Monitoring a Poisson Process in Several Categories subject to Changes in the Arrival Rates

Marlo Brown

Department of Mathematics, Niagara University,
Niagara University, NY 14109
mbrown@niagara.edu

Abstract. We look at a Poisson process in several categories where the arrival rate changes at some unknown integer. For some of these categories the arrival rates increase, while in other categories the arrival rates decrease. The point at which the process changes may be different for each category. We develop procedures for detecting when a change has occurred in at least one of the categories. We provide some numerical results to illustrate the effectiveness of the detection procedures.

Key words: Change-point Detection, Dynamic Programming, Multiple Categories, Poisson processes.

1 Introduction

Detection of changes in the distribution of random variables has become very important in many aspects of life today. When there is an increase in the arrival rates of patients coming to a hospital, it is important to detect this change as soon as possible. This could be due to environmental factors or other issues. This is also important in industry where quality control depends upon being able to detect changes in the process mean as soon as possible.

Several studies were published recently on detecting changes in the intensity of a homogeneous ordinary Poisson process. Among these studies, we mention Peskir and Shiryaev (2002), Herberts and Jensen (2004) and Brown and Zacks (2006a). These papers dealt with a Poisson process which is monitored continuously. Brown and Zacks (2006b) also studied a Poisson process which is monitored only at discrete time points. In that paper, we look at the sequence of random variables X_i where X_i is the number of arrivals that occur in the time interval $(i-1, i]$. Thus, we see only the number of arrivals that happen in each time interval, not exactly where the arrivals occurred within that time interval.

In 2006, Tartakovsky studied the detection of changes in at least one of several categories monitored simultaneously. Tartakovsky gave many applications to invassions in computer systems. There are several instances where we many want to split a Poisson process into several different categories. At some unknown time point, there may be a change in the arrival rate of one or more categories.

In the present paper, we assume that the arrival rates before and after the change are known but the change-point is unknown and may or may not be the same for each category. We use a Bayesian approach, putting Shiryaev (1978) geometric prior on the change-point τ_j for each category j. In Section 2, we discuss the posterior process of calculating the probability that a change has already occurred by time n for each category j. In Section 3, we discuss the optimal stopping rule based on Dynamic Programming procedures. We generalize the Dynamic Programming procedure used for one sequence of variables to one that is optimal for several sequences of variables monitored simultaneously. We assume there is a cost associated with stopping early and a cost associated with late detection. Our goal is to stop as soon as a change happens in at least one of several categories. We develop optimal stopping rules that will minimize the expected cost. In Section 4, we give an explicit formulation of the two step ahead stopping rule. We conclude with a numerical example showing how our method works.

2 The Bayesian Framework

We have a Poisson process in each of k categories. We assume the k Poisson processes are independent. We monitor each process at fixed discrete time points. Let $X_{i,j}$ be the number of arrivals in the interval $(i-1, i]$ in category j. At the end of some unknown interval indexed by an integer τ_j, the rate of arrivals changes from $\lambda_{1,j}$ to $\lambda_{2,j}$. We assume the arrival rates before and after the change, namely $\lambda_{1,j}$ and $\lambda_{2,j}$, are known in each category, but the change-point τ_j is unknown. In some of the categories this change may cause an increase in the arrival rates while in other categories the change may cause a decrease in arrival rates. It is also important to note that we are allowing the change-point to differ from one category to another. Our goal is to find a detection procedure that will detect a change in at least one of the categories as soon as possible. The event $\{\tau_j = 0\}$ represents the case where the change in category j happened prior to the first observed interval. In this case, all the observations in this category are taken from a process with the arrival rate at $\lambda_{2,j}$. The event $\{\tau_j = r\}$ means that the change in category j happens at time r. Define, $T_{r,j} = \sum_{i=1}^{r} X_{i,j}$ as the total number of arrivals before time r in category j, and $T^*_{n-r,j} = \sum_{i=r+1}^{n} X_{i,j}$ as the total number of arrivals after time r in category j. $T_{0,j} = 0$ and $T^*_{0,j} = 0$ for all j. Thus, when $\tau_j = r$, T_{r_j} is Poisson $(r\lambda_{1,j})$ and $T^*_{n-r,j}$ is Poisson $((n-r)\lambda_{2,j})$. For each category j, the likelihood of τ_j is given by

$$L(\tau_j) = \sum_{r=0}^{n-1} I(\tau_j = r)\, \lambda_{1,j}^{T_{r,j}} \lambda_2^{T^*_{n-r,j}} e^{-r\lambda_{1,j}-(n-r)\lambda_{2,j}} + I(\tau_j \geq n)\, e^{-n\lambda_{1,j}} \lambda_{1,j}^{T_{n,j}}. \tag{1}$$

