

Nonparametric Monitoring of Multiple Count Data

Peihua Qiu¹, Zhen He² and Zhiqiong Wang³

¹Department of Biostatistics

University of Florida, Gainesville, United States

²College of Management and Economics

Tianjin University, Tianjin, China

³School of Management

Tianjin University of Technology, Tianjin, China

Abstract

Process monitoring of multiple count data has received considerable attention recently in the statistical process control literature. Most existing methods on this topic are based on parametric modeling of the observed process data. However, the assumed parametric models are often invalid in practice, leading to unreliable performance of the related control charts. In this paper, we first show the consequence of using a parametric control chart in cases when the underlying parametric distribution is invalid. Then, we thoroughly investigate the performance of some parametric and nonparametric control charts in monitoring multiple count data. Our numerical results show that nonparametric methods can provide a more reliable and effective process monitoring in such cases. A real-data example about the crime log of the University of Florida Police Department is used for illustrating the implementation of the related control charts.

Keywords: Distribution-free; Log-linear modeling; Multiple count data; Nonparametric procedures; Poisson distribution; Statistical process control.

1 Introduction

Due to the spreading usage of sensors and the rapid development in information storage technology, it is common to collect and analyze several correlated quality characteristics of a process simultaneously. On-line monitoring of these quality characteristics separately may not be effective in detecting process changes. Therefore, many multivariate control charts have been proposed in the literature to solve this problem more effectively (cf. [Qiu, 2014](#), Chapters 7 and 9). Many of these charts are designed for monitoring multiple continuous quality characteristics. However, multiple count data are also common in practice. Examples include the number of delaminations at different positions of a printed circuit board in manufacturing industry ([Wang et al., 2017](#)), incidences of different types of diseases in epidemiology ([Chen et al., 2015](#)), and numbers of purchases of different products in marketing ([Brijs et al., 2004](#)). Therefore, monitoring of multiple count data is important. But, existing research on this research problem is limited. This paper aims to fill the gap by discussing different strategies to solve the problem.

Traditionally, multiple count data could be modeled by either multivariate binomial/multinomial distributions or multivariate Poisson distributions ([Johnson et al., 1997](#)). [Patel \(1973\)](#) first proposed a T^2 -type chart for monitoring both multiple binomial and multiple Poisson data by approximating the related multivariate binomial or Poisson distribution by a multivariate normal distribution. Some subsequent research in this direction includes [Lu et al. \(1998\)](#), [Chiu and Kuo \(2010\)](#) and [Li et al. \(2014\)](#). Some other existing research handles multiple Poisson count data based on multivariate Poisson models ([Holgate, 1964](#); [Karlis, 2003](#); [Karlis and Meligkotsidou, 2005](#)). Parametric control charts proposed in this direction include [Chiu and Kuo \(2008\)](#), [Lee Ho and Branco Costa \(2009\)](#), [Laungrungrong et al. \(2011\)](#) and [He et al. \(2014\)](#). [Jones et al. \(1999\)](#) and [Skinner et al. \(2003\)](#) proposed two control charts in special cases when different components of the multiple Poisson data are assumed independent. According to [Schmidt and Rodriguez \(2011\)](#), the control charts based on the multivariate Poisson models have two major limitations. One is that they do not allow negative correlation between any pair of two variables, and thus lack generality for solv-

ing real-world problems. To address this issue, [Chen et al. \(2015\)](#) developed a multivariate exponentially weighted moving average (MEWMA) control chart to monitor multiple count data by using the Poisson log-normal distribution. The second limitation is that the multivariate Poisson models cannot describe over-dispersed count data well. To overcome this limitation, [Saghir and Lin \(2014\)](#) proposed a Shewhart-type multivariate control chart based on the Conway-Maxwell-Poisson (COM-Poisson) distribution, which can accommodate over-dispersion in the observed data. Based on the same model as that in [Chen et al. \(2015\)](#), [Das et al. \(2016\)](#) proposed a control chart under the framework of the likelihood ratio test for monitoring over-dispersed multiple count data. Several overviews on monitoring multiple count data can be found in papers such as [Topalidou and Psarakis \(2009\)](#), [Saghir and Lin \(2015\)](#), and [Cozzucoli and Marozzi \(2017\)](#).

All existing methods mentioned above on monitoring multiple count data are parametric in the sense that they rely on a specific parametric distribution. In the literature on monitoring multiple continuous data, it has been well demonstrated that such parametric control charts are not reliable to use in cases when the assumed parametric distributions are invalid (cf. [Qiu, 2014](#), Chapters 8 and 9). In practice, the assumed parametric distributions are rarely valid, especially in multivariate cases, because there are so many different factors that may affect the quality variables and the relationship among them, and such factors are usually not included in the observed data. To address this important issue, there are several multivariate nonparametric control charts developed in the literature, including those discussed in [Qiu and Hawkins \(2001\)](#), [Qiu and Hawkins \(2003\)](#), [Qiu \(2008\)](#), [Zou and Tsung \(2011\)](#), [Boone and Chakraborti \(2012\)](#), [Zou et al. \(2012\)](#), [Chen et al. \(2016\)](#) and [Qiu \(2018\)](#). But, all these methods are suggested for monitoring multiple continuous data, and no nonparametric methods can be found for monitoring multiple count data.

In this paper, we first introduce several representative parametric control charts for monitoring multiple count data in Section 2. Then, we propose a nonparametric control chart for the same purpose in Section 3. This chart is constructed under the framework of data categorization and log-linear modeling that was first discussed in [Qiu \(2008\)](#). Then, the related control charts are compared in Section 4 by a large simulation study regarding their in-control (IC) and out-of-control (OC) performance. A real-data example is used for

illustrating the implementation of the related methods in Section 5. Finally, several remarks conclude the article in Section 6.

2 Parametric Methods for Monitoring Multiple Count Data

In the literature, there have been some methods proposed for monitoring multiple count data, as discussed in Section 1. In this section, we briefly introduce two recent ones that are considered flexible and versatile among the existing methods in this area.

Let X_j , for $j = 1, 2, \dots, p$, be the number of defects or nonconformities of the quality characteristic j , and $\mathbf{X} = (X_1, X_2, \dots, X_p)'$. The first recent method is the MEWMA control chart proposed by [Chen et al. \(2015\)](#). They used the multivariate Poisson log-normal distribution to model the multiple count data. More specifically, they assumed that X_j followed a univariate Poisson distribution with rate parameter θ_j , for $j = 1, 2, \dots, p$. To allow correlation among different components of \mathbf{X} , they further assumed that $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_p)'$ was a random vector and followed a multivariate log-normal distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. In that framework, the IC value of $\boldsymbol{\mu}$, denoted as $\boldsymbol{\mu}_0$, and the IC value of $\boldsymbol{\Sigma}$, denoted as $\boldsymbol{\Sigma}_0$, are assumed known or they can be estimated accurately from an IC dataset. Then, when the related process is IC, the observed data $\mathbf{X}(n)$, for $n \geq 1$, would have mean vector \mathbf{m}_0 and covariance matrix \mathbf{V}_0 , both of which depend on $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$ implicitly. Then, the MEWMA charting statistic is defined as

$$\mathbf{E}_n = \mathbf{R}[\mathbf{X}(n) - \mathbf{m}_0] + (\mathbf{I} - \mathbf{R})\mathbf{E}_{n-1}, \quad \text{for } n \geq 1,$$

where $\mathbf{E}_0 = \mathbf{0}$, \mathbf{I} is the identity matrix, and \mathbf{R} is a smoothing matrix with equal diagonal elements and equal off-diagonal elements as discussed in [Hawkins et al. \(2007\)](#). The chart gives a signal of process distributional shift when

$$T_n^2 = \mathbf{E}_n' \mathbf{W}_n^{-1} \mathbf{E}_n > h_T, \quad (1)$$

where \mathbf{W}_n is the covariance matrix of \mathbf{E}_n , and $h_T > 0$ is a control limit. Once the smoothing matrix \mathbf{R} is chosen properly, h_T is chosen to achieve a given value of the IC average run length, denoted as ARL_0 .

The second parametric chart to introduce is the one based on the COM-Poisson distributional assumption that was discussed in [Saghir and Lin \(2014\)](#). The COM-Poisson distribution generalizes the regular Poisson distribution by introducing a location parameter and a dispersion parameter to allow under-dispersion or over-dispersion in the observed data. In that method, it is assumed that each component of $\mathbf{X}(n)$ has a COM-Poisson distribution and $D(n)$ denotes the summation of all components at the time point n . Namely,

$$D(n) = \sum_{j=1}^p X_j(n).$$

The mean and variance of $D(n)$ can be calculated based on the COM-Poisson distributional assumption. Then, [Saghir and Lin \(2014\)](#) suggested a Shewhart-type control chart with the control limit determined by the mean and standard deviation of $D(n)$. As well demonstrated in the literature, CUSUM charts are usually more effective in detecting persistent shifts than Shewhart charts (cf. [Qiu, 2014](#), Chapter 4). For that reason and for a fair comparison among different charts in Sections 4 and 5 below, we also define a CUSUM chart based on $D(n)$ here. Let $u_{0,D}^+ = u_{0,D}^- = 0$, and

$$\begin{aligned} u_{n,D}^+ &= \max(0, u_{n-1,D}^+ + (D(n) - D_0) - k_D), \\ u_{n,D}^- &= \min(0, u_{n-1,D}^- + (D(n) - D_0) + k_D), \quad \text{for } n \geq 1, \end{aligned}$$

where $k_D > 0$ is an allowance constant, and D_0 is the IC mean of $D(n)$ that can be estimated from an IC data. This CUSUM chart, denoted as DCUSUM hereafter, signals a shift in $\mathbf{X}(n)$ when

$$u_{n,D}^+ > h_D \quad \text{or} \quad u_{n,D}^- < -h_D, \tag{2}$$

where $h_D > 0$ is a control limit chosen to achieve a given ARL_0 value.

3 Nonparametric Monitoring of Multiple Count Data

The performance of the parametric methods described in the previous section depends heavily on the validity of the assumed parametric distributions. In the next section, we will show that their results are unreliable in cases when the assumed distributions are invalid. In practice, parametric distributions are often invalid to describe the observed quality characteristics of a specific process because the quality characteristics are often affected by many different factors and the mechanism of this impact is usually too complicated to describe by a parametric model. Therefore, development of nonparametric control charts is important for monitoring multiple count data. In this section, we propose a nonparametric control chart based on log-linear modeling.

