

Secure Change-Point Detection for Time Series under Homomorphic Encryption

Federico Mazzone*
University of Twente
Enschede, Netherlands
f.mazzone@utwente.nl

Giorgio Micali*
University of Twente
Enschede, Netherlands
g.micali@utwente.nl

Massimiliano Pronesti
IBM Research Europe
Dublin, Ireland
massimiliano.pronesti@ibm.com

Abstract

We introduce the first method for change-point detection on encrypted time series. Our approach employs the CKKS homomorphic encryption scheme to detect shifts in statistical properties (e.g., mean, variance, frequency) without ever decrypting the data. Unlike solutions based on differential privacy, which degrade accuracy through noise injection, our solution preserves utility comparable to plaintext baselines. We assess its performance through experiments on both synthetic datasets and real-world time series from healthcare and network monitoring. Notably, our approach can process one million points within 3 minutes.

Keywords

time series, change-point detection, privacy, homomorphic encryption, ordinal patterns

1 Introduction

The analysis of time series data often requires detecting structural changes, known as *change points*, which indicate significant shifts in the underlying data distribution. The change-point detection (CPD) problem consists of identifying time indices at which such shifts occur. This problem has been studied for over seventy years, dating back to the pioneering work of Page [31]. Since then, CPD has become a central topic of research across many areas of quantitative and applied science: in econometrics and finance [5, 19, 25], network monitoring [26, 27], signal processing [10], genetics [11], climatology [33, 37], quantum optics [34], speech recognition [16], MRI analysis [4], and healthcare [35]. Across these domains, a change in the data distribution may indicate the onset of a disease or a shift in a patient’s condition, a market crash or policy intervention, a sudden variation in climatic patterns, or an anomaly in network traffic. Accurately detecting such changes is essential both for retrospective analysis and for real-time decision-making.

Time-series data are often collected by sensors or client devices and transmitted to external servers or data curators for processing. In many practical scenarios, on-device (local) analysis is either infeasible or undesirable. For instance, IoT monitoring devices typically lack permanent storage and are mainly designed to collect, optionally encrypt, and forward data rather than to run specialized

analytics such as CPD. Even when feasible, embedding task-specific logic would reduce their flexibility, as analysts may not know in advance which time intervals to inspect or which preprocessing steps to apply. Outsourcing the raw data first and performing the analysis off-device offers greater versatility and allows multiple tasks, such as CPD, anomaly detection, or other retrospective analyses, to be executed later. Moreover, many data providers already rely on cloud-based platforms (e.g., Azure Health Data Services or Google Healthcare API) for secure storage, regulatory compliance, or cross-institutional collaboration, making outsourced CPD a natural fit. Such a setting is also common when data originate from multiple parties and a third-party curator is required to coordinate tasks such as timeline alignment, normalization, or data fusion, which cannot be performed locally.

However, outsourcing raises significant privacy concerns, as the transmitted time-series data are often highly sensitive. For example, medical time series such as ECGs or EEGs contain personal health information; network traffic traces may reveal private user behavior or communication patterns; and industrial monitoring data can expose proprietary system activity. Previous work [18, 38] addressed privacy-preserving CPD exclusively in the *central* privacy model [20], which assumes the data curator to be trusted and focuses on providing privacy at the moment the computed change-point index is released to an external party. To do so, they apply differential privacy (DP) by injecting noise into the CPD mechanism, so that only a noisy version of the result is released. This limits what a potential adversary can infer about the underlying time series, at the cost of reduced accuracy.

In many settings, however, the data provider does not trust the curator, placing the problem in the *local* privacy model. In this setting, the client must protect its data before sending it out. To date, no work has addressed privacy-preserving CPD in the local model, and this work aims to close this gap. In particular, **our goal is to enable an honest-but-curious curator** (or computing server) **to perform CPD on the client’s data without compromising the privacy of the underlying time series**. One option would be to use local DP, where the client perturbs the raw time-series values with random noise before sending it to the curator. As we will show, injecting noise directly at the data level severely degrades the accuracy of the CPD task, making this approach impractical in many realistic scenarios (over 30% relative error in the resulting change-point index for a 10,000-point series with $\epsilon = 25$). To avoid this limitation, we instead employ fully homomorphic encryption (FHE), specifically the CKKS scheme [13]. In our solution, the client encrypts its entire time series and sends only ciphertexts to the curator. Thanks to the homomorphic properties of CKKS, the server can then execute the full CPD pipeline directly over encrypted

*The two authors made equal contributions and share first authorship.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2026(2), 127–142
© 2026 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2026-0040>

data, without ever accessing the underlying values, and return the encrypted result to the client. In addition to avoiding the accuracy issues of local DP, this approach also offers a stronger privacy guarantee, as the curator never sees the data, not even in noisy form.

These improved accuracy and privacy guarantees come at the cost of a large computational overhead due to the homomorphic operations. Nonetheless, our choice of FHE over DP is still motivated by the fact that, in many applications, accuracy requirements tend to outweigh runtime constraints. In clinical monitoring, for instance, detecting a physiological regime change (e.g., heart-rate) within a margin of a few seconds can distinguish between early-stage and late-stage patient deterioration. By contrast, runtime can tolerate delays, since such analyses are commonly executed offline or asynchronously with respect to data collection. Similarly, in network-traffic monitoring, accurately spotting sudden changes in packet-rate or latency is important to distinguish normal fluctuations from the onset of a DDoS attack or a routing fault. Missing such a change can delay incident detection or trigger false alarms, while some computational delay is typically acceptable since analyses often run on buffered or batched traffic.

Our Approach and Contributions. Prior works [18, 38] build their approaches on classical detection techniques based on the log-likelihood ratio principle. However, these techniques rely on strong statistical assumptions that rarely hold in practice:

- (1) independence among the data points of the time series, and
- (2) prior knowledge of the data distribution before and after the change point.

To avoid the need for such assumptions, we base our solution on the detection strategy proposed by Bandt and Pompe [6] and Betken et al. [9]. This strategy consists of first transforming the time series into a sequence of ordinal-pattern histograms (i.e., permutations encoding the relative ordering of observations), reducing the problem to an easier mean-change detection task, and then applying standard Cumulative Sum (CUSUM) statistics to solve it on the resulting summaries. Since ordinal patterns provide a non-parametric description of the time series, they are naturally robust to noisy observations and outliers, and capture temporal structure without assuming independence or specific data distributions. As a result, they are well suited for detecting changes in temporal dependence (e.g., frequency-related changes), but are inherently insensitive to changes in marginal properties such as the mean or variance, for which we implement different transformations.

The main challenge in implementing this CPD strategy under FHE lies in performing the required computations efficiently, given both the high cost of homomorphic operations and the large size of typical time-series data. Our framework builds on recent advances in encrypted comparison techniques [29], which enable an efficient homomorphic evaluation of ranking and order-statistics operations. Our approach proceeds in two steps. First, we homomorphically transform the input time series into a sequence of block-level summaries, implementing different transformations for detecting changes in mean, variance, and frequency. Second, we implement the CUSUM statistic under encryption and apply it to these block-level sequences to localize the change point. We evaluate our solution on multiple datasets, including synthetic data from standard toy models and real-world data from healthcare and

network monitoring. Notably, our method can process a time series with 1,000,000 points, detecting mean changes in under two minutes and frequency and variance changes in under three minutes, all while preserving plaintext-level accuracy.

The main contributions of this paper are as follows.

- We present the first solution for secure outsourced CPD on dependent time series, achieving plaintext-level accuracy and practical runtime for detecting mean, variance, and frequency shifts under CKKS.
- We introduce new cryptographic building blocks for CKKS, including an efficient method for partial sums and an improved implementation of the argmax operation.
- We provide a thorough experimental evaluation of our solution on both synthetic and real-world datasets.

2 Background

2.1 System and Threat Model

The most general system model for CPD computation involves the following actors:

- one or more data providers,
- a computing server or data curator, and
- a result recipient.

The ideal functionality is that the data providers send their time-series data to the computing server, which computes the CPD index and delivers the result to the recipient. In this work, we focus on the outsourced-computation setting, hence, with a single data provider, who also coincides with the result recipient. The data provider encrypts its time series locally and sends the resulting ciphertexts to the server, which performs the entire CPD algorithm homomorphically and returns the encrypted output. Only the data provider, who holds the secret key, can decrypt the result.

Threat Model. We consider the computing server to be honest-but-curious, i.e., it follows the protocol but may attempt to extract information about the data provider’s time series from any intermediate computation it observes. In particular, the server is assumed to have full access to all ciphertexts and to all homomorphic computations performed during CPD, and its goal is to infer statistical properties, patterns, or individual data values from them. Security against such an adversary follows directly from the semantic security of the encryption scheme. Note that our threat model does not cover a disclosure of the decrypted result to a third-party recipient. In that case, central-DP mechanisms can be employed to protect the released index.

2.2 Time Series and the CPD Problem

A time series $(X_t)_{t \geq 1}$ is a sequence of random variables indexed by discrete time t . A standard approach in time series analysis is to model the data as a realization of a *stationary process*, where the statistical properties such as mean, variance, and autocorrelation do not change over time [12]. Formally, stationarity means that the joint distribution of $(X_{t_1}, \dots, X_{t_k})$ is the same as that of $(X_{t_1+h}, \dots, X_{t_k+h})$ for all time shifts $h \in \mathbb{Z}$ and all index tuples t_1, \dots, t_k . This assumption enables meaningful inference, as observations from different time intervals can be treated as samples from the same underlying distribution. However, many real-world time




Figure 1: Three common types of change points in time series.

series exhibit stationary behavior only over certain intervals and undergo structural changes at unknown times. In such cases, the global stationarity assumption fails, and detecting these change-points becomes essential for proper modeling and interpretation. Sometimes changes can be recognized visually (see Figure 1 for some examples), but in many cases they are too subtle to detect by eye, or appear in contexts where manual inspection is not feasible. To automate this task, a variety of methods have been proposed, most of which rely on a hypothesis-testing formulation.

Related works in the privacy-preserving context [18, 38] formalize the change-point detection problem by selecting a transformation ψ and testing whether its expectation remains constant over time:

$$\begin{aligned} \mathcal{H}_0 &: \mathbb{E}[\psi(X_t)] = \mu, \quad \forall t \\ \mathcal{H}_1 &: \exists \tau \text{ such that } \mathbb{E}[\psi(X_\tau)] \neq \mathbb{E}[\psi(X_{\tau+1})]. \end{aligned}$$

The choice of ψ determines the type of change being tested. For instance, $\psi(x) = x$ corresponds to changes in mean (location), while $\psi(x) = x^2$ corresponds to changes in variance (volatility). In fact, a classic test for detecting mean shifts is the following:

$$\begin{aligned} \mathcal{H}_0 &: X_t = \mu + \varepsilon_t, \quad \varepsilon_t \text{ i.i.d.}, \quad \forall t \\ \mathcal{H}_1 &: X_t = \begin{cases} \mu_1 + \varepsilon_t, & t = 1, \dots, \tau, \\ \mu_2 + \varepsilon_t, & t = \tau + 1, \dots, n \end{cases} \quad \text{with } \mu_1 \neq \mu_2. \end{aligned}$$

where n is the length of the series.

