Control Chart for Monitoring Processes with Recurrent Shifts (Running Title: Monitoring Processes with Recurrent Shifts)

Peihua Qiu[1], Parichita Nandi[1] and Qinglin Pei[2]

[1]Department of Biostatistics, University of Florida

[2]College of Nursing, University of Florida

Gainesville, FL 32610, USA

Corresponding Author: Peihua Qiu, E-mail: pqiu@ufl.edu

Funding: The authors received no specific funding for this research.

Abstract

Control charts are widely used for sequential monitoring, but most existing methods are designed to detect only the first distributional shift because they were originally developed for manufacturing processes that are stopped after an out-of-control (OC) signal. In many real-world applications, however, such as economic monitoring, environmental surveillance, and disease surveillance, the monitored process continues to operate and may repeatedly alternate between in-control (IC) and OC states. This paper investigates online monitoring of such sequential processes with recurrent shifts (SPRS), where timely detection of both the onset and the termination of OC periods is essential for initiating and terminating interventions efficiently. We introduce new concepts and propose a general online monitoring framework for SPRS. For clarity, the methodology is presented for serially independent univariate normal observations, although the framework can be extended to more general settings. The proposed approach provides a foundation for sequential monitoring of processes exhibiting recurrent distributional shifts.

Key Words: Control charts; Delay times; Multiple shifts; Phase II monitoring; Sequential decision making; Statistical process control.

1 Introduction

Control charts are a primary analytical tool for the online monitoring of sequential processes (Hawkins and Olwell, 1998; Qiu, 2013). They have been widely applied in manufacturing, environmental monitoring, disease surveillance, and many other fields (Montgomery, 2012; Qiu, 2024;

Xie et al., 2024). Most existing control charts are designed to detect the first distributional shift in a sequential process of key quality variables. However, in many applications, the underlying process may experience recurrent distributional shifts, making the timely detection of these repeated changes critically important. This paper introduces new concepts and proposes a new charting scheme for the online (Phase II) monitoring of sequential processes with recurrent shifts (SPRS).

Control charts were originally developed for monitoring manufacturing production processes to detect out-of-control (OC) conditions as early as possible after a process departs from its in-control (IC) state. Since Walter A. Shewhart introduced the first control chart a century ago, numerous control charts have been developed in the literature, including various forms of the Shewhart chart (Shewhart, 1931), the cumulative sum (CUSUM) chart (Page, 1954), the exponentially weighted moving average (EWMA) chart (Roberts, 1959), and the change-point detection (CPD)-based chart (Hawkins et al., 2003). Early control charts were developed under the conventional assumptions that in-control observations are independent and identically distributed (i.i.d.) over time and follow a specified parametric distribution, such as the normal distribution. As sequrntial process monitoring methods have been increasingly applied to fields such as environmental monitoring and disease surveillance, substantial research has been devoted to addressing challenges arising from serially correlated observations (Apley and Tsung, 2002; Runger, 2002), time-varying IC distributions (Qiu, 2024; Xie et al., 2024), and situations in which the IC distribution cannot be adequately characterized by a prespecified parametric model (Chakraborti and Graham, 2019; Qiu, 2018).

The control charts described above are all designed to detect the first distributional shift in a monitored process. In manufacturing, once such a shift occurs in key quality variables, product quality becomes questionable. The production process is then typically halted so that the root causes of the shift can be identified and eliminated before production resumes. Otherwise, subsequent products may fail to meet quality specifications, resulting in defective products and wasted resources. This operational practice explains why conventional control charts are designed primarily to detect the first distributional shift.

However, many real-world processes cannot be easily stopped and may repeatedly alternate between IC and OC states. For example, economic indicators such as the unemployment rate are collected and released sequentially, requiring timely identification of persistent departures from normal economic conditions. Once an OC condition is signaled, the labor market and the national

economy cannot simply be halted. Nevertheless, timely detection of both the onset and the termination of OC periods is crucial for enabling governments to initiate interventions promptly and to terminate them once they are no longer needed. After the process returns to the IC state, it may subsequently enter another OC period, making recurrent shifts a natural feature of the underlying process. Similar data streams with recurrent shifts arise in many other applications, including monitoring air pollutant concentrations in environmental systems and disease incidence rates in infectious disease surveillance.

To the best of our knowledge, no existing research has directly addressed the problem of online monitoring of sequential processes with recurrent shifts (SPRS). In the change-point detection literature, numerous methods have been developed for detecting multiple change points (cf. Gao et al., 2024; Hawkins, 2001; Niu et al., 2016; Zou et al., 2014). Many of these methods accommodate an unknown number of change points and estimate both their number and locations from the observed data (cf. Anastasiou and Fryzlewicz, 2022; Chen et al., 2023; Cho and Fryzlewicz, 2020). However, these methods are retrospective in nature because they are designed for fixed-sample settings in which all observations are available before the analysis is performed. Consequently, they are generally unsuitable for the present sequential decision-making problem, where observations are collected continuously over time and the process status must be determined whenever a new observation becomes available.

Online monitoring of SPRS is particularly challenging because both the starting and ending points of each OC period must be detected, as discussed above. However, these two types of change points have fundamentally different characteristics. For example, when detecting a starting point, the shift magnitude is typically unknown. In contrast, the shift magnitude at an ending point can be reasonably estimated after the corresponding starting point has been signaled and its change point has been estimated. Furthermore, an ending point can only be detected after the associated starting point has been identified. Although the OC process mean and the corresponding shift magnitude can be estimated following the signal of the starting point, the estimation uncertainty introduces additional challenges to accurately detecting the ending point.

In this paper, we introduce new concepts and propose a general framework for the online monitoring of SPRS. For clarity and simplicity, we focus on the setting in which the process observations are univariate, serially independent, and normally distributed. The proposed framework can be extended to more general scenarios, and such extensions will be investigated in future research.

The remainder of the paper is organized as follows. Section 2 presents the proposed methodology in detail. Section 3 evaluates its numerical performance through simulation studies. Section 4 provides a real-data example to illustrate the practical utility of the proposed method. Finally, Section 5 concludes the paper with a discussion.

2 Proposed Method

2.1 Problem formulation

Let X be the quality variable of interest, whose IC distribution is assumed to be $N(\mu_0, \sigma^2)$, where both μ_0 and σ^2 are unknown. Suppose an initial IC dataset $\{X_{-m+1}, X_{-m+2}, \ldots, X_0\}$ is available, consisting of i.i.d. observations from the IC distribution. The process observations for online monitoring are denoted by $\{X_i, i \geq 1\}$, and they follow the multiple-change-point model

$$X_i = \begin{cases} \mu_0 + \varepsilon_i, & \text{when } 1 \leq i \leq \tau_{1,1}, \\ \mu_0 + \delta_1 + \varepsilon_i, & \text{when } \tau_{1,1} + 1 \leq i \leq \tau_{1,2}, \\ \mu_0 + \varepsilon_i, & \text{when } \tau_{1,2} + 1 \leq i \leq \tau_{2,1}, \\ \mu_0 + \delta_2 + \varepsilon_i, & \text{when } \tau_{2,1} + 1 \leq i \leq \tau_{2,2}, \\ \vdots \end{cases} \tag{1}$$

where $\tau_{1,1} < \tau_{1,2} < \tau_{2,1} < \tau_{2,2} < \ldots$ denote the unknown change points, $\{\delta_1, \delta_2, \ldots\}$ are unknown nonzero mean-shift sizes, and $\{\varepsilon_i, i \geq 1\}$ are i.i.d. $N(0, 1)$ random errors.

Figure 1 illustrates the mean function of Model (1), which contains change points at $\{\tau_{1,1}, \tau_{1,2}, \tau_{2,1}, \tau_{2,2}, \ldots\}$. The process observations $\{X_i, i \geq 1\}$ described by Model (1) are IC for $1 \leq i \leq \tau_{1,1}$, become OC with a mean shift of δ_1 for $\tau_{1,1}+1 \leq i \leq \tau_{1,2}$, return to the IC state for $\tau_{1,2}+1 \leq i \leq \tau_{2,1}$, and then become OC again with a mean shift of δ_2 for $\tau_{2,1} + 1 \leq i \leq \tau_{2,2}$, and so on. The primary objective of online process monitoring is to detect these change points as quickly as possible under suitable regularity conditions discussed below.

2.2 Proposed online monitoring framework

In environmental monitoring, disease surveillance, and many other applications, the primary concern is the detection of upward mean shifts in the intensity levels of major pollutants, the incidence

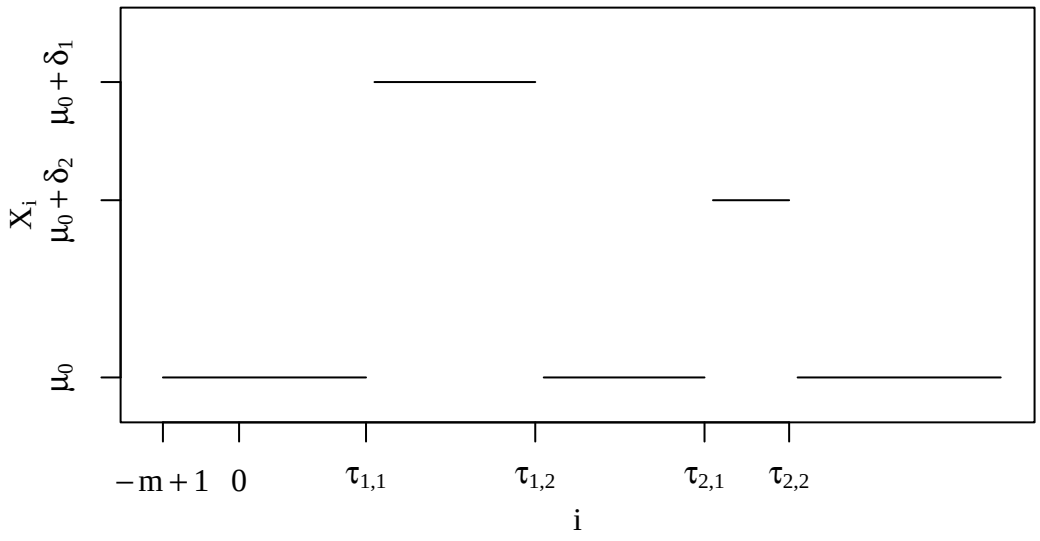


Figure 1: Demonstration of the mean function of the multiple-change-point model (1) with change points at $\{\tau_{1,1}, \tau_{1,2}, \tau_{2,1}, \tau_{2,2}, \ldots\}$. The observations $\{X_i, -m+1 \leq i \leq 0\}$ constitute the initial IC data. The process observations $\{X_i, i \geq 1\}$ are IC for $1 \leq i \leq \tau_{1,1}$, become OC with a mean shift of δ_1 for $\tau_{1,1}+1 \leq i \leq \tau_{1,2}$, return to the IC state for $\tau_{1,2}+1 \leq i \leq \tau_{2,1}$, and then become OC again with a mean shift of δ_2 for $\tau_{2,1}+1 \leq i \leq \tau_{2,2}$, and so forth.

rates of infectious diseases, or other key process performance variables. For these applications, and for simplicity of exposition, the proposed online monitoring framework is presented in the context of detecting recurrent *upward* mean changes (cf., Figure 1). Detection of downward changes or changes in either direction can be handled in a similar manner. In addition, the proposed online monitoring framework is developed using the CUSUM charting technique. Other approaches, such as Shewhart charts, EWMA charts, or CPD charts, can also be incorporated to construct a similar online monitoring framework.

Unlike conventional online process monitoring problems, which focus solely on detecting the first distributional shift, the current research problem involves multiple mean shifts. Each shift has a starting time (e.g., $\tau_{1,1}$) and an ending time (e.g., $\tau_{1,2}$). Timely detection of the starting time is essential for implementing appropriate interventions in applications such as environmental monitoring. Equally important is the timely detection of the ending time, which helps prevent un-

necessary expenditure of human and financial resources. In what follows, we describe our proposed method for detecting the starting and ending times of the first mean shift.