As pointed out by [Qiu \(2008\)](#), the major difficulty in describing the distribution of multiple quality characteristics is due to the complicated relationship among them. If the related variables are categorical, then the log-linear modeling would be a powerful tool for describing such relationship. Therefore, it is natural to categorize the original quality characteristics $\mathbf{X}(n) = (X_1(n), X_2(n), \dots, X_p(n))'$, and then apply the log-linear modeling approach to the categorized data. To this end, let m_j be the IC median of $X_j(n)$, $\mathbf{Y}(n) = (Y_1(n), Y_2(n), \dots, Y_p(n))'$, and

$$Y_j(n) = I(X_j(n) > m_j) \quad \text{for } j = 1, 2, \dots, p, \quad (3)$$

where $I(a)$ is an indicator function that equals 1 if a is “true” and 0 otherwise. Then, $\mathbf{Y}(n)$ is a binary version of $\mathbf{X}(n)$, and the IC distribution of each of its components should be Bernoulli(0.5). For count data, however, it is often difficult to find the exact median value m_j due to the discreteness of the original variable $X_j(n)$. To address this issue, there are several possible solutions. One is to add a small random number to $X_j(n)$ before categorization, to make the distribution of $X_j(n)$ less discrete, as discussed in [Qiu and Li \(2011\)](#). Alternatively, in Equation (3), the IC median m_j can be replaced by the more general IC r th quantile of $X_j(n)$, with r being a real number in $(0, 1)$ that is achievable and as close to 0.5 as possible. We consider using the median in Equation (3) because the resulting joint distribution of $\mathbf{Y}(n)$ could be more efficiently

estimated by the log-linear modeling approach (Agresti, 2013). Theoretically, it can be checked that as long as the IC distribution $F(\mathbf{x})$ of $\mathbf{X}(n)$ has a positive probability mass in any neighborhood of the IC median vector $(m_1, m_2, \dots, m_p)'$, the distribution of $\mathbf{Y}(n)$ would be changed by any mean shift in the original data $\mathbf{X}(n)$. The reason is that $\mathbf{X}(n)$ has a mean shift if and only if it has a median shift, and a median shift in $\mathbf{X}(n)$ is equivalent to a distributional shift of $\mathbf{Y}(n)$, according to Equation (3). Thus, if we are concerned about mean shifts in $\mathbf{X}(n)$, we can just monitor the categorized data $\mathbf{Y}(n)$.

Next, we briefly describe the log-linear modeling of the categorized data $\mathbf{Y}(n)$. For simplicity, the dimension is fixed at $p = 3$, and the modeling can be discussed similarly for cases with $p > 3$. Let $O_{j_1 j_2 j_3}$ be the observed cell count of the (j_1, j_2, j_3) th cell of the three-way contingency table of an IC data, with the three binary variables $Y_1(n)$, $Y_2(n)$ and $Y_3(n)$ defined in Equation (3) as classifiers, for $j_1, j_2, j_3 = 0, 1$. Then, a saturated log-linear model is defined as

$$\log(O_{j_1 j_2 j_3}) = \lambda + \lambda_{j_1}^{Y_1} + \lambda_{j_2}^{Y_2} + \lambda_{j_3}^{Y_3} + \lambda_{j_1 j_2}^{Y_1 Y_2} + \lambda_{j_1 j_3}^{Y_1 Y_3} + \lambda_{j_2 j_3}^{Y_2 Y_3} + \lambda_{j_1 j_2 j_3}^{Y_1 Y_2 Y_3}, \quad \text{for } j_1, j_2, j_3 = 0, 1, \quad (4)$$

where λ is a constant term, $\lambda_{j_1}^{Y_1}$, $\lambda_{j_2}^{Y_2}$ and $\lambda_{j_3}^{Y_3}$ are the main effects of Y_1 , Y_2 and Y_3 , respectively, $\lambda_{j_1 j_2}^{Y_1 Y_2}$, $\lambda_{j_1 j_3}^{Y_1 Y_3}$ and $\lambda_{j_2 j_3}^{Y_2 Y_3}$ are the two-way interaction terms, and $\lambda_{j_1 j_2 j_3}^{Y_1 Y_2 Y_3}$ is the three-way interaction term (cf., Agresti, 2013).

By removing certain terms from the model (4), the resulting models can describe all kinds of possible association among Y_1 , Y_2 and Y_3 . For convenience, let us use a notation that lists the highest-order term(s) of each variable to denote a specific log-linear model. Then, the saturated log-linear model (4) can be denoted as $(Y_1 Y_2 Y_3)$. If the three-way interaction term is excluded from the saturated model, then the conditional association between any two of Y_1 , Y_2 and Y_3 are identical at the two levels of the remaining variable. That is, each pair of variables has homogeneous association. This model is denoted as $(Y_1 Y_2, Y_1 Y_3, Y_2 Y_3)$. Similarly, the model $(Y_1 Y_2, Y_1 Y_3)$ denotes the one with the two-way interaction term $\lambda_{j_2 j_3}^{Y_2 Y_3}$ and the three-way interaction term $\lambda_{j_1 j_2 j_3}^{Y_1 Y_2 Y_3}$ excluded from the saturated model. In that model, Y_2 and Y_3 are assumed conditionally independent given Y_1 . These notations of log-linear models have been used in Agresti (2013).

For a specific categorical dataset, a proper log-linear model should be selected and estimated. To this end, the likelihood ratio test statistic G^2 and the hierarchy principle are routinely used for model selection. The hierarchy principle requires that all lower-order terms should be included in a model if a higher-order interaction term is in the model. In all simulation studies in this paper, we use the backward elimination procedure for model selection. Namely, we start from the saturated model. Then, in each step of model selection, only one term is considered to be deleted by performing a likelihood ratio test with the significance level of 0.05. For instance, when comparing the saturated model $(Y_1Y_2Y_3)$ (denoted as M_1) with the reduced model (Y_1Y_2, Y_1Y_3, Y_2Y_3) (denoted as M_0), the likelihood ratio test statistic is $G^2(M_0|M_1) = -2 \log(l_{M_0}/l_{M_1})$, where l_{M_0} and l_{M_1} denote the likelihood functions of the reduced model M_0 and the saturated model M_1 , respectively. The null distribution of $G^2(M_0|M_1)$ is the $\chi^2(1)$ distribution, based on which the p -value of the test can be calculated for making decisions. This model selection process continues until no terms need to be deleted. Once the final model is determined, its estimation can be achieved by using the iterative weighted least square procedure. It should be pointed out that model selection and estimation in log-linear modeling are easy to implement, because almost all existing statistical software packages have specific functions for that purpose.

Based on the estimated log-linear model described above, the IC joint distribution of $\mathbf{Y}(n)$, denoted as $\{f_{j_1, \dots, j_p}^{(0)} = P(Y_1 = j_1, \dots, Y_p = j_p), j_1, \dots, j_p = 0, 1\}$, can be estimated accordingly. For example, when $p = 3$, $f_{j_1, j_2, j_3}^{(0)}$ can be estimated by $E_{j_1, j_2, j_3} / n_0$, where E_{j_1, j_2, j_3} is the estimated count of the (j_1, j_2, j_3) th cell that can be obtained from the estimated log-linear model and n_0 is the sample size of the IC dataset. Then, we can construct a Phase II control chart as follows. Let

$$g_{j_1, \dots, j_p}(n) = I(Y_1(n) = j_1, \dots, Y_p(n) = j_p), \quad \text{for } j_1, \dots, j_p = 0, 1,$$

$\mathbf{g}(n)$ be a vector of all $g_{j_1, \dots, j_p}(n)$ values for $j_1, \dots, j_p = 0, 1$, and $\mathbf{f}^{(0)}$ be a vector of all $f_{j_1, \dots, j_p}^{(0)}$ values in the corresponding order. By combining the Pearson's χ^2 test statistic with the CUSUM online monitoring

scheme, we can define a CUSUM control chart as follows. First, define

$$\left\{ \begin{array}{ll} \mathbf{S}_n^{\text{obs}} = \mathbf{0}, & \text{if } C_n \leq k_P, \\ \mathbf{S}_n^{\text{exp}} = \mathbf{0}, & \text{if } C_n \leq k_P, \\ \mathbf{S}_n^{\text{obs}} = (\mathbf{S}_{n-1}^{\text{obs}} + \mathbf{g}(n))(C_n - k_P)/C_n, & \text{if } C_n > k_P, \\ \mathbf{S}_n^{\text{exp}} = (\mathbf{S}_{n-1}^{\text{exp}} + \mathbf{f}^{(0)})(C_n - k_P)/C_n, & \text{if } C_n > k_P, \end{array} \right. \quad (5)$$

where $k_P > 0$ is an allowance constant, $\mathbf{S}_0^{\text{obs}} = \mathbf{S}_0^{\text{exp}} = \mathbf{0}$,

$$C_n = \left((\mathbf{S}_{n-1}^{\text{obs}} - \mathbf{S}_{n-1}^{\text{exp}}) + (\mathbf{g}(n) - \mathbf{f}^{(0)}) \right)' \left(\text{diag}(\mathbf{S}_{n-1}^{\text{exp}} + \mathbf{f}^{(0)}) \right)^{-1} \left((\mathbf{S}_{n-1}^{\text{obs}} - \mathbf{S}_{n-1}^{\text{exp}}) + (\mathbf{g}(n) - \mathbf{f}^{(0)}) \right),$$

$\text{diag}(\mathbf{a})$ denotes a diagonal matrix with its diagonal elements equal to the corresponding ones of the vector \mathbf{a} , and the superscripts “obs” and “exp” denote the observed and expected counts, respectively. In Expression (5), C_n is a quantity that measures the overall difference, up to the current time point n , between the observed counts and expected counts of the related p -way contingency table. When $C_n \leq k_P$ (i.e., the data do not show any significant evidence of a distributional shift in $\mathbf{Y}(n)$), we reset $\mathbf{S}_n^{\text{obs}}$ and $\mathbf{S}_n^{\text{exp}}$ to be zero. Otherwise, $\mathbf{S}_n^{\text{obs}}$ is used to record the cumulative observed counts, and $\mathbf{S}_n^{\text{exp}}$ is used to record the cumulative expected counts, in the scale of $(C_n - k_P)/C_n$. Then, the CUSUM charting statistic is defined by

$$u_{n,P} = (\mathbf{S}_n^{\text{obs}} - \mathbf{S}_n^{\text{exp}})' (\text{diag}(\mathbf{S}_n^{\text{exp}}))^{-1} (\mathbf{S}_n^{\text{obs}} - \mathbf{S}_n^{\text{exp}}), \quad (6)$$

and the chart signals a shift when

$$u_{n,P} > h_P, \quad (7)$$

where $h_P > 0$ is a control limit chosen to achieve a given ARL_0 level. The chart (7) is denoted as PCUSUM chart hereafter, since it is derived from the Pearson's χ^2 test. We can check that $u_{n,P}$ in (6) is the conventional Pearson's χ^2 test statistic that measures the difference between the cumulative observed and expected counts

as of the time point n , when $k_p = 0$. Also, it can be checked that $u_{n,p} = \max(0, C_n - k_p)$ when $k_p \neq 0$. Therefore, $u_{n,p}$ is defined in the way that the CUSUM scheme can repeatedly restart when there is little evidence of shifts. By using this CUSUM scheme and the log-linear modeling, we can detect any mean shift in the multivariate count data, without assuming the IC process distribution to follow a parametric form.

