

# Uncertainty Quantification for Data-Driven Change-Point Learning via Cross-Validation

Hui Chen<sup>1\*</sup>, Yinxu Jia<sup>2\*</sup>, Guanghui Wang<sup>3\*†</sup>, Changliang Zou<sup>2\*</sup>

<sup>1</sup>School of Mathematics and Statistics, Jiangsu Normal University

<sup>2</sup>School of Statistics and Data Science, LPMC, KLMDASR, and LEBPS, Nankai University

<sup>3</sup>School of Statistics, Academy of Statistics and Interdisciplinary Sciences, and KLASADS-MOE, East China Normal University

chenhuiah@163.com, yxjia@mail.nankai.edu.cn, ghwang.nk@gmail.com, zoucl@nankai.edu.cn

## Abstract

Accurately detecting multiple change-points is critical for various applications, but determining the optimal number of change-points remains a challenge. Existing approaches based on information criteria attempt to balance goodness-of-fit and model complexity, but their performance varies depending on the model. Recently, data-driven selection criteria based on cross-validation has been proposed, but these methods can be prone to slight overfitting in finite samples. In this paper, we introduce a method that controls the probability of overestimation and provides uncertainty quantification for learning multiple change-points via cross-validation. We frame this problem as a sequence of model comparison problems and leverage high-dimensional inferential procedures. We demonstrate the effectiveness of our approach through experiments on finite-sample data, showing superior uncertainty quantification for overestimation compared to existing methods. Our approach has broad applicability and can be used in diverse change-point models.

## Introduction

Change-point detection plays a crucial role in diverse domains, including machine learning and statistics (Aue and Horváth 2013; Niu, Hao, and Zhang 2016; Aminikhanghahi and Cook 2017; Fearnhead and Rigai 2020; Truong, Oudre, and Vayatis 2020; Cho and Kirch 2021). Consider a scenario where we have collected a sequence of independent data observations from a parametric multiple change-point model. In this model, the data is divided into  $K_n + 1$  pieces, each adhering to a parametric model. Importantly, the parameters differ between successive pieces, representing variations such as the mean or variance of the data, or regression coefficients illustrating the relationship between a response and covariates.

Among the primary challenges in the multiple change-point model is the estimation of the number of change-points. Existing approaches involve selecting an optimal threshold or penalty, where the ideal value may depend on the specific model at hand. For instance, in binary segmentation and its

variants (Venkatraman 1992; Fryzlewicz 2014; Harchaoui, Moulines, and Bach 2008; Li et al. 2015), each step of the partitioning algorithm requires the determination of a threshold to decide when to stop segment partitioning. Nonetheless, identifying the appropriate threshold poses a significant challenge. Similarly, in the case of penalized minimization algorithms (Yao 1988; Bai and Perron 1998; Braun, Braun, and Müller 2000; Hannart and Naveau 2012; Killick, Fearnhead, and Eckley 2012; Zou et al. 2014; Haynes, Eckley, and Fearnhead 2017; Haynes, Fearnhead, and Eckley 2017; Wang, Zou, and Yin 2018; Cho and Kirch 2022; Baranowski, Chen, and Fryzlewicz 2019; Fearnhead, Maidstone, and Letchford 2019), it becomes crucial to carefully select the penalized parameters.

In recent developments, Zou, Wang, and Li (2020) introduced a data-driven selection criterion utilizing cross-validation (CV) to address the issue of threshold or penalty selection. This criterion is versatile, as it can be applied alongside various change-point detection algorithms and is suitable for different parametric change-point models. Building upon this work, Pein and Shah (2021) further refined the cross-validation criterion with a specific focus on detecting changes that involve large magnitudes.

Cross-validation has been widely used as a technique to estimate the prediction error of a model (Allen 1974; Stone 1974; Geisser 1975). The fundamental concept behind cross-validation is to fit and evaluate candidate models on separate datasets to obtain an unbiased performance evaluation. However, traditional cross-validation methods, such as leave-one-out and V-fold variants, have limitations when it comes to model selection. These methods can suffer from overfitting, which leads to inaccurate model selection. Theoretical studies by Shao (1993); Zhang (1993); Yang (2007) have shown that cross-validation may not consistently select the correct model in low-dimensional linear models unless the training-testing split ratio tends to zero. Unfortunately, achieving such ideal split ratios is not useful in real-world applications. Recent advances in this field by Austern and Zhou (2020); Bayle et al. (2020) have focused on studying the asymptotic distribution of the cross-validated risk under certain stability conditions.

The cross-validation criterion in the context of change-point detection problems introduced by Zou, Wang, and Li (2020) demonstrated that the criterion can lead to consistent

\*These authors contributed equally, and are listed in alphabetical order.

†Guanghui Wang is the corresponding author.  
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

selection of the number of change-points. However, achieving consistency between the estimated number of change-points, denoted as  $\hat{K}_{CV}$ , and the true  $K_n$  may require stringent conditions for the change-point detection algorithms. In practical applications, this criterion can be susceptible to slight overfitting. To gain further insights into this phenomenon, Table 1 provides a measure of  $\hat{K}_{CV} - K_n$  for two commonly used change-point detection algorithms: wild binary segmentation (WBS) (Fryzlewicz 2014) and the pruned exact linear time algorithm (PELT) (Killick, Fearnhead, and Eckley 2012). The results indicate a slight overestimation, which can be beneficial as it helps capture important change signals, preventing them from being missed.

The main objective of this paper is to assess the level of uncertainty associated with cross-validation in change-point detection. We are fortunate that the slight overestimation of the cross-validation criterion can be measured and managed in real-world scenarios, with theoretical assurance under less stringent conditions compared to achieving the consistency of  $\hat{K}_{CV}$ . This research accomplishes this by showcasing that the likelihood of overestimation while utilizing the cross-validation criterion can be adequately controlled. Through these findings, we contribute to a better understanding of the cross-validation criterion in change-point detection.

In recent years, there has been a growing interest in quantifying the uncertainty associated with the number of change-points. Frick, Munk, and Sieling (2014) proposed a methodology based on multi-scale statistics to detect change-points while controlling the probability of overestimating at least one change-point. Expanding on this approach, Li, Munk, and Sieling (2016) refined this methodology by incorporating false discovery rate control. More recently, Chen et al. (2023) proposed a data-driven approach that leverages globally symmetric statistics to achieve false discovery rate control. These approaches quantify uncertainty for the number of change-points via error rate control in multiple testing framework. On the other hand, the method in Frick, Munk, and Sieling (2014) is distribution-dependent and cannot keep the probability of overestimation in the case of more complex data.

