JOURNAL OF QUALITY TECHNOLOGY
https://doi.org/10.1080/00224065.2018.1507557

RESEARCH PAPER

Monitoring communication outbreaks among an unknown team of actors in dynamic networks

Ross Sparks[a] and James D. Wilson[b]

[a]Private Bag 17, North Ryde, Sydney, NSW, Australia; [b]Department of Mathematics and Statistics, University of San Francisco, San Francisco, California

ABSTRACT

This article investigates the detection of communication outbreaks among a small team of actors in time-varying networks. We propose monitoring plans for known and unknown teams based on generalizations of the exponentially weighted moving average (EWMA) statistic. For unknown teams, we propose an efficient neighborhood-based search to estimate a collection of candidate teams. This procedure dramatically reduces the computational complexity of an exhaustive search. Our procedure consists of two steps: communication counts between actors are first smoothed using a multivariate EWMA strategy. Densely connected teams are identified as candidates using a neighborhood search approach. These candidate teams are then monitored using a surveillance plan derived from a generalized EWMA statistic. Monitoring plans are established for collaborative teams, teams with a dominant leader, as well as for global outbreaks. We consider weighted heterogeneous dynamic networks, where the expected communication count between each pair of actors is potentially different across pairs and time, as well as homogeneous networks, where the expected communication count is constant across time and actors. Our monitoring plans are evaluated on a test bed of simulated networks as well as on the U.S. Senate co-voting network, which models the Senate voting patterns from 1857 to 2015. Our analysis suggests that our surveillance strategies can efficiently detect relevant and significant changes in dynamic networks.

KEYWORDS

anomaly detection; exponentially weighted moving average; network surveillance; outbreak detection; statistical process control

1. Introduction

In many applications, it is of interest to identify anomalous behavior among the actors in a time-varying network. For example, in online social networks, sudden increased communications often signify illegal behavior such as fraud or collusion (Pandit et al. 2007; Savage et al. 2014). Anomalous changes like these are reflected by local structural changes in the network. The goal of network monitoring is to provide a surveillance plan that can detect such structural changes. Network monitoring techniques have been successfully utilized in a number of applications, including the identification of central players in terrorist groups (Krebs 2002; Porter and White 2012; Reid et al. 2005) and the detection of fraud in online networks (Akoglu and Faloutsos 2013; Chau et al. 2006; Pandit et al. 2007). As available data has become more complex, there has been a recent surge of interest in the development and application of scalable network monitoring methodologies (see Savage et al. (2014) and Woodall et al. (2017) for recent reviews).

In this article, we investigate monitoring the interactions of a fixed collection of n actors $[n] = \{1, \ldots, n\}$ over discrete times $t = 1, \ldots, T$. In general, an interaction is broadly defined and may represent, for example, communications in an online network (Prusiewicz 2008), citations in a co-authorship network (Liu et al. 2005), or gene-gene interactions in a biological network (Parker et al. 2015). We model the interactions of these actors at time t by an $n \times n$ stochastic adjacency matrix $\mathrm{Y}_t = (y_{i,j,t})$, where $y_{i,j,t}$ is the discrete random variable that represents the number of interactions between actor i and actor j at time t. Our goal is to develop a surveillance strategy to detect communication outbreaks among a subset of actors $\Omega_t \subseteq [n]$ at time t.

The identification of outbreaks among a subset of actors Ω_t corresponds to detecting sudden increases in the collection of edges $\{y_{i,j,t} : i, j \in \Omega_t\}$. When the

CONTACT James D. Wilson jdwilson4@usfca.edu Department of Mathematics and Statistics, University of San Francisco, 2130 Fulton St, San Francisco, CA 94117, USA.
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/ujqt.

team is unknown, monitoring can be computationally expensive due to the need for identifying candidate teams. For example, consider a simple case where we know the size of the target team is $n_{\Omega_t} = |\Omega_t|$. An exhaustive monitoring of all teams of size n_{Ω_t} requires a procedure of complexity $\binom{n}{n_{\Omega_t}} \approx n^{n_{\Omega_t}}$, which is infeasible even for moderately sized networks. As social networks are generally large, for example, n is on the order of 1 million for online networks like those representing Facebook or Twitter, exhaustive searches are not practical in real time. To address this challenge, we propose a computationally efficient local surveillance strategy that monitors the interactions of densely connected neighborhoods through time. Our proposed strategy has computational complexity of order n^2 and provides a viable strategy for large networks.

Our surveillance procedure consists of two steps, which can be briefly described as follows. First, we smooth the communication counts across all pairs and time using a multivariate adaptation of the exponentially weighted moving average (EWMA) technique for smoothing Poisson counts. By monitoring the smoothed counts, our strategy is robust to sudden random oscillations in the observed count process. Next, candidate teams are identified locally for each node using a neighborhood-based approach. In particular, at time t, we define a candidate team for node $i \in [n]$ as one that contains larger than expected communication. Surveillance plans for these candidate teams are developed using appropriate generalizations of the multivariate EWMA statistic.

We develop surveillance plans using the above technique in general for *heterogeneous* dynamic networks $\mathcal{Y} = \{\mathrm{Y}_1, \ldots, \mathrm{Y}_T\}$, where we suppose that the expected communication counts are possibly different for each pair and time, namely, $\mathbb{E}[y_{i,j,t}] = \lambda_{i,j,t}$. We consider three situations describing the team Ω_t:

(i) *Collaborative teams*: members of Ω_t communicate with one another far more than they communicate with actors outside of the team.
(ii) *Dominant leader teams*: the members of Ω_t have a dominant leader ν who communicates frequently with members of Ω_t, but the members of Ω_t themselves do not necessarily communicate frequently among themselves.
(iii) *Global outbreaks*: the entire network undergoes a communication outbreak, namely $\Omega_t \equiv [n]$.

Scenarios (i) and (ii) are considered for both unknown and known teams. Each of the scenarios are also considered for homogeneous networks, where $\mathbb{E}[y_{i,j,t}] \equiv \lambda$. By investigating both a test bed of simulated networks as well as a real network describing the U.S. Senate voting patterns, we find that our surveillance strategy can efficiently and reliably detect significant changes in dynamic networks.

1.1. Related work

The most closely related work to our current manuscript is that introduced in Heard et al. (2010). In that article, the authors also consider monitoring changes in communication volume between subgroups of targeted people over time. Their approach evaluates pairwise communication counts and determines whether these have significantly increased using a p-value, which assesses the deviation of the communication rate at time t and what is considered normal behavior. Here, normal behavior is modeled using conjugate Bayesian models for the discrete-valued time series of communications up to time t. While their focus is detecting changes on the entire network, our approach considers detecting communication outbreaks for members of a small team within the dynamic network.

There are other model-based network monitoring approaches that have been recently developed, which we briefly describe here. Azarnoush et al. (2016) proposed a longitudinal logistic model that describes the (binary) occurrence of an edge at time t as a function of time-varying edge attributes in the sequence of networks $\mathrm{G}([n], T)$. Likelihood-ratio tests of the fitted model are used to identify significant changes in $\mathrm{G}([n], T)$. Peel and Clauset (2014) developed a generalized hierarchical random graph model (GHRG) to model $\mathrm{G}([n], T)$. To detect anomalies, the authors used the GHRG as a null model to compare observed graphs in $\mathrm{G}([n], T)$ via a Bayes factor, which is calculated using bootstrap simulation. Wilson et al. (2016) proposed modeling and estimating change in a sequence of networks using the dynamic degree-corrected stochastic block model (DCSBM). In that work, maximum-likelihood estimates of the DCSBM are used for monitoring via Shewhart control charts. Our model is similar to the DCSBM in that edges are modeled as having discrete-valued edge weights, which flexibly model communications in social networks.

The EWMA control chart is a popular univariate monitoring technique. The multivariate EWMA process that we use here is a generalization of the univariate EWMA strategies for Poisson counts

considered in Weiß (2007, 2009), Sparks et al. (2009, 2010), and Zhou et al. (2012). A related multivariate EWMA control chart has previously been successfully applied to space-time monitoring of crime (Kim and O'Kelly 2008; Nakaya and Yano 2010; Neill 2009; Zeng et al. 2004).

A related problem to what we consider here is that of high-dimensional testing and monitoring over large collections of data streams (Liu et al. 2015; Mei 2010; Zou et al. 2015). Generally, the application of these methods involve continuous data that stream in over time and can be thought about as a dynamic process on a graph; whereas the current study investigates changes in the graph structure itself. In each of the cited studies, the authors monitored applications with on the order of hundreds of data streams, which computationally is much easier to handle than a collection of unordered individuals, as we consider here. Our application involves a maximum of 1000 by $1000 = 1{,}000{,}000$ data streams of counts and thus is an order of magnitude more complex than these aforementioned applications. The process of screening for members of a leader's team that we consider here when the team is unknown is reminiscent of the local CUSUM strategy of Mei (2011). Furthermore, our strategy is similar to the adaptive sampling strategy introduced in Liu et al. (2015), which relies on random sampling when the process is in-control and greedy sampling of the population when the process is out-of-control.

Our specified dynamic network model for $\mathcal{Y} = \{\mathrm{Y}_1, \ldots, \mathrm{Y}_T\}$ is related to several well-studied random graph models, which are ubiquitous in social network analysis. For example, when $y_{i,j,t}$ are independent and identically distributed Poisson(λ) random variables, the graph at time t is an Erdős-Rényi random graph model with edge connection probability λ (Erdös and Rényi 1960). On the other hand, when $y_{i,j,t}$ are independent Poisson$(\lambda_{i,j,t})$ random variables, graph t is a weighted variant of the Chung-Lu random graph model (Aiello et al. 2000). Random graph models play an important role in the statistical analysis of relational data. Goldenberg et al. (2010) provide a recent survey about random graph models and their applications.

1.2. Organization of this article

The remainder of this article is organized as follows. In Section 2, we describe how to smooth the observed communication counts using multivariate EWMA smoothing. In Section 3, we develop surveillance strategies for communication outbreaks among small teams of actors in a dynamic network when the target team is known. We consider collaborative teams, dominant leader teams, as well as global outbreaks. Section 4 describes our proposed local search and monitoring approach for unknown target teams. Section 5 investigates the performance of our surveillance strategies on a test-bed of simulated networks. We make recommendations on designing the plans in such a way to minimize false discovery. In Section 6, we further assess the performance of our strategy by applying the plans to the heterogeneous network describing the U.S. Senate voting patterns from the 35th to the 113th Congress. We discuss the advantages of utilizing the square-root transform for our proposed monitoring strategies in Section 7. We conclude with a summary of our findings and discuss directions for future work in Section 8.

2. Temporal EWMA smoothing of interactions

Throughout this work, we are concerned with detecting significant increases in communication among the members of some subset of actors $\Omega_t \subseteq [n]$. Such fluctuations correspond to sudden spikes in the collection of edge weights $\{y_{i,j,t} : i,j \in \Omega_t\}$. In many cases, the communication counts $\{y_{i,j,t} : i,j \in [n], t = 1, \ldots, T\}$ are prone to random fluctuations that arise from noise in the observed process. If not accounted for, direct monitoring of counts may lead to false discovery. To reduce this possibility, we smooth the observed counts using a reflective EWMA strategy (Gan 1993).

To begin, we first obtain a collection of smoothed values $\{\widetilde{y}_{i,j,t} : i,j \in [n], t = 1, \ldots, T\}$ using an EWMA strategy. Fix $\alpha \in [0,1]$, and define

$$\widetilde{y}_{i,j,t} = \alpha\ y_{i,j,t} + (1-\alpha)\ \widetilde{y}_{i,j,t-1}. \qquad [1]$$

Denote the expected value of $\widetilde{y}_{i,j,t}$ by $\widetilde{\lambda}_{i,j,t}$. The expected values of these smoothed counts can be calculated using the following recursion:

$$\widetilde{\lambda}_{i,j,t} = \alpha\ \lambda_{i,j,t} + (1-\alpha)\ \widetilde{\lambda}_{i,j,t-1}.$$

In the above recursion, the initial values are set as $\widetilde{y}_{i,j,0} = \widetilde{\lambda}_{i,j,0} = \lambda_{i,j,1}$. Here, α acts as a smoothing parameter that dictates the temporal memory retained in the stochastic process $\{\widetilde{y}_{i,j,t} : i,j \in [n], t = 1, \ldots, T\}$. Large values of α retain less memory and result in less smoothing. In our applications, we fix α to 0.075 based on the previous analysis and suggestion of Sparks and Patrick (2014).

Notably, the EWMA in [1] will not reflect a change in the observed count process in the scenario that $y_{i,j,t}$

decreases immediately before a significant (anomalous) increase. To avoid this worst-case scenario, we use the reflective boundary EWMA process $\{y^*_{i,j,t} : i,j \in [n], t = 1,\ldots,T\}$, defined by

$$y^*_{i,j,t} = \max\left(\alpha\ \widetilde{y}_{i,j,t} + (1-\alpha)\ y^*_{i,j,t}, \widetilde{\lambda}_{i,j,t}\right) \qquad [2]$$

The reflective boundary EWMA specified in [2] is robust to sudden oscillations in the count process. Our surveillance plans will utilize the smoothed counts from [2] rather than the originally observed counts.

