two distributions before and after T_{split} respectively:

$$P_{\text{left}} = \alpha^{T_{\text{split}}} P_1 + (1 - \alpha^{T_{\text{split}}}) P_2, \quad P_{\text{right}} = P_2, \quad (17)$$

where, $\alpha^{T_{\text{split}}} = \frac{T^*}{T_{\text{split}}}$. Expected value of the log-likelihood ratio $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$ before the change point T^* (i.e., $\forall t < T^*$) is:

$$\begin{aligned} & \mathbb{E}_{x_t} \left[\log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) \right] \\ &= \mathbb{E}_{x_t \sim P_1} \left[\log \left(\frac{\alpha^{T_{\text{split}}} P_1(x_t) + (1 - \alpha^{T_{\text{split}}}) P_2(x_t)}{P_2(x_t)} \right) \right] \\ &= \int_x P_1(x) \log \frac{P_{\text{left}}(x)}{P_2(x)} dx \\ &\stackrel{(a)}{=} \frac{1}{\alpha^{T_{\text{split}}}} \left[\int_x P_{\text{left}}(x) \log \frac{P_{\text{left}}(x)}{P_2(x)} dx \right. \\ &\quad \left. - \int_x (1 - \alpha^{T_{\text{split}}}) P_2(x) \log \frac{P_{\text{left}}(x)}{P_2(x)} dx \right] \\ &\stackrel{(b)}{=} \frac{1}{\alpha^{T_{\text{split}}}} KL(P_{\text{left}} || P_2) + \frac{1 - \alpha^{T_{\text{split}}}}{\alpha^{T_{\text{split}}}} KL(P_2 || P_{\text{left}}) \end{aligned} \quad (18)$$

where, (a) follows from the fact that $P_1(x_t) = (P_{\text{left}}(x_t) - (1 - \alpha^{T_{\text{split}}})P_2(x_t)) / \alpha^{T_{\text{split}}}$, (b) follows from the definition of KL-divergence. For simplicity of notation, let us define $\alpha_2 \triangleq \alpha^{T_{\text{split}}}$, and subsequently $\forall t \geq T^*$, (18) becomes

$$\begin{aligned} \mathbb{E}_{x_t} \left[\log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) \right] &= \frac{1}{\alpha_2} KL(P_{\text{left}} \| P_2) + \frac{1 - \alpha_2}{\alpha_2} KL(P_2 \| P_{\text{left}}) \\ &= f_2(\alpha_2, \alpha_2) \geq 0. \end{aligned} \quad (19)$$

The expected values of the log-likelihood ratio $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$ for all points after T^* (i.e., $\forall t \geq T^*$) can be computed as

$$\begin{aligned} \mathbb{E}_{x_t} \left[\log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) \right] &= \mathbb{E}_{x_t \sim P_2} \left[\log \frac{\alpha^{T_{\text{split}}} P_1(x_t) + (1 - \alpha^{T_{\text{split}}}) P_2(x_t)}{P_2(x_t)} \right] \\ &= -KL(P_2 \| P_{\text{left}}) = -KL(P_2 \| P(\alpha_2)) \leq 0. \end{aligned} \quad (20)$$

As a consequence of the above results, we infer that for any choice of T_{split}

$$\mathbb{E}_{x_t} \log \left(\frac{P_{\text{left}}(x_t)}{P_{\text{right}}(x_t)} \right) = \begin{cases} \geq 0, & \text{for } t < T^* \\ \leq 0, & \text{for } t \geq T^* \end{cases} \quad (21)$$

From the above result, Corollary 1 stated in the paper follows immediately, since the statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$ is a linear function in t with a non-negative slope before T^* , and conversely, a linear function in t with a non-positive slope for all points after T^* . \square