While the ψ -formulation is convenient, it only captures specific aspects of the distribution and ignores the temporal dependence usually present in time-series data, making it inadequate in realistic settings. Hence, in this work we focus on the more general




Figure 2: The six ordinal patterns of order 2 (figure from [9]).

formulation:

$$\begin{aligned} \mathcal{H}_0 &: (X_1, \dots, X_n) \text{ is a stationary process} \\ \mathcal{H}_1 &: \exists \tau \in \{1, \dots, n-1\} \text{ such that} \\ & \quad (X_1, \dots, X_\tau) \stackrel{\mathcal{D}}{\neq} (X_{\tau+1}, \dots, X_n). \end{aligned}$$

Our goal is to privately estimate the location of the change point τ under the alternative hypothesis. To this end, we build upon an algorithm that transforms the time series into a symbolic representation known as *ordinal patterns*.

We refer the reader to review papers by Gupta et al. [22] and Jandhyala et al. [23], and the monographs by Basseville and Nikiforov [8] and Csörgő and Horváth [17] for comprehensive treatments of change-point theory and methods.

2.3 Ordinal Patterns Representation

To estimate the change-point location under the general formulation, we adopt the approach introduced by Bandt and Pompe [7], who proposed to encode the dynamics of a time series as a stream of symbols. These symbols correspond to ordinal patterns, each describing the relative order of values within a local window of the series. Formally, let \mathcal{S}_r denote the set of permutations of $\{0, 1, \dots, r\}$. Given a real-valued time series (X_1, \dots, X_n) and an embedding order $r \geq 1$, we define the ordinal pattern map

$$\Pi : \mathbb{R}^{r+1} \rightarrow \mathcal{S}_r, \quad (X_t, \dots, X_{t+r}) \mapsto \pi, \quad (1)$$

where $\pi = (\pi_0, \dots, \pi_r)$ is the permutation that sorts the values (X_t, \dots, X_{t+r}) , i.e., π_j denotes the rank of X_{t+j} within the window. Sliding the window along the series yields a symbolic sequence of ordinal patterns.

Example. Consider the time series

$$\{x_1, x_2, \dots\} = \{4.2, 3.1, 5.0, 6.3, 2.9, 7.1, 1.8, 3.7, \dots\}$$

with embedding order $r = 2$, so each ordinal pattern is extracted from a window of three consecutive values. The first window $(4.2, 3.1, 5.0)$ is sorted as $(3.1, 4.2, 5.0)$, yielding the ranks $(2, 1, 3)$. The next window $(3.1, 5.0, 6.3)$ gives $(1, 2, 3)$, and so on. Thus, the time series is transformed into a sequence of rank tuples:

$$\{(2, 1, 3), (1, 2, 3), (2, 3, 1), (2, 1, 3), \dots\}.$$

Figure 2 illustrates the six possible ordinal patterns of order $r = 2$.

This symbolic representation offers multiple advantages: it is invariant under strictly increasing transformations of the data (such as rescaling or nonlinear monotonic mappings), and it is robust to outliers, since the ranks depend only on the relative order of the values rather than their magnitudes. Moreover, it induces a probability distribution over \mathcal{S}_r that captures the dependence structure of the process. For each $\pi \in \mathcal{S}_r$, we define the pattern probability

$$p(\pi) := \mathbb{P}(\Pi(X_t, \dots, X_{t+r}) = \pi), \quad (2)$$

which can be estimated from a finite sample via the empirical frequency

$$\hat{p}_n(\pi) = \frac{1}{n} \sum_{t=1}^{n-r} \mathbb{1}(\Pi(X_t, \dots, X_{t+r}) = \pi). \quad (3)$$

This distribution serves as a compact summary of the time series. It can be shown that under stationarity, \hat{p}_n remains constant over time, while a change in the underlying data-generating process induces a shift [9]. This way, CPD in real-valued time series reduces to detecting changes in discrete distributions over the finite alphabet \mathcal{S}_r .

2.4 Turning Rates and CPD Algorithm

Several statistics can be derived from ordinal patterns to detect distributional changes. Among these, Bandt [6] introduced the *turning rate*, which is based on ordinal patterns of order 2. It is defined as the total probability mass assigned to the four up-down patterns:

$$\mathcal{T} = \{(1, 3, 2), (3, 2, 1), (2, 3, 1), (2, 1, 3)\}.$$

The turning rate is then

$$q := \sum_{\gamma \in \mathcal{T}} p(\gamma), \quad (4)$$

where $p(\gamma)$ is the pattern probability from (2). Given observations X_1, \dots, X_n , the series is partitioned into $n_b = \lceil n/m \rceil$ non-overlapping blocks of length m , and a separate estimate $\hat{q}_{j,m}$ is computed for each block. This produces a sequence of local statistics that reflects how the turning rate evolves over time and can therefore be analyzed for structural breaks. Concretely, for each block we compute

$$\hat{q}_{j,m} = \frac{1}{m} \sum_{i=1}^{m-2} \sum_{\gamma \in \mathcal{T}} \mathbb{1}(\Pi(X_{s_{i,j}}, X_{s_{i,j}+1}, X_{s_{i,j}+2}) = \gamma), \quad (5)$$

$$s_{i,j} := m(j-1) + i, \quad j = 1, \dots, n_b.$$

This reformulates the original stationarity testing problem as a mean change detection problem on the sequence $(\hat{q}_{j,m})_{j=1}^{n_b}$. Specifically, the hypotheses become

$$\mathcal{H}_0 : \mathbb{E}[\hat{q}_{j,m}] = \mu_1 \quad \forall j, \quad (6)$$

$$\mathcal{H}_1 : \exists \tau^* \in \{1, \dots, n_b\} \text{ s.t.}$$

$$\mathbb{E}[\hat{q}_{j,m}] = \begin{cases} \mu_1, & j = 1, \dots, \tau^*, \\ \mu_2, & j = \tau^* + 1, \dots, n_b, \end{cases} \quad \mu_1 \neq \mu_2.$$

To evaluate this mean-change hypothesis, Betken et al. [9] use the classical CUSUM statistic applied to $(\hat{q}_{j,m})_j$, which reaches its maximum absolute value at the mean shift location:

$$\tau^* = \operatorname{argmax}_{k=1, \dots, n_b} \left| \sum_{j=1}^k \hat{q}_{j,m} - \frac{k}{n_b} \sum_{j=1}^{n_b} \hat{q}_{j,m} \right|. \quad (7)$$

This provides the estimated change point in the block domain, which can be translated into the original time scale by multiplying it by the block size, i.e., $\tau = m \cdot \tau^*$. We refer to [9] for a rigorous analysis of the statistical properties of the turning rate estimator, while Algorithm 1 summarizes the procedure.

Algorithm 1 Change-Point Detection via Ordinal Patterns

Input: x_1, x_2, \dots, x_n time series, m block size

Output: τ change-point

```

1:  $n_b \leftarrow \lceil n/m \rceil$  ▷ Number of blocks
2: for  $j = 1$  to  $n_b$  do
3:    $\hat{q}_{j,m} \leftarrow 0$  ▷ Initialize block statistic
4:   for  $i = 1$  to  $m - 2$  do
5:      $s \leftarrow (j - 1) \cdot m + i$ 
6:      $\pi \leftarrow \Pi(x_s, x_{s+1}, x_{s+2})$ 
7:     if  $\pi \in \mathcal{T}$  then
8:        $\hat{q}_{j,m} \leftarrow \hat{q}_{j,m} + 1$ 
9:    $\hat{q}_{j,m} \leftarrow \hat{q}_{j,m}/m$ 
10: for  $k = 1$  to  $n_b$  do
11:    $S_k \leftarrow \left| \sum_{j=1}^k \hat{q}_{j,m} - \frac{k}{n_b} \sum_{j=1}^{n_b} \hat{q}_{j,m} \right|$ 
12: return  $\tau \leftarrow m \cdot \operatorname{argmax}_k S_k$  ▷ see Equation 7

```

2.5 FHE and CKKS

FHE is a cryptographic primitive that allows arbitrary computations to be performed directly on encrypted data. This enables a client to outsource computations to untrusted servers without revealing any information about the underlying plaintexts. CKKS, in particular, is an FHE scheme introduced by Cheon et al. [13] and designed to work with vectors of real (floating-point) numbers. CKKS natively supports three operations on ciphertexts: component-wise addition ($X + Y$), component-wise multiplication ($X \cdot Y$), and vector rotations (left $X \ll k$ and right $X \gg k$). Throughout the paper, we denote ciphertexts with uppercase letters (e.g., X is the encryption of x).

Besides vectors, CKKS can also handle matrices, typically by encoding them row-by-row. Additions, multiplications, and rotations can be combined to define operations on matrices. In particular, for a matrix x encrypted as X :

- $\text{SumR}(X)$ sums all rows of x into the first row;
- $\text{ReplR}(X)$ replicates the first row across all rows;
- $\text{TransR}(X)$ transposes the first row into the first column.

Similar operations are available for the columns, i.e. $\text{SumC}(X)$, $\text{ReplC}(X)$, $\text{TransC}(X)$. These algorithms can be implemented recursively to improve efficiency, but require padding the input to the next power of two. Their pseudocode can be found in Appendix A.

Encrypted Ranking and Argmax. Computing the ordinal patterns (or ranking) Π and the argmax of a vector under homomorphic encryption is non-trivial, since these operations inherently involve comparisons (e.g., $X > Y$), which are extremely expensive under CKKS. To compute these operations, we build upon the design of Mazzone et al. [29], which we adapt and optimize to the CPD setting. The core idea is as follows. Given an input vector x , we first encode it into the first row of a null matrix X . Two re-encodings are then produced: one where values are replicated across rows, $X_R = \text{ReplR}(X)$, and one where they are replicated across columns, $X_C = \text{ReplC}(\text{TransR}(X))$. For example, for $x = (x_1, x_2, x_3, x_4)$ we

obtain

$$x_R = \begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ x_1 & x_2 & x_3 & x_4 \\ x_1 & x_2 & x_3 & x_4 \\ x_1 & x_2 & x_3 & x_4 \end{bmatrix} \quad x_C = \begin{bmatrix} x_1 & x_1 & x_1 & x_1 \\ x_2 & x_2 & x_2 & x_2 \\ x_3 & x_3 & x_3 & x_3 \\ x_4 & x_4 & x_4 & x_4 \end{bmatrix}.$$

A component-wise comparison between X_R and X_C provides the result of the comparison of all pairs (x_i, x_j) .

To perform encrypted comparison, we employ the polynomial approximation of the sign function provided by Cheon et al. [14], which is also used in [29]. Specifically, we define two polynomials

$$f(x) = (35x - 35x^3 + 21x^5 - 5x^7)/2^4$$

$$g(x) = (4589x - 16577x^3 + 25614x^5 - 12860x^7)/2^{10}.$$

and compose them a given number of times to obtain $\text{cmp}(x, y) = (f^{d_f}(g^{d_g}(x - y)) + 1)/2$, where d_f, d_g are the number of times each polynomial is composed with itself. These polynomials can be evaluated using a proper composition of homomorphic additions and multiplications. For sufficiently high d_f, d_g , we have $\text{cmp}(x, y) \approx 1$ if $x > y$, $\text{cmp}(x, y) \approx 0$ if $x < y$, and $\text{cmp}(x, y) = 0.5$ if $x = y$. Applying cmp to x_R and x_C results in a comparison matrix where each column contains as many entries close to one as the rank of the corresponding element (up to an offset of 0.5 due to self-comparisons). Hence, summing the rows and correcting by 0.5 provides the encrypted rank (see Algorithm 2).