We put Shiryaev's (1978) geometric prior distribution h on the change-point τ_j for each j.

$$h(\tau_j = i) = \begin{cases} \pi & \text{for } i = 0 \\ (1-\pi)\, p_j\, (1-p_j)^{i-1} & \text{for } i \geq 1 \end{cases}. \tag{2}$$

Thus, for each category j, we obtain the posterior probability that there has been a change by time n is

$$\psi_{n,j} = 1 - \frac{(1-\pi)(1-p_j)^n e^{-n\lambda_{1,j}} \lambda_{1,j}^{T_{n,j}}}{D_{n,j}}, \tag{3}$$

where

$$\begin{aligned} D_{n,j} &= \pi\lambda_2^{T_{n,j}} e^{-n\lambda_{2,j}} + (1-\pi)\sum_{r=1}^{n-1} p_j (1-p_j)^r \lambda_{1,j}^{T_{r,j}} \lambda_2^{T^*_{n-r,j}} e^{-r\lambda_{1,j}-(n-r)\lambda_{2,j}} \\ &\quad + (1-\pi)(1-p_j)^n e^{-n\lambda_{1,j}} \lambda_{1,j}^{T_{n,j}}. \end{aligned} \tag{4}$$

Let $\bar{\psi}_{n,j} = 1 - \psi_{n,j}$ be the probability that there has not been a change by time n in category j. We are interested in the probability that at least one of the categories change. By independence, the negation $q_n = \prod_{j=1}^{k} \overline{\psi}_{n,j}$ is the probability that the change has not happened in any category.

Lemma 1. *For each n and j, the posterior probabilities $\psi_{n,j}$ satisfy the recursive relationship*

$$\psi_{n+1,j}(X_{1,j}, \ldots, X_{n+1,j}, \pi) = \psi_{1,j}(X_{n+1,j}, \psi_n). \tag{5}$$

Proof. Letting $\rho_j = \frac{\lambda_{2,j}}{\lambda_{1,j}}$ and $\delta_j = \lambda_{2,j} - \lambda_{1,j}$ we can express the probability that a change has not occurred by time $n+1$ in category j is

$$\bar{\psi}_{n+1,j} = \frac{(1-\pi)(1-p_j)^{n+1}}{D_{n+1,j}} \tag{6}$$

By making the substitution $(1-\pi)(1-p_j)^n = D_{n,j}\bar{\psi}_{n,j}$, we obtain

$$\bar{\psi}_{n+1,j} = \frac{\bar{\psi}_{n,j}(1-p_j)}{\psi_{n,j}\rho_j^{X_{n+1}} e^{-\delta_j} + \bar{\psi}_{n,j} p \rho_j^{X_{n+1}} e^{-\delta_j} + \bar{\psi}_{n,j}(1-p_j)}. \tag{7}$$

Thus $\psi_{n+1,j}(X_{1,j}, \ldots, X_{n+1,j}, \pi) = \psi_{1,j}(X_{n+1,j}, \psi_{n,j})$.

Therefore, the posterior distribution of τ_j for each j is given by

$$P(\tau_j = i | X_{1,j}, \ldots, X_{n,j}) = \begin{cases} \psi_{n,j} & \text{for } i = 0 \\ \bar{\psi}_{n,j} p_j (1-p_j)^{i-1} & \text{for } i = 1, 2, \ldots \end{cases}. \tag{8}$$

