
Online Heavy-tailed Change-point detection

Abishek Sankararaman*¹

Balakrishnan (Murali) Narayanaswamy¹

¹AWS AI Labs

Abstract

We study algorithms for online change-point detection (OCPD), where samples that are potentially heavy-tailed, are presented one at a time and a change in the underlying mean must be detected as early as possible. We present an algorithm based on clipped Stochastic Gradient Descent (SGD), that works even if we only assume that the second moment of the data generating process is bounded. We derive guarantees on worst-case, finite-sample false-positive rate (FPR) over the family of all distributions with bounded second moment. Thus, our method is the first OCPD algorithm that guarantees finite-sample FPR, even if the data is high dimensional and the underlying distributions are heavy-tailed. The technical contribution of our paper is to show that clipped-SGD can estimate the mean of a random vector and simultaneously provide confidence bounds at all confidence values. We combine this robust estimate with a union bound argument and construct a sequential change-point algorithm with finite-sample FPR guarantees. We show empirically that our algorithm works well in a variety of situations, whether the underlying data are heavy-tailed, light-tailed, high dimensional or discrete. No other algorithm achieves bounded FPR theoretically or empirically, over all settings we study simultaneously.

1 INTRODUCTION

Online change-point detection (OCPD) is a fundamental problem in statistics where instantiations of a random variable are presented one after another and we want to detect

if some parameter or statistic corresponding to the underlying data generating distribution has changed. This problem has been widely studied in machine learning, mathematical statistics and information theory over the past century. In part, this is due to the wide-ranging applications of OCPD to computational biology [Muggeo and Adelfio, 2011], online advertising [Zhang et al., 2017], cyber-security [Osanaie et al., 2016, Kurt et al., 2018, Polunchenko et al., 2012], cloud-computing [Maghakian et al., 2019], finance [Lavielle and Teysiere, 2007], medical diagnostics [Yang et al., 2006, Gao et al., 2018] and robotics [Konidaris et al., 2010]. We refer interested readers to the recent surveys of [Aminikhanghahi and Cook, 2017] and [Xie et al., 2021] for details of applications of OCPD. These surveys build upon the classical texts in change-point detection obtained over the last decade [Basseville et al., 1993, Tartakovsky, 1991, Krichevsky and Trofimov, 1981].

Classical results for OCPD have focused on algorithms that assume known distributions for either one or both of the pre- and post-change data [Wald, 1992, Page, 1954, Shiryaev, 2007, Lorden, 1971, Pollak, 1985, Ritov, 1990, Moustakides, 1986, Tartakovsky, 1991]. In recent years, algorithms have been developed for cases when the pre- and post-change distributions are unknown, but belong to a parametric class such as the exponential family [Lai and Xing, 2010, Fryzlewicz, 2014, Frick et al., 2014, Cho, 2016]. Non-parametric algorithms have been developed in [Padilla et al., 2021, Madrid Padilla et al., 2021] and the references therein, but they only give asymptotic guarantees. The algorithms of [Adams and MacKay, 2007, Lai and Xing, 2010, Maillard, 2019, Alami et al., 2020] have finite-sample guarantees, but either rely on parametric assumptions such as an exponential family, or on tail assumptions such as sub-gaussian distribution families. The works of [Bhatt et al., 2022] and [Li and Yu, 2021] build upon the work in [Niu and Zhang, 2012], and give algorithms for multiple change-points with possibly heavy-tailed data in the *offline* case with all data available up-front.

In many modern applications such as cloud-computing and

* Correspondence to Abishek Sankararaman : abisanka@amazon.com

monitoring, data is known to often be heavy-tailed [Nair et al., 2022, Loiseau et al., 2010, Nizam et al., 2016] and too complex to model with any simple parametric family [Barnett and Onnela, 2016, Hallac et al., 2015, Dartmann et al., 2019]. Given the velocity, variety and volume of modern data streams, performance of change-point detection is measured through false-positive rates in order to combat alert fatigue [Ruff et al., 2021], and algorithms must work for streams that have multiple change points. Motivated by these requirements, we seek an OCPD algorithm that simultaneously meet the following desiderata : it (i) detects multiple change-points, (ii) makes no parametric assumptions on the distribution of data, (iii) works with potentially heavy-tailed data, (iv) works for high-dimensional data streams, and (v) guarantees finite sample FPR.

1.1 MAIN CONTRIBUTIONS

Our paper is the first to give an online algorithm satisfying all the 5 desiderata listed above. Specifically, our algorithm gives finite sample guarantees for FPR and detection-delay without assuming that data comes from a specific parametric family or assuming strong tail conditions - such as that the data have sub-gaussian distributions. No previous algorithm for OCPD simultaneously achieves all desiderata. Our main technical contribution is to provide a `clipped-SGD` algorithm with finite sample confidence bounds for heavy-tailed mean estimation, *that hold for all confidence values simultaneously*, a result of independent interest. We use these bounds to build a OCPD algorithm with finite sample FPR.

We further show good empirical performance across a variety of data streams with heavy-tailed, light-tailed, high dimensional or discrete distributions. However while our algorithm is designed to work across different distributions, we observe theoretically and empirically that when data has additional structure such as being one-dimensional with sub-gaussian tails or is binary, then specialized OCPD algorithms for those cases yield better results than our method. Closing these gaps is an ongoing direction of research.

2 PROBLEM SETUP

At each time t , a random vector $X_t \in \mathbb{R}^d$ is revealed to an OCPD algorithm. X_t has a probability measure and expectation denoted by \mathbb{P}_t and \mathbb{E}_t respectively, and mean $\mathbb{E}_t[X_t] \in \mathbb{R}^d$. Subsequently, using all the samples observed so far - X_1, \dots, X_t - the algorithm outputs a binary decision denoting whether a change in mean has occurred since time $t = 1$ or the last time a change was output by the algorithm, whichever is larger. The goal of the OCPD algorithm is identify the change points as quickly as possible after they occur, with bounded false-positive rate (FPR). The observed datum $(X_t)_{t \geq 1}$ are independent, although not identically distributed with piece-wise constant mean.

Definition 2.1 (Piece-wise constant mean process). Let T be the time horizon (stream-length) and let $Q_T < T$ be the total number of change-points. A set of strictly increasing time-points $1 < \tau_1 < \tau_2 \dots < \tau_{Q_T+1} := T + 1$ are called change-points, if for all $c \in \{1, \dots, Q_T\}$

- $\forall t \in [1, T], X_t \sim \mathbb{P}_t$ independently.
- $\forall t \in [\tau_c, \tau_{c+1})$, the mean $\mathbb{E}_t[X_t] := \theta_c$ of the observation is constant and does not depend on t .
- $\forall c \in [1, Q_T], \theta_c \neq \theta_{c+1}$.

Thus, a piece-wise constant mean process is identified by the quadruple $\mathfrak{M} := (T, Q_T, (\tau_c)_{c=1}^{Q_T}, (\mathbb{P}_t)_{t=1}^T)$. Throughout, we use probability and expectation operators \mathbb{P} and \mathbb{E} , to denote the joint product probability distribution $(\mathbb{P}_t)_{t=1}^T$.

2.1 ASSUMPTIONS

Let \mathcal{P} be a family of probability measures on \mathbb{R}^d such that the probability distributions \mathbb{P}_t , for all t , are from this family, i.e., $\mathbb{P}_t \in \mathcal{P}, \forall t \in [1, T]$. Throughout this paper, we make the following non-parametric assumptions on the family \mathcal{P} .