APPENDIX B PROOF OF THEOREM 1

To prove that the change estimate $\hat{T}_{\text{DR-CUSUM}}$ is (α, β) -accurate, we use similar reasoning to that used in [6] to show that the maximum likelihood estimate \hat{T}_{ML} is (α, β) -accurate. However, the maximum likelihood approach assumes that $P_1(\cdot)/P_2(\cdot)$ can be readily computed. The key distinction of our proof is that we show the change point estimate using DR-CUSUM $\hat{T}_{\text{DR-CUSUM}}$ is (α, β) -accurate even when $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$ can be computed in the unsupervised setting. To prove this result, we consider two separate cases: (a) $T_{\text{split}} \leq T^*$ and (b) $T_{\text{split}} \geq T^*$. For case (a), when $T_{\text{split}} \leq T^*$, we have

$$P_{\text{left}} = P_1, \quad P_{\text{right}} = \alpha^{T_{\text{split}}} P_1 + (1 - \alpha^{T_{\text{split}}}) P_2. \quad (22)$$

The density-ratio CUSUM statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$ is defined $\forall t \in [1, n]$ as:

$$S_{\text{DR}}^{T_{\text{split}}}(t) \triangleq \sum_{j=1}^t \left(\log \left(\frac{P_{\text{left}}(x_j)}{P_{\text{right}}(x_j)} \right) \right) \quad (23)$$

For a given α , let us define the region \mathcal{R} as:

$$\mathcal{R} = [n] \setminus [T^* - \alpha, T^* + \alpha] \quad (24)$$

which essentially consists of all time instances which are at least a distance α from the true change point T^* . We first note the following inequality:

$$\begin{aligned} P[|\hat{T}_{\text{DR-CUSUM}} - T^*| > \alpha] &\leq P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) \end{aligned} \quad (25)$$

which is a direct consequence of the definition of \mathcal{R} and the DRE-CUSUM estimator. From now on, our goal will be to upper bound the probability $P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0)$. To this end, we consider regions $\mathcal{R}^- = [1, T^* - \alpha]$ and $\mathcal{R}^+ = (T^* + \alpha, n]$ as shown in Fig. 7a, such that $\mathcal{R} = \mathcal{R}^- \cup \mathcal{R}^+$. Thereafter, on applying union-bound to the r.h.s. of (25) over regions \mathcal{R}^- and \mathcal{R}^+ , we have

$$\begin{aligned} P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) &\leq P(\max_{t \in \mathcal{R}^-} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) \\ &\quad + P(\max_{t \in \mathcal{R}^+} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) \end{aligned} \quad (26)$$

We have the following two cases to compute $S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*)$:

$$\begin{aligned} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) &= \begin{cases} -\sum_{i=t}^{T^*-1} \log [P_{\text{left}}(x_i)/P_{\text{right}}(x_i)], & t \in \mathcal{R}^- \\ -\sum_{i=T^*}^{t-1} \log [P_{\text{right}}(x_i)/P_{\text{left}}(x_i)], & t \in \mathcal{R}^+ \end{cases} \end{aligned} \quad (27)$$

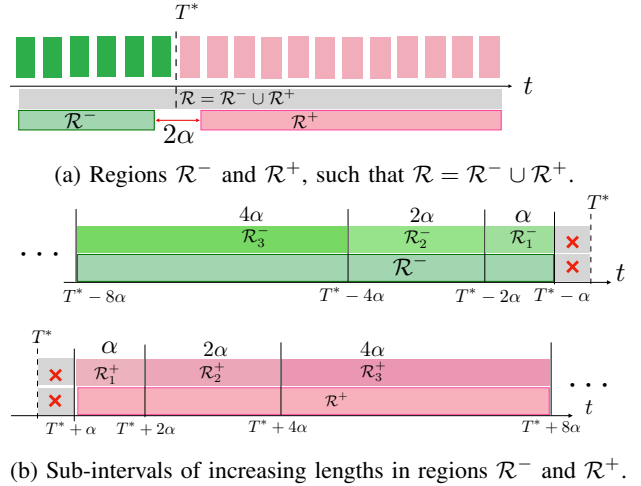


Fig. 7: Region \mathcal{R} split into to sub-regions $\mathcal{R}^- = [1, T^* - \alpha]$ and $\mathcal{R}^+ = (T^* + \alpha, n]$ in Fig. 1(a), which is further segmented to smaller regions as shown in Fig. 1(b).