Detection of the starting time $\tau_{1,1}$ of the first mean shift. To sequentially monitor the observed data $\{X_i, i \geq 1\}$ and detect the first upward mean shift starting at $\tau_{1,1}$, a natural approach is to use the following conventional one-sided CUSUM charting statistic (cf., Qiu 2014, Chapter 4):

$$\widetilde{C}_i^+ = \max\left[0, \widetilde{C}_{i-1}^+ + (X_i - \mu_0)/\sigma - \widetilde{k}_1\right], \text{ for } i \geq 1,$$

where $\widetilde{C}_0^+ = 0$, and $\widetilde{k}_1 > 0$ is a reference value (or allowance constant) pre-specified to be half of the tartget mean shift. Because μ_0 and σ^2 are typically unknown in practice, they are usually estimated from the initial IC data by the sample mean $\widehat{\mu}_0^{(0)}$ and sample variance $(\widehat{\sigma}^2)^{(0)}$. As discussed in Hawkins (1987), the variability in these estimates can deteriorate the IC performance of the resulting control chart, particularly when the initial IC sample size m is small. To mitigate this issue, Hawkins (1987) proposed a self-starting CUSUM chart in which the newly observed data at the current time point are combined with the IC data from previous time points, and the IC parameter estimates are updated accordingly whenever the underlying process is deemed to be IC.

As discussed in the literature (cf., Qiu 2014, Section 4.5.1), the self-starting CUSUM chart has an important limitation: a mean shift may be permanently missed if it is not detected quickly, because OC observations collected after the shift can contaminate the estimates of the IC parameters. To address this issue, we propose the following modified one-sided self-starting CUSUM chart for detecting the first upward mean shift at $\tau_{1,1}$:

$$C_i^+ = \max\left(0, C_{i-1}^+ + Z_i - k_1\right), \text{ for } i \geq 1, \tag{2}$$

where $C_0^+ = 0$, $k_1 > 0$ is a pre-specified reference value, and Z_i is defined below. The chart signals a shift at time i if

$$C_i^+ > h_{1,1}, \tag{3}$$

where $h_{1,1} > 0$ is a control limit.

It is well known that the CUSUM chart (2)-(3) has a *restarting property*: its charting statistic C_i^+ is reset to 0 whenever the cumulative sum $C_{i-1}^+ + Z_i$ falls below the pre-specified constant k_1. This indicates that the data up to time i contain little evidence of a mean shift. Chatterjee and Qiu (2009) introduced the concept of *Spring Length* to characterize this property, defined at time

i as

$$T_i = \begin{cases} 0, & \text{if } C_i^+ = 0, \\ b, & \text{if } C_i^+ \neq 0, \ldots, C_{i-b+1}^+ \neq 0, C_{i-b}^+ = 0. \end{cases}$$

From this definition and the fact that $C_0^+ = 0$, we have $T_i \leq i$. Using this concept, it is unlikely that a mean shift occurred at or before time $i - T_i$. Therefore, at the current time i, we update the estimates of the IC parameters using observations only up to time $i - T_i$, thereby reducing potential contamination from OC observations.

Based on this idea, Z_i in (2) is defined as

$$Z_i = \Phi^{-1}\left[\mathcal{T}_{i-T_{i-1}-2+m}\left(\sqrt{\frac{i-T_{i-1}-2+m}{i-T_{i-1}-1+m}} \times \frac{X_i - \widehat{\mu}_0^{(i-T_{i-1}-1)}}{(\widehat{\sigma})^{(i-T_{i-1}-1)}}\right)\right], \tag{4}$$

where $\Phi^{-1}(\cdot)$ is the inverse cumulative distribution function (cdf) of the $N(0,1)$ distribution, $\mathcal{T}_{i-T_{i-1}-2+m}(\cdot)$ is the cdf of the $t_{i-T_{i-1}-2+m}$ distribution, and $\widehat{\mu}_0^{(i-T_{i-1}-1)}$ and $(\widehat{\sigma})^{(i-T_{i-1}-1)}$ are the estimates of the IC mean and standard deviation computed using data up to time $i - T_{i-1} - 1$. In Equation (4), the spring length T_{i-1} is used instead of T_i because T_i cannot be determined until C_i^+ has been computed. This substitution has little effect on the statistic Z_i, since it is straightforward to check that $T_i \leq t_{i-1} + 1$. It can be verified that $(X_i - \widehat{\mu}_0^{(i-T_{i-1}-1)})/(\widehat{\sigma})^{(i-T_{i-1}-1)}$ follows a $t_{i-T_{i-1}-2+m}$ distribution. The transformation on the right-hand side of (4) then transforms this quantity to Z_i which has a $N(0,1)$ distribution. In practice, when $i - T_{i-1} - 2 + m$ is sufficiently large (e.g., ≥ 50), the $t_{i-T_{i-1}-2+m}$ distribution is well approximated by the $N(0,1)$ distribution. In such cases, the transformation in (4) can be omitted, and we can directly use $(X_i - \widehat{\mu}_0^{(i-T_{i-1}-1)})/(\widehat{\sigma})^{(i-T_{i-1}-1)}$ in (2).

To compute the updated estimates of the IC process mean μ_0 and variance σ^2 at different time points, the following recursive formulas are useful:

$$\begin{aligned} \widehat{\mu}_0^{(i)} &= \frac{m+i-1}{m+i}\widehat{\mu}_0^{(i-1)} + \frac{1}{m+i}X_i, \\ \left(\widehat{\sigma}^2\right)^{(i)} &= \frac{m+i-1}{m+i}\left(\widehat{\sigma}^2\right)^{(i-1)} + \frac{m+i-1}{(m+i)^2}\left(\widehat{\mu}_0^{(i-1)} - X_i\right)^2, \text{for } i \geq 1, \end{aligned} \tag{5}$$

where $\widehat{\mu}_0^{(0)}$ and $(\widehat{\sigma}^2)^{(0)}$ are the sample mean and sample variance of the initial IC data. From the construction of the CUSUM chart (2)–(3), the estimates of μ_0 and σ^2 do not need to be updated at every observation time. Instead, they only need to be updated at times when the charting statistic is restarted. Nevertheless, the recursive formulas in (5) remain useful because the current estimates

$\widehat{\mu}_0^{(i)}$ and $\left(\widehat{\sigma}^2\right)^{(i)}$ can be obtained efficiently from the most recently updated estimates at a previous observation time by applying the recursive updates a few times.

In the CUSUM chart (2)-(3), the allowance constant k_1 is typically chosen to be half of the target mean shift, and the control limit $h_{1,1}$ is computed using existing algorithms (cf. Qiu 2014, Chapter 4) so that a pre-specified IC average run length, denoted $ARL_{IC}^{(1)}$, is achieved. In this notation, the subscript "IC" indicates that the ARL is evaluated under the IC process, and the superscript "(1)" indicates that the chart is designed for detecting the first upward mean shift. The chart's detection performance can be measured by the OC average run length, defined as the expected number of observation times from the start of the shift to the signal time of the chart. In this paper, we refer to this quantity as the OC average delay (AD) of the signal, denoted $AD_{OC}^{(1)}$, to avoid confusion with the detection of the ending time $\tau_{1,2}$ discussed below.

Detection of the ending time $\tau_{1,2}$ of the first mean shift. Once the CUSUM chart (2)-(3) signals at time $\xi_{1,1}$ the occurence of the first shift, we must promptly detect the ending time $\tau_{1,2}$, which corresponds to a downward shift of size $-\delta_1$, and verify that the underlying process has returned to the IC state (cf., Figure 2). To accomplish this, we first estimate the starting time $\tau_{1,1}$ and the magnitude δ_1 of the detected shift. For this purpose, the following change-point detection procedure (cf. Qiu 2014, Chapter 6) can be applied to the process observations $\{X_r, -m+1 \le r \le \xi_{1,1}\}$:

$$\begin{aligned} \widehat{\tau}_{1,1} &= \arg \min_{0 \le r \le \xi_{1,1}-1} \frac{(m+r)(\xi_{1,1}-r)}{m+\xi_{1,1}} \left[\overline{X}_{(-m+1):r} - \overline{X}_{(r+1):\xi_{1,1}}\right]^2, \\ \widehat{\delta}_1 &= \overline{X}_{(\widehat{\tau}_{1,1}+1):\xi_{1,1}} - \overline{X}_{(-m+1):\widehat{\tau}_{1,1}}, \end{aligned} \tag{6}$$

where $\overline{X}_{r_1:r_2}$ denotes the sample mean of $\{X_r, r_1 \le r \le r_2\}$ for any $-m+1 \le r_1 \le r_2 \le \xi_{1,1}$.

A natural approach to detect the downward shift at $\tau_{1,2}$ after the signal time $\xi_{1,1}$ of the CUSUM chart (2)-(3) is to sequentially monitor the process observations $\{X_i, i \ge \xi_{1,1}+1\}$ using a one-sided CUSUM chart designed for detecting downward shifts (cf. Qiu 2014, Chapter 4):

$$\widetilde{C}_i^- = \min\left(0, \widetilde{C}_{i-1}^- + \widetilde{Z}_i + \widehat{\delta}_1/2\right), \quad \text{for } i \ge \xi_{1,1}+1, \tag{7}$$

where $\widetilde{C}_{\xi_{1,1}}^- = 0$, and $\widetilde{Z}_i$ is a transformed observation at time i, defined below. The chart signals a downward shift at time i if

$$\widetilde{C}_i^- < -\widetilde{h}_{1,2}, \tag{8}$$

where $-\widetilde{h}_{1,2} < 0$ is the control limit.

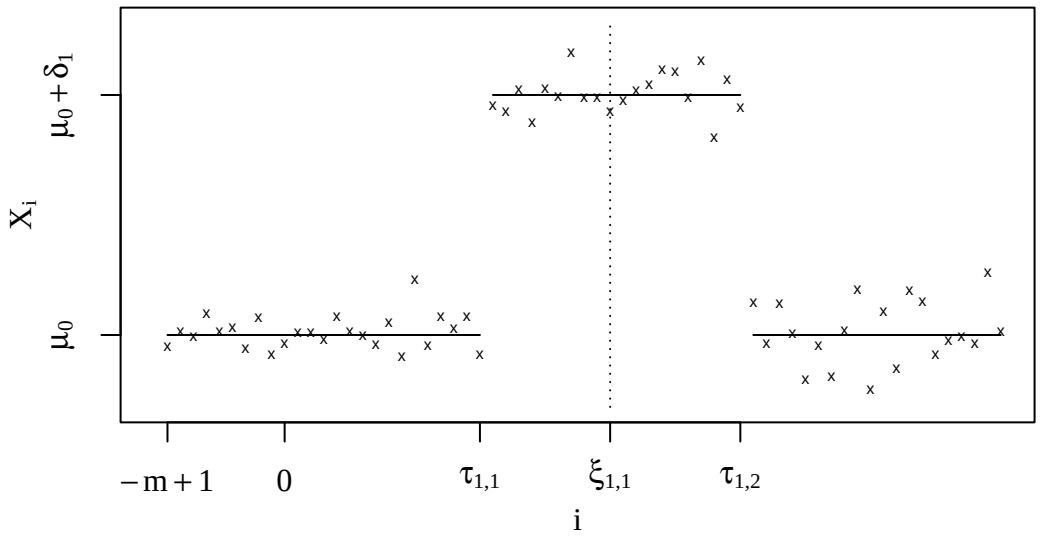


Figure 2: Demonstration of the first mean shift (solid lines), the process observations (small crosses), and the signal time $\xi_{1,1}$ of the CUSUM chart (2)-(3) (vertical dotted line).