Lemma 2. *For each $n \geq 1$, and $j = 1, \ldots, k$, the predictive density of $X_{n,j}$ is*

$$f_{n,j}(x) = \psi_{n-1,j} \frac{\lambda_{2,j}^x e^{-\lambda_{2,j}}}{x!} + \bar{\psi}_{n-1,j} p_j \frac{\lambda_{2,j}^x e^{-\lambda_{1,j}}}{x!} + \bar{\psi}_{n-1,j} \left(1 - p_j\right) \frac{\lambda_{1,j}^x e^{-\lambda_{1,j}}}{x!}. \tag{9}$$

Proof. The density of $X_{n,j}$, given τ_j, is

$$f_{n,j}(x \mid \tau) = I(\tau_j = 0) \frac{\lambda_{2,j}^x e^{-\lambda_{2,j}}}{x!} + I(\tau_j = 1) \frac{\lambda_{2,j}^x e^{-\lambda_{1,j}}}{x!} + I(\tau_j > 1) \frac{\lambda_{1,j}^x e^{-\lambda_{1,j}}}{x!}. \tag{10}$$

Taking the expected value of $f_{n,j}(x \mid \tau)$ with respect to the posterior distribution of τ_j, given $X_{i,j}, i = 1, \ldots, n$, we obtain (9).

At each time point, we calculate for each category j, the probability that a change has not occurred in that category to be $\bar{\psi}_{n,j}$. We define $\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}$ to be the vector: $\left\{\bar{\psi}_{n1}, \bar{\psi}_{n2}, \ldots, \bar{\psi}_{nk}\right\}^t$.

Lemma 3. *The sequence $\{q_n, \boldsymbol{\psi}_n, n \geq 1\}$ is a supermartingale.*

Proof. We calculate $E\left(q_{n+1} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)$ as follows.

$$E\left(q_{n+1} | \boldsymbol{\psi}_{\boldsymbol{n}}\right) = E\left(\prod_{i=1}^{k} \bar{\psi}_{n+1,i} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = \prod_{i=1}^{k} E\left(\bar{\psi}_{n+1,i} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = \prod_{i=1}^{k} \left(1 - p_i\right)^k \bar{\boldsymbol{\psi}}_{\boldsymbol{n}} = q_n \prod_{i=1}^{k} \left(1 - p_i\right). \tag{11}$$

where the conditional expectation in (11) is taken with respect to the predictive distribution of $X_{n+1,l}$ for all l. Thus the sequence $\{q_n, \boldsymbol{\psi}_n, n \geq 1\}$ is a supermartingale.

For notational purposes, we let $\bar{P}_k = \prod_{i=1}^{k} \left(1 - p_i\right)$. Hence, $E\left(q_{n+1} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = q_n \bar{P}_k$.

3 Optimal Stopping Rule

In Section 2, we calculated at each integer n, the probability that a change has not occurred in category j to be $\bar{\psi}_{n,j}$. Thus we found the probability that the change has not occurred in any of the k categories to be $q_n = \prod_{j=1}^{k} \bar{\psi}_{n,j}$. We assume there is a cost associated with stopping early and a cost associated with a delay in stopping. We assume without loss of generality that the cost associated with a false alarm is 1, and the cost per time unit to stop late is c. Thus, the risk of stopping at time n is q_n. The risk of continuing at time n is given by $c\left(1 - q_n\right) + E\left(R_{n+1} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)$. Therefore, the risk at time n is

$$R_n = \min\left(q_n, c\left(1 - q_n\right) + E\left(R_{n+1} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)\right). \tag{12}$$

According to (12) we would opt to stop and declare there is a change whenever the risk $R_n = q_n$. Equation (12) is equivalent to $R_n = q_n + \left[c - q_n\left(c + 1\right) + E\left(R_{n+1} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)\right]^-$. The stopping rule specifies to stop sampling and declare a change in at least one category the first time that $q_n \leq \frac{c + E\left(R_{n+1} | \boldsymbol{\psi}_{\boldsymbol{n}}\right)}{c+1}$. To find an optimal stopping rule, we first consider a truncated rule. We stop after n^* observations if we have not stopped before. Let $R_n^{(j)}$ be the risk at time n, when only j more observations are allowed. Note that $R_n^{(0)} = q_n$, and $R_n^{(1)} = q_n + \left[c - q_n\left(c + 1\right) + E\left(q_{n+1} | \bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)\right]^-$.

