

Research Article

Improving EWMA Plans for Detecting Unusual Increases in Poisson Counts

R. S. Sparks,¹ T. Keighley,¹ and D. Muscatello²

¹ CSIRO Mathematical and Information Sciences, Locked Bag 17, North Ryde NSW 1670, Australia

² Centre for Epidemiology and Research, NSW Health Department, Locked Mail Bag 961, North Sydney NSW 2059, Australia

Correspondence should be addressed to R. S. Sparks, ross.sparks@csiro.au

Received 11 November 2008; Revised 27 April 2009; Accepted 19 June 2009

Recommended by Chin Lai

Automated public health records provide the necessary data for rapid outbreak detection. An adaptive exponentially weighted moving average (EWMA) plan is developed for signalling unusually high incidence when monitoring a time series of nonhomogeneous daily disease counts. A Poisson transitional regression model is used to fit background/expected trend in counts and provides “one-day-ahead” forecasts of the next day’s count. Departures of counts from their forecasts are monitored. The paper outlines an approach for improving early outbreak data signals by dynamically adjusting the exponential weights to be efficient at signalling local persistent high side changes. We emphasise outbreak signals in steady-state situations; that is, changes that occur after the EWMA statistic had run through several in-control counts.

Copyright © 2009 R. S. Sparks et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. Introduction

Early detection of disease outbreaks is essential for the efficient control of acute public health risks. For example, identifying the source of, and restricting exposure to, contaminated food products can only proceed once there is enough information for an outbreak data signal. Automated detection systems that accumulate sufficient information for earlier identification of an outbreak data signal are common now days. Identifying unusual disease incidence is normally carried out by monitoring data streams such as daily, weekly, or monthly counts or rates.

Statistical process control (SPC) methods, including control charts, are increasingly being applied to public health surveillance [1, 2]. Unusual incidence is automatically signalled when the control chart of the disease time series exceeds a threshold “control limit”. However, disease incidence time series can exhibit characteristics that increase the complexity of detection of unusual incidences. Counts typically vary due to known influences

such as school holidays, seasons, day of the week, and normal transitional (autocorrelation or autoregressive) influences. Subtracting expected behaviour from the counts reduces predictable variation effects in the time series, making it easier to detect unusual behaviour. Traditional monitoring of standardised differences between observed and modelled counts “standardised forecast errors” assumes normality of the forecast errors. However, Buckeridge et al. [2] indicated that conventional process control methods applied to syndromic data result in thresholds that produce unacceptably high false alarm rates when thresholds are not adjusted for the violation of distributional and other assumptions. The focus of the paper is detecting unexpected aberrations in heterogeneous time series of counts while trying to control for false alarms.

Exponentially weighted moving average (EWMA) control charts have been useful for monitoring industrial production processes (e.g., see Lucas and Saccucci [3]). More recently they have been used in public health surveillance (see review article by Woodall [1]) and for monitoring network counts (see Lambert and Liu [4]). Lambert and Liu used an EWMA plan of pseudo residuals (see Zucchini and MacDonald [5]) to monitor network counts (note that a surveillance methodology with all parameters specified is defined as a plan in this paper). We propose an alternative approach that avoids conversion to pseudo residuals. Lambert and Liu [4] used the EWMA to update means every minute, or on hourly cycles, whereas we use transitional Poisson regression models to update one-day-ahead forecast mean daily rates similar to that of Burr et al. [6]. We compared the CUSUM rather than EWMA of pseudoforecast errors to our proposed plans, because CUSUM and EWMA have similar performance in terms of early detection, but the CUSUM plan can be optimised [7].

Grigg and Spiegelhalter [8] used generalised linear models to estimate expected outcomes and correct counts by adjusting for risk factors. Their adjusted scores (called pseudo-observations) facilitate the temporal comparisons to baseline values thus allowing confounders to the comparison to be removed. Similarly we adjust for the influences of known variation caused by day of the week influences, seasonal influences and usual transitional influences, thus giving our plans a better chance of detecting unusual causes of variation. Unlike Grigg and Spiegelhalter [8] a primary aim of our work is to improve the early detection of outbreak data signals; a major goal of automated real-time public health surveillance. We offer a practical way of dynamically changing the exponential weights, giving the plan a better chance of detecting predictable (yet relatively small) (unusual) epidemic footprints. The strategy used is similar to that applied by Capizzi and Masarotto [9] and Nembhard and Kao [10]. The traditional method for overcoming the heterogeneity in trend is to use one-step-ahead forecasts from regression models that conform to usual, background disease count behaviour and then calculate control charts from the standardised forecast residuals (e.g., see Parker [11], Miller et al. [12], Kleinschmidt et al. [13], Wang et al. [14], Farrington et al. [15], Cowling et al. [16], Stoto et al. [17], and Kaninda et al. [18]). The forecast residuals are the difference between new count observations and the forecast count predicted by the regression model. In other words, regression models are used to define usual or expected behaviour, and the forecast residuals to detect unusual variation. We included explanatory variables that accounted for

- (i) the annual seasonal influence,
- (ii) day of the week influences,
- (iii) usual transitional (day-to-day) carry over influences of the lag counts.

EWMA plans are used to accumulate memory of local trends. The benefit of accumulating memory is having sufficient power to signal a change. The exponential weights λ control the amount of memory retained in the EWMA plan. Lower values retain a longer memory.

Monitoring performance is measured in terms of out-of-control average time to signal. We used the average number of days from the day of onset of outbreak to the day the outbreak signal to define performance. The average time to signal for a step change of δ in mean (denoted by ATS_δ) is used to measure a plan's early detection performance. An in-control ATS_0 of 100 is used for monitoring daily counts, which translates to just over 3 false alarms per annum.

Most of the work on control chart efficiency in the literature on surveillance has been carried out for EWMA plans in their initial state where the out-of-control behaviour is generated starting the EWMA at its in-control value. In practice, process changes seldom start when EWMA statistics equals the in-control value. Therefore, all ATS_δ values are found using the more realistic steady-state situation, where the EWMA statistic has worked through some in-control background data prior to the introduction of outbreak data. Steady-state situations were simulated by running the EWMA plans through 25 days of in-control counts (denoted the burn-in period) before introducing a change in the mean count representing outbreak data.