Assumption 2.2. There exists a convex compact set $\Theta \subset \mathbb{R}^d$ known to the algorithm, such that for all $\mathbb{P} \in \mathcal{P}$, $\mathbb{E}_{X \sim \mathbb{P}}[X] \in \Theta$. In words, the mean of all the distributions in the family belong to a known bounded set Θ such that $\max_{\theta_1, \theta_2 \in \Theta} \|\theta_1 - \theta_2\| := G$.

Assumption 2.3. There exists $\sigma > 0$ known to the algorithm, such that for all $\mathbb{P} \in \mathcal{P}$ and $\theta \in \Theta$, $\mathbb{E}_{X \sim \mathbb{P}}[\|X - \mathbb{E}_{X \sim \mathbb{P}}[X]\|_2^2] \leq \sigma^2$. In words, the second moment is uniformly bounded for all distributions in \mathcal{P} .

These assumptions are very general and encompass a wide range of families such as any bounded distribution, the set of sub-Gaussian distributions and heavy-tailed distributions that do not have finite higher moments. We seek algorithms that work without knowing the length of the data stream, the number of change-points and that do not make any assumptions on the underlying distributions generating the samples, beyond Assumptions 2.2 and 2.3.

2.2 PERFORMANCE MEASURES

Any OCPD algorithm is measured by two performance metrics - (i) False-positive rate and (ii) Detection delay. We set notation to define these measures.

Notation 2.4. For every $1 \leq r \leq s < T$, we denote by $X_{r:s} := (X_r, X_{r+1}, \dots, X_s)$ to be the set of observed vectors from time r to time s , with both end-points r and s inclusive.

Definition 2.5 (OCPD algorithm). A sequence of measurable functions $\mathcal{A} := (\mathcal{A}_t)_{t \geq 1}$ is called an OCPD algorithm

if for every time $t \geq 1$, $\mathcal{A}_t \in \{0, 1\}$ and is measurable with respect to the sigma algebra generated by $X_{1:t}$. The interpretation is that if $\mathcal{A}_t = 1$ for some t , then the algorithm has detected a change at time t and if $\mathcal{A}_t = 0$, no change is detected at time t .

Notation 2.6. For an OCPD algorithm \mathcal{A} and for all $t \in [T]$, denote by $R^{(\mathcal{A})}(t) \in \mathbb{N}$ to be the random variable denoting the number of detections made till time t , i.e., $R^{(\mathcal{A})}(t) = \sum_{s=1}^t \mathcal{A}_s$.

Notation 2.7. For an OCPD algorithm \mathcal{A} , and every $r \in \mathbb{N}$ and, denote by $t_r^{(\mathcal{A})}$ as the stopping time

$$t_r^{(\mathcal{A})} := \min(\inf\{t \in [0, T] \text{ s.t. } R^{(\mathcal{A})}(t) \geq r\}, T + 1),$$

where the inf of an empty set is defined to be ∞ . In words, $t_r^{(\mathcal{A})}$ is the stopping time when the OCPD algorithm detects a change for the r th time, or $T + 1$, whichever is larger.

Definition 2.8 (False Positive Detection). The r th detection of an OCPD algorithm \mathcal{A} is said to be a False Positive, if there exists no change-point between the $r - 1$ th and the r th detection. Formally, denote by the indicator (random) variable $\chi_r^{(\mathcal{A})} = \mathbf{1}\{\exists c \in [1, Q_T] \text{ s.t. } \tau_c \in (t_{r-1}^{(\mathcal{A})}, t_r^{(\mathcal{A})})\}$ to denote if the r th detection of \mathcal{A} is a false-positive. Note that by definition, on the event that $R^{(\mathcal{A})}(T) < r$, $\chi_r^{(\mathcal{A})} = 0$.

Definition 2.9 (False Positive Rate (FPR)). An OCPD algorithm \mathcal{A} is said to have false-positive rate bounded by $\delta \in (0, 1)$ if

$$\sup_{\mathfrak{M}} \mathbb{E} \left[\frac{\sum_{r=1}^T \chi_r^{(\mathcal{A})}}{R^{(\mathcal{A})}(T)} \mathbf{1}\{R^{(\mathcal{A})}(T) > 0\} \right] \leq \delta. \quad (1)$$

In words, an OCPD algorithm \mathcal{A} has bounded false positive rate, if for every piece-wise constant mean process \mathfrak{M} , the expected fraction of false-positives made by the algorithm \mathcal{A} is bounded by δ . In Equation (1), we take the sum till T because that is the maximum number of possible change points detected. If an algorithm only detects $s < T$ change points, then by definition $\chi_r^{(\mathcal{A})} = 0$ for all $r > s$.

Definition 2.10 (Worst-case Detection Delay). For $n \in \mathbb{N}$ and $\Delta > 0$, let $X_1, X_2, \dots, X_n, X_{n+1}, \dots$ be an infinite stream with the following distribution. For every $t < n$, $X_t \stackrel{\text{ind}}{\sim} \mathbb{P}_t$ with $\mathbb{E}_{X \sim \mathbb{P}_t}[X] = \theta_1 \in \Theta$ and for every $t \geq n$, $X_t \stackrel{\text{ind}}{\sim} \mathbb{P}_t$ with $\mathbb{E}_{X \sim \mathbb{P}_t}[X] = \theta_2 \in \Theta$ with $\|\theta_1 - \theta_2\| = \Delta$. Let $\mathfrak{M}^{(n, \Delta)}$ denote all such infinite piece-wise constant mean process. An algorithm \mathcal{A} is said to have worst-case detection delay $\mathcal{D}(\Delta, n, \delta')$, if

$$\sup_{\mathfrak{M}^{(n, \Delta)}} \mathbb{P} \left[\inf\{t > n: \mathcal{A}_t = 1\} - n \geq \mathcal{D}(\Delta, n, \delta') \right] \leq \delta' \quad (2)$$

holds for all $n \in \mathbb{N}$, $\Delta > 0$ and $\delta' \in (0, 1)$.

In words, the detection delay function $\mathcal{D}(\Delta, n, \delta')$ is such that for every admissible process $\mathfrak{M}^{(n, \Delta)}$ that has a single change-point at time n with jump magnitude Δ , algorithm \mathcal{A} detects the change-point before time $n + \mathcal{D}(\Delta, n, \delta')$, with probability at-least $1 - \delta'$. Note that the delay metric is measured on data streams with exactly one change-point. Defining detection delay for streams with multiple change-points is ambiguous as there could be missed detections, with only a subset of the change-points being detected [Alami et al., 2020], [Maillard, 2019]. The main question this paper studies is

For each $\delta \in (0, 1)$, does there exists an OCPD algorithm with FPR bounded by δ and having small worst-case detection-delay that only makes Assumptions 2.2 and 2.3 ?

Observe that it is trivial to achieve a FPR of 0 for example the constant function where $\mathcal{A}(\cdot) = 0$, i.e., an algorithm that never detects change-point at all. However, this algorithm has a worst-case detection-delay of ∞ , i.e., $\mathcal{D}(\Delta, n, \delta') = +\infty$ for all $\Delta > 0$, $n \in \mathbb{N}$ and $\delta' \in (0, 1)$. Thus, the challenge is to design an algorithm that satisfies the FPR constraint of δ while having small, finite worst-case detection delay, without making parametric assumptions on the underlying data generating distributions.

