



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

INSTITUT FÜR STATISTIK



Michael Höhle, Michaela Paul and Leonhard Held

Statistical approaches to the surveillance of infectious diseases for veterinary public health

Technical Report Number 014, 2007
Department of Statistics
University of Munich

<http://www.stat.uni-muenchen.de>



Statistical approaches to the surveillance of infectious diseases for veterinary public health

Michael Höhle^{a,*} Michaela Paul^b Leonhard Held^b

^a*Department of Statistics, University of Munich, Germany*

^b*Biostatistics Unit, Institute of Social and Preventive Medicine, University of Zurich, Switzerland*

Abstract

This paper covers the aspect of using statistical methodology for the monitoring of routinely collected surveillance data in veterinary public health. An account of the Farrington algorithm and Poisson cumulative sum schemes for the detection of aberrations is given with special attention devoted to the occurrence of seasonality and spatial aggregation of the time series. Modelling approaches for retrospective analysis of surveillance counts are described. To illustrate the applicability of the methodology in veterinary public health, data from the surveillance of rabies among fox in Hesse, Germany, are analysed.

Key words: Veterinary public health, outbreak detection, rabies, GIS

1 Introduction

Understanding and controlling infectious diseases is a problem that needs to be addressed by a collaboration of scientists in human and veterinary medicine, mathematical modelling, economy and social sciences. This paper covers the aspect of using statistical methodology for the monitoring of routinely collected surveillance data, which has become increasingly important as the

* Dept. of Statistics, University of Munich, Ludwigstr. 33, 80539 Munich, Germany
Email address: hoehle@stat.uni-muenchen.de (Michael Höhle).

amount of data gathered through automatic surveillance increases. Statistical surveillance methods have predominately originated from applications in human public health surveillance, but they find increased utilization in veterinary public health. The aim of this paper is to explain concepts behind prospective and retrospective statistical surveillance and illustrate use and potential in veterinary epidemiology.

Examples from human epidemiology include the monitoring of notifiable diseases, congenital malformations, surgical outcomes and bioterrorism syndromes (Widdowson et al., 2003; Chen, 1978; Steiner et al., 2000; Bravata et al., 2004). In veterinary epidemiology applications often correspond to the human equivalent, e.g. Kosmider et al. (2006) monitor salmonella in livestock reports and Carpenter et al. (2007) look at abortions in dairy cattle. One issue in adapting surveillance from humans to animals is the fundamental differences in terms: Animal surveillance has to deal with diverse species and completely different conditions of living (e.g. production, wild, or companion animal) resulting in different entities of interest (e.g. individual or herd) and different professionals interacting with the population. As a consequence the possibility and cost of investigation and control strategy depends heavily on character of living. However, as much as the human and veterinary surveillance differ, zoonoses like salmonellosis, rabies or emerging zoonoses (e.g. avian flu) underline the need for a comparative and co-operative approach in surveillance.

Data quality is a major practical concern in the analysis of surveillance data, e.g. lack of clear case definition, under-reporting or reporting delays and lack of denominator data complicate the statistical analysis. In this paper the focus is however on the statistical challenges of analysing the resulting univariate and multivariate time series containing daily, weekly or monthly counts.

The first part of the article deals with prospective surveillance for count data. One important question in this process is deciding on the best level of aggregation – to this end a simple scheme for hierarchical disease detection is introduced. Our presentation thus emphasizes the time series nature of the data as an alternative to spatial and spatio-temporal cluster detection methods, e.g. scan statistics (Kulldorff, 2001; Rogerson, 2001). The second part describes a stochastic model for the analysis of multivariate surveillance data. This model can be used to detect temporal and spatio-temporal dependencies in multivariate time series of surveillance counts. Applicability of the presented methods is illustrated through data from the monitoring database of the WHO rabies surveillance program (WHO Collaboration Centre for Rabies Surveillance and Research, 2007).

2 Prospective surveillance

In this section, statistical methods for univariate and hierarchical disease surveillance are discussed with a focus on outbreak detection for count data with seasonality. Broader surveys of outbreak detection methods can be found in Farrington and Andrews (2003); Sonesson and Bock (2003); Lawson and Kleinman (2005); Buckeridge et al. (2005).

2.1 Univariate surveillance

For many surveillance problems univariate time series of counts are readily available. If not some preprocessing is performed, e.g. by aggregating geo-referenced outbreak data to an appropriate level or by time-wise aggregation of event time data. We denote the resulting univariate time series by $\{y_t; t = 1, 2, \dots\}$. Online outbreak detection can be seen as a classification task: based on the observed values y_1, \dots, y_n it is to be decided if there is an aberration at time n or not. In what follows two classes of methods that address the problem are described.

