Epidemic detection using CUSUM

Georgios Fellouris[1], George V. Moustakides[2] and Yajun Mei[3]

[1] Department of Statistics
Columbia University,
1255 Amsterdam Avenue
New York, NY 10027, USA
(e-mail: gaf2106@columbia.edu)

[2] Department of Electrical and Computer Engineering
University of Patras
26500 Rio , GREECE
(e-mail: moustaki@upatras.gr)

[3] School of Industrial and Systems Engineering
Georgia Institute of Technology
765 Ferst Drive NW
Atlanta, GA 30332-0205 , USA
(e-mail: ymei@isye.gatech.edu)

Abstract. We consider the problem of detecting a proportional change in the intensity of a counting process with a continuous compensator. We prove that if the counting process does not have explosions, the CUSUM test is optimal in a Lorden sense. In particular, the CUSUM test minimizes the worst-case conditional expected number of events that occur after the change and before raising an alarm given the worst possible history up to the time of the change. We apply this result to the detection of the onset of epidemics and we discuss modifications of the optimal detection rule when the processes of interest are monitored in discrete times.
Keywords. CUSUM, Change-Detection, Epidemics, Counting Process.

1 Introduction and motivation

Since its introduction by Page (1954) the Cumulative Sums or CUSUM test has been one of the most popular sequential schemes for detecting abrupt changes in the dynamics of a stochastic system. A strong theoretical argument in its favor was provided by Moustakides (1986), who proved that the CUSUM is the optimal detection rule according to the criterion suggested by Lorden(1971) in the case of independent and identically distributed random variables before and after the change. The optimality of the CUSUM rule was later extended for other probabilistic models for the observations (for example, Shiryaev (1996), Moustakides(2004)).

As it was discussed by Mei et. al (2009), a natural field of application for the techniques of quickest detection -and the CUSUM in particular- is biosurveillance. The goal in biosurveillance is to detect rapidly and accurately the onset of an epidemic, that is, when the *incidence rate* of a particular disease in a certain population has increased significantly above the standard rate. Thus, in order to detect the outbreak of a disease it is important to take into account not only the number of cases but also the corresponding population size. We do this in a counting process framework.

2 Problem formulation and main result

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space on which we define the stochastic processes $\{N_t, B_t, D_t\}_{t \geq 0}$. N_t counts the number of cases in the population up to time t, whereas B_t and D_t represent the number of births and deaths in the population up to time t. We assume that all these processes are observable and we denote by $\{\mathcal{F}_t\}$ the corresponding filtration:

$$\mathcal{F}_t = \sigma(N_s, B_s, D_s : 0 \leq s \leq t), \quad t \geq 0.$$

We model the emergence of an epidemic as an abrupt change in the *intensity* of the counting process $\{N_t\}$, which occurs at some *unknown* and *deterministic* time $\tau \in [0, \infty]$ ($\tau = \infty$ corresponds to the case of no change). In other words, we parametrize the underlying probability measure $\mathbb{P}$ by the unknown time of the change τ and we assume that

$$M_t = N_t - \int_0^t \lambda_s ds, \quad t \geq 0$$

is a $(\mathbb{P}_\tau, \mathcal{F}_t)$-*martingale*, where

$$\lambda_t = \begin{cases} \ell_t\,, & \text{when } t \leq \tau \\ \rho\,\ell_t\,, & \text{when } \ t > \tau \end{cases} \quad , \quad \ell_t = P_0 + B_t - D_t, \quad t \geq 0.$$

P_0 is the initial population size at time $t = 0$ and $\rho \neq 1$ a known constant. Thus, we assume that the intensity of $\{N_t\}$ is proportional to the population size with a constant of proportionality that changes at the onset of the epidemic τ. The goal is to detect τ based on the sequentially acquired observations, thus any $\mathcal{F}_t$-stopping time is a potential detection rule.

The CUSUM detection rule has the following form:

$$\mathcal{S}_c = \inf\{t \geq 0 : y_t \geq c\}, \quad \text{where} \quad y_t = u_t - m_t \quad \text{and}$$

$$u_t = (\log \rho)\; N_t + (1-\rho) \int_0^t \ell_s\, ds \quad , \quad m_t = (-y) \wedge \inf_{0 \leq s \leq t} u_s, \quad 0 \leq t < \infty.$$

y is the initialization for $(y_t)_{t \geq 0}$ and can take any value in the interval $[0, c]$. The standard choice is $y = 0$, whereas a positive value for y leads to faster detection.

In this work, we establish the exact optimality of $\mathcal{S}_c$ under a *modified* version of Lorden's criterion. Thus, we consider the following performance measure for an arbitrary detection rule $\mathcal{T}$:

$$\mathcal{J}[\mathcal{T}] = \sup_{\tau \geq 0}\; \mathit{esssup}\; \mathbb{E}_\tau \left[(N_\mathcal{T} - N_\tau)^+ | \mathcal{F}_\tau \right],$$

where $\mathbb{E}_\tau[\cdot]$ is the expectation with respect to $\mathbb{P}_\tau$. $\mathcal{J}[\mathcal{T}]$ represents the worst-case conditional expected delay of $\mathcal{T}$ given the worst possible history up to the time of the change. The difference with the original performance measure suggested by Lorden is that here we measure the detection delay in terms of the *number of cases* - not the actual *time*- that occur between the change (or onset of epidemic) and the issuing of the alarm.