3 ONLINE ROBUST MEAN ESTIMATION

The central workhorse of our change-point detection algorithm is heavy-tailed online mean estimation. Suppose X_1, X_2, \dots are a sequence of independent random vectors, with the means $\mathbb{E}_t[X_t] = \theta^* \in \Theta$ being a constant independent of time t . Let $(\hat{\theta}_t)_{t \geq 1}$ be a sequence of random variables such that $\hat{\theta}_t$ is an estimate of θ based on the samples X_1, \dots, X_t defined through clipped-SGD algorithm described as follows. For a given non-negative sequence $(\eta_t)_{t \geq 1}$ and $\lambda > 0$, the estimate $\hat{\theta}_0 \in \Theta$ is arbitrary, $\hat{\theta}_t$ for each $t \geq 1$, is given by

$$\hat{\theta}_t := \prod_{\Theta} (\hat{\theta}_{t-1} - \eta_t \text{clip}(X_t - \hat{\theta}_{t-1}, \lambda)), \quad (3)$$

where, \prod_{Θ} is the projection operator onto the convex compact set Θ and for every $x \in \mathbb{R}^d$ and $\lambda > 0$,

$$\text{clip}(x, \lambda) = x \min \left(1, \frac{\lambda}{\|x\|} \right). \quad (4)$$

Our main result on the convergence of the estimator $\hat{\theta}_t$ to the true θ^* with increasing number of samples t is the following.

Theorem 3.1. *For all times $t \geq 1$, when clipped SGD in Equation (3) is run with $\lambda = 2G$ and $\eta_t = \frac{2}{m(t+\gamma)}$ for $\gamma = \max \left(120\lambda\sigma(\sigma + 1), \frac{320\sigma^2}{m} + 1 \right)$, then for every $t \geq 1$*

and every $\delta \in (0, 1)$,

$$\mathbb{P} \left[\|\hat{\theta}_t - \theta^*\|_2^2 \geq \mathcal{B}(t, \delta) \right] \leq \frac{\delta}{t(t+1)},$$

where

$$\mathcal{B}(t, \delta) := C_t \left[\frac{\gamma^2 G^2}{(t+1)^2} + \left(\frac{16\sigma^2}{\lambda} + 4\sigma^2 \right) \frac{1}{2m^2(t+1)} + \frac{96\lambda^2 \ln \left(\frac{2t^2(t+1)}{\delta} \right) \sigma(\sigma+1)}{m(t+\gamma)\sqrt{t+1}} \right], \quad (5)$$

$$\text{and } C_t = \max \left(\frac{1024\sigma^4}{G^2 m^2 \lambda^2}, \frac{8\lambda \sqrt{\ln \left(\frac{2t^2(t+1)}{\delta} \right)}}{\gamma^2 G} \right).$$

Corollary 3.2. *There exists an universal constant $A > 0$ such that for all $t \geq 1$, when clipped SGD in Equation (3) is run with parameters in Theorem 3.1*

$$\mathbb{P} \left[\|\hat{\theta}_t - \theta^*\| \geq A \max \left(\frac{\sigma^3}{\sqrt{t}}, \frac{\sigma \sqrt{\ln \left(\frac{t^3}{\delta} \right)}}{\sqrt{t}} \right) \right] \leq \frac{\delta}{t(t+1)},$$

holds for every $\delta \in (0, 1)$.

Proof is in Appendix in Section B and uses tools from [Bubeck, 2015], [Gorbunov et al., 2020, Tsai et al., 2022] and [Victor, 1999].

Remark 3.3. Compared to [Tsai et al., 2022], we do not need the failure probability δ in the input and we can give simultaneous confidence intervals for all failure probabilities δ . In contrast, the algorithm of [Tsai et al., 2022] requires $\delta \in (0, 1)$ as an input and only guarantees that the estimate mean is close to the true mean, upto error probability of δ . However, the bound in Theorem 3.1 is off by logarithmic factors compared to [Tsai et al., 2022]. Concretely, $C_t = O(1)$ for the algorithm of [Tsai et al., 2022], while it is $O(\log(t/\delta))$ for us. This is the price to have confidence intervals hold for all failure probabilities simultaneously as opposed to just having one single failure probability.

Remark 3.4. Compared to the setting of [Tsai et al., 2022], our setting is *weaker* as we assume that the domain Θ is compact with finite diameter G . This is what enables us to use an appropriately tuned learning rate and clipping parameter to make the algorithm any-time and obtain confidence intervals at all failure probabilities simultaneously. It is an open question whether the assumption that Θ is compact can be relaxed and if we can still make guarantees confidence interval holding for all failure probabilities δ for all t for heavy tailed distributions.

Remark 3.5. The constants in Theorem 3.1 are not optimal. In Section 5, we suggest an alternative set of constants that work well empirically across variety of settings.

Remark 3.6. There have been significant recent advances in robust mean estimation [Diakonikolas et al., 2020, Lugosi and Mendelson, 2021, Depersin and Lecué, 2022, Cherapanamjeri et al., 2020, Diakonikolas et al., 2022], that are known to provide near optimal error bounds. However, unlike our method, none of these algorithms can give confidence bounds for all confidence values simultaneously.

Remark 3.7. Theorem 4.3 in [Devroye et al., 2016] proves that it is impossible to get a finite sample confidence bound to hold for all $\delta \in (0, 1)$. Our result does not contradict this since the restriction on the allowable δ is *implicit* in Theorem 3.1. Equation (5) gives that, for every $t \in \mathbb{N}$, as $\delta \searrow 0$, $\mathcal{B}(t, \delta) \nearrow \infty$. However, from Assumption 2.2, if $\mathcal{B}(t, \delta) \geq G$, then the statement of Theorem 3.1 is vacuous. Thus, Theorem 3.1 gives a non-vacuous bounds only for $\delta \in (\delta_{min}^{(t)}, 1)$ where $\delta_{min}^{(t)} := \inf_{\delta > 0} \{\mathcal{B}(t, \delta) < G\}$.

3.1 UNIFORM OVER TIME BOUND

As a corollary of Theorem 3.1, we get the following bound that holds uniformly over all time.

Corollary 3.8. *There exists an universal constant $A > 0$ such that, when clipped SGD in Equation (3) is run with parameters in Theorem 3.1,*

$$\mathbb{P} \left[\exists t \in \mathbb{N} : \|\hat{\theta}_t - \theta^*\| \geq A \max \left(\frac{\sigma^3}{\sqrt{t}}, \frac{\sigma \sqrt{\ln \left(\frac{t^3}{\delta} \right)}}{\sqrt{t}} \right) \right] \leq \delta,$$

holds for every $\delta \in (0, 1)$.

The proof follows by taking an union bound over all $t \geq 1$, i.e., summing over $t \geq 1$ on both the LHS and RHS of Corollary 3.8 and noticing that $\sum_{t \geq 1} \frac{1}{t(t+1)} = 1$. The bounds in Theorem 3.1 and Corollary 3.8 are *dimension free*, i.e., the term d does not appear in the bounds. The moment bound σ plays the role of dimension. In particular, suppose that all distributions in the family \mathcal{P} have covariance matrices bounded in the positive semi-definite sense by $\Sigma \in \mathbb{R}^{d \times d}$. In this case, by definition, $\sigma^2 \leq \text{Trace}(\Sigma)$ and plays the role of dimension.