Algorithm 2 Rank from [29]

Input: X encryption of $x = (x_1, \dots, x_N)$

Output: R encryption of the ranking of x

- 1: $X_R \leftarrow \text{ReplR}(X)$
 - 2: $X_C \leftarrow \text{ReplC}(\text{TransR}(X))$
 - 3: $C \leftarrow \text{cmp}(X_R, X_C)$
 - 4: $R \leftarrow \text{SumR}(C) + (0.5, \dots, 0.5)$
 - 5: **return** R
-

To compute the argmax of N encrypted values, we first compute their ranking using Algorithm 2. We then apply an indicator function ind_N , defined as $\text{ind}_N(x) = 1$ if $x = N$, and 0 otherwise. This function can be implemented as a product of two comparisons. Applying the indicator function to the output of the ranking algorithm returns a one-hot encoding of the argmax of the input vector (see Algorithm 3). In Section 3.3, we describe a more efficient argmax design that avoids computing the full ranking.

Algorithm 3 Argmax from [29]

Input: X encryption of $x = (x_1, \dots, x_N)$

Output: A encryption of $\text{argmax}(x)$ as one-hot encoding

- 1: $R \leftarrow \text{Rank}(X)$ ▷ see Algorithm 2
 - 2: $A \leftarrow \text{ind}_N(R)$
 - 3: **return** A
-

As a side note, the above algorithms may not behave correctly when $x_i = x_j$ for $i \neq j$. While Mazzone et al. [29] propose an additional tie-handling mechanism, we can avoid its cost here, since ties usually occur with low probability in time-series data.

3 Secure CPD

We now present our approach for estimating the change point τ of a time series under encryption. The approach can be divided into two steps:

- (1) a block-wise summarization of the time series data, which turns the CPD problem into a mean-shift detection problem, and
- (2) an application of the CUSUM statistic to solve the mean-CPD problem.

The exact summarization function depends on the type of change we want to detect. Mean and variance are used to detect shifts in the corresponding properties, while the turning rate statistic is used for changes in frequency (see Section 2.4). We first describe the matrix encoding we use for the input data, then we provide a secure implementation of these three summarization functions and the CUSUM statistic, and finally we put them together to build CPD algorithms for detecting changes in mean, variance, and frequency.

3.1 Block-Based Matrix Encoding

Given a time series $x = (x_1, \dots, x_n)$ and a block size m for our summarization function, we define $n_b = \lceil n/m \rceil$ to be the number of blocks of the series. We make the following assumptions on the block size to simplify the description of the algorithm:

- (1) m divides n , ensuring all blocks have the same size, and
- (2) m is a power of two, allowing a straightforward application of the recursive SumR, SumC, ReplR, ReplC, TransR, TransC algorithms.

Both assumptions can be lifted with minor adjustments to the algorithms, as discussed in Section 3.4.

A naive implementation of the plaintext algorithms using the tools from Section 2.5 would result in many nested loops of homomorphic operations and a large number of ciphertexts to manage. As we will see in the following sections, we can significantly reduce the computational cost by fully exploiting the SIMD capabilities of CKKS. For this purpose, we represent the time series as a matrix:

$$\begin{bmatrix} x_1 & x_2 & \dots & x_m \\ x_{m+1} & x_{m+2} & \dots & x_{2m} \\ \vdots & \vdots & & \vdots \\ x_{n-m+1} & x_{n-m+2} & \dots & x_n \end{bmatrix}$$

where each row corresponds to a block of m elements. We also design the summarization functions so that each block's summary is stored in the first element of the corresponding row. The CUSUM computation is then designed accordingly to take as input the first column of this matrix representation (see Figure 3). The reason behind this choice will become clear during the description of the different components.

3.2 Summarization Functions

Implementing efficient block-wise mean and variance is relatively straightforward, while the ordinal pattern representation is more challenging.

Mean. Computing the mean of each block can be done efficiently with a SumC operation, which recursively sums all elements of each




Figure 3: Schematic representation of our secure CPD pipeline. Here $n = 9$, $m = 3$, and the change point is detected at the second block.

block into the first column of the matrix, and a scalar multiplication by $1/m$ (see Algorithm 4).

Algorithm 4 Mean

Input: X encryption of $x = (x_1, \dots, x_n)$, m block size
Output: S encryption of $s^\top = (s_1, \dots, s_{n_b})^\top$ block-wise mean of x

- 1: $S \leftarrow \text{SumC}(X)$
- 2: $S \leftarrow (1/m) \cdot S$
- 3: **return** S

Variance. Computing the variance is more complex, but it can still be done efficiently. We recall that for each block $x_{km+1}, \dots, x_{(k+1)m}$ we need to compute

$$\frac{1}{m-1} \sum_{i=1}^m (x_{km+i} - \bar{x}_k)^2$$

where \bar{x}_k is the mean of the given block. First, we compute the mean as in Algorithm 4 and we replicate it over the columns with a ReplC . Then, we subtract it from the original time series matrix, square the result, sum it block-wise with a SumC , and multiply it by $\frac{1}{m-1}$ (see Algorithm 5).

Algorithm 5 Variance

Input: X encryption of $x = (x_1, \dots, x_n)$, m block size
Output: S encryption of $s^\top = (s_1, \dots, s_{n_b})^\top$ block-wise variance of x

- 1: $M \leftarrow \text{Mean}(X, m)$ ▷ Algorithm 4
- 2: $M \leftarrow \text{ReplC}(M)$
- 3: $S \leftarrow 1/(m-1) \cdot \text{SumC}((X - M)^2)$
- 4: **return** S

Turning Rates. For each block we create a sliding window of size three and count how many triplets have rank in \mathcal{T} . A straightforward implementation would consist of mirroring Algorithm 1. For each block $j = 1, \dots, n_b$:

- (1) **Triplet extraction:** extract $m-2$ triplets of consecutive values (x_i, x_{i+1}, x_{i+2}) by applying plaintext bitmasks and ciphertext rotations that zero out irrelevant entries and isolate each triplet.
- (2) **Ordinal pattern computation:** compute the ordinal pattern of each triplet using Algorithm 2, resulting in $R_i = (r_i, r_{i+1}, r_{i+2})$.

- (3) **Pattern membership check:** determine whether the pattern belongs to \mathcal{T} . This can be done for instance by evaluating

$$M_i = \frac{1}{12} (\|R_i - (1, 2, 3)\|_2^2 + \|R_i - (3, 2, 1)\|_2^2)$$

which equals one if $R_i \in \mathcal{T}$ and zero otherwise, since the only patterns not in \mathcal{T} are $(1, 2, 3)$ and $(3, 2, 1)$ (one can easily verify this by checking all six possible permutations).

- (4) **Turning rate estimation:** the number of matching patterns is computed by summing all M_i ; then it is normalized by the number of triplets, yielding an encrypted estimate of the turning rate:

$$\text{Enc}(\hat{q}_{j,m}) = \text{Enc}(\#\text{matches}/(m-2)) .$$

This implementation employs $n_b(m-2)$ ciphertexts and requires as many Rank operations, making it particularly inefficient. Each call to Rank uses only $3^2 = 9$ slots, leaving most ciphertext slots unused (a ciphertext typically has 2^{13} to 2^{17} slots). Parallelizing multiple ranking operations within one ciphertext would be a valid optimization, but a more efficient approach is possible. To verify whether a triplet (x_i, x_{i+1}, x_{i+2}) belongs to \mathcal{T} , we do not actually need all 9 comparisons $x_i > x_j$ performed in Rank. It is enough to check whether the triplet is monotonic, increasing or decreasing, and declare membership to \mathcal{T} otherwise. This happens if and only if

$$x_i < x_{i+1} < x_{i+2} \quad \text{OR} \quad x_i > x_{i+1} > x_{i+2} .$$

Hence, we potentially need only four comparisons instead of nine. Furthermore, we observe that $(x_i < x_{i+1}) = 1 - (x_i > x_{i+1})$, and similarly $(x_{i+1} < x_{i+2}) = 1 - (x_{i+1} > x_{i+2})$. Hence, membership of a triplet in \mathcal{T} can be verified using only two comparisons:

$$c_i := (x_i > x_{i+1}), \quad c_{i+1} := (x_{i+1} > x_{i+2}) .$$

In particular, the triplet (x_i, x_{i+1}, x_{i+2}) is not a member of \mathcal{T} if and only if either c_i, c_{i+1} are both zero or both one. The membership condition corresponds to the XNOR $\overline{c_i \oplus c_{i+1}}$, which we can write in arithmetic form as

$$\overline{c_i \oplus c_{i+1}} = c_i + c_{i+1} - 2c_i c_{i+1} = 1 - (c_i + c_{i+1} - 1)^2 . \quad (8)$$

Finally, we note that the next triplet $(x_{i+1}, x_{i+2}, x_{i+3})$ requires us to compute $c_{i+1} = (x_{i+1} > x_{i+2})$ again, which we can recycle from the current triplet:

$$\begin{array}{lll} (x_i, x_{i+1}, x_{i+2}) & \text{requires} & (x_i > x_{i+1}) \quad \text{and} \quad (x_{i+1} > x_{i+2}) \\ (x_{i+1}, x_{i+2}, x_{i+3}) & \text{requires} & (x_{i+1} > x_{i+2}) \quad \text{and} \quad (x_{i+2} > x_{i+3}) \\ (x_{i+2}, x_{i+3}, x_{i+4}) & \text{requires} & (x_{i+2} > x_{i+3}) \quad \text{and} \quad (x_{i+3} > x_{i+4}) \\ \vdots & & \vdots \\ \vdots & & \vdots \end{array}$$

Thus, to evaluate all triplets within a block $x_{km+1}, \dots, x_{(k+1)m}$, it is enough to compute all the adjacent comparisons:

$$(x_{km+1} > x_{km+2}), \dots, (x_{(k+1)m-1} > x_{(k+1)m}) .$$

This can be done efficiently in one batch by $C = \text{cmp}(X, X \ll 1)$, which operates on every block concurrently. To determine membership we proceed as described in Equation 8 by computing

$$F = 1 - (C + (C \ll 1) - 1)^2$$

which produces a matrix

$$f_{k,i} = \begin{cases} 1 & \text{if } r_{k,i} \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases}$$

where $r_{k,i} = \Pi(x_{km+i}, x_{km+i+1}, x_{km+i+2})$. Finally, by summing each column with SumC we can count the number of triplets in \mathcal{T} per block, resulting in a column vector $(q_1, \dots, q_{n_b})^\top$ such that

$$q_k = \sum_{i=1}^m f_{k,i}.$$

The turning rates are then computed by a plaintext multiplication by $1/(m-2)$. Algorithm 6 summarizes these steps in pseudocode.

Algorithm 6 TurningRates

Input: X encryption of $x = (x_1, \dots, x_n)$, m block size

Output: Q encryption of $q^\top = (q_1, \dots, q_{n_b})^\top$ block-wise turning rates of x

- 1: $C \leftarrow \text{cmp}(X, X \ll 1)$
 - 2: $F \leftarrow 1 - (C + (C \ll 1) - 1)^2$ ▷ pattern membership
 - 3: $Q \leftarrow \frac{1}{m-2} \cdot \text{SumC}(F)$ ▷ block-wise turning rate
 - 4: **return** Q
-

3.3 CUSUM Statistic

The block-wise summaries are now stored in the first column of a matrix as s_1, \dots, s_{n_b} and we need to compute

$$\operatorname{argmax}_{k=1, \dots, n_b} \left| \sum_{j=1}^k s_j - \frac{k}{n_b} \sum_{j=1}^{n_b} s_j \right|. \quad (9)$$

There are four main operations we need to consider here:

- the total sum $s_T := \sum_{j=1}^{n_b} s_j$,
- the partial sum $s_p^{(k)} := \sum_{j=1}^k s_j$ for each k ,
- the absolute value, and
- the argmax computation.

