
Locally Private Parametric Methods for Change-Point Detection

Anuj Kumar Yadav
EPFL
Switzerland

Cemre Cadir
EPFL
Switzerland

Yanina Shkel
EPFL
Switzerland

Michael Gastpar
EPFL
Switzerland

Abstract

We study parametric change-point detection, where the goal is to identify distributional changes in time series, under local differential privacy. In the non-private setting, we derive substantially improved finite-sample accuracy guarantees for a change-point detection algorithm based on the generalized log-likelihood ratio test, using martingale methods. In the private setting, we propose two locally differentially private algorithms based on randomized response and binary mechanisms, and analyze their performance across different privacy regimes, both theoretically and empirically. As part of our analysis, we establish a structural result on strong data processing inequality (SDPI) coefficients for Jeffreys-Rényi divergence, which may also be of independent interest. We corroborate our theoretical results through empirical evaluation on synthetic data.

1 Introduction

1.1 Motivation

Consider a scenario in which a city health office is responsible for the detection of disease outbreaks and the enforcement of appropriate public health measures. To accomplish this task, they rely on daily patient-admission records provided by a city-wide hospital network. The task is to detect an outbreak by tracking hospital admissions for designated symptoms and identifying statistically significant changes in incidence. This process of identifying distributional shifts in the underlying data is commonly referred to as Change-Point Detection (CPD). CPD has diverse applications in many areas including but not limited to finance [11], climate and environmental systems [2], communication networks [5] and industrial fault detection [7].

Differential privacy (DP) [4] is a widely adopted central notion of privacy. Standard DP assumes the presence of a trusted curator who has access to the raw data and applies a privacy-preserving algorithm before releasing the output. While this assumption is appropriate in some settings, it may be undesirable when the analyst responsible for producing or releasing the output cannot be completely trusted with sensitive data. Local differential privacy (LDP) [1, 17] strengthens the privacy guarantee by requiring each data owner to locally privatize their data prior to sharing, thereby protecting the data at the source and removing reliance on a trusted curator. Motivated by this stronger privacy model, we study change-point detection under LDP. The change-point detection (CPD) problem has been studied in both parametric [13, 10, 15] and non-parametric settings [14, 8]. In the parametric setting, the pre-change and post-change distributions, denoted by P_0 and P_1 , are assumed to be known. In contrast, in the non-parametric setting, these distributions are unknown, and methods rely on discrepancy statistics or other distribution-free procedures [6]. Offline CPD estimates changes after observing a batch of data points of fixed (small) size, whereas online CPD detects changes sequentially as data arrive. We focus on the offline parametric CPD setting.

1.2 Our Contributions

• **Non-private accuracy guarantees.** We derive finite-sample performance bounds for an offline change-point detection algorithm based on the generalized log-likelihood ratio test (GLRT). Using martingale concentration techniques, we establish two accuracy bounds (Theorem 3.1). Both of our bounds significantly improve upon existing results in the literature. (Section 3)

• **Characterization of SDPI coefficients.** In Section 4, we present novel characterizations of strong data processing inequality (SDPI) coefficients for Jeffreys–Rényi divergence (of order $\rho = \infty$). We show that the SDPI coefficients are achieved by input distributions supported on at most two atoms. Leveraging this result, we explicitly characterize the SDPI coefficients for symmetric channels, which are later used to analyze theoretical performance of our private CPD algorithms.

• **LDP-based CPD algorithms.** We propose two CPD algorithms based on randomized response and binary mechanisms. Using our characterization of SDPI coefficients, we derive finite-sample accuracy bounds for both private CPD algorithms. Our analysis also shows that the algorithm based on the Chernoff-optimal binary mechanism outperforms the algorithm based on the randomized response in the high-privacy regime, while the converse holds in the low-privacy regime. (Section 5)

For detailed proofs and experiments, please refer to the arXiv version of the paper [16].

