Int. J. Anal. Appl. (2026), 24:145

International Journal of Analysis and Applications

Bootstrap EWMA and CUSUM Charts for Monitoring Mean Shifts in Poisson INAR Processes

Wandee Wanishsakpong, Jeeraporn Thaithanan, Sawaporn Hinsheranan*

Department of Statistics, Faculty of Science, Kasetsart University, Bangkok, Thailand

**Corresponding author:* sawaporn.h@ku.th

Abstract. Autocorrelated count data commonly arise in applications where classical control charts assuming independence may not be appropriate. This paper develops bootstrap-calibrated one-sided EWMA and CUSUM charts for detecting upward mean shifts in Poisson INAR(p) processes. Two bootstrap procedures are considered: a Discrete (D) approach that refits the model in each replication to account for parameter estimation variability, and a Model-Based (MB) approach that conditions on a single fitted in-control model. Detection performance is evaluated using expected detection delay and detection reliability. Simulation results show that EWMA charts generally provide faster detection, whereas CUSUM charts may achieve slightly higher detection reliability. As the shift magnitude increases, performance differences diminish. The D approach performs better for small shifts, while the MB approach provides a computationally efficient alternative with comparable performance for moderate and large shifts. A real-data example using road-traffic injury counts illustrates the practical relevance of the method.

1. Introduction

Statistical process control (SPC) techniques are commonly used to monitor a change in industrial, environmental, and public health processes. When the monitored characteristic is a count, such as the number of system failures, disease cases, or traffic accidents, Shewhart c and u charts are commonly employed. These classical methods rely on the assumption that observations are independent and identically distributed (i.i.d.). In many real applications, however, count data exhibit serial dependence due to temporal aggregation, behavioral persistence, or system dynamics. It is well documented that applying independence-based control charts to autocorrelated counts can lead to distorted false-alarm behavior and degraded detection performance [1,2].

To model dependence in count time series, integer-valued autoregressive (INAR) processes [3,4] have become a standard framework. INAR models preserve the discrete and non-negative

Received: Mar. 13, 2026.

2020 *Mathematics Subject Classification.* 62F40; 62P30.

Key words and phrases. Statistical process control; Bootstrap method; Poisson INAR process, Autocorrelated count data.

https://doi.org/10.28924/2291-8639-24-2026-145
ISSN: 2291-8639

nature of counts through the binomial thinning operator and have been applied in areas such as epidemiology and disease surveillance [5, 6]. Within SPC, several control charts have been proposed for INAR(1) processes, including moving-average charts, one-sided cumulative sum (CUSUM) charts [7], and exponentially weighted moving average (EWMA) charts [2, 8]. These studies show that INAR-based monitoring can outperform classical c charts. However, in practice the model parameters are unknown and need to be estimated, making chart performance sensitive to estimation variability.

The effect of parameter estimation on SPC is well recognized, as estimation error can substantially distort run-length distributions, especially for dependent and count processes [9]. To address this issue, bootstrap-based control charts have been widely studied, including bootstrap Shewhart charts [10], bootstrap CUSUM and EWMA charts [11], percentile-based charts [12, 16], and model-based charts [13,14]. Bootstrap resampling also plays an importance role in guaranteed-performance charts, where control limits are calibrated to satisfy conditional ARL constraints with high probability [15]. However, most existing bootstrap charts are developed for independent or Gaussian data. Direct resampling under dependence is invalid, while classical sieve bootstrap methods [17], though effective for continuous-valued processes, can generate negative or non-integer values and are therefore unsuitable for count time series.

Bootstrap methods that preserve both discreteness and thinning-based dependence in INAR processes remain limited. Existing work has focused mainly on forecasting and parameter inference [5, 14, 17], whereas SPC studies for Poisson INAR processes typically rely on analytical or simulation-based EWMA and CUSUM charts under known or estimated parameters [1,2,7]. Bootstrap control charts for non-normal or autocorrelated data usually adopt continuous or residual-based approximations [10,11,13], which are not directly applicable to integer-valued autoregressive models.

Moreover, closed-form run-length distributions for EWMA and CUSUM charts under INAR(p) models are limited, particularly when parameters are estimated. This motivates bootstrap calibration, whereby control limits are chosen so that the empirical in-control average run length (ARL_0) from bootstrap-generated series matches a target value. In this paper, we develop a bootstrap framework for one-sided EWMA and CUSUM charts designed to detect positive mean shifts in Poisson INAR(p) processes. Two sieve bootstrap approaches are proposed: a Discrete (D) bootstrap that regenerates integer-valued series from the fitted Poisson INAR(p) model, preserving both discreteness and dependence, and a computationally simpler Model-Based (MB) bootstrap that conditions on Phase I parameter estimates. Control limits are calibrated via bootstrap to achieve the desired ARL_0. The performance of the proposed charts is assessed through extensive simulation studies using expected detection delay, detection rate, false-alarm rate, and non-detection rate. Finally, the methodology is illustrated with daily road-traffic accident data from Thailand.