2.1.1 Farrington Method

Core of the method by Farrington et al. (1996) is to predict the observed value y_n using a set of reference values taken from the observed values y_1, \dots, y_{n-1} . To handle long term trends and seasonality, only values from a window of size $2w + 1$ around time n up to b years back in time are taken. Thus, the set of reference values consists of recent values with similar conditions as at time n and can formally be defined as

$$R(w, b) = \left(\bigcup_{i=1}^b \bigcup_{j=-w}^w y_{n-i \cdot r + j} \right),$$

where r is the period of the observations, e.g. for monthly data r is 12. Thus no observations from the current year are used. Poisson regression with overdispersion is then used to model the $(2w+1)b$ reference values, i.e. for $y_t \in R(w, b)$

$$E(y_t) = \mu_t, \quad \text{with} \quad \log \mu_t = \alpha + \beta t \quad \text{and} \quad \text{Var}(y_t) = \phi \mu_t.$$

Based on the estimated model a one-sided $(1 - \kappa) \cdot 100\%$ prediction interval for y_n can be formed. The classical way to compute such a prediction interval is based on the normal distribution, however, as the skewness of the Poisson distribution with mean μ is $1/\sqrt{\mu}$, for low valued μ this is a bad approximation. Therefore a $\frac{2}{3}$ -power transformation is applied to normalize the distribution

before computing the interval. The resulting back-transformed upper limit of the prediction interval for y_n is then

$$U_n = \hat{\mu}_n \left\{ 1 + \frac{2}{3} z_{1-\kappa} \cdot \sqrt{\frac{\hat{\phi} \hat{\mu}_n + \text{Var}(\hat{\mu}_n)}{\hat{\mu}_n^2}} \right\}^{3/2},$$

where $\hat{\mu}_n = \exp(\hat{\alpha} + n\hat{\beta})$ and $z_{1-\kappa}$ is the $100(1 - \kappa)\%$ quantile of the standard normal distribution. Subsequently, if $y_n > U_n$ an alarm is sounded. To ease exposition some details of the algorithm have been left out in the above description; e.g. the linear trend is only included if it is significant at the 5% level and a second round of estimation is performed with observations weighted by their inverse residuals. The latter corrects for possible past outbreaks in the reference values. Furthermore, protection against preposterous alarms is made by post-processing alarms and only reporting those where enough cases have been seen.

Big virtue of the Farrington procedure is its simple yet flexible modelling depending on only one user specified parameter κ , which is of great advantage when applying the method to multiple surveillance time series. One shortcoming is that only a moving window of historical values is taken for estimation with no values taken from the current year. A simple extension would be to include seasonal terms into the linear predictor of the Poisson regression

$$\log \mu_t = \alpha + \beta t + \sum_{s=1}^S \left[\gamma_s \sin\left(\frac{2\pi}{r} s \cdot t\right) + \delta_s \cos\left(\frac{2\pi}{r} s \cdot t\right) \right], \quad (1)$$

where r is a known period (e.g. 12 for monthly data). Reference values could then consist of historical values or all values within a moving window of b years. However, sequential estimation becomes more complicated and careful selection of the parameter S is required. Only when a few series, where performance is vital, are of interest such modelling becomes advantageous.

Another shortcoming is the recomputing of prediction-intervals at every time instance. As a consequence, sustained shifts become hard to detect as deviations are not accumulated. Performance on such shifts can be improved by exploiting methods from statistical process control (SPC), which accumulate information. The cumulative sum (CUSUM) scheme described in the following section is such a method.

2.1.2 CUSUM likelihood ratio detectors

Statistical process control has its root in the quality control of manufactured goods, but has since found numerous application in health care (Woodall, 2006) by providing a more formal setting for surveillance methods. In our

treatment we shall put special emphasis on the count data character of the time series – a feature which is less typical for SPC methods.

Central in SPC is the online detection of change-points. With n observations and known change-point τ one assumes the following model:

$$y_t|z_t, \tau \sim \begin{cases} f_{\theta_0}(\cdot|z_t) & \text{for } t = 1, \dots, \tau - 1 \text{ (in-control)} \\ f_{\theta_1}(\cdot|z_t) & \text{for } t = \tau, \tau + 1, \dots \text{ (out-of-control)} \end{cases}$$

where z_t denotes known covariates at time t and f_θ is e.g. the Poisson probability function with its mean μ being a function of θ and z_t . Objective of online change-point detection is to use the observations y_1, \dots, y_n to decide at time n whether a change-point has occurred during $1, \dots, n$. One approach to do this is to use the so called cumulative sum (CUSUM) likelihood ratio detector (Lai, 1995; Frisén, 2003)

$$N = \min \left\{ n \geq 1 : \max_{1 \leq \tau \leq n} \left[\sum_{t=\tau}^n \log \left\{ \frac{f_{\theta_1}(y_t|z_t)}{f_{\theta_0}(y_t|z_t)} \right\} \right] \geq c \right\}. \quad (2)$$

For a specific n the above detector computes the log likelihood ratio (LR) statistic for testing the hypothesis that all observations originate from the in-control distribution against the alternative that they from change-point τ on stem from the out-of-control distribution. Maximizing the LR statistic for each possible change-point $1 \leq \tau \leq n$ means finding the maximum likelihood estimator, $\hat{\tau}$, for the most likely location of the change-point. If $\text{LR}(\hat{\tau})$ is above a pre-specified threshold c , then there is enough information at time n to say that a change-point happened at $\hat{\tau}$. Otherwise, no decision is made and the monitoring continues at time $n + 1$.