3. Monitoring a known team of actors

We begin by considering the simplest case when the target team Ω_t is known *a priori*. This scenario arises, for example, in the surveillance of the communication among a known active group of terrorists in a large terrorist network. We develop surveillance plans for collaborative and dominant leader teams, as well as global changes, where the entire network undergoes a communication outbreak. For each of these scenarios, we describe monitoring a homogeneous sequence of networks $\mathcal{Y}$, where the collection of expected communications $\{\lambda_{i,j,t} : i,j \in [n], t = 1,\ldots,T\}$ are such that $\lambda_{i,j,t} \equiv \lambda$ for all i, j and t, and further describe how to extend the plans in this regime to the more general heterogeneous case, where expected communications are possibly different across time and actor pairs.

In both this section and Section 4, we will make use of two tunable parameters – $\alpha \in [0,1]$: a smoothing parameter that controls the extent to which a proposed EWMA statistic has temporal memory, and $h(\cdot,\cdot)$: threshold functions that are chosen to control false discovery of the proposed monitoring plan. We fix $\alpha = 0.075$ based on previous analysis conducted in Sparks and Patrick (2014). The threshold functions $h(\cdot,\cdot)$ are chosen via simulation of the monitored process. We describe how these are chosen in detail in the Appendix.

Throughout this and the following section, let $\widetilde{y}_{i,j,t}$ and $y^*_{i,j,t}$ be the EWMA and reflective boundary EWMA defined in [1] and [2], respectively. Further, we denote $n_{\Omega_t} = |\Omega_t|$ as the number of individuals in the team.

3.1. Ω_t is a collaborative team

We first consider monitoring for outbreaks among a collaborative team Ω_t, wherein all members of Ω_t are expected to communicate regularly. An outbreak in a collaborative team is reflected by a large average number of communications between members $i,j \in \Omega_t$. To detect such outbreaks, we analyze the mean, μ_{Ω_t}, of the smoothed interactions in the collection defined as

$$\mu_{\Omega_t} = \mathbb{E}\left[\sum_{i\in\Omega_t}\sum_{j\in\Omega_t}\widetilde{y}_{i,j,t}\right] = \sum_{i\in\Omega_t}\sum_{j\in\Omega_t}\widetilde{\lambda}_{i,j,t} \qquad [3]$$

In the case that $\mathcal{Y}$ is homogeneous, note that $\mu_{\Omega_t} = n^2_{\Omega_t}\lambda$. We use a group-EWMA (GEWMA) statistic to identify outbreaks among the actors in Ω_t. The GEWMA_t process is defined by the following recursion:

$$\text{GEWMA}_t = \max\left(\alpha\sum_{i\in\Omega_t}\sum_{j\in\Omega_t}\widetilde{y}_{i,j,t} + (1-\alpha)\ \text{GEWMA}_{t-1},\ \mu_{\Omega_t}\right), \qquad [4]$$

where the initial value $\text{GEWMA}_1 = \sum_{i\in\Omega_t}\sum_{j\in\Omega_t}\widetilde{y}_{i,j,1}$.

For homogeneous networks, we use the GEWMA_t process from [4] and flag an outbreak within the team Ω_t when

$$\sqrt{\text{GEWMA}_t} - n_{\Omega_t}\sqrt{\lambda} > h_G(\lambda, n_{\Omega_t}), \qquad [5]$$

where $h_G(n_{\Omega_t},\lambda)$ is designed to give the plan a low false-discovery rate. Importantly, the square-root transform of the GEWMA_t process in [5] stabilizes the variance of the process (see Bartlett 1936). We find from simulation, which we discuss in detail in Section 7, that the threshold $h_G(n_{\Omega_t},\lambda)$ is *not* a function of λ nor a function of the group size nor network size in the case of our application. Hence, even in the heterogeneous case, we can use a plan with the threshold $h_G(n_{\Omega_t})$. We describe how to choose the value $h_G(n_{\Omega_t})$ in the Appendix. Thus, for heterogeneous networks, we flag an outbreak in the team Ω_t when

$$\sqrt{\text{GEWMA}_t} - \sqrt{\sum_{i\in\Omega_t}\sum_{j\in\Omega_t}\widetilde{\lambda}_{i,j,t}} > h_G(n_{\Omega_t}). \qquad [6]$$

In practice, a target team Ω_t may purposefully reduce their communication levels prior to, say, planning a crime, which may hamper early detection when using the GEWMA_t statistic defined in [4]. To avoid this scenario, one can alternatively use a reflective boundary GEWMA statistic defined as

$$\text{GEWMA}^*_t = \sum_{i\in\Omega_t}\sum_{j\in\Omega_t} y^*_{i,j,t}, \qquad [7]$$

and apply an analogous plan as defined in [6].

3.2. Ω_t has a dominant leader

We now consider the scenario in which the target team Ω_t has a known dominant leader $\nu \in [n]$. We expect that ν will have a high level of communication with the members of Ω_t, but unlike the collaborative team setting, the members of Ω_t do not necessarily significantly interact with one another. In this case, an outbreak is signaled when there is either a significant rate of communications between ν and the members of Ω_t or by a significant rate of interactions among the members of Ω_t. As we primarily need to be concerned with the communications between a single actor and a collection of actors, we develop a monitoring strategy that exploits sparsity in the interactions among the members of Ω_t. At time t, we monitor only the collection of actors that (a) significantly communicate with the dominant leader ν and (b) significantly communicate with one another. That is, we identify the dominant leader team Ω_t by following two steps. First, we identify the team $W_{\nu,t}$ that contains all individuals in $[n]$ with a significant number of interactions with ν, namely

$$W_{\nu,t} = \left\{ i \neq \nu \in [n] : \sqrt{y^*_{\nu,i,t} + y^*_{i,\nu,t}} - \sqrt{\widetilde{\lambda}_{\nu,i,t} + \widetilde{\lambda}_{i,\nu,t}} > k \right\}. \quad [8]$$

Next we refine the team $W_{\nu,t}$ to include only those members who share a significant number of communications. We set

$$\Omega_t = \left\{ i,j \in W_{\nu,t} : \sqrt{y^*_{i,j,t}} - \sqrt{\widetilde{\lambda}_{i,j,t}} > k \text{ or } \sqrt{y^*_{j,i,t}} - \sqrt{\widetilde{\lambda}_{j,i,t}} > k \right\}. \quad [9]$$

The value k is a suitable constant that helps identify members of the target group and is chosen to control the size of the team Ω_t. We consider the choice of k in our simulation study in Section 5. To monitor Ω_t, we use the dominant leader EWMA (DEWMA) statistic, defined as

$$\text{DEWMA}_{\nu,t} = \sum_{i \in W_{\nu,t}} \left(y^*_{i,\nu,t} + y^*_{\nu,i,t} \right) + \sum_{i \in \Omega_t} \sum_{j \in \Omega_t} y^*_{i,j,t}. \quad [10]$$

When ν is known, we can use the DEWMA statistic from [10] to flag outbreaks in a dominant leader team. In the case that $\mathcal{Y}$ is homogeneous, we flag an outbreak when

$$\sqrt{\text{DEWMA}_{\nu,t}} - \sqrt{2 n_{W_{\nu,t}} \lambda + n^2_{\Omega_t} \lambda} > h_D(n_{\lambda,\Omega_t}, \lambda). \quad [11]$$

Above, $h_D(n_{\Omega_t}, \lambda)$ is chosen to control false discovery. Once again, simulations suggest that the square-root transformation rids the dependence of the threshold $h_D(n_{\Omega_t}, \lambda)$ on λ. In addition, this threshold is not a function of the group size or the sparsity of the graph, for example, sparsity can change from one time point to the next and the threshold will remain the same. Thus, we use the following general surveillance plan for heterogeneous networks when ν is known

$$\sqrt{\text{DEWMA}_{\nu,t}} - \sqrt{\sum_{i \in W_{\nu,t}} \left(\widetilde{\lambda}_{i,\nu,t} + \widetilde{\lambda}_{\nu,i,t} \right) + \sum_{i \in \Omega_t} \sum_{j \in \Omega_t} \widetilde{\lambda}_{i,j,t}} > h_D(n_{\Omega_t}), \quad [12]$$

We note that, when the team and dominant leader are both unknown, the plan in [12] is complicated by the fact that we must estimate ν and Ω_t. We discuss our strategy to handle this in Section 4.

3.3. Global outbreaks

We now consider the case when there is a significant increase in the number of interactions among every pair of actors in the network, that is, when $\Omega_t \equiv [n]$ for all t. One can generally detect this anomaly early by monitoring the aggregated interactions over the target network. To monitor the network for a global outbreak, one can directly extend the GEWMA$_t$ statistic from [4] to the entire network. Note that , in the case that $\Omega_t \equiv [n]$, we have from [3] that $\mu_{[n]} = \sum_{i \in [n]} \sum_{j \in [n]} \widetilde{\lambda}_{i,j,t}$. Following our previous development of the GEWMA$_t$ statistic in [4], we define the total-EWMA (TEWMA) statistic using the following recursion

$$\text{TEWMA}_t = \max\left(\alpha \sum_{i \in [n]} \sum_{j \in [n]} \widetilde{y}_{i,j,t} + (1-\alpha)\ \text{TEWMA}_{t-1},\ \mu_{[n]} \right), \quad [13]$$

where $\text{TEWMA}_1 = \sum_{i \in [n]} \sum_{j \in [n]} \widetilde{y}_{i,j,1}$, and $\alpha \in [0,1]$ are chosen to smooth the TEWMA process. Using the statistic in [13], we flag a global outbreak in homogeneous networks when

$$\sqrt{\text{TEWMA}_t} - n\sqrt{\lambda} > h_T(\lambda, n). \quad [14]$$

The threshold $h_T(n, \lambda)$ is designed to give the plan a low enough false-discovery rate and is chosen in the same manner as plan [5]. In general, we flag an outbreak in heterogeneous networks when

$$\sqrt{\text{TEWMA}_t} - \sqrt{\sum_{i \in [n]} \sum_{j \in [n]} \widetilde{\lambda}_{i,j,t}} > h_T. \quad [15]$$

To avoid issues arising from sudden oscillations in counts, we can instead use the reflected-boundary TEWMA statistic

$$\text{TEWMA}_t^* = \sum_{i\in[n]}\sum_{j\in[n]} y_{i,j,t}^*, \qquad [16]$$

and apply the plan given in [15].

4. Monitoring of an unknown team of actors

In many applications, Ω_t is *not* known *a priori*. In this situation, there are two primary difficulties that one must address. First, the unknown team must be efficiently estimated. An exhaustive search for an anomalous team has complexity of order $n^{n_{\Omega_t}}$; thus, it is important to employ scalable approaches for estimation. When Ω_t is known, the GEWMA_t and $\text{DEWMA}_{\nu,t}$ statistics are invariant to variations in the communication means. However, when Ω_t is unknown, these statistics are no longer invariant to heterogeneous communication rates through time. Thus, the second complication comes in adapting the monitoring plan for a changing mean in heterogeneous networks. In this section, we describe a local search strategy to identify densely connected teams on which our proposed statistics can be used for monitoring. Because the global outbreak plan in [15] is invariant to mean changes, we only need to consider the scenarios when Ω_t is either a collaborative team or a dominant leader team.

4.1. Estimating unknown teams

Here, we describe our local search strategy to estimate collaborative teams as well as teams with a dominant leader.

4.1.1. Collaborative teams

When the target team is unknown and collaborative, we propose monitoring a collection of densely connected teams $\mathbf{\Omega}_{C,t} := \{\widehat{\Omega}_{\ell,t} : \ell \in [n]\}$ at each time t. We define a candidate team $\widehat{\Omega}_{\ell,t}$ as one in which all constituent members significantly interact. In particular, for each $\ell \in [n]$ and each time t, we identify the candidate team

$$\widehat{\Omega}_{\ell,t} = \left\{ i \in [n] : \sqrt{y_{i,\ell,t}^*} - \sqrt{\widetilde{\lambda}_{i,\ell,t}} > k, \text{ or } \sqrt{y_{\ell,i,t}^*} - \sqrt{\widetilde{\lambda}_{\ell,i,t}} > k \right\}. \qquad [17]$$

Above, k is a suitable constant with good detection properties and is chosen via simulation. Our specification of each candidate team $\widehat{\Omega}_{\ell,t}$ is motivated by empirical properties of real networks. One can view $\widehat{\Omega}_{\ell,t}$ structurally as a hub with center node ℓ. Hub structures commonly arise in sparse social and biological networks as well as the well-studied scale-free family of networks (Barabási and Albert 1999; Tan et al. 2014). Thus, if the unknown team is suspected to be a collaborative team, we propose monitoring at most n densely connected teams.