4 Numerical Performance Assessment

We evaluate the numerical performance of the related methods that were discussed in the previous sections here. In different numerical examples, we assume that $p = 3$, and the true IC distribution of (X_1, X_2, X_3) belongs to one of the following four cases:

Case I $X_1 \sim \text{Bin}(10, 0.25)$, $X_2 \sim \text{Bin}(10, 0.50)$, $X_3 \sim \text{Bin}(10, 0.75)$, and X_1 , X_2 and X_3 are independent;

Case II $X_1 \sim \text{Poisson}(3)$, $X_2 \sim \text{Poisson}(5)$, $X_3 \sim \text{Poisson}(8)$, and X_1 , X_2 and X_3 are independent;

Case III $X_1 \sim \text{ZIP}(4, 0.20)$, $X_2 \sim \text{ZIP}(6, 0.20)$, $X_3 \sim \text{ZIP}(10, 0.20)$, and X_1 , X_2 and X_3 are independent;

Case IV $X_1 \sim \text{ZIP}(4, 0.20)$, $X_2 \sim \text{ZIP}(6, 0.20)$, $X_3 = X_1 + \delta$, $\delta \sim \text{ZIP}(6, 0.20)$, and X_1 , X_2 and δ are independent.

In Cases III and IV, $\text{ZIP}(\eta, \pi)$ denotes the zero-inflated Poisson distribution with parameters η and π . More specifically, the probability mass function of $X \sim \text{ZIP}(\eta, \pi)$ is defined as

$$P(X = 0) = \pi + (1 - \pi) \exp(-\eta), \quad P(X = x) = (1 - \pi) \eta^x \exp(-\eta) / x!, \quad \text{for } x > 0.$$

In the above four cases, Case II denotes the case when the conventional Poisson distribution with independent components assumption is valid, Cases I and III represent two cases when the Poisson distribution assumption is violated, and Case IV represents a case when some components are correlated. Also, the

binomial, Poisson and zero-inflated Poisson distributions are three representative discrete distributions with under-dispersion, equal-dispersion and over-dispersion, respectively.

We first compare the IC performance of the two parametric charts MEWMA and DCUSUM (cf., Equations (1) and (2)) that were described in Section 2 with some nonparametric charts. Besides the PCUSUM chart discussed in Section 3, we also consider another two representative nonparametric control charts. The first one is the CUSUM version of the multivariate sign chart proposed by Boone and Chakraborti (2012), which is denoted as SNCUSUM. Note that Boone and Chakraborti (2012) originally suggested a Shewhart-type control chart based on the sign statistic and the control limit of the chart was determined based on the asymptotic distribution of the charting statistic. The robustness of the IC performance of this Shewhart chart cannot be guaranteed, particularly in cases with small subgroup sizes. To address this issue and make a fair comparison with the PCUSUM chart, we change the Shewhart chart into a CUSUM chart and use the bootstrap method (Chatterjee and Qiu, 2009) to find its control limit. The second alternative nonparametric CUSUM chart is the one based on the antiranks (Qiu and Hawkins, 2003), denoted as ARCUSUM. Since we usually do not know the shift direction in practice, we choose to use the first and last antiranks in this chart for fair comparisons. The smoothing matrix used in the MEWMA chart is chosen to be the same as that in Chen et al. (2015), the allowance constants in the four CUSUM charts are chosen to be 0.5, the subgroup size of the SNCUSUM chart is fixed at 10, and certain IC parameters used in all five charts are estimated from an IC dataset of size 500. The actual ARL_0 values of the charts are calculated from 10,000 replicated simulations and presented in Table 1. In the table, we would like to point out that the MEWMA chart cannot be used in Case I, because the chart is based on the multivariate Poisson log-normal distributional assumption, which implies that the variance of any component of $\mathbf{X}(n)$ cannot be smaller than its mean (Aitchison and Ho, 1989). But, Case I denotes a case of under-dispersion, as mentioned above. So, that multivariate Poisson log-normal distributional assumption cannot be valid in that case.

From Table 1, we can see that the actual ARL_0 values of the nonparametric charts PCUSUM, ARCUSUM and SNCUSUM are all very close to the nominal ARL_0 values in all cases considered. When

Table 1. The actual ARL_0 values and their standard errors (in parentheses) of the five control charts when the nominal ARL_0 values are fixed at 200 or 500.

Nominal ARL_0	Chart	Case I	Case II	Case III	Case IV
200	MEWMA	-	198.1 (2.05)	164.7 (1.83)	218.6 (2.01)
	DCUSUM	249.9 (2.87)	211.7 (2.54)	250.9 (2.91)	279.6 (3.23)
	SNCUSUM	200.0 (2.28)	200.0 (2.25)	200.0 (2.32)	200.0 (2.34)
	ARCUSUM	200.1 (3.25)	199.9 (3.31)	200.1 (2.93)	199.9 (2.99)
	PCUSUM	199.9 (2.87)	199.9 (2.93)	199.9 (2.81)	200.0 (2.72)
500	MEWMA	-	473.1 (3.67)	429.0 (3.25)	559.7 (3.31)
	DCUSUM	675.9 (7.30)	514 (5.60)	630.3 (6.53)	761.1 (7.73)
	SNCUSUM	500.0 (5.05)	500.0 (5.51)	500.0 (5.11)	499.9 (5.06)
	ARCUSUM	499.9 (6.43)	500.0 (6.50)	499.9 (5.92)	500.0 (6.03)
	PCUSUM	499.3 (5.80)	500.0 (5.69)	499.9 (5.75)	500.0 (5.65)

further comparing the three nonparametric control charts in terms of the standard error, the SNCUSUM chart has the best performance and the PCUSUM chart performs better than the ARCUSUM chart. As a comparison, the actual ARL_0 values of the parametric charts MEWMA and DCUSUM deviate significantly from the nominal ARL_0 values in most cases considered. If the actual ARL_0 value of a chart is substantially larger than the nominal value, then the chart would be too conservative in detecting process distributional shifts in the sense that a real shift would not be detected as quickly as expected. Thus, we could have the situation when the process produces many defective products without notice. In the case when the actual ARL_0 value of a chart is substantially smaller than the nominal value, then the chart would give many false signals. Thus, the production process would be unnecessarily stopped too often and too soon. Much human resource and production efficiency would be wasted as a consequence. Therefore, the parametric charts should be used with care in practice, and their distributional assumptions should be adequately checked in advance.

Next, we compare the OC performance of the related control charts. From Table 1, we can see that the MEWMA and DCUSUM charts have unreliable IC performance in various cases considered. In such cases, their shift detection power might be irrelevant because a good shift detection power could be due to an overly small actual ARL_0 value. To make a relatively fair comparison about their OC performance,

we use the bootstrap procedure to compute the control limits of the MEWMA and DCUSUM charts so that their actual ARL_0 values reach the nominal ARL_0 value, and the related control charts are denoted as MEWMA(b) and DCUSUM(b), respectively. Thus, the MEWMA(b) and DCUSUM(b) charts can also be regarded as nonparametric control charts because computation of their control limits does not depend on the assumed parametric distributions. The bootstrap sample size used in these two charts is chosen to be 500 that is the same as that of the IC dataset. For illustration purposes, the original versions of the MEWMA and DCUSUM charts are also considered here. In this example, the true process distribution is assumed to be the standardized version with mean 0 and standard deviation 1 of one of the Cases I-IV. In each case, two shift scenarios are considered: one is that a shift occurs only in the first component with the shift size changing from -1.0 to 1.0, and the other is that a shift occurs in all three components with the same size changing from -0.4 to 0.4. “delta” is used to denote the magnitude of the considered shift in the following figures. Because performance of different CUSUM charts with a same allowance constant may not be comparable (cf., [Qiu, 2008](#)), we choose to compare their optimal performance when detecting a specific shift. Namely, for detecting a given shift, we search the allowance constant of a CUSUM chart such that the ARL_1 value reaches the minimum when its ARL_0 value is fixed at the nominal level. For the MEWMA chart, we still use the same smoothing parameters as those in [Chen et al. \(2015\)](#), because it was shown in that paper that these smoothing parameters were robust and could give optimal or close to optimal performance across different settings. Also, that chart is not considered in Case I for the reason given above about Table 1.

Based on 10,000 replications, the calculated ARL_1 values of the PCUSUM, DCUSUM, MEWMA, DCUSUM(b) and MEWMA(b) charts are shown in Figures 1-2, respectively, in the two shift scenarios. These two figures are used for demonstrating the unreliability of the parametric charts DCUSUM and MEWMA and for comparing the PCUSUM chart with the modified parametric charts DCUSUM(b) and MEWMA(b). As shown in Figures 1-2, when the true process distribution is Poisson (i.e., plot (b) in both Figures), DCUSUM and MEWMA perform similarly to DCUSUM(b) and MEWMA(b), respectively. This is expected because the Poisson distribution assumption is valid in such cases. In the other cases when

the true process distribution is binomial or zero-inflated Poisson, DCUSUM and MEWMA are unreliable. Regarding DCUSUM(b) and MEWMA(b), they have a similar performance to PCUSUM. The PCUSUM chart seems to be more sensitive to small shifts while the DCUSUM(b) and MEWMA(b) charts perform better for detecting larger shifts.