We further segment regions \mathcal{R}^- and \mathcal{R}^+ , such that the interval lengths double in length as we move away from change point T^* as shown in Fig. 7b. Assuming finite samples in the time-series data, the total number intervals in \mathcal{R}^- , and \mathcal{R}^+ are $\log_2 \left(\frac{T^*}{\alpha} \right)$, and $\log_2 \left(\frac{n - T^*}{\alpha} \right)$, respectively. Therefore, we can

simplify (26) further as:

$$\begin{aligned}
& P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T^* \text{split}}(t) - S_{\text{DR}}^{T^* \text{split}}(T^*) > 0) \\
& \stackrel{(a)}{\leq} P\left(\max_{t \in \mathcal{R}^-} - \sum_{i=t}^{T^*-1} \log \frac{P_{\text{left}}(x_i)}{P_{\text{right}}(x_i)} > 0\right) \\
& \quad + P\left(\max_{t \in \mathcal{R}^+} - \sum_{i=T^*}^{t-1} \log \frac{P_{\text{right}}(x_i)}{P_{\text{left}}(x_i)} > 0\right) \\
& \stackrel{(b)}{\leq} \underbrace{\sum_{i=1}^{\log_2(\frac{T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^-} - \sum_{i=t}^{T^*-1} \log \left(\frac{P_{\text{left}}(x_i)}{P_{\text{right}}(x_i)}\right) > 0\right)}_{S_1} \\
& \quad + \underbrace{\sum_{i=1}^{\log_2(\frac{n-T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^+} - \sum_{i=T^*}^{t-1} \log \left(\frac{P_{\text{right}}(x_i)}{P_{\text{left}}(x_i)}\right) > 0\right)}_{S_2}
\end{aligned} \tag{28}$$

where, (a) follows from substituting (27) in (26), and (b) follows from applying union bound over segments in regions \mathcal{R}^- and \mathcal{R}^+ depicted in Fig. 7b. We individually upper bound the terms S_1 , and S_2 in (28). The term S_1 in (28), when $T_{\text{split}} \leq T^*$ can be simplified as follows:

$$\begin{aligned}
& \sum_{i=1}^{\log_2(\frac{T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^-} - \sum_{i=t}^{T^*-1} \log \left(\frac{P_{\text{left}}(x_i)}{P_{\text{right}}(x_i)}\right) > 0\right) \\
& \stackrel{(a)}{\leq} \sum_{i=1}^{\log_2(\frac{T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^-} - \sum_{i=t}^{T^*-1} \log \left(\frac{P_{\text{left}}(x_i)}{P_{\text{right}}(x_i)}\right) \right. \\
& \quad \left. + (T^* - t) KL(P_1 || P(1 - \alpha_1)) \right. \\
& \quad \left. - (T^* - t) KL(P_1 || P(1 - \alpha_1)) > 0\right) \\
& \stackrel{(b)}{\leq} \sum_{i=1}^{\log_2(\frac{T^*}{\alpha})} P(\max_{t \in \mathcal{R}_i^-} | \sum_{j=t}^{T^*-1} u(x_j) | > 2^{i-1} \alpha KL(P_1 || P(1 - \alpha_1))) \\
& \stackrel{(c)}{\leq} \sum_{i=1}^{\log_2(\frac{T^*}{\alpha})} \frac{2 \exp(-2^{i-2} \alpha (KL(P_1 || P(1 - \alpha_1)))^2 / A^2)}{1 - 2 \exp(-2^{i-2} \alpha (KL(P_1 || P(1 - \alpha_1)))^2 / A^2)} \\
& \stackrel{(d)}{\leq} \sum_{i=1}^{\log_2(\frac{T^*}{\alpha})} 4 \exp(-2^{i-2} \alpha (KL(P_1 || P(1 - \alpha_1)))^2 / A^2)
\end{aligned} \tag{29}$$

