Taylor & Francis
Taylor & Francis Group

Quality Engineering

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/lqen20

Foundations of network monitoring: Definitions and applications

Nathaniel T. Stevens, James D. Wilson, Anne R. Driscoll, Ian McCulloh, George Michailidis, Cecile Paris, Kamran Paynabar, Marcus B. Perry, Mostafa Reisi-Gahrooei, Srijan Sengupta & Ross Sparks

To cite this article: Nathaniel T. Stevens, James D. Wilson, Anne R. Driscoll, Ian McCulloh, George Michailidis, Cecile Paris, Kamran Paynabar, Marcus B. Perry, Mostafa Reisi-Gahrooei, Srijan Sengupta & Ross Sparks (2021) Foundations of network monitoring: Definitions and applications, Quality Engineering, 33:4, 719-730, DOI: 10.1080/08982112.2021.1974033

To link to this article: https://doi.org/10.1080/08982112.2021.1974033

Published online: 14 Oct 2021.

Submit your article to this journal

Article views: 95

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=lqen20

QUALITY ENGINEERING
2021, VOL. 33, NO. 4, 719–730
https://doi.org/10.1080/08982112.2021.1974033

Taylor & Francis
Taylor & Francis Group

Check for updates

THE PAST, PRESENT, AND FUTURE OF NETWORK MONITORING: A PANEL DISCUSSION

Foundations of network monitoring: Definitions and applications

Nathaniel T. Stevens[a*], James D. Wilson[b*], Anne R. Driscoll[c], Ian McCulloh[d,e**], George Michailidis[f**], Cecile Paris[g**], Kamran Paynabar[h**], Marcus B. Perry[i**], Mostafa Reisi-Gahrooei[f**], Srijan Sengupta[j**], and Ross Sparks[g**]

[a]University of Waterloo, Waterloo, Ontario, Canada; [b]University of Pittsburgh, Pittsburgh, Pennsylvania, USA; [c]Virginia Polytechnic Institute and State University, Blacksburg, Virginia, USA; [d]Accenture Federal Services, Washington, District of Columbia;[e]Johns Hopkins University, Baltimore, Maryland [f]University of Florida, Gainesville, Florida, USA; [g]CSIRO, Sydney, NSW, Australia; [h]Georgia Institute of Technology, Atlanta, Georgia, USA; [i]University of Alabama, Tuscaloosa, Alabama, USA; [j]North Carolina State University, Raleigh, North Carolina, USA

ABSTRACT

In this article, the panelists broadly discuss the definition of network monitoring, and how it may be similar to or different from network surveillance and network change-point detection. The discussion uncovers ambiguity and contradictions associated with these terms and we argue that this lack of clarity is detrimental to the field. The panelists also describe existing and emerging applications of network monitoring, which serves to illustrate the wide applicability of the tools and research associated with the field.

KEYWORDS

graphs; network monitoring; network science; statistical process monitoring; surveillance

Question

The problem of detecting change in network data goes by different names in different fields, and some names reflect different nuances of the problem. With the aim of formulating an encompassing definition for this area of work, discuss what comes to mind when you read the phrases "network monitoring", "network surveillance", or "network change-point detection". What is your preferred moniker for this discipline, and why?

McCulloh

Perhaps no other field of science suffers from more ambiguous and conflicting definitions than network science, known in different disciplines as social network analysis, or graph theory, or cyber-physical networks, or neural networks. The oldest definition, graph theory, encompasses the mathematical foundations for the analysis of entities and their relationships. Unfortunately, to the layman, a graph is often confused with a chart or data visualization, thus networks gain more traction in business applications. So much of the terminology used is propelled by what is understood by the nontechnical business leader that drives the applications and use-cases for science.

Many different disciplines have contributed to the scientific literature surrounding networks, often unaware of the advances in those other disciplines. For example, it is very reasonable that a physicist exploring network science in the 21st century would be unaware of Moreno's publication proposing social physics in 1934 involving social interactions among elementary students and launching the field of modern-day social network analysis (Moreno 1934).

The growth of data and modern awareness of networks has increased the number of people conducting research in the field and has given rise to greater interest in network problems. For example, Barabasi and Albert's concept of preferential attachment (Barabási and Albert 1999) published in 1999 was proposed in research correspondence between Nobel laureates Herbert Simon, Allen Newell, and their colleague Harrison White in the 1950s. At that time, however, the availability of data, computational resources, and applied problems left their work in relative obscurity for nearly a half century.