2 Parametric offline change-point detection with local differential privacy

In the offline setting, the parametric CPD problem is defined by a dataset $\mathcal{D} \triangleq \{x_1, x_2, \dots, x_n\}$ of size $n > 1$, the pre-change distribution $P_0 \in \Delta(\mathcal{X})$, and the post-change distribution $P_1 \in \Delta(\mathcal{X})$, where $\mathcal{X}$ is the sample space. The data points in $\mathcal{D}$ are initially sampled i.i.d. from P_0 ; at some unknown change-point $k^* \in (1, n]$, an event triggers and the underlying distribution becomes P_1 . The data analyst has access to $(\mathcal{D}, P_0, P_1)$. The goal of the analyst is to output an estimate of the change-point, denoted by $\hat{k} \in (1, n]$. We assume that the dataset $\mathcal{D}$ contains at most one true change-point.

LDP is a well-studied privacy framework. It guarantees that a strong adversary cannot distinguish between two possible inputs by knowing the privatized output. Let $(\mathcal{D}, P_0, P_1)$ denote the true dataset, the pre- and the post-change distributions at the source. In the LDP based CPD framework, each data point $x_i \in \mathcal{D}$ is privatized independently at the source via an ε -LDP mechanism W^ε . Let $\tilde{\mathcal{D}}$ denote the privatized dataset which contains the noisy data points y_i for all $i \in [n]$. Q_0 and Q_1 denote the distributions induced through the ε -LDP mechanism W^ε , by the pre- and post-change distributions P_0 and P_1 , respectively. The source then sends the tuple $(\tilde{\mathcal{D}}, Q_0, Q_1)$ to the data analyst. The analyst with access to the privatized dataset and induced pre- and post- change distributions estimates the change-point of the true dataset $\mathcal{D}$.

3 Non-Private Change-point detection

Inspired by the CUSUM approach [13] which is commonly studied for the online setting, our analysis is based on the GLRT described in Algorithm 1 (also studied in [18]).

Algorithm 1 Offline CPD $(\mathcal{D}, \mathcal{X}, P_0, P_1)$

- 1: **Input:** Dataset $\mathcal{D}$, Alphabet $\mathcal{X}$, Distributions P_0, P_1
 - 2: **for** $k = 1$ **to** n **do**
 - 3: Compute likelihood ratio $l(\mathcal{D}, k) = \sum_{i=k}^n \log \frac{P_1(x_i)}{P_0(x_i)}$
 - 4: **end for**
 - 5: **Output:** $\hat{k} = \arg \max_{k \in [n]} l(\mathcal{D}, k)$
-

The sensitivity s of the GLRT statistic w.r.t P_0 and P_1 is given by the maximum difference in the likelihood ratio for two neighboring databases $\mathcal{D}$ and $\mathcal{D}'$ (which differ in only one data point) i.e. $s \triangleq \max_{x \in \mathcal{X}} \log \frac{P_1(x)}{P_0(x)} - \min_{x \in \mathcal{X}} \log \frac{P_1(x)}{P_0(x)} = D_\infty^J(P_0, P_1)$. We also denote $C \triangleq \min\{D_{\text{KL}}(P_0 \| P_1), D_{\text{KL}}(P_1 \| P_0)\}$. Now, we state our finite-sample accuracy bounds for the offline change-point estimator in Algorithm 1.

Theorem 3.1. Let $k^* \in (1, n]$ be the true change-point of the dataset $\mathcal{D} \in \mathcal{X}^n$, with data points drawn from the distributions $P_0, P_1 \in \Delta(\mathcal{X})$. Then, the Offline CPD estimator in Algorithm 1 is (α, β) -accurate for any $\alpha \in [n]$ and

$$\beta \leq 2 \min \left\{ \sum_{i=1}^{i^*} \exp \left(\frac{-2^{i-1} \alpha C^2}{s^2} \right); \exp(-\alpha I_{ch}(P_0, P_1)) \right\}, \quad (1)$$

where $i^* = \lceil \log_2 \left(\frac{n-1}{\alpha} \right) \rceil$.

Although Algorithm 1 was previously studied in [18], our analysis yields significantly sharper bounds. To derive accurate performance bounds for our privatized algorithms we will need to quantify the information contraction of Jeffreys-Rényi divergence of order infinity, through an ε -LDP mechanism. It is therefore essential to characterize the SDPI coefficients associated with $W_{Y|X}^\varepsilon$. For the rest of the paper, we restrict ourselves to finite alphabets $\mathcal{X}$ and $\mathcal{Y}$.