where, (a) follows by adding and subtracting $(T^* - t) KL(P_1 || P(1 - \alpha_1))$, (b) follows from the fact that for any interval $\mathcal{R}_i^- = (T^* - 2^i \alpha, T^* - 2^{i-1} \alpha]$, we have $T^* - t \geq 2^{(i-1)} \alpha$ and from definition $u(x_j) = -\log[P_{\text{left}}(x_j)/P_{\text{right}}(x_j)] + KL(P_1 || P(1 - \alpha_1))$, for any x_j such that $u(\cdot)$ has zero mean, (c) follows directly from application of Corollary 2 in [6], which states that: For $S_k = \sum_{i \in [k]} u_i$ for $k \in [m]$, where u_1, \dots, u_m are i.i.d random variables with mean zero and strictly bounded by

interval of length L , we have

$$\Pr \left[\max_{k \in [m]} |S_k| > \lambda_1 + \lambda_2 \right] \leq \frac{2 \exp(-2\lambda_1^2 / (mL^2))}{1 - 2 \exp(-2\lambda_2^2 / (mL^2))} \tag{30}$$

where, $\lambda_1, \lambda_2 > 0$. The concentration inequality in (30) stems by further upper bounding Ottaviani's inequality using Hoeffding's inequality under the assumption that u_j can only take values from interval of bounded length L . Furthermore, in step (c) in (29), as the individual terms represent probabilities, we can upper bound the terms:

$$\begin{aligned}
& \frac{2 \exp(-2^{i-2} \alpha (KL(P_1 || P(1 - \alpha_1)))^2 / A^2)}{1 - 2 \exp(-2^{i-2} \alpha (KL(P_1 || P(1 - \alpha_1)))^2 / A^2)} \leq 1 \\
& \implies \exp(-2^{i-2} \alpha (KL(P_1 || P(1 - \alpha_1)))^2 / A^2) < \frac{1}{4} \tag{31}
\end{aligned}$$

Following similar steps, we can bound the term S_2 in (28) when $T_{\text{split}} \leq T^*$ as

$$\begin{aligned}
& \sum_{i=1}^{\log_2(\frac{n-T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^+} - \sum_{j=T^*}^{t-1} \log \left(\frac{P_{\text{right}}(x_j)}{P_{\text{left}}(x_j)}\right) > 0\right) \\
& \stackrel{(a)}{\leq} \sum_{i=1}^{\log_2(\frac{n-T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^+} - \sum_{j=T^*}^{t-1} \log \left(\frac{P_{\text{right}}(x_j)}{P_{\text{left}}(x_j)}\right) \right. \\
& \quad \left. + (t - T^*) f_1(\alpha_1, 1 - \alpha_1) \right. \\
& \quad \left. - (t - T^*) f_1(\alpha_1, 1 - \alpha_1) > 0\right) \\
& \stackrel{(b)}{\leq} \sum_{i=1}^{\log_2(\frac{n-T^*}{\alpha})} P\left(\max_{t \in \mathcal{R}_i^+} | \sum_{j=T^*}^{t-1} u(x_j) | > 2^{(i-1)} \alpha f_1(\alpha_1, 1 - \alpha_1)\right) \\
& \stackrel{(c)}{\leq} \sum_{i=1}^{\log_2(\frac{n-T^*}{\alpha})} \frac{2 \exp(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2)}{1 - 2 \exp(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2)} \\
& \stackrel{(d)}{\leq} \sum_{i=1}^{\log_2(\frac{n-T^*}{\alpha})} 4 \exp(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2)
\end{aligned} \tag{32}$$

The steps (a,b,c,d) in (32) follow the same logical reasoning to steps (a,b,c) in (29). Step (e) follows from the same logical reasoning as (31), which is.

$$\begin{aligned}
& \frac{2 \exp(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2)}{1 - 2 \exp(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2)} \leq 1 \\
& \implies \exp(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2) < \frac{1}{4} \tag{33}
\end{aligned}$$