Despite the growth of scholarly network publications and the emergence of academic conferences on

CONTACT Nathaniel T. Stevens nstevens@uwaterloo.ca University of Waterloo, Waterloo, Ontario, Canada; James D. Wilson wilsonj41@upmc.edu University of Pittsburgh, Pittsburgh, Pennsylvania, USA.
*Discussion Editors.
**Discussion Panelists.

© 2021 Taylor & Francis Group, LLC

network science, the term network still finds multiple meanings and multiple applications. In my experience, applications are as diverse as computer/information technology (IT) networks, social networks, or protein interaction networks in biology. The terms "network monitoring" or "network surveillance" in an IT network treats the network as a fixed, tangible object and sensors are placed to detect certain signatures across the network that may indicate a threat. This is very different from monitoring a social network for organizational changes or social media for the introduction of propaganda campaigns.

To my knowledge the earliest use of the term "network monitoring" was McCulloh, Carley, and Webb (2007) where statistical process monitoring methods were applied to intelligence data on the Al-Qaeda terrorist organization. This research, which began in 2004 was initiated at the U.S. Military Academy at West Point following the National Research Council report on Network Science and established the Army's Network Science Center. At that time, network monitoring was the application of statistical process monitoring methods to highly dependent social network data collected on adversarial and extremist networks. Also referred to as social network change detection (SNCD), these methods enabled detection of small changes in media and online information networks that provided reliable early warnings to large-scale, coordinated attacks in Iraq and Afghanistan. Military data scientists were able to recognize that the change in an adversarial organization did not occur at the time of attack, but rather as the organization began to plan and resource the attack. The application of network monitoring allowed analysts to detect organizational changes quickly, before attacks occurred, enabling them to operate within the enemy's decision cycle and gain advantage.

These methods have also been applied in the context of cyber network security. Rather than simply looking for gross changes in IP traffic or some other metric, social network methods allow analysts to model the structure of the cyber network and detect small changes resulting from adversarial attacks or other threats in a zero-trust framework. In this manner, network monitoring does not mean surveillance of specific actors or content, but rather the application of statistical process monitoring methods to graph-level social network measures that may indicate structural changes in the network. Given the ambiguity and wide range of terms and definitions, my preferred term of choice is "statistical network monitoring" to indicate the type of analysis or methods and the use of the more common lay term for dependent data.

Michailidis

Networks are complex objects consisting of nodes and edges and the terms network monitoring, network surveillance and network change-point detection depend both on the nature of the network under consideration and its function. For example, there are physical and/or engineered systems with network structure, such as protein-protein interaction networks in biology, or computer, communications, power and road networks. The nodes in such networks correspond to physical entities (e.g., proteins, computers or other network devices, cities) and the edges are also physical entities or express physical interactions (e.g., a transmission line in a power network, a road connecting two destinations in a road network, or the interaction process between binding regions of proteins in the respective biological network). Further, such networks carry out specific functions (e.g., power from generating stations to consumers, or information in the form of packets from one computer to another). In the case of engineered systems, networks are designed with specific operating characteristics in mind (e.g., amount of traffic to be carried per unit of time, degree of resilience in fulfilling its function in the presence of various interruptions, etc.).

There are also other types of networks, such as social and even "association" networks. Nodes in social networks correspond to various types of social actors and edges represent various forms of interactions (e.g., physical, organizational, digital, etc.), whereas nodes in "association" networks can represent both physical entities similar to those in physical or social networks, but also properties of them (e.g., expression levels of genes, blood oxygen level dependent response of brain regions as measured by functional magnetic resonance imaging technologies, or some measure of liquidity of commercial and investment banks), while the edges correspond to some measure of statistical association between the nodes. Such networks may have multiple functions, including connecting people and exchanging information through a social media platform, to developing hypotheses to be validated through follow-up experimentation in the case of changes in the gene association network under disease progression.

Network monitoring for physical networks usually refers to systematic processes put in place by their owners/administrators to ensure that all its

components operate according to operating specifications and the network as a whole fulfills its function. Taking a computer network as an example, its administrator is interested in identifying in a timely manner crashed, frozen or overloaded servers, failing/failed switches and routers, and so forth. The end goal is to intervene to fix failures or initiate preventive maintenance actions. In the longer term, network monitoring can also aid in capacity planning to upgrade the capabilities of the network, by improving the underlying infrastructure (e.g., install new routing software, or higher throughput switches in a computer network) and also expanding the capacity of the edges (e.g., add lanes in a road network segment). On the other hand, network monitoring for other types of networks usually refers to identifying stable patterns in its characteristics and also transient or long-lasting deviations from them, as well as emerging trends. For example, in a network that captures voting patterns amongst legislators, network monitoring may reveal patterns of ideological cohesion of members of both the same political party and across parties, and also realignments over time. At the technical level, network monitoring relies on techniques of outlier, anomaly and change-point detection.