We address each operation as follows.

Total Sum. Computing the total sum s_T is straightforward. Since the s_i are stored as elements of the first column, it is enough to execute a SumR and the resulting s_T will end up in the top-left entry of the matrix. In preparation to the argmax computation over k elements, we replicate s_T across the first row using a ReplC operation, and then multiply element-wise by the plaintext vector

$$\left(\frac{1}{n_b}, \frac{2}{n_b}, \dots, \frac{n_b}{n_b} \right)$$

to scale it according to Equation 9.

Partial Sums. Computing the partial sum $s_p^{(k)}$ for each k is more challenging. A naive solution would compute each $s_p^{(k)}$ individually by masking all elements but the first k and summing them with a SumC, but this approach would split the operations over n_b ciphertexts. Instead, we can process all the partial sums in one go.

- (1) First, we replicate the input column vector over all columns using ReplC.

- (2) We then multiply the resulting matrix by an upper-triangular bitmask $T = (t_{i,j})$, where $t_{i,j} = 1$ if $j \geq i$ and 0 otherwise. This mask ensures that the first column contains only s_1 , the second column contains s_1 and s_2 , and in general the k -th column contains s_1, \dots, s_k .
- (3) Finally, a SumR aggregates the elements in each column, resulting in the vector of partial sums $(s_p^{(1)}, s_p^{(2)}, \dots, s_p^{(n_b)})$ in the first row of the matrix.

Absolute Value. The absolute value is not a polynomial function, hence it is not straightforward to implement under CKKS. The usual way is to compute it as $|s| = \sqrt{s^2}$, using an iterative method to approximate the square root [15]. Fortunately, we do not need to compute the absolute values explicitly. Since $\operatorname{argmax}_k |s_k| = \operatorname{argmax}_k s_k^2$, we can square the values instead and compare them directly, avoiding the expensive square root approximation.

Argmax. The argmax can be computed using the approach described in Section 2.5 (see Algorithm 3). This involves computing the comparison matrix $(s_i > s_j)$ for all i, j , summing over the rows to obtain the rankings, and applying an indicator function to select the desired rank. However, in cases where we are specifically interested in the argmax, we can significantly improve this approach. We observe that the column of the comparison matrix corresponding to the maximum entry contains only ones (and a 0.5 on the diagonal), while all other columns contain at least one zero. This means that instead of computing a full ranking and applying an expensive indicator function, we can directly isolate the argmax position by multiplying all values across the rows, that is computing a product down each column. Only the column corresponding to the maximum will result in a non-zero product, producing a one-hot encoding of the argmax.

This method allows us to replace the indicator function with a much cheaper recursive multiplication over rows, requiring only a logarithmic number of homomorphic rotations and multiplications (think of it as the multiplicative variant of SumR). To give a practical idea of the speed-up, consider the case of $n_b = 256$. In our experiments, the cmp operation costs around $v = \rho = 25$ multiplications (v) and circuit depth (ρ). Hence, the baseline method (comparison followed by ind) requires $2v + 1 = 51$ extra multiplications and $\rho + 1 = 26$ extra circuit depth, for a total of 76 multiplications and 51 circuit depth. On the other hand, our optimized method only adds $\log_2(n_b) = 8$ multiplications and depth on top of the comparison, for a total of 33 multiplications and circuit depth, resulting in a reduction of $\sim 57\%$ in the number of multiplications and $\sim 35\%$ in circuit depth.

We put these four operations together in Algorithm 7 to compute the CUSUM statistic provided in Equation 9 under encryption.

3.4 CPD Algorithms

We now combine the summarization algorithms with the CUSUM statistic to build our full CPD pipeline. Depending on the type of change to detect (mean, variance, or frequency) we use the appropriate summarization function:

- mean: uses block-wise mean (see Algorithm 4),
- variance: uses block-wise variance (see Algorithm 5),

Algorithm 7 CUSUM

Input: S encryption of $s^\top = (s_1, \dots, s_{n_b})^\top$
Output: C encryption of a one-hot encoding of the change-point location τ (block-wise)

Partial and total sums

- 1: $S_T \leftarrow \text{ReplC}(\text{SumR}(S))$
- 2: $S_P \leftarrow \text{SumR}(\text{ReplC}(S) \cdot \text{triu}(1))$
 $\quad \triangleright \text{triu}(1)$ is the upper-triangular binary mask
- 3: $U \leftarrow (S_P - (1/n_b, 2/n_b, \dots, n_b/n_b) \cdot S_T)^2$

Our revised argmax

- 4: $U_R \leftarrow \text{ReplR}(U)$
 - 5: $U_C \leftarrow \text{ReplC}(\text{TransR}(U))$
 - 6: $C = \text{cmp}(U_R, U_C) + \text{diag}(0.5)$
 $\quad \triangleright \text{diag}(0.5)$ is the diagonal matrix with 0.5 entries
 - 7: **for** $i = 0, \dots, \log n_b - 1$ **do**
 - 8: $C \leftarrow C \cdot (C \ll m \cdot 2^i)$
 - 9: **return** C
-

- frequency: uses block-wise turning rate distribution (see Algorithm 6).

The output summary is then processed with the encrypted CUSUM algorithm (Algorithm 7) to obtain the final result. See Algorithm 8 for the complete pseudocode of the encrypted CPD pipeline. Note that the resulting change-point index c_b is expressed in block units. To recover the index in the original time series, it must be scaled by the block size: $c = c_b \cdot m$.

Algorithm 8 CPD

Input: X encryption of $x = (x_1, \dots, x_n)$, m block size, t change type (mean, variance, frequency)

Output: C encryption of a one-hot encoding of the change-point location τ (block-wise)

- 1: **if** $t = \text{mean}$ **then**
 - 2: $Q \leftarrow \text{Mean}(X)$ \triangleright see Algorithm 4
 - 3: **else if** $t = \text{variance}$ **then**
 - 4: $Q \leftarrow \text{Variance}(X)$ \triangleright see Algorithm 5
 - 5: **else if** $t = \text{frequency}$ **then**
 - 6: $Q \leftarrow \text{TurningRates}(X)$ \triangleright see Algorithm 6
 - 7: $C \leftarrow \text{CUSUM}(Q)$ \triangleright see Algorithm 7
 - 8: **return** C
-

Two technical details still need to be addressed though: the normalization bounds for approximate comparisons, and the support for general values of m (lifting the assumptions we made in Section 3.1).

Normalizing Input of cmp. The homomorphic evaluation of `cmp` relies on a polynomial approximation of the sign function, approximated in the interval $[-1, 1]$. Hence, to make this work correctly, we need to ensure that all inputs passed to the `cmp` function are normalized in $[0, 1]$. Since the cost of evaluating the `cmp` depends on the precision of the approximation, obtaining the tightest possible bounds on its inputs is important for efficiency.

There are two `cmp` operations in our algorithm: one in the turning rate computation and one in the CUSUM algorithm. The first `cmp` is fine as it is directly applied to the input data, which we can

simply normalize to $[0, 1]$ under the (mild) assumption that upper and lower bounds of the series are known (e.g., heart rate between 0 and 600 bpm, EEG between -200 and $+200\mu\text{V}$). The second one (CUSUM) is more challenging. In this case, the input to the `cmp` are the values

$$\left| \sum_{j=1}^k s_j - \frac{k}{n_b} \sum_{j=1}^{n_b} s_j \right| \quad \text{for } k = 1, \dots, n_b.$$

We are providing some bounds on $\Delta_k := \sum_{j=1}^k s_j - \frac{k}{n_b} \sum_{j=1}^{n_b} s_j$. First, as the time series values x_j are assumed to be normalized in $[0, 1]$, then the s_j must also be in $[0, 1]$:

- for the mean, this is trivial;
- for the variance, you can actually see that $s_j \in [0, 1/4]$;
- for the turning rates, s_j is the number of triplets in \mathcal{T} divided by the number of total triplets, hence again $s_i \in [0, 1]$.

Since $s_T \geq s_p^{(k)}$, we have that

$$\Delta_k = s_p^{(k)} - \frac{k}{n_b} s_T \leq s_p^{(k)} - \frac{k}{n_b} s_p^{(k)} = s_p^{(k)} \left(1 - \frac{k}{n_b}\right).$$

As $s_j \in [0, 1]$, we have that $s_p^{(k)} \leq k$, and thus

$$\Delta_k \leq k \left(1 - \frac{k}{n_b}\right) = k - \frac{k^2}{n_b}.$$

The right hand side attains maximum at $k = n_b/2$, hence

$$\Delta_k \leq \frac{n_b}{2} - \frac{(n_b/2)^2}{n_b} = \frac{n_b}{2} - \frac{n_b}{4} = \frac{n_b}{4}.$$

Now, for the lower bound, we can rewrite $-\Delta_k$ as:

$$-\Delta_k = \frac{k}{n_b} \sum_{j=1}^{n_b} s_j - \sum_{j=1}^k s_j = \left(\frac{k}{n_b} - 1\right) \sum_{j=1}^k s_j + \frac{k}{n_b} \sum_{j=k+1}^{n_b} s_j.$$

Since $s_j \in [0, 1]$ and $k \leq n_b$, we have that $\left(\frac{k}{n_b} - 1\right) \sum_{j=1}^k s_j$ is a negative term and $\frac{k}{n_b} \sum_{j=k+1}^{n_b} s_j$ is a positive term. Hence, $-\Delta_k$ is maximized when $s_j = 0$ for $j \leq k$ and $s_j = 1$ for $j > k$. So,

$$-\Delta_k \leq \left(\frac{k}{n_b} - 1\right) \sum_{j=1}^k 0 + \frac{k}{n_b} \sum_{j=k+1}^{n_b} 1 = \frac{k}{n_b} (n_b - k) = k - \frac{k^2}{n_b}.$$

Following the same argument as above, we have that $-\Delta_k \leq n_b/4$. Putting everything together leads to $|\Delta_k| \leq n_b/4$, which provides a bound we can use to normalize Δ_k before the `argmax` is computed.

Handling Blocks of General Sizes. In Section 3.1 we assumed that the block size m divides the total data length n and is a power of two. If m does not divide n , then the last block has length $m' < m$. In this case, the series is padded with $m - m'$ zeros, and the summarization functions require some minor adjustments. For mean and variance, the normalization by m becomes a multiplication by $(\frac{1}{m}, \dots, \frac{1}{m}, \frac{1}{m'})^\top$. For turning rates, a masking is applied to F to mask out the first $m' - 2$ elements in the last row; and similarly, the normalization step becomes a multiplication by $(\frac{1}{m-2}, \dots, \frac{1}{m-2}, \frac{1}{m'-2})^\top$.

If m is not a power of two, the straightforward solution is to pad each row of the matrix with zeros until a power of two is reached and then adopting similar corrections as above on all rows. Although, this way we might end up with up to 50% of unused slots. To avoid this, all recursive operations on matrices `SumR`, `SumC`,

ReplR, ReplC, TransR, TransC can be easily adapted to work recursively on input sizes that are not power of two, using appropriate branching, with their cost going from $\lceil \log(m) \rceil$ to $2\lceil \log(m) \rceil$. See for instance Section 3.1 and Appendix B of [28].