2. Methods

2.1. **Poisson INAR(p) model.** Let X_t be a non-negative integer-valued time series observed at time t. A Poisson integer-valued autoregressive process of order p , INAR(p), is defined by

$$X_t = \sum_{i=1}^{p} \alpha_i \circ X_{t-i} + \varepsilon_t, \qquad t \geq 1, \tag{2.1}$$

where α_i are autoregressive coefficients $(0 \leq \alpha_i \leq 1)$ and $\circ$ denotes the binomial thinning operator [18]. Specifically,

$$\alpha_i \circ X_{t-i} = \sum_{j=1}^{X_{t-i}} Y_{i,j}, \tag{2.2}$$

where $Y_{i,j} \sim \text{Bernoulli}(\alpha_i)$ are mutually independent random variables, independent of the innovation process. The innovations ε_t are i.i.d. Poisson random variables with mean μ_ε, that is, $\varepsilon_t \sim \text{Poisson}(\mu_\varepsilon)$. Under the standard stability condition $\sum_{i=1}^{p} \alpha_i < 1$, the INAR($p$) process is strictly stationary and ergodic [4]. When the innovations are Poisson distributed, the marginal distribution of X_t is also Poisson with mean

$$\mu = E(X_t) = \frac{\mu_\varepsilon}{1 - \sum_{i=1}^{p} \alpha_i} \tag{2.3}$$

and the autocorrelation structure coincides with that of a real-valued AR(p) process with the same autoregressive coefficients [3,7].

In this study, the in-control process is assumed to follow an INAR(p) model with parameter vector $\theta = (\alpha_1, \ldots, \alpha_p, \mu_\varepsilon)$. Given Phase I observations, the model parameters can be estimated by conditional least squares or conditional likelihood-based methods, which are standard for INAR models and yield consistent estimators under mild regularity conditions [3,4,7]. Let $\widehat{\theta} = (\widehat{\alpha}_1, \ldots, \widehat{\alpha}_p, \widehat{\mu}_\varepsilon)$ denote the Phase I estimates. The implied estimate of the in-control mean is $\widehat{\mu} = \frac{\widehat{\mu}_\varepsilon}{1 - \sum_{i=1}^{p} \widehat{\alpha}_i}$.

To generate INAR data from a given parameter vector (θ), we use the thinning recursion

$$X_t = \sum_{i=1}^{p} \alpha_i \circ X_{t-i} + \varepsilon_t, \qquad \varepsilon_t \sim \text{Poisson}(\mu_\varepsilon), \tag{2.4}$$

with a sufficiently long burn-in period to approximate stationarity. This simulation mechanism forms the basis of the bootstrap procedures described in Section 2.3.

2.2. **EWMA and CUSUM charts.** Exponentially weighted moving average (EWMA) and cumulative sum (CUSUM) charts are well known to outperform Shewhart charts for detecting small to moderate shifts in the process mean and are generally more robust to departures from normality [19]. For this reason, CUSUM and EWMA charts have been developed to detect mean shift in autocorrelated process. In this study we focus on control charts designed to detect upward shifts in the mean of a Poisson INAR process. Let μ_0 denote the in-control mean, which is unknown in

practice and approximated by the Phase I estimate $\widehat{\mu}$. The monitoring statistics are constructed to be sensitive to positive deviations from μ_0.

2.2.1. *EWMA chart.* For a smoothing parameter λ, $(0 < \lambda \leq 1)$, the one-sided EWMA statistic for detecting a positive shift is defined recursively by

$$Z_t = \max(\lambda X_t + (1-\lambda)Z_{t-1},\, \mu_0), \qquad Z_0 = \mu_0. \tag{2.5}$$

The EWMA chart will flag an alarm if Z_t exceeds the upper control limit H.

The performance of the EWMA chart depends on the a smoothing parameter λ. A small value of λ gives more weight on past onservation, whereas a large value of λ give more weight to the current observation. When $\lambda = 1$, EWMA chart becomes equivalent to the Shewhart chart. To detect a small shift, the recommended value is $0.1 \leq \lambda \leq 0.3$ [19,20]. In this study we set $\lambda = 0.2$, as it provides a reasonable balance for detecting small to moderate shifts considered in the simulation.

2.2.2. *CUSUM chart.* To detect a positive mean shift, an upper one-sided CUSUM statistic for detecting an upward mean shift in a count process is defined as

$$C_t = \max\{0,\, X_t - (\mu_0 + k) + C_{t-1}\}, \qquad C_0 = 0, \tag{2.6}$$

where k is the reference value that determines the magnitude of shifts desired to detect. A alarm will flag if C_t exceeds an upper control limit H.

In CUSUM design, k is chosen in relation to the mean shift of interest. Typically, k is set to one-half of the shift expressed in standard deviation units [19]. However, for a Poisson count process, a shift in the mean also implies a shift in the variance, since the mean and variance of the process are not separable. Therefore, in this study, a mean shift is defined as a relative (multiplicative) change in the in-control process mean rather than as a fixed absolute deviation. The reference value is set as $k = c\mu_0$, with $c = 0.5$. Scaling the reference value with μ_0 allows the CUSUM chart to adapt to shifts at different levels of the process mean.

The control limit H for both the EWMA and CUSUM charts is determined via bootstrap calibration to achieve a desired in-control average run length, ARL_0. Details of the calibration procedure are provided in Section 2.4.

2.3. **Bootstrap for autocorrelated count data.** The classical bootstrap is designed for i.i.d. data and is not directly applicable to autocorrelated time series, particularly for discrete counts. For dependent data, sieve bootstrap methods approximate the process by a parametric model, typically an AR(p) representation, and generate resamples from the fitted model [17]. While effective for continuous-valued series, these approaches do not preserve discreteness and may produce invalid realizations when applied to integer-valued processes. In this study, we adapt the sieve bootstrap framework to Poisson INAR processes to preserve both discreteness and serial dependence, and develop bootstrap-calibrated EWMA and CUSUM charts for detecting positive mean shifts. Two bootstrap schemes are considered: a Discrete INAR bootstrap (D), which explicitly incorporates

Phase I estimation variability, and a Model-Based INAR bootstrap (MB), which conditions on a single fitted in-control model. These two approaches are described below.