Without covariates and pre-specified θ_0 and θ_1 the above detector can be written in the well known CUSUM recursive form

$$l_0 = 0, \quad l_n = \max \left(0, l_{n-1} + \log \left\{ \frac{f_{\theta_1}(y_n)}{f_{\theta_0}(y_n)} \right\} \right), \quad n \geq 1 \quad (3)$$

where the first alarm is given at time $N = \min\{n : l_n \geq c\}$. This detector can be shown to be optimal (in some technical sense) for the detection of a shift from θ_0 to θ_1 . Evaluating the performance of the proposed schemes is a question about which performance criterion to consider. Typical choices are location parameters of the run length distribution, e.g. the average run lengths (ARLs) $ARL_0 = E(N|\tau = \infty)$, i.e. expected waiting time until the first false alarm, or $ARL_1 = E(N|\tau = 0)$, i.e. expected time until detection of the change, when this change occurs immediately. Alternatives include the conditional expected delay $E(N - \tau|\tau, N \geq \tau)$ or the probability of false alarm within the first m time points $P(N \leq m|\tau = \infty)$. Frisén (2003) gives a thorough treatment of the various criteria and available optimality results.

One can also use the classification framework to discuss performance using sensitivities, specificities and ROC curves (Kleinmann and Abrams, 2006).

Lucas (1985) covers the CUSUM likelihood ratio method for the Poisson distribution with constant parameters $\mu_0 = \theta_0$ and $\mu_1 = \theta_1$. In a surveillance context main interest is in detecting upward changes, thus typically $\theta_1 > \theta_0$. Dividing by an appropriate constant in (3) one obtains for this upward detection the equivalent form $S_n = \max(0, S_{n-1} + (y_n - k))$, where $k = (\mu_1 - \mu_0)/(\log(\mu_1) - \log(\mu_0))$ and $S_0 = 0$. A common way to select μ_0 is to use a period known to be in-control to estimate μ_0 . In addition, μ_1 is the change, which is to be detected quickly. Note that if the in-control assumption in the μ_0 estimation is violated, e.g. if the training data contain outbreaks, an incorrect in-control parameter is estimated.

When the in-control mean is time varying due to long term and short term trends (such as seasonality) matters become more complicated. Two methods operating with time-varying parameters for count data are the method by Rossi et al. (1999) and Rogerson and Yamada (2004a). Letting $\mu_{0,t}$ be the time varying in-control mean the first suggests a transformation to normality by looking at

$$x_t = \frac{y_t - 3\mu_{0,t} + 2\sqrt{\mu_{0,t} \cdot y_t}}{2\sqrt{\mu_{0,t}}} \quad (4)$$

and applying a Gaussian CUSUM to these transformed values. However, simulation studies show that in case of low-counts the resulting ARLs are far away from the anticipated ARLs as computed for the Gaussian CUSUM (Rogerson and Yamada, 2004a; Höhle and Paul, 2007). This lead Rogerson and Yamada (2004a) to propose time-varying control parameters of the Poisson CUSUM in order to obtain a specific in-control ARL_0 value of γ . Their CUSUM is

$$S_n = \max(0, S_{n-1} + h_n(x_n - k_n)), \quad \text{with} \quad k_n = \frac{\mu_{1,n} - \mu_{0,n}}{\log(\mu_{1,n}) - \log(\mu_{0,n})},$$

and $\mu_{1,n}$ being a multiple of standard deviations larger than $\mu_{0,n}$. The factor $h_n = c/c_n$ scales the contribution of $(x_n - k_n)$ at each time point. The threshold c_n is determined at each time point as the threshold of a time-constant Poisson CUSUM with reference value k_n having ARL_0 equal to γ , which can be computed e.g. by the algorithm of Hawkins (1992). Finally, c is the threshold of an ordinary Poisson CUSUM with constant in-control mean parameter μ_0 , e.g. selected as the mean of the training period observations. An alarm is given if $S_n \geq c$.

As an alternative Höhle and Paul (2007) suggest using (2) directly, which makes allowance for quite flexible models for $\mu_{1,t}$ at the cost of losing the recursive computation in (3) and having to compute performance criterion by Monte-Carlo simulation. Big advantage of the CUSUM methods is their ability for optimal detection of change-points from $\mu_{0,t}$ to $\mu_{1,t}$, however relying on these models being adequate. Care thus has to be exercised in order to specify




Fig. 1. Hierarchical time series structure. Each time series $y_{X,t}$ is formed by aggregating its immediate descendants in the graph, e.g. $y_t = \sum_{i=1}^{l_0} y_{i,t}$.

a reasonable in-control model such as (1). When monitoring massive amounts of time series such care is not always possible. Also, for extremely low counts it can be beneficial to monitor the number of zero count periods before a period with one or more counts by the geometric distribution (Bourke, 1992).