Lemma 4.

$$\lim_{n \to \infty} q_n = 0 \ \textit{a.s.}. \tag{13}$$

Proof. Indeed, according to equation (11), $E\{q_n\} = \left(1 - \pi\right) \bar{P}_k^n$, $n \geq 1$. Thus, $\limsup_{n \to \infty} E\{q_n\} = 0$. Since $\bar{\psi}_{n,j} \geq 0$ w.p. 1 for all $n \geq 0$ and $1 \leq j \leq k$, $q_n \geq 0$. Hence we obtain (13).

The risk looking one step ahead is $R_n^{(1)} = q_n + \left[c - q_n\left(c + 1\right) + \bar{P}_k q_n\right]^-$. Therefore, the one step ahead procedure is to stop sampling and declare a change in at least one category when $q_n < Q^*$, where $Q^* = \frac{c}{c+1-\bar{P}_k}$. Therefore, the one step ahead stopping variable is $N^{(1)} = \min\left\{n : q_n < Q^*\right\}$.

Lemma 5. *The one step ahead stopping random variable $N^{(1)} < \infty$ with probability 1.*

Proof. Since by Lemma 4, $\lim_{n\to\infty} \bar{\psi}_{n,j} = 0$ for all j. Thus, there exists an n such that $\bar{\psi}_{n,j} < Q^*$ for some j with probability 1. Therefore, $q_n < Q^*$ with probability 1 for some finite n. Hence, the one step ahead stopping random variable $N^{(1)} < \infty$ w.p. 1.

Now we look at "j-step look ahead" stopping rules. Define for $j \geq 0$, $M_n^{(j)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)$ recursively as the expected risk of continuing at the next observation.

$$M_n^{(j)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = E\left[\left[c - q_{n+1}\left(c+1-\bar{P}_k\right) + M_{n+1}^{(j-1)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n+1}}\right)\right]^- |\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right], \tag{14}$$

and $M_n^{(0)} \equiv 0$. Therefore the risk looking j steps ahead is $R_n^{(j)} = q_n + \left[c - q_n\left(c+1-\bar{P}_k\right) + M_n^{(j)}\right]^-$. The "j step ahead" procedure is to stop sampling and declare that a change has occurred in at least one of the categories the first time that $q_n < b_n^{(j)}$ where $b_n^{(j)}$ is defined as the "j step ahead" boundary $b_n^{(j)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = Q^* + \frac{M_n^{(j-1)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)}{\left(c+1-\bar{P}_k\right)}$. To calculate these boundary functions, we first calculate $M_n^{(l)}$.

$$M_n^{(l)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = \sum_{i_j}\prod_{j=1}^{k} P\left(X_{n+1,j} = i_j\right)\left[c - \left(c+1-\bar{P}_k\right)\prod_{j=1}^{k}\bar{\psi}_{n,j}\left(i_j\right) + M_{n+1}^{(l-1)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)\right]^-. \tag{15}$$

The sum in equation (15) is taken over those points for which $q_{n+1} > b_{n+1}^{(l-1)}$.

Lemma 6. *$\bar{\psi}_{n+1,l}\left(\bar{\psi}_{n,l}, X_{n+1,l}\right)$ is a decreasing function of $X_{n+1,l}$ when $\lambda_{1,l} < \lambda_{2,l}$ and increasing when $\lambda_{1,l} > \lambda_{2,l}$.*

Proof. Consider the function f defined by $f(X) = \psi_n \rho^X e^{-\delta} + \bar{\psi}_n p \rho^X e^{-\delta} + \psi_n(1-p)$ for $X = 0, 1, 2, \ldots$. Taking differences, $f(X+1) - f(X) = \left(\psi_n + \bar{\psi}_n p\right)\rho^X e^{-\delta}(\rho - 1)$. Thus, f is positive for $\rho > 1$ and negative for $\rho < 1$. Since $\bar{\psi}_{n+1,l} = \frac{\bar{\psi}_{n,l}(1-p_l)}{f\left(X_{n+1,l}\right)}$ where $\delta_l = \lambda_{2,l} - \lambda_{1,l}$ and $\rho_l = \frac{\lambda_{2,l}}{\lambda_{1,l}}$, the result follows.