2.3.1. *Discrete INAR bootstrap.* The Discrete INAR bootstrap (D) is designed to preserve the integer-valued nature and thinning-based dependence of the INAR process while explicitly accounting for parameter-estimation uncertainty arising from Phase I data. It is a fully parametric bootstrap based on the fitted Poisson INAR(p) model and incorporates model refitting within each bootstrap replication.

Let $(X_1, \ldots, X_n)$ denote the Phase I observations, and let the initial in-control fit yield the parameter estimates

$$\widehat{\theta} = (\widehat{\alpha}_1, \ldots, \widehat{\alpha}_p, \widehat{\mu}_\varepsilon), \qquad \widehat{\mu} = \frac{\widehat{\mu}_\varepsilon}{1 - \sum_{i=1}^{p} \widehat{\alpha}_i}.$$

For each bootstrap replication $b = 1, \ldots, B$, the D bootstrap proceeds in three steps.

First, a bootstrap Phase I sample $(X_1^{*(b)}, \ldots, X_n^{*(b)})$ is generated from the initially fitted model $\widehat{\theta}$, representing an in-control realization under the estimated INAR dynamics. Second, the INAR(p) model is refitted to this bootstrap Phase I sample, yielding new parameter estimates

$$\widehat{\theta}^{*(b)} = (\widehat{\alpha}_1^{*(b)}, \ldots, \widehat{\alpha}_p^{*(b)}, \widehat{\mu}_\varepsilon^{*(b)})$$

and the corresponding in-control mean

$$\widehat{\mu}^{*(b)} = \frac{\widehat{\mu}_\varepsilon^{*(b)}}{1 - \sum_{i=1}^{p} \widehat{\alpha}_i^{*(b)}}.$$

Third, $X_t^{(b)}$, an in-control bootstrap series used for calibration in replication b, is generated from $\widehat{\theta}^{*(b)}$ using the thinning recursion

$$X_t^{(b)} = \sum_{i=1}^{p} \widehat{\alpha}_i^{*(b)} \circ X_{t-i}^{(b)} + \varepsilon_t^{(b)}, \qquad \varepsilon_t^{(b)} \sim \text{Poisson}(\widehat{\mu}_\varepsilon^{*(b)}),$$

with a suitable burn-in period to approximate stationarity. Using the generated series, the EWMA or CUSUM statistic is computed using $\widehat{\mu}^{*(b)}$ as the in-control reference, and the corresponding run length $R_b(H)$ is obtained for a candidate control limit H.

By combining bootstrap resampling of Phase I data, model refitting, and regeneration of in-control bootstrap-generated series, the D bootstrap captures both the discrete data-generating mechanism and the variability induced by parameter estimation. Its main advantage lies in its fidelity to the underlying Poisson INAR model. The primary drawback is increased computational cost, as model fitting is required within each bootstrap replication.

2.3.2. *Model-based INAR bootstrap.* Alternatively, the Model-Based INAR bootstrap (MB) provides a simpler procedure than the D approach by conditioning on a single fitted in-control model. Instead of refitting the model parameters in each replication, the parameters of a Poisson INAR(p) model are estimated from the Phase I data and then kept fixed across all bootstrap replications.

The Phase I data are used to estimate $\widehat{\theta} = (\widehat{\alpha}_1, \ldots, \widehat{\alpha}_p, \widehat{\mu}_\varepsilon)$ and the corresponding mean $\widehat{\mu}$. For each bootstrap replication b, an in-control bootstrap series is generated directly from this fitted model,

$$X_t^{(b)} = \sum_{i=1}^{p} \widehat{\alpha}_i \circ X_{t-i}^{(b)} + \varepsilon_t^{(b)}, \qquad \varepsilon_t^{(b)} \sim \text{Poisson}(\widehat{\mu}_\varepsilon),$$

using a burn-in period to approximate stationarity. The EWMA or CUSUM statistic is then computed using $\widehat{\mu}$ as the in-control reference, and the run length $R_b(H)$ is recorded. This MB bootstrap is closely related to model-based INAR bootstrap procedures previously proposed for forecasting and inference in count time series [5,7,14]. From an SPC perspective, it approximates the conditional in-control run-length distribution given the Phase I parameter estimates, which is a natural target when designing a control chart for a specific, already-estimated process [10,15].

Compared with the D bootstrap, the MB approach is computationally simpler, as model parameters are estimated only once from the Phase I data. However, because MB conditions on a single set of estimated in-control parameters, it does not account for additional uncertainty arising from parameter estimation. As a result, its performance may be sensitive to estimation error or model misspecification, particularly when the Phase I sample size is limited.

2.4. **Calibration of control limits.** All control charts are calibrated to achieve a pre-specified in-control average run length, denoted by ARL_0. For a given chart type (EWMA or CUSUM), fixed design parameters (such as the smoothing parameter λ or the reference value k), and a candidate control limit H, let $R_b(H)$ denote the run length obtained from the b-th bootstrap replication. The bootstrap estimate of the in-control ARL is then given by

$$\widehat{\text{ARL}}_0(H) = \frac{1}{B} \sum_{b=1}^{B} R_b(H), \tag{2.7}$$

where B denotes the number of bootstrap replications.