4 Experimental Evaluation

We evaluate our encrypted CPD approach on both synthetic and real-world time series, reporting runtime, memory, and accuracy.

4.1 Experimental Setup

We use the CKKS implementation from the OpenFHE library [1],¹ with a scaling factor ranging between 44 and 50 bits. The ring dimension is set to 2^{17} , and all parameters are chosen in accordance with the Homomorphic Encryption Standard to guarantee 128-bit security [2, 3]. Our implementation is publicly available at <https://github.com/FedericoMazzone/secure-change-point-detection>. To compute ranking and argmax, we employ the f, g approximation of the sign function described in Section 2.5, with composition degrees $d_f = 2$ and $d_g = 4$. Since the multiplicative depth of our circuit is bounded by 65, bootstrapping is not required, and CKKS is used as a leveled scheme. Following Betken et al. [9], we set the block size to $m = \lfloor \sqrt{n} \rfloor$. All experiments are executed on a Linux machine equipped with an AMD EPYC 7763 64-Core Processor running at 2.45GHz, and 512 GB of RAM.

4.2 Synthetic Data

To evaluate our method under controlled conditions, we use synthetic data in which a change point is explicitly introduced. Given a length n and a distribution Dist , we generate a series (X_1, \dots, X_n) with a change point at τ as follows.

Mean Shift. Given $\mu_1 \neq \mu_2$ and $\sigma = \sigma_1 = \sigma_2$:

$$\begin{aligned} X_t &\sim \text{Dist}(\mu_1, \sigma^2) & \text{for } t = 1, \dots, \tau \\ X_t &\sim \text{Dist}(\mu_2, \sigma^2) & \text{for } t = \tau + 1, \dots, n \end{aligned}$$

Variance Shift. Given $\mu = \mu_1 = \mu_2$ and $\sigma_1 \neq \sigma_2$:

$$\begin{aligned} X_t &\sim \text{Dist}(\mu, \sigma_1^2) & \text{for } t = 1, \dots, \tau \\ X_t &\sim \text{Dist}(\mu, \sigma_2^2) & \text{for } t = \tau + 1, \dots, n \end{aligned}$$

Frequency Shift. Given μ, σ , and $\phi_1 \neq \phi_2$, we simulate data from an autoregressive process of order one (AR(1)):

$$\begin{aligned} X_t &= \phi_1 X_{t-1} + \varepsilon_t, & \text{for } t = 1, \dots, \tau \\ X_t &= \phi_2 X_{t-1} + \varepsilon_t, & \text{for } t = \tau + 1, \dots, n \end{aligned}$$

where $\varepsilon_t \sim \text{Dist}(\mu, \sigma^2)$ forms a sequence of mutually uncorrelated variables $(\varepsilon_t)_t$, the parameters $\phi_1, \phi_2 \in (-1, 1)$ control the strength of the temporal dependence, and X_0 is fixed.

In our experiments, the change point is placed at $\tau = \lfloor n/2 \rfloor$, Dist is selected from Gaussian, Uniform, Laplace, and Student- t distributions, and the AR(1) parameters are set to $\phi_1 = 0.3$ and $\phi_2 = 0.7$.

4.3 Real-World Data

We showcase the practical applicability of our approach on three real-world datasets. They all contain frequency shifts, which represent the most common type of change-points observed in practice.

Sleep Phase Change Detection in EEG. Electroencephalogram (EEG) signals measure brain activity and are widely used in sleep studies to identify transitions between sleep stages. These transitions are typically reflected by changes in the signal frequency. For our evaluation, we use data from the CAP Sleep Database by Terzano et al. [36], available through PhysioBank.² Following Bandt [6], we consider the overnight recording of subject 5 from the C4–P4 channel. The extracted series contains 60,000 samples (about four minutes at 256 Hz) and is annotated with expert-provided sleep stage labels, which we use as ground truth.

Heartbeat Interval Time Series. Inter-beat interval (IBI) signals measure the time between successive heartbeats and are commonly used as a non-invasive proxy for autonomic nervous system activity. For our assessment, we use data from the Meditation Task Force dataset [32], available via PhysioNet.³ During meditation, changes in the IBI series indicate shifts in heart rate frequency, linked to transitions in relaxation and stress levels. Specifically, we analyze a one-hour meditation recording from a single subject, which yields an IBI time series of length 8,083.

Network Traffic Anomaly Detection. Network flow time series are widely used to detect unusual activity patterns in communication networks. For our experiments, we use the CESNET-TimeSeries24 dataset [24], a large-scale collection of time series extracted from IP flow records in the CESNET3 academic ISP network.⁴ The dataset comprises 40 weeks of monitoring over more than 275,000 IP addresses and includes volumetric, ratio-based, and temporal statistics (e.g., number of bytes, packets, unique destination IPs, average flow duration). Changes in frequency in these series correspond to anomalies in traffic behavior. We employ the per-IP time series of the inbound-to-outbound packet ratio (`dir_ratio_packets`), aggregated over 10-minute intervals, yielding sequences of 40,298 time points.

4.4 Empirical Results

Runtime. In Figure 4, we report the runtime of our solution on synthetic data with changes in mean, variance, and frequency, for time series of increasing length (from 1,000 to 1,000,000 points). The reported values are averages over 4 runs, using a Gaussian distribution $\mathcal{N}(0, 1)$ as base. We generate the three types of changes by shifting the mean to 1, increasing the variance to 2, and setting $\phi_1 = 0.3, \phi_2 = 0.7$ for the frequency change. As expected, frequency detection takes longer than mean and variance detection, since computing the turning rate statistic requires an additional comparison, which is the bottleneck of our method. Specifically, detecting changes in mean requires between 37.91s and 101.52s, variance detection between 44.05s and 136.20s, and frequency detection between 69.38s and 179.66s. Overall, the runtime seems to increase logarithmically with the data size, as expected by the logarithmic scaling of the matrix operations in CKKS.

²<https://physionet.org/content/capslpdb/1.0.0/>

³<https://physionet.org/physiobank/database/meditation/data/>

⁴<https://zenodo.org/records/13382427>

¹<https://github.com/openfheorg/openfhe-development>

Table 1: Total and per-component runtime, communication size, and memory consumption of our approach on synthetic data.

| Type | Length | Client | | Server | | | | |
|-----------|-----------|----------------|------------|-----------------|------------------|------------|-----------|-------------|
| | | Encryption (s) | Comm. (GB) | Turn. Rates (s) | Partial Sums (s) | Argmax (s) | Total (s) | Memory (GB) |
| mean | 1,000 | 0.51 | 0.1 | - | 8.22 | 20.29 | 37.91 | 8.6 |
| mean | 10,000 | 0.55 | 0.1 | - | 11.70 | 24.34 | 48.54 | 11.3 |
| mean | 100,000 | 3.01 | 0.3 | - | 15.47 | 35.06 | 68.99 | 16.1 |
| mean | 1,000,000 | 10.69 | 1.3 | - | 16.52 | 56.98 | 101.52 | 31.6 |
| variance | 1,000 | 0.52 | 0.1 | - | 7.16 | 19.87 | 44.05 | 8.7 |
| variance | 10,000 | 0.56 | 0.1 | - | 10.83 | 22.81 | 57.17 | 11.3 |
| variance | 100,000 | 2.98 | 0.3 | - | 14.78 | 31.37 | 82.67 | 16.4 |
| variance | 1,000,000 | 11.17 | 1.4 | - | 16.66 | 53.35 | 136.20 | 32.3 |
| frequency | 1,000 | 0.85 | 0.1 | 38.16 | 6.09 | 20.34 | 69.38 | 13.2 |
| frequency | 10,000 | 0.86 | 0.1 | 40.11 | 8.63 | 23.54 | 78.27 | 17.4 |
| frequency | 100,000 | 4.92 | 0.5 | 54.67 | 11.69 | 32.95 | 107.16 | 25.6 |
| frequency | 1,000,000 | 17.46 | 2.2 | 108.80 | 12.70 | 51.20 | 179.66 | 48.9 |




Figure 4: Runtime of our approach for increasing time series length.

Table 2: Performance of plaintext frequency-change detection.

| Length | Runtime (s) | Memory (MB) |
|-----------|-------------|-------------|
| 1,000 | 0.013 | 0.026 |
| 10,000 | 0.018 | 0.292 |
| 100,000 | 0.152 | 2.491 |
| 1,000,000 | 1.551 | 24.467 |

The runtime provided in Figure 4 refers to the total processing time of the encrypted time series on the server-side. For a runtime breakdown and memory cost, we refer to Table 1, where we also include the encryption costs on the client-side. We note that the client’s encryption time and ciphertext size scale linearly with the number of data points, with the time ranging from 0.5-0.9s for 1,000 points to 10.7-17.5s for 1,000,000 points, and the ciphertext size going up to 2.2GB. In practice, the encrypted data would be streamed

in ~0.1GB ciphertext chunks, without waiting for all points to be collected. Table 2 reports the runtime and memory consumption of the plaintext baseline for frequency-change detection, which shows that the homomorphic-encryption layer adds between two and three orders of magnitude of runtime overhead and between three and six orders of magnitude of memory overhead to the underlying algorithm.

Ablation Study. To assess how the individual components of our approach improve over their naive implementations or existing solutions, we consider a time series with 10,000 points over which we want to perform frequency-shift detection. These points are encrypted into one ciphertext as 100 blocks of 100 elements each. Starting with the turning rates computation, our Algorithm 6 performs 1 SIMD comparison, $\lceil \log_2 100 \rceil + 2 = 9$ homomorphic rotations, and 2 homomorphic multiplications, leading to a runtime of 40.11s, with the comparison being the bottleneck of the computation (>90% of the cost). On the other hand, the naive approach described in Section 3.2 would extract each triplet of each block and process it separately, requiring 9800 SIMD comparisons (one for each triplet ranking), in addition to other relatively smaller costs, for an estimated runtime of ~100.45h. Improving on the naive strategy, one could pack multiple ranking computations in the same ciphertext, requiring in this case 2 SIMD comparisons, $4\lceil \log_2 100 \rceil + 4 = 32$ rotations, and 3 multiplications, for a cost of 86.85s (thus, a 2.17× speed-up for our design). The actual cost to integrate this method in the overall CPD pipeline would be higher, as the output encoding would not match what we need for computing the CUSUM statistic (see Section 3.1).

As for the partial sums, our method (line 2 of Algorithm 7) requires $2\lceil \log_2 100 \rceil = 14$ rotations and 1 multiplication, leading to a runtime of 8.63s. On the other hand, the naive solution described in Section 3.3 would require 99 multiplications to mask out each non-trivial sub-vector, $\lceil \log_2 100 \rceil$ rotations for each partial sum, hence 700 in total, resulting in a runtime of ~460s (thus, a 53.3× speed-up for our design). Lastly, as already discussed in Section 3.3, our argmax design improves over the one by Mazzone et al. [29], more than halving the number of homomorphic multiplications

Table 3: Accuracy of our approach on synthetic data with $n = 40,000$ on different distributions.

| Change Type | Distributional change | Ground-Truth | Our Approach (rel. error) | Plaintext Baseline (rel. error) |
|-------------|---|--------------|---------------------------|---------------------------------|
| mean | $\mathcal{N}(0, 1)$ to $\mathcal{N}(1, 1)$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |
| mean | $\mathcal{U}(0, 1)$ to $\mathcal{U}(1, 1)$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |
| variance | $\mathcal{N}(0, 1)$ to $\mathcal{N}(0, 2)$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |
| variance | $\mathcal{U}(0, 1)$ to $\mathcal{U}(0, 2)$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |
| frequency | $\mathcal{N}(0, 1)$ with $\phi = 0.3 \rightarrow 0.7$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |
| frequency | Lap(0, 4) with $\phi = 0.3 \rightarrow 0.7$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |
| frequency | t_5 with $\phi = 0.3 \rightarrow 0.7$ | 20,000 | 20,000 (0.00%) | 20,000 (0.00%) |

Table 4: Performance of our approach on real-world datasets for detecting frequency changes.

| Dataset | Runtime (s) | Ground-Truth | Our Approach (rel. error) | Plaintext Baseline (rel. error) |
|------------|-------------|--------------|---------------------------|---------------------------------|
| EEG | 91.19 | 46,500 | 46,360 (0.30%) | 46,604 (0.22%) |
| Meditation | 78.93 | 4,049 | 4,005 (1.09%) | 4,005 (1.09%) |
| Network | 89.19 | 23,752 | 23,800 (0.20%) | 23,800 (0.20%) |

(from 76 to 33), and leading to a runtime reduction from 54.16s to 23.54s (2.30× speed-up).