Section 2 introduces the EWMA for homogeneous counts. Section 2.1 examines how the EWMA plan can be improved when the in-control count follows a homogeneous Poisson distribution. Section 3 extends the EWMA plan for nonhomogeneous background (in-control) counts, and examines their improvement when parameters are known. We discuss how the performance of EWMA plans can be improved when counts are nonhomogeneous Poisson distributed. We consider early detection of both seasonal epidemic footprints, such as the start of the influenza season, and unusual increases in counts. Seasonal epidemic footprints are expected each year and therefore are not unusual. However, early detection is important because seasonal epidemic footprints vary in size and timing. Early detection of the start of the epidemic footprint defines when closer monitoring is required for informed decision making, such as, increasing the use of agency staff to cope with increased patient demand for hospital emergency department services. A simulated example is used to demonstrate the efficiency of the recommended improvement strategy and demonstrate that this plan is significantly better than using the "optimal" adaptive CUSUM of pseudo residuals in terms of early signals for the same false alarm rate. Section 4 presents an example of application that demonstrates the value of our improved plans relative to alternatives. The paper concludes with a discussion in Section 5.

2. EWMA Plans for Homogeneous Counts

Let the daily count on day t of a certain disease be denoted by y_t . The EWMA statistic for these counts is given by $EWMA_0 = \mu_0$ (in-control or background means are constant with time therefore homogeneous counts have means that do not depend on t) and

$$EWMA_t = \lambda y_t + (1 - \lambda)EWMA_{t-1}, \quad (2.1)$$

where λ (the exponential weight) is a suitable constant $0 < \lambda < 1$. The λ value determines the level of memory included in the EWMA plan. At time t , $y_{t-\ell}$, the count for time $t - \ell$ is given

Table 1: Threshold values $h(\mu_0, \lambda, 100)$ for a range of μ_0 and λ values.

λ	μ_0										
	1	2	3	3.5	4	4.5	5	5.5	6	6.5	7
	$h(\mu_0, \lambda, 100)$										
0.05	1.232	2.336	3.410	3.939	4.466	4.991	5.515	6.038	6.562	7.085	7.608
0.07	1.319	2.450	3.547	4.088	4.625	5.159	5.692	6.224	6.756	7.287	7.818
0.10	1.435	2.602	3.732	4.288	4.839	5.387	5.932	6.476	7.018	7.560	8.100
0.15	1.604	2.827	4.005	4.583	5.154	5.720	6.283	6.843	7.401	7.957	8.511
0.20	1.758	3.033	4.254	4.849	5.438	6.020	6.597	7.171	7.741	8.309	8.875
0.25	1.905	3.229	4.489	5.101	5.705	6.301	6.891	7.477	8.058	8.637	9.213
0.30	2.039	3.412	4.710	5.339	5.957	6.566	7.168	7.765	8.357	8.946	9.532
0.35	2.142	3.565	4.904	5.549	6.182	6.805	7.419	8.027	8.631	9.230	9.827

a weight of $(1 - \lambda)^\ell$, therefore a burn-in of 25 observations is adequate for the EWMA output series to reduce to a steady-state for values of $\lambda > 0.05$, because $(1 - \lambda)^{25}$ is close to zero.

Suppose outbreak data cause the mean to increase from μ_0 to μ_1 . An important consideration is the selection of a λ -value that detects this shift earliest.

For Poisson counts with mean μ , the EWMA plan signals a significant shift in mean from the in-control mean of μ_0 whenever

$$\text{EWMA}_t > h(\mu_0, \lambda, \text{ATS}_0), \quad (2.2)$$

where $h(\mu_0, \lambda, \text{ATS}_0)$ is a function of μ_0 , λ , and ATS_0 (the in-control average run length). Values for $h(\mu_0, \lambda, \text{ATS}_0 = 100)$ were found using simulation techniques (for more details on how the control limits were found see Sparks et al. [19]). Table 1 provides values for $h(\mu_0, \lambda, 100)$ for a range of μ_0 and λ values.

The EWMA_t statistic is easy to interpret because it provides a local estimate of the mean for daily counts (the average incidence expected per day). Improving EWMA charts for early detection is now investigated in steady-state situations.

2.1. Improving EWMA Plans for Homogeneous Poisson Counts When Parameters Are Known

Sparks [7] and Shu and Jiang [20] optimised adaptive CUSUM plans for detecting shifts when data are approximately normally distributed. The same approach is followed here. Let the $\lambda_{\text{opt}}(\mu_0, \mu_1)$ value be the exponential weight that provides the earliest signal (the smallest ATS_δ) for the shift from μ_0 to μ_1 (i.e., $\delta = \mu_1 - \mu_0$). An operating model for interpolating $\lambda_{\text{opt}}(\mu_0, \mu_1)$ values for values of $0.05 \leq \mu_0 \leq 30$ is provided in (2.3). The operating model is used for interpolating $\lambda_{\text{opt}}(\mu_0, \mu_1)$ values for any given values for μ_0 and μ_1 . When $\mu_0 > 30$ the normal approximation is adequate and standard EWMA plans can be applied.

Table 2: “Optimal” λ values for a range of μ_0 and μ_1 values ($\hat{\lambda}_{\text{opt}}(\mu_0, \mu_1)$).

$\mu_1 = \eta \times \mu_0$	μ_0												
η	1	1.5	2	2.5	3	3.5	4	4.5	5	5.5	6	6.5	7
1.3				0.05	0.06	0.07	0.07	0.08	0.08	0.09	0.09	0.10	0.10
1.4	0.05	0.06	0.07	0.08	0.09	0.11	0.12	0.13	0.14	0.15	0.16	0.17	0.18
1.5	0.06	0.08	0.10	0.11	0.13	0.15	0.16	0.18	0.19	0.21	0.23	0.24	0.26
1.6	0.07	0.10	0.12	0.14	0.16	0.19	0.21	0.23	0.25	0.27	0.30	0.32	0.34
1.7	0.08	0.11	0.14	0.17	0.20	0.23	0.25	0.28	0.31	0.34			
1.8	0.10	0.13	0.16	0.20	0.23	0.27	0.30	0.33					
2.0	0.12	0.16	0.21	0.26	0.30	0.35							
2.5	0.18	0.25	0.32										
3.0	0.23	0.34											