Thus, there exists a region T_l such that if $q_{n+1}(i_1, \ldots, i_k) > b_{n+1}^{(l-1)}$, when $(i_1, \ldots, i_k) \in T_l$.

Lemma 7. *For the predictive distribution specified above,*
1. $M_n^{(l)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) \leq M_n^{(l-1)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) \leq 0$ for all $l \geq 1$.
2. $T_l \subset T_{l+1}$ for all $l \geq 1$.
3. $b_n^{(l)} \leq b_n^{(l-1)}$ for all $l \geq 1$.

Proof. The proof is by induction on l. First, $M_n^{(1)} \leq 0$ by the definition. Suppose that $M_n^{(k)} \leq M_n^{(k-1)}$ for all $k \leq l-1$. Now let's look at $M_n^{(l)} - M_n^{(l-1)}$ Subtracting we obtain.

$$E\left[c - q_{n+1}\left(c+1-\bar{P}_k\right) + M_{n+1}^{(l-1)}\right]^- - E\left[c - q_{n+1}\left(c+1-\bar{P}_k\right) + M_{n+1}^{(l-2)}\right]^- \leq 0. \tag{16}$$

This is true since $M_{n+1}^{(l-1)} < M_{n+1}^{(l-2)}$ by the inductive hypothesis. Thus, the $M_n^{(l)}$ functions are decreasing in l. To prove part 2, let $x \in T_l$. Therefore, by equation (3.14), $\left[c - \left(c+1-\bar{P}_k\right)\prod_{i=1}^{k}\bar{\psi}_{n,l}(i_i) + M_{n+1}^{(l-1)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right)\right] \leq 0$. This implies that

$$q_{n+1}(X_n, x) > Q^* + \frac{M_{n+1}^{(l-1)}}{c+1-(1-p)^k} > Q^* + \frac{M_{n+1}^{(l)}}{c+1-(1-p)^k}. \tag{17}$$

Thus $x \in T_{l+1}$. Similarly we can show that $b_n^{(l)} < b_n^{(l-1)}$ by inequality (17).

Since the $b_n^{(l)}$ are monotone decreasing and bounded below by 0, $\lim_{l\to\infty} b_n^{(l)} = b_n^{\infty}$. Therefore the optimal stopping rule is to stop sampling at the first $n \geq 1$ such that $q_n < b_n^{\infty}$.

Lemma 8. *We have the following convergence of $M_n^{(l)}$ and $b_n^{(l)}$.*
1. For a fixed $l \geq 1$, $\lim_{n\to\infty} M_n^{(l)} = 0$ a. s..
2. For a fixed $l \geq 2$, $\lim_{n\to\infty} b_n^{(l)} = Q^*$ a. s..

Proof. The proof is by induction on l. When $l = 1$, by equation (15),

$$M_n^{(1)}\left(\bar{\boldsymbol{\psi}}_{\boldsymbol{n}}\right) = \sum_{i_1,\ldots,i_k} \prod_{j=1}^{k} P_{\psi_j}\left(X_{n+1,j} = i_j\right) \left[c - \left(c + 1 - \bar{P}_k\right) \prod_{j=1}^{k} \psi_{n,j}\left(i_j\right)\right]^{-}. \tag{18}$$

By Lemma 4, $\bar{\psi}_{n,j} = 0$ as $n \to \infty$ for all j. Suppose there is at least one category where the arrival rate increases. In that category j, we can find N_j such $\bar{\psi}_{n,j}(0) < Q^*$ for all $n > N_j$. Thus, $\bar{\psi}_{n+1,j}(i_j) < Q$ for all i_j. Thus T_1 is empty, and $M_n^{(1)} = 0$ for all $n > N_j$. Now suppose for all the categories, the arrival rates decrease. For each category j, given any $\epsilon > 0$ there exists an I_j such that $P\left(X_{n+1,j} > i_j\right) < \epsilon$ for all $i_j > I_j$ and $0 \leq \psi_{n,j} \leq 1$. Since for all categories, $\bar{\psi}_{n,j} = 0$ as $n \to \infty$, there exists an N_j such that $\bar{\psi}_{n+1,j}(i_j) < Q_j$ for all $i_j < I_j$ and $n > N_j$. Thus, $T_1 = \{(i_1,\ldots,i_k) : i_j > I_j, j = 1,\ldots,k\}$, and $|M_n^{(1)}| < \prod_{j=1}^{k} P\left(X_{n+1,j} > I_j\right) < \epsilon^k$. Thus $M_n^{(1)} = 0$ as $n \to \infty$. Similarly by induction one can show that $M_n^{(l)} = 0$ as $n \to \infty$ for all $l > 1$. By (15) it immediately follows that $b_n^{(j)} = Q^*$ as $n \to \infty$.