Network surveillance for physical networks shares many of the objectives of network monitoring, but a key distinction is that it can be done by organizations that are not the network owners/administrators and also possibly in a covert manner. For example, whereas the administrator of a road network may deploy loop detectors and traffic cameras to monitor traffic flows, a third party with access to geolocation information from drivers' mobile devices can accomplish the same task. In the former case, the signal monitored corresponds to actual vehicle flows, whereas in the latter to a proxy that represents the evolving position of the mobile device. Nevertheless, both signals can provide accurate real time information for the status of the road network. Similar examples arise for computer and communication networks. On many occasions, the network administrator has a much more granular view than a third party engaged in network surveillance.

Network change-point detection broadly aims to identify changes over time in network characteristics, in the form of deviations from and return to baseline patterns, or emergence of new operational regimes. In a physical/engineering network whose function is to move traffic through it, change-point analysis usually focuses on the status of different flows between consecutive, or between pairs of source and destination nodes. On the other hand, in other types of networks, change-point analysis and detection usually focuses on changes in the structure of the topology, and in the case of "association" networks, in the parameters of the posited data generating mechanism. For example, suppose one has access to a sequence of observed edges for the same set of nodes. One type of change-point analysis may focus on understanding changes in the global clustering coefficient, a measure of the degree to which nodes in the network tend to cluster together. Another change-point analysis starts by assuming that the observed network over time is a realization from a stochastic block model network generating mechanism; then, the focus of the analysis becomes to identify at what points in time the parameters of the posited model changed and in what ways.

Sengupta

To start on a lighthearted note, I am reminded of Shakespeare's famous quote from Romeo and Juliet: "What's in a name? That which we call a rose, by any other name would smell as sweet."

To me, the phrase "network monitoring" indicates the classical two-phase process monitoring framework applied to network data (Woodall and Montgomery 2014; Woodall et al. 2017), and in our own work we have preferred this moniker (Zhao et al. 2018a, 2018b; Kodali et al. 2020). In Phase I, the user collects a sample of time-varying networks that represent the in-control state, and uses this sample to gain an understanding of the in-control behavior. In Phase II, the user observes networks successively over time, and the goal is to determine whether there is a significant deviation from the in-control state. To make this determination at a given point in time, the user is only allowed to use the Phase I sample and the Phase II data collected until that specific time point.

To me, "network surveillance" is *almost* synonymous with "network monitoring", and several prominent researchers have used it (Jeske et al. 2018a, 2018b). The term "surveillance" is commonly used in geospatial public health surveillance (Declich and Carter 1994; Teutsch and Churchill 2000; Patil and Taillie 2003; Kulldorff et al. 2006), where, after a deviation is detected, it is important to also identify the specific area which is the source of the deviation. This task of localizing the deviation is an intrinsic part of surveillance. Whereas in monitoring, the analogous task of root cause analysis is important, but considered a separate, downstream task – not an intrinsic part of monitoring (Zhou, Chen, and Shi 2004; Dey and Stori 2005). Therefore, it might be

preferable to use "network surveillance" when localizing the part of the network that causes the deviation is considered an intrinsic task, and to use "network monitoring" otherwise.

Some researchers use the phrase "network change-point detection" to indicate a related but somewhat different problem. Here, typically there is no Phase I and Phase II like traditional statistical process monitoring. Instead, the entire stream of network data is available to the user, and the goal is to retrospectively estimate time points when the model changed. This problem definition has been used in several recent statistical papers on network change-point detection, such as (Barnett and Onnela 2016; Bhattacharjee, Banerjee, and Michailidis 2018; Wang, Yu, and Rinaldo 2021).

On a related note, it is important to distinguish between monitoring of network data and monitoring of data over a network (Sengupta and Woodall 2018). The first case arises when we have a time series of networks and the network structure itself, i.e., nodes and edges, is the time series variable of interest, e.g., social networks. The second case arises when the network itself is fixed over time, and we observe certain node-level or edge-level variables, and we are interested in monitoring these variables over time, e.g., the transmission of electric power via the power grid network.