For disease counts that are Poisson distributed, a model for interpolating “optimal” exponential weights (denoted by $\lambda_{\text{opt}}(\mu_0, \mu_1)$) that deliver approximately the smallest ATS_δ for $\text{ATS}_0 = 100$ and $0.05 \leq \lambda_0 \leq 0.35$ is

$$\begin{aligned} \hat{\lambda}_{\text{opt}}(\mu_0, \mu_1) = & 0.02130805 + 0.02945567\mu_1 - 0.03252782\mu_0 + 0.00659441z_t \\ & + 0.00009397627\mu_0\mu_1 + 0.08497906z_t\mu_1 - 0.1052998z_t\mu_0 - 0.0001734491z_t\mu_0\mu_1, \end{aligned} \tag{2.3}$$

where z_t is an indicator variable that takes on a value of 1 if $(\mu_1 - \mu_0)/\mu_0 > 0.25$ and zero otherwise. For example, improving EWMA plans for detecting shifts from 1 to 2.25, 2 to 4, 3 to 5.25, and 4 to 6.5 would be achieved approximately by selecting $\lambda_{\text{opt}}(\mu_0, \mu_1)$ equal 0.175, 0.2, 0.225, and 0.225, respectively. The restriction of the model (2.3) to $\lambda \leq 0.35$ is not serious because the ATS_δ values for shifts that lead to $\lambda_{\text{opt}}(\mu_0, \mu_1) \geq 0.35$ are generally 2 or less, and therefore there is little opportunity to improve the EWMA performance for such large shifts. Table 2 reports $\hat{\lambda}_{\text{opt}}(\mu_0, \mu_1)$ for all combinations of $\mu_0 = 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5, 5.5, 6, 6.5,$ or 7 and $\mu_1 = c\mu_0$ where $c = 1.3, 1.4, 1.5, 1.6, 1.7, 1.8, 2, 2.5,$ or 3. The correlation between the “true” values (derived from simulation) and the model fitted values (expected) using (2.3) is 0.984 with the standard error for the regression model is given by 0.0188. The standard error of 0.0188 is adequate because the EWMA plan with $\lambda \in \lambda_{\text{opt}}(\mu_0, \mu_1) \pm 2 \times 0.0188$ delivers close to the lowest ATS_δ value. The change in ATS_δ is only slight within this range. In practice, improving EWMA plans for shifts from μ_0 to μ_1 is achievable by selecting $\lambda = \lambda_{\text{opt}}(\mu_0, \mu_1)$ only when μ_0 and μ_1 are known or can be reliably estimated. The minimum value of $\hat{\lambda}_{\text{opt}}(\mu_0, \mu_1)$ is taken as 0.05. Placing this lower bound on $\hat{\lambda}_{\text{opt}}(\mu_0, \mu_1)$ reduces the importance of early detection of small shifts that may not be worth investigating.

3. Plans for Nonhomogeneous Counts

For nonhomogeneous counts y_t , the EWMA statistic of (2.1) needs to be standardised in some way to remove the influence of nonhomogeneity, that is, changing μ_0 with time t , denoted by

$\mu_0(t)$. In addition, the optimisation process for the plan involves changing the exponential weights dynamically for each day t to λ_t ($\neq \lambda$). One possible standardisation is to divide both sides of inequality (2.2) by $h(\mu_0(t), \lambda_t, ATS_0)$ to adjust for changing $\mu_0(t)$ and λ_t with time t . For known moving means $\mu_0(t)$, this standardisation leads to the adaptive EWMA for nonhomogeneous Poisson counts given by

$$AEWMA_t = \frac{\lambda_t y_t}{h(\mu_0(t), \lambda_t, ATS_0)} + (1 - \lambda_t) AEWMA_{t-1}. \quad (3.1)$$

Statistic (3.1) is similar to the adaptive control statistic used in Sparks [7] for normally distributed data. Sparks developed a CUSUM plan, that is, nearly invariant to changes in the reference value. Similarly $AEWMA_t$'s control limit is nearly invariant to changes in $\mu_0(t)$ and λ_t . That is, the statistic $AEWMA_t$ signals whenever

$$AEWMA_t > 1 \quad (3.2)$$

with average time between false alarms close to ATS_0 . The invariant property of plans are very useful, because plans do not need to be adjusted for changes in $\mu_0(t)$ and λ_t . The multiplicative adjustment a to the threshold of 1 for $AEWMA_t$ needed to achieve $ATS_0 = 100$ is examined. The control limit of 1 is based on homogeneous Poisson count control limits with known means, and these may need to be adjusted for nonhomogeneous Poisson counts where means are unknown and are estimated from the data. The adjustments for specific models are available for particular examples in Tables 3 and 4. Note that a values are very close to 1 for all models considered, that is, almost no adjustment was required. In addition, it will be demonstrated later that estimating λ_t to be temporally "optimal" makes the plan perform better than fixing λ at 0.1.

A disadvantage with $AEWMA_t$ is that it has lost the simple interpretation of $EWMA_t$ estimating the local mean. In addition, it is difficult to establish whether $AEWMA_t$ is trending above or below what is expected. It is therefore helpful to draw the time series line of the EWMA smoothed value of $\mu_0(t)/h(\mu_0(t), \lambda, ATS_0)$ on the time series plot of $AEWMA_t$ (represented by a dashed line in all plots). If $AEWMA_t$ is trending above or below the dashed line then counts are higher or lower than expected, respectively.

If future shifts are known in advance for homogeneous counts then the λ value can be selected to minimise the EWMA plan's out-of-control ATS_0 as outlined in Section 2.

3.1. Efficient Plans for Nonhomogeneous Poisson Counts When Means Are Unknown

An adaptive EWMA plan using $\lambda = \lambda_{opt}(\mu_0, \mu_1)$ should work significantly better than guessing in advance a value for λ and using it unchanged. In practice, efficient estimates of $\mu_0(t)$ and $\mu_1(t)$ are required before an improvement is possible. In addition, we are less interested in detecting small shifts in mean ($\delta = \mu_1 - \mu_0$), for example, we may decide not to respond to any increase in mean of δ_{min} or less. Designing plans for $\mu_1(t) \geq \mu_0(t) + \delta_{min}$ has the advantage of delaying signals of increases in means less than δ_{min} . We achieve this by setting the lower bound for $\min(\mu_1(t)) = \mu_0(t) + \delta_{min}$, thus λ never fall below $\lambda_{opt}(\mu_0, \min(\mu_1(t)))$. In

this paper, we arbitrarily set $\min(\mu_1(t)) = \delta_{\min} = 0.35\sqrt{\mu_0(t)}$. Therefore, the application of the methodology involves the following steps.