4 Strong Data Processing Inequalities (SDPI) coefficients

Consider a channel $W_{Y|X} : \mathcal{X} \rightarrow \mathcal{Y}$. For any two input distributions (P_0, P_1) in $\Delta(\mathcal{X})$, let Q_0 and Q_1 denote the distributions induced in $\Delta(\mathcal{Y})$ by passing P_0 and P_1 through $W_{Y|X}$, respectively, i.e., $Q_i(y) = \sum_{x \in \mathcal{X}} P_i(x) W_{Y|X}(y|x)$, $i \in \{0, 1\}$. The data processing inequality (DPI) [3] states that processing cannot increase the statistical distinguishability between distributions. Specifically, for any divergence D that is monotone under stochastic transformations, $D(Q_0, Q_1) \leq D(P_0, P_1)$. The strengthening of DPI is commonly formalized via SDPI coefficients which quantify the contraction of divergences induced by a stochastic channel. Next, we establish results on the SDPI coefficients for Jeffreys-Rényi divergences associated with a channel W . It was shown in [12] that for the class of f -divergences, the defining supremum of the SDPI coefficients is achieved by distributions supported on at most two atoms. We prove a similar result for Jeffreys-Rényi divergence.

Theorem 4.1. For Jeffreys-Rényi divergence of order $\rho = \infty$ i.e., $D_\infty^J(\cdot, \cdot)$, the SDPI coefficient of a channel W , $\eta_\rho^J(W)$, is achieved by binary input distributions.

Next, we derive the SDPI coefficients for ∞ -order Jeffreys-Rényi divergence for a class of symmetric channels via Theorem 4.1.

Theorem 4.2. Let W be a q -ary symmetric channel. Then, the SDPI coefficient of W for Jeffreys-Rényi divergence of order infinity is $\eta_\infty^J(W) = \frac{|v-u|}{v+u}$. Furthermore, it is achieved in the limit by uniform binary distributions.

5 Private Change-point Detection with Local Differential Privacy (LDP)

Recall that in Thm. 3.1 the error probability of the non-private CPD estimator (Algorithm 1) admits a sharp upper bound in terms of Chernoff information between the input distributions (P_0, P_1) . Motivated by this bound, we use Chernoff information between the induced output distributions, (Q_0, Q_1) , as a proxy for algorithmic performance in the private setting. We identify ε -LDP mechanisms that maximize induced Chernoff information in low and high privacy regimes.

Theorem 5.1. Let (P_0, P_1) be any two PMFs. Then, there exists a positive $\varepsilon^*(P_0, P_1)$ such that for all $\varepsilon \geq \varepsilon^*$, the randomized response mechanism maximizes the Chernoff information $I_{ch}(Q_0, Q_1)$ between the induced output distributions (Q_0, Q_1) , among all ε -LDP mechanisms.

In [9], the authors show that for any (P_0, P_1) , the binary mechanism maximizes any f -divergence between (Q_0, Q_1) for sufficiently small ε . Based on this result, we have the following corollary.

Corollary 5.2. Let (P_0, P_1) be any two PMFs. There exists a positive threshold $\varepsilon'(P_0, P_1)$ such that for all $\varepsilon \leq \varepsilon'$, the binary mechanism maximizes $I_{ch}(Q_0, Q_1)$ over all ε -LDP mechanisms.

With the induced Chernoff information as a performance proxy, Theorem 5.1 and Corollary 5.2 yield a mechanism-level ordering in terms of Chernoff information contraction. In particular, the

binary mechanism incurs less contraction of Chernoff information than randomized response in the high-privacy regime (small ε), while the converse holds in the low-privacy regime (large ε).