Substituting (29) and (32) in (28) we get,

$$\begin{aligned}
& P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) \\
& \leq \sum_{i=1}^{\frac{T^*}{\alpha}} 4 \exp\left(-2^{i-2} \alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / A^2\right) \\
& \quad + \sum_{i=1}^{\log_2\left(\frac{n-T^*}{\alpha}\right)} 4 \exp\left(-2^{i-2} \alpha f_1^2(\alpha_1, 1 - \alpha_1) / A^2\right) \\
& \stackrel{(a)}{=} \sum_{i=1}^{\frac{T^*}{\alpha}} 4 \exp\left(-\alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / 2A^2\right)^{2^{i-1}} \\
& \quad + \sum_{i=1}^{\log_2\left(\frac{n-T^*}{\alpha}\right)} 4 \exp\left(-\alpha f_1^2(\alpha_1, 1 - \alpha_1) / 2A^2\right)^{2^{i-1}} \\
& \stackrel{(b)}{\leq} \sum_{i=1}^{\frac{T^*}{\alpha}} 4 \exp\left(-\alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / 2A^2\right)^i \\
& \quad + \sum_{i=1}^{\log_2\left(\frac{n-T^*}{\alpha}\right)} 4 \exp\left(-\alpha f_1^2(\alpha_1, 1 - \alpha_1) / 2A^2\right)^i \\
& \stackrel{(c)}{\leq} \frac{4 \exp\left(-\alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / 2A^2\right)}{1 - \exp\left(-\alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / 2A^2\right)} \\
& \quad + \frac{4 \exp\left(-\alpha f_1^2(\alpha_1, 1 - \alpha_1) / 2A^2\right)}{1 - \exp\left(-\alpha f_1^2(\alpha_1, 1 - \alpha_1) / 2A^2\right)} \quad (34)
\end{aligned}$$

From (31) and (33), the following inequalities hold:

$$\begin{aligned}
& \exp\left(-\alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / 2A^2\right) < \frac{1}{4} \\
& \exp\left(-\alpha f_1^2(\alpha_1, 1 - \alpha_1) / 2A^2\right) < \frac{1}{4} \quad (35)
\end{aligned}$$

Thereby substituting in (34) we get,

$$\begin{aligned}
& P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) \\
& \leq \frac{16}{3} \left[\exp\left(-\alpha (KL(P_1 \| P(1 - \alpha_1)))^2 / 2A^2\right) \right. \\
& \quad \left. + \exp\left(-\alpha f_1^2(\alpha_1, 1 - \alpha_1) / 2A^2\right) \right] \\
& \leq \frac{32}{3} \exp\left(-\alpha C^2 / 2A^2\right) \quad (36)
\end{aligned}$$

where, $C = \min\{KL(P_1 \| P(1 - \alpha_1)), f_1(\alpha_1, 1 - \alpha_1)\}$. The analysis when $T_{\text{split}} \geq T^*$ follows in a similar manner, wherein we can show the following for the DR-CUSUM statistic $S_{\text{DR}}^{T_{\text{split}}}(t)$:

$$P(\max_{t \in \mathcal{R}} S_{\text{DR}}^{T_{\text{split}}}(t) - S_{\text{DR}}^{T_{\text{split}}}(T^*) > 0) \leq \frac{32}{3} \exp\left(-\alpha C^2 / 2A^2\right) \quad (37)$$

where, $C = \min\{f_2(\alpha_2, \alpha_2), KL(P_2 \| P(\alpha_2))\}$. This completes the proof of Theorem 1.

APPENDIX C

OVERVIEW ON DENSITY RATIO ESTIMATORS (DRE)

In this section, we present a brief overview of data driven approaches for density ratio estimation (DRE) [33]. The basic problem statement is the following: we are given samples from two distributions P_{left} and P_{right} and the goal is to estimate the density ratio functional $w(x) = P_{\text{left}}(x) / P_{\text{right}}(x)$. The main idea behind DRE is to use a parametric model for the density ratio functional, and then to learn the parameters using only samples from P_{left} and P_{right} in a principled manner.