In the special case when the samples $(X_t)_{t \geq 1}$ are i.i.d. with sub-gaussian distributions with mean θ^* and covariance matrix Σ , [Abbasi-yadkori et al., 2011, Maillard, 2019, Chowdhury et al., 2022] show that for all $\delta \in (0, 1)$,

$$\mathbb{P} \left[\exists t \in \mathbb{N} : \left\| \frac{1}{t} \sum_{s=1}^t X_s - \theta^* \right\| \geq \sqrt{2\lambda_{max}(\Sigma) \left(1 + \frac{1}{t} \right) \ln \left(\frac{(t+1)^d}{\delta} \right)} \right] \leq \delta, \quad (6)$$

holds, where $\lambda_{max}(\Sigma)$ is the highest eigen-value of the covariance matrix Σ . Thus, for the special case of sub-gaussian

distributions, Equation (6) has a better dependence on time t compared to our Corollary 3.8. The improved dependence on time arises as Equation (6) is based on the construction of a self normalized martingale and using the martingale stopping theorem to obtain uniform over time bounds while Corollary 3.8 is based on a simple union bound.

However, the bound in Equation (6) is un-satisfactory for high dimensional problems since it is not dimension free and depends on the scale of the problem through the term $d\lambda_{max}(\Sigma)$ which by definition is larger than $\text{Trace}(\Sigma)$. In many high dimensional settings, $d\lambda_{max}(\Sigma)$ is much larger than $\text{Trace}(\Sigma)$ and thus algorithms and bounds depending explicitly on d is un-desirable [Wainwright, 2019, Lugosi and Mendelson, 2019]. It is an open problem whether a bound can be established that is both dimension free and has the improved time-dependence of the kind in Equation (6), even for sub-gaussian random vectors. For heavy-tailed distributions, [Wang and Ramdas, 2022, 2023] establish confidence bounds with sharp dependence on time by extending the martingale recipe followed by [Maillard, 2019]. However, it is unclear if trivially extending their approach to high dimensions can yield dimension free bounds.

4 CHANGE-POINT DETECTION ALGORITHM

Our algorithm is described in Algorithm 1 and is based on the following idea. A change point is detected in the time-interval $[r, t]$ if there exists $r < s < t$ such that confidence interval around the estimated mean of the observations $X_{r:s}$ is separated from the confidence interval around the estimated mean of the observations $X_{s+1:t}$. Further, in order to accommodate multiple change-points, the algorithm *restarts* after every change detection, similar to [Alami et al., 2020]. It is known that standard empirical mean is a poor estimator when the underlying distributions can potentially be heavy-tailed, as its confidence interval under only assumptions in 2.3 is wide [Lugosi and Mendelson, 2019]. To attain better confidence intervals, we use the `clipped-SGD` Algorithm 1 that gives a confidence interval for the estimated mean for every failure probability $\delta \in (0, 1)$ simultaneously. Having multiple confidence intervals is crucial as we show that adaptively testing different intervals of times at different carefully chosen confidence intervals (Line 8 of Algorithm 1) leads to the bounded FPR guarantee.

4.1 CONNECTIONS TO GLR

Restating our algorithm, a change point is detected in a time-interval $[t_0, t]$ if

$$\exists s \in (t_0, t) \text{ s.t. } \|\hat{\theta}_{t_0:s} - \hat{\theta}_{s+1:t}\|^2 \geq \mathcal{C}(t_0, s, t, \delta),$$

where the function $\mathcal{C}(\cdot)$ is given in Line 8 of Algorithm 1. In the above re-statement, the estimates $\hat{\theta}_{t_0:s}$ and $\hat{\theta}_{s+1:t}$

are robust estimates of the mean based on the set of observations $\{X_{t_0}, \dots, X_s\}$ and $\{X_{s+1}, \dots, X_t\}$ respectively. The `IMPROVED-GLR` of [Maillard, 2019] uses a detector that is structurally similar to the above equation except that they (i) use the empirical mean as they are dealing with sub-gaussian random variables, and (ii) use a function $\mathcal{C}(\cdot)$ derived from the Laplace method that gives confidence bounds with better dependence on time, but is not dimension free. In contrast, we use the robust mean estimator given by clipped-SGD and the function $\mathcal{C}(\cdot)$ is derived from the confidence guarantees that only require the existence of the second moment and make no other tail assumptions and yields dimension free bounds. The cost however is that the confidence bound derived from clipped SGD has a weaker dependence on time compared to that obtained by the Laplace’s method [Maillard, 2019].

4.2 FALSE-POSITIVE GUARANTEE

We will prove the following result on Algorithm 1. For a given process \mathfrak{M} , and every $r \in \mathbb{N}$, denote by the deterministic time $\tau_c^{(r)} := \inf\{\tau_c : \tau_c > r\}$ be the first change-point after time r .

Theorem 4.1 (False Positives). *When Algorithm 1 is run with parameters $\lambda = 2G$, $\eta_t = \frac{2}{m(t+\gamma)}$ for $\gamma = \max\left(120\lambda\sigma(\sigma+1), \frac{320\sigma^2}{m} + 1\right)$ and $\delta \in (0, 1)$,*

$$\sup_{\mathfrak{M}, r} \mathbb{P}[\exists t \in [r, \tau_c^{(r)}], \text{ s.t. } \mathcal{A}_t = 1 | \mathcal{A}_r = 1] \leq \delta,$$

holds almost-surely.

Proof is in Appendix in Section C.1. This result states that with probability at-most δ , a true change-point *does not* lie between any two consecutive detections made by the algorithm. This theorem implies the following lemma.

Lemma 4.2. *Under the conditions of Theorem 4.1, the FPR condition in Equation (1) holds.*

The proof is in the Appendix in Section C.2. We emphasize that the guarantee in 4.2 is a *worst-case guarantee*. In other words, no matter the underlying distribution, as long as Assumptions 2.2 and 2.3 are met, Algorithm 1 will not have more than a δ fraction of false-positives.

4.3 WORST-CASE DETECTION DELAY GUARANTEE

Lemma 4.3. *If Algorithm 1 is run with the parameters from Theorem 4.1, then for every $n \in \mathbb{N}$, $\Delta > 0$ and $\delta' \in (0, 1)$*

$$\mathcal{D}(n, \Delta, \delta') \geq \inf \left\{ d \in \mathbb{N} : \Delta^2 \geq \mathcal{B}\left(n - 1, \frac{\delta'}{2}\right) + \right.$$