4.1.2. Dominant leader teams

When the dominant leader ν and target team Ω_t are unknown, we monitor a collection of candidate dominant leader teams $\mathbf{\Omega}_{D,t} := \{\widehat{\Omega}_{\nu,t} : \nu \in [n]\}$ at each time t. Like the identification of dominant leader teams in Section 3, we identify a collection of candidate dominant leader teams that have a significantly large rate of communication. First, for a fixed leader $\nu \in [n]$, we identify a team $\widehat{W}_{\nu,t}$ by finding all individuals in $[n]$ with a significant number of interactions with ν given by

$$\widehat{W}_{\nu,t} = \left\{ i \neq \nu \in [n] : \sqrt{y_{\nu,i,t}^* + y_{i,\nu,t}^*} - \sqrt{\widetilde{\lambda}_{\nu,i,t} + \widetilde{\lambda}_{i,\nu,t}} > k \right\}. \qquad [18]$$

We next refine the team $\widehat{W}_{\nu,t}$ to include only those members who share a significant number of interactions. Namely, we specify the team $\widehat{\Omega}_{\nu,t}$ as

$$\widehat{\Omega}_{\nu,t} = \left\{ i,j \in \widehat{W}_{\nu,t} : \sqrt{y_{i,j,t}^*} - \sqrt{\widetilde{\lambda}_{i,j,t}} > k \text{ or } \sqrt{y_{j,i,t}^*} - \sqrt{\widetilde{\lambda}_{j,i,t}} > k \right\} \qquad [19]$$

The value k is a suitable constant that helps identify members of the target group with larger than expected communications with the dominant leader ν. We note that, rather than a normal standardized score to identify Ω_t, we use a 'signal-to-noise' team identification scheme in [18], as this strategy can efficiently avoid unusual changes that involve very low communication levels. In the case that the team is unknown, thresholds depend on both the team size as well as the total expected communication when the network is in-control.

4.2. Adapting the plans for heterogeneous networks

Once the candidate teams $\mathbf{\Omega}_{C,t} = \{\widehat{\Omega}_{\ell,t} : \ell \in [n]\}$ and $\mathbf{\Omega}_{D,t} = \{\widehat{\Omega}_{\nu,t} : \nu \in [n]\}$ have been estimated for each time t, we can develop a monitoring plan. For $\ell, \nu \in [n]$, define the following local GEWMA and DEWMA statistics:

$$\mathrm{GEWMA}^{*}_{\ell,t} = \sum_{i\in\widehat{\Omega}_{\ell,t}}\sum_{j\in\widehat{\Omega}_{\ell,t}} y^{*}_{i,j,t} \quad [20]$$

$$\mathrm{DEWMA}^{*}_{\nu,t} = \sum_{i\in\widehat{W}_{\nu,t}} \left(y^{*}_{i,\nu,t} + y^{*}_{\nu,i,t}\right) + \sum_{i\in\widehat{\Omega}_{\nu,t}}\sum_{j\in\widehat{\Omega}_{\nu,t}} y^{*}_{i,j,t}. \quad [21]$$

When the observed network is homogeneous, one can readily monitor collaborative and dominant leader teams by using plans [6] and [11], respectively, for the local GEWMA and DEWMA statistics in [20] and [21]. When the network is heterogeneous, we develop an adaptive plan for surveillance as follows. Note that, for a fixed candidate collaborative team $\widehat{\Omega}_{\ell,t}$, the plan in [6] can be re-expressed as

$$\sqrt{\mathrm{GEWMA}^{*}_{\ell,t}/h^{2}_{G}\left(\lambda_{i,j,t}, n_{\widehat{\Omega}_{\ell,t}}\right)} - \sqrt{\sum_{i\in\widehat{\Omega}_{\ell,t}}\sum_{j\in\widehat{\Omega}_{\ell,t}} \lambda_{i,j,t}/h^{2}_{G}\left(\lambda_{i,j,t}, n_{\widehat{\Omega}_{\ell,t}}\right)} > 1 \quad [22]$$

The threshold in plan [22] no longer depends on the observed data. We exploit this property and define an adaptive plan using the local adaptive group-EWMA (AGEWMA) statistic:

$$\mathrm{AGEWMA}_{\ell,t} = \mathrm{GEWMA}^{*}_{\ell,t}/h^{2}_{G}\left(\widetilde{\lambda}_{i,j,t}, n_{\widehat{\Omega}_{\ell,t}}\right). \quad [23]$$

For an unknown team Ω_t, a communication outbreak is flagged when

$$\sqrt{\mathrm{AGEWMA}_{\ell,t}} - \sqrt{\sum_{i\in\widehat{\Omega}_{\ell,t}}\sum_{j\in\widehat{\Omega}_{\ell,t}} \widetilde{\lambda}_{i,j,t}/h^{2}_{G}\left(\widetilde{\lambda}_{i,j,t}, n_{\widehat{\Omega}_{\ell,t}}\right)} > 1, \quad [24]$$

for any $\ell \in [n]$. Here, the team must be re-estimated at each time period t. This adaptive plan in [24] has the same in-control ATS value used to design the homogeneous plans for all $\lambda_{i,j,t}$.

We can use a similar adaptive plan to identify communication outbreaks in candidate dominant leader teams. Define the local adaptive dominant leader - EWMA (ADEWMA) statistic by

$$\mathrm{ADEWMA}_{\nu,t} = \sum_{i\in\widehat{W}_{\nu,t}} \left(\frac{y^{*}_{i,\nu,t}}{h_D\left(\widetilde{\lambda}_{i,\nu,t}, n_{\widehat{\Omega}_{\nu,t}}\right)} + \frac{y^{*}_{j,\nu,t}}{h_D\left(\widetilde{\lambda}_{j,\nu,t}, n_{\widehat{\Omega}_{\nu,t}}\right)}\right) + \sum_{i\in\widehat{\Omega}_{\nu,t}}\sum_{j\in\widehat{\Omega}_{\nu,t}} \frac{y^{*}_{i,j,t}}{h_D\left(\widetilde{\lambda}_{i,j,t}, n_{\widehat{\Omega}_{\nu,t}}\right)}. \quad [25]$$

Using an analogous argument as above for the adaptive GEWMA plan, we flag a communication outbreak among dominant leader teams when

$$\sqrt{\mathrm{ADEWMA}_{\nu,t}} - \sqrt{\sum_{i\in\widehat{W}_{\nu,t}} \left(\frac{\widetilde{\lambda}^{*}_{i,\nu,t}}{h_D\left(\widetilde{\lambda}_{i,\nu,t}, n_{\widehat{\Omega}_{\nu,t}}\right)} + \frac{\widetilde{\lambda}^{*}_{j,\nu,t}}{h_D\left(\widetilde{\lambda}_{j,\nu,t}, n_{\widehat{\Omega}_{\nu,t}}\right)}\right)} + \sqrt{\sum_{i\in\widehat{\Omega}_{\nu,t}}\sum_{j\in\widehat{\Omega}_{\nu,t}} \frac{\widetilde{\lambda}^{*}_{i,j,t}}{h_D\left(\widetilde{\lambda}_{i,j,t}, n_{\widehat{\Omega}_{\nu,t}}\right)}} > 1 \quad [26]$$

for any $\nu \in [n]$. There are two distinct scenarios in which an outbreak will be flagged by the plan [26]. In the first scenario, an outbreak is detected if the team size of any candidate team significantly increases. This is likely to happen when, for instance, a leader of organized crime is trying to recruit a team. In the second scenario, an outbreak is detected when the number of interactions within any candidate team significantly increases. This can occur in two ways: (i) when individuals within the same team interact more with individuals outside of their current group or (ii) members of the group interact significantly more frequently among themselves. Combinations of (i) and (ii) may also flag communication outbreaks.

5. Simulation study

We now access the utility of our proposed surveillance plans on a test bed of simulated networks. We consider two types of communication outbreaks among small target teams. In the first scenario, we simulate a collaborative team outbreak where every actor in a small and unknown team is involved in the outbreak. In the second scenario, the target team has an unknown dominant leader whose communication levels with the remaining team undergoes an outbreak. For each of these cases, we investigate the effectiveness of the GEWMA and DEWMA strategies.

For each simulation, we generate 100 in-control networks followed by 500 networks that have undergone an outbreak. We record the time to signal - the number of networks after the change until a signal is flagged - of the DEWMA and GEWMA plans and repeat the experiment 10,000 times for the collaborative team outbreak and 1000 times for the dominant leader outbreak. To evaluate the performance of a plan, we record the average time to signal (ATS)

over the collection of simulations. We present the results for all simulations in Tables 3 through 15 in the Appendix.

5.1. Collaborative team outbreaks

Tables 1 through 10 outline the detection properties of simulated collaborative team outbreaks for networks of size $n = 100$. To simulate an outbreak, we select a fixed but hidden team $\Omega \subseteq \{1, \ldots, 100\}$. In the first 100 in-control networks, communication counts among the nodes in Ω have mean λ. In the remaining networks, the nodes in Ω have an increased mean communication count of $(1+\delta)\lambda$. We simulate networks with target teams of size $n_\Omega = 6, 7, 8, 9,$ and 10. For each time series of networks, we estimate candidate collaborative teams and dominant leader teams via [17] and [19] and then apply the GEWMA and DEWMA plans from [5] and [11], respectively.

Table 1. Comparison of the CUSUM, EWMA, and TEWMA plans. Each plan was trained to have an in-control ATS of 100. Results are the average number of time steps until the plan is flagged from 1000 simulations.

Mean out-of-control data	CUSUM w/ alt. mean rate 7	EWMA	TEWMA
Poisson mean 5			
6.0	16.64	15.18	15.41
6.5	9.68	9.21	9.65
7.0	6.60	6.43	7.01
7.5	4.95	4.97	5.56
8.0	4.05	4.05	4.59
8.5	3.39	3.46	3.94
9.0	2.96	3.02	3.46
10.0	2.30	2.45	2.82
Poisson mean 40			
43	14.59	14.01	13.84
44	9.88	9.63	9.91
45	7.24	7.24	7.71
46	5.70	5.85	6.32
47	4.72	4.77	5.33
48	3.96	4.18	4.65
49	3.44	3.66	4.12
50	3.09	3.28	3.74
Poisson mean 80			
84	14.74	14.96	14.66
85	11.05	11.05	11.37
86	8.62	8.94	9.15
87	6.97	7.36	7.62
88	5.80	6.22	6.72
89	4.99	5.38	5.85
90	4.44	4.73	5.26
91	3.97	4.24	4.76
92	3.63	3.87	4.33
Poisson mean 200			
206	15.18	15.92	15.47
207	12.42	13.13	13.11
208	10.44	11.02	11.27
209	8.95	9.45	9.91
210	7.78	8.19	8.52
211	6.83	7.28	7.72
212	6.11	6.53	6.96
213	5.54	5.88	6.45
214	4.95	5,35	5.85
215	4.64	4.99	5.43

5.1.1. The GEWMA$_t$ plan

In the first part of our study, we simulate homogeneous target networks with mean communication counts of either $\lambda = 0.20$ or 0.70. We investigate significance thresholds k between 0.05 and 0.40 in increments of 0.05. Table 4 explores changes in communication counts in a team of size 6. Table 1 reveals that $k = 0.40$ provides the best performance for both λ values.

We extend the first simulation to seek the best plan for detecting the collaborative team Ω, when $n_\Omega = 6$ and $n = 100$. We investigate significance thresholds of k between 0.40 and 0.70 for expected communication rates of $\lambda = 0.20, 0.40$ and 0.70. Together, Tables 4 and 5 indicate that $k = 0.60$ is the best choice for all λ and Ω involving 6 of the 100 actors. Furthermore, we find that the performance of the GEWMA plan strongly depends on an appropriate choice of k; the detection performance of the GEWMA plan is dramatically improved for $k = 0.60$.

We repeat the collaborative team outbreak simulation for a target team of size 7, 8, 9, and 10. In each simulation, we seek the best significance threshold k for homogeneous networks with mean communication $\lambda = 0.20, 0.40$, and 0.70. We report the ATS over 10,000 simulations for each of these settings in Tables 4 through 9. Our results suggest that $k = 0.50$ is the best choice for all λ when n_Ω is 8, 9, or 10, while $k = 0.50$ or 0.60 is most suitable for networks where the target team is of size 7. This result suggests that there is an inverse relationship between the optimal value of k and the size of the target team. This is helpful in deciding the choice of k for the GEWMA plan, and it appears that $k = 0.50$ is a robust choice for the outbreaks considered in this study.

5.1.2. The DEWMA$_{\nu,t}$ plan

Tables 10 through 13 report the results of the DEWMA surveillance plan on the collaborative team outbreaks described above for target teams of size 6, 7, 8, and 9. For each setting, $k = 0.45$ tends to be the best choice for significance threshold. The only exception is in the case that the team is of size 9 and the mean communication is $\lambda = 0.70$, in which case $k = 0.40$ is the better choice.

Table 2. The in-control ATS for each network in the Congressional co-voting network application for thresholds set to 0.513 and 0.399. The ATS remains stable across time for each threshold. These results reveal that the threshold is not a function of network or party size. The first 40 congresses are shown, though the results are consistent for the remaining congresses.