- (1) Estimate $\mu_0(t)$ or $\mu_1(t)$ using one-day-ahead forecasts (details on how this is achieved under different situations are given in the next two subsections).
- (2) Use the estimates in step (1) to estimate $\lambda_{\text{opt}}(\mu_0, \mu_1)$, and in turn calculate $h(\mu_0(t), \lambda_{\text{opt}}(\mu_0, \mu_1), \text{ATS}_0)$.
- (3) Calculate AEWMA_t in (3.1) using the estimates of $\lambda = \lambda_{\text{opt}}(\mu_0, \mu_1)$ and $h(\cdot) = h(\mu_0(t), \lambda_{\text{opt}}(\mu_0, \mu_1), \text{ATS}_0)$ obtained in step (2).
- (4) Signal an unusual epidemic footprint whenever $\text{AEWMA}_t > a$ where a is the adjustment value discussed earlier.

Let the one-day-ahead model forecast for mean made at time $t - 1$ be denoted by $\hat{\mu}_{t/t-1}$. The next two subsections describe the improved EWMA design procedure for different purposes.

3.1.1. Detecting Unexpected Increases in Counts

The one-day-ahead forecast is used to remove the natural variation that comes from day of the week, holidays, seasonal and transitional influences, that is, $\mu_0(t)$ is estimated by $\hat{\mu}_{t/t-1}$. This process improves the detection of unusual outbreak data signals such as might be caused by bioterrorism or an epidemic footprint. The smoothed day ahead model forecast error ew_t , constrained to be positive (since we are only interested in shifts on the high side) is defined by

$$ew_0 = 0, \quad ew_t = \max(0, \theta(y_t - \hat{\mu}_{t/t-1}) + (1 - \theta)ew_{t-1}) \quad \text{for } t > 0, \quad (3.3)$$

where $0 < \theta < 1$ is a suitable smoothing weight. Note that ew_t estimates $\delta_t = \mu_1(t) - \mu_0(t)$. The estimate of $\mu_1(t) = \mu_0(t) + \delta_t$ is therefore

$$\hat{\mu}_1(t) = \max\left(\hat{\mu}_{t/t-1} + 0.35\sqrt{\hat{\mu}_{t/t-1}}, ew_{t-1} + \hat{\mu}_{t/t-1}\right). \quad (3.4)$$

The estimates of $\mu_0(t)$ and $\mu_1(t)$ given by $\hat{\mu}_{t/t-1}$ and $\hat{\mu}_1(t)$ above, respectively, are used in step (2) to estimate $\lambda_{\text{opt}}(\cdot)$ from (2.3).

3.1.2. Detecting Seasonal Epidemics Early

Suppose the aim of surveillance is identifying the start of a seasonal epidemic as early as possible, for example, the start of the annual influenza season. An epidemic footprint is expected each year, but the start time varies from season to season and varies in magnitude. Assume we have prior information that the seasonal epidemic event is rare in a certain period of the year (e.g., summer for influenza). Let the average (nearly homogenous) mean during the nonepidemic period be denoted by $\bar{\mu}_{\text{ne}}$ (used to estimate $\mu_0(t)$). Assume we are not interested in the early detection of any epidemic with a mean less than $\bar{\mu}_{\text{ne}} + c\sqrt{\bar{\mu}_{\text{ne}}}$ for some fixed $c > 0$. An improved adaptive EWMA can be constructed by estimating $\mu_0(t)$ by $\bar{\mu}_{\text{ne}}$ and

$\mu_1(t)$ by $\max(\bar{\mu}_{ne} + c\sqrt{\bar{\mu}_{ne}}, ew_{t-1} + \hat{\mu}_{t/t-1})$, and use these estimates in step (2) to estimate $\lambda_{opt}(\cdot)$ from (2.3). The resulting plan should be close to the “optimal” plan for detecting shifts from $\bar{\mu}_{ne}$ to $\mu_1(t)$. The closeness depends on how well $\hat{\mu}_{t/t-1} + ew_t$ estimates $\mu_1(t)$. In other words, if $\hat{\mu}_{t/t-1} + ew_t$, for all days t , are reasonable forecasts of the means during epidemics, then the improved plan should be efficient at detecting the epidemic footprint early.

3.2. Efficiency of the Plans

A simulation study is used to evaluate the practical efficiency of the plans. Tables 3 and 4 evaluate the average time to signal for two types of outbreak data: step changes and parabolic changes, respectively.

In-control background counts: were generated using means equal to $\mu_0(t)(= \mu_{tT})$ given by

$$\log(\mu_{tT}) = \alpha_{0\kappa(t)T} + \alpha_{1T} \cos(\omega t) + \alpha_{2T} \sin(\omega t) + \alpha_{3T} s_T + \sum_{j=1}^p \beta_{jT} \log(y_{t-j} + 1) \quad (3.5)$$

with the following models.

Model 1. $\alpha_{0iT} = 0.14$ for $i = 1, 2, \dots, 5$ (Monday to Friday), $\alpha_{06T} = 0.2$ (Saturday), $\alpha_{06T} = 0.22$ (Sunday), $\alpha_{1T} = 0.09$, $\alpha_{2T} = -0.09$, $\alpha_{3T} = 0$.

Model 2. $\alpha_{0iT} = 0.4$ for $i = 1, 2, \dots, 5$ (Monday to Friday), $\alpha_{06T} = 0.7$ (Saturday), $\alpha_{06T} = 0.8$ (Sunday), $\alpha_{1T} = 0.4$, $\alpha_{2T} = -0.4$, $\alpha_{3T} = 0$.

Model 3. $\alpha_{0iT} = 0.7$ for $i = 1, 2, \dots, 5$ (Monday to Friday), $\alpha_{06T} = 1$ (Saturday), $\alpha_{06T} = 1.1$ (Sunday), $\alpha_{1T} = 0.45$, $\alpha_{2T} = -0.45$, $\alpha_{3T} = 0$.

Model 4. $\alpha_{0iT} = 0.9$ for $i = 1, 2, \dots, 5$ (Monday to Friday), $\alpha_{06T} = 1.8$ (Saturday), $\alpha_{06T} = 2.1$ (Sunday), $\alpha_{1T} = 0.6$, $\alpha_{2T} = -0.6$, $\alpha_{3T} = 0$.