4 Explicit Formulation of the two step Boundary

In this section we provide an algorithm for calculating the two step ahead boundary. To calculate the two step ahead procedure, we first calculate $M_n^{(1)}$. The expectation is taking with respective to the predictive distribution where $P_{\psi_j}\left(X_{n+1,j} = i_j\right)$ is the predictive probability that the $(n+1)^{st}$ observation in category j will be i_j. Since the Poisson processes in each category are independent, $\prod_{i=1}^{k} P_{\psi_j}\left(X_{n+1,j} = i_j\right)$ is the predictive probability that $X_{n+1} = (i_1,\ldots,i_k)$. We only sum over $(i_1,\ldots,i_k)$ for which $q_{n+1} > Q^*$. By Lemma 6, we conclude there exists a region T_1, where $\prod_{j=1}^{k} \bar{\psi}_{n,j}(i_j) > Q^*$, when $(i_1,\ldots,i_k) \in T_1$. Hence, the sum in (1) is taking over T_1. In this section, we find an algorithm for finding the region T_1 which is needed to calculate the two step ahead boundary.

Lemma 9. *If $\bar{\psi}_{n+1,j}\left(0, \bar{\psi}_{n,j}\right) < Q^*$ and $\lambda_{1,j} < \lambda_{2,j}$ for some j, then $M_n^{(1)} = 0$ and the one step ahead boundary Q^* is optimal.*

Proof. Since $\lambda_{1,j} < \lambda_{2,j}$, by Lemma 6, the posterior probabilities that a change has not occurred in category j, $\bar{\psi}_{n+1,j}\left(i_j, \bar{\psi}_{n,j}\right)$ are decreasing for all i_j. Thus, $\bar{\psi}_{n+1,j} < Q^*$. Hence for all i_j, $q_{n+1} < \bar{\psi}_{n+1,j} < Q^*$. Therefore, T_1 is empty, and $M_n = 0$. This makes the one step ahead boundary Q^* optimal in this case.

In the next theorem, we develop an algorithm for finding T_1.

Theorem 41 *The following algorithm finds T_1.*
1. For category 1, find I_1 such that $\bar{\psi}_{n,1}(i_1) > Q^*$ for all $i_1 \leq I_1$ if $\lambda_{1,1} < \lambda_{2,1}$, or for $i_1 \geq I_1$ if $\lambda_{1,1} > \lambda_{2,1}$.
2. For each i_1 in the range above, we can calculate the i_1 subset of T_1 using $Q_{i_1} = \frac{Q^*}{\bar{\psi}_{n,i}(i_1)}$.

Proof. We prove this by induction on k. If $k = 1$, we are only looking at one category. By Lemma 6, we find the 1 dimensional region that satisfies the condition. If $k > 1$, we find a $k-1$ dimensional region $T_1^{(k-1)}$ such that $\prod_{j=1}^{k-1} \bar{\psi}_{n+1,j}(i_j) > Q^*$ for all $\{i_1,\ldots,i_{k-1}\} \in T_1^{(k-1)}$. For each member $x \in T_1^{(k-1)}$, one can calculate $Q^*(x) = \frac{Q^*}{\bar{\psi}_{n+1,j}(i_j)}$. Then find a region $A_x \in Z$ such that $\bar{\psi}_{n+1,k}(i_j) > Q^*(x)$ for all $i_k \in A_x$. Thus the k dimensional region T_1 is $\bigcup_{x \in T_1^{(k-1)}} (x \cap A_x)$.