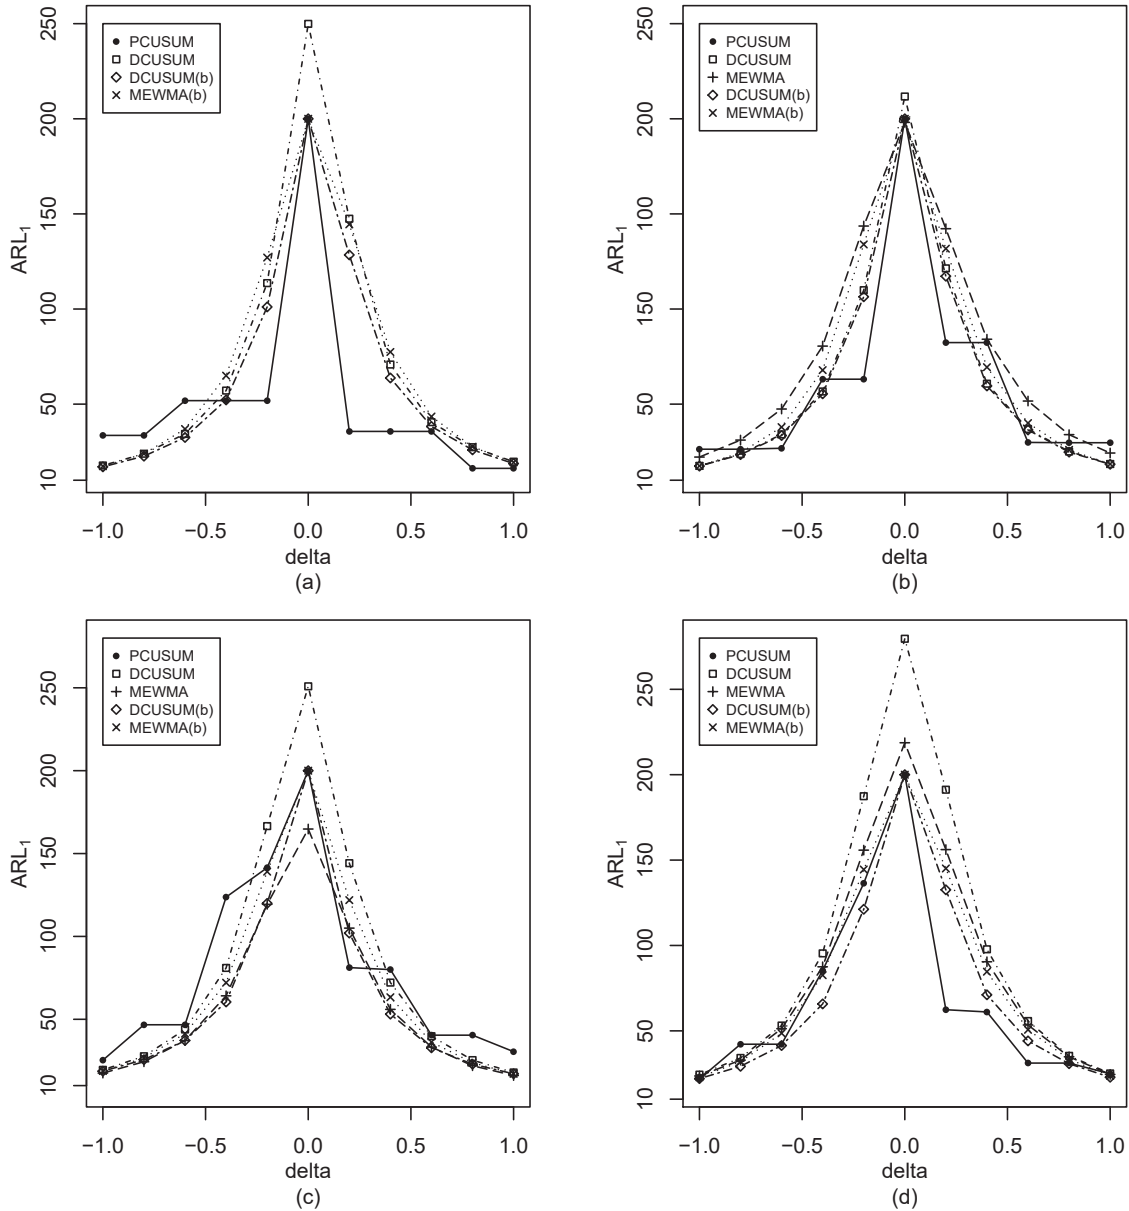


Figure 1. Calculated ARL_1 values of the control charts when the nominal ARL_0 is fixed at 200, and the actual IC process distribution is the standardized version of the one in Case I (plot (a)), Case II (plot (b)), Case III (plot (c)), and Case IV (plot (d)). Shift of size “delta” occurs only in the first component of $X(n)$.

Figures 3-4 show the calculated ARL_1 values of the three nonparametric control charts PCUSUM,

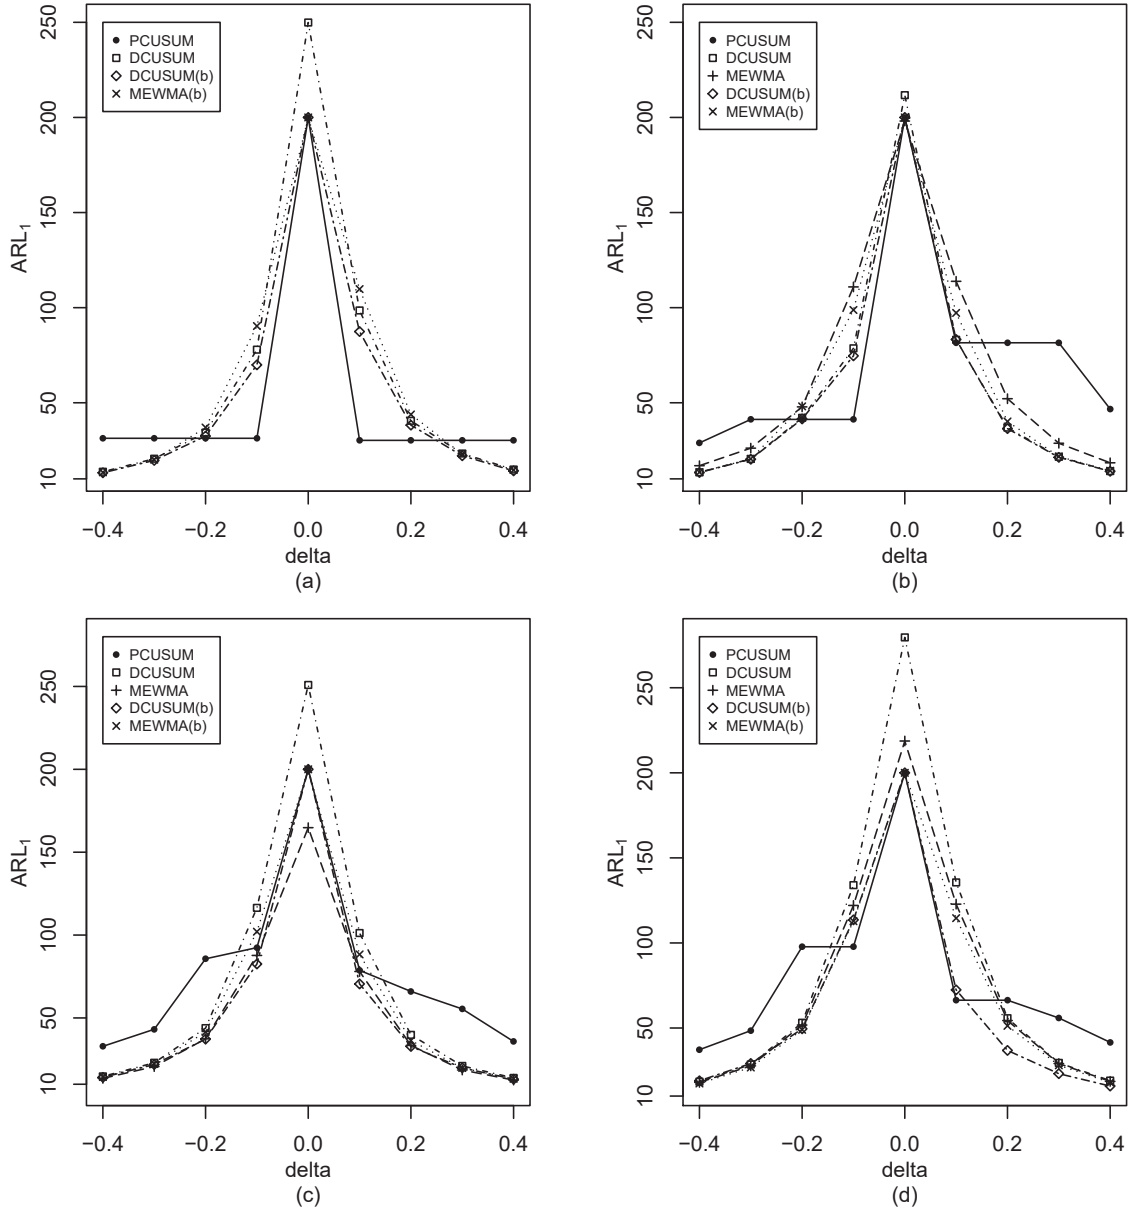


Figure 2. Calculated ARL_1 values of the control charts when the nominal ARL_0 is fixed at 200, and the actual IC process distribution is the standardized version of the one in Case I (plot (a)), Case II (plot (b)), Case III (plot (c)), and Case IV (plot (d)). Shift of size “delta” occurs in all components of $\mathbf{X}(n)$.

ARCUSUM and SNCUSUM in the two shift scenarios, respectively, based on 10,000 replicated simulations.

From the figures, we can see that PCUCUM outperforms ARCUSUM and SNCUSUM in most cases, and ARCUSUM outperforms SNCUSUM in most cases as well.

Next, we study the standard deviation of the OC run length, denoted as $SDRL_1$, of various control charts. In cases when the shift occurs in the first component of $\mathbf{X}(n)$ (i.e., the first shift scenario) with the