Our work is inspired by the ‘‘cross-validation with confidence’’ method recently introduced by Lei (2020). Their method incorporates cross-validation to establish a framework for quantifying the uncertainty related to selected models or tuning parameters. In a similar vein, our work aims to provide a measure of uncertainty for cross-validation estimates in the context of change-point detection.

We focus on effectively quantifying the difference  $\hat{K}_{CV} - K_n$  with a certain confidence level. The key idea is to find a lower bound,  $\hat{K}_{\min}$ , for  $K_n$  that provides confidence. To control the overestimation of  $\hat{K}_{CV}$ , we can utilize the difference,  $\hat{K}_{CV} - \hat{K}_{\min}$ . In order to find  $\hat{K}_{\min}$ , we conduct a sequence of hypothesis tests for each candidate number of change-points. The null hypothesis is that a specified model estimated from a candidate  $r$  has the smallest prediction error among the candidates larger than  $r$ . This test is conducted sequentially and individually for each candidate until we accept the null hypothesis for the first time. To obtain a valid critical value, we implement a high-dimensional mean testing

Algorithm	Measure	$n$					
		300	400	500	600	700	800
WBS	$\hat{K}_{CV} - K_n$	0.7	0.8	0.5	0.5	0.4	0.4
	$U$	0.9	0.9	0.5	0.5	0.4	0.4
	$P_+$	2.0	1.0	1.5	1.5	1.5	1.0
PELT	$\hat{K}_{CV} - K_n$	0.4	0.3	0.3	0.3	0.3	0.3
	$U$	0.5	0.4	0.4	0.3	0.3	0.3
	$P_+$	1.0	1.0	0.0	0.0	0.0	0.0

Table 1: Average of  $\hat{K}_{CV} - K_n$  and  $U = \hat{K}_{CV} - \hat{K}_{\min} > 0$ , and empirical overestimation probability  $P_+ := \mathbb{P}(\hat{K}_{CV} - K_n > U)$  (in %), for two change-point detection algorithms WBS and PELT, across different sample sizes  $n$ . The nominal level is  $\alpha = 10\%$ . The data is generated from Gaussian distribution with a signal-to-noise ratio 1 and  $K_n = 5$ ; see the experimental section.

procedure that tests for prediction error discrepancy for a new sample.

Table 1 presents the average values of  $\hat{K}_{CV} - K_n$  and  $U = \hat{K}_{CV} - \hat{K}_{\min}$ , and the empirical overestimation probability  $P_+ := \mathbb{P}(\hat{K}_{CV} - K_n > U)$  based on 500 replications. The results are obtained for two change-point detection algorithms WBS and PELT, across different sample sizes  $n$ . It is observed that the data-driven cross-validation estimate  $\hat{K}_{CV}$  indicates slight overestimation and we can control the overestimation.

The main contributions and advantages of our proposal are as follows:

- **Uncertainty quantification for cross-validation criterion:** Our proposal quantifies the uncertainty of cross-validation in change-point detection (Zou, Wang, and Li 2020). By utilizing hypothesis testing and model comparison techniques, we can control the probability of overestimation within a specified level.
- **Weaker requirements for control:** Our research demonstrates that achieving control over the probability of overestimation can be achieved with weaker requirements on the change-point detection algorithms compared to those needed for consistency, as evidenced by the work of Zou, Wang, and Li (2020). This finding highlights the robustness and reliability of the cross-validation criterion in change-point detection.
- **Versatility:** Our method is versatile and can be used in conjunction with various change-point detection algorithms. It is also suitable for different change-point models, making it adaptable to a wide range of applications and scenarios.

**Notations.** Denote  $\mathbf{1}(\cdot)$  as the indicator function. Let  $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$  and  $\{\mathbf{y}_1, \dots, \mathbf{y}_n\}$  be two sets of  $d$ -variate vectors. Define the norm  $\|\mathbf{x}\| = \sqrt{\mathbf{x}^\top \mathbf{x}}$ . For any interval  $(a, b]$  with  $0 \leq a < b \leq n$ , denote  $\bar{\mathbf{x}}_{a,b} = (b - a)^{-1} \sum_{i=a+1}^b \mathbf{x}_i$ . Let  $\mathcal{T}_r = (\tau_1, \dots, \tau_r)$  be a set of  $r$  points such that  $0 < \tau_1 <$

$\dots < \tau_r < n$ . We denote

$$\mathcal{C}_{\mathbf{xy}}(\mathcal{T}_r) = \sum_{j=0}^r (\tau_{j+1} - \tau_j) \bar{\mathbf{x}}_{\tau_j, \tau_{j+1}}^\top \bar{\mathbf{y}}_{\tau_j, \tau_{j+1}}$$

where  $\tau_0 = 0$  and  $\tau_{r+1} = n$ . Note that  $\mathcal{C}_{\mathbf{x}}^2 = \mathcal{C}_{\mathbf{xx}}$ .

## Methodology

### The Problem

Suppose we have a sequence of independent data observations  $\boldsymbol{\xi} = \{\boldsymbol{\xi}_1, \dots, \boldsymbol{\xi}_n\}$  from a parametric multiple change-point model:

$$\boldsymbol{\xi}_i \sim \mathcal{G}(\cdot | \boldsymbol{\beta}_k^*), \tau_k^* < i \leq \tau_{k+1}^*, 0 \leq k \leq K_n, 1 \leq i \leq n, \quad (1)$$

where  $K_n$  is the true number of change-points,  $\tau_k^*$ 's are the locations of change-points with the convention of  $\tau_0^* = 0$  and  $\tau_{K_n+1}^* = n$ ,  $\boldsymbol{\beta}_k^*$  is a  $d$ -dimensional parameter vector of interest and  $\mathcal{G}(\cdot | \boldsymbol{\beta}_k^*)$  represents the model structure of the segment  $k$  satisfying  $\boldsymbol{\beta}_k^* \neq \boldsymbol{\beta}_{k+1}^*$ . In this model, we use the notation  $\boldsymbol{\xi}_i$  which could either be a  $d$ -variate random vector or  $(Y_i, \mathbf{X}_i)$  with  $Y_i$  and  $\mathbf{X}_i$  being respectively the response variable and  $d$ -variate explanatory variable.

**Problem** (Uncertainty quantification for change-point detection via cross-validation)

Learn the change-points specific to the task at hand by employing the cross-validation criterion proposed by Zou, Wang, and Li (2020), and obtain an estimated number of change-points  $\widehat{K}_{CV}$ ;

Aim to demonstrate that the probability of overestimation, defined as  $\mathbb{P}(\widehat{K}_{CV} - K_n > U)$ , can be effectively controlled under a prescribed level  $0 < \alpha < 1$ . Here  $U$  is a quantity to be determined.