2.2 Multivariate and hierarchical surveillance

Practical surveillance typically means monitoring multiple time series simultaneously, e.g. series for different diseases and serotypes, age groups or distinct geographical regions. A naive way to perform multivariate monitoring of m time series $\{y_{i,t}; i = 1, \dots, m\}$ is to apply the univariate method of choice to each time series separately and report all alarms. However, this approach ignores any correlations between the time series and thus leads to inferior detection. Multivariate change-point detection methods such as multivariate CUSUMs take correlations into account, but they are only developed for the continuous case – Rogerson and Yamada (2004b) investigate one of several proposals for a multivariate CUSUM in a surveillance context. However, the more detailed the partition due to serotype, age or region, the rarer the cases. Continuous approximations are thus unsound, but methods for handling correlated multivariate count data time series, e.g. Tourneret et al. (2002), are still scarce. As a consequence, operating with multiple univariate detectors is still the pragmatic choice.

One issue also not dealt with in the literature is the important question of choosing the appropriate level of aggregation for the surveillance. Figure 1 shows a hierarchy of time series resulting from aggregation of three levels, e.g. spatial aggregation with top level being the total aggregated number of cases in a country and lower levels representing the series for administrative regions and districts.

The simple, yet very effective idea is now to monitor all m time series independently by univariate methods. To keep in-control alarm rates comparable between levels one would use higher thresholds at the lower levels. A subsequent plot showing all generated alarms provides a good overview of aggregation




Fig. 2. Monthly number of rabies cases in Hesse, Germany.

effects and outbreak sizes.

2.3 Results

As an example of surveillance in veterinary public health, Fig. 2 shows the 1985-2006 time series of monthly incidence of rabies among fox in the federal state of Hesse, Germany. These data are part of the monitoring database kept by the Collaboration Centre for Rabies Surveillance and Research (WHO Collaboration Centre for Rabies Surveillance and Research, 2007).

A drastic decrease in the number of cases is seen as a consequence of the oral rabies vaccination program started in 1985 using Tübingen baits. However, several set-backs for the boarder region between Hesse and Bavaria occurred as mentioned in Müller et al. (2005) and shown in the following analyses. To illustrate seasonality of the time series we proceed as in Harnos et al. (2006) and divide the monthly cases by the respective yearly average and compute monthly means of this detrended time series as shown in Fig. 3. Strong seasonality of the rabies data can be seen – the increased incidence of up to 1.5 times the yearly average in spring and autumn corresponds to the mating season and dispersal of young foxes (Thulke et al., 2000).

Figure 4 shows the result of online surveillance for the Hesse time series, beginning from January 1998. To obtain the time varying $\mu_{0,t}$ values for each time point t , a seasonal Poisson model as in (1) is fitted to the observed values from January 1985 up to December 1997. As observations close to t are considered more informative, a weighted Poisson regression is performed, where each observation is weighted according to its distance to t in time. This thusly fitted model is then used to compute the predicted in-control mean for




Fig. 3. Seasonality in the rabies cases illustrated by the monthly average relative to yearly average for each month.




Fig. 4. Surveillance of the Hesse data using the CUSUM LR method.

all future y_t . Lines show the resulting $\mu_{0,t}$ values and the resulting value of the LR(n) statistic for a 50% increase, i.e. with $\mu_{1,t} = \frac{3}{2} \cdot \mu_{0,t}$. With a threshold of $c = 4.0$ the probability for a false alarm within the monitored 98 timepoints is calculated to be 0.05 – with this threshold the first outbreak is detected March 2000. After the alarm the detector is reset by re-estimating the Poisson model now using values up to March 2000 as historical values. Subsequent predictions for the in-control mean are then used to continue monitoring from April 2000 until the next alarm. This procedure continues until the end of the monitoring period is reached – resulting alarms (triangles) and in-control mean, $\mu_{0,t}^{re}$, obtained after re-fitting following each alarm, are shown in Fig. 4.