For instance, kernels, feed-forward or convolutional neural network (CNNs) can be used to model the density ratio estimator ($\hat{w}(x)$) [33] [18] [26]. Kernel based estimators are typically preferred low-dimensional data, while NNs have been shown to give better performance otherwise [26]. We next describe one such principled approach [33] (Kullback-Leibler Importance Estimation Procedure (KLIEP)) for DRE. Given the estimated density ratio $\hat{w}(x)$ (output of DRE model), if we know the true density value $P_{\text{right}}(x)$, we can estimate density $P_{\text{left}}(x)$ as follows,

$$\hat{P}_{\text{left}}(x) = \hat{w}(x) P_{\text{right}}(x) \quad (38)$$

For training the DRE model, KLIEP suggests minimizing the KL divergence between $P_{\text{left}}(\cdot)$ and estimated density $\hat{P}_{\text{left}}(\cdot)$, which simplifies to:

$$\begin{aligned}
KL(P_{\text{left}} \| \hat{P}_{\text{left}}) & \stackrel{(i)}{=} \int_x P_{\text{left}}(x) \log \frac{P_{\text{left}}(x)}{\hat{w}(x) P_{\text{right}}(x)} dx \\
& = KL(P_{\text{left}} \| P_{\text{right}}) - \int_x P_{\text{left}}(x) \log \hat{w}(x) dx \\
& = KL(P_{\text{left}} \| P_{\text{right}}) - \mathbb{E}_{x \sim P_{\text{left}}} [\log(\hat{w}(x))] \quad (39)
\end{aligned}$$

where, (i) follows from definition of KL divergence and by using (38). In (39), the first term is a constant w.r.t. DRE model parameters. Hence, the minimization of the KL-divergence in (39) is equivalent to maximization of $\mathbb{E}_{x \sim P_{\text{left}}} [\log(\hat{w}(x))]$. Furthermore, for the estimated \hat{P}_{left} to be a valid density, it must satisfy

$$\int_x \hat{P}_{\text{left}}(x) = \int_x \hat{w}(x) P_{\text{right}}(x) = \mathbb{E}_{x \sim P_{\text{right}}} [\hat{w}(x)] = 1 \quad (40)$$

Using Lagrange parameter λ to satisfy the constraints, we obtain the constrained optimization for the DRE as follows,

KLIEP Objective:

$$\max_{\hat{w}} \left(\mathbb{E}_{x \sim P_{\text{left}}} [\log(\hat{w}(x))] - \lambda (\mathbb{E}_{x \sim P_{\text{right}}} [\hat{w}(x)] - 1) \right) \quad (41)$$

By replacing the expectations by the sample means over P_{left} and P_{right} , one can then use gradient ascent based optimization to find \hat{w} . We want to highlight that there are several other approaches for density ratio estimation [18], [23], [42]. For instance, an unconstrained optimization (Least Squares Importance Fitting (LSIF)) is obtained in [18] (also see [27]) by minimizing the least square loss between the actual and estimated density ratios.

LSIF Objective:

$$\min_{\hat{w}} \left(\mathbb{E}_{x \sim P_{\text{left}}} \hat{w}(x)^2 - 2\mathbb{E}_{x \sim P_{\text{right}}} \hat{w}(x) \right) \quad (42)$$

Any model (neural network or kernels) for density ratio estimation can be leveraged in principle, for the design of our DRE-CUSUM change detection framework.

Implementing DRE models

In practice, we can only determine the empirical approximation to the objective functions in (41), (42) (for both KLIEP and LSIF). Given, samples $X_{[1:T_{\text{split}}]} \sim P_{\text{left}}$ and $X_{[T_{\text{split}}:n]} \sim P_{\text{right}}$. To estimate $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$, we sample N_1 samples from $X_{[1:T_{\text{split}}]}$ and N_2 samples from $X_{[T_{\text{split}}:n]}$ in each iteration, for training a density ratio estimator model denoted by \hat{w} . For implementation, KLIEP objective in (40) is empirically computed as follows:

$$\frac{1}{N_1} \sum_{x \sim P_{\text{left}}} \log(\hat{w}(x)) - \lambda \left(\frac{1}{N_2} \left(\sum_{x \sim P_{\text{right}}} \hat{w}(x) \right) - 1 \right) \quad (43)$$

Likewise, the LSIF objective in (42) can be empirically computed as:

$$\frac{1}{N_1} \sum_{x \sim P_{\text{left}}} \hat{w}(x)^2 - \frac{2}{N_2} \sum_{x \sim P_{\text{right}}} \hat{w}(x) \quad (44)$$