Selecting an appropriate value for  $U$  is crucial to effectively control the level of overestimation. We approach this problem by formulating a series of hypothesis testing problems as follows:

$$H_{0,r} : K_n = r, \quad H_{1,r} : K_n > r, \quad 0 \leq r \leq p_n, \quad (2)$$

where  $p_n$  represents an upper bound on the number of change-points. We initiate the process with  $r = 0$  and test the hypothesis (2) using the desired level  $\alpha$ . If we reject  $H_{0,r}$ , we increment  $r$  by 1 and proceed to the next hypothesis. This iterative process continues until we accept  $H_{0,r}$  for the first time. The resulting value of  $r$  is denoted as  $\widehat{K}_{\min}$ . We will show that, by choosing  $U = \widehat{K}_{CV} - \widehat{K}_{\min} \geq 0$ ,

$$\mathbb{P}(\widehat{K}_{CV} - K_n > U) \leq \alpha.$$

The null hypothesis being tested is that the true number of change-points in the model is equal to  $r$ . To test this hypothesis, we compare the prediction error discrepancy between a model with  $r$  change-points and other models with  $s$  change-points with  $s > r$ . The prediction error discrepancy is denoted as  $\Delta_{r,s}$ .

If  $r$  represents the optimal number of change-points, we expect the model with  $r$  change-points to have a lower prediction error compared to models with a larger number of

change-points. In other words, we anticipate that the maximum value of  $\Delta_{r,s}$  when  $s > r$  will be less than or equal to 0.

Formally, we consider the following hypothesis:

$$H_{0,r} : \max_{s>r} \Delta_{r,s} \leq 0, \quad H_{1,r} : \max_{s>r} \Delta_{r,s} > 0. \quad (3)$$

This testing problem in fact involves a sequence of  $(p_n - r - 1)$ -dimensional vectors  $(\Delta_{r,s} : s > r)$ , for  $r = 0, 1, \dots$ , where  $p_n$  can diverge to infinity as the sample size  $n \rightarrow \infty$ . Hence some high-dimensional inferential tools may be leveraged for designing a valid testing procedure.

If the null hypothesis  $H_{0,r}$  is rejected, it suggests that adding more change-points to the model could potentially improve prediction performance. In such a case, it becomes necessary to consider models with a greater number of change-points for a more accurate prediction.

### Testing Procedure

Consider the general parametric model (1). To estimate the underlying model and evaluate the prediction error based on different data, we employ a sample-splitting strategy. That is, we use one part of the data to train the model and another part to compute the prediction error. Inspired by the order-preserved splitting strategy proposed by Zou, Wang, and Li (2020), we partition the data into training set  $\boldsymbol{\xi}^{\text{tr}}$  and validation set  $\boldsymbol{\xi}^{\text{te}}$  with  $n_{\text{tr}}$  and  $n_{\text{te}}$  observations respectively, according to whether the observed index being odd or even, such that the two data sets share a similar change-point pattern as much as possible.

Next, we can apply some commonly used change-point detection algorithm  $\mathcal{A}(\cdot)$  to learn  $r$  change-points based on the data set  $\boldsymbol{\xi}^{\text{tr}}$ , and obtain the set of change-points  $\widehat{\mathcal{T}}_r = (\widehat{\tau}_{r,1}, \dots, \widehat{\tau}_{r,r})$ . In the meantime, we obtain the estimated parameters, i.e.,  $\widehat{\mathcal{B}} = \{\widehat{\boldsymbol{\beta}}_0, \dots, \widehat{\boldsymbol{\beta}}_r\}$ . The resulting model is denoted as  $\widehat{\mathcal{U}}_r^{\text{tr}} = \{\widehat{\mathcal{T}}_r, \widehat{\mathcal{B}}\}$ . Correspondingly, for  $\boldsymbol{\xi}_i \in \boldsymbol{\xi}^{\text{te}}$ , we introduce the prediction error  $\mathbb{E}\{\mathcal{O}(\boldsymbol{\xi}_i; \widehat{\mathcal{U}}_r^{\text{tr}}) | \boldsymbol{\xi}^{\text{tr}}\}$ , where

$$\mathcal{O}(\boldsymbol{\xi}_i; \widehat{\mathcal{U}}_r^{\text{tr}}) = \sum_{j=0}^r \ell(\boldsymbol{\xi}_i, \widehat{\boldsymbol{\beta}}_j) \mathbf{1}(\widehat{\tau}_{r,j} < i \leq \widehat{\tau}_{r,j+1}),$$

here  $\ell(\cdot, \cdot)$  is a loss function. Take the classic quadratic loss as an example. Under the mean shift model or the linear regression model with structural breaks, we have  $\ell(\boldsymbol{\xi}_i, \widehat{\boldsymbol{\beta}}_j) = \|\boldsymbol{\xi}_i - \widehat{\boldsymbol{\beta}}_j\|^2$  or  $\ell(\boldsymbol{\xi}_i, \widehat{\boldsymbol{\beta}}_j) = (Y_i - \mathbf{X}_i^\top \widehat{\boldsymbol{\beta}}_j)^2$ , respectively. Our method can be used beyond the quadratic loss, such as the quantile loss or other robust loss functions.

For any  $s > r$ , define the discrepancy of the loss for  $\boldsymbol{\xi}_i$  between Model  $r$  and  $s$  by

$$\delta_{r,s}^{(i)} = \mathcal{O}(\boldsymbol{\xi}_i; \widehat{\mathcal{U}}_r^{\text{tr}}) - \mathcal{O}(\boldsymbol{\xi}_i; \widehat{\mathcal{U}}_s^{\text{tr}}), \quad i \in \mathcal{I}_{\text{te}},$$

where  $\mathcal{I}_{\text{te}}$  is the index set of  $\boldsymbol{\xi}^{\text{te}}$ . Correspondingly, we define the discrepancy of the prediction error between Model  $r$  and  $s$ ,  $\Delta_{r,s} = \sum_{i \in \mathcal{I}_{\text{te}}} \mathbb{E}(\delta_{r,s}^{(i)} | \boldsymbol{\xi}^{\text{tr}}) / n_{\text{te}}$ .

We adopt the test statistic

$$\mathcal{D} = \max_{s>r} \sqrt{n_{\text{te}}} \widehat{\Delta}_{r,s} / \widehat{\sigma}_{r,s},$$

where  $\widehat{\Delta}_{r,s}$  and  $\widehat{\sigma}_{r,s}^2$  are the sample mean and variance of the sequence  $\{\delta_{r,s}^{(i)} : i \in \mathcal{I}_{te}\}$ , respectively.