As a second step taking spatial aggregation into account we consider the three

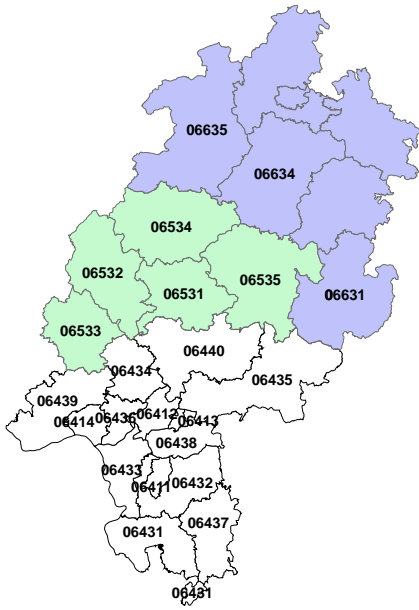


Fig. 5. Map of Hesse, Germany. The shaded regions indicate the three different administrative regions (south, middle, north), the numbers are the unique identifiers (Gemeindeschlüssel) of the 26 (14+5+7) districts of Hesse.

levels of aggregation arising from the administrative division in Hesse, see Fig. 5. However, due to an administrative reform data on district level only available from 1990. Figure 6 shows surveillance of all time series of the hierarchy beginning from January 1998. To avoid careful tuning, the Farrington method with $w = 2$ and $b = 4$ is used.

One promptly sees that outbreak during 2000 for entire Hesse is due to problems in the region south – especially district 06435 (Main-Kinzig-Kreis) located at the border to Bavaria. Investigations showed that there were problems with synchronising aerial and hand distribution of the oral rabies vaccination in this and the neighbouring city of Offenbach (06413) (Müller et al., 2005). Figure 7 summarizes the results in a so called alarm plot, which shows a problem in district 06532 (Lahn-Dill-Kreis), causing alarms for region middle at the turn of 1999.

3 Retrospective analysis of surveillance data

The focus of prospective surveillance, described in the previous section, is on outbreak detection. In contrast, retrospective surveillance tries to explain temporal and spatio-temporal patterns in the data through statistical modelling. Following Held et al. (2005), one possible model approach for the analysis of multivariate surveillance data is presented.

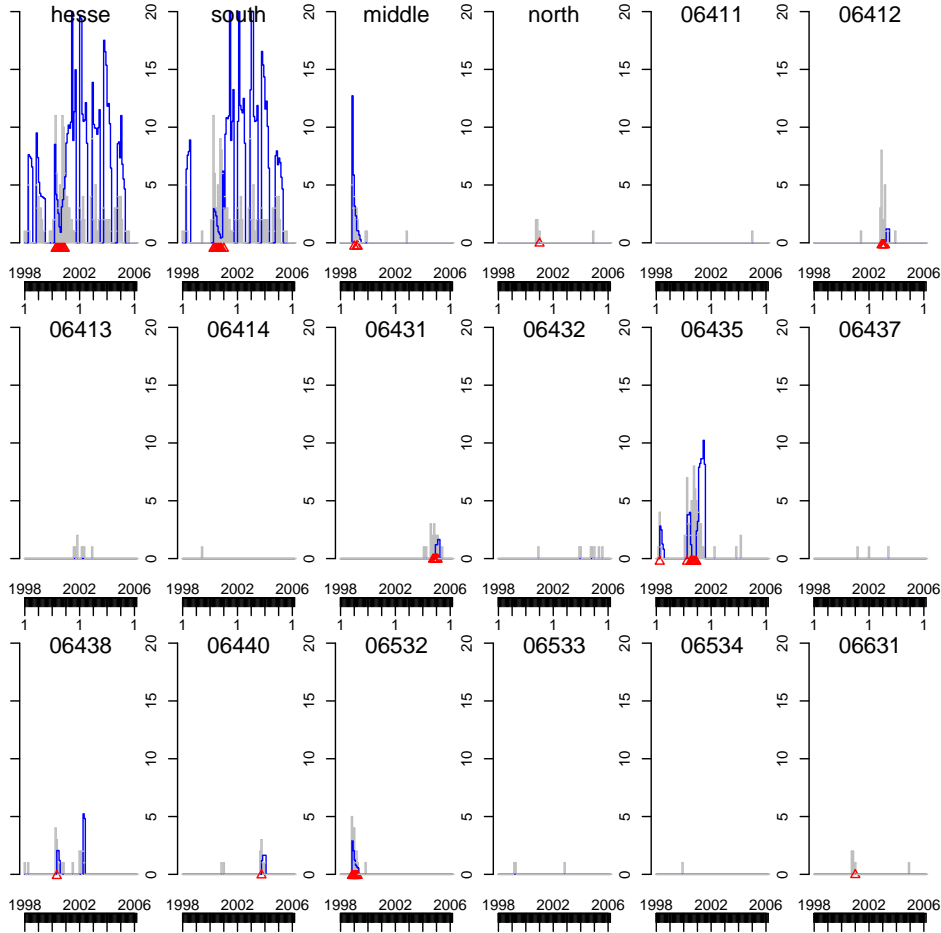


Fig. 6. Surveillance results for each of the $m = 30$ time series having at least one case in the monitored period. Top-left plot is of all cases in Hesse followed by the three administrative regions and the districts. For better visualization the y -axis is truncated at the value 20.