For Models 1–4, $\beta_{jT} = 0.1, 0.06, 0.02, 0.01$ for $j = 1, 2, 3, 4$, respectively, (i.e., $p = 4$). The value T in (3.5) relates to the moving window size in the model fitting process and therefore T can be ignored in the data simulation process. After generating the counts, these models are assumed hidden and refitted using the simulated data. The parameter estimates are assumed to move slowly by estimating them using a moving window of T observations. Therefore, no knowledge of the fixed regression parameters are assumed, and the parameters are assumed to change slowly over time with (say) process improvements, for example, improve data collection or slow drifts in population sizes. The moving window can be thought of as allowing the monitoring plan to adapt to changes in the process in the same way as control charts are updated as the process improves. The moving window relates more to good practice in updating the plan and is unrelated to the simulated data.

Simulated examples of unusual outbreak data: in-control counts were generated as above and used as training data to fit the model for the first day of the outbreak. An epidemic footprint was generated by adding counts simulated from a Poisson distribution with mean δ_t to the in-control generating counts on day t . The result is Poisson counts with means $\mu_0(t) + \delta_t$.

Table 3: Outbreak data signals involving a step change in means (a is the multiplicative adjustment needed from the theoretical threshold).

	Model number								
	2	3	4	2	3	4	2	3	4
	Poisson adaptive EWMA $\lambda = 0.1$			Poisson adaptive EWMA λ dynamic			Adaptive CUSUM of pseudo residuals		
θ (see (3.3))				0.075	0.1	0.1			
a	1.0075	1.0076	1.0144	1.0055	1.0050	1.0200	2.104	1.65	1.372
δ	ATS (Average time to signal/detection in days)								
0.0	102.38	100.96	97.23	100.21	104.70	99.50	102.90	100.09	100.11
0.5	25.22	29.21	35.11	24.75	29.89	35.32	44.31	47.89	52.95
1.0	12.59	15.34	19.85	12.44	15.40	20.49	26.33	28.99	33.71
1.5	7.79	9.72	13.21	7.73	9.73	13.41	17.71	19.71	24.01
2.0	5.50	6.99	9.83	5.40	6.90	9.35	12.87	14.19	17.83
2.5	4.25	5.46	6.98	4.20	5.34	6.69	9.85	11.11	13.66
3.0	3.46	4.40	5.26	3.37	4.31	5.16	7.80	8.61	10.95
3.5	2.88	3.65	4.38	2.84	3.58	4.19	6.61	7.21	9.07
4.0	2.49	3.06	3.73	2.47	3.05	3.60	5.70	6.15	7.57
5.0	2.03	2.46	2.89	2.01	2.45	2.68	4.46	4.72	5.60

Two outbreak data used for the results in Tables 3 and 4, respectively, are the following.

- (i) A step change of $\delta_t = \delta$ in mean was simulated to mimic increasing disease counts that result, for example, from health consequences to step changes in environmental pollution. The simulated epidemic footprint was continued until all plans signalled.
- (ii) A parabolic change in mean simulated to mimic say seasonal increases in say influenza counts. Here the mean increases of

$$\mu_1(t + \tau) - \mu_0(t + \tau) = \delta_\tau = \Delta(L - \tau) \times \tau, \tag{3.6}$$

where $\tau = 1, 2, \dots, L$, were generated as an unusual epidemic starting from day $t + 1$. The L indicates how long the epidemic lasts. L and Δ are selected to keep the epidemics of the same size, that is, Δ is varied so that $\sum_{\tau=0}^L \delta_\tau$ is constant for all selections of L making all outbreaks roughly equally important. We use

$$(L, \Delta) = (8, 0.5), (15, 0.075), (22, 0.023715), (29, 0.010345), (36, 0.0054055), (43, 0.00317). \tag{3.7}$$

A signal could only be given for its epidemic footprint duration, and therefore, it could be missed by the monitoring plans. Therefore the proportion of missed detections is recorded (see Table 4).

The model in (3.5) was fitted using the first three years of simulated data to deliver one-day-ahead forecast in-control (usual) behaviour (expected counts). Thereafter forecasts were made using the fitted model for a moving window of $T = 1095$ days of data. The fitted values were used to select λ values that delivered smaller ATS_δ values for the “unknown”

Table 4: Seasonal outbreak data involving a parabolic change in means (a is the multiplicative adjustment needed from the theoretical threshold).

	Model number					
	1	2	3	1	2	3
	Dynamic Poisson adaptive EWMA			Poisson adaptive EWMA		
θ in (3.3)				0.075	0.1	0.1
a	1.0076	1.0075	1.0076	1.0056	1.0055	1.0050
$(L, j\Delta)$	$j = 0.5$	$j = 0.5$	$j = 1.0$	$j = 0.5$	$j = 0.5$	$j = 1.0$
	Overall average number of extra victims in the outbreak data					
$\sum et = 1^L \delta_{et}$	42	42	84	42	42	84
	ATS (Pr(no outbreak data signal))					
(0,0)	102.38 (0.00)	100.96 (0.00)	97.23 (0.00)	99.50 (0.00)	100.19 (0.00)	103.03 (0.00)
$(43, j \times 0.00317)$	22.05 (0.16)	14.73 (0.03)	16.11 (0.07)	18.56 (0.14)	15.31 (0.03)	17.00 (0.07)
$(36, j \times 0.0054055)$	18.93 (0.14)	12.36 (0.02)	13.57 (0.06)	15.72 (0.14)	12.59 (0.02)	14.17 (0.06)
$(29, j \times 0.01035)$	14.62 (0.12)	9.66 (0.01)	10.88 (0.03)	12.85 (0.13)	10.37 (0.01)	11.61 (0.03)
$(22, j \times 0.023715)$	11.20 (0.11)	7.54 (0.00)	8.43 (0.03)	9.81 (0.11)	7.60 (0.00)	8.65 (0.03)
$(15, j \times 0.075)$	7.44 (0.10)	4.75 (0.00)	5.03 (0.02)	6.66 (0.08)	4.93 (0.00)	5.69 (0.02)
$(8, j \times 0.5)$	3.50 (0.04)	2.40 (0.00)	2.81 (0.00)	3.33 (0.05)	2.43 (0.00)	2.83 (0.00)

simulated epidemic footprints. In Table 3, new data were recursively simulated until the outbreak footprint was signalled, and the time to signal (in days) was recorded. This same process is repeated 10 000 times to give the average time to signal (i.e., ATS_δ or ATS_0). Tables 3 and 4 compare the improved adaptive EWMA to the adaptive EWMA with $\lambda = 0.1$ and the adaptive CUSUM plan (Sparks [7] and Shu and Jiang [20]) applied to pseudo-residuals given by

$$\Phi^{-1} \left(\sum_{i=0}^{y_t} \exp_{(-u_i)} \frac{(-\mu_i)(\mu_i)^i}{i!} \right), \quad \text{where } \Phi(a) = \int_{-\infty}^a \frac{e^{-z^2/2}}{\sqrt{2\pi}} dz. \quad (3.8)$$