To approximate the null distribution of  $\mathcal{D}$ , we use a multiplier bootstrap approach. For  $b = 1, \dots, B$ , we generate independent and identically distributed (iid) random variables  $e_i \sim N(0, 1)$ , for  $i \in \mathcal{I}_{te}$ . Let

$$\mathcal{D}_b = \max_{s>r} n_{te}^{-1/2} \sum_{i \in \mathcal{I}_{te}} (\delta_{r,s}^{(i)} - \widehat{\Delta}_{r,s}) e_i / \widehat{\sigma}_{r,s}.$$

We obtain the critical value

$$c_{r,\alpha} = \inf \left\{ t \in \mathbb{R} : B^{-1} \sum_{b=1}^B \mathbf{1}(\mathcal{D}_b > t) \leq \alpha \right\}, \quad (4)$$

and then reject  $H_{0,r}$  if  $\mathcal{D} > c_{r,\alpha}$ . If we reject  $H_{0,r}$ , we increment  $r$  by 1 and proceed to the next hypothesis. This iterative process continues until we accept  $H_{0,r}$  for the first time. The smallest value of  $r$  for which we do not reject  $H_{0,r}$  is called  $\widehat{K}_{\min}$ .

### Refinement with Cross-Validation

To improve the stability and mitigate the power loss due to sample-splitting, we introduce the V-fold cross-validation. Define the discrepancy of the loss for  $\xi_i$  between Model  $r$  and  $s$  by

$$\delta_{r,s}^{(i)} = \mathcal{O}(\xi_i; \widehat{\mathcal{U}}_r^{(-v)}) - \mathcal{O}(\xi_i; \widehat{\mathcal{U}}_s^{(-v)}), \quad i \in \mathcal{I}_v, \quad v = 1, \dots, V,$$

where  $\mathcal{I}_v$  is the index set of the  $v$ -th fold, and  $\widehat{\mathcal{U}}_r^{(-v)}$  denotes the model fitted by the data  $\xi^{(-v)}$  with the  $v$ -th fold removing from the full data. Accordingly, the discrepancy of the prediction error between Model  $r$  and  $s$  for the  $v$ -th fold is

$$\Delta_{r,s}^{(v)} = n^{-1} V \sum_{i \in \mathcal{I}_v} \mathbb{E}(\delta_{r,s}^{(i)} \mid \xi^{(-v)}).$$

Denote  $\Delta_{r,s} = V^{-1} \sum_{v=1}^V \Delta_{r,s}^{(v)}$ . We propose the following steps to implement the hypothesis testing for (3). To construct the test statistic, we introduce the estimate of  $\Delta_{r,s}$ ,  $\widehat{\Delta}_{r,s} = V^{-1} \sum_{v=1}^V \widehat{\Delta}_{r,s}^{(v)}$ , here  $\widehat{\Delta}_{r,s}^{(v)} = n^{-1} V \sum_{i \in \mathcal{I}_v} \delta_{r,s}^{(i)}$ .

We denote the group-wise centered discrepancy of predictive loss between Model  $r$  and  $s$  by  $\widetilde{\delta}_{r,s}^{(i)} = \delta_{r,s}^{(i)} - \widehat{\Delta}_{r,s}^{(v)}$ ,  $i \in \mathcal{I}_v$ . Let  $\widehat{\sigma}_{r,s}$  be the sample standard deviation of  $\{\widetilde{\delta}_{r,s}^{(i)} : 1 \leq i \leq n\}$ . We propose the statistic

$$\widetilde{\mathcal{D}} = \max_{s>r} \sqrt{n} \widehat{\Delta}_{r,s} / \widehat{\sigma}_{r,s}.$$

Then we reject  $H_{0,r}$  if  $\widetilde{\mathcal{D}} > c_{r,\alpha}$ , where  $c_{r,\alpha}$  is the critical value of the test.

To derive  $c_{r,\alpha}$ , we also generate iid standard Gaussian random variables  $e_i$ ,  $1 \leq i \leq n$ , accordingly

$$\widetilde{\mathcal{D}}_b = \max_{s>r} n^{-1/2} \sum_{v=1}^V \sum_{i \in \mathcal{I}_v} (\widetilde{\delta}_{r,s}^{(i)} - \widehat{\Delta}_{r,s}^{(v)}) e_i / \widehat{\sigma}_{r,s}$$

for  $b = 1, \dots, B$ . We obtain the critical value  $c_{r,\alpha}$  according to (4) with  $\mathcal{D}_b$  replaced by  $\widetilde{\mathcal{D}}_b$ .

## Theoretical Analysis

In this section, we present theoretical results of our method under the mean shift model. That is,  $d$ -variate observations  $\xi_i$ 's follow

$$\xi_i = \mu_i + \varepsilon_i, \quad i = 1, \dots, n,$$

where  $\varepsilon_i$  is a  $d$ -dimensional random vector with mean zero, and  $\mu$  is piecewise constant with change-points  $\tau_1^* < \dots < \tau_{K_n}^*$ ; that is  $\mu_i \neq \mu_{i+1}$  if and only if  $i = \tau_k^*$  for some  $k$  and  $i = 1, \dots, n-1$ . For  $k = 0, \dots, K_n$ ,  $\beta_k^* = \mathbb{E}(\xi_i)$ ,  $\Sigma_k = \text{cov}(\varepsilon_i)$  for  $\tau_k^* < i \leq \tau_{k+1}^*$ . Denote the maximum and minimum eigenvalues of  $\Sigma_k$  by  $\bar{\sigma}(\Sigma_k)$  and  $\underline{\sigma}(\Sigma_k)$  for  $k = 0, \dots, K_n$ . Let  $\bar{\sigma} = \max\{\bar{\sigma}(\Sigma_0), \dots, \bar{\sigma}(\Sigma_{K_n})\}$  and  $\underline{\sigma} = \min\{\underline{\sigma}(\Sigma_0), \dots, \underline{\sigma}(\Sigma_{K_n})\}$ . The maximal and minimal distance between change points are denoted as  $\bar{\lambda}_n = \max_{0 \leq k \leq K_n} (\tau_{k+1}^* - \tau_k^*)$  and  $\underline{\lambda}_n = \min_{0 \leq k \leq K_n} (\tau_{k+1}^* - \tau_k^*)$ . Let  $\gamma_n^{(k)} = \|\beta_{k-1}^* - \beta_k^*\|^2$  for  $1 \leq k \leq K_n$  and the minimal signal strength be  $\underline{\gamma}_n = \min_{1 \leq k \leq K_n} \|\beta_{k-1}^* - \beta_k^*\|^2$ . Without loss of generality, we use the order-preserved splitting strategy to split the data based on odd and even indices. For convenience, assume that the sample size  $n$  is an even number. We denote

$$\xi_i^{\text{tr}} = \mu_i^{\text{tr}} + \varepsilon_i^{\text{tr}}, \quad i \in \mathcal{I}_{\text{tr}}, \quad \xi_i^{\text{te}} = \mu_i^{\text{te}} + \varepsilon_i^{\text{te}}, \quad i \in \mathcal{I}_{\text{te}}.$$