Algorithm 2 Offline RR – CPD

```

1: Input:  $\mathcal{D}, \mathcal{X}, P_0, P_1$ .
2: Let  $Q_0, Q_1$  be the RR output distributions of  $P_0, P_1$ .
3: Apply RR to each  $x_i$  to obtain  $y_i$ ; set  $\tilde{\mathcal{D}} = \{y_i\}_{i=1}^n$ .
4: for  $k = 1$  to  $n$  do
5:    $l(\tilde{\mathcal{D}}, k) = \sum_{i=k}^n \log \frac{Q_1(y_i)}{Q_0(y_i)}$ .
6: end for
7: Output:  $\hat{k} = \arg \max_{k \in [n]} l(\tilde{\mathcal{D}}, k)$ .
```

Algorithm 3 Offline BM – CPD

```

1: Input:  $\mathcal{D}, \mathcal{X}, P_0, P_1$ .
2: Let  $Q_0^\tau, Q_1^\tau$  be the output distributions of  $P_0, P_1$  under  $W_{\tau}^b$ .
3: Compute  $\tau^* \in \arg \max_{\tau} I_{ch}(Q_0^\tau, Q_1^\tau)$ .
4: Apply  $W_{\tau^*}^b$  to each  $x_i$  to obtain  $y_i \in \{0, 1\}$ ; set  $\tilde{\mathcal{D}} = \{y_i\}_{i=1}^n$ .
5: for  $k = 1$  to  $n$  do
6:    $l(\tilde{\mathcal{D}}, k) = \sum_{i=k}^n \log \frac{Q_1^{\tau^*}(y_i)}{Q_0^{\tau^*}(y_i)}$ .
7: end for
8: Output:  $\hat{k} = \arg \max_{k \in [n]} l(\tilde{\mathcal{D}}, k)$ .
```

5.1 Private Offline CPD with randomized response

Definition 5.3 (Randomized Response mechanism). Let $|\mathcal{X}| = |\mathcal{Y}| = q$. Given $X = x$, the RR mechanism $W_{Y|X}^r$ outputs $Y = x$ with probability $e^\varepsilon / (e^\varepsilon + q - 1)$ and each other symbol $y \neq x$ with probability $1 / (e^\varepsilon + q - 1)$.

In the next theorem, we provide an upper bound on the error probability of Algorithm 2, for accurately estimating the change-point while satisfying ε -LDP constraint.

Theorem 5.4. Let $k^* \in (1, n]$ be the true change-point of the dataset $\mathcal{D} \in \mathcal{X}^n$, with data points drawn from the distributions $P_0, P_1 \in \Delta(\mathcal{X})$. Then, the Offline RR – CPD estimator in Algorithm 2 satisfies ε -LDP and (α, β_r) -accurate for any $\alpha \in [n]$ and

$$\beta_r \leq 2 \min \left\{ \sum_{i=1}^{i^*} \exp \left(\frac{-2^{i-1} \alpha C_r^2}{s_r^2} \right); \left(1 - \frac{C_r}{2} \right)^{\frac{\alpha}{2}} \right\}, \quad (2)$$

where $s_r = \min \{2\varepsilon; \tanh(\varepsilon/2)s\}$ and $C_r = 2 \left(\frac{e^\varepsilon - 1}{e^\varepsilon + q - 1} \right)^2 d_{TV}^2(P_0, P_1)$. Here, $i^* = \lceil \log_2 \left(\frac{n-1}{\alpha} \right) \rceil$ and $\tanh(\varepsilon/2) = \frac{e^\varepsilon - 1}{e^\varepsilon + 1}$.

5.2 Private Offline CPD with binary mechanism

Definition 5.5 (binary mechanism). The binary mechanism W_{τ}^b is defined by a channel $W_{\tau, Y|X}^b : \mathcal{X} \rightarrow \{0, 1\}$ in two steps:

- (i) *Quantization*. The quantizer $Z_\tau : \mathcal{X} \rightarrow \{0, 1\}$ is a deterministic mapping parameterized by $\tau > 0$ such that it partitions $\mathcal{X}$ into $\mathcal{S}_\tau$ and $\mathcal{S}_\tau^c$ as follows: $\mathcal{S}_\tau := \{x \in \mathcal{X} : P_0(x) \geq \tau P_1(x)\}$. The elements in $\mathcal{S}_\tau$ are mapped to 0, while the elements in $\mathcal{S}_\tau^c$ are mapped to 1. Let $V = Z_\tau(X)$ where $V \in \{0, 1\}$.
- (ii) *Privatization (binary randomized response)*. Given $V = v$, output $Y = v$ with probability $e^\varepsilon / (1 + e^\varepsilon)$ and $Y = 1 - v$ otherwise.