Congress	0.513 Threshold	0.399 Threshold	Democrats	Republicans	Other
35	402.3	100.2	47	21	
36	397.3	102.3	41	27	
37	396.3	102.1	27	43	
38	399.8	100.6	10	44	
39	398.7	101.7	11	48	
40	400.1	100.2	12	57	
41	402.1	99.6	13	67	
42	399.3	96.8	18	57	
43	399.1	101.8	21	58	
44	399.9	100.5	34	47	1
45	399.2	101.8	38	43	1
46	401.7	99.8	45	35	1
47	401.2	101.1	39	43	1
48	400.1	100.9	38	40	
49	399.7	99.9	36	45	
50	402.1	101.0	37	39	
51	400.5	100.3	38	53	
52	398.7	101.9	44	47	2
53	402.2	100.5	48	43	3
54	398.3	101.2	40	44	4
55	400.6	102.5	38	46	10
56	399.1	101.3	27	56	8
57	398.2	100.7	32	57	1
58	402.4	101.8	33	60	
59	401.3	99.5	33	60	
60	397.8	100.4	32	63	
61	400.3	100.4	40	62	
62	398.8	99.1	53	56	
63	399.8	100.4	56	44	1
64	399.7	100.7	58	42	
65	402.0	100.5	62	49	
66	402.8	99.3	50	51	
67	401.6	97.4	39	66	
68	398.2	100.1	43	57	2
69	402.7	102.3	43	61	1
70	401.5	100.8	48	53	1
71	399.6	101.7	44	64	1
72	402.6	102.8	52	50	1
73	399.2	99.9	63	36	1
74	398.0	98.8	72	25	3
75	400.7	101.0	82	16	4

5.1.3. Comparison of the GEWMA$_t$ and DEWMA$_{\nu,t}$ plans

In comparing the results for the GEWMA and DEWMA plans on the collaborative team outbreak simulation, we find that, in general, the GEWMA plan outperforms the DEWMA plan. In particular, the GEWMA strategy detects the collaborative team sooner than its counterpart. For example, when $\delta = 1$ and $\lambda = 0.2$, the strategy based on GEWMA$_t$ in Table 2 had an ATS equal to 11.62 ($k = 0.60$) whereas the technology based on DEWMA$_{\nu,t}$ in Table 7 had an ATS equal to 12.90 ($k = 0.45$). Similarly, when $\delta = 0.50$ and $\lambda = 0.70$; the GEWMA$_t$ strategy had an ATS equal to 8.54 ($k = 0.50$); whereas, the DEWMA$_{\nu,t}$ plan had an ATS of 8.87 ($k = 0.40$).

Table 3. Comparison of results from the application of the square transformation after the sum (our proposed method) and before the sum of the counts on a network of 50 actors with in-control homogeneous means across all actors.

Poisson mean 0.001		
Threshold	0.531	2.6446
Delta (shift)	Sum then square root	Square root then sum
0	441.359	407.373
0.0001	121.959	134.363
0.0002	49.146	55.492
0.0003	27.079	28.726
0.0004	16.917	18.853
0.0005	11.939	13.202
0.0006	9.867	9.773
0.0007	7.499	7.504
0.0008	6.611	6.354
Poisson mean 0.005		
Threshold	0.531	2.6446
Delta (shift)	Sum then square root	Square root then sum
0	439.47	407.129
0.00025	105.368	120.705
0.0005	38.082	47.608
0.00075	21.454	23.374
0.00100	13.037	15.309
0.00125	9.648	10.372
0.0015	7.934	8.007
0.002	5.316	5.305
0.0025	4.205	3.994
0.003	3.528	3.240
Poisson mean 0.025		
Threshold	0.531	8.20532
Delta (shift)	Sum then square root	Square root then sum
0	417.04	400.98
0.0005	115.799	116.017
0.001	47.450	43.729
0.002	15.444	14.958
0.003	9.016	8.378
0.004	5.948	5.926
0.005	4.512	4.687
0.006	3.643	3.758
0.007	3.185	3.199
0.008	2.781	2.855
Poisson mean 0.05		
Threshold	0.531	13.685
Delta (shift)	Sum then square root	Square root then sum
0	399.97	400.38
0.001	77.833	68.323
0.002	26.04	24.305
0.003	13.295	13.371
0.004	9.227	9.118
0.005	6.826	6.959
0.006	5.483	5.606
0.007	4.542	4.616
0.008	3.946	3.996
0.010	3.116	3.187
Poisson mean 0.5		
Threshold	0.531	23.9
Delta (shift)	Sum then square root	Square root then sum
0	420.17	401.50
0.003	78.298	97.075
0.005	37.338	43.236
0.007	21.641	24.998
0.009	15.139	16.774
0.012	9.533	10.162
0.015	7.026	7.560

(*Continued*)

Table 3. Continued.

Poisson mean 0.5		
Threshold Delta (shift)	0.531 Sum then square root	23.9 Square root then sum
0.018	5.905	6.011
0.021	4.832	5.073
0.024	4.196	4.190
Poisson mean 1.0		
Threshold Delta (shift)	0.531 Sum then square root	46.5 Square root then sum
0	400.08	398.185

(*Continued*)

Table 3. Continued.

Poisson mean 1.0		
Threshold Delta (shift)	0.531 Sum then square root	46.5 Square root then sum
0.005	62.545	57.482
0.010	21.455	20.675
0.015	11.793	11.774
0.020	7.954	8.463
0.025	5.757	6.245
0.030	4.842	5.135
0.035	4.014	4.357
0.040	3.417	3.738
0.050	2.728	3.069

Table 4. Collaborative team ATS performance for GEWMA$_t$ with $n_{\Omega_t} = 6$.

	Communication outbreaks in team of size 6 from a network of size 100															
λ	0.2								0.7							
k	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4
δ	ATS															
0.5	70.60	75.33	76.30	74.83	74.11	72.61	53.71	43.36	55.29	57.67	58.23	55.59	47.71	33.97	23.70	17.41
1.0	50.25	56.20	57.65	48.34	45.86	34.02	20.27	16.08	27.69	33.63	37.48	34.24	24.05	14.30	9.19	6.59
2.0	28.15	33.89	34.26	28.32	19.64	12.62	8.05	6.16	15.57	18.32	19.46	16.79	11.60	6.60	4.29	3.09
3.0	19.12	23.48	23.51	19.35	12.67	7.87	5.20	3.95	10.32	12.23	12.96	11.01	7.77	4.50	2.89	2.12
4.0	14.70	15.10	17.46	14.37	9.91	5.78	3.74	2.94	8.23	9.57	10.18	8.22	5.69	3.52	2.24	1.80
5.0	12.06	14.30	13.88	11.30	7.56	4.68	3.13	2.39	6.70	7.89	8.17	6.60	4.69	2.91	1.88	1.41
6.0	10.31	12.15	12.14	9.35	6.29	3.96	2.60	2.11	5.72	6.68	7.03	5.68	3.89	2.46	1.69	1.07
7.0	8.97	10.55	10.33	7.77	5.43	3.41	2.28	1.89	5.22	5.84	5.96	4.83	3.37	2.13	1.46	1.00
8.0	8.05	9.41	9.30	7.25	4.81	3.01	2.01	1.76	4.40	5.26	5.40	4.34	3.02	1.94	1.21	1.00

Table 5. Collaborative team ATS performance for GEWMA$_t$ with $n_{\Omega_t} = 6$ pt. 2.

	Communication outbreaks in team of size 6 from a network of size 100														
λ	0.2					0.4					0.7				
k	0.4	0.45	0.5	0.6	0.7	0.4	0.45	0.5	0.6	0.7	0.4	0.45	0.5	0.6	0.7
δ	ATS														
0.5	43.36	42.98	41.90	**39.75**	50.45	43.36	24.01	21.39	**20.42**	21.61	17.41	14.30	12.98	**12.23**	13.64
1.0	16.08	12.91	11.87	**11.62**	12.64	9.28	7.78	7.24	**6.93**	7.74	6.59	5.50	5.10	**5.02**	5.56
2.0	6.16	5.29	4.91	**4.70**	5.40	4.17	3.55	3.38	**3.28**	3.65	3.09	2.71	2.51	**2.49**	2.71
3.0	3.95	3.37	3.18	**3.11**	3.52	2.74	2.42	2.27	**2.23**	2.47	2.12	1.92	1.84	**1.83**	1.96
4.0	2.94	2.63	2.46	**2.43**	2.66	2.16	1.95	**1.81**	1.82	1.95	1.80	1.55	**1.39**	**1.39**	1.58
5.0	2.39	2.13	2.06	**2.03**	2.18	1.87	1.72	1.57	**1.52**	1.68	1.41	1.10	**1.06**	**1.06**	1.19
6.0	2.11	1.89	1.84	**1.75**	1.86	1.66	1.32	**1.19**	1.21	1.40	1.07	1.02	**1.00**	**1.00**	1.03
7.0	1.89	1.68	1.60	**1.56**	1.72	1.31	1.08	**1.05**	**1.05**	1.14	1.00	1.00	1.00	1.00	1.00
8.0	1.76	1.46	1.36	**1.35**	1.54	1.11	1.02	**1.00**	**1.00**	1.04	1.00	1.00	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

Table 6. Collaborative team ATS performance for GEWMA$_t$ with $n_{\Omega_t} = 7$

	Communication outbreaks in team of size 7 from a network of size 100														
λ	0.2					0.4					0.7				
λ	TEWMA	GEWMA				TEWMA	GEWMA				TEWMA	GEWMA			
k		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7
δ	ATS														
0.25											57.26	45.36	40.46	**38.98**	43.89
0.5	53.45	41.97	34.38	**32.10**	42.06	43.36	30.71	17.07	**16.95**	19.56	34.25	13.34	10.70	**10.69**	12.26
1.0	31.84	12.16	**9.83**	9.96	11.58	22.00	9.73	**6.12**	6.16	6.90	16.08	5.44	**4.48**	4.52	5.02
2.0	14.93	4.986	**4.26**	4.26	4.80	9.65	4.16	2.98	**2.96**	3.30	6.68	2.65	**2.25**	2.30	2.54
3.0	8.99	3.29	**2.86**	2.91	3.16	5.89	2.81	2.08	**2.06**	2.26	4.23	1.92	**1.72**	1.74	1.83
4.0	6.25	2.57	**2.18**	2.22	2.47	4.19	2.11	1.72	**1.72**	1.84	3.08	1.53	**1.18**	1.22	1.33
5.0	4.83	2.06	**1.68**	1.78	2.01	3.31	1.88	1.36	**1.34**	1.59	2.49	1.07	**1.01**	1.02	1.09
6.0	3.95	1.83	**1.37**	1.39	1.79	2.75	1.63	**1.06**	1.08	1.24	2.12	1.00	1.00	1.00	1.01
7.0	3.36	1.68	**1.15**	1.20	1.59	2.39	1.34	**1.00**	1.01	1.06	1.89	1.00	1.00	1.00	1.00
8.0	2.94	1.45	**1.02**	1.05	1.10	2.11	1.10	1.00	1.00	1.00	1.39	1.00	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

Table 7. Collaborative team ATS performance for GEWMA$_t$ with $n_{\Omega_t} = 8$.

	Communication outbreaks in team of size 8 from a network of size 100														
λ	0.2					0.4					0.7				
	TEWMA	GEWMA				TEWMA	GEWMA				TEWMA	GEWMA			
k		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7
δ	ATS														
0.25											48.55	**30.51**	31.99	33.17	38.11
0.5	44.73	30.51	27.60	**27.06**	32.37	34.98	16.79	**14.07**	14.50	17.12	25.38	9.89	**9.26**	9.54	10.91
1.0	24.32	9.89	**8.60**	8.70	10.34	16.21	6.38	**5.60**	5.71	6.42	11.20	4.36	**4.05**	4.13	4.60
2.0	10.44	4.36	**3.87**	3.95	4.49	6.90	3.30	**2.71**	2.73	3.11	4.89	2.92	**2.09**	2.14	2.41
3.0	6.36	2.92	2.62	**2.61**	2.97	4.33	2.13	**1.93**	1.94	2.15	3.19	2.20	**1.55**	1.62	1.77
4.0	4.58	2.20	**2.04**	2.10	2.33	3.19	1.78	**1.58**	1.59	1.73	2.38	1.88	**1.06**	1.11	1.33
5.0	3.62	1.88	**1.77**	1.79	1.91	2.55	1.37	**1.14**	1.18	1.41	2.01	1.69	**1.00**	1.00	1.04
6.0	3.02	1.69	**1.49**	1.50	1.69	2.15	1.09	**1.03**	**1.03**	1.14	1.68	1.43	1.00	1.00	1.00
7.0	2.57	1.43	**1.20**	1.25	1.47	1.88	1.01	1.00	1.00	1.02	1.52	1.19	1.00	1.00	1.00
8.0	2.29	1.19	**1.06**	1.09	1.24	1.71	1.00	1.00	1.00	1.00	1.34	1.02	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

Table 8. Collaborative team ATS performance for GEWMA$_t$ with $n_{\Omega_t} = 9$.

	Communication outbreaks in team of size 9 from a network of size 100														
λ	0.2					0.4					0.7				
	TEWMA	GEWMA				TEWMA	GEWMA				TEWMA	GEWMA			
k		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7
δ	ATS														
0.25											38.20	28.93	25.89	26.56	33.81
0.5	38.23	24.63	**22.25**	23.50	30.47	27.38	12.45	**12.34**	13.17	15.19	20.06	9.07	**8.54**	8.62	9.97
1.0	18.64	8.29	**7.58**	8.00	9.50	12.01	4.93	**4.86**	5.14	6.01	8.54	3.94	**3.71**	3.86	4.07
2.0	7.91	3.83	**3.53**	3.65	4.15	5.27	2.48	**2.47**	2.58	2.91	3.87	2.06	**1.90**	2.02	2.14
3.0	4.89	2.56	**2.45**	2.46	2.82	3.39	1.82	**1.80**	1.87	2.01	2.52	1.54	**1.38**	1.48	1.77
4.0	3.56	2.04	**1.89**	1.95	2.17	2.52	**1.40**	**1.40**	1.47	1.70	2.00	1.05	**1.01**	1.05	1.33
5.0	2.84	1.80	**1.63**	1.66	1.81	2.05	**1.05**	**1.05**	1.09	1.33	1.64	1.00	1.00	1.00	1.04
6.0	2.39	1.51	**1.30**	1.40	1.56	1.76	**1.00**	**1.00**	1.01	1.08	1.44	1.00	1.00	1.00	1.00
7.0	2.08	1.21	**1.07**	1.14	1.38	1.58	1.00	1.00	1.00	1.00	1.27	1.00	1.00	1.00	1.00
8.0	1.86	1.07	**1.02**	1.03	1.14	1.43	1.00	1.00	1.00	1.00	1.09	1.00	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

Table 9. Collaborative team ATS performance for GEWMA$_t$ with $n_{\Omega_t} = 10$.