The quadratic loss function for the model  $\widehat{\mathcal{U}}_r^{\text{tr}}$  for  $j = 0, \dots, r$  and  $\widehat{\tau}_{r,j} < i \leq \widehat{\tau}_{r,j+1}$  is of the form

$$\ell(\xi_i^{\text{te}}, \widehat{\beta}_j) = \|\xi_i^{\text{te}} - \widehat{\xi}_{\widehat{\tau}_{r,j}, \widehat{\tau}_{r,j+1}}^{\text{tr}}\|^2.$$

For other change-point models, we can use scores and transform the problem to the mean domain (Zou, Wang, and Li 2020).

Let  $\Delta_{r,s} = \mathbb{E}(\widehat{\Delta}_{r,s} \mid \xi^{\text{tr}})$  and  $\sigma_{r,s}^2 = \mathbb{E}(\widehat{\sigma}_{r,s}^2 \mid \xi^{\text{tr}})$ , where  $\widehat{\Delta}_{r,s}$  and  $\widehat{\sigma}_{r,s}^2$  have been defined in the methodology section.

Let  $\eta_{r,s}^{(i)} = \sigma_{r,s}^{-1} (\delta_{r,s}^{(i)} - \Delta_{r,s})$  for  $i \in \mathcal{I}_{te}$  and  $0 \leq r < s \leq p_n$ .

We need the following assumption to facilitate our theoretical demonstration.

**Assumption 1 (Moments)** (i) *There exists some  $\vartheta > 2$  such that  $\mathbb{E}(|\eta_{r,s}^{(i)}|^\vartheta) < \infty$  for every  $i \in \mathcal{I}_{te}$ ;*

(ii) *There exists some constants  $\nu \in (0, 1/2)$  and  $C > 0$  such that*

$$(M_{n,3}^3 \vee M_{n,4}^2 \vee B_n)^2 \log^{7/2} p_n \leq C n^{1/2-\nu}, \quad (5)$$

where  $M_{n,k} = \max_{i,r} \max_{s>r} \mathbb{E}^{1/k}(|\eta_{r,s}^{(i)}|^k)$  for  $k = 3, 4$  and  $B_n = \max_{i,r} \mathbb{E}^{1/4}\{\max_{s>r}(\eta_{r,s}^{(i)})^4\}$ .

Assumption 1 is motivated by Lei (2020) and puts restrictions on the moments of  $\delta_{r,s}^{(i)}$ . Lei (2020) in fact assumed sub-exponential tail conditions. Both kinds of conditions are common in developing high-dimensional central limit theorems and justifying the validity of the procedures; see Chernozhukov, Chetverikov, and Kato (2017, 2019) for details.

**Theorem 1** *Suppose that Assumptions 1 holds and for  $s > K_n$ ,*

$$\begin{aligned} & \{\mathcal{C}_{\xi^{\text{tr}}}^2(\widehat{\mathcal{T}}_{K_n}) - \mathcal{C}_{\xi^{\text{tr}}}^2(\widehat{\mathcal{T}}_s)\} \\ & - 2\{\mathcal{C}_{\mu^{\text{tr}} \xi^{\text{tr}}}(\widehat{\mathcal{T}}_{K_n}) - \mathcal{C}_{\mu^{\text{tr}} \xi^{\text{tr}}}(\widehat{\mathcal{T}}_s)\} \leq 0 \end{aligned} \quad (6)$$

holds with probability one. Then  $\mathbb{P}(\widehat{K}_{CV} - K_n > U) \leq \alpha + o(1)$ .

**Remark 1** Theorem 1 shows that the probability of overestimation  $\mathbb{P}(\widehat{K}_{CV} - K_n > U)$  can be effectively controlled for the sample-splitting procedure. The main step of the proof is to show that the hypothesis  $H_{0,r}$  with the true value of  $r = K_n$  will not be rejected. The condition (6) puts some requirements on the change-point detection algorithm when the model is overfitting. It is weaker than the condition (10) imposed in Zou, Wang, and Li (2020) to achieve consistent selection of the number of change-points. Since we are testing the discrepancy of two prediction errors conditional on the training set by (3), our method eliminates the randomness from the testing set. On the contrary, the cross-validation procedure to render consistency requires that the randomness of the training set overrides the randomness from the testing set (Zou, Wang, and Li 2020). To see this, consider the univariate sequence with  $K_n = 0$  (no change-point). Suppose that the number of change-points is overestimated and  $s = 1$ . In this case, it can be verified that the second term of (6) is 0. Some simple algebra shows that the first term of (6) is of the form

$$C_{\xi^{tr}}^2(\widehat{T}_0) - C_{\xi^{tr}}^2(\widehat{T}_1) = -\frac{\widehat{\tau}(n_{tr} - \widehat{\tau})}{n_{tr}} (\bar{\xi}_{0,\widehat{\tau}}^{tr} - \bar{\xi}_{\widehat{\tau},n_{tr}}^{tr})^2 \leq 0,$$

where  $\widehat{\tau}$  is the estimated location of the change-point under  $s = 1$ . Hence, Condition (6) holds immediately. Furthermore, assume there exists one true change-point  $\tau^*$ . Suppose we identify a change-point at  $\widehat{\tau}$  for  $s = 1$ , and find two change-points  $\widehat{\tau}_1$  and  $\widehat{\tau}_2$  with  $\widehat{\tau}_1 < \widehat{\tau}_2$  for  $s = 2$ . Consider  $\tau^* \leq \widehat{\tau} \leq \widehat{\tau}_1$ . Condition (6) holds by observing that

$$(\widehat{\tau}_1^{-1} - \widehat{\tau}^{-1})\{\gamma_n^{(1)}\}^2 + C_{\xi^{tr}}^2(\widehat{T}_1) - C_{\xi^{tr}}^2(\widehat{T}_2) \leq 0.$$

In contrast, to achieve consistency, Condition (10) in Zou, Wang, and Li (2020) generally requires that  $-\{C_{\xi^{tr}}^2(\widehat{T}_{K_n}) - C_{\xi^{tr}}^2(\widehat{T}_s)\}$  is of larger order than  $\log \log n$ .