Driscoll

"Network change-point detection" is my least favorite for describing the field of detecting changes in network data. Change-point methods are commonly used in statistical process monitoring during Phase I. These methods help establish an understanding of baseline behavior to prepare for real-time monitoring in Phase II. It is important to use a term for analyzing network data that is broad enough for the entire field of study and I find the term change-point detection to be too specific, applying only to a subset of applications.

In many publications, the terms "network surveillance" and "network monitoring" seem to be used interchangeably. I have used both terms to address network-change research and initially didn't think that the terms had different meanings. Upon further consideration, I find network surveillance to refer to a specific application of network monitoring. One dictionary definition of surveillance is, "close observation, especially of a suspected spy or criminal." The term "network surveillance" seems to relate more closely to the use of monitoring methods to detect or prevent negative outcomes. Savage et al. (2014) listed application areas including the detection of important and influential network participants, the detection of clandestine organizational structures, and the detection of fraudulent or predatory activity. The former two application areas are broader and would fall under the general term "network monitoring"; using monitoring methods to detect fraudulent or predatory activity would more closely align with the term "network surveillance". Some other applications of network surveillance would include identifying key players in terrorist networks and detection of spammers. I would argue that much of the research in the field up to this point develops techniques to apply to network surveillance applications.

It is important to use a broad term that will encourage researchers to develop tools that can be used in a broad range of application areas, not just to detect negative outcomes. In many cases, the same monitoring technique can successfully be applied to a wide variety of applications, so by using the term "network monitoring" we are not limiting the impact of our methods to one subject area. In my opinion, "network monitoring" is the most all-encompassing of the three terms suggested to describe detecting changes in network data, and so it is my preferred moniker for this discipline.

Perry

Network science is a vast multidisciplinary field of study spanning such disciplines as physics, social science, computer science, statistics, and more. It often involves the study of complex systems by representing these systems as networks, e.g., see Newman (2018). The area of network monitoring is a smaller subset of the network science discipline and involves the monitoring and detection of important changes in one or more aspects of network systems over time. Although there are several monikers, my preferred is "statistical network monitoring", so as to distinguish between those monitoring strategies that are mostly descriptive and those that are inferential. The former involves the calculation of a sample network charting statistic but without any real attempt to accurately quantify the uncertainty in the statistic, e.g., see Tambayong (2014). Although more challenging, the latter does make attempts to accurately quantify the uncertainty in the charting statistic, and consequently, is more in line with the philosophies of statistical process monitoring (SPM), e.g., see Azarnoush et al. (2016) and Perry (2020). That is, control limits for the charting statistic can be established such that the false alarm

rate of the monitoring scheme can be adequately controlled at some user-specified level.

Reisi and Paynabar

Characterizing the scope and objectives of this field is crucial to identifying an appropriate and informative terminology. The problem of detecting change in network data mainly focuses on identifying a subset of network data that has spatially or temporally different behavior from a dominant normal set. This can be studied in both static, where a single snapshot of the system is considered, and dynamic networks, where a sequence of snapshots is analyzed over time. The objective of change detection for each type is different (Ranshous et al. 2015). In the former, which is mainly referred to as "network anomaly detection", the main objective is to identify a subset of nodes that illustrates different interactions from that of the overall network (Akoglu, McGlohon, and Faloutsos 2010; Eberle and Holder 2006; Noble and Cook 2003). For example, by analyzing the Enron network generated from email communications between 1998 to 2002, Akoglu, McGlohon, and Faloutsos (2010) identified a set of employees with abnormal behavior. These employees were those who played a role in the famous Enron scandal.

In the latter, however, a user is interested in identifying time intervals during which the network structure or node interactions globally or locally changes from an expected dynamic behavior (Woodall et al. 2017). For example, in the Enron network, one may investigate to identify the onset of the scandal by analyzing a sequence of networks which represent weekly email communications of employees (Dong, Chen, and Wang 2020; Gahrooei and Paynabar 2018). The phrases "network monitoring", or "network change-point detection" are appropriate for these cases. In most change detection applications of dynamic networks, in addition to identifying the time of change, one would be interested in pinpointing the portions of the network affected by assignable causes. This is known as change diagnosis in the statistical process monitoring (SPM) literature, which utilizes techniques from both areas of dynamic and static networks (Woodall et al. 2017; Ranshous et al. 2015). "Network surveillance'" is mainly used for change detection for computer and internet networks over time (Jeske et al. 2018a) and may be perceived as less general.