For the CUSUM chart (7)-(8), the quantities $\widehat{\mu}_0^{(\widehat{\tau}_{1,1})}+\widehat{\delta}_1$ and $\left(\widehat{\sigma}^2\right)^{(\widehat{\tau}_{1,1})}$ can be regarded as the initial estimates of $\mu_1=\mu_0+\delta_1$ and σ^2, denoted by $\widetilde{\mu}_1^{(0)}$ and $(\widetilde{\sigma}^2)^{(0)}$, respectively. The transformed observation $\widetilde{\widetilde{Z}}_i$ is then defined as: for $i \geq \xi_{1,1}+1$,

$$\widetilde{\widetilde{Z}}_i = \Phi^{-1}\left[\mathcal{T}_{i-\xi_{1,1}-\widetilde{T}_{i-1}+m+\widehat{\tau}_{1,1}-2}\left(\sqrt{\frac{i-\xi_{1,1}-\widetilde{T}_{i-1}+m+\widehat{\tau}_{1,1}-2}{i-\xi_{1,1}-\widetilde{T}_{i-1}+m+\widehat{\tau}_{1,1}-1}}\times\frac{X_i-\widetilde{\mu}_1^{(i-\xi_{1,1}-\widetilde{T}_{i-1}-1)}}{(\widetilde{\sigma})^{(i-\xi_{1,1}-\widetilde{T}_{i-1}-1)}}\right)\right],$$

where $\widetilde{T}_{i-1}$ is the spring length of the chart (7)-(8) at time $i-1$, defined analogously to T_{i-1}, and $\widetilde{\mu}_1^{(i-\xi_{1,1}-T_{i-1}-1)}$ and $(\widetilde{\sigma}^2)^{(i-\xi_{1,1}-T_{i-1}-1)}$ are computed in the same recursive manner as in (5), after the estimates of μ_0 are replaced by those of μ_1 and the initial values are replaced by $\widetilde{\mu}_1^{(0)}=\widehat{\mu}_0^{(\widehat{\tau}_{1,1})}+\widehat{\delta}_1$ and $(\widetilde{\sigma}^2)^{(0)}=\left(\widehat{\sigma}^2\right)^{(\widehat{\tau}_{1,1})}$.

After applying the change-point detection procedure in (6), the process becomes OC starting at approximately $\widehat{\tau}_{1,1}$. Although process monitoring typically cannot go back in time after the chart signals at $\xi_{1,1}$, the charting statistic $\widetilde{C}_i^-$ defined in (7) can be modified by incorporating both the change-point detection results and the observed data over the time interval $[\widehat{\tau}_{1,1}+1, \xi_{1,1}]$. Motivated by this idea, we propose the following CUSUM charting statistic for detecting the downward shift

at $\tau_{1,2}$:

$$C_i^- = \min\left(0, C_{i-1}^- + \widehat{Z}_i + \widehat{\delta}_1/2\right), \text{ for } i \geq \xi_{1,1} + 1, \tag{9}$$

and the chart signals a shift when

$$C_i^- < -h_{1,2}, \tag{10}$$

where $-h_{1,2} < 0$ is the control limit. Here, $C_{\xi_{1,1}}^-$ is computed from the recursive formula in (9) as i varies over $[\widehat{\tau}_{1,1} + 1, \xi_{1,1}]$, with the initial value $C_{\widehat{\tau}_{1,1}}^- = 0$. The quantity $\widehat{Z}_i$ is defined as

$$\widehat{Z}_i = \Phi^{-1}\left[\mathcal{T}_{i-\widehat{T}_{i-1}-2+m}\left(\sqrt{\frac{i-\widehat{T}_{i-1}-2+m}{i-\widehat{T}_{i-1}-1+m}} \times \frac{X_i - \widehat{\mu}_1^{(i-\widehat{\tau}_{1,1}-T_{i-1}-1)}}{(\widehat{\sigma})^{(i-T_{i-1}-1)}}\right)\right], \text{ for } i \geq \widehat{\tau}_{1,1} + 1.$$

In computing $\widehat{Z}_i$, the variance estimator $(\widehat{\sigma}^2)^{(i-T_{i-1}-1)}$ is obtained using the recursive formula in (5), except that estimates of μ_0 are replaced by those of μ_1 for $i \geq \widehat{\tau}_{1,1}+1$. Likewise, $\widehat{\mu}_1^{(i-\widehat{\tau}_{1,1}-T_{i-1}-1)}$ is computed recursively in the same manner as in (5), with initial value $\widehat{\mu}_1^{(0)} = \widehat{\mu}_0^{(\widehat{\tau}_{1,1})} + \widehat{\delta}_1$. Finally, $\widehat{T}_{i-1}$ denotes the spring length of the chart (9)-(10) at time $i-1$.

For the CUSUM chart (9)-(10) (or (7)-(8)), its control limit $-h_{1,2}$ (or $-\widetilde{h}_{1,2}$) can be computed using existing algorithms in the same manner as for the CUSUM chart (2)–(3), so that a prespecified ARL under the OC process with mean $\mu_0 + \delta_1$ is attained. We denote this ARL as $ARL_{OC}^{(1)}$, where the subscript "OC" indicates evaluation under the OC process and the superscript "(1)" indicates that the chart is designed for detecting the ending time of the first mean shift. The detection performance of the chart can be quantified by the average delay until it signals that the process has returned to the IC state, denoted $AD_{IC}^{(1)}$.

Remark 1. The CUSUM chart (9)-(10) (or (7)-(8)) is applied only after the chart (2)–(3) signals and the change-point detection procedure (6) has been carried out. Intuitively, the event $\xi_{1,1} \in [\tau_{1,2} + 1, \tau_{2,1}]$ is much less likely than the event $\xi_{1,1} \in [\tau_{1,1} + 1, \tau_{1,2}]$, because the process observations in $[\tau_{1,2}+1, \tau_{2,1}]$ are IC, making it unlikely for the chart to signal in that interval. This expectation is confirmed by the numerical results presented in Figure 6 in Section 3.

Detection of all subsequent shifts. Once the starting and ending times of the first shift have been detected, subsequent shifts can be identified in a similar manner. For example, after the CUSUM chart (9)-(10) signals at $\xi_{1,2}$, the change point $\tau_{1,2}$ can be estimated using a procedure analogous to (6). Then, the start time $\tau_{2,1}$ of the second shift can be detected using a CUSUM chart similar to (2)-(3), after appropriately updating the initial estimates of μ_0 and σ^2 using the observations collected between the estimated change point $\widehat{\tau}_{1,2}$ and the signal time $\xi_{1,2}$. The

allowance constant of this chart, denoted by k_2, is chosen to be half of the magnitude of the second target shift. Its control limit, denoted $h_{2,1}$, is selected to achieve the pre-specified IC average run length, $ARL_{IC}^{(2)}$, and its detection performance is measured by the average delay to signal, denoted $AD_{OC}^{(2)}$.

As discussed in Subsection 2.1 and the first paragraph in Subsection 2.2, the proposed method is intended for settings in which the monitored process is in IC most of the time and only occasionally becomes OC due to upward mean shifts. Given this structure of the monitoring problem, the sequence of control charts should be implemented according to the following rules: (i) monitoring begins with a chart designed to detect an upward mean shift (cf. (2)–(3)); and (ii) upward and downward charts (cf. (8)–(9)) are applied alternately following each signal.

The proposed online monitoring framework described above is summarized in the flowchart shown in Figure 3.

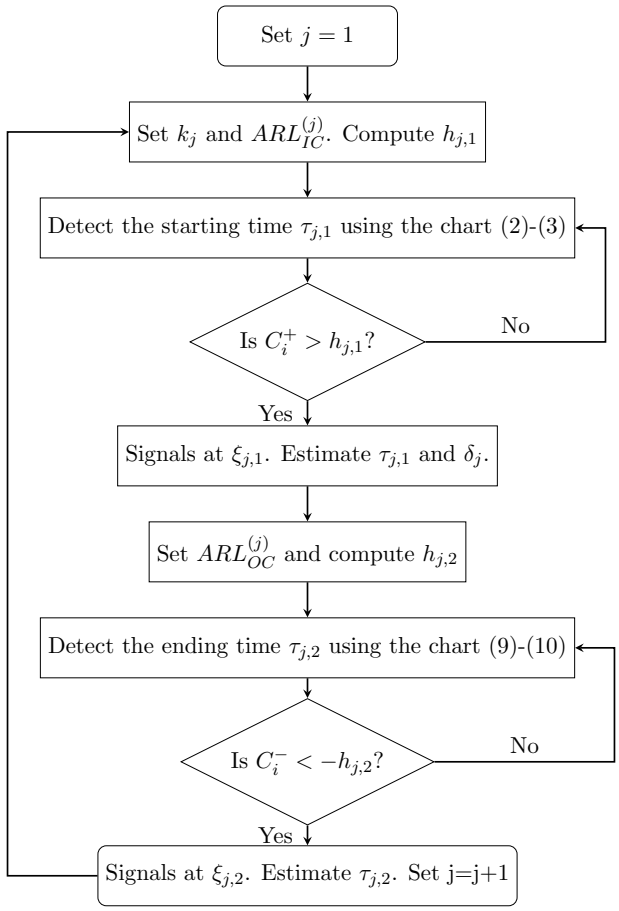


Figure 3: Flowchart of the proposed online monitoring framework.

Remark 2. The proposed method described above assumes that all shift sizes $\delta_1, \delta_2, \ldots$ in Model (1) are positive. When all shift sizes are negative, the method can still be applied by replacing the one-sided CUSUM chart (2)–(3) with the corresponding one-sided CUSUM chart for detecting downward mean shifts (cf. Qiu, 2013, Chapter 4), and replacing the chart (9)–(10) with the corresponding one-sided CUSUM chart for detecting upward mean shifts. When the shift sizes may be arbitrary in direction, the proposed method can be modified as follows. Specifically, the one-sided CUSUM chart (2)–(3) is replaced by the corresponding two-sided CUSUM chart for detecting mean shifts in either direction, while the chart (9)–(10) is replaced by a one-sided CUSUM chart designed to detect a mean shift in the direction opposite to the estimated shift size at the starting time of the related mean shift(cf. (6)).

3 Simulation Studies

In this section, we evaluate the numerical performance of the proposed approach using Monte Carlo simulations. Unless otherwise specified, all simulation studies are conducted under the following setup. Each monitored process consists of 3,000 observations, with the first 50 observations serving as the initial IC data. Process monitoring begins at the 51st observation. Two shifts are introduced over the intervals $[\tau_{1,1}+1, \tau_{1,2}] = [161, 220]$ and $[\tau_{2,1}+1, \tau_{2,2}] = [331, 390]$, respectively, with corresponding shift magnitudes δ_1 and δ_2. IC observations are generated independently from a $N(0, \sigma^2)$ distribution, whereas OC observations are generated independently from $N(\delta_1, \sigma^2)$ and $N(\delta_2, \sigma^2)$ over the two respective intervals.

3.1 Numerical performance of the proposed method

For the charts used to detect the starting and ending times of mean shifts, the nominal levels of $ARL_{IC}^{(j)}$ and $ARL_{OC}^{(j)}$ are both set at 50 for each j. Performance metrics, including $\{(AD_{OC}^{(j)}, AD_{IC}^{(j)}), j \geq 1\}$, are evaluated based on 5,000 replicated simulation runs. As an illustration, consider chart (2)-(3) for detecting the onset of the first mean shift. A total of 5,000 independent process-monitoring simulations are conducted. In each run, the delay time $\xi_{1,1} - \tau_{1,1}$ of the signal is recorded, and runs with $\xi_{1,1} \leq \tau_{1,1}$ are excluded. The average of the retained delay times provides an estimate of $AD_{OC}^{(1)}$ and the corresponding standard error can also be computed. The remaining performance

measures in $\{(AD_{OC}^{(j)}, AD_{IC}^{(j)}), j \geq 1\}$ are obtained in a similar manner.

For the cases where $\sigma = 0.2$ or 0.5, $\delta_1 = \delta_2 = \delta \in \{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, and $k_1 = k_2 = k \in \{0.25, 0.5, 0.75\}$, the computed values of $AD_{OC}^{(1)}$ and $AD_{OC}^{(2)}$ for the control charts designed to detect the onset of the first two shifts (cf., (2)-(3)) are reported in Table 1. The table also includes results for $\delta = 0$; in this case, $AD_{OC}^{(1)}$ and $AD_{OC}^{(2)}$ reduce to the actual $ARL_{IC}^{(1)}$ and $ARL_{IC}^{(2)}$, respectively. Table 1 shows that: (i) all actual $ARL_{IC}^{(1)}$ and $ARL_{IC}^{(2)}$ values are within 10% of their nominal levels; (ii) both $AD_{OC}^{(1)}$ and $AD_{OC}^{(2)}$ decrease as δ increases or σ decreases; and (iii) smaller values of k yield smaller $AD_{OC}^{(j)}$ for $j = 1, 2$ when detecting small shifts (e.g., $\delta = 0.2$), whereas larger values of k perform better for larger shifts (e.g., $\delta = 1.0$). Conclusion (i) confirms that the proposed charts maintain reliable in-control performance. Conclusion (ii) is intuitively expected. Conclusion (iii) is consistent with the well-known result in the SPC literature that smaller k values are more effective for detecting small shifts, while larger k values are preferable for detecting large shifts (cf., Qiu 2014, Chapter 4).