3.1 Multivariate modelling

As in Sect. 2.2, let $y_{i,t}$, $i = 1, \dots, m$, $t = 1, \dots, n$, denote a multivariate time series of counts. The main feature of the model proposed by Held et al. (2005) is the additive decomposition of the incidence into an endemic and an epidemic component with rates $\eta_{i,t}$ and $\nu_{i,t}$, respectively. In the simplest case, the observed counts $y_{i,t}$ in region i at time t are assumed to be Poisson distributed with mean

$$\mu_{i,t} = \eta_{i,t} + \nu_{i,t}.$$

Endemic incidence is persistent with a stable temporal pattern. The endemic component $\nu_{i,t}$ thus may include terms for long-term trends and seasonality and is basically modelled as in (1). Region-specific intercepts allow for different

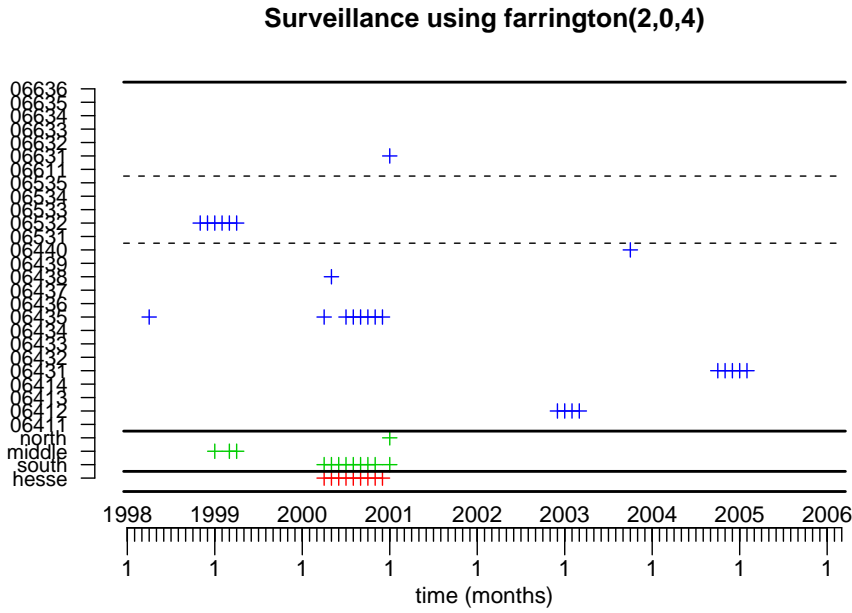


Fig. 7. Alarm plot resulting from the hierarchical surveillance of the $m = 1 + 3 + 26$ time series using the Farrington method. The solid lines divide the series according to their hierarchy (at bottom is lvl. zero), dotted lines group descendants.

incidence levels in the m regions. All in all, $\nu_{i,t}$ is specified as

$$\log \nu_{i,t} = \alpha_i + \beta t + \sum_{s=1}^S \left[\gamma_s \sin \left(\frac{2\pi}{r} s \cdot t \right) + \delta_s \cos \left(\frac{2\pi}{r} s \cdot t \right) \right].$$

The epidemic component $\eta_{i,t}$ should be able to explain occasional outbreaks and capture spatio-temporal dependence caused by the spread of the disease across regions. One possible approach is to let the number of previous counts in the region and in neighbouring regions enter as autoregressive covariates in the epidemic component, i.e.

$$\eta_{i,t} = \lambda y_{i,t-1} + \phi \sum_{j \sim i} y_{j,t-1}$$

where $j \sim i$ denotes all regions adjacent to region i .

In many applications the Poisson assumption of equal mean and variance is not realistic. To adjust for possible overdispersion, Held et al. (2005) suggest a negative binomial model with additional dispersion parameter $\psi > 0$, where the mean remains the same but the variance increases to $\mu_{i,t} + \mu_{i,t}^2/\psi$.

3.2 Results

To illustrate the modelling approach of Held et al. (2005) we consider the monthly number of rabies cases in Hesse and Bavaria. The data are analysed on two aggregation levels: state level and district level. Data on district level are not available until 1990, hence we only use data for 1990 – 2006. As described in Sect. 2.3, there is a strong seasonality and the number of cases is decreasing as a consequence of the vaccination program. Therefore, seasonal terms and a linear time trend are always included.

Separate univariate analyses of the counts in Hesse and in Bavaria showed that incidence levels and the slopes of the linear trend differ whereas seasonality is similar in both federal states. Likelihood-ratio tests suggest to use $S = 2$ seasonal terms. Furthermore, there is evidence for overdispersion because the negative binomial models result in a significant increase in terms of maximised log-likelihood compared to the corresponding Poisson models.

Table 1 shows results for the joint analysis on state level. In all models, $S = 2$ seasonal terms and a state-specific linear trend are included in the endemic component. Inclusion of an autoregressive parameter λ leads to a pronounced increase of the likelihood. The maximum likelihood (ML) estimate of λ is 0.69 (0.05), clearly indicating a temporal dependence after adjustment for seasonal effects. Inclusion of the autoregressive parameter ϕ does improve the fit only slightly and is not required according to the AIC model choice criterion.