The control limit H is treated as an unknown constant and is determined numerically so that $\widehat{\text{ARL}}_0(H)$ matches the desired target ARL_0. For the one-sided EWMA and CUSUM charts considered in this study, the function $\widehat{\text{ARL}}_0(H)$ is monotone increasing in H. This property allows the calibrated control limit $\widehat{H}$ to be obtained efficiently using a bisection search algorithm. Specifically, two initial values H_L and H_U are selected such that

$$\widehat{\text{ARL}}_0(H_\text{L}) < \text{ARL}_0 \quad \text{and} \quad \widehat{\text{ARL}}_0(H_\text{U}) > \text{ARL}_0.$$

At each iteration of the bisection search, the midpoint $H_\text{m} = (H_\text{L} + H_\text{U})/2$ is evaluated by computing $\widehat{\text{ARL}}_0(H_m)$ from B bootstrap-generated in-control series of length L. If $\widehat{\text{ARL}}_0(H_\text{m}) < \text{ARL}_0$, the lower bound is updated to $H_\text{L} = H_m$; otherwise, the upper bound is set to $H_\text{U} = H_\text{m}$. This procedure is repeated until

$$\left|\widehat{\text{ARL}}_0(H_m) - \text{ARL}_0\right|$$

falls below a pre-specified tolerance [21,22]. In this study, calibration is carried out in two stages: an initial pilot stage using smaller bootstrap samples to localize the control limit, followed by a

refinement stage using larger samples to reduce Monte Carlo error in $\widehat{\mathrm{ARL}}_0(H)$. This strategy improves computational efficiency while maintaining calibration accuracy [10, 15]. The above calibration procedure is applied separately for each chart type (EWMA or CUSUM) and bootstrap method (Discrete or Model-Based). As a result, all charts are calibrated to operate at the same desired in-control ARL_0, while allowing the control limits to reflect differences in bootstrap approximation and the handling of parameter-estimation variability.

3. Simulation Design

A simulation study is conducted to evaluate the performance of the proposed bootstrap EWMA and CUSUM charts for detecting upward mean shifts in Poisson INAR processes. For simplicity, simulations are carried out under a Poisson INAR(1) model, although the proposed methodology extends directly to higher-order INAR(p) models. Data are generated from the Poisson INAR(1) model $X_t = \alpha \circ X_{t-1} + \varepsilon_t$, with innovation means $\mu_\varepsilon = 2$ and 5, representing low and moderate count levels. The autoregressive parameter is set to $\alpha = 0.1, 0.3$, and 0.5, corresponding to low to moderately strong dependence. These parameter settings yield in-control process means ranging approximately from 2 to 4 for the low-count scenario and from 5 to 10 for the moderate count scenario. Each simulated series is divided into a Phase I in-control period, used for parameter estimation and bootstrap calibration, and a Phase II monitoring period following a known change point τ. Mean shifts are introduced in Phase II through changes in μ_ε, which directly control the marginal mean while preserving the thinning-based dependence structure.

In each replication, model parameters $(\widehat{\alpha}, \widehat{\mu}_\varepsilon)$ are estimated by conditional maximum likelihood and used for both the Discrete (D) and Model-Based (MB) bootstrap procedures. Control limits are calibrated via bootstrap using the method described in Section 2.4, employing between 1,000 and 5,000 bootstrap replications within a bisection search to achieve the target ARL_0. After calibration, the charts are applied to Phase II data. At time τ, a positive mean shift is introduced by increasing the innovation mean according to

$$\mu'_\varepsilon = (1 + r)\mu_\varepsilon, \tag{3.1}$$

where r denotes the relative shift size, ranging from 0.1 to 2, covering small to large mean increases. Since the marginal mean of the INAR(1) process satisfies $\mu = \mu_\varepsilon/(1-\alpha)$, this implies a multiplicative change in the process mean, $\mu' = (1 + r)\mu$, while preserving the thinning-based dependence structure.

Detection performance is evaluated using 5,000 Monte Carlo replications for each combination of parameters and bootstrap method. Let T_A denote the alarm time and L the length of the monitoring period. Four performance measures are considered:

(1) Expected Detection Delay (EDD) measures the timeliness of detection conditional on a true signal: $\mathrm{EDD} = E[T_A - \tau \mid T_A \geq \tau]$.
(2) False Alarm Rate (FA) is the probability of signaling before the change point: $\mathrm{FA} = P(T_A < \tau)$.

(3) True Detection Rate (DT) is the probability of correctly detecting the shift during the monitoring period: DT $= P(\tau \leq T_A \leq L)$.

(4) Non-Detection Rate (ND) is the probability of failing to detect the shift within the monitoring period: ND $= P(T_A > L)$.

The events FA, DT, and ND are mutually exclusive and satisfy FA + DT + ND = 1 [23]. These measures are estimated from the simulation study. These measures provide an overall assessment of detection speed and reliability under different simulation scenarios.

4. Results

This section compares the detection performance of the proposed bootstrap monitoring charts: CUSUM_{D}, CUSUM_{MB}, EWMA_{D}, and EWMA_{MB}. The comparison of detection speed (EDD) is illustrated using line graphs, whereas detection reliability (FA, DT and ND) is summarized using stacked bar charts across different levels of dependence and shift magnitude. Results are presented for both low and moderate count scenarios.