For the scenarios considered in Table 1, the computed values of $AD_{IC}^{(1)}$ and $AD_{IC}^{(2)}$ for the two charts designed to detect the ending of the shifts (cf. (9)–(10)) are reported in Table 2. Note that allowance constants are not used in these charts, because the shift magnitudes are estimated in advance and incorporated into their construction. As before, when $\delta = 0$, the reported $AD_{IC}^{(1)}$ and $AD_{IC}^{(2)}$ values reduce to the actual $ARL_{OC}^{(1)}$ and $ARL_{OC}^{(2)}$, respectively. Table 2 indicates that these two charts also maintain reliable in-control performance under the considered settings. Moreover, both $AD_{IC}^{(1)}$ and $AD_{IC}^{(2)}$ decrease as δ increases or σ decreases.

The density histograms of the estimated change point, $\widehat{\tau}_{1,1}$, $\widehat{\tau}_{1,2}$, $\widehat{\tau}_{2,1}$, and $\widehat{\tau}_{2,2}$, based on 5,000 replicated process-monitoring simulations under the settings $k = 0.5$, $\delta = 0.4$ or 1.0, and $\sigma = 0.2$ or 0.5, are presented in Figure 4. The results show that the change-point estimates are more accurate when δ is larger and/or σ is smaller, which is intuitively expected.

For the estimated shift sizes $\widehat{\delta}_1$ and $\widehat{\delta}_2$, we assess their performance using the relative estimation errors $(\widehat{\delta}_1 - \delta_1)/\delta_1$ and $(\widehat{\delta}_2 - \delta_2)/\delta_2$. Under the settings $k = 0.5$, $\delta = 0.4$ or 1.0, and $\sigma = 0.2$ or 0.5, the relative estimation errors based on 5,000 replicated process-monitoring simulations are displayed in Figure 5. The results indicate that the shift-size estimates are quite accurate when the noise level of the observed data is relatively low (i.e., $\sigma = 0.2$), particularly when the true shift sizes are large (i.e., $\delta_1 = \delta_2 = 1.0$). When the noise level is higher (i.e., $\sigma = 0.5$), the estimates

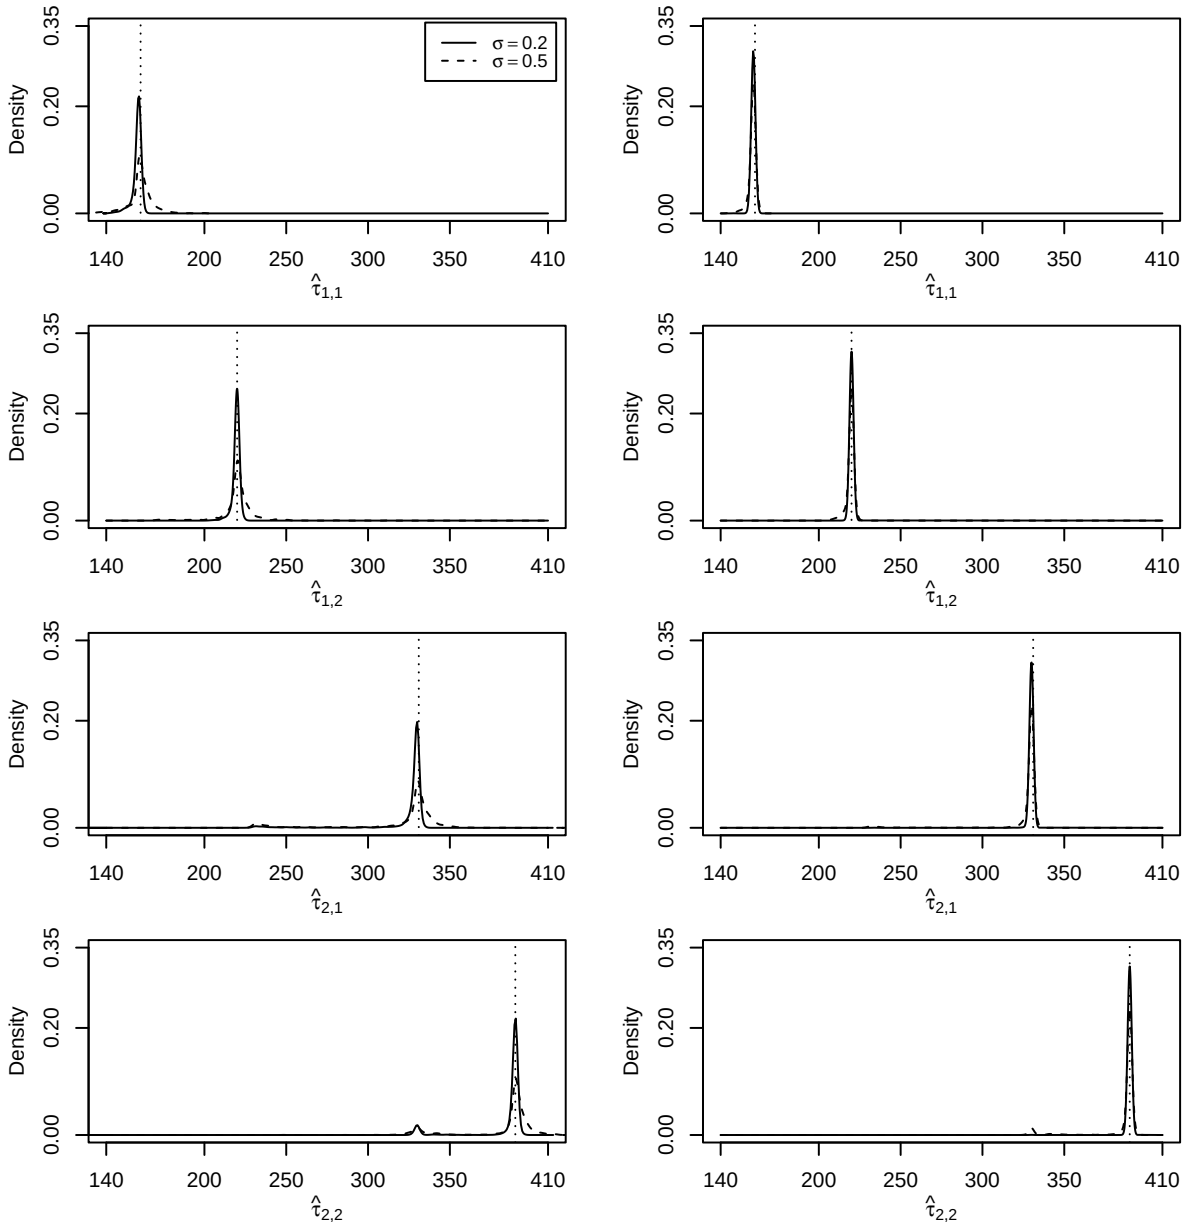


Figure 4: Density histograms of the estimated change points $\widehat{\tau}_{1,1}$ (1st row), $\widehat{\tau}_{1,2}$ (2nd row), $\widehat{\tau}_{2,1}$ (3rd row), and $\widehat{\tau}_{2,2}$ (4th row), based on 5,000 replicated process-monitoring simulations with $k = 0.5$, $\delta = 0.4$ (left column) or 1.0 (right column), and $\sigma = 0.2$ or 0.5. The true change points are $\tau_{1,1} = 160$, $\tau_{1,2} = 220$, $\tau_{2,1} = 330$, and $\tau_{2,2} = 390$, shown by vertical dotted lines in the plots.

Table 1: Computed values of $AD_{OC}^{(1)}$ and $AD_{OC}^{(2)}$, along with their standard errors (in parenthesis), are presented for scenarios in which the nominal levels of $ARL_{IC}^{(1)}$, $ARL_{OC}^{(1)}$, $ARL_{IC}^{(2)}$, and $ARL_{OC}^{(2)}$ are all set to 50, $\sigma = 0.2$ or 0.5, two shifts are introduced with equal magnitudes $\delta_1 = \delta_2 = \delta \in \{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, and the allowance constants for detecting the onset of the shifts are specified as $k_1 = k_2 = k \in \{0.25, 0.5, 0.75\}$.

	σ	k	$\delta = 0$	0.2	0.4	0.6	0.8	1.0
$AD_{OC}^{(1)}$	0.2	0.25	52.3 (0.79)	5.0 (0.05)	1.5 (0.02)	0.7 (0.01)	0.3 (0.01)	0.1 (0.00)
		0.5	52.4 (0.76)	5.8 (0.07)	1.2 (0.02)	0.4 (0.01)	0.1 (0.00)	0.0 (0.00)
		0.75	50.0 (0.72)	7.5 (0.09)	1.2 (0.02)	0.3 (0.01)	0.1 (0.00)	0.0 (0.00)
	0.5	0.25	53.5 (0.80)	24.8 (0.41)	6.3 (0.08)	3.1 (0.04)	1.9 (0.02)	1.3 (0.02)
		0.5	53.2 (0.75)	29.1 (0.46)	7.1 (0.10)	3.0 (0.04)	1.6 (0.02)	1.0 (0.01)
		0.75	50.9 (0.71)	32.7 (0.50)	8.6 (0.13)	3.4 (0.05)	1.7 (0.03)	1.0 (0.02)
$AD_{OC}^{(2)}$	0.2	0.25	51.3 (0.30)	5.2 (0.06)	1.6 (0.02)	0.8 (0.01)	0.5 (0.01)	0.2 (0.01)
		0.5	50.4 (0.32)	5.6 (0.10)	1.6 (0.02)	0.6 (0.01)	0.2 (0.01)	0.0 (0.00)
		0.75	51.2 (0.32)	10.0 (0.09)	1.7 (0.02)	0.5 (0.01)	0.1 (0.00)	0.0 (0.00)
	0.5	0.25	52.7 (0.32)	22.5 (0.41)	8.1 (0.08)	3.5 (0.04)	2.2 (0.03)	1.5 (0.02)
		0.5	52.9 (0.33)	29.7 (0.48)	9.7 (0.18)	3.8 (0.05)	2.1 (0.03)	1.3 (0.02)
		0.75	53.1 (0.36)	34.4 (0.47)	12.9 (0.13)	4.6 (0.08)	2.3 (0.04)	1.2 (0.03)

Table 2: Computed values of $AD_{IC}^{(1)}$ and $AD_{IC}^{(2)}$, along with their standard errors (in parenthesis), are presented for scenarios considered in Table 1.

	σ	$\delta = 0$	0.2	0.4	0.6	0.8	1.0
$AD_{IC}^{(1)}$	0.2	51.9 (0.31)	1.5 (0.02)	0.6 (0.01)	0.2 (0.01)	0.0 (0.00)	0.0 (0.00)
	0.5	52.3 (0.30)	8.4 (0.17)	3.5 (0.05)	2.0 (0.03)	1.3 (0.02)	0.0 (0.00)
$AD_{IC}^{(2)}$	0.2	51.4 (0.31)	2.0 (0.04)	0.8 (0.02)	0.3 (0.01)	0.0 (0.01)	0.0 (0.00)
	0.5	52.5 (0.43)	8.2 (0.25)	3.6 (0.06)	2.2 (0.04)	1.4 (0.03)	0.0 (0.00)

exhibit greater variability and can deviate substantially from the true values, especially when the true shift sizes are small (i.e., $\delta_1 = \delta_2 = 0.4$), which is intuitively reasonable.