Our main contribution can be summarized in the following theorem:

Theorem 1. *Let c be chosen so that $\mathbb{E}_\infty[\mathcal{S}_c] = \gamma$, where γ is a fixed design parameter. If $\mathbb{P}_0(N_\infty = \infty) = 1$, then:*

$$\mathcal{J}[\mathcal{S}_c] = \inf_{\mathcal{T}} \mathcal{J}[\mathcal{T}],$$

where the infimum is taken over detection rules such that $\mathbb{E}_\infty\left[N_\mathcal{T}\right] \geq \gamma$.

Thus, we prove that $\mathcal{S}_c$ is the best detection rule among stopping times for which the expected number of events till a false alarm is at least γ. The condition $\mathbb{P}_0(N_\infty = \infty) = 1$ guarantees that $\mathcal{S}_c$ terminates almost surely and implies that the number of cases -after the change- increases with the horizon of observations. Therefore, the optimality of $\mathcal{S}_c$ does not necessarily hold for right-censored counting processes where only a finite number of events can be observed.

Moreover, note that the above optimality property is valid independently of the particular stochastic model for ℓ_t. Thus, it is straightforward to generalize the above model by allowing ℓ_t to be a function of other observed stochastic processes which may explain the evolution of $\{N_t\}$.

3 The case of discrete-time observations

Suppose now that the processes $\{N_t, B_t, D_t\}$ are observed only at the discrete times $0 = t_0 < t_1 < \ldots < t_n < \ldots$, which may not be deterministic and equidistant. Since

$$u_{t_n} - u_{t_{n-1}} = (\log \rho)\ (N_{t_n} - N_{t_{n-1}}) + (1-\rho) \int_{t_{n-1}}^{t_n} \ell_s \, ds\,, \quad n \in \mathbb{N}$$

we cannot in general recover $u_{t_n} - u_{t_{n-1}}$ based on $B_{t_n} - B_{t_{n-1}}$ and $D_{t_n} - D_{t_{n-1}}$, the observed number of births and deaths in $(t_{n-1}, t_n]$.

We suggest approximating the integral $\int_{t_{n-1}}^{t_n} \ell_s \, ds$ with the corresponding conditional expectation, i.e. $P_0(t_n - t_{n-1}) + \hat{B}_n - \hat{D}_n$, where

$$\hat{B}_n = \mathbb{E}\Big[\int_{t_{n-1}}^{t_n} B_s \, ds \Big| B_{t_n} - B_{t_{n-1}}\Big] \quad , \quad \hat{D}_n = \mathbb{E}\Big[\int_{t_{n-1}}^{t_n} D_s \, ds \Big| D_{t_n} - D_{t_{n-1}}\Big], \quad n \in \mathbb{N}.$$

Thus, the (stochastic) model for the evolution of the population plays an important role in the construction of an efficient discrete-time approximation for $\mathcal{S}_c$, unlike the continuous-time case where the stochastic models for $\{B_t\}$ and $\{D_t\}$ were not affecting the implementation of $\mathcal{S}_c$.

Then, we can approximate $\{u_{t_n}\}$ and $\{m_{t_n}\}$ with

$$\hat{u}_n = (\log \rho)\ N_{t_n} + (1-\rho)\Big[P_0\, t_n + \sum_{j=1}^{n} (\hat{B}_j - \hat{D}_j)\Big] \quad , \quad \hat{m}_n = \min_{j=0,\ldots,n} \hat{u}_j \quad , \quad n \in \mathbb{N}$$

and apply the CUSUM rule as follows

$$\hat{\mathcal{S}}_{\hat{c}} = \inf\{t_n : \hat{y}_n \geq \hat{c}\}, \quad \text{where} \quad \hat{y}_n = \hat{u}_n - \hat{m}_n.$$

Again, $\hat{c}$ should be chosen so that $\mathbb{E}_\infty[\hat{\mathcal{S}}_{\hat{c}}] = \gamma$.

References

Lorden, G. (1971): "Procedures for reacting to a change in distribution," *Ann. Math. Stat.*, *41*, 520–527.

Mei, Y. , Han, S.W. and Tsui, K.L (2010): " Early detection of a change in Poisson rate after accounting for population size effects ", submitted for publication.

Moustakides, G.V (1986): "Optimal stopping times for detecting changes in distributions," *Annals of Statistics*, *14*, 1379–1387.

Moustakides, G.V. (2004): "Optimality of the CUSUM procedure in continuous time," *Annals of Statistics*, *32*, 302–315.

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

Shiryaev, A.N. (1996): "Minimax optimality of the method of cumulative sums (CUSUM) in the case of continuous time," *Russ. Math. Surv.*, *51*, 750–751