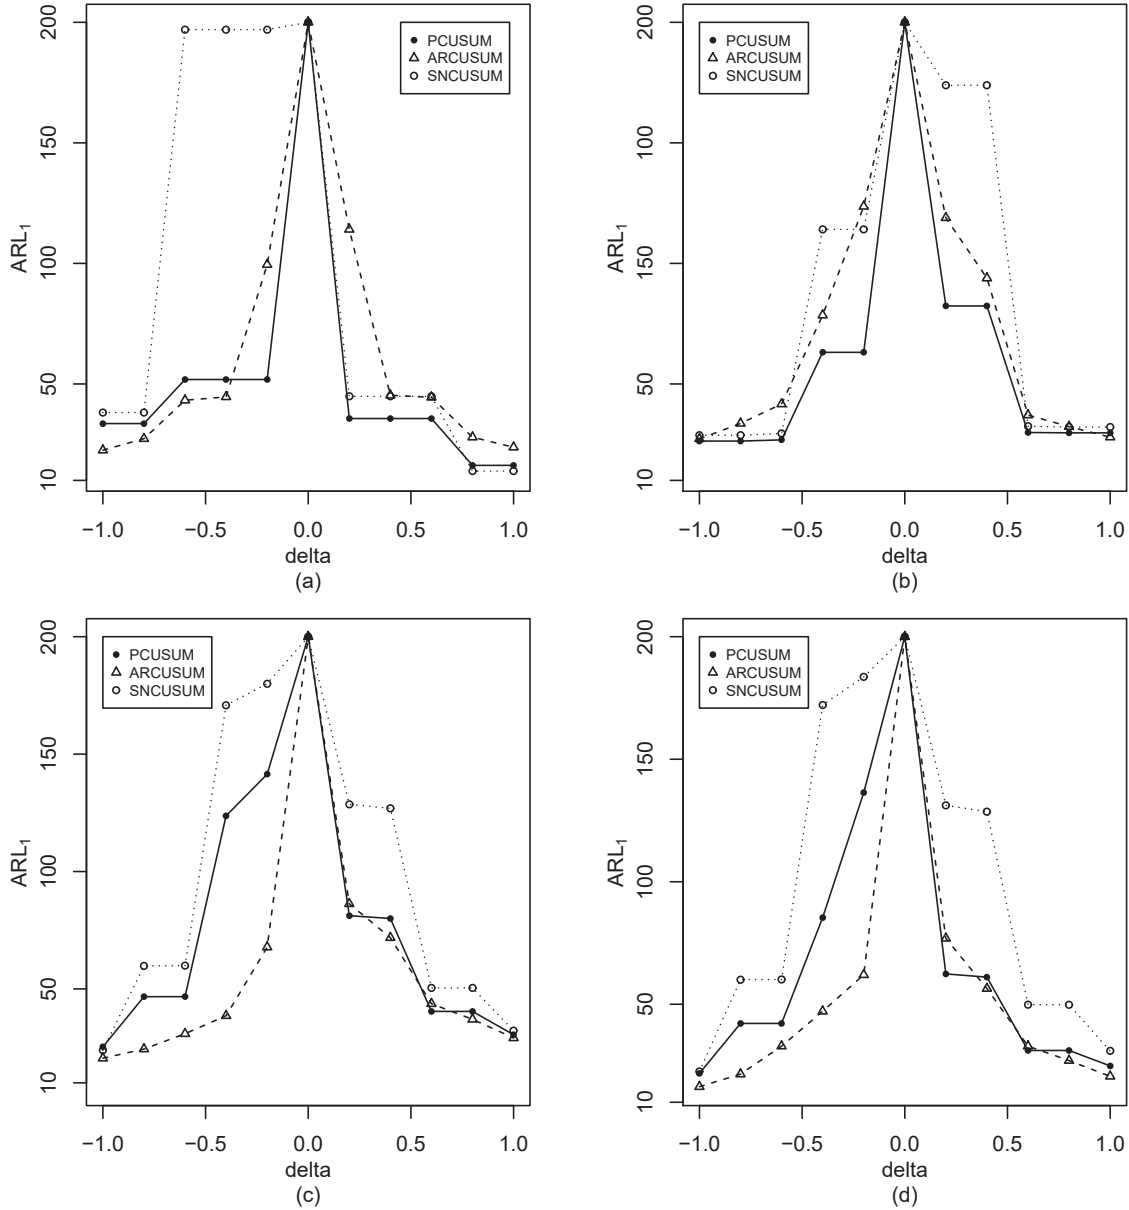


Figure 3. Calculated ARL_1 values of the control charts when the nominal ARL_0 is fixed at 200, and the actual IC process distribution is the standardized version of the one in Case I (plot (a)), Case II (plot (b)), Case III (plot (c)), and Case IV (plot (d)). Shift of size “delta” occurs only in the first component of $\mathbf{X}(n)$.

size “delta” changes from 0.2 to 1.0 and other setupa are the same as those in Figures 1-2, the calculated $SDRL_1$ values of the charts PCUSUM, ARCUSUM, SNCUSUM, DCUSUM(b) and MEWMA(b) are shown in Figure 5. From the plots in the figure, it can be seen that PCUSUM has relatively small $SDRL_1$ values, compared to the other four charts, in most cases considered. For downward shifts in the first shift scenario and for shifts in the second shift scenario, results are similar and thus omitted here.

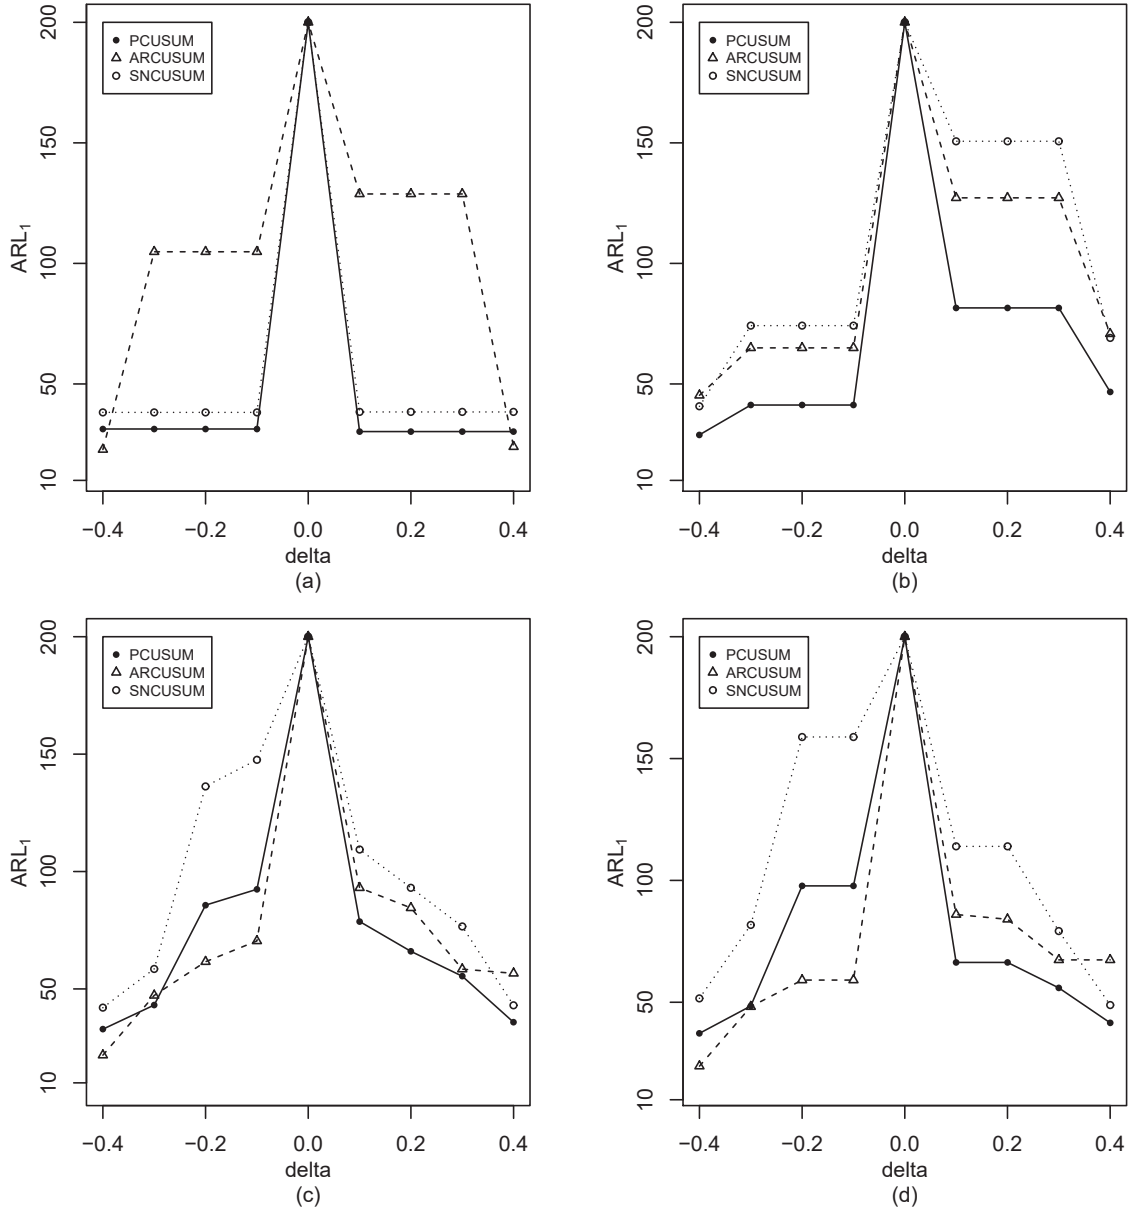


Figure 4. Calculated ARL_1 values of the control charts when the nominal ARL_0 is fixed at 200, and the actual IC process distribution is the standardized version of the one in Case I (plot (a)), Case II (plot (b)), Case III (plot (c)), and Case IV (plot (d)). Shift of size “delta” occurs in all components of $\mathbf{X}(n)$.

As a summary of the above simulation results, we can have the following conclusions. (i) The parametric charts DCUSUM and MEWMA are unreliable to use in cases when their distribution assumptions are violated. (ii) Among the three nonparametric charts ARCUSUM, SNCUSUM and PCUSUM, in terms of ARL_1 , PCUSUM outperforms SNCUSUM in all cases considered, and it outperforms ARCUSUM when detecting most upward shifts in all four cases and when detecting downward shifts in Cases I and II. In terms