Algorithm 2 Training density ratio estimator (DRE) models

INPUT: where $X_{[1:T_{\text{split}}]} \sim P_{\text{left}}$ and $X_{[T_{\text{split}}:n]} \sim P_{\text{right}}$

OUTPUT: Estimate of $P_{\text{left}}(\cdot)/P_{\text{right}}(\cdot)$

for number of epochs **do**

1. Sample a mini-batch of N_1 samples from $X_{[1:T_{\text{split}}]}$ and N_2 samples from $X_{[T_{\text{split}}:n]}$.
2. Compute objective in (43), which is

$$\frac{1}{N_1} \sum_{x \sim P_{\text{left}}} \log(\hat{w}(x)) - \lambda \left(\frac{1}{N_2} \left(\sum_{x \sim P_{\text{right}}} \hat{w}(x) \right) - 1 \right)$$

3. Update parameters of the DRE model (for example, neural network weights) to maximize the above objective using gradient descent.

end for

We summarize the training of DRE model using the KLIEP objective in (43) in Algorithm 2, although one could use adapt Algorithm 2 to LSIF objective by optimizing the parameters of the DRE model to minimize (44). We next describe: the architecture of DRE models, the training of the DRE models, and details on synthetic dataset generation used for the results in the main paper.

A. DRE Architectures used in the paper

For the scope of the experiments using DRE modeled using neural network, the neural network architectures are tabulated in Table III. For the kernel-based DRE, we use the package provided in [8].

Details of experiment to test robustness of DRE-CUSUM: In the Fig. 5b (in the paper) we have the 10-dimensional time-series data $X_{[1:1000]}$ with change point $T^* = 350$, such that $X_{[1:350]} \sim P_1$ and $X_{[350:1000]} \sim P_2$, where $P_1 \sim \mathcal{N}(\vec{\mu}_1, I)$ and $P_2 \sim \mathcal{N}(\vec{\mu}_2, I)$. Entries of mean-vector $\vec{\mu}_1$ is sampled from $\text{Unif.}[-1, 1]$, while $\vec{\mu}_2$ is varied by adding (small) increments to $\vec{\mu}_1$. We set $T_{\text{split}} = 500$. Consequently, for this example, we have $P_{\text{left}} = \alpha^{T_{\text{split}}} P_1 + (1 - \alpha^{T_{\text{split}}}) P_2$, where $\alpha^{T_{\text{split}}} = 0.7$ (i.e. = 350/500). However, we have $P_{\text{right}} = P_2$. In each epoch, we sample a minibatch of data from both P_{left} and P_{right} , and train the DRE model with either using (43) or (44) as the training objective. Per table III, we use a feed-forward neural network with 10- input nodes. The width of the three hidden layers are 256, 512, 128, respectively, and the output of the neural network DRE is the estimated density ratio $\hat{w}(x)$ corresponding to the input sample x . During training, we set the size of the mini-batch to be 64, and train the neural network DRE for 500 iterations. We train the neural network DRE using KLIEP objective in (43). Post-training, we compute the density ratio $\hat{w}(x)$, $\forall x \in X_{[1:1000]}$, and subsequently plot the DRE-CUSUM statistic which is depicted in Fig. 5b in the paper. We refer the readers to [this link](#) (which was provided in the paper) which contains the code to generate the results. *Details of experiment in Table 1 of the paper:* We generate a 50-dimensional Gaussian time-series data of length 2000 with change points at time-instances $t = \{150, 200, 450, 525, 700, 725, 1200\}$. We set the co-variance matrix across all segments in the time-series data to be the same, and is generated as follows: each entry in standard deviation vector σ is sampled from a distribution $\text{Unif.}[1, 3]$. We obtain covariance matrix $\Sigma = \sigma^T \sigma I$. However, we vary the mean vector across the different segments in the time-series data, such that the entries of the mean vector across different segments (ordered) are sampled from: (i) $\text{Unif.}[-1, 1]$, (ii) $\text{Unif.}[-2, 2]$, (iii) $\text{Unif.}[-3, 3]$, (iv) $\text{Unif.}[-4, 4]$, (v) $\text{Unif.}[-3, 3]$, (vi) $\text{Unif.}[-10, 10]$, (vii) $\text{Unif.}[-20, 20]$, (viii) $\text{Unif.}[-1, 1]$. The architecture used in this experiment is a feed-forward neural network with 50 input nodes. The hidden layer widths (from input layer towards output) are 256, 512, 128, and the neural network based DRE is trained for 500 iterations. The experiment is performed using both KLIEP and LSIF objectives.