On district level, we restrict our attention to 12 districts in the boarder region between Hesse and Bavaria (nine districts in Hesse and three districts in Bavaria). Two districts have been defined to be adjacent if they share a common border. Results for several models are shown in Tab. 2. As above, we chose the negative binomial model and included a linear trend and $S = 1$ seasonal terms in the linear predictor, higher terms for seasonality did not lead to a significant improvement in the likelihood. Again, there is evidence for temporal dependence. In addition, there exists spatial dependence: the autoregressive parameter ϕ that captures the influence of neighbouring districts contributes markedly to a better fit. This finding is consistent with the spread of the disease across districts described in Müller et al. (2005).

4 Discussion

Surveillance in veterinary public health is an important contribution for the detection and control of diseases in veterinary epidemiology. An important criterion for selecting the appropriate surveillance method is the number of time series to monitor. If performance is premium making the cost of tuning second-rank, the Rogerson and Yamada (2004a) or the direct LR-CUSUM

Table 1

Summary of ML estimates (standard errors) of different negative binomial models for the rabies data in Hesse and Bavaria (state level), $\log L$ denotes the maximised log likelihood, p is the number of parameters and $AIC = -2 \log L + 2p$.

$\hat{\lambda}_{ML}$ (se)	$\hat{\phi}_{ML}$ (se)	$\hat{\psi}_{ML}$ (se)	$\log L$	p	AIC
-	-	1.04 (0.12)	-829.1	9	1676.2
0.69 (0.05)	-	4.76 (0.96)	-702.8	10	1425.6
0.68 (0.05)	0.022 (0.019)	4.87 (0.99)	-701.9	11	1425.8

Table 2

Summary of ML estimates (standard errors) of different models for the rabies data in 12 districts in the boarder region of Hesse and Bavaria, $\log L$ denotes the maximised log likelihood, p is the number of parameters and $AIC = -2 \log L + 2p$.

$\hat{\lambda}_{ML}$ (se)	$\hat{\phi}_{ML}$ (se)	$\hat{\psi}_{ML}$ (se)	$\log L$	p	AIC
-	-	0.22 (0.02)	-1359.2	16	2750.5
0.57 (0.05)	-	0.82 (0.12)	-1167.8	17	2369.7
0.55 (0.05)	0.041 (0.008)	0.91 (0.14)	-1146.6	18	2329.3

method from Sect. 2.1.2 should be used. When there is uncertainty about the correct parameter of the alternative so called generalized likelihood ratio (GLR) detectors can be used, which estimate the unknown parameter (Höhle and Paul, 2007).

Multivariate surveillance of count time series is still an area of active research. Our approach of multiple univariate surveillance ignores correlations between the series. With an appropriate time series model explaining seasonality the remaining correlation is often negligible or caused by auto-regression due to disease transmission. One possibility is thus to let the out-of-control model contain an autoregressive component for transmission within and between units as suggested in Sect. 3.2.

A retrospective analysis of surveillance data may give clues about temporal and spatio-temporal patterns of the disease considered. In our example, considered in Sect. 3.2, we were able to identify a significant spatio-temporal autoregressive coefficient, which gives evidence of spatio-temporal interaction of fox rabies. However, the model considered is quite simplistic. In current work we consider more realistic model approaches for the analysis of spatio-temporal surveillance data.

All methods covered in this paper are available in `surveillance` (Höhle, 2008) – a software package for surveillance methods using the free software environment for statistical computing and graphics "R".

5 Acknowledgments

We thank Christoph Staubach, Federal Research Institute for Animal Health, Germany, for helpful comments and for providing us with the rabies data. Financial support was given by the German Science Foundation (DFG, 2003-2006) and the Swiss National Science Foundation (SNF, since 2007). The first author also thanks the Munich Center of Health Sciences for financial support.