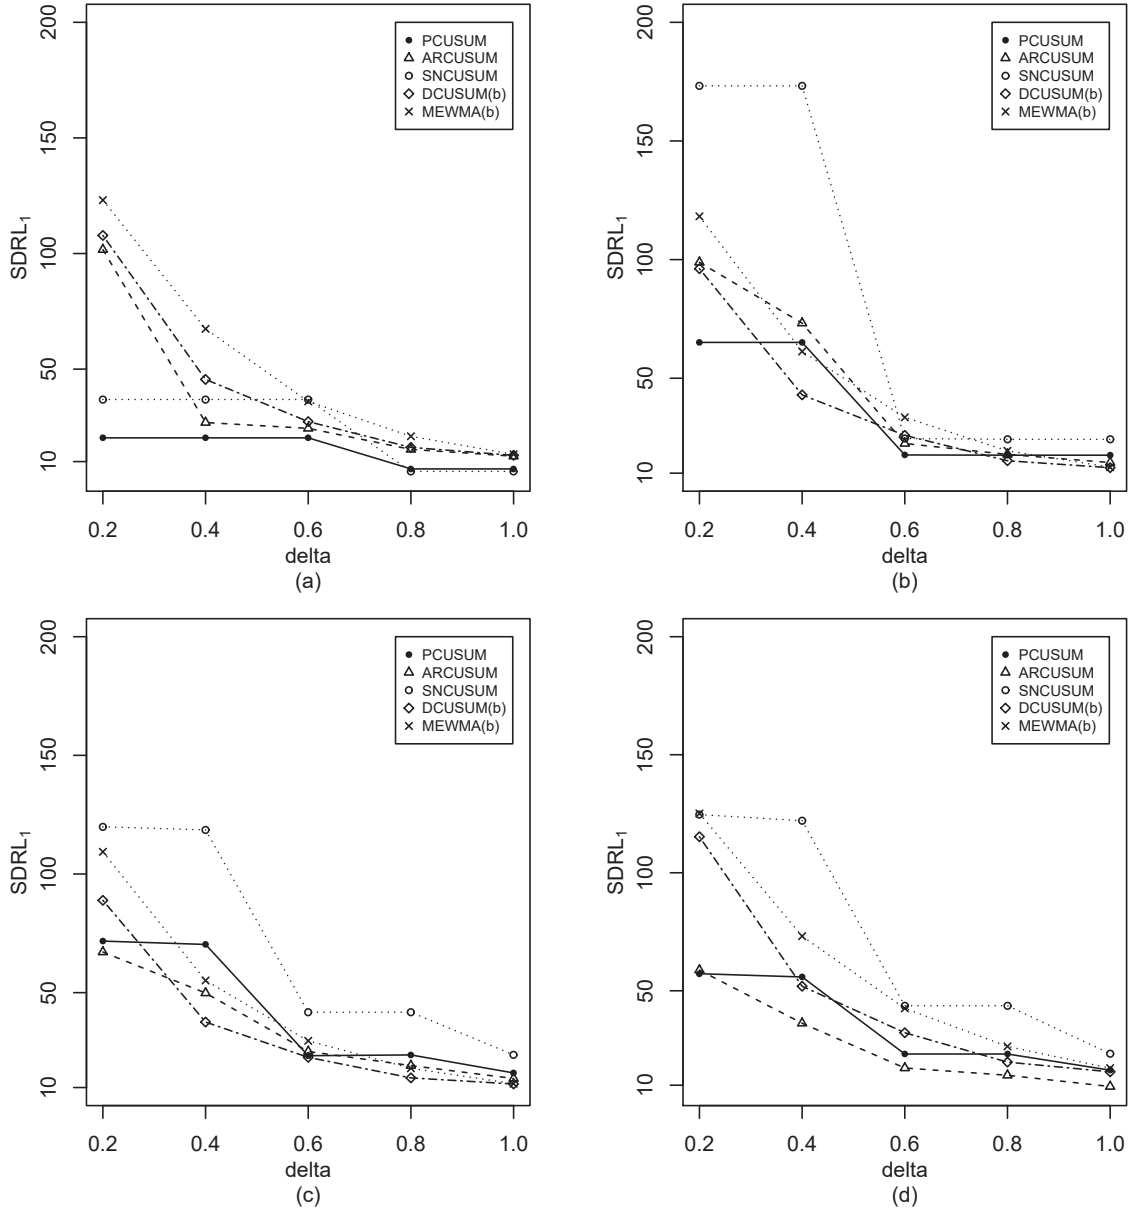


Figure 5. The $SDRL_1$ values of the five nonparametric control charts when the nominal ARL_0 is fixed at 200, and the actual IC process distribution is the standardized version of the one in Case I (plot (a)), Case II (plot (b)), Case III (plot (c)), and Case IV (plot (d)). Shift of size “delta” occurs only in the first component of $\mathbf{X}(n)$.

of $SDRL_1$, similar conclusions can be made. Therefore, the PCUSUM chart is preferred in general among the three nonparametric control charts considered here for monitoring multivariate count data.

As mentioned earlier, certain IC parameters (e.g., the medians m_j) of the PCUSUM chart need to be estimated from an IC dataset. Thus, its performance depends on the size n_0 of the IC dataset. To study the

impact of n_0 on the performance of the PCUSUM chart, let us consider Case I and compute the optimal ARL_1 values of the PCUSUM chart when n_0 changes from 100, 500, 1000 to 5000 and all other parameters are chosen to be the same as those in the example of Figure 1. The results are presented in Figure 6, where the y-axis is in natural logarithm scale to better distinguish the difference. From the plots, we can see that the results are better when n_0 is chosen larger, as expected, and the results do not change much when $n_0 \geq 1000$.

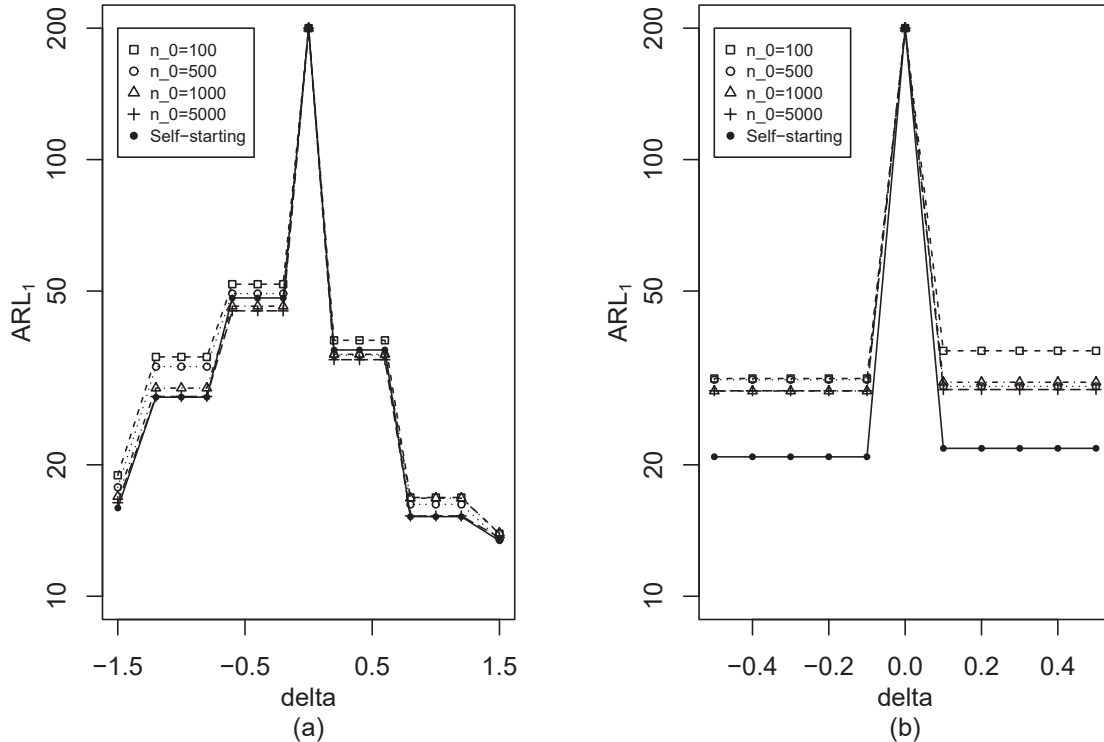


Figure 6. Calculated optimal ARL_1 values of the PCUSUM chart when $n_0 = 100, 500, 1000$ and 5000 , and of the self-start version of the PCUSUM chart in Case I. Shift of size “delta” occurs only in the first component of $\mathbf{X}(n)$ in plot (a), and occurs in all components of $\mathbf{X}(n)$ in plot (b).

In certain applications, it might be difficult to have 1000 or even 500 IC observations before online monitoring. In such cases, a natural approach to overcome that difficulty is to use a self-starting control chart (Hawkins, 1987; Capizzi and Masarotto, 2010). The basic idea behind that chart is that the new observation collected at the current time point can be combined with the IC dataset once we confirm that the related process is IC at the current time point. So, the size of the IC dataset can potentially increase during online process monitoring and the IC parameters can potentially be estimated more and more accurately. To use a self-starting chart, a small number of IC observations is still needed in advance, to

calculate initial estimates of the IC parameters. Several existing studies, including those in [Hawkins \(1987\)](#), [Hawkins and Maboudou-Tchao \(2007\)](#) and [Sullivan and Jones \(2002\)](#), showed that the results would be quite stable if we could collect 20-50 IC observations in advance, depending on the dimensionality of the process observations. It was shown that a dozen initial IC observations were usually good enough in univariate cases and we needed 30-50 initial IC observations when the dimensionality of the process observations was up to 15. We constructed the self-starting version of the PCUSUM chart, and performed a big simulation study about its numerical performance. Our results are overall consistent with those in the existing studies mentioned above. In the setup of the previous example, its optimal ARL_1 values are shown in [Figure 6](#) by the solid lines, in cases when 30 initial IC observations are used. We can see that its performance is about the same as that of the original PCUSUM chart with 5,000 IC observations in the scenario of plot (a), and even much better than the latter in the scenario of plot (b).

5 A Real-Data Application

In this section, we apply the related control charts to a real-data example about the crime logs in 2016 at the University of Florida. Three types of common crimes are considered in this example, including Driving Under the Influence (X_1), Narcotics Violation (X_2) and Larceny/Theft (X_3). Daily counts of these crimes can be found at the website of the University of Florida Police Department (<http://www.police.ufl.edu/wp-content/asp/crimelog/default2.asp>). We then use the first half of the data (more specifically, the first 185 observations) as the IC data, and the remaining as the test data for online monitoring. Both the IC data and the test data are shown in the 3-D plot in [Figure 7](#). We can see that the test data have a larger mean than the IC data. The IC data are also shown in the first row of [Figure 8](#). The corresponding histograms of the three variables are shown in the second row, along with their density curves (solid lines) and the density curves of the Poisson distributions with the same means (dashed lines). From [Figure 8\(d\)-\(f\)](#), we can see that the marginal distributions of X_1 , X_2 and X_3 are quite different from the corresponding Poisson distributions for

the IC data. We then use the Pearson's Chi-square goodness-of-fit test and the Fisher's index of dispersion test (Fisher et al., 1922) to formally test whether each variable in the IC data follows a Poisson distribution. Both tests conclude that X_2 and X_3 are significantly different from Poisson (Pearson's test: p-value=0.000 for X_2 and p-value=0.001 for X_3 ; Fisher's test: p-value=0.000 for X_2 and p-value=0.002 for X_3). As for X_1 , the Pearson's test gives the p-value of 0.077, while the Fisher's test gives the p-value of 0.000. Therefore, we can conclude that the joint distribution of $\mathbf{X}(n) = (X_1(n), X_2(n), X_3(n))'$ cannot be multivariate Poisson, because a joint Poisson distribution implies that all marginal distributions are Poisson.

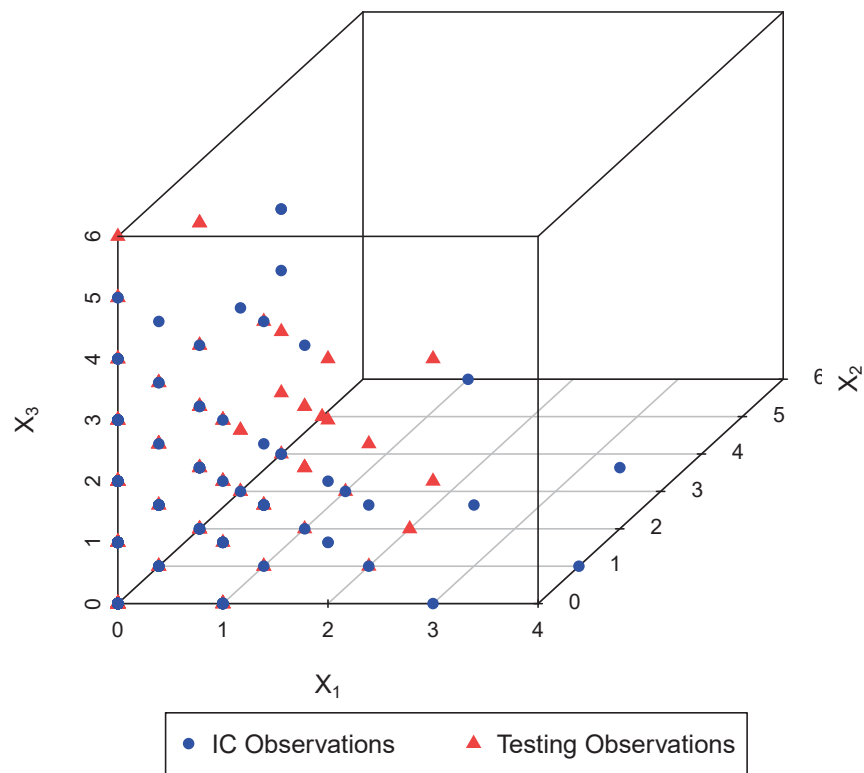


Figure 7. A 3-D scatter plot of the daily counts of the three types of crimes in year 2016 at the University of Florida.