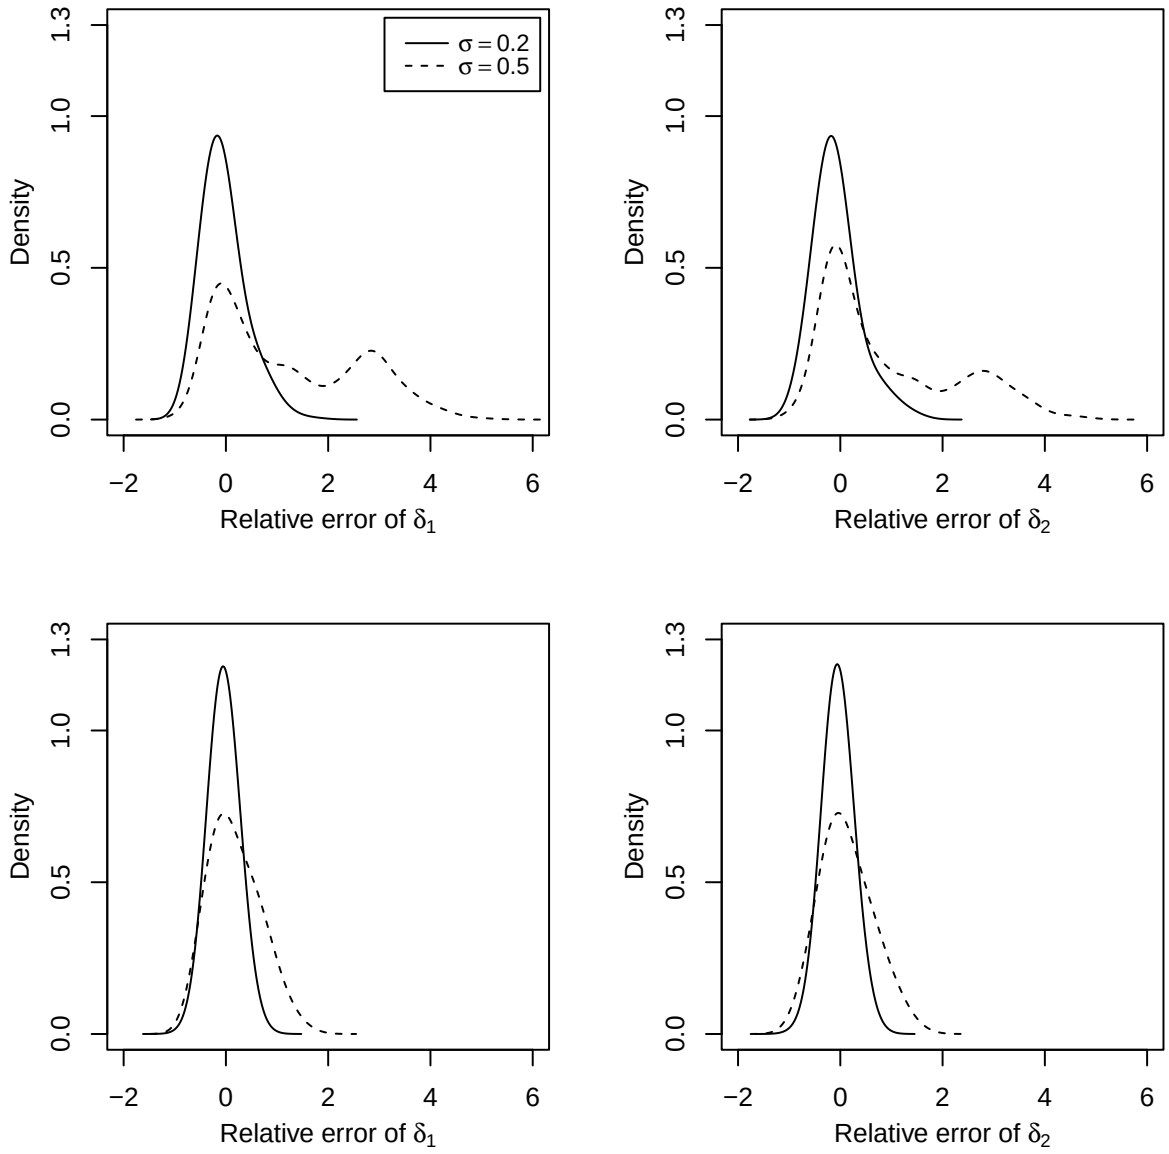


Figure 5: Relative estimation errors of the estimated shift sizes $\widehat{\delta}_1$ and $\widehat{\delta}_2$, based on 5,000 replicated process-monitoring simulations under the settings $k = 0.5$, $\delta = 0.4$ (1st row) or 1.0 (2nd row), and $\sigma = 0.2$ or 0.5.

In Remark 1 of Subsection 2.2, it is noted that the event in which the signal time $\xi_{1,1}$ of the first chart (2)-(3) falls within the interval $[\tau_{1,2} + 1, \tau_{2,1}]$ is much less likely than the event that $\xi_{1,1}$ falls within $[\tau_{1,1} + 1, \tau_{1,2}]$. To verify this claim, we consider settings where $k = 0.25$, 0.5, or 0.75, σ ranges from 0.1 to 1.5, and $\delta = 0.4$ or 1.0. Based on 5,000 replicated process-monitoring

simulations, we compute the ratio of the estimated probabilities

$$R(\sigma) = \frac{P(\xi_{1,1} \in [\tau_{1,2}+1, \tau_{2,1}])}{P(\xi_{1,1} \in [\tau_{1,1}+1, \tau_{1,2}])},$$

and present the results in Figure 6. When $\delta = 0.4$, the left panel of Figure 6 shows that $R(\sigma) = 0$ for $\sigma \leq 0.5$, and remains very small when $\sigma > 0.5$. When $\delta = 1.0$, the right panel indicates that $R(\sigma) = 0$ for all cases considered. These results confirm that the former event is indeed much less likely than the latter.

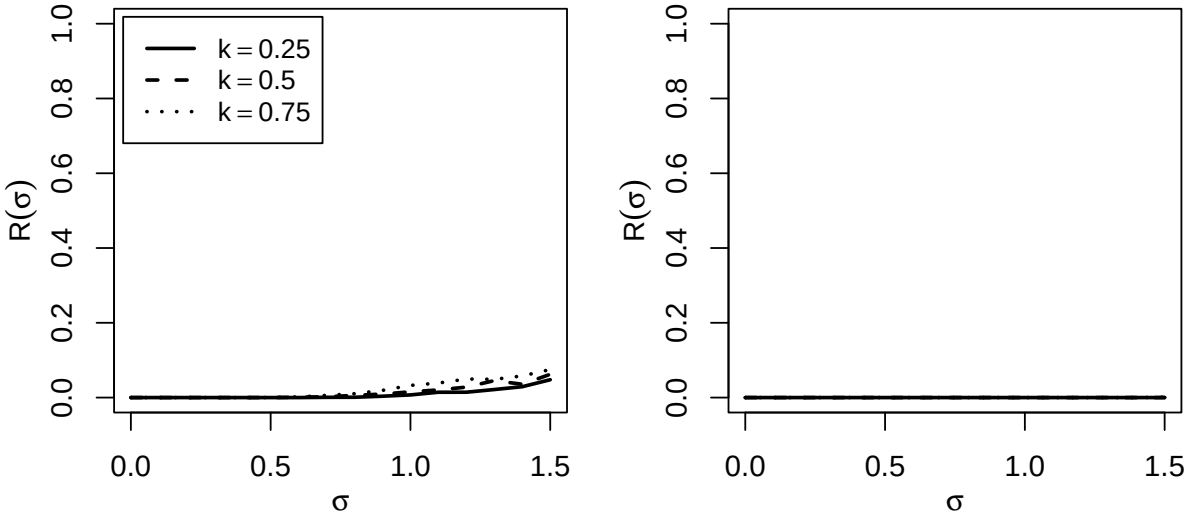


Figure 6: Plots of $R(\sigma) = P(\xi_{1,1} \in [\tau_{1,2}+1, \tau_{2,1}])/P(\xi_{1,1} \in [\tau_{1,1}+1, \tau_{1,2}])$ when $k = 0.25$, 0.5, or 0.75, σ changes from 0.1 to 1.5, and $\delta = 0.4$ (left panel) or 1.0 (right panel).

All simulation settings considered above involve only two recurrent shifts with equal durations. To evaluate the performance of the proposed method in a more general setting, we next consider an additional setting involving four recurrent shifts with different durations. Specifically, the four OC intervals are set to be

$$[\tau_{1,1}+1, \tau_{1,2}] = [101, 160], \quad [\tau_{2,1}+1, \tau_{2,2}] = [241, 280],$$
$$[\tau_{3,1}+1, \tau_{3,2}] = [351, 420], \quad [\tau_{4,1}+1, \tau_{4,2}] = [461, 540].$$

Thus, the four corresponding shift durations are 60, 40, 70, and 80 observation times, respectively. In addition, the IC periods between consecutive OC periods have different lengths, consisting of 80, 70, and 40 observations, respectively. When $\sigma = 0.2$, $\delta = 0.4$, and $k = 0.5$, with all other settings remaining the same as those in the preceding example, simulation results based on $5,000$ Monte Carlo replications are presented in Table 3.

Table 3: Simulation results for the proposed method under the considered four-recurrent-shift setting. The numbers in parentheses are the corresponding standard errors.

j	Shift Duration	$AD_{OC}^{(j)}$	$AD_{IC}^{(j)}$	$\widehat{\delta}_j$	$\widehat{\tau}_{j,1}$	$\widehat{\tau}_{j,2}$
1	60	1.1 (0.03)	1.8 (0.04)	0.42 (0.01)	99.4 (0.08)	159.8 (0.05)
2	40	1.0 (0.03)	1.4 (0.03)	0.43 (0.01)	239.8 (0.05)	279.8 (0.05)
3	70	1.1 (0.03)	1.3 (0.03)	0.39 (0.01)	350.0 (0.05)	419.9 (0.05)
4	80	1.1 (0.03)	1.2 (0.03)	0.45 (0.01)	459.2 (0.14)	539.5 (0.48)

As shown in Table 3, the proposed method consistently achieves small detection delays across all four recurrent shifts. The estimated shift magnitudes are close to the true value of $\delta = 0.4$, and the estimated change-point locations are also close to their corresponding true values. These results demonstrate that the proposed method maintains good performance in more general settings involving more than two recurrent shifts and IC and OC intervals of different durations.

3.2 Comparison with alternative approaches

Next, we compare the performance of the proposed method, denoted by NEW, with four alternative approaches: (i) the conventional one-sided CUSUM for detecting upward mean shifts, denoted by CUSUM-UP ("UP" for upward); (ii) the conventional one-sided CUSUM for detecting mean shifts in alternating directions, denoted by CUSUM-ALT ("ALT" for alternating directions); (iii) a simplified version of NEW, denoted by NEW1; and (iv) another simplified version of NEW, denoted by NEW2. These methods are briefly described below.

CUSUM-UP: This approach employs the conventional one-sided CUSUM chart, with the charting statistic defined as

$$\mathcal{C}_i^+ = \max\left(0, \mathcal{C}_{i-1}^+ + \frac{X_i - \widehat{\mu}_0^{(0)}}{\widehat{\sigma}^{(0)}} - k_{\mathcal{C}}\right), \quad \text{for } i \geq 1,$$

where $\mathcal{C}_0^+ = 0$, $\widehat{\mu}_0^{(0)}$ and $\widehat{\sigma}^{(0)}$ are estimated IC mean and standard deviation obtained from the initial IC data, and $k_{\mathcal{C}} > 0$ is a pre-specified allowance constant. The chart signals an upward mean shift whenever $\mathcal{C}_i^+ > h_{\mathcal{C}}$, where $h_{\mathcal{C}} > 0$ is the control limit. The decision rule is defined as follows. The first time that $\mathcal{C}_i^+$ exceeds $h_{\mathcal{C}}$ is taken as the signal of the first upward mean

shift. The first subsequent time that $\mathcal{C}_i^+$ falls below $h_{\mathcal{C}}$ is regarded as the signal of the end of that detected shift, and the procedure continues in this manner. The parameters $k_{\mathcal{C}}$ and $h_{\mathcal{C}}$ are chosen to match k_1 and $h_{1,1}$ in (2)-(3), respectively, so that the pre-specified nominal level of $ARL_{IC}^{(1)}$ is achieved.

CUSUM-ALT: This approach begins with the upward CUSUM chart used in CUSUM-UP and switches to a downward CUSUM chart once the previous chart signals. The downward charting statistic is defined as

$$\mathcal{C}_i^- = \min\left(0, \mathcal{C}_{i-1}^- + \frac{X_i - \widehat{\mu}_0^{(0)}}{\widehat{\sigma}^{(0)}} + k_{\mathcal{C}}\right), \quad \text{for } i \geq 1,$$

where $\mathcal{C}_0^- = 0$. The chart signals when $\mathcal{C}_i^- < -h_{\mathcal{C}}$. The parameters $k_{\mathcal{C}}$ and $h_{\mathcal{C}}$ here are chosen to be the same as those in CUSUM-UP. After the downward CUSUM chart signals, the procedure continues by alternating between the upward and downward charts.