**Proposition 1** (i) If  $\alpha < 0.5$ ,  $\lim_{B \rightarrow \infty} \mathbb{P}(\widehat{K}_{CV} \geq \widehat{K}_{\min}) = 1$ .

(ii) Suppose that the  $\Sigma_j$ 's are positive-definite matrices and there exists a positive integer  $m \geq 2$  such that  $\mathbb{E}(\|\Sigma_j^{-1/2} \varepsilon_i\|^{2m}) < \infty$  for  $i \in (\tau_j^*, \tau_{j+1}^*]$ ,  $j = 0, \dots, K_n$ . If the jump sizes  $\|\beta_{k-1}^* - \beta_k^*\}$ 's satisfy

$$\frac{\lambda_n \gamma_n}{K_n \bar{\sigma} \lambda_n^{2/m}} \rightarrow \infty, \tag{7}$$

then we have  $\lim_{n \rightarrow \infty} \mathbb{P}(\widehat{K}_{CV} \geq K_n) = 1$ .

Proposition 1 ensures that  $U \geq 0$  and  $\widehat{K}_{CV} - K_n \geq 0$ , which makes our control of overestimation probability more useful. Assumptions in Proposition 1–(ii) are from Zou, Wang, and Li (2020), which put some requirements on the minimal signal strength and distance between two change-points to avoid underestimation.

## Experiments

In this section, we demonstrate the performance of the proposed methods through numerical studies. The experiments

are run on a personal computer with an Intel Core i7-10700 CPU, 8GB of memory, a 64-bit operating system and R software version 4.2.1.

## Synthetic Data

**Data.** Take the model  $\xi_i = \mu_i + \sigma \varepsilon_i$  as an example. We set the signal vector  $\theta = (-1, 1)$ , and let  $\theta_{k-1}$  denote the  $k$ -th element of  $\theta$ . Take  $\mu_i = \theta_{\text{mod}(\mathcal{J}_0+j, 2)}$  for  $\tau_j^* < i \leq \tau_{j+1}^*$ ,  $j = 0, \dots, K_n$ , where  $\mathcal{J}_0$  is an integer randomly sampling from  $\{1, 2\}$  and  $\text{mod}(\cdot, \cdot)$  is the modulo operator. The signal-to-noise (SNR) is defined as  $\text{sd}(\mu_i\text{'s})/\sigma$ , where  $\text{sd}(\mu_i\text{'s})$  is the standard deviation of the signals  $\{\mu_1, \dots, \mu_n\}$ .

The error terms (i)  $\varepsilon_i \stackrel{\text{iid}}{\sim} N(0, 1)$ ; (ii)  $\varepsilon_i \stackrel{\text{iid}}{\sim} \sqrt{0.6}t(5)$ ,  $i = 1, \dots, n$  are considered, where  $t(5)$  is  $t$  distribution with 5 degrees of freedom. The change-points  $\tau_j^* = j \lfloor n/(K_n + 1) \rfloor + \text{Uniform}(-a, a)$  with  $a = \lfloor n^{1/4} \rfloor$  for  $j = 1, \dots, K_n$ , where  $\text{Uniform}(a, b)$  is the continuous uniform distribution with support  $[a, b]$ . All the simulation results are based on 500 replications and the bootstrap sample size is  $B = 500$ .

**Methods.** In the experiments, we evaluate the performance of our sample-splitting-based method (SS) for quantifying the uncertainty of overestimation of cross-validation criterion. We also investigate the refinement of our method using multiple folds cross-validation (RC) as discussed in the methodology section. For illustration purposes, we conduct the experiments using a three-fold cross-validation approach.

In this study, we examine the cross-validation criterion in Zou, Wang, and Li (2020) in combination with the WBS and PELT algorithms. We use the implementations of these algorithms provided in the R packages "wbs" (Baranowski and Fryzlewicz 2019) and "changepoint" (Killick, Haynes, and Eckley 2016; Killick and Eckley 2014). The method we propose, which utilizes sample-splitting along with the PELT algorithm, is referred to as SS-PELT. The other methods are named following a similar convention. The approach introduced by Frick, Munk, and Sieling (2014) (abbreviated as SMUCE) is taken into account as a benchmark method. SMUCE is implemented in the R package `stepR` (Pein et al. 2019).

**Performance Measures.** A key metric is the empirical overestimation probability  $P_+$ . It is essential for  $P_+$  to be lower than the specified nominal level  $\alpha$ , demonstrating controlled overestimation. Additionally, we will emphasize the average values of  $\widehat{K} - K_n$ , which should be slightly larger than 0 (with smaller values being more desirable), where  $\widehat{K}$  denotes the estimates  $\widehat{K}_{\text{SMUCE}}$  and  $\widehat{K}_{\text{CV}}$  for SMUCE and CV. This emphasizes that  $\widehat{K}$  is a reliable estimator of the true number of change-points  $K_n$  and prevents underestimation.

**Results.** Table 2 presents a comparative analysis between our proposed method and the benchmark approach SMUCE under normal and  $t(5)$  error term distributions. Our method consistently maintains  $P_+$  within the specified  $\alpha$  for both distributions. However, the benchmark, while controlling  $P_+$  for normal errors, exhibits instances of  $\widehat{K} - K_n$  being smaller than 0, indicating slight underestimation. Notably, for error terms with a  $t(5)$  distribution, the benchmark results in significantly higher  $P_+$  than the nominal level.




Figure 1: Slope heuristics (left panel) of cell line GM05296 with the red solid line denoting the sample mean in each segmentation (right panel).

Error	$K_n$	Method	$\alpha = 5\%$		$\alpha = 10\%$		$\alpha = 20\%$	
			$P_+$	$\widehat{K} - K_n$	$P_+$	$\widehat{K} - K_n$	$P_+$	$\widehat{K} - K_n$
$N(0, 1)$	25	SMUCE	0.0	0.0	0.0	0.0	0.0	0.0
		SS-PELT	0.0	0.6	1.0	0.6	3.0	0.6
	35	SMUCE	0.0	-0.2	0.0	-0.1	1.0	-0.1
		SS-PELT	1.0	0.5	2.0	0.5	3.0	0.5
$t(5)$	25	SMUCE	63.5	1.8	77.5	2.4	83.0	3.5
		SS-PELT	5.5	0.3	9.0	0.3	11.5	0.3
	35	SMUCE	66.0	1.8	71.0	2.3	77.5	2.8
		SS-PELT	5.0	0.7	11.5	0.7	19.5	0.7

Table 2: Performance study of SMUCE and our method SS-PELT, with  $n = 1000$  and  $\text{SNR} = 1.2$ . Here the estimates  $\widehat{K}_{\text{SMUCE}}$  and  $\widehat{K}_{\text{CV}}$  for SMUCE and CV are denoted by  $\widehat{K}$ . For SMUCE,  $P_+$  (in %) is the empirical version of  $\mathbb{P}(\widehat{K}_{\text{SMUCE}} > K_n)$ .