Accuracy. We evaluate the accuracy of our method against both the ground truth and the plaintext baseline. The detection error is measured using the *relative error*, defined as the normalized deviation of the estimated change point $\hat{\tau}$ from the true change point τ :

$$\text{error} := |\hat{\tau} - \tau| / \tau. \quad (10)$$

On synthetic data, our approach matches the plaintext baseline exactly across different types of distributional changes, as shown in Table 3 for $n = 40,000$. For short synthetic time series, a small relative error is introduced due to the limited number of samples. However, both the plaintext baseline and our approach produce identical estimates and therefore the same relative error. This behavior is a limitation of the statistical method and not of the encrypted computation. When applied to real-world data, our approach achieves accuracy comparable to the plaintext baseline, with only a small deviation from the ground truth (Table 4). Overall, the relative error remains minimal, typically within one percentage point, and may partly reflect uncertainty in the human-annotated ground-truth labels. Figure 6 illustrates representative real-world time series together with the corresponding turning rate, CUSUM statistic, and detected change point.

4.5 Comparison with DP-Based Solutions

The only privacy-preserving CPD solutions available in the literature [18, 38] work in the *central* DP model, where the raw data are accessible to a trusted curator that applies a randomized algorithm to produce a privacy-preserving output. In contrast, our solution operates entirely on encrypted data without requiring access to plaintext or any trusted server. This is closer to the *local* DP model, where the data curator applies noise to the data before outsourcing it to a (non-trusted) computing server. These two paradigms are fundamentally different, both in terms of trust assumptions and privacy guarantees, making a direct comparison meaningless. To

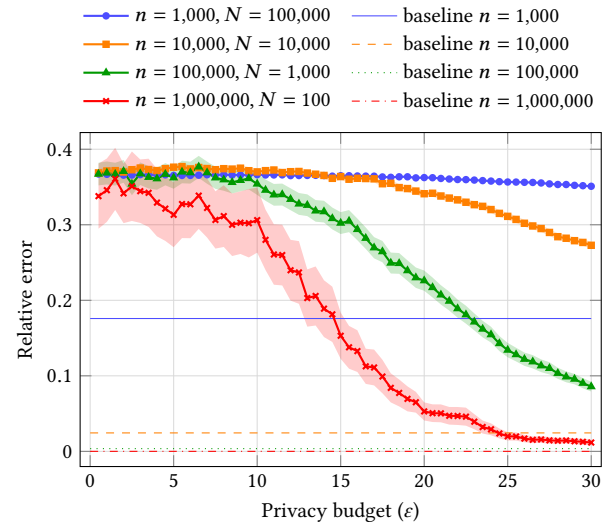



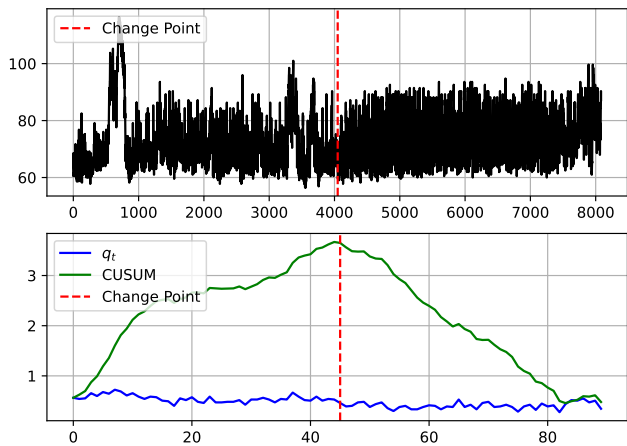
Figure 5: Relative error of the local-DP approach on four AR(1) series. The shaded areas show ± 2 SEM around the mean relative error. The horizontal dashed lines indicate the relative error of our approach as a baseline.

compare our approach against DP, we opted for implementing a simple local-DP solution ourselves.

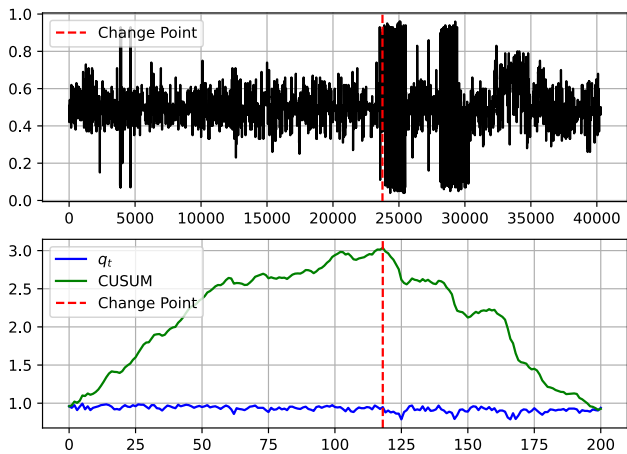
Comparison with Local-DP. Our local-DP solution operates by adding calibrated noise to a clipped version of the time series before applying a standard CUSUM procedure. The resulting privacy guarantees follow from the post-processing property of DP. For a more detailed description, we refer to Appendix B. We assess our approach against the local-DP solution on the AR(1) synthetic data, and report the results in Figure 5. The parameter δ is set to $1/n^2$ as commonly done in DP applications, and the error is computed via Equation 10 and averaged over a large number of experiments (N) to account for its



(a) Sleep Phase Change Detection in EEG



(b) Heartbeat Interval Time Series



(c) Network Traffic Anomaly Detection

Figure 6: Real-world time-series data with the detected change point (vertical dashed line).

high variability. For the clipping bound, we set $M = 1$, which empirically provides a reasonable trade-off between input information loss due to clipping and an overly large sensitivity. The runtime and memory performance of the local-DP solution are of the same order of magnitude as the plaintext baseline, with client runtime being 0.1-9.3ms, server runtime going from 3.7ms to 1.54s, and memory usage upper bounded by ~ 40 MB on both sides. However, the accuracy degrades significantly as ϵ decreases, with relative error going below 20% only for $\epsilon > 13$ (for 1,000,000 points), making local-DP impractical in scenarios where accuracy is required.

Central DP. We also include an empirical evaluation of [18] and [38]. These experiments, detailed in Appendix C, are meant to provide an indicative sense of the accuracy trade-offs under central DP guarantees and are only presented for completeness.

It is important to note that both these algorithms assume that the data consists of independent observations and that the distributions before and after the change point are known in advance. For this reason, evaluations are performed on synthetic data designed to satisfy these assumptions. As remarked throughout our work, in real-world applications, time series exhibit temporal dependence, and pre- and post-change distributions are typically unknown and must be estimated. Even when independence approximately holds, estimating the distributions introduces additional statistical noise and model misspecification.

It is also worth noting that the local DP mechanism we consider introduces noise proportional to $1/\sqrt{\epsilon}$, which is substantially larger than the $1/\epsilon$ noise scale required under central DP mechanisms. This discrepancy arises because, in our setting, privacy is enforced directly on the raw time series (effectively privatizing the identity function), and noise must be added to each individual data point. The total privacy noise thus accumulates over time, degrading the overall utility. In contrast, central DP schemes typically add noise only to aggregate statistics (e.g., means or sums), allowing for much lower noise magnitudes while still satisfying the same (ϵ, δ) -DP guarantees.

5 Conclusion

In this paper, we presented the first fully encrypted pipeline for change-point detection under homomorphic encryption. Our solution combines ordinal-pattern-based estimators with recent advances in encrypted comparison, enabling the detection of structural changes in sensitive time series with plaintext-level accuracy and without exposing the underlying data. This establishes, for the first time, the feasibility of CPD in cryptographically private settings, a direction that had remained unexplored despite the central role of CPD in statistics.

As future work, our framework points to a general methodology for outsourced statistical analysis of time series via ordinal patterns. In this direction, it will be possible to incorporate more refined statistics, such as permutation entropy or other complexity-based measures, into the encrypted domain. A particularly challenging but promising extension is the treatment of multiple change points, which remains difficult even in the plaintext setting.

Finally, this work assumes operation under the alternative hypothesis, i.e., that a change point is known to exist. Extending our

approach to support full hypothesis testing under encryption, including the null hypothesis of no change, would require securely estimating decision thresholds or p -values, and represents an important open problem. Addressing this challenge would further broaden the applicability of encrypted CPD to real-world scenarios where the presence of structural changes is not guaranteed.

Acknowledgments

We would like to thank Florian Hahn and Luca Mariot for their valuable insights. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The authors used generative AI-based tools to revise the text, improve flow and correct any typos, grammatical errors, and awkward phrasing.