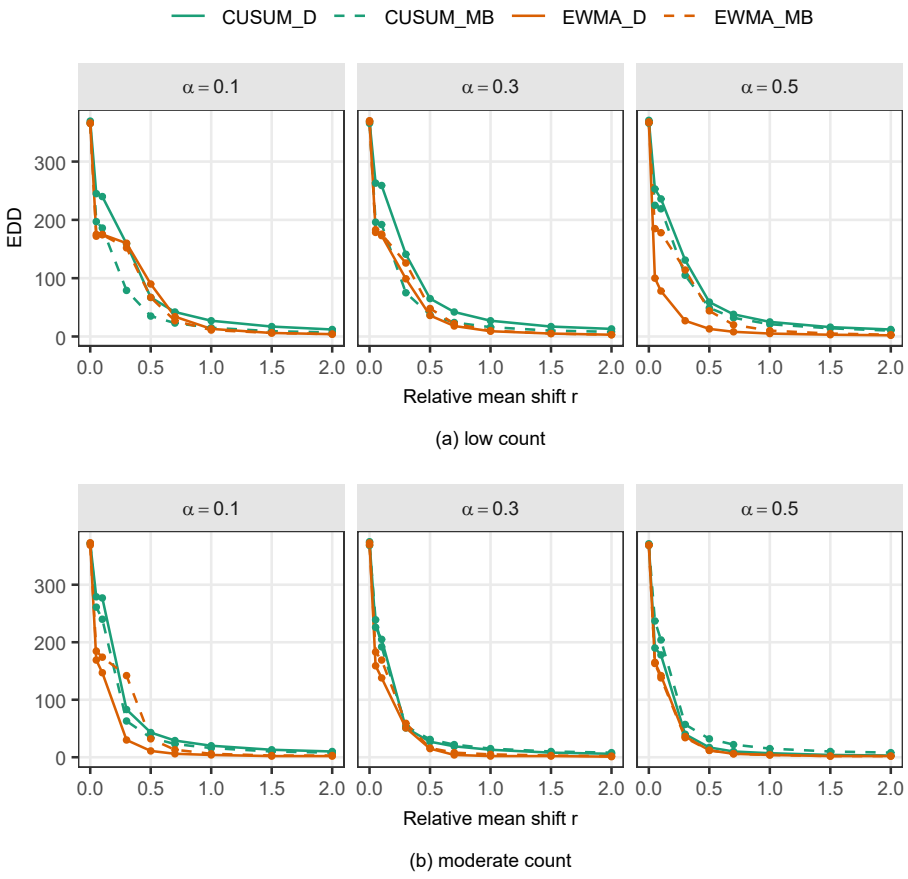

Figure 1. Expected detection delay (EDD) across charts and scenarios.

The conditional expected detection delay (EDD) for the two scenarios is shown in Figure 1. Across all charts, EDD decreases as the shift magnitude increases. Considering the count levels, processes with moderate counts tend to be detected faster than those with low counts. In the low-count scenario, the charts perform quite differently under moderate to strong dependence. EWMA_{D} produces the shortest delays for detecting small to moderate shifts, whereas the CUSUM charts give the longest delays under these conditions. Under weak dependence, the performance of CUSUM_{MB} improves for very small shifts, but this advantage diminishes as r increases. In the moderate count scenario, the EWMA charts generally give slightly shorter delays than the CUSUM charts. At low dependence levels, EWMA_{D} continues to provide the fastest detection for small and moderate shifts. However, under strong dependence, EWMA_{D} and EWMA_{MB} show very similar EDD values, indicating that the benefit of refitting in the bootstrap becomes less substantial at higher count levels. In contrast, both CUSUM charts remain slightly slower for small shifts.

Detection reliability for small to moderate shifts $(r = 0.1, 0.3, 0.5)$ is shown in Figure 2, where differences between charts are most evident. Results for larger shifts are omitted, as all charts behave similarly. Across all scenarios, false alarm rates (FA) are very small, indicating that the bootstrap calibration maintains stable in-control behavior. Differences in performance between charts therefore depend mainly on the detection rates (DT) and non-detection rates (ND).