5 Appendix

We provide a numerical example demonstrating how this method works. In this example, we have two categories both with a change happening at time 10. In category one, the arrival rates before and after the change increases from an arrival rate of 1 to 2, whereas in category two, the arrival rates decreases from an arrival rate of 2 to 1. Assuming the cost for stopping late is $c = 0.06$, and the prior probability that there is a change at the next trial is $p_1 = p_2 = 0.01$, the one step boundary is calculated as $Q^* = 0.7509$. For each n, we calculate the probability that a change has not occurred in each category and the two step boundary b_n^2.

Table 6.1

n	$X_{1,n}$	$X_{2,n}$	$\bar{\psi}_{n,1}$	$\bar{\psi}_{n,2}$	q_n	$b_n^{(2)}$
1	1	1	0.9853	0.9731	0.9588	0.5806
2	1	1	0.9818	0.9509	0.9336	0.6045
3	1	2	0.9792	0.9594	0.9395	0.5988
4	1	1	0.9773	0.9330	0.9119	0.6247
5	0	0	0.9878	0.8166	0.8067	0.6891
6	3	4	0.9377	0.9613	0.9015	0.6309
7	0	0	0.9724	0.8787	0.8545	0.6660
8	1	5	0.9723	0.9875	0.9601	0.5789
9	1	0	0.9722	0.9413	0.9151	0.6214
10	1	1	0.9721	0.9097	0.8843	0.6473
11	2	1	0.9456	0.8695	0.8222	0.6862
12	4	4	0.7135	0.9733	0.6944	0.7305

From looking at the Table 6.1, both the one step and the two step ahead stopping rule stops after observation 12. This shows that the detection is quick. The table also shows the convergence of the two step ahead boundary to Q^*. We have provided 3 runs of simulations both using the 1 and 2 step procedure. In each of the simulations, category 1 increased from $\lambda_{1,1} = 1$ to $\lambda_{1,2} = 2$, which category 2 decreased from $\lambda_{2,1} = 2$ to $\lambda_{2,2} = 1$. Doing 1000 simulations for each of various cases of arrival rates before and after the change, we obtain the following results for probability of false alarms and expected delay.

Table 6.2

		1 Step Ahead			2 step ahead		
τ_1	τ_2	Alarms	Delay	Cost	Alarm	Delay	Cost
10	10	0.278	2.97	0.456	0.227	3.15	0.4157
10	15	0.254	4.14	0.502	0.237	4.33	0.4968
15	10	0.253	4.38	0.516	0.207	4.85	0.4980

In the above table we have provided examples where the change happens at the same spot ($\tau = 10$) for both categories, and where it happens at different spots for each category. The best results happen when the change happens at the same time for both categories. The procedure seems to stop quicker for increases in the arrival rate than for decreases.

References

Brown, M. and Zacks, S. (2006a) *Note on optimal Stopping for Possible Change in The Intensity of an Ordinary Poisson Process*. Probability and Statistics Letters, 76, 1417-1425.

Brown, M. and Zacks, S. (2006b) *On the Bayesian Detection of a Change in the Arrival Rate of a Poisson Process Monitored at Discrete Epochs* International Journal for Statistics and Systems., 1, 1-13.

Herberts, T. and Jensen, U. (2004) *Optimal Detection of a Change Point in a Poisson Process for Different Observation Schemes*. Scand. J. Stat., 31, 347-366.

Peskir, G. and Shiryaev A.N. (2002) *Solving the Poisson Disorder Problem*. Advances in Finance and Stochastics: Essays in Honor of Dieter Sonderman, pp. 295-312, Springer, New York.

Shiryaev, A.N. (1978) *Statistical Sequential Analysis: Optimal Stopping Rules*, Translations Math. Monographs, Vol. 38, Amer. Math. Soc., Providence, R.I.

Tartakovsky, A., Rozovskii, B., Blazek, R., and Kim, H. (2006) *Sequential Change-Point Detection Methods with Applications to Rapid Detection of Intrusions in Information Systems*. Sequential Methodology, 3, 252-293.

Zacks, S. (1983). *Survey of classical and Bayesian Approaches to the change-point problem: Fixed sample and sequential procedures for testing and estimation*. Recent Advances in Statistics (M.H. Rizvi, J. Rastogi and D. Siegmund, Ed.), Academic Press, New York.