APPENDIX D

ADDITIONAL EXPERIMENTAL RESULTS ON VIDEO DATASETS

We conducted additional experiments using DRE-CUSUM on real-world video data in 2012-Dataset [11]. In particular, the objective was to perform activity detection (in particular, detect the entry/exit of a person) in the sequence of video frames. We present the results on pedestrian and overpass video sequences present in the 2012-Dataset [11]. In this experiment, with a time-series of 240 frames, a person is present in frames 0 – 100. As shown in Fig. 8a, we first set $T_{\text{split}} = 120$. We observe slope changes at frames 65 and 100. It can be noted that, the video frames 65 – 100 belong to

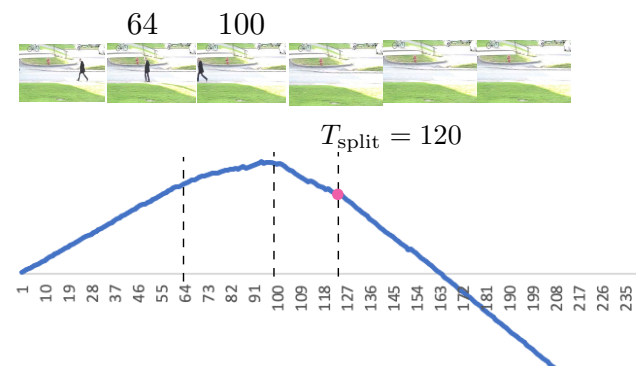
<i>Experiment</i>	<i>DRE model</i>	<i>Architecture details</i>
Synthetic datasets	Feedforward neural network DRE	4 dense layers Hidden layer activation: Sigmoid Final layer activation: Softplus
Real-world datasets (USC, HASC)	Kernel based DRE	Kernel type: Gaussian [8]
Video datasets	Convolutional neural network DRE	4 convolutional layers Hidden layer activation: Sigmoid Final layer activation: Softplus

TABLE III: Neural network DRE architecture details.

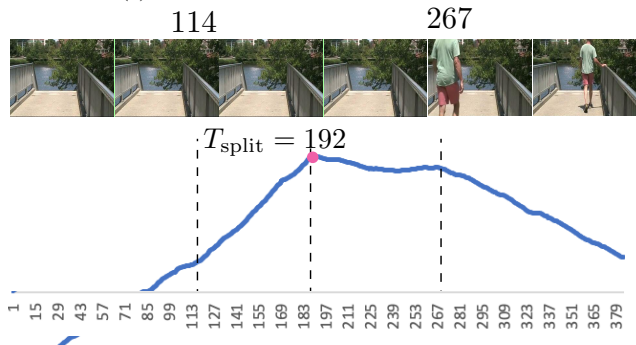
the transition period when the person gradually exists and is no more present in the video. As we can observe, the DRE-CUSUM statistic is able to detect both the beginning and end of these transition frames.

In the second experiment, we consider a time-series of 385 frames. The person appears in the 260th frame. We first set $T_{\text{split}} = 192$ and obtain the corresponding DRE-CUSUM statistic as shown in Fig. 8b. Slope changes at are observed at instances corresponding to frames 192 (i.e. T_{split}), and 267. The slope change at around frame 120 corresponds to a false alarm (upon visual inspection no change is observed).

the hidden layers of the convolutional neural network based DRE, we apply max-pooling, and the KLIEP objective is used to train the parameters of the neural network. We train the neural network DRE for 2000 iterations.



(a) DRE-CUSUM on Pedestrian dataset.



(b) DRE-CUSUM algorithm on Overpass dataset.

Fig. 8: Video event detection using DRE-CUSUM for detecting entry and exit instances of a person.

Additional Architectural details: In general for event detection experiments, the architecture in Table III is suitable. In