Because it is confirmed above that the IC data do not follow a multivariate Poisson distribution, only the five nonparametric charts PCUSUM, ARCUSUM, SNCUSUM, DCUSUM(b) and MEWMA(b) are considered in this example. When implementing the related control charts, the nominal ARL_0 is fixed at 200 for all five charts, the allowance constant is chosen to be 0.1 for the PCUSUM, ARCUSUM, SNCUSUM and DCUSUM(b) charts, the subgroup size in the SNCUSUM chart is fixed at 10, and the smoothing parameters

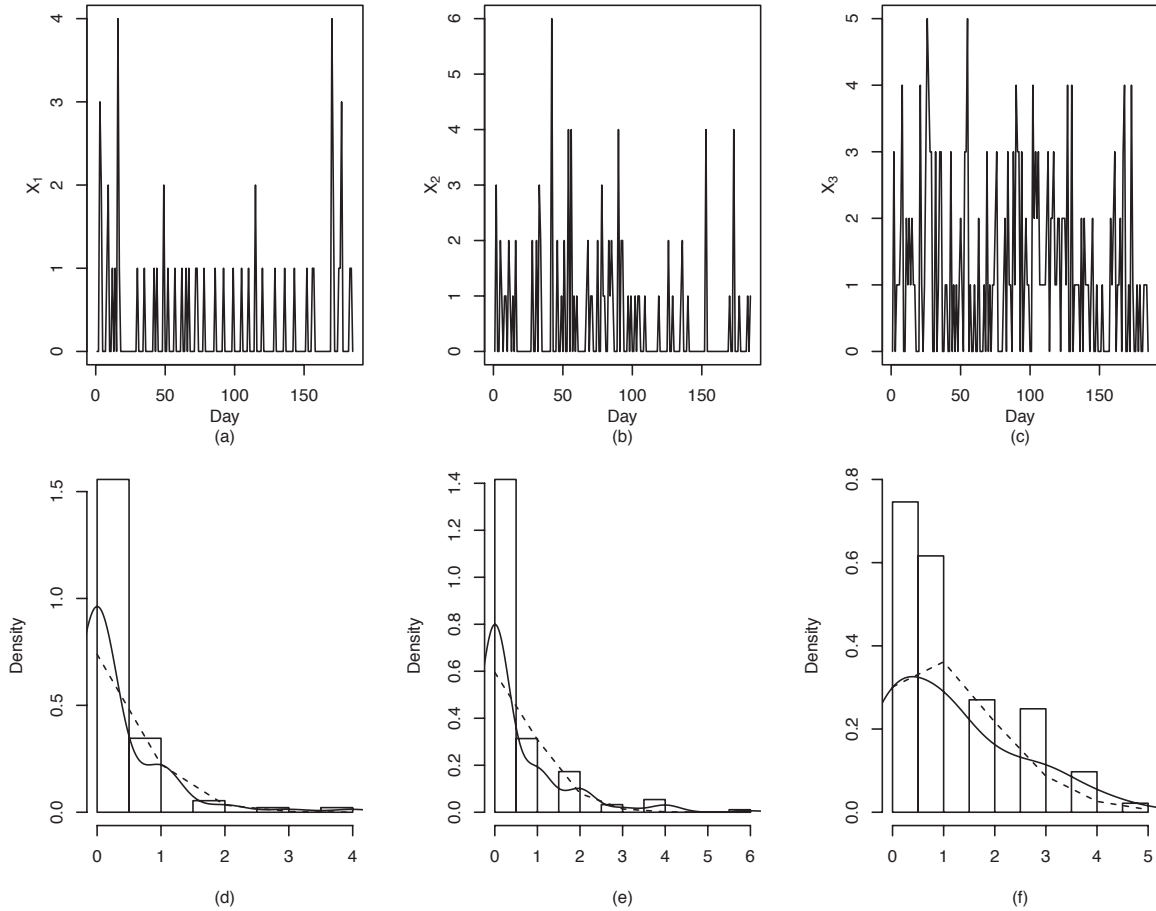


Figure 8. Daily counts ((a)-(c)) and histograms ((d)-(f)) of the three types of crimes in the first 185 days in year 2016 at the University of Florida. The dashed lines in plots (d)-(f) denote the density curves of the data, and the solid lines denote the density curves of the Poisson distributions with the same means.

of the MEWMA(b) chart are chosen to be the same as those used in [Chen et al. \(2015\)](#). The five control charts are presented in Figure 9, where the dashed horizontal lines denote their control limits. From the plots, we can see that the PCUSUM, ARCUSUM, SNCUSUM, DCUSUM(b) and MEWMA(b) charts give signals at the 2nd, 2nd, 180th, 82nd and 70th testing observations, respectively. Therefore, the PCUSUM and ARCUSUM charts detect the distributional shifts earlier than the remaining 3 charts in this example.

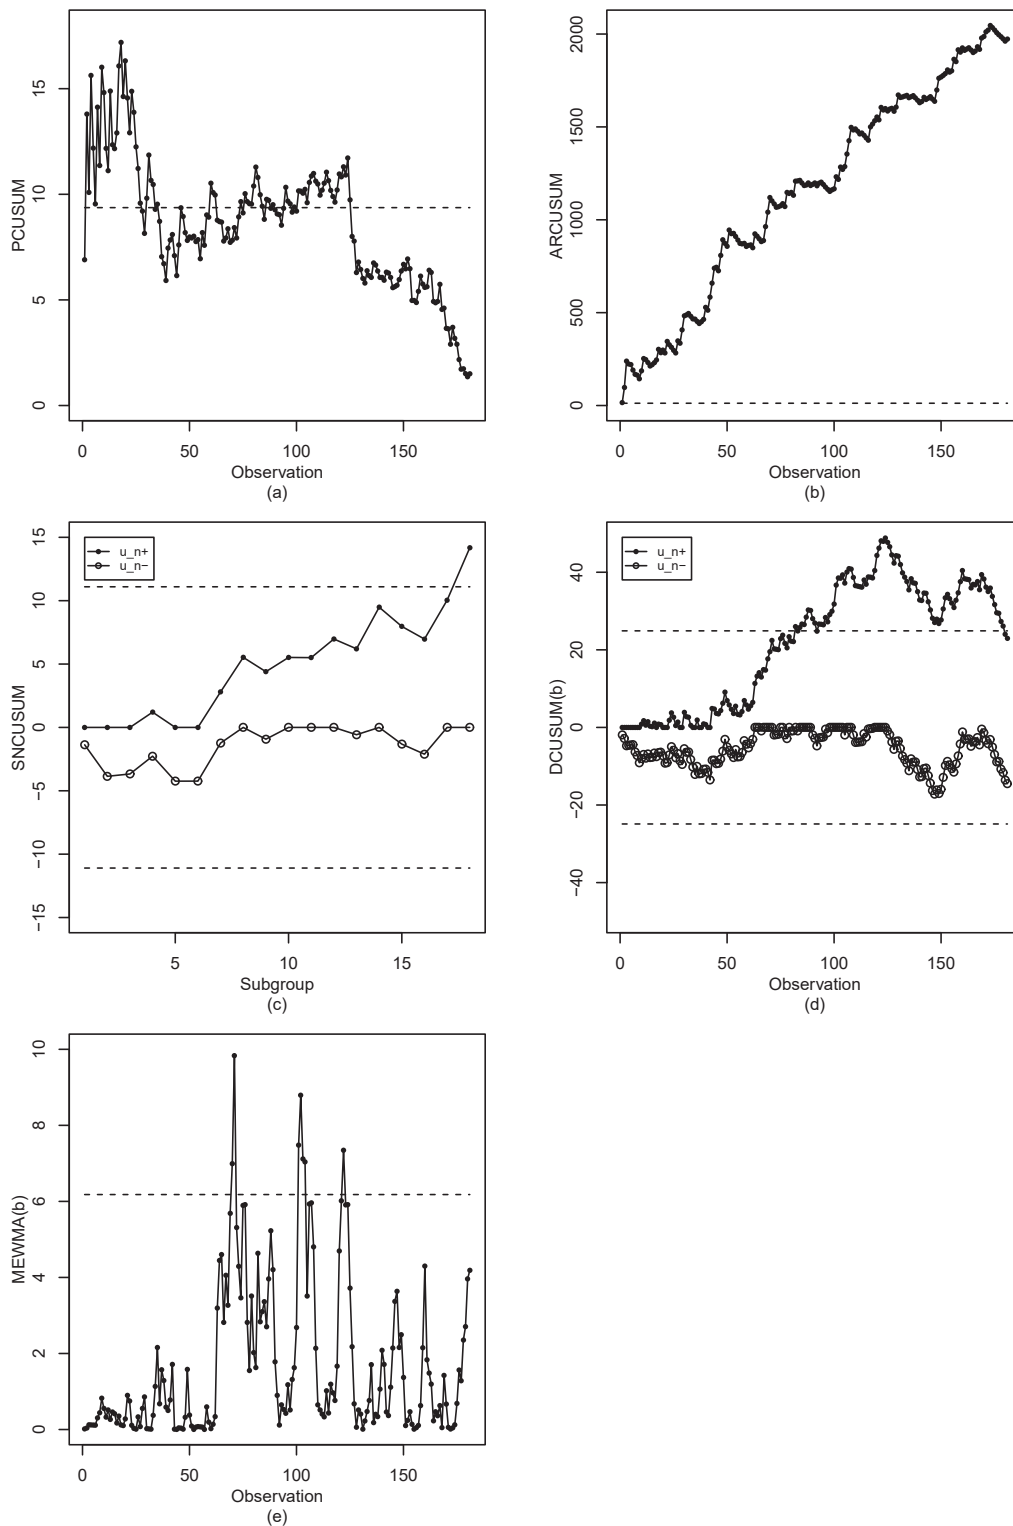


Figure 9. The PCUSUM, ARCUSUM, SNCUSUM, DCUSUM(b) and MEWMA(b) charts for monitoring the daily counts of the three types of crimes at the University of Florida during the second half of the year in 2016. The dashed horizontal lines are the control limits of the related control charts.