	Communication outbreaks in team of size 10 from a network of size 100														
λ	0.2					0.4					0.7				
	TEWMA	GEWMA				TEWMA	GEWMA				TEWMA	GEWMA			
k		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7		0.4	0.5	0.6	0.7
δ	ATS														
0.25						52.79	43.36	**37.06**	43.40	50.28	35.88	23.48	**22.57**	23.55	28.32
0.50	30.94	21.23	**19.82**	21.63	25.36	21.69	12.45	**10.84**	11.64	15.91	10.80	8.30	**7.27**	7.89	8.89
1.00	14.38	7.25	**7.02**	7.33	8.60	9.42	4.93	**4.67**	4.80	6.63	4.62	3.64	**3.43**	3.66	4.07
2.00	6.12	3.39	**3.26**	3.47	3.87	4.12	2.47	**2.35**	2.45	2.81	3.03	1.94	**1.90**	1.96	2.14
3.00	3.89	2.37	**2.20**	2.39	2.63	2.70	1.82	**1.76**	1.81	1.93	2.09	1.35	**1.23**	1.38	1.63
4.00	2.88	1.90	**1.83**	1.86	1.99	2.09	1.40	**1.27**	1.38	1.65	1.26	**1.00**	**1.00**	1.02	1.14
5.00	2.32	1.64	**1.51**	1.62	1.77	1.73	1.04	**1.02**	1.04	1.23	1.41	1.01	1.00	1.00	1.00
6.00	1.99	1.30	**1.18**	1.28	1.49	1.51	1.00	1.00	1.00	1.04	1.22	1.00	1.00	1.00	1.00
7.00	1.75	1.06	**1.03**	1.09	1.25	1.35	1.00	1.00	1.00	1.00	1.08	1.00	1.00	1.00	1.00
8.00	1.57	1.04	**1.00**	1.02	1.09	1.22	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

5.1.4. Is the methodology fit for purpose?

In order to judge whether the technology is fit for purpose, we consider the monitoring of a crime. To be effective, we would like our strategy to flag the planning of a crime within 7 days. We assume the following specifications of team behavior:

1. In order to plan a crime, team members should call each other at least 0.5 per day during the planning phase. We consider this to be the lowest level of communication necessary to plan a crime.

Table 10. Collaborative team ATS performance for DEWMA$_{\nu,t}$ with $n_{\Omega_t} = 6$.

	Communication outbreaks in team of size 6 from a network of size 100								
λ	0.2			0.4			0.7		
k	0.4	0.45	0.5	0.4	0.45	0.5	0.4	0.45	0.5
δ	ATS								
0.5				25.14	22.27		14.94	13.50	14.57
1.0	13.80	**12.90**	13.56	8.02	**7.78**	8.0	5.78	**5.25**	5.69
2.0	5.51	**5.34**	5.50	3.60	**3.57**	3.74	2.80	**2.66**	2.78
3.0	360	**3.37**	3.55	2.49	**2.41**	2.51	1.99	**1.92**	1.97
4.0	2.75	**2.55**	2.63	1.99	**1.91**	1.98	1.65	**1.60**	1.61
5.0	2.23	**2.12**	2.20	1.71	**1.61**	1.69	1.15	**1.14**	1.20
6.0	1.98	**1.91**	1.94	1.42	**1.38**	1.40	**1.01**	**1.01**	1.03
7.0	1.75	**1.71**	1.75	1.19	**1.18**	**1.18**	1.00	1.00	1.00
8.0	1.58	**1.51**	1.54	1.04	**1.03**	**1.03**	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

2. The planning stage of the crime would result in at least a doubling of their usual communication intensity during this planning stage.
3. The usefulness specification is that detection should be well within 7 days of the start (i.e., the out-of-control ATS <7).

The last specification allows law enforcement agencies enough time for appropriate detective work to be carried out and potentially avoid catastrophic events such as terrorism. The optimal plan for $\lambda = 0.4$ and 0.7 pass the usefulness test by flagging within 7 days on average for all groups (e.g., with $\lambda = 0.4$, k = 0.6, the GEWMA$_t$ statistics detect the outbreak on average in 6.93 days). On the other hand, when the overall communication in the network is relatively sparse ($\lambda = 0.2$), this fit-for-purpose test is only met for collaborative teams having 8 or more members.

5.2. Dominant leader team outbreaks

We now investigate the performance of the GEWMA and DEWMA plans when the outbreak occurs among a fixed but unknown dominant leader team in a homogeneous dynamic network. We simulate the networks with the same specifications as the collaborative team study in Section 5.1, except now the outbreak only occurs on a fixed subset of communications in the team (rather than throughout the entire team as in the collaborative team scenario). In particular, we consider four different dominant leader teams where a communication outbreak occurs on the directed edges shown in Figure 1. In each of these four teams, team member 6 is assumed to be the dominant leader and communicates with all other members of the team.

We assess the performance of the GEWMA and DEWMA plans on these dominant leader outbreaks and report the results in Tables 14 and 15. Our results suggest that again the choice of k plays an important role in establishing the best performing monitoring strategy. Furthermore, across all values of λ k, and n_Ω, we found that the DEWMA method outperformed the GEWMA strategy in this simulation study. Both methods witness improved performance as the signal-to-noise ratio (δ) increases. Our results provide empirical evidence that the DEWMA plan is an effective strategy when the target team has a dominant leader or when the team is more sparsely connected than a collaborative team.

5.3. Heterogeneous networks with no outbreak

We now assess the performance of the ADEWMA plan from [26] on heterogeneous networks that undergo no outbreak, but whose size changes through time. Without loss of generality, we fix the mean communication count between node i and j at time t as $\lambda_{i,j,t} = a|i-j| + 0.90$, for a fixed constant $a < 0$. This specification gives a higher likelihood of communication between nodes that are close to one another in the ordering of the nodes. To vary the size of the network through time, we fix lower (m_L) and upper bounds (m_H) and select the size of the t-th network n_t by randomly drawing a discrete value uniformly from the interval $[m_L, m_H]$.

As there is no outbreak in our simulated collection of networks, we seek a plan that identifies no change for some fixed number of time steps. By investigating this aspect of the ADEWMA plan, we can better understand how to control the number of false discoveries under a null model where no outbreak is present. For our current study, we seek an ADEWMA plan that delivers an ATS of 100. We note that one could alternatively seek an ATS of 370 to match the standard 3-sigma strategy of Shewhart control charts, but the choice is arbitrary. We vary the values of a, m_L, and m_H and identify the threshold adjustment that acquires the desired ATS over 1000 simulations.

Table 11. Collaborative team ATS performance for $DEWMA_{\nu,t}$ with $n_{\Omega_t} = 7$.

	Communication outbreaks in team of size 7 from a network of size 100								
λ	0.2			0.4			0.7		
k	0.4	0.45	0.5	0.4	0.45	0.5	0.4	0.45	0.5
δ	ATS								
0.5			38.72	18.44	**16.89**	19.57	11.97	**11.36**	12.56
1.0	10.56	**10.52**	11.18	6.72	**6.66**	7.42	**4.78**	4.82	5.57
2.0	4.67	**4.61**	4.90	3.20	**3.19**	3.44	**2.44**	2.46	2.76
3.0	3.12	**3.09**	3.21	2.26	**2.18**	2.36	**1.84**	1.86	1.96
4.0	2.43	**2.35**	2.47	**1.85**	**1.85**	1.91	1.33	**1.29**	1.60
5.0	1.96	**1.93**	1.99	**1.45**	**1.45**	1.60	**1.04**	**1.04**	1.20
6.0	1.83	**1.73**	1.81	1.17	**1.14**	1.33	1.01	**1.00**	1.04
7.0	1.59	**1.54**	1.62	1.04	**1.03**	1.13	1.00	1.00	1.00
8.0	1.36	**1.32**	1.40	**1.00**	**1.00**	1.04	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

Table 12. Collaborative team ATS performance for $DEWMA_{\nu,t}$ with $n_{\Omega_t} = 8$.

	Communication outbreaks in team of size 7 from a network of size 100								
λ	0.2			0.4			0.7		
k	0.4	0.45	0.5	0.4	0.45	0.5	0.4	0.45	0.5
δ	ATS								
0.25							32.18	35.76	37.24
0.5	32.78	32.88	33.92	15.66	**15.29**	16.32	**10.07**	10.16	10.34
1.0	**9.54**	9.64	10.26	**6.09**	6.20	6.32	4.36	**4.24**	4.62
2.0	**4.14**	4.17	4.32	**2.84**	2.90	3.02	2.42	**2.31**	2.35
3.0	2.78	**2.74**	2.94	**2.01**	2.03	2.14	1.70	**1.68**	1.76
4.0	2.12	**2.10**	2.35	1.66	**1.65**	1.78	1.21	**1.19**	1.28
5.0	**1.86**	**1.86**	1.90	1.32	**1.28**	1.40	**1.01**	**1.01**	1.02
6.0	1.68	**1.63**	1.71	1.05	**1.04**	1.09	1.00	1.00	1.00
7.0	1.37	**1.36**	1.50	**1.00**	**1.00**	1.01	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

Table 13. Collaborative team ATS performance for $DEWMA_{\nu,t}$ with $n_{\Omega_t} = 9$.

	Communication outbreaks in team of size 9 from a network of size 100								
λ	0.2			0.4			0.7		
k	0.4	0.45	0.5	0.4	0.45	0.5	0.4	0.45	0.5
δ	ATS								
0.25				54.90	**48.93**	52.01	28.93	**26.74**	29.78
0.5	22.27	21.86	25.7	14.26	**13.10**	14.06	**8.87**	8.94	9.33
1.0	8.47	**8.14**	9.05	5.51	**5.38**	5.76	**3.91**	4.04	4.16
2.0	3.81	**3.80**	3.96	2.73	**2.60**	2.84	**2.14**	2.15	2.19
3.0	**2.59**	**2.59**	2.84	1.94	**1.91**	2.01	**1.52**	1.58	1.70
4.0	**2.04**	**2.04**	2.13	**1.48**	**1.48**	1.68	**1.04**	1.10	1.13
5.0	1.80	**1.79**	1.86	**1.19**	**1.19**	1.26	1.00	1.00	1.00
6.0	1.49	**1.48**	1.62	**1.03**	**1.03**	1.02	1.00	1.00	1.00
7.0	1.25	**1.19**	1.32	1.00	1.00	1.00	1.00	1.00	1.00
8.0	1.09	**1.06**	1.10	1.00	1.100	1.00	1.00	1.00	1.00

Note: Bolded values indicate the best ATS performance across simulations

The threshold adjustments and calculated ATS are provided in Table 16.

The simulation results in Table 16 reveal that the ADEWMA plan with threshold 0.984 has an in-control ATS closest to the desired value of 100 when $m_H > 135$. On the other hand, when $m_H \leq 135$ selecting a threshold of 1 delivers the best plan. These results suggest that the ADEWMA plan is robust to large changes in the size of the network from one time to the next. In many applications (like our application in Section 6), the size n_t is likely to have a small variation over time. We find that, in these situations, the ADEWMA plan witnesses an improvement in overall robustness.

5.4. Comparison of the TEWMA, EWMA, and CUSUM

We now assess the utility of the TEWMA monitoring plan in [15] by comparing it with the EWMA plan for

Table 14. Dominant leader team outbreaks involving teams of size 6 to 9.

	DEWMA				GEWMA				DEWMA				GEWMA			
λ	0.2								0.4							
n_Ω	6	7	8	9	6	7	8	9	6	7	8	9	6	7	8	9
k	0.45				0.6				0.45				0.6			
δ	ATS															
0.25											62.66	54.82			98.60	83.27
0.50	49.48	39.94	34.19	33.78	72.00	63.46	44.72	39.01	27.42	22.54	20.24	15.94	32.74	24.59	21.99	17.97
1.00	15.99	13.19	11.29	10.51	16.91	14.32	13.24	11.44	9.99	8.49	7.27	6.59	9.03	8.42	7.42	6.73
2.00	6.26	5.65	5.09	4.62	6.27	5.68	5.43	4.65	4.29	3.86	3.46	3.16	4.18	3.73	3.67	3.16
3.00	4.25	3.77	3.31	2.39	4.17	3.68	3.48	3.00	2.90	2.56	2.42	2.19	2.83	2.59	2.42	2.21
4.00	3.20	2.83	2.62	1.73	3.18	2.73	2.66	2.38	2.28	2.09	1.90	1.70	2.17	2.06	1.92	1.79
5.00	2.66	2.28	2.11	1.32	2.55	2.39	2.22	1.95	1.94	1.70	1.56	1.36	1.84	1.71	1.59	1.42
6.00	2.16	2.00	1.64	1.06	2.19	2.06	1.86	1.74	1.65	1.48	1.32	1.16	1.61	1.48	1.30	1.18
7.00	1.89	1.78	1.43	1.00	2.01	1.72	1.67	1.52	1.43	1.29	1.10	1.04	1.34	1.29	1.13	1.07
8.00	1.64	1.36	1.24	1.00	1.70	1.61	1.48	1.35	1.22	1.04	1.01	1.00	1.17	1.06	1.00	1.00

Table 15. Dominant leader team outbreaks involving teams of size 6 to 9.