For each simulation run different three thousand Poisson counts were generated using the above model. The first 1095 days (3 years) were used as training data to fit the model that provides a forecast for day 1096, and thereafter a moving window of 1095 days of data fitted the model to give the day ahead forecasts. Additional counts for the epidemic were added to the counts starting from day 1096 onwards (going for L days for the parabolic increase in means). The number of days from 1095 to the plan’s first out-of-control signal is taken as the time to detection (in days). Repeat the process 10 000 times to estimate of ATS_δ . The ATS_δ values are reported in Tables 3 and 4 for the step change and parabolic change, respectively. The plans in Tables 3 and 4 all are designed to deliver in-control ATS_0 of 100, so plans are comparable. The CUSUM of the pseudoforecast errors has the multiplier adjustment needed to deliver a plan with $ATS_0 \approx 100$, a , not close to one, thus indicating the normal approximation of pseudo residuals is inadequate to use EWMA plans based on normality without adjustment. However, more importantly a is not approximately constant across all models, which means that

- (1) if the conventional CUSUM is applied to the pseudoforecast errors (without adjustment), the plan will deliver too many false alarms,
- (2) the threshold needs an adjustment for each change in the mean distribution.

On the other hand, the Poisson adaptive EWMA plans need very little adjustments for the design control limit of one (a ranging from 1.0075 to 1.02). From the results in Table 3, we conclude that the improved Poisson adaptive EWMA is always, on average, more efficient than the Poisson adaptive EWMA with $\lambda = 0.1$ for step changes. For Models 2 in Table 3, improved Poisson adaptive EWMA plans have smaller ATS_δ values for all $\delta > 0$. For Models 3 and 4 in Table 3, improved Poisson adaptive EWMA plans have smaller ATS_δ values for all $\delta \geq 2$. In Table 4, the improved Poisson adaptive EWMA's advantage only holds for Models 2 and 3. In addition, for Models 2 and 3, the improved Poisson adaptive EWMA is likely to miss fewer parabolic unusual changes (not evident from Table 4 because the gains are small). For low counts (Model 1), there is generally insufficient information to improve the λ exponential weights, and so we recommend using $\lambda = 0.1$ for low counts.

In summary, the dynamic adaptive Poisson EWMA plan signals, on average, earlier than the adaptive Poisson EWMA plans with $\lambda = 0.1$ for most step changes in mean of 2 or more (Table 3), while for parabolic changes the results are mixed. This knowledge allows users to revise the dynamic plan by setting $\lambda_{\text{opt}}(\mu_0, \mu_1) = 0.1$ when means trend lower. Notice that the Poisson adaptive EWMA signals step changes on average earlier than the "optimal" adaptive CUSUM of the pseudoforecast residuals. The similar results hold for parabolic changes but these results are not reported. Therefore, the "optimal" adaptive CUSUM of the pseudoforecast residuals plan is not recommended based on these simulation results.

4. An Application

The application involves influenza A daily counts in New South Wales (NSW), Australia ranging from the end of 2003 to the end of 2005. Temporal trends in the daily disease counts are found in the first plot of Figures 1 and 2. Figures 1 and 2 examine the detection of unusual epidemic footprints and early detection of the start of the flu season, respectively. The second plot in Figures 1 and 2 is the "optimal" adaptive EWMA plan ($AEWMA_t$), and the last plot is a time series of the "optimal" CUSUM of the one-day-ahead pseudoforecast errors. These are compared in terms of early detection.

The adaptive plan is designed to give a false alarm every 100 days under the assumption that counts are Poisson distributed with the means given by their one-day-ahead forecasts. We consider models of the form of (3.5) except that a school holiday factor was included and the number of lags in Poisson transition regression model (3.5) was taken as 4 (the transitional model that gave the smallest AIC value). Model (3.5) is used to forecast the mean for the day ahead. Similar approaches to modelling disease counts can be found in Burr et al. [6], Bjornstad et al. [21], and Christensen et al. [22]. Moving windows of one, two, and three years of data were tried. In addition, various discounting weights, for example, 0.996^t , giving more recent observations greater weight in determining regression parameters were tried. The forecasting plan that produced the smallest forecast errors was (see [23])

- (i) a moving window of 730 days of data (2 years),
- (ii) discounted weights 0.996^t as the number of days t moves away from the day being forecasted ($t = 1, \dots, 730$).

The autocorrelations of standardised forecast errors for influenza were insignificant (see Figure 3). Therefore applying standard control assuming independence is reasonable provided that the normality assumption is not violated.

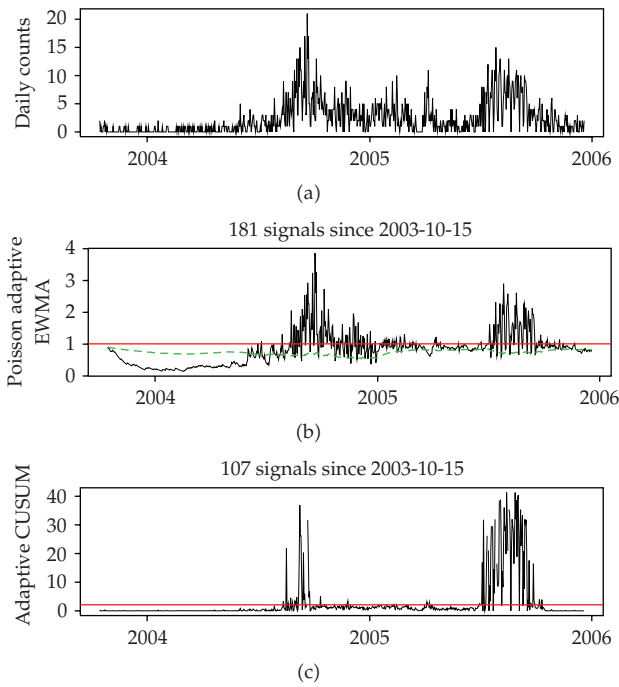


Figure 1: EWMA plans optimised for detecting seasonal epidemic footprints applied to influenza A daily counts in NSW.