6 Concluding Remarks

SPC for multiple count data remains a challenging problem. Most existing methods for monitoring multiple count data are based on certain parametric models, among which the most commonly used one is the multivariate Poisson distribution. In practice, however, the assumed parametric models are rarely valid, due to the complicated impact of various factors (e.g., the environment, weather, and so forth) on the count data. We have shown in the paper that the parametric control charts are unreliable to use in such cases because their actual ARL_0 values could be substantially different from a nominal level. In this paper, we carefully compare the performance of some representative parametric and nonparametric control charts in monitoring multiple count data. We suggest that the related parametric model should be checked carefully before a parametric chart is used. In cases when we are unsure whether the assumed parametric model is valid, or when we do not know which parametric chart is appropriate, we suggest using a nonparametric chart instead. To this end, the PCUSUM chart and its self-starting version can give reasonably good results in general, based on our intensive numerical study in various different cases.

There are a number of issues that have not been discussed thoroughly in the current paper. For instance, we focus on mean shifts only in this paper. Although we believe that the PCUSUM chart can also detect shifts in the scale parameters, we do not know how effective it is for that purpose. Also, the allowance constant k_p needs to be determined in advance and it might be selected adaptively, as discussed in Section 4.5 of [Qiu \(2014\)](#). These issues and some others will be addressed in our future research.

Acknowledgments

The authors thank the editors and referees for many helpful comments and suggestions, which improved the quality of the paper greatly. This research is supported in part by the National Science Foundation grant DMS-1405698 in US and by the Natural Sciences Foundation of China grants 71532008 and 7161147003.

Biographical Sketches

Peihua Qiu received his PhD in statistics from the Department of Statistics at the University of Wisconsin at Madison in 1996. He worked as a senior research consulting statistician of the Biostatistics Center at the Ohio State University during 1996-1998. Then, he worked as an assistant professor (1998-2002), an associate professor (2002-2007), and a full professor (2007-2013) of the School of Statistics at the University of Minnesota. He is an elected fellow of the American Statistical Association, an elected fellow of the Institute of Mathematical Statistics, an elected member of the International Statistical Institute, a senior member of the American Society for Quality, and a lifetime member of the International Chinese Statistical Association. He served as associate editor for a number of top journals in statistics, including Journal of the American Statistical Association, Biometrics, and Technometrics. He was the editor of Technometrics during 2014-2016 and has been the founding chair of the Department of Biostatistics at the University of Florida since July 1, 2013. Peihua Qiu has made substantial contributions in the areas of jump regression analysis, image processing, statistical process control, survival analysis, and reliability. So far, he has published over 100 research papers in top refereed journals, including Technometrics, Journal of the American Statistical Association, Annals of Statistics, Annals of Applied Statistics, Journal of the Royal Statistical Society (Series B), Biometrika, Biometrics, IEEE Transactions on Pattern Analysis and Machine Intelligence, and IIE Transactions. His research monograph titled Image Processing and Jump Regression Analysis (2005, Wiley) won the inaugural Ziegel prize in 2007, for its contribution in bridging the gap between jump regression analysis in statistics and image processing in computer science. His second book titled Introduction to Statistical Process Control was published in 2014 by Chapman & Hall/CRC.

Zhen He is a professor in the College of Management and Economics, Tianjin University. He received his PhD in Management Science and Engineering from Tianjin University, China, in 2001. He is the recipient of the Outstanding Research Young Scholar Award of the National Natural Science Foundation of China. His research interests include statistical quality control, design of experiment, and Six Sigma management.

Zhiqiong Wang received his PhD degree in 2018 from the College of Management and Economics of the Tianjin University in China. He is currently an assistant professor of the School of Management at the Tianjin University of Technology. His major research interests include quality control and management, change-point detection, and various quality-related applications. This paper was written during his one-year visit at the Department of Biostatistics of the University of Florida.

References

- Agresti, A. (2013). *Categorical data analysis*. John Wiley & Sons,.
- Aitchison, J. and Ho, C. (1989). The multivariate Poisson-log normal distribution. *Biometrika*, 76(4):643–653.
- Boone, J. and Chakraborti, S. (2012). Two simple Shewhart-type multivariate nonparametric control charts. *Applied Stochastic Models in Business and Industry*, 28(2):130–140.
- Brijs, T., Karlis, D., Swinnen, G., Vanhoof, K., Wets, G., and Manchanda, P. (2004). A multivariate Poisson mixture model for marketing applications. *Statistica Neerlandica*, 58(3):322–348.
- Capizzi, G. and Masarotto, G. (2010). Self-starting CUSCORE control charts for individual multivariate observations. *Journal of Quality Technology*, 42(2):136—151.
- Chatterjee, S. and Qiu, P. (2009). Distribution-free cumulative sum control charts using bootstrap-based control limits. *The Annals of Applied Statistics*, 3(1):349–369.
- Chen, N., Li, Z., and Ou, Y. (2015). Multivariate exponentially weighted moving-average chart for monitoring Poisson observations. *Journal of Quality Technology*, 47(3):252—263.
- Chen, N., Zi, X., and Zou, C. (2016). A distribution-free multivariate control chart. *Technometrics*, 58(4):448–459.

- Chiu, J. E. and Kuo, T. I. (2008). Attribute control chart for multivariate Poisson distribution. *Communications in Statistics-Theory and Methods*, 37(1):146–158.
- Chiu, J. E. and Kuo, T. I. (2010). Control charts for fraction nonconforming in a bivariate binomial process. *Journal of Applied Statistics*, 37(10):1717–1728.
- Cozzucoli, P. C. and Marozzi, M. (2017). Monitoring multivariate Poisson processes: a review and some new results. *Quality Technology & Quantitative Management*.
- Das, D., Zhou, S., Chen, Y., and Horst, J. (2016). Statistical monitoring of over-dispersed multivariate count data using approximate likelihood ratio tests. *International Journal of Production Research*, 54(21):6579–6593.
- Fisher, R., Thornton, H., and Mackenzie, W. (1922). The accuracy of the plating method of estimating the density of bacterial populations. *Annals of Applied Biology*, 9(3-4):325–359.
- Hawkins, D. M. (1987). Self-starting cusum charts for location and scale. *The Statistician*, pages 299–316.
- Hawkins, D. M., Choi, S., and Lee, S. (2007). A general multivariate exponentially weighted moving-average control chart. *Journal of Quality Technology*, 39(2):118—125.
- Hawkins, D. M. and Maboudou-Tchao, E. M. (2007). Self-starting multivariate exponentially weighted moving average control charting. *Technometrics*, 49(2):199–209.
- He, S., He, Z., and Wang, G. A. (2014). CUSUM control charts for multivariate Poisson distribution. *Communications in Statistics-Theory and Methods*, 43(6):1192–1208.
- Holgate, P. (1964). Estimation for the bivariate Poisson distribution. *Biometrika*, 51(1-2):241–287.
- Johnson, N. L., Kotz, S., and Balakrishnan, N. (1997). *Discrete multivariate distributions*. Wiley New York.
- Jones, L. A., Woodall, W. H., and Conerly, M. D. (1999). Exact properties of demerit control charts. *Journal of Quality Technology*, 31(2):207—215.

- Karlis, D. (2003). An EM algorithm for multivariate Poisson distribution and related models. *Journal of Applied Statistics*, 30(1):63–77.
- Karlis, D. and Meligkotsidou, L. (2005). Multivariate Poisson regression with covariance structure. *Statistics and Computing*, 15(4):255–265.
- Laungrungrong, B., Borrór, C. M., and Montgomery, D. C. (2011). EWMA control charts for multivariate Poisson-distributed data. *International Journal of Quality Engineering and Technology*, 2(3):185–211.
- Lee Ho, L. and Branco Costa, A. F. (2009). Control charts for individual observations of a bivariate Poisson process. *The International Journal of Advanced Manufacturing Technology*, 43(7):744–755.
- Li, J., Tsung, F., and Zou, C. (2014). Multivariate binomial/multinomial control chart. *IIE Transactions*, 46(5):526–542.
- Lu, X. S., Xie, M., Goh, T. N., and Lai, C. D. (1998). Control chart for multivariate attribute processes. *International Journal of Production Research*, 36(12):3477–3489.
- Patel, H. (1973). Quality control methods for multivariate binomial and Poisson distributions. *Technometrics*, 15(1):103–112.
- Qiu, P. (2008). Distribution-free multivariate process control based on log-linear modeling. *IIE Transactions*, 40(7):664–677.
- Qiu, P. (2014). *Introduction to statistical process control*. Chapman & Hall/CRC: Boca Raton, FL.
- Qiu, P. (2018). Some perspectives on nonparametric statistical process control. *Journal of Quality Technology*, 50(1):49–65.
- Qiu, P. and Hawkins, D. (2001). A rank-based multivariate CUSUM procedure. *Technometrics*, 43(2):120–132.

- Qiu, P. and Hawkins, D. (2003). A nonparametric multivariate cumulative sum procedure for detecting shifts in all directions. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 52(2):151–164.
- Qiu, P. and Li, Z. (2011). On nonparametric statistical process control of univariate processes. *Technometrics*, 53(4):390–405.
- Saghir, A. and Lin, Z. (2014). Control chart for monitoring multivariate COM-Poisson attributes. *Journal of Applied Statistics*, 41(1):200–214.
- Saghir, A. and Lin, Z. (2015). Control charts for dispersed count data: an overview. *Quality and Reliability Engineering International*, 31(5):725–739.
- Schmidt, A. M. and Rodriguez, M. A. (2011). Modelling multivariate counts varying continuously in space. *Bayesian Statistics*, 9:611–638.
- Skinner, K. R., Montgomery, D. C., and Runger, G. C. (2003). Process monitoring for multiple count data using generalized linear model-based control charts. *International Journal of Production Research*, 41(6):1167–1180.
- Sullivan, J. H. and Jones, L. A. (2002). A self-starting control chart for multivariate individual observations. *Technometrics*, 44(1):24–33.
- Topalidou, E. and Psarakis, S. (2009). Review of multinomial and multiattribute quality control charts. *Quality and Reliability Engineering International*, 25(7):773–804.
- Wang, Z., Li, Y., and Zhou, X. (2017). A statistical control chart for monitoring high-dimensional Poisson data streams. *Quality and Reliability Engineering International*, 33:307—321.
- Zou, C. and Tsung, F. (2011). A multivariate sign EWMA control chart. *Technometrics*, 53(1):84–97.
- Zou, C., Wang, Z., and Tsung, F. (2012). A spatial rank-based multivariate EWMA control chart. *Naval Research Logistics*, 59(2):91–110.