	DEWMA				GEWMA			
λ	0.7							
n_Ω	6	7	8	9	6	7	8	9
k	0.45				0.6			
δ	ATS							
0.25	53.61	46.84	40.94	33.88	98.12	80.02	67.48	44.12
0.50	16.95	13.82	12.72	10.51	18.12	15.48	14.72	10.94
1.00	6.77	6.01	5.25	4.62	6.38	5.88	5.47	4.81
2.00	3.28	2.90	2.66	2.38	3.04	2.84	2.74	2.46
3.00	2.28	2.03	1.93	1.73	2.17	1.98	2.00	1.83
4.00	1.82	1.70	1.51	1.32	1.76	1.66	1.56	1.39
5.00	1.53	1.21	1.16	1.06	1.40	1.31	1.19	1.09
6.00	1.23	1.09	1.02	1.00	1.13	1.01	1.01	1.00
7.00	1.02	1.00	1.00	1.00	1.02	1.00	1.00	1.00

Table 16. The ADEWMA plans for heterogeneous networks with no outbreak.

m_L	m_H	Threshold adjustment	ATS	a
100	135	1.005	102.9	−0.0030
115	135	1.0037	103.1	−0.0030
110	150	0.982	104.0	−0.0060
130	150	0.984	102.2	−0.0060
100	175	0.984	100.9	−0.0050
115	175	0.984	102.2	−0.0050
135	175	0.982	103.3	−0.0050
155	175	0.982	101.6	−0.0050
135	250	0.985	99.4	−0.0035
200	250	0.983	102.4	−0.0035
150	275	0.985	100.9	−0.0030
215	275	0.985	103.6	−0.0030
305	315	0.981	101.6	−0.0027
250	350	0.986	99.3	−0.0025

Poisson counts in [2] as well as the CUSUM plan for Poisson counts described in Sparks et al. (2010). We set $\alpha = 0.075$ for the TEWMA plan. All plans below were trained to have approximately in-control ATS = 100. We report the results in Table 1. We note that the same threshold was used for all TEWMA statistics, but different thresholds were needed for the EWMA and the CUSUM statistic when the in-control mean changed. Thus, the main advantage of the TEWMA statistic is that the threshold is invariant of the mean counts. We see from Table 8, however, that its detection properties are not as good as the EWMA or CUSUM plans. Its primary advantage, therefore, is on its scalability for large networks.

6. Application to U.S. Congressional voting

We now apply the GEWMA monitoring plan from [6] to investigate the dynamic relationship between Republican and Democratic senators in the U.S. Congress. We analyze the voting habits of each U.S. senator according to his or her vote (yay, nay, or abstain) on each bill that went to Congress. We investigate these voting habits from 1857 (Congress 35) to 2015 (Congress 113).

We generated a dynamic network to model the co-voting patterns among U.S. Senators in the following manner. We first collected the raw roll call voting data for each bill from http://voteview.com. For each Congress, we generate a new network, where the senators of that Congress are the nodes and the edge weight between two senators is the number of bills for which those two senators voted concurrently in that Congress. We restrict our analysis to Republican and Democrat senators only (thus ignoring the Independent party and other affiliations).

Predictable behavior is regarded as in-control. To model in-control behavior, we use a logistic regression model to predict whether two senators will vote the same on a newly submitted bill. We fit a logistic model to estimate the probability that a senator (Senator A) would vote the same as another senator (Senator B) using the following predictors: (a) the political affiliation of each senator (Senators A and B), (b) which party had a majority in the Congress,

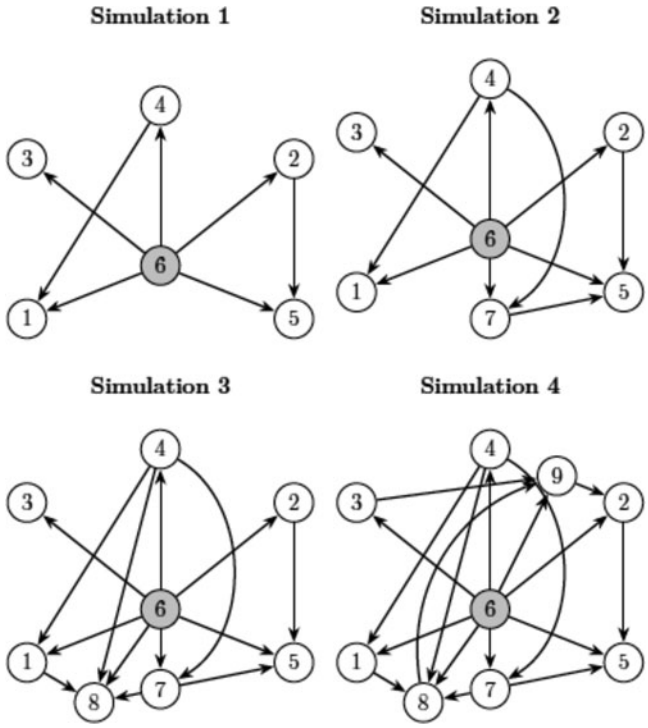

Figure 1. Dominant leader target teams for the simulation study. Teams are of size 6, 7, 8, and 9 among a network of size 100. For each simulation, a communication outbreak occurs only on the directed edges shown. In each simulation, node 6 is the dominant leader and communicates with every member of the team.

(c) the proportion of that majority, and (d) the proportion of representation of Senator A's political affiliation. The expected number of votes from Senator A to Senator B was calculated by multiplying the predicted probability from the logistic regression by the total number of votes for that senator. This count was assumed to be Poisson distributed with in-control mean given by this expected count.

In this application we are interested in both unusually high counts and unusually low counts. Therefore we run two one-sided charts. In particular, for a target team Ω_t we analyze the GEWMA$_t$ statistic from [4], as well as the lower GEWMA (L−GEWMA$_t$) statistic defined by

$$\text{L}-\text{GEWMA}_t = \min\Bigg(\alpha \sum_{i\in\Omega_t}\sum_{j\in\Omega_t}\widetilde{y}_{i,j,t} + (1-\alpha)\ \text{L}-\text{GEWMA}_{t-1}, \mu_{\Omega_t}\Bigg),$$

where α was fixed to be 0.075. The plans are trained using simulation to deliver an in-control false-alarm rate of 200. The GEWMA and L-GEWMA curves were calculated from two sources, (i) the likelihood of Republicans voting with Democrats and (ii) the likelihood of Democrats voting with Republicans. We do not expect our co-voting patterns to remain in-control and predictable; thus, we are particularly interested in identifying sustained periods of unusual behavior.

The GEWMA and L-GEWMA curves are plotted in Figure 4. These plots reveal several interesting trends in the Congressional co-voting network. First, the tendency for Republican and Democratic senators to vote with one another has been significantly low beginning from Congress 103. This finding supports the political polarization theory observed in Moody and Mucha (2013), who noted that the Republican and Democrat schism began around the time of Bill Clinton's first term as president (Congress 103). Second, there was a sustained coherence of voting between opposing political parties between Congress 85 (1957) and Congress 100 (1987). During this time, the likelihood of one party concurrently voting with the other opposing party was significantly high. Much of this time period coincides with the so-called"Rockefeller Republican"era (1960 - 1980) in which Republican party members were known to hold particularly moderate views like the former governor of New York, Nelson Rockefeller (Rae 1989; Smith 2014). This finding was also identified using network surveillance techniques in Wilson et al. (2016).

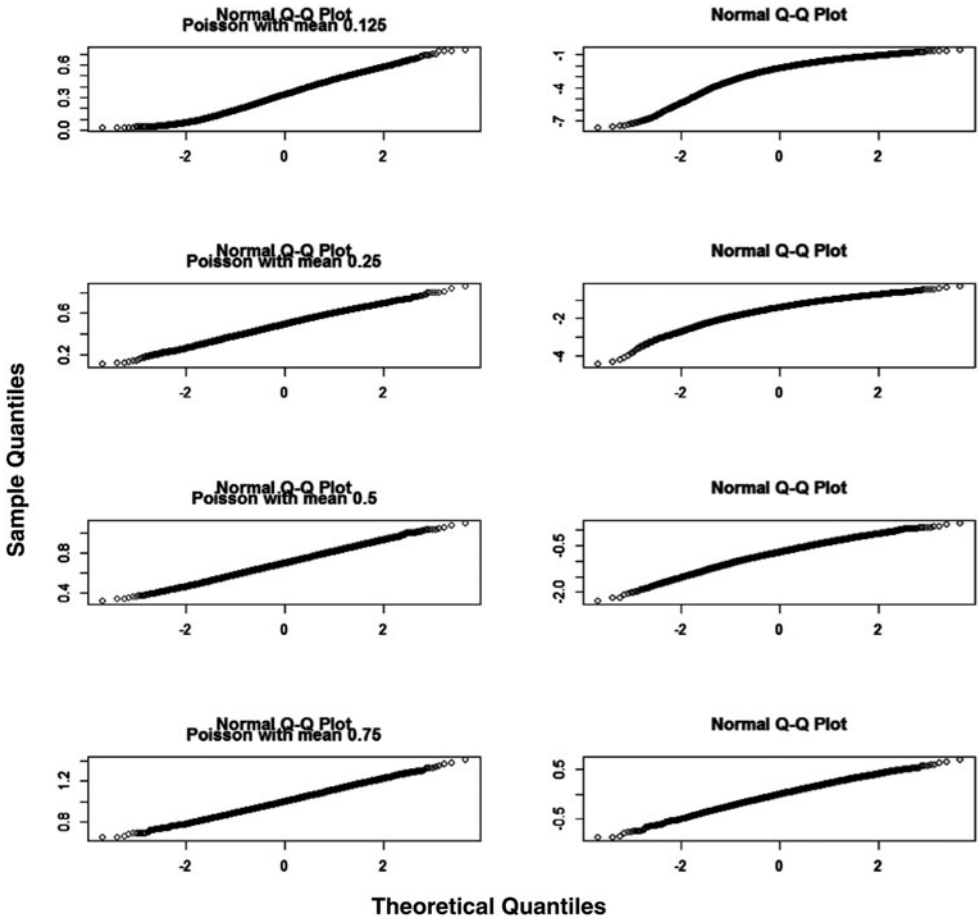

Figure 2. QQ plots for smoothed EWMA counts generated from Poisson data with small mean, indicated in each plot. The square-root transformation is on the left and the corresponding log transformation on the right. This figure suggests that the square-root transformation provides a better Normal approximation than the log transformation when the mean number of counts is small.

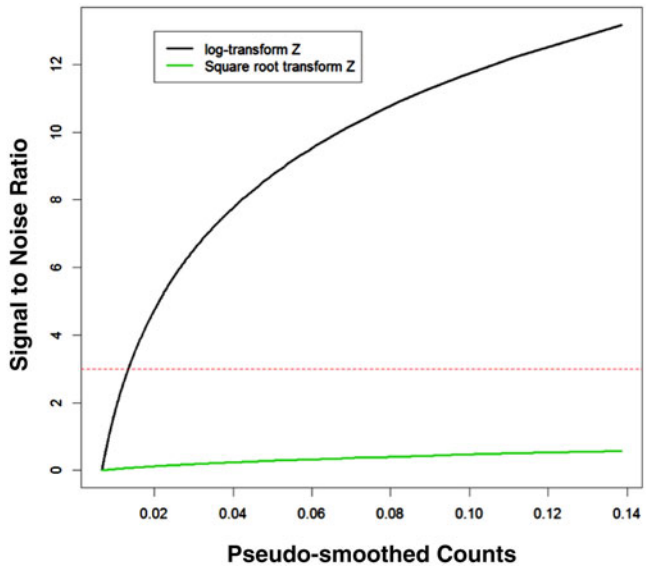

Figure 3. Signal-to-noise statistics for the log- and square-root transformed smoothed EWMA counts when no structural change has been introduced and data are generated from a log-Normal distribution. These results reveal that the square-root transformations are well suited to networks with sparse communication.