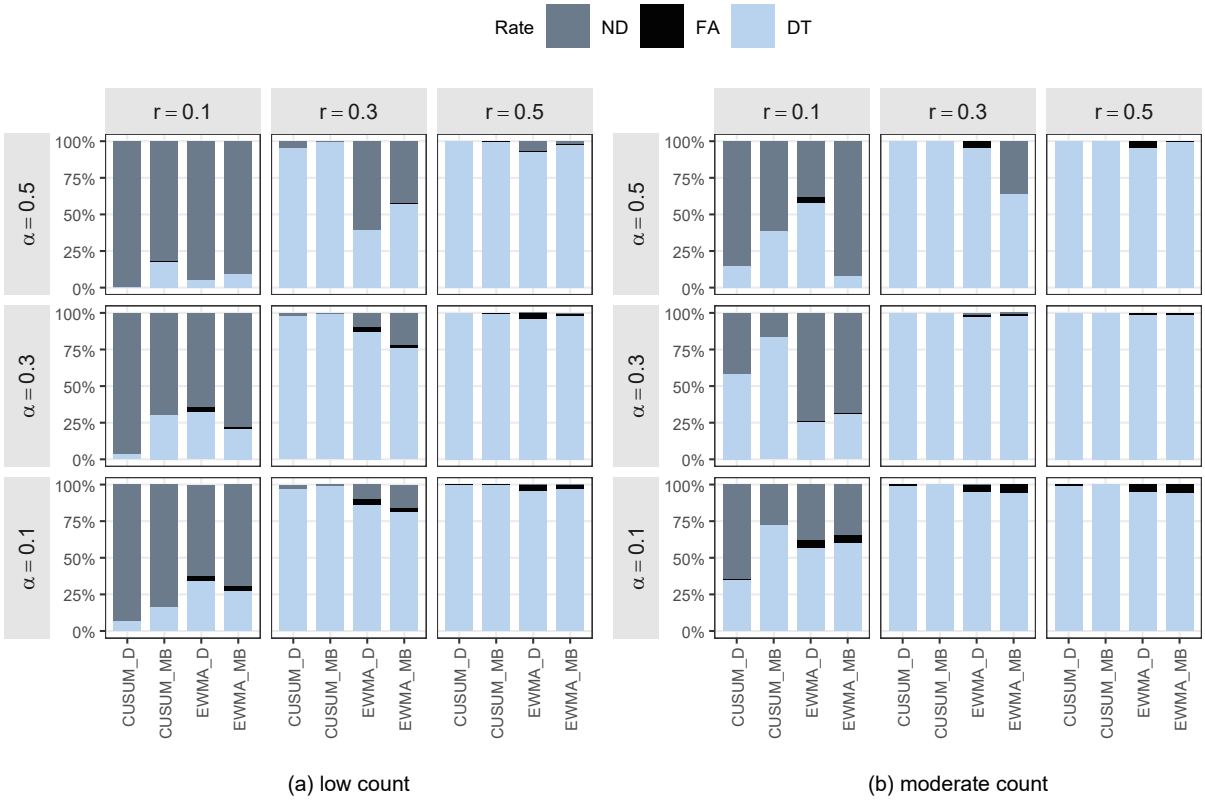

FIGURE 2. Detection reliability across charts and scenarios.

For very small shifts ($r = 0.1$) under weak dependence, EWMA_{D} produces a higher detection rate (DT) than the other charts for detecting shifts in the low count process, whereas CUSUM_{MB} outperforms the others in the moderate count process. Under strong dependence, however, CUSUM_{MB} yields the highest DT in the low-count scenario, while EWMA_{D} achieves the highest DT in the moderate count scenario. Nevertheless, all methods perform poorly in detecting such small shifts as the dependence level increases, especially in the low-count scenario. These results suggest that stronger serial dependence can affect detection performance, with missed detections remaining common across methods.

As the shift size increases ($r = 0.3$ and 0.5), detection performance improves substantially in both scenarios. Missed detections decrease as the shift size increases. In the low count process, the CUSUM charts outperform the EWMA charts. At $r = 0.3$, , EWMA_{D} produces a higher DT than EWMA_{MB}. However, this pattern reverses when the process is strongly dependent. In the moderate-count scenario, the CUSUM charts remain slightly superior to the EWMA charts. The discrete and model-based approaches perform similarly under low dependence. Under strong dependence, however, EWMA_{MB} shows weaker performance by yielding a higher ND. When $r = 0.5$, the differences in performance between the charts become small.

5. Application to Real Data

The proposed control charts are illustrated using two real datasets on the daily number of road-traffic injuries in Phuket and Chiang Mai, representing processes with low and high autocorrelation, respectively. The data are publicly available from the Injury Data Collaboration Center (IDCC) in Thailand. Data during the Songkran holiday were excluded because injury counts are unusually high during this period. Each dataset was divided into two phases: Phase I for parameter estimation and control limit calibration, and Phase II for evaluating detection performance. As no obvious change was observed in the series, a change point was selected arbitrarily, and synthetic relative mean shifts (r), ranging from 0.01 to 2, were introduced into Phase II for evaluation purposes. Figure 3 shows the corresponding time series, with the dashed line indicating the start of Phase II.

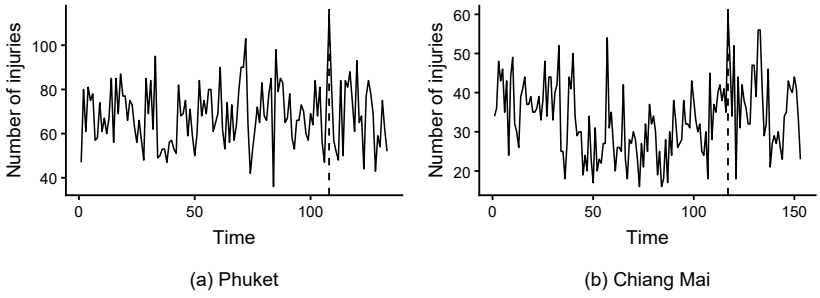

Figure 3. Daily number of road-traffic injuries in Phuket and Chiang Mai.

For Phuket (low dependence), the change point was set at 6 August 2025. Phase I includes 107 observations (18 April–5 August 2025), with an estimated mean of 68.08, standard deviation of 12.97, and a Poisson INAR(1) autoregressive coefficient of 0.113. Phase II consists of 26 observations (6–31 August 2025). For Chiang Mai (high dependence), the change point was set at 1 July 2025. Phase I contains 116 observations (1 March–30 June 2025), with an estimated mean of 32.64, standard deviation of 9.07, and an autoregressive coefficient of 0.450. Phase II includes 37 observations (1 July–7 August 2025).

Figures 4 and 5 illustrate the one-sided EWMA and CUSUM charts for detecting a relative mean shift of $r = 0.3$. Although the control limits under the D approach are slightly lower than those under the MB approach, Table 1 shows that detection performance is nearly identical between the two bootstrap implementations. For Phuket, small shifts ($r \leq 0.1$) are unlikely to be detected during the Phase II monitoring period, resulting in NA values for all procedures. Once the shift becomes moderate ($r \geq 0.3$), all charts detect the change immediately, with zero detection delay. For Chiang Mai, differences appear only for very small shifts. When $r = 0.01$ and $r = 0.05$, the CUSUM charts exhibit delayed detection, whereas both EWMA charts signal immediately. For $r \geq 0.1$, all charts detect the shift without delay. Overall, these results indicate that performance differences are confined to extremely small shifts, while for moderate and larger shifts the monitoring schemes behave equivalently across both dependence levels.