References

- [1] Ahmad Al Badawi, Jack Bates, Flavio Bergamaschi, David Bruce Cousins, Saroja Erabelli, Nicholas Genise, Shai Halevi, Hamish Hunt, Andrey Kim, Yongwoo Lee, Zeyu Liu, Daniele Micciancio, Ian Quah, Yuriy Polyakov, Saraswathy R.V., Kurt Rohloff, Jonathan Saylor, Dmitriy Suponitsky, Matthew Triplett, Vinod Vaikuntanathan, and Vincent Zucca. 2022. OpenFHE: Open-Source Fully Homomorphic Encryption Library. In *Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography* (Los Angeles, CA, USA) (WAHC'22). Association for Computing Machinery, New York, NY, USA, 53–63. <https://doi.org/10.1145/3560827.3563379>
- [2] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. 2018. *Homomorphic Encryption Security Standard*. Technical Report. HomomorphicEncryption.org, Toronto, Canada.
- [3] Martin R. Albrecht, Rachel Player, and Sam Scott. 2015. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology* 9, 3 (2015), 169–203. <https://doi.org/doi:10.1515/jmc-2015-0016>
- [4] John A. D. Aston and Claudia Kirch. 2012. Evaluating stationarity via change-point alternatives with applications to fMRI data. *The Annals of Applied Statistics* 6, 4 (2012), 1906–1948. <https://doi.org/10.1214/12-AOAS565>
- [5] Juschan Bai and Pierre Perron. 1998. Estimating and Testing Linear Models with Multiple Structural Changes. *Econometrica* 66, 1 (1998), 47–78. <https://doi.org/10.2307/2998540>
- [6] Christoph Bandt. 2020. Order patterns, their variation and change points in financial time series and Brownian motion. *Stat Papers* 61 (2020), 1565–1588. <https://doi.org/10.1007/s00362-020-01171-7>
- [7] Christoph Bandt and Bernd Pompe. 2002. Permutation Entropy: A Natural Complexity Measure for Time Series. *Phys. Rev. Lett.* 88 (Apr 2002), 174102. Issue 17. <https://doi.org/10.1103/PhysRevLett.88.174102>
- [8] Michèle Basseville and Igor V. Nikiforov. 1993. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, Englewood Cliffs, NJ.
- [9] Annika Betken, Giorgio Micali, and Johannes Schmidt-Hieber. 2025. Ordinal pattern-based change point detection. *TEST* 34 (2025), 927–980. <https://doi.org/10.1007/s11749-025-00983-9>
- [10] Duncan A. J. Blythe, Paul von Bunau, Frank C. Meinecke, and Klaus-Robert Müller. 2012. Feature Extraction for Change-Point Detection Using Stationary Subspace Analysis. *IEEE Transactions on Neural Networks and Learning Systems* 23, 4 (2012), 631–643. <https://doi.org/10.1109/TNNLS.2012.2185811>
- [11] J. V. Braun, R. K. Braun, and H.-G. Müller. 2000. Multiple changepoint fitting via quasilielihood, with application to DNA sequence segmentation. *Biometrika* 87, 2 (2000), 301–314. <https://doi.org/10.1093/biomet/87.2.301>
- [12] Peter J. Brockwell and Richard A. Davis. 1991. *Time Series: Theory and Methods* (2nd ed.). Springer. <https://doi.org/10.1007/978-1-4419-0320-4>
- [13] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yongsoo Song. 2017. Homomorphic Encryption for Arithmetic of Approximate Numbers. In *Advances in Cryptology – ASIACRYPT 2017*, Tsuyoshi Takagi and Thomas Peyrin (Eds.). Springer International Publishing, Cham, 409–437.
- [14] Jung Hee Cheon, Dongwoo Kim, and Duhyeon Kim. 2020. Efficient Homomorphic Comparison Methods with Optimal Complexity. In *Advances in Cryptology – ASIACRYPT 2020*, Shihō Moriai and Huaxiong Wang (Eds.). Springer International Publishing, Cham, 221–256.
- [15] Jung Hee Cheon, Dongwoo Kim, Duhyeon Kim, Hun Hee Lee, and Keewoo Lee. 2019. Numerical Method for Comparison on Homomorphically Encrypted Numbers. In *Advances in Cryptology – ASIACRYPT 2019*, Steven D. Galbraith and Shihō Moriai (Eds.). Springer International Publishing, Cham, 415–445.
- [16] M. F. R. Chowdhury, S.-A. Selouani, and D. O’Shaughnessy. 2011. Bayesian on-line spectral change point detection: a soft computing approach for on-line ASR. *International Journal of Speech Technology* 15, 1 (2011), 5–23. <https://doi.org/10.1007/s10772-011-9116-2>
- [17] Miklós Csörgő and Lajos Horváth. 1997. *Limit Theorems in Change-Point Analysis*. John Wiley & Sons, Chichester.
- [18] Rachel Cummings, Sara Krehbiel, Yajun Mei, Rui Tuo, and Wanrong Zhang. 2018. Differentially Private Change-Point Detection. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2018/file/f19ec2b84181033bf4753a5a51d5d608-Paper.pdf
- [19] Laurie Davies, Christian Höhenrieder, and Walter Krämer. 2012. Recursive computation of piecewise constant volatilities. *Computational Statistics & Data Analysis* 56, 11 (2012), 3623–3631. <https://doi.org/10.1016/j.csda.2010.06.027> 1st issue of the Annals of Computational and Financial Econometrics Sixth Special Issue on Computational Econometrics.
- [20] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.
- [21] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407. <https://doi.org/10.1561/04000000042>
- [22] Muktesh Gupta, Rajesh Wadhvani, and Akhtar Rasool. 2024. Comprehensive analysis of change-point dynamics detection in time series data: A review. *Expert Systems with Applications* 248 (2024), 123342. <https://doi.org/10.1016/j.eswa.2024.123342>
- [23] Venkata Jandhyala, Stergios Fotopoulos, Ian MacNeill, and Pengyu Liu. 2013. Inference for single and multiple change-points in time series. *Journal of Time Series Analysis* 34, 4 (2013), 423–446. <https://doi.org/10.1111/jtsa.12035> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/jtsa.12035>
- [24] Jan Koumar, Lukas Hynek, Tomáš Čejka, et al. 2025. CESNET-TimeSeries24: Time Series Dataset for Network Traffic Anomaly Detection and Forecasting. *Scientific Data* 12 (2025), 338. <https://doi.org/10.1038/s41597-025-04603-x>
- [25] Marc Lavielle and Gilles Teyssière. 2007. *Adaptive Detection of Multiple Change-Points in Asset Price Volatility*. Springer Berlin Heidelberg, Berlin, Heidelberg, 129–156. https://doi.org/10.1007/978-3-540-34625-8_5
- [26] C. Lévy-Leduc and F. Roueff. 2009. Detection and localization of change-points in high-dimensional network traffic data. *The Annals of Applied Statistics* 3, 2 (2009), 637–662. <https://doi.org/10.1214/08-AOAS232>
- [27] A. Lung-Yut-Fong, C. Lévy-Leduc, and O. Cappé. 2012. Distributed detection/localization of change-points in high-dimensional network traffic data. *Statistics and Computing* 22, 2 (2012), 485–496. <https://doi.org/10.1007/s11222-011-9240-5>
- [28] Federico Mazzone, Trevor Brown, Florian Kerschbaum, Kevin H Wilson, Maarten Everts, Florian Hahn, and Andreas Peter. 2025. Privacy-Preserving Vertical K-Means Clustering. *arXiv preprint arXiv:2504.07578* (2025).
- [29] Federico Mazzone, Maarten Everts, Florian Hahn, and Andreas Peter. 2025. Efficient ranking, order statistics, and sorting under CKKS. In *Proceedings of the 34th USENIX Security Symposium (USENIX Security 2025)*. <https://www.usenix.org/system/files/usenixsecurity25-mazzone.pdf>
- [30] Ilya Mironov. 2017. Rényi Differential Privacy. *2017 IEEE 30th Computer Security Foundations Symposium (CSF)* (2017), 263–275. <https://api.semanticscholar.org/CorpusID:9386213>
- [31] E. S. Page. 1955. A test for a change in a parameter occurring at an unknown point. *Biometrika* 42 (1955), 523–527. <https://doi.org/10.2307/2333401>
- [32] C.-K Peng, Joseph E Mietus, Yanhui Liu, Gurucharan Khalsa, Pamela S Douglas, Herbert Benson, and Ary L Goldberger. 1999. Exaggerated heart rate oscillations during two meditation techniques. *International Journal of Cardiology* 70, 2 (1999), 101–107. [https://doi.org/10.1016/S0167-5273\(99\)00066-2](https://doi.org/10.1016/S0167-5273(99)00066-2)
- [33] J. Reeves, J. Chen, X. L. Wang, R. Lund, and Q. Q. Lu. 2007. A review and comparison of changepoint detection techniques for climate data. *Journal of Applied Meteorology and Climatology* 46, 6 (2007), 900–915. <https://doi.org/10.1175/JAM2493.1>
- [34] Robert Schmidt, Cornelius Krasselt, and Christian von Borczyskowski. 2012. Change point analysis of matrix dependent photoluminescence intermittency of single CdSe/ZnS quantum dots with intermediate intensity levels. *Chemical Physics* 406 (2012), 9–14. <https://doi.org/10.1016/j.chemphys.2012.02.018> Single molecule spectroscopy: Current status and perspectives.
- [35] M. Staudacher, S. Telser, A. Amann, H. Hinterhuber, and M. Ritsch-Marte. 2005. A new method for change-point detection developed for on-line analysis of the heart beat variability during sleep. *Physica A: Statistical Mechanics and its Applications* 349, 3 (2005), 582–596. <https://doi.org/10.1016/j.physa.2004.10.026>
- [36] M. G. Terzano, L. Parrino, A. Sherieri, R. Chervin, S. Chokroverty, C. Guilleminault, M. Hirshkowitz, M. Mahowald, H. Moldofsky, A. Rosa, R. Thomas, and A. Walters. 2001. Atlas, rules, and recording techniques for the scoring of cyclic alternating pattern (CAP) in human sleep. *Sleep Medicine* 2, 6 (Nov. 2001), 537–553. [https://doi.org/10.1016/S1526-8959\(01\)00066-2](https://doi.org/10.1016/S1526-8959(01)00066-2)

doi.org/10.1016/S1389-9457(01)00149-6 Erratum in: Sleep Med. 2002 Mar;3(2):185. PMID: 14592270.

- [37] Jan Verbesselt, Rob Hyndman, Glenn Newnham, and Darius Culvenor. 2010. Detecting trend and seasonal changes in satellite image time series. *Remote Sensing of Environment* 114, 1 (2010), 106–115. <https://doi.org/10.1016/j.rse.2009.08.014>
- [38] Wanrong Zhang, Sara Krehbiel, Rui Tuo, Yajun Mei, and Rachel Cummings. 2021. Single and Multiple Change-Point Detection with Differential Privacy. *Journal of Machine Learning Research* 22, 29 (2021), 1–36. <http://jmlr.org/papers/v22/19-770.html>

A Encrypted Matrix Operations

Here we provide some well-known algorithms for performing operations on encrypted square matrices of size N , which must be a power of two.

Algorithm 9 SumR

Input: X encryption of a square matrix of size N .

Output: X encryption of a row vector.

- 1: **for** $i = 0, \dots, \log N - 1$ **do**
 - 2: $X \leftarrow X + (X \ll N \cdot 2^i)$
 - 3: $X \leftarrow X \cdot (1^N \parallel 0^{N(N-1)})$
 - 4: **return** X
-

Algorithm 10 SumC

Input: X encryption of a square matrix of size N .

Output: X encryption of a column vector.

- 1: **for** $i = 0, \dots, \log N - 1$ **do**
 - 2: $X \leftarrow X + (X \lll N \cdot 2^i)$
 - 3: $X \leftarrow X \cdot (1 \parallel 0^{N(N-1)})^N$
 - 4: **return** X
-

Algorithm 11 ReprR

Input: X encryption of a row vector of size N .

Output: X encryption of a square matrix.

- 1: **for** $i = 0, \dots, \log N - 1$ **do**
 - 2: $X \leftarrow X + (X \gg N \cdot 2^i)$
 - 3: **return** X
-

Algorithm 12 ReprC

Input: X encryption of a column vector of size N .

Output: X encryption of a square matrix.

- 1: **for** $i = 0, \dots, \log N - 1$ **do**
 - 2: $X \leftarrow X + (X \ggg 2^i)$
 - 3: **return** X
-

Algorithm 13 TransR

Input: X encryption of a vector x encoded as a row.

Output: X encryption of the vector x encoded as a column.

- 1: **for** $i = 1, \dots, \lceil \log N \rceil$ **do**
 - 2: $X \leftarrow X + (X \ggg N(N-1)/2^i)$
 - 3: $X \leftarrow X \cdot (1 \parallel 0^{N-1})^N$
 - 4: **return** X
-

Algorithm 14 TransC

Input: X encryption of a vector x encoded as a column.

Output: X encryption of the vector x encoded as a row.