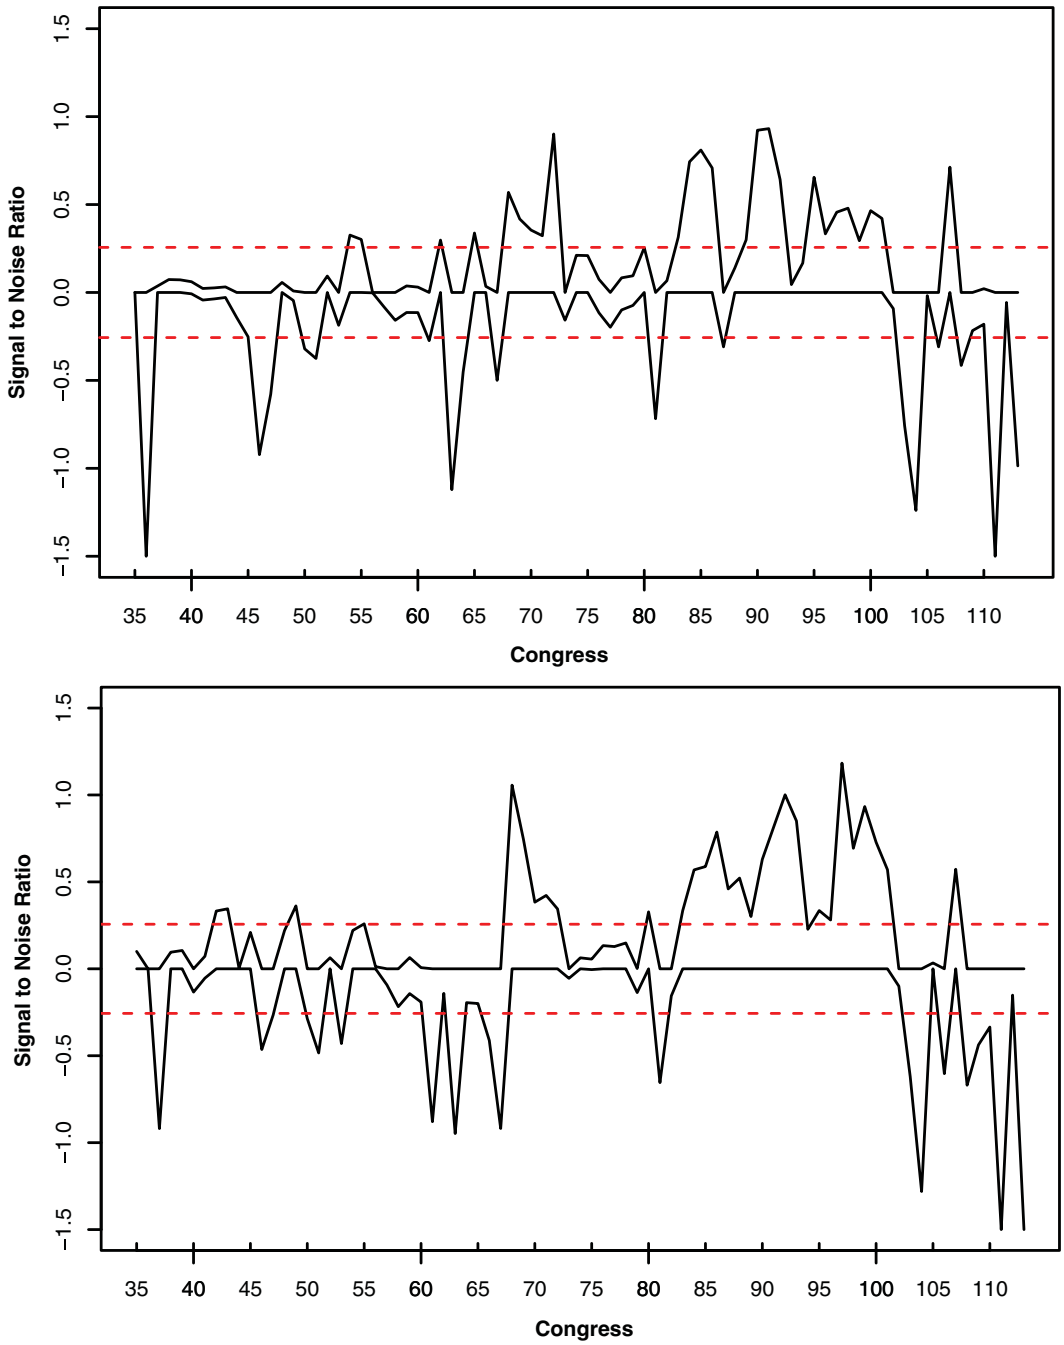

Figure 4. GEWMA and L-GEWMA control charts for monitoring (left): the likelihood of Democratic senators to vote with Republican senators, and (right): the likelihood of Republican senators to vote with Democratic senators. Red dotted lines mark the control limits of the GEWMA signal-to-noise value for each Congress. In each plot, the upper curve represents the GEWMA statistic and the lower curve represents the L-GEWMA statistic over time.

7. The square-root transformation

For each of the EWMA plans introduced in Sections 3, 4, and 5, we propose monitoring the square root of the defined EWMA statistic rather than the original statistic itself. There are three important reasons for considering the square-root transformation, which we describe here. The motivation of the square-root transform becomes particularly evident for sparse Poisson-weighted networks. Indeed, in this scenario, we find that (i) the square-root transform provides a better suited Normal approximation than the log-transformed counts, (ii) the square-root transform has fewer false discoveries, and (iii) the square-root transformed counts are robust to changes in network size and mean counts. We discuss these three reasons and provide empirical support for

these claims. Finally, we investigate by simulation the transformation of the EWMA statistic as well as the original count statistic, motivating our proposed statistics.

7.1. On the normal approximation for Poisson counts

We first consider the Normal approximation of the square-root transformed EWMA statistics arising from Poisson data with low counts. This scenario is relevant for our application to Congressional voting, and is representative of sparse networks that are often observed in practice. We generate Poisson stochastic processes with means 0.125, 0.25, 0.5, and 0.75, and calculate the EWMA statistic under the square root. We compare the normality of these statistics with the natural log-transformed EWMA data for comparison. Quantile-quantile (QQ) plots for each of the transformations are presented in Figure 2.

This figure demonstrates that, with EWMA smoothing with weights, the square-root transformation leads to a better normal approximation than the log transformation for Poisson data with low average counts. Notably in Figure 2, the upper tail for the square root transformation is always very close to a Normal distribution.

7.2. On false discoveries

Another primary reason for utilizing the square-root transformation is that it avoids flagging unusual outbreaks in sparse communication regions of the observed communication matrix. That is, the square-root transformed version of the EWMA efficiently reduces false discovery over EWMA data with no transformation or under the log transformation. We illustrate this with an example.

Assume that the observed counts are approximately log-normal distributed with mean equal to -5 and variance 1, thus favoring the log transformation. In sparse networks like these, it is desirable to signal a change only when the counts are reasonably large, e.g., large enough to efficiently plan a crime. We apply EWMA smoothing to these counts as is carried out in the article using the smoother for log counts $\log y_t$:

$$\mathrm{EWMA}_t = 0.1 \log y_t + 0.9\mathrm{EWMA}_{t-1}$$

The signal to noise ratio for the log counts is reported in Figure 3 for both the square-root and log-transformed smoothed counts. For this process, the signal-to-noise ratio is given by

$$Z_t = \sqrt{19}(\mathrm{EWMA}_t + 5)$$

Figure 2 plots $\exp(\mathrm{EWMA}_t)$ against Z_t. This figure represents"pseudo smoothed counts" versus the signal-to-noise ratio Z_t. Using a 3-sigma Shewhart chart, then the signal-to-noise ratio for the log transform flags an unusual change for very low counts and these counts are far too low to plan a crime, whereas the square-root transform is nowhere near flagging these as unusual. The square-root transform does what we want it to do in these circumstances. Similarly, in the voting application, we consider in Section 6, we are interested in changes in which there are enough votes before flagging a signal.

7.3. Robust ATS and thresholds

Our third reason for considering the square-root transform is due to the fact that, for the known subgroup outbreaks, including the outbreaks involving the whole network, our EWMA plans deliver the identical plan with the same threshold and in-control ATS. That is, the ATS of the square root plans do not depend on:

- the in-control mean rates of the individual counts over time.
- sample changes to the size of the network. For example, the same threshold applied independent of the make up of the Senate or the number of seats the Republicans or Democrats held for in-control ATS of 100.

We do not prove these results theoretically, but we demonstrated this through simulation within the bounds considered in the article. For our political application, we calculate the ATS for the thresholds 0.513 and 0.399 for the EWMA plan to investigate whether or not the thresholds are a function of the number of Democrats and Republicans in the network or the size of the total network. These values are reported in Table 2. This table demonstrates that the test statistic in our application is invariant of changes in the number of Democratic and Republican Senators in the senate at any time point. This property together with the statistic threshold being independent of the mean counts makes this statistic extremely useful in simplifying the process of applying these charts.

7.4. On transforming the EWMA statistic vs. transforming the original counts

The simulation below is of a 50 by 50 communication network of individuals, which is assumed to

have an in-control homogeneous communication mean across the members of the network. We compare the monitoring results of the square root of the sum of counts (our current approach) against the sum of the square root of the individual counts. Our simulations (provided in Table 3 reveal two important outcomes:

1. Transforming the counts before aggregating does not have a homogeneous threshold value and therefore is difficult to implement.
2. The approach of transforming the counts before aggregating generally results in later detections of outbreaks.

Therefore, we conclude that the approach taken in the article - transforming the EWMA statistics directly - is the preferred method due to the additional computational effort needed to identify possibly heterogeneous threshold values for the plan.

8. Discussion

This article introduces novel and computationally feasible surveillance plans for identifying communication outbreaks in dynamic networks. In the worst-case scenario when the target team is unknown, the proposed method monitors at most n^2 candidate teams, which dramatically improves the computational memory needed for an exhaustive search. Our new plan uses a general multivariate EWMA approach to accumulate temporal memory of communication counts. The approach can easily be extended to situations with more than one communication channel. Plans were extended to handle networks with heterogeneous mean counts (as in the application) and the value of our proposed plans was further demonstrated with simulated applications.

In our simulation study, we found that our new approach is able to effectively identify outbreaks even when the outbreak covers a small number of communications (<1 percent of total communications). These results suggest that the technology will be particularly useful in crime management, as crime is typically committed by gangs of a small size (Morgan and Shelley, 2014). Furthermore, we believe that law enforcement agencies would value our proposed technique as it could be used to help gain insights on persons of interest (e.g., it could be applied to juvenile crime rings as a preventative tool to help reduce repeat offenders).

We found that, when the outbreak is global across all communications of the targeted people, using the TEWMA_t plan is the best approach and this plan is invariant of the distribution of communication counts in the target network. If the communication outbreak involves a small subgroup of the targeted people, then the group-EWMA (GEWMA_t) plan has the best performance. As the size of the outbreak group is seldom known in advance, applying these plans simultaneously in a single plan may offer a more robust means to detect the full range of potential outbreaks.

Our proposed technique motivates several areas of future research. For example, future work should explore the potential of extending this approach to cover geographic dimensions (see Carley et al. (2013)) to account for the spatial nature of observed dynamic systems. Furthermore, one can explore other ways of estimating the target team for monitoring. New approaches could involve defining people in the targeted network with either increased connectivity or historically a high connectivity. The target group itself could be regarded as varying according to whether they achieve a certain level of connectivity with the leaders or average connectivity within the target group. In principle, one could also estimate teams of individuals that are most densely connected at time t using a community detection or extraction algorithm on the network Y_t (Lancichinetti et al. 2010; Wilson 2017; Wilson et al. 2014; Zhao et al. 2011). Alternatively, one could identify candidate teams in a network with statistically significant edges using a p-value technique like that developed in Wilson et al. (2013).

This article arbitrarily selected the temporal smoothing parameter $\alpha = 0.075$. Future research effort could be devoted to selecting an appropriate value for the multivariate temporal smoothing. We believe that this effort should be devoted either to establishing an appropriate robust choice for α or to alternatively varying the choice of α for each communication count so as to exploit local trends in the network such as the work done in Capizzi and Masarotto (2003). Finally, the methods presented in this work implicitly assume that the investigator knows what type of outbreak he or she is searching for in the data. In practice, however, it is often the case that the investigator does not know the type of outbreak that will occur in the system being monitored. In such cases, we recommend either applying each of the DEWMA, GEWMA, and TEWMA monitoring plans or using some consensus of these plans.

The latter approach would rely on developing some robust plan, which is akin to the three-CUSUM plan developed in Sparks (2000). In either case, simulations must account for the use of some combination of methods and plan thresholds chosen accordingly. We will investigate the application of robust plans in future work.

ORCID

Ross Sparks http://orcid.org/0000-0001-5852-5334
James D. Wilson http://orcid.org/0000-0002-2354-935X

References

Aiello, W., F. Chung, and L. Lu. 2000. A random graph model for massive graphs. Paper presented at the Thirty-Second Annual ACM Symposium on Theory of Computing, Portland, OR, pp. 171–180.

Akoglu, L., and C. Faloutsos. 2013. Anomaly, event, and fraud detection in large network datasets. Paper presented at the Sixth ACM International Conference on Web Search and Data Mining, Rome, Italy, pp. 773–774.

Azarnoush, B., K. Paynabar, J. Bekki, and G. Runger. 2016. Monitoring temporal homogeneity in attributed network streams. *Journal of Quality Technology* 48 (1):28–43. doi: 10.1080/00224065.2016.11918149

Barabási, A.-L., and R. Albert. 1999. Emergence of scaling in random networks. *Science* 286 (5439):509–12. PMC: 10521342

Bartlett, M. 1936. The square root transformation in analysis of variance. *Supplement to the Journal of the Royal Statistical Society* 3 (1):68–78. doi:10.2307/2983678

Capizzi, G., and G. Masarotto. 2003. An adaptive exponentially weighted moving average control chart. *Technometrics* 45 (3):199–207. doi: 10.1198/004017003000000023

Carley, K. M., J. Pfeffer, H. Liu, F. Morstatter, and R. Goolsby. 2013. Near real time assessment of social media using geo-temporal network analytics. Paper presented at the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), Niagara Falls, Canada, pp. 517–524.