NEW1: This approach is identical to NEW, except that we set $C_{\xi_{1,1}}^- = 0$ in (9) rather than computing its value using the observations in $[\widehat{\tau}_{1,1} + 1, \xi_{1,1}]$. This variant is included in the comparison to demonstrate that incorporating the observed data between the estimated change point and the signal time of the previous chart improves the performance of the subsequent process monitoring.

NEW2: This approach is identical to NEW, except that the IC mean and variance estimates are updated at each time point, rather than using the spring-length mechanism described in (4) to mitigate potential contamination from OC observations. This variant is included to demonstrate that incorporating the spring-length concept in updating the IC mean and variance estimates improves the performance of NEW.

When $k = 0.5$ and the other settings are the same as those in Table 1, the calculated values of $AD_{OC}^{(1)}$ and $AD_{IC}^{(1)}$ for the five approaches CUSUM-UP, CUSUM-ALT, NEW1, NEW2, and NEW, based on 5,000 replicated process-monitoring simulations are presented in Figure 7, where δ_1 varies over $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, and $\sigma = 0.2$ (left column) or 0.5 (right column). When $\delta_1 = 0$, $AD_{OC}^{(1)}$ and $AD_{IC}^{(1)}$ reduce to $ARL_{IC}^{(1)}$ and $ARL_{OC}^{(1)}$, respectively.

From the figure, several conclusions can be drawn. First, for detecting the starting point of the first shift, CUSUM-UP and CUSUM-ALT exhibit identical performance, as indicated by their

overlapping dotted and short-dashed lines for $AD_{OC}^{(1)}$. Their performance is slightly inferior to that of NEW1, NEW2, and NEW, which also perform identically in this setting. This difference arises because CUSUM-UP and CUSUM-ALT estimate the IC parameters using only the initial IC data, whereas NEW1, NEW2, and NEW are self-starting procedures that update IC parameter estimates during process monitoring. Second, for detecting the ending point of the first shift, CUSUM-UP and CUSUM-ALT do not provide reliable in-control performance, as their $ARL_{OC}^{(1)}$ values deviate substantially from the nominal value of 50. In contrast, NEW1, NEW2, and NEW demonstrate much more reliable in-control performance, with NEW achieving the best performance among the five approaches. Third, CUSUM-UP and CUSUM-ALT have limited power for detecting the ending point of the first shift. NEW1, NEW2, and NEW perform substantially better in this regard, with NEW again exhibiting the strongest overall performance.

The corresponding results for $AD_{OC}^{(2)}$ and $AD_{IC}^{(2)}$ for the five approaches in detecting the second shift are presented in Figure 8. In this setting, CUSUM-UP and CUSUM-ALT have little power to detect the starting point of the second shift, primarily because they often fail to reliably detect the ending point of the first shift, as shown in Figure 7. All other conclusions drawn from Figure 7 remain applicable here.

This example illustrates that (i) the naive approaches, CUSUM-UP and CUSUM-ALT, are ineffective for detecting multiple shifts, (ii) modifying the value of $C_{\xi_{1,1}}^{-}$ in (9) using observations from $[\widehat{\tau}_{1,1}+1, \xi_{1,1}]$ is beneficial, and (iii) incorporating the spring-length concept when updating the IC mean and variance estimates can further improve the performance of NEW.

Next, we compare the proposed method with the representative multiple change-point detection approach based on dynamic programming developed by Hawkins (2001). As discussed in Section 1, like other multiple change-point detection methods in the literature, this approach is retrospective and is applied to a sequence of observations with a fixed sample size. We consider two versions of this approach: (i) MCPD1, which is applied to the first 400 observations of the monitored sequence, and (ii) MCPD2, which is applied to the first 1,000 observations. Because MCPD1 and MCPD2 do not provide sequential signals of mean shifts, the comparison focuses on the accuracy of the estimated change-point locations rather than on signal times or detection delays. Specifically, for a true change point τ and its estimate $\widehat{\tau}$, we consider the absolute error $|\widehat{\tau}-\tau|$, with smaller values indicating more accurate estimation.

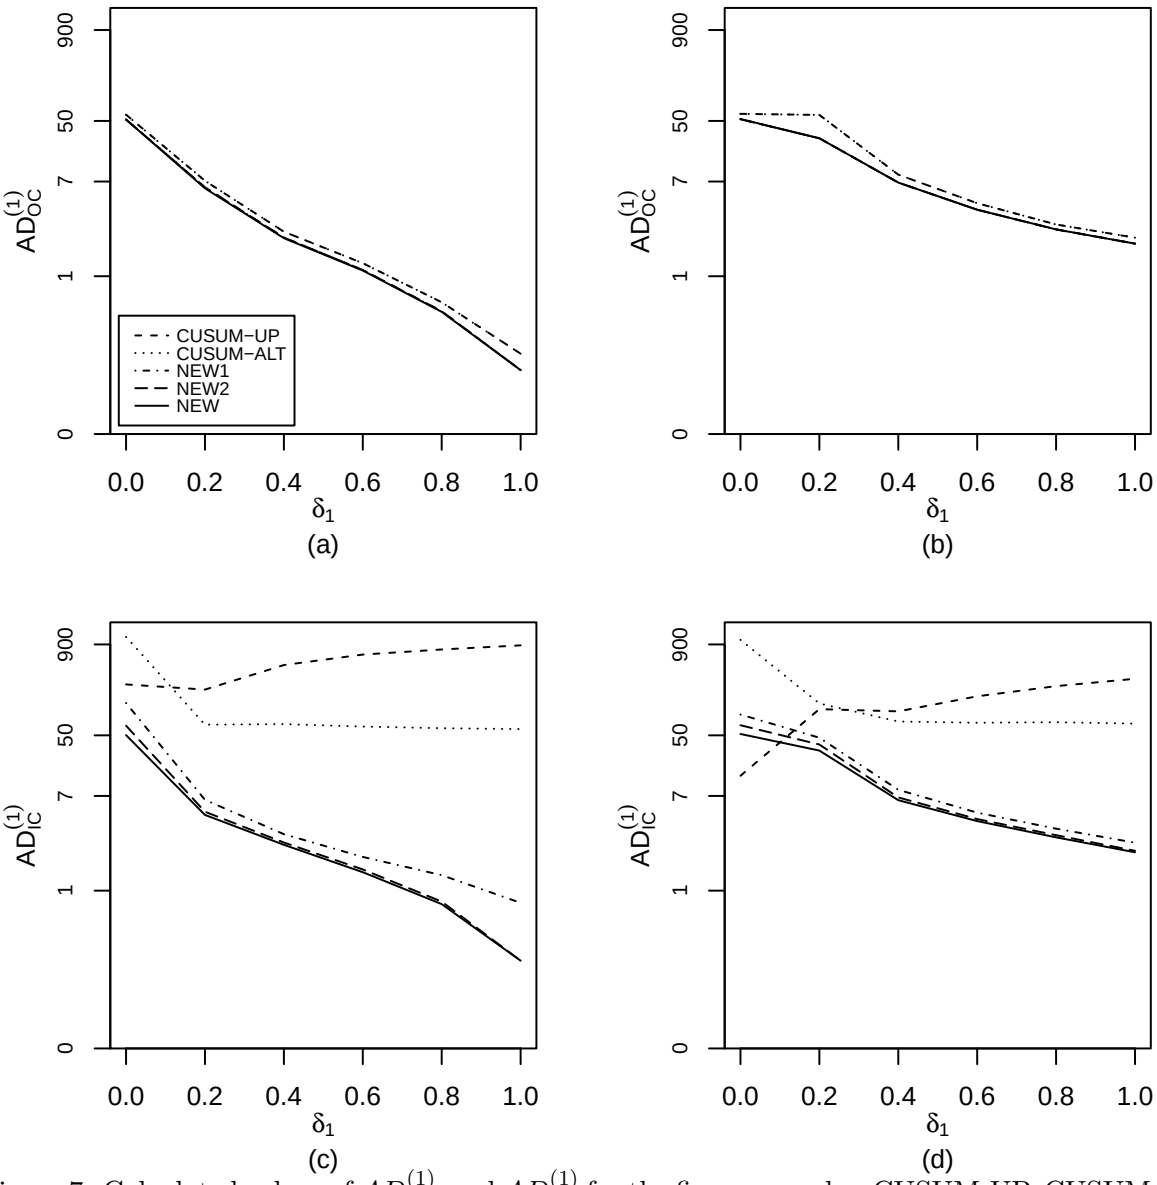


Figure 7: Calculated values of $AD_{OC}^{(1)}$ and $AD_{IC}^{(1)}$ for the five approaches CUSUM-UP, CUSUM-ALT, NEW1, NEW2, and NEW to detect the first shift, based on 5,000 replicated process-monitoring simulations, when $k = 0.5$, δ_1 varies over $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, and $\sigma = 0.2$ (left column) or 0.5 (right column).

When $k = 0.5$, $\delta_1 = \delta_2 = \delta \in 0.2, 0.4, 0.6, 0.8, 1.0$, and all other settings are the same as those used in Figure 7, the mean absolute errors (MAEs) of the estimated change-point locations obtained by NEW, MCPD1, and MCPD2, based on 5,000 Monte Carlo replications, are presented in Figure 9.