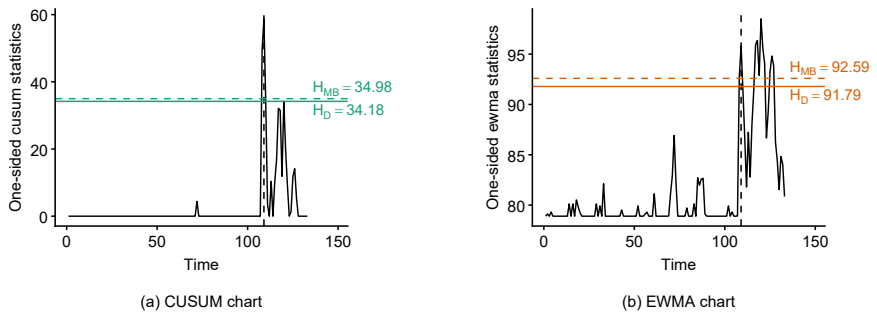

FIGURE 4. CUSUM and EWMA charts for detecting a mean shift in Phuket.

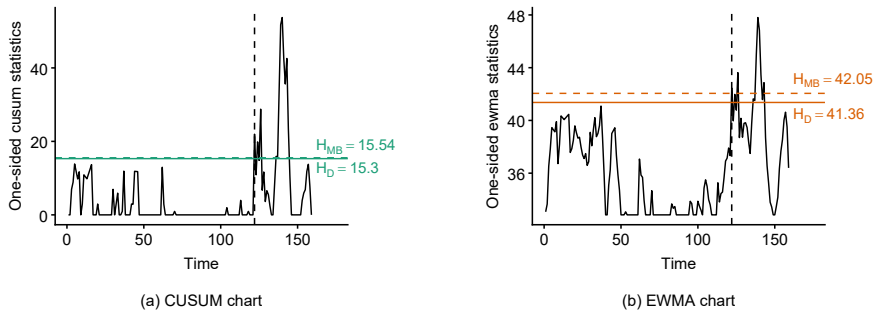

FIGURE 5. CUSUM and EWMA charts for detecting a mean shift in Chiang Mai.

Table 1. Expected detection delay (EDD) for injury count data.

Province	Chart	Relative shift magnitude (r)							
		0.01	0.05	0.1	0.3	0.5	1	1.5	2
Phuket	$CUSUM_D$	NA	NA	NA	0	0	0	0	0
	$CUSUM_{MB}$	NA	NA	NA	0	0	0	0	0
	$EWMA_D$	NA	NA	NA	0	0	0	0	0
	$EWMA_{MB}$	NA	NA	NA	0	0	0	0	0
Chiang Mai	$CUSUM_D$	17	17	0	0	0	0	0	0
	$CUSUM_{MB}$	17	17	0	0	0	0	0	0
	$EWMA_D$	0	0	0	0	0	0	0	0
	$EWMA_{MB}$	0	0	0	0	0	0	0	0

6. Conclusion

This paper developed and evaluated bootstrap-calibrated one-sided EWMA and CUSUM charts for detecting upward mean shifts in Poisson INAR(p) processes with unknown parameters. Two bootstrap approaches were examined: a Discrete (D) bootstrap that refits the model within each replication to reflect parameter-estimation variability, and a Model-Based (MB) bootstrap that conditions on a single fitted in control model. Performance was assessed using detection speed, measured by the expected detection delay (EDD), and detection reliability, summarized through false-alarm, detection, and non-detection rates. Across all scenarios, detection performance depends on the shift magnitude, with EDD decreasing rapidly as the shift size increases. The count level and serial dependence of the process also influence performance, particularly for small shifts in low-count or strongly dependent processes, where very small shifts are difficult to detect regardless of the chart used.

Differences between EWMA and CUSUM charts are most apparent for very small shifts. Simulation results indicate that EWMA charts generally provide faster detection, reflected in lower expected detection delay (EDD), whereas CUSUM charts may achieve slightly higher detection rates (DT). However, as the shift becomes moderate to large, differences in performance between the charts become small. These results highlight a trade-off between detection speed and detection reliability. The methods that respond more quickly might not always yield the highest detection rate. Therefor, evaluating both detection delay and detection reliability provides a more comprehensive assessment of monitoring performance.

Comparing the two bootstrap procedures, the Discrete (D) approach tends to provide more reliable performance for very small shifts, particularly in low-count or strongly dependent processes. In contrast, the Model-Based (MB) approach, which conditions on a single estimated model, may

exhibit slightly weaker detection reliability under these conditions. However, as the shift magnitude increases, the two approaches show only small differences and perform quite similarly for moderate and large shifts. The D bootstrap may be preferable for detecting small shifts in low-count or strongly dependent processes, where parameter estimation variability can influence detection reliability. The MB approach offers a computationally efficient alternative and performs similarly to the D approach for moderate and large shifts. However, because the MB approach relies on a single fitted model, it may be more sensitive to parameter misspecification, which could affect detection performance.

Finally, although control limits are calibrated to achieve a common ARL_0, detection performance also depends on the design parameters of the charts. The EWMA smoothing parameter λ determines the weights assigned to past and current observations, while the CUSUM reference value k determines the shift size to which the chart is most sensitive. Different choices of λ and k may affect small-shift performance, particularly in highly dependent or low-count processes. Extensions to higher-order INAR(p) models, alternative count distributions accommodating overdispersion, and systematic sensitivity analysis of chart design parameters would be considered for future research.

Acknowledgments: The author would like to thank the Injury Data Collaboration Center (IDCC), Department of Disease Control, Ministry of Public Health, Thailand, for providing the data used in this study. This work was supported by the International SciKU Branding (ISB), Faculty of Science, Kasetsart University, and the Department of Statistics, Faculty of Science, Kasetsart University, Thailand.