References

- Bourke, P., 1992. Performance of cumulative sum schemes for monitoring low count-level processes. *Metrika* 39, 365–384.
- Bravata, D. M., McDonald, K. M., Smith, W. M., Rydzak, C., Szeto, H., Buckeridge, D. L., Haberland, C., Owens, D. K., 2004. Systematic review: Surveillance systems for early detection of bioterrorism-related diseases. *Annals of Internal Medicine* 140 (11), 910–922.
- Buckeridge, D., Burkom, H., Campbell, M., Hogan, W., Moore, A., 2005. Algorithms for rapid outbreak detection: a research synthesis. *Journal of Biomedical Informatics* 38 (2), 99–113.
- Carpenter, T. E., Chriel, M., Greiner, M., 2007. An analysis of an early-warning system to reduce abortions in dairy cattle in Denmark incorporating both financial and epidemiologic aspects. *Preventive Veterinary Medicine* 78, 1–11.
- Chen, R., 1978. A surveillance system for congenital malformations. *Journal of the American Statistical Association* 73, 323–327.
- Farrington, C., Andrews, N., 2003. Monitoring the health of populations. Oxford University Press, Ch. 8 Outbreak Detection: Application to Infectious Disease Surveillance, pp. 203–231.
- Farrington, C., Andrews, N., Beale, A., Catchpole, M., 1996. A statistical algorithm for the early detection of outbreaks of infectious disease. *Journal of the Royal Statistical Society, Series A* 159, 547–563.
- Frisén, M., 2003. Statistical surveillance: Optimality and methods. *International Statistical Review* 71 (2), 403–434.
- Harnos, A., Reiczigel, J., Solymosi, N., Vattay, G., 2006. Analysis of the effect of immunization in rabies time series. *Journal of Theoretical Biology* 240, 72–77.
- Hawkins, D. M., 1992. Evaluation of average run lengths of cumulative sum charts for an arbitrary data distribution. *Communications in Statistics. Simulation and Computation* 21 (4), 1001–1020.
- Held, L., Höhle, M., Hofmann, M., 2005. A statistical framework for the analysis of multivariate infectious disease surveillance data. *Statistical Modelling* 5, 187–199.

- Höhle, M., 2008. surveillance: An R package for the monitoring of infectious diseases. Computational Statistics, To appear.
- Höhle, M., Paul, M., 2007. Poisson regression charts for the monitoring of surveillance time series. Submitted.
- Kleinmann, K. P., Abrams, A. M., 2006. Assessing surveillance using sensitivity, specificity and timeliness. *Statistical Methods in Medical Research* 15, 445–464.
- Kosmider, R., Kelly, L., Evans, S., Gettinby, G., 2006. A statistical system for detecting salmonella outbreaks in british livestock. *Epidemiology and Infection* 134, 952–960.
- Kulldorff, M., 2001. Prospective time periodic geographical disease surveillance using a scan statistic. *Journal of the Royal Statistical Society, Series A* 164, 61–72.
- Lai, T., 1995. Sequential changepoint detection in quality control and dynamical systems. *Journal of the Royal Statistical Society, Series B* 57, 613–658.
- Lawson, A., Kleinman, K. (Eds.), 2005. *Spatial and Syndromic Surveillance for Public Health*. Wiley.
- Lucas, J. M., 1985. Counted data CUSUMs. *Technometrics* 27 (2), 129–144.
- Müller, T., Selhorst, T., Pötsch, C., 2005. Fox rabies in Germany – an update. *Eurosurveillance* 10, 229–231.
- Rogerson, P., Yamada, I., 2004a. Approaches to syndromic surveillance when data consist of small regional counts. *Morbidity and Mortality Weekly Report* 53, 79–85.
- Rogerson, P., Yamada, I., 2004b. Monitoring change in spatial patterns of disease: comparing univariate and multivariate cumulative sum approaches. *Statistics in Medicine* 23, 2195–2214.
- Rogerson, P. A., 2001. Monitoring point patterns for the development of space-time clusters. *Journal of the Royal Statistical Society, Series A* 164, 87–96.
- Rossi, G., Lampugnani, L., Marchi, M., 1999. An approximate CUSUM procedure for surveillance of health events. *Statistics in Medicine* 18, 2111–2122.
- Sonesson, C., Bock, D., 2003. A review and discussion of prospective statistical surveillance in public health. *Journal of the Royal Statistical Society, Series A*, 5–12.
- Steiner, S. H., Cook, R. J., Farewell, V. T., Treasure, T., 2000. Monitoring surgical performance using risk-adjusted cumulative sum charts. *Biostatistics* 1 (4), 441–452.
- Thulke, H.-H., Tischendorf, L., Staubach, C., Selhorst, T., Jeltsch, F., Müller, T., Schlüter, H., Wissel, C., 2000. The spatio-temporal dynamics of a post-vaccination resurgence of rabies in foxes and emergency vaccination planning. *Preventive Veterinary Medicine* 47, 1–21.
- Tourneret, J.-Y., Ferrari, A., Letac, G., 2002. Changepoint detection in multivariate Poisson distributions. In: *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*. Orlando, May 13 – 17, pp. 1573–1576.
- WHO Collaboration Centre for Rabies Surveillance and Research, 2007. WHO

- Rabies - Bulletin - Europe. <http://www.rbe.fli.bund.de/>.
- Widdowson, M.-A., Bosman, A., van Straten, E., Tinga, M., Chaves, S., van Eerden, L., van Pelt, W., 2003. Automated, laboratory-based system using the internet for disease outbreak detection, the Netherlands. *Emerging Infectious Diseases* 9 (9), 1046–1052.
- Woodall, W., 2006. The use of control charts in health-care and public-health surveillance. *Journal of Quality Technology* 38 (2), 89–104.