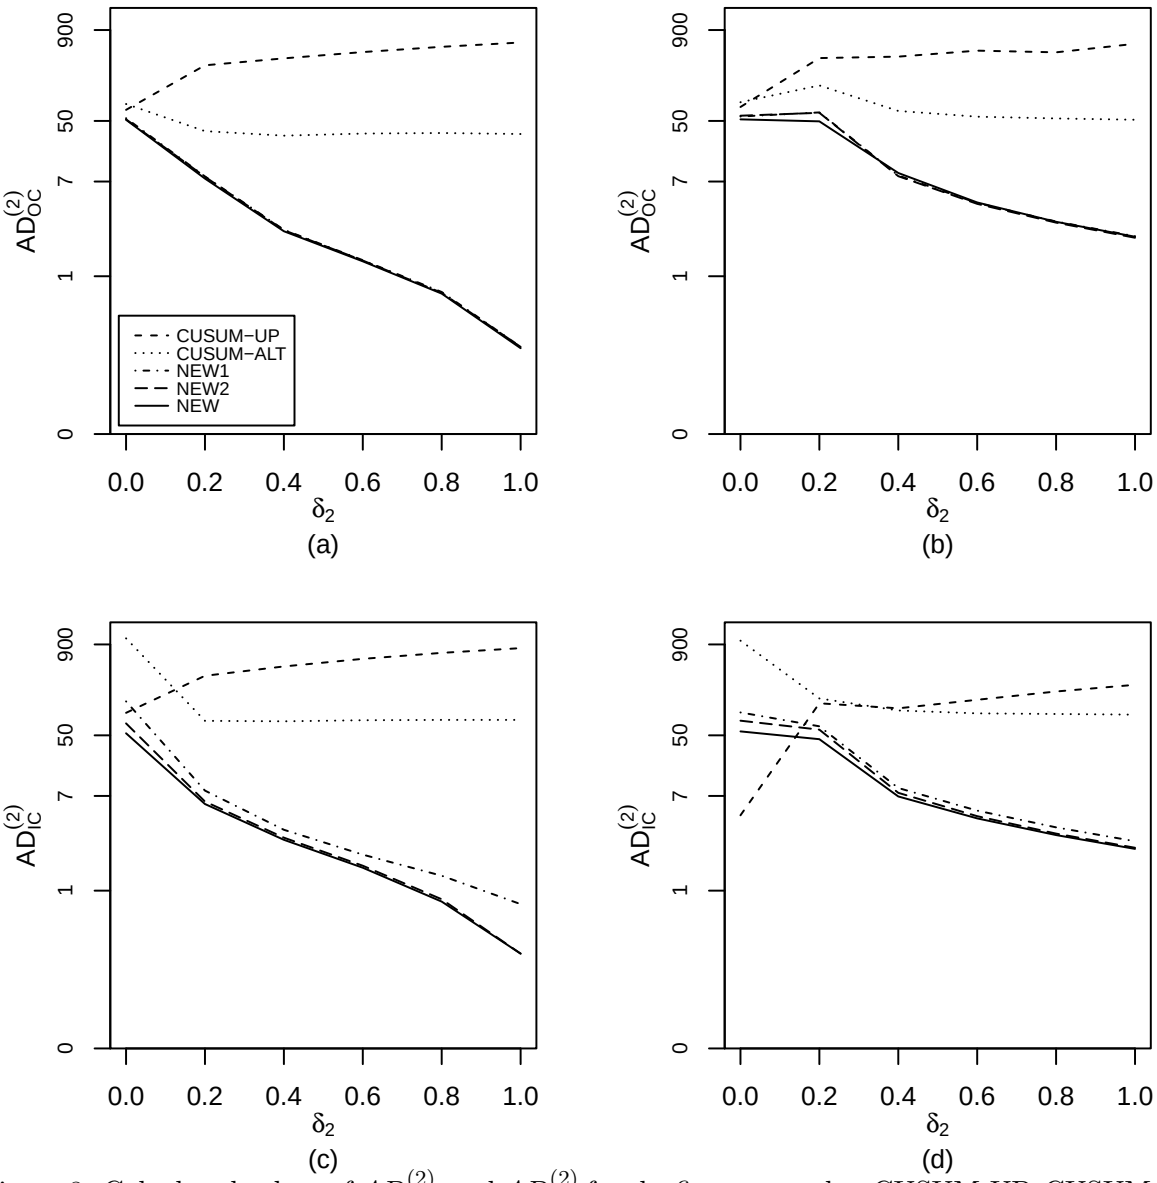


Figure 8: Calculated values of $AD_{OC}^{(2)}$ and $AD_{IC}^{(2)}$ for the five approaches CUSUM-UP, CUSUM-ALT, NEW1, NEW2, and NEW to detect the second shift, based on 5,000 replicated process-monitoring simulations, when $k = 0.5$, δ_2 varies over $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$, and $\sigma = 0.2$ (left column) or 0.5 (right column).

From Figure 9, the proposed method demonstrates competitive change-point estimation accuracy relative to the retrospective multiple change-point detection approaches. When $\sigma = 0.2$, the MAEs are generally small for all three methods. When $\sigma = 0.5$, the proposed method, NEW, substantially outperforms MCPD1 and MCPD2 when $\delta = 0.2$, with similar performance as δ increases. Overall, these results indicate that NEW maintains good change-point localization accuracy, par-

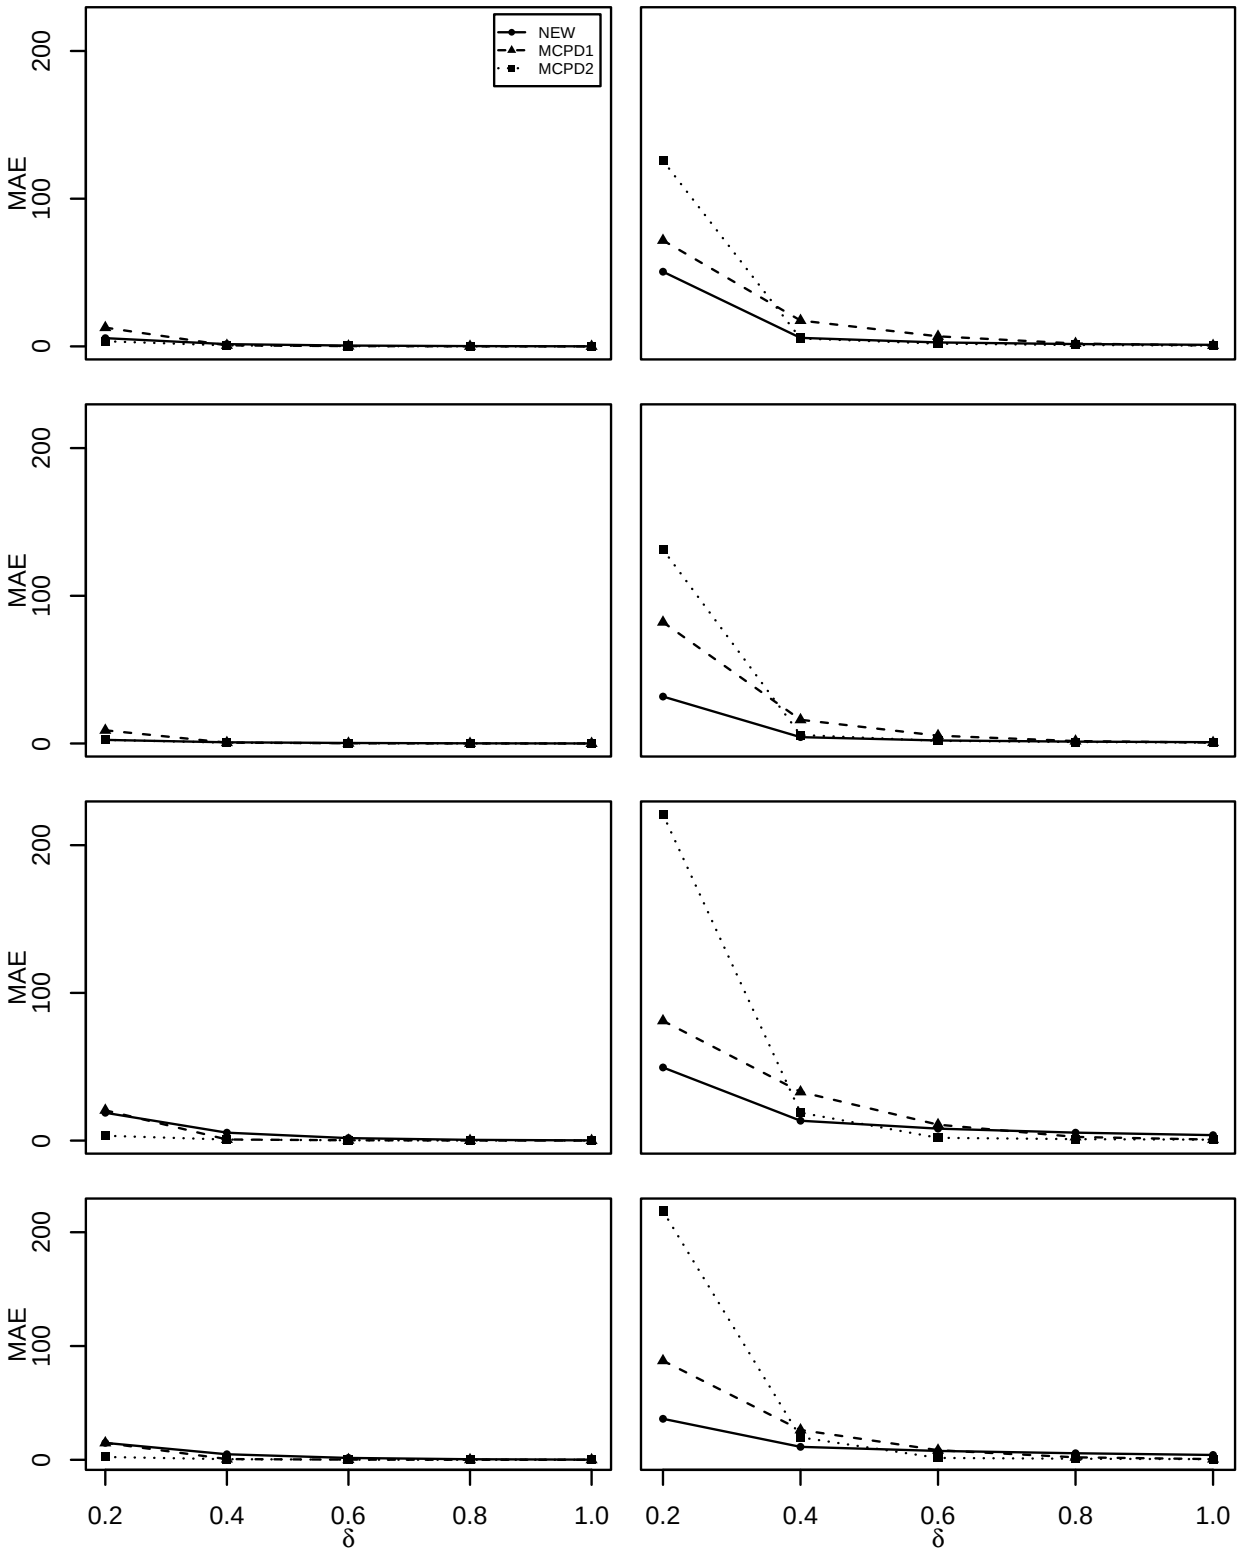


Figure 9: Mean absolute errors (MAEs) of the estimated change-point locations obtained by NEW, MCPD1, and MCPD2 when $\delta_1 = \delta_2 = \delta \in \{0.2, 0.4, 0.6, 0.8, 1.0\}$ and $\sigma = 0.2$ (left column) or 0.5 (right column). The four rows correspond to the results for the four change points $\tau_{1,1}, \tau_{1,2}, \tau_{2,1}$, and $\tau_{2,2}$, respectively.

ticularly under more challenging settings involving smaller shifts and greater data variability.

4 An Application

To demonstrate the practical utility of the proposed monitoring procedure, we applied it to a human activity recognition dataset available on Kaggle (https://www.kaggle.com/datasets/uciml/human-activity-recognition-with-smartphones/data). The data were collected from an experiment in which a participant alternated between walking and stationary activities while wearing a smartphone positioned at the waist. The smartphone's embedded accelerometer and gyroscope sensors were used to capture motion signals. Specifically, the sensors recorded three-axis linear acceleration and three-axis angular velocity at a constant sampling rate of 50 Hz, corresponding to 50 observations per second for each axis.

From the raw signals, a number of informative features can be extracted. In this study, we focus on the feature "Max Acceleration (Y-axis)," defined as the maximum Y-axis acceleration within each sliding window of 2.56 seconds. The sliding windows are constructed with 50% overlap between consecutive windows. The original observations of this variable exhibit serial correlation. Because the proposed method is designed for independent, normally distributed data, a preprocessing step is performed using piecewise ARIMA modeling followed by a rank-based inverse normal transformation (Van der Waerden, 1952). Specifically, for each phase corresponding to walking or stationary activity, an ARIMA model is first fitted to the original observations. A rank-based inverse normal transformation is then applied to the resulting residuals, and the transformed data are subsequently subjected to a location-scale transformation so that their mean matches the mean of the original observations in the corresponding phase, while their standard deviation is set equal to the pooled sample standard deviation calculated from the original observations across all four phases. After this preprocessing, potential shifts in the original data across different phases are preserved. The preprocessed observations of "Max Acceleration (Y-axis)" are shown in Figure 10. As seen in the figure, there are four alternating phases of walking and stationary activities. The p-values of the Ljung–Box tests for the four phases are 0.3653, 0.4062, 0.6294, and 0.5619, respectively, indicating no significant autocorrelation; the p-values of the Shapiro–Wilk tests are all close to 1.0, suggesting no evidence against normality.