Table 3 assesses the performance of our methods combined with different change-point detection algorithms across various settings. All methods demonstrate the capability to control  $P_+$  at different nominal levels  $\alpha$ .

### Real Data

Change-point detection is useful for detecting changes in a wide range of applications, such as bioscience, image analysis, and so on. One area where change-point detection is commonly applied is in DNA copy number analysis. For example, in array Comparative Genomic Hybridization (array CGH) data analysis, change-point detection is used to identify segments of the genome where there is a change in the copy number of DNA. The copy number refers to the number of copies of a specific DNA segment in a cell. By detecting these changes, researchers can identify regions of

the genome that may be associated with genetic diseases or abnormalities.

Consider the array comparative genomic hybridization (CGH) data from the coriell dataset available in the R package `DNAcopy` in Seshan and Olshen (2019). Here the goal is to quantify the changes in terms of DNA copy number. We focus on the chromosomal aberration locations detection and consider cell line GM05296 for illustration, which can be treated as a univariate mean shift detection problem (see the left panel in Figure 1). There are totally 2112 observations after removing na values.

Here we refer to Lavielle (2005) and Haynes, Eckley, and Fearnhead (2017), which suggest a method for determining the optimal number of change-points in a dataset by looking at how the cost changes as the number of change-points increases. The ‘‘elbow’’ (red triangular point) in the right panel of Figure 1 is used to locate the optimal number of change-points. We note that the cost should decrease more when detecting true changes, and in this case, the true number of change-points is determined to be 5.

We investigate the performance of  $\widehat{K} - K_n$  for both our proposed method (CV-WBS) and the benchmark approach SMUCE. The result of CV-WBS is 1. The benchmark yields a result of 18 under a target level  $\alpha$  of 0.05, 0.1 and 0.2. Though both methods consistently yield nonnegative values for  $\widehat{K} - K_n$ , CV-WBS provides significantly smaller values compared to the benchmark, indicating slight overestimation.

### Discussion

In this study, we have presented a novel approach for quantifying uncertainty of cross-validation criterion in learning multiple change-points. Through the utilization of the model comparison framework and advanced high-dimensional inferential tools, we have established a robust methodology that allows us to control the probability of overestimation within a specified level. Our findings contribute to a better understanding of the reliability and robustness of the cross-validation criterion in change-point detection. Experimental

SNR	$K_n$	Method	$\alpha = 5\%$		$\alpha = 10\%$		$\alpha = 20\%$		$\widehat{K}_{CV} - K_n$
			$P_+$	$U$	$P_+$	$U$	$P_+$	$U$	
0.9	15	SS-WBS	3.0	0.9(1.6)	6.0	0.7(1.6)	7.0	0.7(1.5)	0.6(1.4)
		SS-PELT	0.0	0.9(1.3)	2.5	0.8(1.3)	2.0	0.4(0.9)	0.4(0.9)
		RC-PELT	0.0	0.5(1.0)	0.0	0.4(0.6)	0.0	0.3(0.6)	0.3(0.9)
25	15	SS-WBS	4.5	3.2(3.7)	10.5	2.5(3.2)	15.5	2.1(3.2)	1.5(2.5)
		SS-PELT	1.0	3.6(3.9)	2.5	2.3(2.4)	7.0	1.7(2.6)	0.6(1.2)
		RC-PELT	0.0	2.0(2.1)	0.0	1.2(2.0)	0.0	0.8(1.6)	0.5(1.1)
1.2	15	SS-WBS	0.3	0.7(1.4)	3.0	0.6(1.2)	5.0	0.5(1.2)	0.6(1.4)
		SS-PELT	0.7	0.4(1.0)	1.5	0.4(0.8)	2.3	0.4(1.1)	0.3(1.0)
		RC-PELT	0.0	0.4(0.8)	0.0	0.3(0.8)	0.0	0.3(0.7)	0.4(1.0)
25	15	SS-WBS	4.5	1.3(2.3)	4.0	1.0(1.9)	7.0	0.9(1.8)	0.8(1.7)
		SS-PELT	0.0	0.9(1.5)	1.0	0.7(1.3)	3.0	0.6(1.3)	0.4(1.0)
		RC-PELT	0.0	0.6(1.1)	0.0	0.5(0.9)	0.0	0.4(0.8)	0.4(0.9)

Table 3: Performance study of the average value of  $U$ ,  $\widehat{K}_{CV} - K_n$  and the empirical probability  $P_+$  (in %), with  $n = 1000$  under Gaussian distribution. Standard deviations are listed in the brackets.

results demonstrate the superiority of our proposed approach in providing accurate uncertainty quantification for overestimation compared to existing methods in finite samples.

One limitation of our method is its assumption of parametric and low-dimensional change-point models. Adapting the method to high-dimensional or nonparametric models and exploring its theoretical properties would be a valuable avenue for future research. Additionally, the assumption of independent observations may not be suitable for handling dependent data. Developing customized cross-validation frameworks and inferential tools for dependent data is an important direction to consider. By addressing these limitations and extending the method to encompass high-dimensional, non-parametric, and possibly correlated models, we can enhance its practical applicability and gain deeper insights into uncertainty quantification for cross-validation criterion in change-point detection in diverse real-world scenarios.

### Acknowledgments

The authors would like to express their sincere appreciation to the anonymous reviewers for their valuable comments and constructive feedback. Zou was supported by the National Key R&D Program of China (Grant Nos. 2022YFA1003703, 2022YFA1003800) and the National Natural Science Foundation of China (Grant Nos. 11925106, 12231011, 11931001, 11971247). Chen was supported by the National Natural Science Foundation of China (Grant No. 12201257). Wang was supported by the National Natural Science Foundation of China (Grant No. 32030063) and the Natural Science Foundation of Shanghai (Grant No. 23ZR1419400).

### References

Allen, D. M. 1974. The relationship between variable selection and data augmentation and a method for prediction.

*Technometrics*, 16(1): 125–127.

Aminikhanghahi, S.; and Cook, D. J. 2017. A survey of methods for time series change point detection. *Knowledge and Information Systems*, 51(2): 339–367.