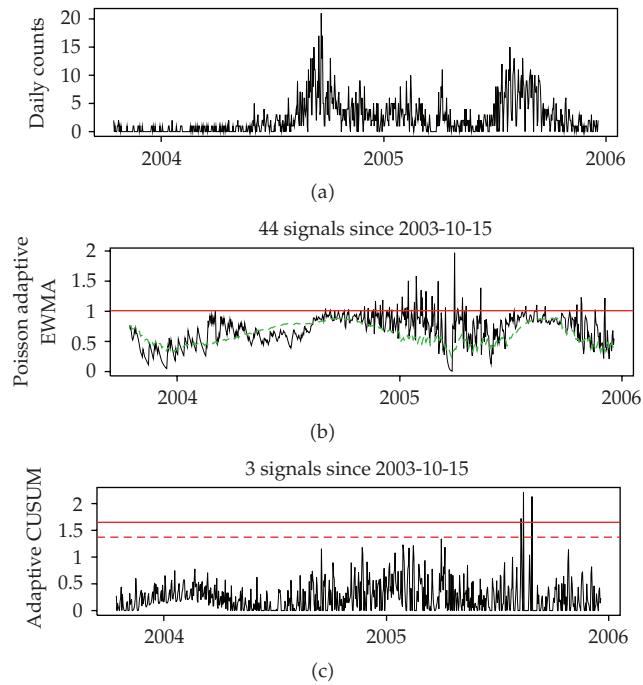


Figure 2: EWMA plans optimised for detecting unexpectedly high influenza A counts in NSW.




Figure 3: Autocorrelation function of the model standardised residuals for influenza A daily counts in NSW.

The forecasts, together with the smoothed forecast errors, are used to optimise the adaptive EWMA plan as outlined previously. After each signal, EWMA plans reset the EWMA statistic to its initial value, and therefore follow-up signals indicate that either the epidemic footprint persisted, or that the counts continued to be higher than expected. The “optimal” adaptive CUSUM plans for the one-day-ahead pseudoforecast errors z_t used in Tables 3 and 4 were compared to the dynamic adaptive Poisson EWMA plan.

4.1. Early Seasonal Epidemic Footprint Detection

The first plot in Figure 1 indicates daily counts for influenza A in NSW, Australia. We assumed that a seasonal increase in influenza counts can start any time from beginning of May to the end of October. The model is fitted using a moving window of two years (730 days). As such, control limits adjust, for example, as either disease control improves or the population size changes. This adjustment is similar to revising the control limits as the process improves in manufacturing. This process helps signal small epidemic footprints after counts reduce due to improved policy or control.

The second plot in Figure 1 presents the Poisson adaptive plan with dynamic exponential weights designed for earlier detection of epidemic footprints. The dashed line in Figure 1 indicates the expected trend for the dynamic Poisson adaptive EWMA. The dynamic Poisson adaptive EWMA line trends above the dashed line when the counts are trending higher than expected.

Two epidemic footprints were detected in each of years 2004 and 2005 (see Figure 1 with time on the x -axis). Note that the Poisson adaptive EWMA flags the first epidemic footprint in 2004 in mid-July 2004 (second plot in Figure 1), much earlier than the “optimal” adaptive CUSUM of the pseudoforecast residuals plan (second plot in Figure 1). However, the second epidemic footprint is signalled a few days earlier by the “optimal” adaptive CUSUM. The major difficulty with the “optimal” adaptive CUSUM is what adjustment (a)

to use for nonnormality. The value $a = 2.104$ (see Table 3) was selected because it was appropriate for early detection of the first epidemic footprint, but it was not appropriate for the second. This point highlights the difficulty with using the “optimal” adaptive CUSUM in practical circumstances. The Poisson adaptive EWMA plan, on the other hand, is very robust to changes in the process mean.

4.2. Detection of Unusual Counts

The second plot in Figure 2 presents the dynamic adaptive plans for early detection of unusual high influenza A incidence after correcting for usual seasonal, day of the week, school holiday, and transitional influences. The dynamic Poisson adaptive EWMA plan signals unusual increase in counts many more times than the “optimal” adaptive CUSUM plan independent of adjustment of normality (both $a = 1.372$ and $a = 1.65$ are recorded in Figure 2). Again this highlights the difficulty with applying the adaptive CUSUM plan to pseudo residuals because of the difficulty knowing what adjustment to use.

A laboratory diagnosing patients as having influenza A was found to produce many false positives in early 2005. Poisson adaptive EWMA plan signalled these whereas the “optimal” adaptive CUSUM plan fails to signal (see southern hemisphere summer of 2004/2005 in Figure 2).

Further analysis (not reported) demonstrated that the Poisson adaptive EWMA plan for $\lambda = 0.1$ misses the first unusual increase in Figure 2, thereafter the performance of both Poisson adaptive EWMA plans was almost identical. The dynamic Poisson adaptive EWMA plan has an advantage over the Poisson adaptive EWMA plan with $\lambda = 0.1$ in this example.

5. Discussion

In summary, we offer a dynamic Poisson adaptive EWMA plan for early detection of epidemic footprints when disease counts follow a Poisson regression model. Results demonstrate that in the case of step changes in means, the dynamic Poisson adaptive EWMA detects outbreaks earlier than the CUSUM plan of pseudoforecast errors. Results for parabolic changes demonstrate the superiority of the dynamic Poisson adaptive EWMA over adaptive EWMA with $\lambda = 0.1$ for only counts with means not too low. We recommend the adaptive EWMA with $\lambda = 0.1$ when counts are low (e.g., average means less than 2).

If the in-control mean changes dramatically from season to season, then the adaptive CUSUM of pseudo residuals plan threshold adjust would need revising from season to season. This aspect makes the implementation of the adaptive CUSUM of pseudo residuals plan more difficult.