Algorithm 1 Online Clipped-SGD Change Point Detection

1: **Input:** $(\eta_t)_{t \geq 1}$, $\lambda > 0$, $\theta_0 \in \Theta$, $\delta \in (0, 1)$ the FPR guarantee
 2: $r \leftarrow 1$
 3: $\hat{\theta}_{t,t-1} \leftarrow \theta_0$, for all $t \geq 1$.
 4: Set Num-change-points $\leftarrow 0$
 5: **for** each time $t = 1, 2, \dots$, **do**
 6: Receive sample X_t
 7: $\hat{\theta}_{s,t} \leftarrow \prod_{\theta}(\hat{\theta}_{s,t-1} - \eta_{t-s} \text{clip}(X_t - \hat{\theta}_{s,t-1}, \lambda))$, for every $r \leq s \leq t$.
 8: **if** $\exists s \in (r, t)$ such that $\|\hat{\theta}_{r:s} - \hat{\theta}_{s+1:t}\|_2^2 > \mathcal{B}\left(s-r, \frac{\delta}{2(t-r)(t-r+1)}\right) + \mathcal{B}\left(t-s-1, \frac{\delta}{2(t-r)(t-r+1)}\right)$ $\{B(\cdot, \cdot)$ is defined in Equation (5) **then**
 9: Set $\mathcal{A}_t \leftarrow 1$ {Change point detected}
 10: $r \leftarrow t + 1$
 11: Set Num-change-points \leftarrow Num-change-points + 1 {Increment number of change-points detected}
 12: **else**
 13: Set $\mathcal{A}_t \leftarrow 0$
 14: **end if**
 15: **end for**




Figure 1: Figure (a) plots the heat-map of $\mathcal{D}(n, \Delta, \delta')$ from Lemma 4.3 for fixed $\delta' = 0.1$. The white cells represent infinity. Figure (b) plots the 90th quantile ($\delta' = 0.1$) of the observed delay for Pareto distribution $d = 32$ over 30 runs. As can be seen, the observed detection delay in (b) is much smaller than the worst case delay in (a).

$$\mathcal{B}\left(d, \frac{\delta'}{2}\right) + \mathcal{B}\left(n-1, \frac{\delta}{2(n+d+1)(n+d)}\right) + \mathcal{B}\left(d, \frac{\delta}{2(n+d+1)(n+d)}\right), \quad (7)$$

where $\mathcal{D}(\cdot)$ and $\mathcal{B}(\cdot)$ are in Eqns (2) and (5) respectively.

Proof is in the Appendix in Section D. Lemma 4.3 is an *upper bound on the worst case delay*. In other words, for any pre- and post-change distribution with norm of the means differing by Δ , Algorithm 1 will detect this change within delay of $\mathcal{D}(n, \Delta, \delta')$ with probability at-least $1 - \delta'$.

For many specific choices of pre- and post-change distribution families however, we expect the observed detection delay to be much smaller than predicted by Lemma 4.3. This bound is conservative as it is worst-case over all distributions. In Figure 1a we plot the bound in Lemma 4.3

for a fixed $\delta' = \delta = 0.1$ as n and Δ varies. We use the constants given in Section 5.1 to plot Figure 1a. In Figure 1b, we plot the empirically observed detection delay for a sequence of 32 dimensional Pareto distributed random vectors with shape parameter 2.01. As can be seen in Figure 1, the observed detection delay is much smaller than that indicated by Lemma 4.3, which is a worst case over all distributions.

Remark 4.4. In the special case when the observations are Bernoulli random variables, the R-BOCPD algorithm of [Alami et al., 2020] gives a smaller detection delay compared to ours – our detection delay bound in 4.3 has additional poly-logarithmic factors of $\log(n/\delta)$ and sub-optimal constants compared to R-BOCPD. However, our bound holds for *any* family of distributions, including high-dimensional and heavy tailed ones, while R-BOCPD can only be applied for Bernoulli distributions.

Corollary 4.5 (Un-detectable Change). *If $\Delta \leq \mathcal{O}\left(\frac{\log(\frac{n}{\delta})}{\sqrt{n}}\right)$, then $\mathcal{D}(n, \Delta, \delta') = \infty$ for all $\delta' \in (0, 1)$, i.e., the change is un-detectable by Algorithm 1.*

Remark 4.6. The undetectable region consists of the grey/white areas of Figure 1a. However, since Lemma 4.3 is only an upper-bound, the fact that $\mathcal{D}(n, \Delta, \delta') = \infty$ does *not imply* that our algorithm cannot detect the change (cf. Figure 1b).

Remark 4.7. In the case of sub-gaussian, exponential families, [Maillard, 2019] give a lower bound for changes that not detectable by *any* algorithm. When Algorithm 1 is applied to sub-gaussian random variables from an exponential family, the detection-delay bound in Lemma 4.3 is sub-optimal by poly-logarithmic factors in $\log(n/\delta)$ compared to the lower bound. However, Algorithm 1 and the delay bound in Lemma 4.3 holds for any class of distributions subject to Assumptions 2.3 and 2.2, while the bounds in

[Maillard, 2019] only applies to sub-gaussian observations from a known exponential family.

4.4 CHANGE-POINT LOCALIZATION

In practice, it is also crucial to identify the location where the change point occurred. In this section we describe how to modify Algorithm 1 to also output the estimate of the location of change in addition to just detecting the existence of a change. Recall that for every $r \in \mathbb{N}$, $\tau_r^{(A)} \in \mathbb{N} \cup \{\infty\}$ is the stopping time denoting the r th time, Algorithm \mathcal{A} detects a change point. We modify Algorithm 1 by additionally outputting for every $r \in \mathbb{N}$, a time interval $[s_{r;1}^{(A)}, s_{2;r}^{(A)}] \subseteq [\tau_{r-1}^{(A)}, \tau_r^{(A)}]$ such that this is an interval that contains a change-point τ_c .

In order to do so, we need an additional definition. For every $r < s < t$ and $\delta \in (0, 1)$, denote by $\mathfrak{B}(r, s, t, \delta) \in \{0, 1\}$ as the indicator variable that

$$\mathfrak{B}(r, s, t, \delta) = \mathbf{1}\left(\|\hat{\theta}_{r:s} - \hat{\theta}_{s+1:t}\|_2^2 > \mathfrak{B}_1 + \mathfrak{B}_2\right), \quad (8)$$

where $\mathfrak{B}_1 = \mathcal{B}\left(s - r, \frac{\delta}{2(t-r)(t-r+1)}\right)$ and $\mathfrak{B}_2 = \mathcal{B}\left(t - s - 1, \frac{\delta}{2(t-r)(t-r+1)}\right)$. The estimates of the location of change in a time-interval $[r, t]$ is all those time instants $s \in [r, t]$ such that $\mathfrak{B}(r, s, t, \delta) = 1$. Line 12 in Algorithm 1 in Section A in the Appendix, precisely defines the estimator. The empirical performance of this method is shown in Figure 3. We observe that this produces an accurate and sharp estimate of the change-point location in simulations.

5 EXPERIMENTS

In this section we give numerical evidence to show that Algorithm 1 can be applied across variety of settings. Line 8 of Algorithm 1 relies on confidence bounds for high-dimensional estimation where the global constants are not optimized. This is an artifact of the proof analysis in robust estimation [Lugosi and Mendelson, 2019, Vershynin, 2018]. Thus, we modify the absolute constants used in Theorem 4.1 as follows. We use $\gamma = \max\left(4\lambda\sigma(\sigma + 1), \frac{8\sigma^2}{m} + 1\right)$ with the color red highlighting the changes from the definition in Theorem 4.1. The constant C_t is modified as follows

$C_t = \max\left(\frac{0.5\sigma^4}{G^2 m^2 \lambda^2}, \frac{1\lambda\sqrt{\ln\left(\frac{2t^2(t+1)}{\delta}\right)}}{\gamma^2 G}\right)$. In addition, we use the following definition of $\mathcal{B}(\cdot, \cdot)$

$$\mathcal{B}(t, \delta) := C_t \left[\frac{\gamma^2 G^2}{t+1} + \left(\frac{2\sigma^2}{\lambda} + 1\sigma^2 \right) \frac{1}{2m^2(t+1)} + \frac{2\lambda^2 \ln\left(\frac{2t^2(t+1)}{\delta}\right) \sigma(\sigma + 1)}{m(t+\gamma)\sqrt{t+1}} \right], \quad (9)$$

where C_t and γ are the modified values stated above. Further, in all simulations we assume $\Theta = \mathbb{R}^d$ to be the whole plane.