- 1: **for** $i = 1, \dots, \lceil \log N \rceil$ **do**
 - 2: $X \leftarrow X + (X \lll N(N-1)/2^i)$
 - 3: $X \leftarrow X \cdot (1^N \parallel 0^{N(N-1)})$
 - 4: **return** X
-

B Differential Privacy Background

We first report the standard definition of (ϵ, δ) -Differential Privacy, and Rényi Differential Privacy (RDP), which generalizes standard DP and is particularly useful for analyzing the privacy properties of mechanisms involving the composition of multiple queries. For a more comprehensive discussion on these concepts, we refer to Mironov (2017) [30] and Dwork and Roth (2014) [21].

Definition B.1. Let $\epsilon > 0$ and $\delta > 0$. A randomized mechanism $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}$ satisfies (ϵ, δ) -Differential Privacy (DP) if, for all adjacent datasets $S_1, S_2 \in \mathcal{Z}^n$ (i.e., differing by at most one element) and for all measurable subsets $\mathcal{S} \subseteq \mathcal{R}$, we have

$$\mathbb{P}[\mathcal{A}(S_1) \in \mathcal{S}] \leq e^\epsilon \mathbb{P}[\mathcal{A}(S_2) \in \mathcal{S}] + \delta.$$

The parameter ϵ controls the privacy loss, with smaller values ensuring stronger privacy guarantees, while δ accounts for the probability of violating pure ϵ -DP.

Next, the Rényi Differential Privacy is a relaxation of (ϵ, δ) -DP based on Rényi divergence, which facilitates tighter privacy accounting, especially under composition.

Definition B.2. For two probability distributions f and g supported over \mathcal{R} , the Rényi divergence of order $\beta > 1$ is defined as

$$D_\beta(f \| g) := \frac{1}{\beta - 1} \log \mathbb{E}_{x \sim g} \left[\left(\frac{f(x)}{g(x)} \right)^\beta \right].$$

Definition B.3. (β, ϵ) -Rényi Differential Privacy (RDP). A randomized mechanism $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathcal{R}$ is said to be ϵ -Rényi differentially private of order β (denoted as (β, ϵ) -RDP) if

$$\sup_{S_1 \sim S_2} D_\beta(\mathcal{A}(S_1) \| \mathcal{A}(S_2)) \leq \epsilon. \quad (11)$$

Properties of Rényi Differential Privacy

We conclude this section by summarizing three fundamental properties of RDP [30]:

- (1) **Post-processing:** If \mathcal{A} is (β, ϵ) -RDP and g is a randomized mapping, then $g \circ \mathcal{A}$ is also (β, ϵ) -RDP.
- (2) **Adaptive Composition:** If $\mathcal{A}_1 : \mathcal{Z}^n \rightarrow \mathcal{X}_1$ is (β, ϵ_1) -RDP, and $\mathcal{A}_2 : \mathcal{X}_1 \times \mathcal{Z}^n \rightarrow \mathcal{X}_2$ is (β, ϵ_2) -RDP, then the combined

mechanism

$$(\mathcal{A}_1(\cdot), \mathcal{A}_2(\mathcal{A}_1(\cdot), \cdot))$$

satisfies $(\beta, \varepsilon_1 + \varepsilon_2)$ -RDP.

(3) **From RDP to DP:** If \mathcal{A} is (β, ε) -RDP, then it satisfies

$$\left(\varepsilon + \frac{\log(1/\delta)}{\beta - 1}, \delta \right)\text{-DP}$$

for any $0 < \delta < 1$.

B.1 Releasing Time Series

Assumption Let X_t admit a probability density $f \in C_c^0(\mathbb{R}) \cap L^\infty(\mathbb{R})$, i.e. f is continuous, bounded, and compactly supported.

As a consequence, X_t is almost surely bounded, meaning there exists $M > 0$ such that

$$\mathbb{P}(|X_t| < M) = 1.$$

Hence, the support of the data vector $\mathbb{X} := (X_1, \dots, X_n)$ is given by

$$\mathcal{Z}^n = [-M, M]^n.$$

Note that Assumption 1 is an analogous of the truncation condition imposed by Zhang et al (Section 2.1 in [38]).

Definition B.4 (Additive Gaussian mechanism). For $\sigma > 0$, let $(\varepsilon_j)_{j=1}^n$ be i.i.d. $\mathcal{N}(0, \sigma^2)$ random variables. We define the mechanism $\mathcal{A} : \mathcal{Z}^n \rightarrow \mathbb{R}^n$ by

$$\mathcal{A}(\mathbb{X}) := \mathbb{X} + \boldsymbol{\varepsilon} \quad (12)$$

where $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_n)$ and $\mathbf{1} = (1, \dots, 1)^\top$. In particular, each coordinate satisfies

$$(\mathcal{A}(\mathbb{X}))_i \sim \mathcal{N}(X_i, \sigma^2),$$

independently across i .

PROPOSITION B.5 (GAUSSIAN RDP, EQUAL VARIANCES). Let $\beta > 1$ and $P = \mathcal{N}(\mu_1, \sigma^2)$, $Q = \mathcal{N}(\mu_2, \sigma^2)$. Then

$$D_\beta(P \| Q) = \frac{\beta}{2\sigma^2} (\mu_1 - \mu_2)^2.$$

COROLLARY B.6 (RDP OF \mathcal{A}). For neighboring datasets \mathbb{X}, \mathbb{X}' differing only at index k , the mechanism (12) satisfies

$$D_\beta(\mathcal{A}(\mathbb{X}) \| \mathcal{A}(\mathbb{X}')) = \frac{\beta}{2\sigma^2} (X_k - X'_k)^2.$$

Since $|X_k - X'_k| \leq 2M$, we obtain the uniform bound

$$D_\beta(\mathcal{A}(\mathbb{X}) \| \mathcal{A}(\mathbb{X}')) \leq \frac{2\beta M^2}{\sigma^2}.$$

PROOF. For two neighboring datasets \mathbb{X}, \mathbb{X}' that differ only in the k -th entry, it follows that

$$\begin{aligned} D_\beta(\mathcal{A}(\mathbb{X}) \| \mathcal{A}(\mathbb{X}')) &= \frac{1}{\beta - 1} \log \mathbb{E}_{x_i \sim \mathcal{N}(x_i, \sigma^2)} \left(\frac{\prod_i \mathcal{N}(x_i, \sigma^2)}{\prod_i \mathcal{N}(x'_i, \sigma^2)} \right) \\ &= \frac{1}{\beta - 1} \log \mathbb{E}_{x_k \sim \mathcal{N}(x_k, \sigma^2)} \left(\frac{\mathcal{N}(x_k, \sigma^2)}{\mathcal{N}(x'_k, \sigma^2)} \right) \\ &= D_\beta((\mathcal{A}(\mathbb{X}))_k \| (\mathcal{A}(\mathbb{X}'))_k). \end{aligned}$$

□

THEOREM B.7 ((ε, δ)-DP GUARANTEE). For any $\delta \in (0, 1)$, the mechanism \mathcal{A} defined in (12) is (ε, δ) -differentially private with

$$\varepsilon = \frac{\Delta^2}{2\sigma^2} + \frac{\sqrt{2}\Delta}{\sigma} \sqrt{\log\left(\frac{1}{\delta}\right)}, \quad (13)$$

where $\Delta := \sup_{\mathbb{X} \sim \mathbb{X}'} |X_k - X'_k| \leq 2M$ denotes the global ℓ_2 -sensitivity. In particular,

$$\varepsilon \leq \frac{2M^2}{\sigma^2} + \frac{2\sqrt{2}M}{\sigma} \sqrt{\log\left(\frac{1}{\delta}\right)}.$$

PROOF. From Corollary B.6, the mechanism is $(\beta, \varepsilon_{\text{RDP}}(\beta))$ -RDP with

$$\varepsilon_{\text{RDP}}(\beta) = \frac{\beta\Delta^2}{2\sigma^2}.$$

The standard RDP-to-DP conversion gives, for any $\beta > 1$,

$$\varepsilon = \varepsilon_{\text{RDP}}(\beta) + \frac{\log(1/\delta)}{\beta - 1}.$$

Optimizing in β yields $\beta^* = 1 + \sqrt{\frac{2\sigma^2 \log(1/\delta)}{\Delta^2}}$, which implies the stated bound. □

COROLLARY B.8. Let $(X_t)_{t=1, \dots, n}$ be a time series satisfying Assumption 1. For any $\varepsilon > 0$ and any $\delta \in (0, 1)$, the mechanism \mathcal{A} defined in (12) with $\sigma_{\text{DP}} = \frac{\sqrt{2}M}{\varepsilon} (\sqrt{\log(1/\delta)} + \sqrt{\log(1/\delta) + \varepsilon})$ is (ε, δ) -differentially private.

PROOF. Let $L := \log(1/\delta)$. Solving $\varepsilon = \Delta^2/(2\sigma^2) + (\sqrt{2}\Delta/\sigma)\sqrt{L}$ for σ gives

$$\sigma = \frac{\Delta}{\sqrt{2}\varepsilon} (\sqrt{L} + \sqrt{L + \varepsilon})$$

With $\Delta = 2M$ yields σ_{DP} as stated in the Corollary. □

B.2 Local DP-CPD

Corollary B.8 specifies the amount of noise required for the mechanism \mathcal{A} to satisfy (ε, δ) -differential privacy. This guarantee holds under Assumption 1, which imposes bounded sensitivity on the input time series. When this assumption is not fulfilled—i.e., when the time series may contain unbounded or heavy-tailed values—it is sufficient to apply a componentwise Huber-type clipping transformation:

$$\phi_M(X_t) = \begin{cases} X_t, & \text{if } |X_t| \leq M, \\ M, & \text{if } X_t > M, \\ -M, & \text{if } X_t < -M, \end{cases}$$

for some truncation level $M > 0$.

We then define the transformed process $Y_t = \phi_M(X_t)$, and apply a CUSUM-type statistic to $\{Y_t\}$. Since differential privacy is preserved under post-processing, the overall procedure remains (ε, δ) -DP. That is, if an algorithm is (ε, δ) -DP, then any deterministic function applied to its output also satisfies (ε, δ) -DP, see [21].

C Central DP CPD simulations

This section reports results for Algorithm 3 of Zhang et al. [38] and Algorithm 1 of [18]. We evaluate both methods on synthetic time series with a single change in the data-generating distribution. For a series of length n and a fixed fraction $\tilde{\tau} \in (0, 1)$, the true

change-point is $\tau = \lfloor \tilde{\tau}n \rfloor$, and the observations are independent with a distributional break at τ :

$$X_t \sim \mathcal{N}(\mu_1, \sigma_1^2) \quad \text{for } t = 1, \dots, \tau$$

$$X_t \sim \text{Lap}(\mu_2, \sigma_2) \quad \text{for } t = \tau + 1, \dots, n$$

where the post-change Laplace law is parameterized so that σ_2 denotes its standard deviation (equivalently, scale $b = \sigma_2/\sqrt{2}$).

Following the protocol of Section 4.5, for each setting we simulate N independent series $\{X_t^{(i)}\}_{t=1}^n$, apply both methods, and record the estimated change $\hat{k}^{(i)}$. Performance is summarized by the mean absolute relative error. All remaining experimental choices (privacy calibration, grids over ϵ , etc.) match Section 4.5.

We report the accuracy of both methods in the case of change of variance in Figure 7, and in the case of change of both mean and variance in Figure 8.




Figure 7: Change of variance: $\sigma_1 \neq \sigma_2$ and $\mu_1 = \mu_2$.




Figure 8: Change of mean and variance: $\sigma_1 \neq \sigma_2$ and $\mu_1 \neq \mu_2$.