Chau, D. H. S. Pandit, and C. Faloutsos. 2006. Detecting fraudulent personalities in networks of online auctioneers. In *Knowledge discovery in databases: PKDD 2006*, ed. J. Fürnkranz, T. Scheffer, and M. Spiliopoulou, 103–114. Berlin, Germany: Springer.

Erdös, P., and A. Rényi. 1960. On the evolution of random graphs. *Publications of the Mathematical Institute of Hungarian Academy of Sciences* 5:17–61.

Gan, F. 1993. Exponentially weighted moving average control charts with reflecting boundaries. *Journal of Statistical Computation and Simulation* 46 (1–2):45–67. doi: 10.1080/00949659308811492

Goldenberg, A., A. X. Zheng, S. E. Fienberg, and E. M. Airoldi. 2010. A survey of statistical network models. *Foundations and Trends® in Machine Learning* 2 (2): 129–233. doi: 10.1561/2200000005

Heard, N. A., D. J. Weston, K. Platanioti, and D. J. Hand. 2010. Bayesian anomaly detection methods for social networks. *The Annals of Applied Statistics* 4 (2):645–62. doi: 10.1214/10-AOAS329

Kim, Y., and M. O' Kelly. 2008. A bootstrap based space–time surveillance model with an application to crime occurrences. *Journal of Geographical Systems* 10 (2):141–65. doi: 10.1007/s10109-008-0058-4

Krebs, V. E. 2002. Mapping networks of terrorist cells. *Connections* 24 (3):43–52.

Lancichinetti, A., F. Radicchi, and J. Ramasco. 2010. Statistical significance of communities in networks. *Physical Review E, Covering Statistical, Nonlinear, and Soft Matter Physics* 81:046110.doi: 10.1103/PhysRevE.81.046110

Liu, K., Y. Mei, and J. Shi. 2015. An adaptive sampling strategy for online high-dimensional process monitoring. *Technometrics* 57 (3):305–19. doi: 10.1080/00401706.2014.947005

Liu, X., J. Bollen, M. L. Nelson, and H. Van de Sompel. 2005. Co-authorship networks in the digital library research community. *Information Processing & Management* 41 (6):1462–80. doi: 10.1016/j.ipm.2005.03.012

Mei, Y. 2010. Efficient scalable schemes for monitoring a large number of data streams. *Biometrika* 97 (2):419–33. doi:10.1093/biomet/asq010

Mei, Y. 2011. Quickest detection in censoring sensor networks. Paper presented at the 2011 IEEE International Symposium on Information Theory Proceedings (ISIT), Saint Petersburg, Russia, pp. 2148–2152.

Morgan, A., K., and W. W. Shelley. 2014. Juvenile street gangs. *The Encyclopedia of Criminology and Criminal Justice*, 1–8. Wiley Online Library.

Moody, J., and P. J. Mucha. 2013. Portrait of political party polarization. *Network Science* 1 (01):119–21. doi: 10.1017/nws.2012.3

Nakaya, T., and K. Yano. 2010. Visualising crime clusters in a space-time cube: An exploratory data-analysis approach using space-time kernel density estimation and scan statistics. *Transactions in GIS* 14 (3):223–39. doi: 10.1111/j.1467-9671.2010.01194.x

Neill, D. B. 2009. Expectation-based scan statistics for monitoring spatial time series data. *International Journal of Forecasting* 25 (3):498–517. doi: 10.1016/j.ijforecast.2008.12.002

Pandit, S., D. H. Chau, S. Wang, and C. Faloutsos. 2007. Netprobe: A fast and scalable system for fraud detection in online auction networks. Paper presented at the 16th International Conference on World Wide Web, Banff, Canada, pp. 201–210.

Parker, K. S., J. D. Wilson, J. Marschall, P. J. Mucha, and J. P. Henderson. 2015. Network analysis reveals sex-and antibiotic resistance-associated antivirulence targets in clinical uropathogens. *ACS Infectious Diseases* 1 (11): 523–32. doi: 10.1021/acsinfecdis.5b00022

Peel, L., and, and A. Clauset. 2014. Detecting change points in the large-scale structure of evolving networks. *arXiv:* 1403.0989.

Porter, M. D., and G. White. 2012. Self-exciting hurdle models for terrorist activity. *The Annals of Applied Statistics* 6 (1):106–24. doi: 10.1214/11-AOAS513

Prusiewicz, A. 2008. A multi-agent system for computer network security monitoring. In *Agent and multi-agent systems: Technologies and applications*, ed. G. S. Jo and L. Jain, 842–849. Incheon, Korea: Springer.

Rae, N. C. 1989. *The decline and fall of the liberal republicans: From 1952 to the present*. New York, NY: Oxford University Press.

Reid, E., J. Qin, Y. Zhou, G. Lai, M. Sageman, G. Weimann, and H. Chen. 2005. Collecting and analyzing the presence of terrorists on the web: A case study of jihad websites. In *Intelligence and security informatics*, ed. P. B. Kantor, G. Muresan, F. S. Roberts, Daniel D. Zeng, F.-Y. Wang, H. Chen, and R. C. Merkle, 402–411. Atlanta, GA: Springer.

Savage, D., X. Zhang, X. Yu, P. Chou, and Q. Wang. 2014. Anomaly detection in online social networks. *Social Networks* 39:62–70. doi: 10.1016/j.socnet.2014.05.002

Smith, R. N. 2014. *On his own terms: A life of Nelson Rockefeller*. New York, NY: Random House.

Sparks, R., and E. Patrick. 2014. Detection of multiple outbreaks using spatio-temporal EWMA-ordered statistics. *Communications in Statistics-Simulation and Computation* 43 (10):2678–701. doi: 10.1080/03610918.2012.758741

Sparks, R. S. 2000. CUSUM charts for signalling varying location shifts. *Journal of Quality Technology* 32 (2): 157–71. doi: 10.1080/00224065.2000.11979987

Sparks, R. S., T. Keighley, and D. Muscatello. 2009. Improving EWMA plans for detecting unusual increases in Poisson counts. *Advances in Decision Sciences* 2009: 512356. doi: 10.1155/2009/512356

Sparks, R. S., T. Keighley, and D. Muscatello. 2010. Early warning CUSUM plans for surveillance of negative binomial daily disease counts. *Journal of Applied Statistics* 37 (11):1911–29. doi: 10.1080/02664760903186056

Tan, K. M., P. London, K. Mohan, S.-I. Lee, M. Fazel, and D. Witten. 2014. Learning graphical models with hubs. *The Journal of Machine Learning Research* 15 (1):3297–331.

Weiß, C. H. 2007. Controlling processes of Poisson counts. *Quality and Reliability Engineering International* 23 (6): 741–54. doi: 10.1002/qre.875

Weiß, C. H. 2009. EWMA monitoring of correlated processes of Poisson counts. *Quality Technology and Quantitative Management* 6 (2):137–53. doi: 10.1080/16843703.2009.11673190

Wilson, J., S. Bhamidi, and A. Nobel. 2013. Measuring the statistical significance of local connections in directed networks. Paper presented at Neural Information Processing Systems: Workshop on the Frontiers of Network Analysis: Methods, Models and Applications, Lake Tahoe, NV.

Wilson, J. D. 2017. Community extraction in multilayer networks with heterogeneous community structure. *Journal of Machine Learning Research* 18 (149):1–49.

Wilson, J. D., N. T. Stevens, and W. H. Woodall. 2016. Modeling and estimating change in temporal networks via a dynamic degree corrected stochastic block model. *arXiv*:1605.04049.

Wilson, J. D., S. Wang, P. J. Mucha, S. Bhamidi, and A. B. Nobel. 2014. A testing based extraction algorithm for identifying significant communities in networks. *The Annals of Applied Statistics* 8 (3):1853–91. doi: 10.1214/14-AOAS760

Woodall, W. H., M. J. Zhao, K. Paynabar, R. Sparks, and J. D. Wilson. 2017. An overview and perspective on social network monitoring. *IISE Transactions* 49 (3):354–65. doi: 10.1080/0740817X.2016.1213468

Zeng, D., W. Chang, and H. Chen. 2004. A comparative study of spatio-temporal hotspot analysis techniques in security informatics. Paper presented at the 7th International IEEE Conference on Intelligent Transportation Systems, Sophia Antipolis, France, pp. 106–111.

Zhao, Y., E. Levina, and J. Zhu. 2011. Community extraction for social networks. *Proceedings of the National Academy of Sciences of the United States of America* 108 (18):7321–6. doi:10.1073/pnas.1006642108

Zhou, Q., C. Zou, Z. Wang, and W. Jiang. 2012. Likelihood-based EWMA charts for monitoring Poisson count data with time-varying sample sizes. *Journal of the American Statistical Association* 107 (499):1049–62. doi: 10.1080/01621459.2012.682811

Zou, C., Z. Wang, X. Zi, and W. Jiang. 2015. An efficient online monitoring method for high-dimensional data streams. *Technometrics* 57 (3):374–87. doi: 10.1080/00401706.2014.940089

Appendix

Specification of threshold values

Simulation methods were used to estimate the thresholds for the $\text{DEWMA}_{\nu,t}$ and GEWMA_t plans so as to deliver an in-control ATS of approximately 100. The thresholds for both the collaborative team and the dominant leader team were established in the identical manner. To avoid redundancy, we will describe the simulation procedure to determine thresholds in the collaborative team scenario.

For the $\text{DEWMA}_{\nu,t}$ plan, we simulated networks of size $n = 100, 125, 150, \ldots, 375, 400$. For each network, we fixed the temporal memory as $\alpha = 0.10$ and generated homogeneous networks with mean counts equal to $\lambda = 0.01, 0.02, 0.03, \ldots, 0.10, 0.15, 0.20, \ldots, 0.95, 1.0$. For each combination, the thresholds $h_D(\lambda, n)$ are estimated to obtain the fixed ATS.

In practice, if the team is known in advance (and hence no search is needed), the threshold is invariant of the mean rate and fairly robust to changes in the team size and the network size as explained in Section 7. In the case that we do need to identify team members, one must account for the mean rate of communication across time. The adaptive approach allows for this to occur across the network as well as temporal changes like day of the week and seasonal influences. We use a one-step-ahead forecast to establish expected communication counts using a Poisson regression. In particular, the λ values were used to build the following regression model:

$$\begin{aligned}\log(h_D(\lambda,n)) = {} & \beta_0 + \beta_1 n + \beta_2 n^2 + \beta_3 n^3 + \beta_4 \lambda + \beta_5 \lambda^2 \\ & + \beta_6\ \mathbb{I}(\lambda < 0.95) + \beta_7\ \mathbb{I}(\lambda < 0.95)\lambda \\ & + \beta_8 \log(\lambda) + \beta_9 n \log(\lambda) + \beta_{10} n\lambda + \beta_{11} n\lambda^2 \\ & + \text{error}.\end{aligned}$$

Once fitted, the above regression model was used to estimate the thresholds for the DEWMA$_{\nu,t}$ plan for homogeneous networks with mean count λ and size n. The above fitted model delivers an in-control ATS within 100 ± 15 for the range of $100 \le n \le 400, 0.01 \le \lambda \le 1.0$ and $\alpha = 0.10$. The standard error of the model was 0.0043 and the correlation between the model-fitted values and the corresponding actual simulated $h_D(\lambda, n)$ values was 0.9996.

For the GEWMA$_t$ plan, we estimated the threshold $h_G(\lambda, n)$ in a similar way as above. We generated networks of size $n = 100, 125, 150, \ldots, 975, 1000$, fixed $\alpha = 0.10$, and simulated homogeneous networks with mean counts $\lambda = 0.01,\ 0.02, 0.3,\ \ldots, 0.1, 0.15, 0.2, \ldots, 0.95, 1.0$. For each combination, we estimated the threshold $h_G(\lambda, n)$ through simulation, and then used these estimates to build the following regression model:

$$\begin{aligned}1/h_G(\lambda,m) = {} & \beta_0 + \beta_1 \log(\lambda) + \beta_2 n + \beta_3 n^2 + \beta_4 n^3 + \beta_5 \lambda \\ & + \beta_6 \lambda^2 + \beta_7 \lambda^3 + \beta_8 \log(n) + \beta_9 \log(\lambda) n \\ & + \beta_{10} \log(\lambda) n^2 + \beta_{11} \log(\lambda) n^3 + \beta_{12} n\lambda \\ & + \beta_{13} n^2 \lambda + \beta_{14} n^3 \lambda + \beta_{15} \lambda^4 + \beta_{15} \lambda \log(n) \\ & + \beta_{16} \lambda^5 + \beta_{17} \lambda^2 \log(n) + \beta_{18} \lambda^3 \log(n) + \text{error}\end{aligned}$$

The above model estimates the thresholds for the GEWMA$_t$ for homogeneous counts and obtains an in-control ATS of 100 ± 7 for $100 \le n \le 1000, 0.01 \le \lambda \le 1$ and $\alpha = 0.10$. The standard error of the model was 0.0007 and the correlation between the model-fitted values and the corresponding actual simulated $h_D(\lambda)$ values was 0.99999.

Simulation study results

Below, we provide tables for the simulation results described in Section 5.