5.1 SYNTHETIC SIMULATIONS

Here, demonstrate that Algorithm 1 with choice of hyper-parameters in Equation (9) is practical and can be applied across a variety of data generating distributions – either heavy-tailed, or high-dimensional or both and still obtains bounded false-positive rates and a much lower detection delay compared to what the conservative bound in Lemma 4.3 would indicate.

5.1.1 Setup

In Figure 2, we construct synthetic situations and introduce change-points with each change lasting 400 time-units. In all experiments, we choose the family of distributions \mathfrak{M} such that $\sigma = 1$, $G = 12$. At each time t , a sample is drawn from the appropriate distribution that we detail below and presented to the change-point algorithm. The true-change points and the median detection times along with the 95 percentile upper and lower confidence bands are show in Figure 2. These are estimated by averaging 30 independent runs for each setting in Figure 2.

Heavy-tailed distribution: In Figures 2a, 2b and 2i, the sample at every time-point is drawn from a Pareto distribution with shape-parameter 2.01. This implies that the third central moment of the distribution is infinity. The mean of the samples in the time-durations $t \in [0, 400) \cup [800, 1200)$ is 0 in all figures and the mean at times $t \in [400, 800) \cup [1200, 1600)$ is $\Delta = 0.5, 1, 1$ respectively in Figures 2a, 2b and 2i. In Figure 2c, 2d and 2j, we consider the observation at time t to be 32 dimensional isotropic random vector with norm having Pareto distributions with shape parameter 2.01. The mean vector at times $[0, 400) \cup [800, 1200)$ is $0 \in \mathbb{R}^{32}$ and at times $t \in [400, 800) \cup [1200, 1600)$ is $\frac{\Delta}{\sqrt{32}}[1, \dots, 1] \in \mathbb{R}^{32}$, where $\Delta = 0.5, 1, 1$ respectively in Figures 2c, 2d and 2j respectively.

Gaussian distribution: In Figures 2e and 2f the sample at every time-point is drawn from a unit variance Gaussian distribution. The mean of the samples in the time-durations $t \in [0, 400) \cup [800, 1200)$ in all three figures is 0 and the mean at times $t \in [400, 800) \cup [1200, 1600)$ in the two figures 2e and 2f are $\Delta = 0.5$ and $\Delta = 1$ respectively. In Figures 2g and 2h we consider the observation at time t to be 32 dimensional isotropic gaussian random vector with co-variance on each axis being $1/\sqrt{32}$. The mean vector at times $[0, 400) \cup [800, 1200)$ is $0 \in \mathbb{R}^{32}$ and at times $t \in [400, 1600) \cup [1200, 1600)$ is $\frac{\Delta}{\sqrt{32}}[1, \dots, 1] \in \mathbb{R}^{32}$.

Bernoulli distribution: In Figures 2k and 2l, the data was $\{0, 1\}$ valued Bernoulli random variable with means at times $[0, 400) \cup [800, 1200)$ was 0.7 and 0.85 respectively in the two figures, and the means at times $[400, 800) \cup [1200, 1600)$ are 0.3, 0.15 respectively in the two figures.

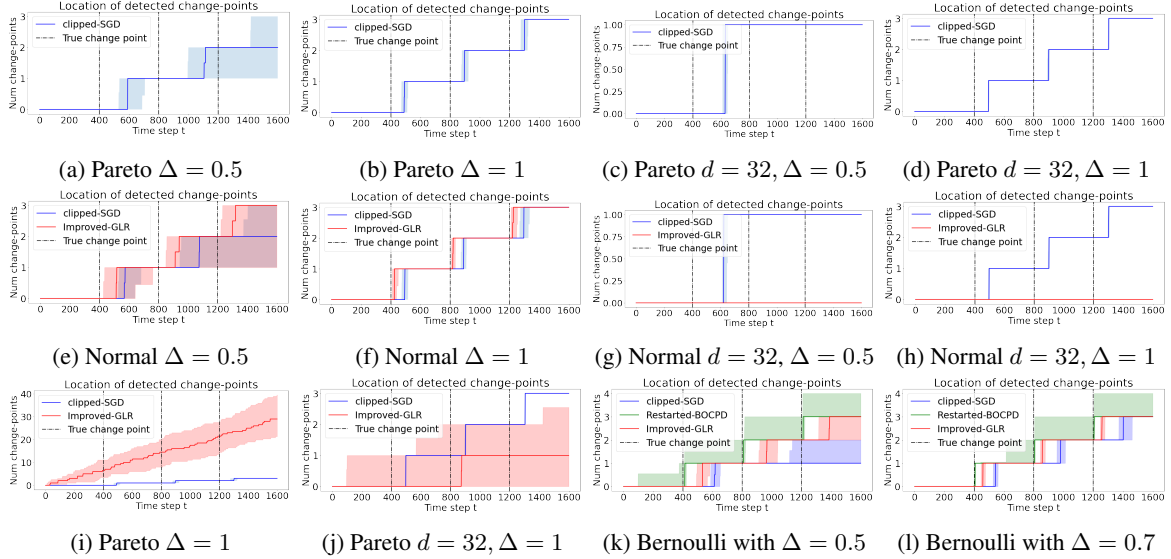


Figure 2: Empirical performance of Algorithm 1 in a variety of scenarios. Exact details of each plot in Section 5.1.




Figure 3: Plots showing that by Algorithm 1 can detect and localize change-points across a variety of settings.

5.1.2 Baselines

We consider two OCPD algorithms, the `Improved-GLR` of [Maillard, 2019] and `R-BOCPD` of [Alami et al., 2020]. We chose these two as baselines since they have been empirically demonstrated to be state-of-art, and are the only other algorithms to possess finite sample, non-asymptotic FPR guarantees. The `Improved-GLR` can be applied to any distribution, while its theoretical guarantees only hold for sub-gaussian distributions. The `R-BOCPD` algorithm is only applicable to detecting change-points in binary data and thus we only apply it to Bernoulli distributed data.

5.1.3 Results

We observe in Figure 2 that our algorithm is the only one to attain bounded FPR across heavy-tailed, Gaussian, high dimensional and Bernoulli distribution.

We can observe from Figures 2h and 2j that under Pareto distribution, the `Improved-GLR` algorithm has a large number of False Positive detections. Intuitively this occurs because the `Improved-GLR` algorithm assumes sub-gaussian tails and thus large deviations that are typical for the heavy-tailed Pareto distributions are mistaken for a

change. (See also Figure 2). In contrast, from Figures 2a, 2b, 2c, 2d and 2j, we see that our algorithm consistently attains bounded false-positive rates and finite detection delay guarantees across choices of Δ and dimension d .

On gaussian distributed data, both our algorithm 1 and the `Improved-GLR` obtains similar performance in-terms of false-positive rates. However, the detection delay of GLR is much smaller compared to ours, as seen by the fact that the median detection time of our algorithm is larger than the 95th percentile detection time of `Improved-GLR`. In Bernoulli distributed data, all methods attain similar False-positive guarantees – however, the specialized algorithm of `R-BOCPD` is superior in terms of detection delay compared to ours and the `Improved-GLR`.