In short, it would be difficult to suggest one name that encompasses all aforementioned cases. As the term "anomaly detection" is common across different disciplines, we would like to propose using "static network anomaly detection" and "dynamic network anomaly detection (and diagnosis)."

Editors' comments

The reader could be forgiven for getting lost in these (at times) contradictory viewpoints and definitions of *network monitoring*, *network surveillance*, and *network change-point detection*. In the panelists' responses we see variation in the use of these terms depending on who is performing the investigation, and if

- the investigation is prospective or retrospective,
- the context is to investigate change in network data or data over physical networks,
- the networks are static or dynamic,
- the procedures are descriptive or inferential,
- the intent is to uncover the root cause of changes,
- global or local changes are of interest,
- detecting malicious threats is important.

We think Dr. McCulloh says it best: the field *suffers* from ambiguous and conflicting definitions. This confusion makes it difficult to define what the field of network monitoring is, and what it isn't. We feel this is more than a pedantic argument about semantics; when a discipline is not well-defined, it becomes unclear who the stakeholders are, who values the research, and what research is relevant. This has implications for deciding where to disseminate research and which problems to work on. It also has implications for ease and productivity of both interdisciplinary and intradisciplinary collaboration.

Dr. Driscoll advocates for the adoption of a term that broadly encompasses the many different nuances and facets associated with detecting change in networks. We agree with this sentiment, but we also see the need for a unified collection of terms that clearly and specifically address some key characteristics of the problem. However, we also acknowledge that given existing confusions surrounding terminology, one should avoid introducing new terms unless they are informative and helpful.

Although there is disagreement in subtleties associated with the terms *network monitoring*, *network surveillance*, and *network change-point detection*, there appears to be some consensus in that the term *network monitoring* is the most general and therefore the more preferable term to broadly refer to the body of work devoted to identifying change in networks. We subscribe to this viewpoint, but we propose that the

Table 1. A list of adjectives that may be used to modify the term "network monitoring".

Adjective	Description
online vs. offline	Used to distinguish contexts and methods in which changes are identified prospectively vs. retrospectively.
statistical vs. descriptive	Used to distinguish contexts and methods in which the uncertainty of one's decision is vs. is not emphasized.
diagnostic	Used to describe contexts and methods in which identifying the cause of the change is as important as identifying the change itself.
static vs. dynamic	Used to distinguish contexts and methods in which one wishes to identify change in a single snapshot vs. a stream of snapshots of a network.
local vs. global	Used to distinguish contexts and methods in which one wishes to identify local vs. global change in a network.
physical vs. association	Used to distinguish contexts involving physical vs. association networks.

term be prefaced by adjectives that clearly indicate the context of its use. For instance, when performing *online* (i.e., prospective and in real-time) network monitoring versus *offline* (i.e., retrospective) network monitoring as described by Dr. Michailidis, we suggest that these contexts be referred to as *online network monitoring* and *offline network monitoring*, respectively. Similarly, *statistical network monitoring* versus *descriptive network monitoring* might be used to distinguish contexts and methods in which uncertainty quantification is and is not emphasized, as suggested by Dr. Perry. Such adjectives may also be combined, so as to generate maximally descriptive terms. For instance, *offline statistical network monitoring* might be used to refer to that which the term *network change-point detection* commonly refers. Table 1 contains a list of adjectives that we propose be used with "network monitoring", and that give rise to a unified collection of terms that may be used to accurately describe the many applications and contexts discussed by the panelists.

Question

What are the predominant applications to which network monitoring is applied now, and what do you see as emerging application areas? Are there other applications for which network monitoring could be used that have not been explored?

Michailidis

Network analysis and monitoring has become ubiquitous across science and engineering. This is due to a number of factors, including new measurement technologies (see, e.g., the emergence of various Omics technologies in molecular biology, or social media platforms), and the ability of organizations to access, store and process vast amounts of data (see, e.g., geolocation and data obtained from sensors). Another driving factor has been an emerging focus on a systems view. Systems are entities with interrelated and interdependent components, wherein changes in some of them affect other components and the system as a whole. Networks offer a powerful paradigm for their analysis and monitoring. Hence, once new measuring technologies are in place for a new application domain, network analytics would soon follow.