Most of the effort has been placed on sequential hypothesis testing, but the performance of plans is also influenced by recursive estimation of model parameters. More effort in modelling would improve the performance of the Poisson adaptive EWMA plan. There is also value in monitoring model lack of fit for each moving window of data (see Zeileis [24, 25] and Zeileis and Kleiber [26]).

Acknowledgments

The authors are grateful to Dr. Petra Graham for writing R programs for fitting the models in the application. In addition, the authors thank the referees for their valuable comments that have improved the presentation of this paper.

References

- [1] W. H. Woodall, "The use of control charts in health-care and public-health surveillance," *Journal of Quality Technology*, vol. 38, no. 2, pp. 89–104, 2006.
- [2] D. L. Buckeridge, H. Burkom, M. Campbell, W. R. Hogan, and A. W. Moore, "Algorithms for rapid outbreak detection: a research synthesis," *Journal of Biomedical Informatics*, vol. 38, no. 2, pp. 99–113, 2005.
- [3] J. M. Lucas and M. S. Saccucci, "Exponentially weighted moving average control schemes: properties and enhancements. With discussion," *Technometrics*, vol. 32, no. 1, pp. 1–29, 1990.
- [4] D. Lambert and C. Liu, "Adaptive thresholds: monitoring streams of network counts," *Journal of the American Statistical Association*, vol. 101, no. 473, pp. 78–88, 2006.
- [5] W. Zucchini and I. L. MacDonald, "Illustrations of the use of pseudo-residuals in assessing the fit of a model," in *Proceedings of the 14th International Workshop on Statistical Modelling*, H. Friedl, A. Berghold, and G. Kauermann, Eds., pp. 409–416, Graz, Austria, 1999.
- [6] T. Burr, T. Graves, R. Klamann, S. Michalak, R. Picard, and N. Hengartner, "Accounting for seasonal patterns in syndromic surveillance data for outbreak detection," *BMC Medical Informatics and Decision Making*, vol. 6, 2006.
- [7] R. S. Sparks, "CUSUM charts for signalling varying location shifts," *Journal of Quality Technology*, vol. 32, no. 2, pp. 157–171, 2000.
- [8] O. Grigg and D. Spiegelhalter, "A simple risk-adjusted exponentially weighted moving average," *Journal of the American Statistical Association*, vol. 102, no. 477, pp. 140–152, 2007.
- [9] G. Capizzi and G. Masarotto, "An adaptive exponentially weighted moving average control chart," *Technometrics*, vol. 45, no. 3, pp. 199–207, 2003.
- [10] H. B. Nembhard and M. S. Kao, "Adaptive forecast-based monitoring for dynamic systems," *Technometrics*, vol. 45, no. 3, pp. 208–219, 2003.
- [11] R. A. Parker, "Analysis of surveillance data with Poisson regression: a case study," *Statistics in Medicine*, vol. 8, no. 3, pp. 285–294, 1989.
- [12] B. Miller, H. Kassenborg, W. Dunsmuir, et al., "Syndromic surveillance for influenzalike illness in an ambulatory care network," *Emerging Infectious Diseases*, vol. 10, no. 10, pp. 1806–1811, 2004.
- [13] I. Kleinschmidt, B. L. Sharp, G. P. Y. Clarke, B. Curtis, and C. Fraser, "Use of generalized linear mixed models in the spatial analysis of small-area malaria incidence rates in KwaZulu Natal, South Africa," *American Journal of Epidemiology*, vol. 153, no. 12, pp. 1213–1221, 2001.
- [14] L. Wang, M. F. Ramoni, K. D. Mandl, and P. Sebastiani, "Factors affecting automated syndromic surveillance," *Artificial Intelligence in Medicine*, vol. 34, no. 3, pp. 269–278, 2005.
- [15] C. P. Farrington, N. J. Andrews, A. D. Beale, and M. A. Catchpole, "A statistical algorithm for the early detection of outbreaks of infectious disease," *Journal of the Royal Statistical Society, Series A*, vol. 159, no. 3, pp. 547–563, 1996.
- [16] B. J. Cowling, I. O. L. Wong, L.-M. Ho, S. Riley, and G. M. Leung, "Methods for monitoring influenza surveillance data," *International Journal of Epidemiology*, vol. 35, no. 5, pp. 1314–1321, 2006.
- [17] M. A. Stoto, M. Schonlau, and L. T. Mariano, "Syndromic surveillance: is it worth the effort?" *Chance*, vol. 17, no. 1, pp. 19–24, 2004.
- [18] A.-V. Kaninda, F. Belanger, R. Lewis, et al., "Effectiveness of incidence thresholds for detection and control of meningococcal meningitis epidemics in northern Togo," *International Journal of Epidemiology*, vol. 29, no. 5, pp. 933–940, 2000.
- [19] R. S. Sparks, C. Carter, P. L. Graham, et al., "A strategy for understanding the sources of variation in syndromic surveillance for bioterrorism and public health incidence," to appear in *IIE Transactions on Healthcare Systems Engineering*.
- [20] L. Shu and W. Jiang, "A Markov chain model for the adaptive CUSUM control chart," *Journal of Quality Technology*, vol. 38, no. 2, pp. 135–147, 2006.
- [21] O. N. Bjornstad, B. F. Finkenstadt, and B. T. Grenfell, "Dynamics of measles epidemics: estimating scaling of transmission rates using a time series SIR model," *Ecological Monographs*, vol. 72, no. 2, pp. 169–184, 2002.
- [22] A. Christensen, H. Melgaard, J. Iwersen, and P. Thyregod, "Environmental monitoring based on a hierarchical Poisson-Gamma model," *Journal of Quality Technology*, vol. 35, no. 3, pp. 275–285, 2003.

- [23] P. Graham and R. S. Sparks, "Adaptation of generalized linear models to diseases that occur infrequently (mean rate 1–15 counts/day)," CSIRO Technical Paper CMIS 06/34, 2006.
- [24] A. Zeileis, "A unified approach to structural change tests based on ML scores, F statistics, and OLS residuals," *Econometric Reviews*, vol. 24, no. 4, pp. 445–466, 2005.
- [25] A. Zeileis, "Implementing a class of structural change tests: an econometric computing approach," *Computational Statistics & Data Analysis*, vol. 50, no. 11, pp. 2987–3008, 2006.
- [26] A. Zeileis and C. Kleiber, "Validating multiple structural change models—a case study," *Journal of Applied Econometrics*, vol. 20, no. 5, pp. 685–690, 2005.