In Table 1, we summarize Figure 2 by measuring *regret*. For any OCPD algorithm \mathcal{A} , we can define a function $R^{(\mathcal{A})} : [T] \rightarrow \mathbb{N}$ where $R^{(\mathcal{A})}(t) = \sum_{s \leq t} \mathcal{A}_s$ is the total number of change-points detected upto time t . Similarly, for any $t \in [T]$, the ground-truth function $R^*(t) = \max\{c : \tau_c \leq t\}$ is the number of true changes till time t . The regret of algorithm \mathcal{A} is defined as $\sum_{t=1}^T |R^{(\mathcal{A})}(t) - R^*(t)|$. This measure is non-negative and is 0 if and only if the output of the algorithm matches the ground truth. In Table 1, we

Distribution	d	Δ	Algorithm 1	Improved GLR [Maillard, 2019]	R-BOCPD [Alami et al., 2020]
Normal	1	1	274 \pm 38	64 \pm 45	N/A
	32	1	300 \pm 6	2400 \pm 0	
	1	0.5	694 \pm 191	356 \pm 150	
	32	0.5	1427 \pm 14	2400 \pm 1	
Pareto	1	1	296 \pm 35	19913 \pm 8143	N/A
	32	1	302 \pm 7	1616 \pm 921	
	1	0.5	868 \pm 365	1891 \pm 663	
	32	0.5	1431 \pm 14	1667 \pm 653	
Bernoulli	-	0.7	515 \pm 49	181 \pm 23	23 \pm 479
	-	0.5	1509 \pm 53	1466 \pm 762	63 \pm 380

Table 1: Quantitative summary of Figure 2 by comparing regret, where lower is better. Our method achieves lower regret across variety of settings of distribution, dimension and change magnitude.




Figure 4: Performance of change-point detection of Algorithm 1 and the Improved-GLR on real data.

give the median value of regret along with 95% confidence interval. We observe in Table 1 that our method achieves lower regret across a variety of situations - whether the data is heavy-tailed, light tailed high dimensional or discrete.

5.1.4 Change-point localization

In Figure 3, we demonstrate sharpness of change-point localization (detailed in Algorithm 1). The setting in Figure 3 is identical to that of Figure 2 with the boundary of the shaded region representing the 5th quantile for the starting point and the 95th quantile for the ending point of the change location interval output in Line 12 of Algorithm 1. The localization region is biased towards the right, which is expected since our algorithm is designed to minimize false positives even in the worst-case. Mathematically proving that the localization is sharp is an interesting future work.

5.2 REAL-DATA

In Figure 4 we show the performance of Algorithm 1 and the Improved-GLR on the well-log dataset [Ó Ruanaidh

and Fitzgerald, 1996]. This dataset consists of 4050 measurements in the range $[6 \times 10^4, 10^5]$ of nuclear-magnetic-response taken during drilling of a well. The data are used to interpret the geophysical structure of the rock surrounding the well. The variations in mean reflect the stratification of the earth’s crust. We process the data by dividing it by $10^{4.5}$ and run Algorithm 1 with $G = 10$, $\sigma = 1$ and Improved GLR with $\sigma = 1$. The detected change-points are shown in Figure 4. Figure 4 shows that the performance of Algorithm 1 is at-least as good as that of Improved-GLR in terms of false-positive detections.

6 CONCLUSIONS

We introduced a new method based on clipped-SGD, to detect change-points with guaranteed finite-sample FPR, without parametric or tail assumptions. The key technical contribution is to give an anytime online mean estimation algorithm, that provides a confidence bound for the mean at all confidence levels simultaneously. We also give a finite-sample, high probability bound on the detection delay as a function of the gap between the means and number of pre-change observations. We further corroborate empirically that ours is the only algorithm to detect change-points with bounded FPR, across multi-dimensional heavy tailed, gaussian or binary-valued data streams.

Our work opens several interesting directions for future investigation. Obtaining sharp confidence intervals for estimating the mean of a random vector without the existence of variance was shown recently in [Cherapanamjeri et al., 2022]. Extending the tools from therein to further relax the second moment assumption we considered is a natural direction of future work. Another open question is to see if the Laplace method used in [Maillard, 2019] can be extended to the high-dimensional to get dimension free confidence bounds. Further, we observe in simulations that our method attains ‘sharp’ localization empirically. Understanding the three-way trade-off between sharpness of localization, FPR and detection delay is an important area of future work.

REFERENCES

- Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011. URL https://proceedings.neurips.cc/paper_files/paper/2011/file/e1d5be1c7f2f456670de3d53c7b54f4a-Paper.pdf.
- Ryan Prescott Adams and David JC MacKay. Bayesian online changepoint detection. *arXiv preprint arXiv:0710.3742*, 2007.
- Réda Alami, Odalric Maillard, and Raphael Féraud. Restarted bayesian online change-point detector achieves optimal detection delay. In *International conference on machine learning*, pages 211–221. PMLR, 2020.
- 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.
- Ian Barnett and Jukka-Pekka Onnela. Change point detection in correlation networks. *Scientific reports*, 6(1):18893, 2016.
- Michele Basseville, Igor V Nikiforov, et al. *Detection of abrupt changes: theory and application*, volume 104. prentice Hall Englewood Cliffs, 1993.
- Sujay Bhatt, Guanhua Fang, and Ping Li. Offline change detection under contamination. In *Uncertainty in Artificial Intelligence*, pages 191–201. PMLR, 2022.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Yeshwanth Cherapanamjeri, Samuel B Hopkins, Tarun Kathuria, Prasad Raghavendra, and Nilesh Tripuraneni. Algorithms for heavy-tailed statistics: Regression, covariance estimation, and beyond. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 601–609, 2020.
- Yeshwanth Cherapanamjeri, Nilesh Tripuraneni, Peter Bartlett, and Michael Jordan. Optimal mean estimation without a variance. In *Conference on Learning Theory*, pages 356–357. PMLR, 2022.
- Haeran Cho. Change-point detection in panel data via double cusum statistic. 2016.
- Sayak Ray Chowdhury, Patrick Saux, Odalric-Ambrym Maillard, and Aditya Gopalan. Bregman deviations of generic exponential families. *arXiv preprint arXiv:2201.07306*, 2022.
- Guido Dartmann, Houbing Song, and Anke Schmeink. *Big data analytics for cyber-physical systems: machine learning for the internet of things*. Elsevier, 2019.
- Jules Depersin and Guillaume Lecué. Robust sub-gaussian estimation of a mean vector in nearly linear time. *The Annals of Statistics*, 50(1):511–536, 2022.
- Luc Devroye, Matthieu Lerasle, Gabor Lugosi, and Roberto I Oliveira. Sub-gaussian mean estimators. 2016.
- Ilias Diakonikolas, Daniel M Kane, and Ankit Pensia. Outlier robust mean estimation with subgaussian rates via stability. *Advances in Neural Information Processing Systems*, 33:1830–1840, 2020.
- Ilias Diakonikolas, Daniel M Kane, Ankit Pensia, and Thanasis Pittas. Streaming algorithms for high-dimensional robust statistics. In *International Conference on Machine Learning*, pages 5061–5117. PMLR, 2022.
- Klaus Frick, Axel Munk, and Hannes Sieling. Multiscale change point inference. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, pages 495–580, 2014.
- Piotr Fryzlewicz. Wild binary segmentation for multiple change-point detection. 2014.
- Zhen Gao, Guoliang Lu, Peng Yan, Chen Lyu, Xueyong Li, Wei Shang, Zhaohong Xie, and Wanming Zhang. Automatic change detection for real-time monitoring of eeg signals. *Frontiers in physiology*, 9:325, 2018.
- Eduard Gorbunov, Marina Danilova, and Alexander Gasnikov. Stochastic optimization with heavy-tailed noise via accelerated gradient clipping. *Advances in Neural Information Processing Systems*, 33:15042–15053, 2020.
- David Hallac, Jure Leskovec, and Stephen Boyd. Network lasso: Clustering and optimization in large graphs. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, pages 387–396, 2015.
- George Konidaris, Scott Kuindersma, Roderic Grupen, and Andrew Barto. Constructing skill trees for reinforcement learning agents from demonstration trajectories. *Advances in neural information processing systems*, 23, 2010.
- Raphail Krichevsky and Victor Trofimov. The performance of universal encoding. *IEEE Transactions on Information Theory*, 27(2):199–207, 1981.
- Bariş Kurt, Çağatay Yıldız, Taha Yusuf Ceritli, Bülent Sankur, and Ali Taylan Cemgil. A bayesian change point model for detecting sip-based ddos attacks. *Digital Signal Processing*, 77:48–62, 2018.