Aue, A.; and Horváth, L. 2013. Structural breaks in time series. *Journal of Time Series Analysis*, 34(1): 1–16.

Austern, M.; and Zhou, W. 2020. Asymptotics of cross-validation. arXiv:2001.11111.

Bai, J.; and Perron, P. 1998. Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1): 47–78.

Baranowski, R.; Chen, Y.; and Fryzlewicz, P. 2019. Narrowest-over-threshold detection of multiple change points and change-point-like features. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(3): 649–672.

Baranowski, R.; and Fryzlewicz, P. 2019. *wbs: Wild Binary Segmentation for Multiple Change-Point Detection*. R package version 1.4.

Bayle, P.; Bayle, A.; Janson, L.; and Mackey, L. 2020. Cross-validation Confidence Intervals for Test Error. In *Advances in Neural Information Processing Systems*, volume 33, 16339–16350. Curran Associates, Inc.

Braun, J. V.; Braun, R. K.; and Müller, H. G. 2000. Multiple changepoint fitting via quasilikelihood, with application to DNA sequence segmentation. *Biometrika*, 87(2): 301–314.

Chen, H.; Ren, H.; Yao, F.; and Zou, C. 2023. Data-driven selection of the number of change-points via error rate control. *Journal of the American Statistical Association*, 118(542): 1415–1428.

Chernozhukov, V.; Chetverikov, D.; and Kato, K. 2017. Central limit theorems and bootstrap in high dimensions. *The Annals of Probability*, 45(4): 2309–2352.

- Chernozhukov, V.; Chetverikov, D.; and Kato, K. 2019. Inference on causal and structural parameters using many moment inequalities. *The Review of Economic Studies*, 86(5): 1867–1900.
- Cho, H.; and Kirch, C. 2021. Data segmentation algorithms: Univariate mean change and beyond. *Econometrics and Statistics*.
- Cho, H.; and Kirch, C. 2022. Two-stage data segmentation permitting multiscale change points, heavy tails and dependence. *Annals of the Institute of Statistical Mathematics*, 74(4): 653–684.
- Fearnhead, P.; Maidstone, R.; and Letchford, A. 2019. Detecting changes in slope with an  $L_0$  penalty. *Journal of Computational and Graphical Statistics*, 28(2): 265–275.
- Fearnhead, P.; and Rigaiil, G. 2020. Relating and Comparing Methods for Detecting Changes in Mean. *Stat*, e291.
- Frick, K.; Munk, A.; and Sieling, H. 2014. Multiscale change point inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(3): 495–580.
- Fryzlewicz, P. 2014. Wild binary segmentation for multiple change-point detection. *The Annals of Statistics*, 42(6): 2243–2281.
- Geisser, S. 1975. The predictive sample reuse method with applications. *Journal of the American Statistical Association*, 70(350): 320–328.
- Hannart, A.; and Naveau, P. 2012. An improved Bayesian information criterion for multiple change-point models. *Technometrics*, 54(3): 256–268.
- Harchaoui, Z.; Moulines, E.; and Bach, F. 2008. Kernel Change-point Analysis. In *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.
- Haynes, K.; Eckley, I. A.; and Fearnhead, P. 2017. Computationally efficient changepoint detection for a range of penalties. *Journal of Computational and Graphical Statistics*, 26(1): 134–143.
- Haynes, K.; Fearnhead, P.; and Eckley, I. A. 2017. A computationally efficient nonparametric approach for changepoint detection. *Statistics and Computing*, 27: 1293–1305.
- Killick, R.; and Eckley, I. 2014. changepoint: An R package for changepoint analysis. *Journal of Statistical Software*, 58(3): 1–19.
- Killick, R.; Fearnhead, P.; and Eckley, I. A. 2012. Optimal detection of changepoints with a linear computational cost. *Journal of the American Statistical Association*, 107(500): 1590–1598.
- Killick, R.; Haynes, K.; and Eckley, I. 2016. *changepoint: An R package for changepoint analysis*. R package version 2.2.2.
- Lavielle, M. 2005. Using penalized contrasts for the change-point problem. *Signal processing*, 85(8): 1501–1510.
- Lei, J. 2020. Cross-validation with confidence. *Journal of the American Statistical Association*, 115(532): 1978–1997.
- Li, H.; Munk, A.; and Sieling, H. 2016. FDR-control in multiscale change-point segmentation. *Electronic Journal of Statistics*, 10(1): 918–959.
- Li, S.; Xie, Y.; Dai, H.; and Song, L. 2015. M-Statistic for Kernel Change-Point Detection. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Niu, Y. S.; Hao, N.; and Zhang, H. 2016. Multiple change-point detection: a selective overview. *Statistical Science*, 31(4): 611–623.
- Pein, F.; Hotz, T.; Sieling, H.; and Aspelmeier, T. 2019. *stepR: Multiscale change-point inference*. R package version 2.0-4.
- Pein, F.; and Shah, R. D. 2021. Cross-validation for change-point regression: pitfalls and solutions. arXiv:2112.03220.
- Seshan, V. E.; and Olshen, A. 2019. *DNAcopy: DNA copy number data analysis*. R package version 1.60.0.
- Shao, J. 1993. Linear model selection by cross-validation. *Journal of the American statistical Association*, 88(422): 486–494.
- Stone, M. 1974. Cross-validated choice and assessment of statistical predictions. *Journal of the Royal Statistical Society: Series B (Methodological)*, 36(2): 111–133.
- Truong, C.; Oudre, L.; and Vayatis, N. 2020. Selective review of offline change point detection methods. *Signal Processing*, 167: 107299.
- Venkatraman, E. S. 1992. *Consistency results in multiple change-point problems*. Ph.D. thesis, Stanford University.
- Wang, G.; Zou, C.; and Yin, G. 2018. Change-point detection in multinomial data with a large number of categories. *The Annals of Statistics*, 46(5): 2020–2044.
- Yang, Y. 2007. Consistency of cross validation for comparing regression procedures. *The Annals of Statistics*, 35(6): 2450–2473.
- Yao, Y.-C. 1988. Estimating the number of change-points via Schwarz’ criterion. *Statistics and Probability Letters*, 6(3): 181–189.
- Zhang, P. 1993. Model selection via multifold cross validation. *The Annals of Statistics*, 21(1): 299–313.
- Zou, C.; Wang, G.; and Li, R. 2020. Consistent selection of the number of change-points via sample-splitting. *The Annals of Statistics*, 48(1): 413–439.
- Zou, C.; Yin, G.; Feng, L.; and Wang, Z. 2014. Nonparametric maximum likelihood approach to multiple change-point problems. *The Annals of Statistics*, 42(3): 970–1002.