Conflicts of Interest: The author declares that there are no conflicts of interest regarding the publication of this paper.

References

[1] C.H. Weiß, M.C. Testik, CUSUM Monitoring of First-Order Integer-Valued Autoregressive Processes of Poisson Counts, J. Qual. Technol. 41 (2009), 389–400. https://doi.org/10.1080/00224065.2009.11917793.

[2] M. Zhang, G. Nie, Z. He, X. Hou, The Poisson INAR(1) One-Sided EWMA Chart with Estimated Parameters, Int. J. Prod. Res. 52 (2014), 5415–5431. https://doi.org/10.1080/00207543.2014.907517.

[3] E. McKenzie, Some Simple Models for Discrete Variate Time Series, J. Am. Water Resour. Assoc. 21 (1985), 645–650. https://doi.org/10.1111/j.1752-1688.1985.tb05379.x.

[4] M.A. Al-Osh, A.A. Alzaid, First-order Integer-valued Autoregressive (INAR(1)) Process, J. Time Ser. Anal. 8 (1987), 261–275. https://doi.org/10.1111/j.1467-9892.1987.tb00438.x.

[5] M. Cardinal, R. Roy, J. Lambert, On the Application of Integer-valued Time Series Models for the Analysis of Disease Incidence, Stat. Med. 18 (1999), 2025–2039. https://doi.org/10.1002/(sici)1097-0258(19990815)18:15<2025::aid-sim163>3.3.co;2-4.

[6] A.C. Rakitzis, P. Castagliola, P.E. Maravelakis, Cumulative Sum Control Charts for Monitoring Geometrically Inflated Poisson Processes: An Application to Infectious Disease Counts Data, Stat. Methods Med. Res. 27 (2016), 622–641. https://doi.org/10.1177/0962280216641985.

[7] C.H. Weiß, EWMA Monitoring of Correlated Processes of Poisson Counts, Qual. Technol. Quant. Manag. 6 (2009), 137–153. https://doi.org/10.1080/16843703.2009.11673190.

[8] C.H. Weiß, Detecting Mean Increases in Poisson INAR(1) Processes with EWMA Control Charts, J. Appl. Stat. 38 (2010), 383–398. https://doi.org/10.1080/02664760903406520.

[9] W.A. Jensen, L.A. Jones-Farmer, C.W. Champ, W.H. Woodall, Effects of Parameter Estimation on Control Chart Properties: A Literature Review, J. Qual. Technol. 38 (2006), 349–364. https://doi.org/10.1080/00224065.2006.11918623.

[10] L.A. Jones, The Statistical Design of EWMA Control Charts with Estimated Parameters, J. Qual. Technol. 34 (2002), 277–288. https://doi.org/10.1080/00224065.2002.11980158.

[11] S. Chatterjee, P. Qiu, Distribution-Free Cumulative Sum Control Charts Using Bootstrap-Based Control Limits, Ann. Appl. Stat. 3 (2009), 349–369. https://doi.org/10.1214/08-AOAS197.

[12] J.Y. Chiang, Y. Lio, H. Ng, T.R. Tsai, T. Li, Robust Bootstrap Control Charts for Percentiles Based on Model Selection Approaches, Comput. Ind. Eng. 123 (2018), 119–133. https://doi.org/10.1016/j.cie.2018.06.012.

[13] G. Capizzi, G. Masarotto, Bootstrap-Based Design of Residual Control Charts, IIE Trans. 41 (2009), 275–286. https://doi.org/10.1080/07408170802120059.

[14] L. Bisaglia, M. Gerolimetto, Model-Based INAR Bootstrap for Forecasting INAR(p) Models, Comput. Stat. 34 (2019), 1815–1848. https://doi.org/10.1007/s00180-019-00902-1.

[15] A. Gandy, J.T. Kvaløy, Guaranteed Conditional Performance of Control Charts via Bootstrap Methods, Scand. J. Stat. 40 (2013), 647–668. https://doi.org/10.1002/sjos.12006.

[16] S. Panda, M. Wang, Bootstrap-Based Control Chart for Percentiles of the Generalized Lognormal Distribution with Reliability Applications, Qual. Reliab. Eng. Int. 41 (2025), 1329–1349. https://doi.org/10.1002/qre.3722.

[17] P. Bühlmann, P. Buhlmann, Sieve Bootstrap for Time Series, Bernoulli 3 (1997), 123–148. https://doi.org/10.2307/3318584.

[18] A.A. Alzaid, M. Al-osh, An Integer-Valued pth-Order Autoregressive Structure (INAR(p)) Process, J. Appl. Probab. 27 (1990), 314–324. https://doi.org/10.2307/3214650.

[19] D.C. Montgomery, Introduction to Statistical Quality Control, John Wiley & Sons, 2019.

[20] J.M. Lucas, M.S. Saccucci, Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements, Technometrics 32 (1990), 1–12. https://doi.org/10.1080/00401706.1990.10484583.

[21] R. Chen, Z. Li, J. Zhang, A Generally Weighted Moving Average Control Chart for Monitoring the Coefficient of Variation, Appl. Math. Model. 70 (2019), 190–205. https://doi.org/10.1016/j.apm.2019.01.034.

[22] D. Zago, G. Capizzi, P. Qiu, An Improved Bisection-Type Algorithm for Control Chart Calibration, Stat. Comput. 35 (2025), 81. https://doi.org/10.1007/s11222-025-10609-7.

[23] M. Frisén, Methods and Evaluations for Surveillance in Industry, Business, Finance, and Public Health, Qual. Reliab. Eng. Int. 27 (2011), 611–621. https://doi.org/10.1002/qre.1204.