- Tze Leung Lai and Haipeng Xing. Sequential change-point detection when the pre-and post-change parameters are unknown. *Sequential analysis*, 29(2):162–175, 2010.
- Marc Lavielle and Gilles Teyssiere. Adaptive detection of multiple change-points in asset price volatility. *Long memory in economics*, pages 129–156, 2007.
- Mengchu Li and Yi Yu. Adversarially robust change point detection. *Advances in Neural Information Processing Systems*, 34:22955–22967, 2021.
- Patrick Loiseau, Paulo Gonçalves, Guillaume Dewaele, Pierre Borgnat, Patrice Abry, and Pascale Vicat-Blanc Primet. Investigating self-similarity and heavy-tailed distributions on a large-scale experimental facility. *IEEE/ACM Transactions on Networking*, 18(4):1261–1274, 2010.
- Gary Lorden. Procedures for reacting to a change in distribution. *The annals of mathematical statistics*, pages 1897–1908, 1971.
- Gábor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190, 2019.
- Gabor Lugosi and Shahar Mendelson. Robust multivariate mean estimation: the optimality of trimmed mean. 2021.
- Oscar Hernan Madrid Padilla, Yi Yu, Daren Wang, and Alessandro Rinaldo. Optimal nonparametric change point analysis. 2021.
- Jessica Maghakian, Joshua Comden, and Zhenhua Liu. On-line optimization in the non-stationary cloud: Change point detection for resource provisioning. In *2019 53rd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2019.
- Odalric-Ambrym Maillard. Sequential change-point detection: Laplace concentration of scan statistics and non-asymptotic delay bounds. In *Algorithmic Learning Theory*, pages 610–632. PMLR, 2019.
- George V Moustakides. Optimal stopping times for detecting changes in distributions. *the Annals of Statistics*, 14(4):1379–1387, 1986.
- Vito MR Muggeo and Giada Adelfio. Efficient change point detection for genomic sequences of continuous measurements. *Bioinformatics*, 27(2):161–166, 2011.
- Jayakrishnan Nair, Adam Wierman, and Bert Zwart. *The fundamentals of heavy tails: Properties, emergence, and estimation*, volume 53. Cambridge University Press, 2022.
- Yue S Niu and Heping Zhang. The screening and ranking algorithm to detect dna copy number variations. *The annals of applied statistics*, 6(3):1306, 2012.
- Farhana Nizam, Shudarshon Chaki, Shamim Al Mamun, M Shamim Kaiser, et al. Attack detection and prevention in the cyber physical system. In *2016 International Conference on Computer Communication and Informatics (ICCCI)*, pages 1–6. IEEE, 2016.
- J. J. K. Ó Ruanaidh and W. J. Fitzgerald. *Numerical Bayesian Methods Applied to Signal Processing*. Springer, 1996.
- Opeyemi Osanaiye, Kim-Kwang Raymond Choo, and Mqhele Dlodlo. Change-point cloud ddos detection using packet inter-arrival time. In *2016 8th Computer Science and Electronic Engineering (CEECE)*, pages 204–209. IEEE, 2016.
- Oscar Hernan Madrid Padilla, Yi Yu, Daren Wang, and Alessandro Rinaldo. Optimal nonparametric multivariate change point detection and localization. *IEEE Transactions on Information Theory*, 68(3):1922–1944, 2021.
- Ewan S Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, 1954.
- Moshe Pollak. Optimal detection of a change in distribution. *The Annals of Statistics*, pages 206–227, 1985.
- Aleksey S Polunchenko, Alexander G Tartakovsky, and Nitis Mukhopadhyay. Nearly optimal change-point detection with an application to cybersecurity. *Sequential Analysis*, 31(3):409–435, 2012.
- Ya’acov Ritov. Decision theoretic optimality of the cusum procedure. *The Annals of Statistics*, pages 1464–1469, 1990.
- Lukas Ruff, Jacob R Kauffmann, Robert A Vandermeulen, Grégoire Montavon, Wojciech Samek, Marius Kloft, Thomas G Dietterich, and Klaus-Robert Müller. A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5):756–795, 2021.
- Albert N Shiryaev. *Optimal stopping rules*, volume 8. Springer Science & Business Media, 2007.
- AG Tartakovsky. *Sequential methods in the theory of information systems*, 1991.
- Che-Ping Tsai, Adarsh Prasad, Sivaraman Balakrishnan, and Pradeep Ravikumar. Heavy-tailed streaming statistical estimation. In *International Conference on Artificial Intelligence and Statistics*, pages 1251–1282. PMLR, 2022.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- H Victor. A general class of exponential inequalities for martingales and ratios. *The Annals of Probability*, 27(1): 537–564, 1999.

- Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- Abraham Wald. *Sequential tests of statistical hypotheses*. Springer, 1992.
- Hongjian Wang and Aaditya Ramdas. Catoni-style confidence sequences for heavy-tailed mean estimation. *arXiv preprint arXiv:2202.01250*, 2022.
- Hongjian Wang and Aaditya Ramdas. Huber-robust confidence sequences. In *International Conference on Artificial Intelligence and Statistics*, pages 9662–9679. PMLR, 2023.
- Liyan Xie, Shaofeng Zou, Yao Xie, and Venugopal V Veeravalli. Sequential (quickest) change detection: Classical results and new directions. *IEEE Journal on Selected Areas in Information Theory*, 2(2):494–514, 2021.
- Ping Yang, Guy Dumont, and John Mark Ansermino. Adaptive change detection in heart rate trend monitoring in anesthetized children. *IEEE transactions on biomedical engineering*, 53(11):2211–2219, 2006.
- Jie Zhang, Zhi Wei, Zhenyu Yan, MengChu Zhou, and Abhishek Pani. Online change-point detection in sparse time series with application to online advertising. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 49(6):1141–1151, 2017.