We then apply the five methods—CUSUM-UP, CUSUM-ALT, NEW1, NEW2, and NEW—to

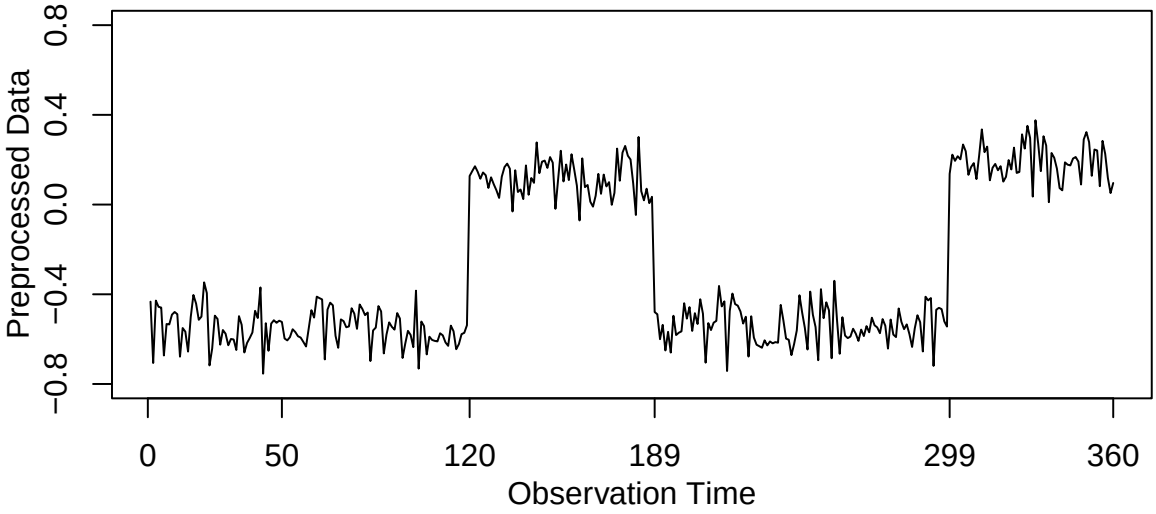


Figure 10: Preprocessed observations of "Max Acceleration (Y-axis)" for a participant alternating between walking and stationary activities, obtained from a human activity recognition experiment.

this dataset. The first 50 observations are treated as the IC data. For the proposed method NEW, we set $k = 0.75$ and $ARL_{IC}^{(j)} = ARL_{OC}^{(j)} = 200$ for each j. The other methods are configured accordingly, as described in Section 3.2. The resulting control charts are presented in Figure 11. As shown in the figure, both CUSUM-UP and CUSUM-ALT successfully detect the first change point around the 120th observation, corresponding to the transition from stationary to walking activity. However, they fail to detect the remaining two change points, likely due to the inherent memory structure of their charting statistics. Methods NEW1 and NEW2 accurately detect the first and third change points around the 120th and 299th observations, respectively, but both signal the second change point prematurely, likely in response to short-term fluctuations rather than a true structural change. In contrast, the proposed method NEW successfully detects all three change points in this example. Its charting statistic resets more rapidly after each detected shift, enabling the monitoring scheme to more effectively track multiple structural changes over time. This reduced memory effect allows for faster adaptation to new regimes while maintaining stability within relatively homogeneous activity segments.

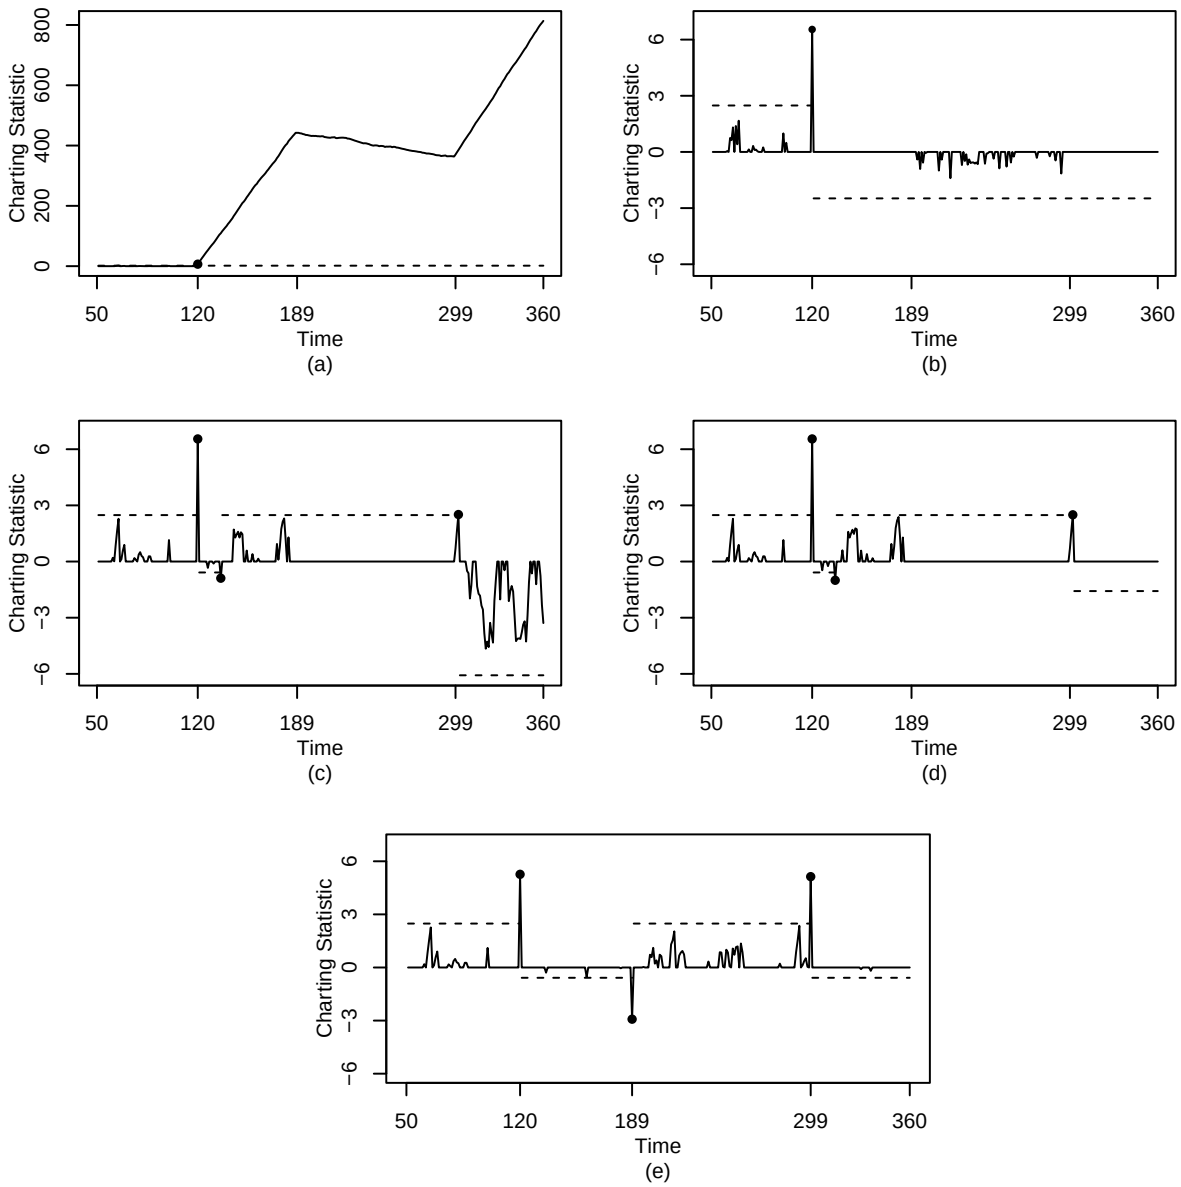


Figure 11: Control charts CUSUM-UP, CUSUM-ALT, NEW1, NEW2, and NEW are displayed in plots (a)-(e), respectively, for monitoring the data shown in Figure 10, starting from the 51st observation. In each plot, horizontal dashed lines indicate the control limits, and the dark points denote the signal times.

5 Concluding Remarks

This paper investigates the problem of sequential process monitoring in the presence of recurrent mean shifts, a setting that arises in many modern applications involving continuous data streams. Traditional control charts are primarily designed to detect a single distributional change and are

generally inadequate for scenarios in which the underlying process may repeatedly transition between IC and OC states. To address this limitation, we propose a new online monitoring framework that sequentially detects both the starting and ending times of multiple process mean shifts through a coordinated sequence of control charts.

The proposed method incorporates several key features. First, it introduces a structured framework that alternates between detecting the starting and ending times of upward mean shifts, allowing effective tracking of recurring changes in the process. Second, it integrates self-starting mechanisms with a spring-length-based updating scheme (cf., Expression (4)) to mitigate contamination from OC observations, thereby improving estimation of IC parameters during monitoring. Third, it combines control charting with change-point estimation to enhance the detection of shift ending times. Numerical studies demonstrate that the proposed approach performs well across a variety of settings.

Despite these advantages, several limitations of the current work suggest directions for future research. The proposed framework is developed under the assumptions that observations are independent and normally distributed, and that shifts occur only in the mean and in a single direction. Extending the methodology to accommodate serial dependence, non-normal distributions, and more complex shift patterns would substantially broaden its applicability. In addition, incorporating adaptive or data-driven selection of tuning parameters, such as allowance constants and control limits, could further enhance the method's robustness and practical usability.

In summary, this work provides a new perspective and practical tools for monitoring sequential processes with recurrent shifts. We hope it will stimulate further research on sequential decision-making methods for complex and evolving processes encountered in modern applications.

Acknowledgments

The authors sincerely thank the Editor and the two referees for their many constructive comments and valuable suggestions, which have greatly improved the quality of the paper.

Data Availability Statement

The data that support the findings of this study are available on request from the corresponding author.

Disclosure of Interest

The author has no relevant financial or non-financial competing interests to report.

References

Anastasiou, A., and Fryzlewicz, P. (2022), "Detecting multiple generalized change-points by isolating single ones," *Metrika*, **85**, 141–174.

Apley D.W., and Tsung, F. (2002), "The autoregressive T^2 chart for monitoring univariate autocorrelated processes," *Journal of Quality Technology*, **34**, 80–96.

Chakraborti, S., and Graham, M. (2019), "Nonparametric (distribution-free) control charts: An updated overview and some results," *Quality Engineering*, **31**, 523–544.

Chatterjee, S., and Qiu, P. (2009), "Distribution-free cumulative sum control charts using bootstrap-based control limits," *Annals of Applied Statistics*, **3**, 349–369.

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**, 1415–1428.

Cho, H., and Fryzlewicz, P. (2020), "Multiple change point detection under serial dependence: Wild contrast maximisation and gappy Schwarz algorithm," *Journal of Time Series Analysis*, **45**, DOI:10.1111/jtsa.12722.

Gao, Z., Xiao, X., Fang, Y.-P., Rao, J., and Mo, H. (2024), "A selective review on information criteria in multiple change point detection," *Entropy*, **26**, 50.

Hawkins, D.M. (2001), "Fitting multiple change-point models to data," *Computational Statistics and Data Analysis*, **37**, 323-341.

Hawkins, D.M., and Olwell, D.H. (1998), *Cumulative Sum Charts and Charting for Quality Improvement*, New York: Springer.

Hawkins, D.M., Qiu, P., and Kang, C.W. (2003), "The changepoint model for statistical process control," *Journal of Quality Technology*, **35**, 355–366.

Montgomery, D.C. (2012), *Introduction to Statistical Quality Control*, New York: John Wiley & Sons.

Niu, Y.S., Hao, N., and Zhang, H. (2016), "Multiple change-point detection: A selective overview," *Statistical Science*, **31**, 611–623.

Page, E.S. (1954), "Continuous inspection scheme," *Biometrika*, **41**, 100–115.

Qiu, P. (2013), *Introduction to Statistical Process Control*, Boca Raton, FL: Chapman Hall/CRC.

Qiu, P. (2018), "Some perspectives on nonparametric statistical process control," *Journal of Quality Technology*, **50**, 49–65.

Qiu, P. (2024), *Statistical Methods for Dynamic Disease Screening and Spatio-Temporal Disease Surveillance*, Boca Raton, FL: Chapman Hall/CRC.

Roberts, S.V. (1959), "Control chart tests based on geometric moving averages," *Technometrics*, **1**, 239–250.

Runger, G. (2002), "Assignable causes and autocorrelation: control charts for observations or residuals?" *Journal of Quality Technology*, **34**, 165–170.

Shewhart, W.A. (1931), *Economic Control of Quality of Manufactured Product,* New York: D. Van Nostrand Company.

Van der Waerden, B.L. (1952). "Order tests for the two-sample problem and their power," *Indagationes Mathematicae*, **14**, 453–458.

Xie, X., Qian, N., and Qiu, P. (2024), "Online monitoring of air quality using PCA-based sequential learning," *Annals of Applied Statistics*, **18**, 600–625.

Zou, C., Yin, G., Feng, L., and Wang, Z. (2014), "Nonparametric maximum likelihood approach to multiple change-point problems," *Annals of Statistics*, **42**, 970–1002.