Therefore, $W_{\tau, Y|X}^b = Z_{\tau, V|X} \cdot W_{Y|V}$. In the next theorem, we provide an upper bound on the error probability of Algorithm 3 for accurately estimating the change-point while satisfying ε -LDP.

Theorem 5.6. Let $k^* \in (1, n]$ be the true change-point of the dataset $\mathcal{D} \in \mathcal{X}^n$, with data points drawn from the distributions $P_0, P_1 \in \Delta(\mathcal{X})$. Then, the Offline BM – CPD estimator in Algorithm 3 satisfies ε -LDP and (α, β_b) -accurate for any $\alpha \in [n]$ and

$$\beta_b \leq 2 \min \left\{ \sum_{i=1}^{i^*} \exp \left(\frac{-2^{i-1} \alpha \tilde{C}_b^2}{s_b^2} \right); \left(1 - \frac{C_b}{2} \right)^{\frac{\alpha}{2}} \right\}, \quad (3)$$

where $s_b = \min \{2\varepsilon; \tanh(\varepsilon/2)s\}$, $C_b = 2 \tanh^2(\varepsilon/2) d_{TV}^2(P_0, P_1)$, and $\tilde{C}_b = 2 \tanh^2(\varepsilon/2) \left| \sum_{x \in \mathcal{S}_{\tau^*}} (P_0(x) - P_1(x)) \right|^2$.

References

- [1] Raef Bassily and Adam Smith. Local, private, efficient protocols for succinct histograms. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 127–135, 2015.
- [2] Bryson C. Bates, Richard E. Chandler, and Adrian W. Bowman. Trend estimation and change point detection in individual climatic series using flexible regression methods. *Journal of Geophysical Research: Atmospheres*, 117(D16), 2012.
- [3] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, USA, 2006.
- [4] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
- [5] Christoph Füllner, Anuj K. Yadav, and René Bonk. Dsp-assisted in-vivo health monitoring for ≥ 50 g pon. In *Optical Fiber Communication Conference (OFC) 2025*, page W3J.4. Optica Publishing Group, 2025.
- [6] Zaïd Harchaoui, Eric Moulines, and Francis Bach. Kernel change-point analysis. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc., 2008.
- [7] Douglas M. Hawkins, Peihua Qiu, and Chang Wook Kang. The changepoint model for statistical process control. *Journal of Quality Technology*, 35(4):355–366, 2003.
- [8] Mark Holmes, Ivan Kojadinovic, and Jean-François Quessy. Nonparametric tests for change-point detection à la gombay and horváth. *Journal of Multivariate Analysis*, 115:16–32, 2013.
- [9] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. Extremal mechanisms for local differential privacy. *Journal of Machine Learning Research*, 17(17):1–51, 2016.
- [10] George V. Moustakides. Optimal Stopping Times for Detecting Changes in Distributions. *The Annals of Statistics*, 14(4):1379 – 1387, 1986.
- [11] Peter Nystrup, Bo William Hansen, Henrik Madsen, and Erik Lindström. Detecting change points in vix and s&p 500: A new approach to dynamic asset allocation. *Journal of Asset Management*, 17(5):361–374, 2016.
- [12] Or Ordentlich and Yury Polyanskiy. Strong data processing constant is achieved by binary inputs. *IEEE Transactions on Information Theory*, 68(3):1480–1481, 2022.
- [13] E. S. Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, 1954.
- [14] A. N. Pettitt. A non-parametric approach to the change-point problem. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(2):126–135, 1979.
- [15] 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.
- [16] Anuj Kumar Yadav, Cemre Cadir, Yanina Shkel, and Michael Gastpar. Locally private parametric methods for change-point detection. *arXiv preprint arXiv:2602.13619*, 2026.
- [17] Mengmeng Yang, Taolin Guo, Tianqing Zhu, Ivan Tjuawinata, Jun Zhao, and Kwok-Yan Lam. Local differential privacy and its applications: A comprehensive survey. *Computer Standards & Interfaces*, page 103827, 2023.
- [18] Wanrong Zhang, Sara Krehbiel, Rui Tuo, Yajun Mei, and Rachel Cummings. Single and multiple change-point detection with differential privacy. *Journal of Machine Learning Research*, 22(29):1–36, 2021.