Reisi and Paynabar

Network monitoring has been predominantly applied to social and organizational behavior analysis (Woodall et al. 2017). For example, network monitoring methods have been widely applied to the Enron email corpus to identify sudden changes in communications among Enron employees during the company's scandal (Dong, Chen, and Wang 2020; Gahrooei and Paynabar 2018). Network monitoring has also been applied to communication networks of terrorist groups (e.g., Al-Qaeda) to detect changes in their level of communications as early as possible before it turns into a crisis (McCulloh, Carley, and Webb 2007). Other examples of applying network monitoring to social behavior analysis can be found in (McCulloh and Carley 2011).

Analysis of financial networks is another application in which network monitoring methods have been utilized. In this context, by monitoring interactions among financial institutes over time, financial shocks may be detected. For instance, Ebrahimi et al. (2021) developed a framework to model and monitor the dynamics of financial interbank lending networks and showed that such a network monitoring framework could have raised alarms to the public prior to the key events of the 2007–2009 financial crisis. Brunetti et al. (2019) also studied two interbank market networks: correlation networks based on publicly traded bank returns and physical networks based on interbank lending transactions. Both networks depicted change in their level of interconnectedness during the financial crisis.

Detecting changes in computer and information networks is another major application of network monitoring. In computer networks, the goal is to detect malicious intrusion attempts, including spam campaigns, phishing, denial of service attacks, or any

other anomaly as quickly as possible with the minimum number of false alarms. In information networks, metrics that indicate the health of the network, such as traffic flow, are monitored for detection of such abnormalities (Jeske et al. 2018a).

Emerging applications of network monitoring can be found in transportation, neuroscience, cybersecurity, and analysis of spatiotemporal data. In transportation, traffic patterns at the network level are monitored to detect changes due to extreme events. Ilbeigi (2019) modeled New York taxi data as a weekly sequence of traffic networks and monitored the main topological features of the network to detect changes in traffic patterns before and after hurricane Sandy. Due to the size and specific structure of road transportation networks, novel network monitoring algorithms that are tailored to traffic networks are needed.

One recent application of network monitoring is in the area of neuroscience, more specifically, for analysis of functional brain networks that are defined based on functional magnetic resonance imaging (FMRI) or electroencephalogram (EEG) data. Functional brain networks are analyzed to understand the dynamics and modular behavior of the brain for normal people and patients suffering from neurological diseases such as epilepsy and Parkinson's disease (Bassett et al. 2011; Lynall et al. 2010; Bassett and Bullmore 2006). Network monitoring can be integrated with these modeling methods to detect changes in the functional brain network dynamics. For example, early detection and identification of the onset of seizure attacks may help with better understanding of how brain functionality changes during these attacks.

Cybersecurity is another emerging application of network monitoring that has not been broadly studied in the literature. The growing interconnectivity within the systems by sensors and Internet of Things (IOT) devices requires new advances in faults and intrusion detection to improve the resiliency and functionality of the cyber-physical systems under complex operational conditions. Network monitoring can play a significant role in these advances by integrating the data-driven and model-based methods. For instance, Li et al. (2021) integrated the state-space modeling and group lasso techniques to identify and localize cyberattacks in smart grids. Finally, network monitoring methods can be used for monitoring spatio-temporal data streams. That is, one may transform the spatiotemporal data into a sequence of networks (for example, based on the spatial correlation of data streams), and use network tools for detection of changes that are difficult to detect in the original space. For instance, Kan and Yang (2017) transformed real time image data into a sequence of networks to be monitored for detecting abrupt changes.

Sengupta

Cybersecurity and social network analysis appear to be two primary application areas. I see infrastructural networks, such as power grid networks and traffic networks, to be emerging application areas. One very promising area where network monitoring could be used is neuroscience, specifically in studying the dynamics of human brain networks (Stevens et al. 2009). Network monitoring has the potential to become a powerful and practical tool to detect or even predict abnormalities in the brain network toward diagnosis of diseases.

McCulloh

Network methods offer a powerful framework for the analysis of highly dependent data, in situations where the independent and identically distributed (IID) assumption does not hold. As a result, the most common applications of network monitoring are in areas where data is highly dependent. In my work, this is typically organizational change or changes in communication patterns among online communities. Increasingly, these methods are applied within zero-trust cyber networks to identify anomalous patterns that may signal malicious behavior.

Given the wide range of network applications, there are no doubt many other areas where network monitoring may be applied. For example, applications in life sciences, physical infrastructure networks, and even transportation networks likely apply some form of network monitoring today.

As applications for Artificial Intelligence (AI) continue to grow, there may be increased opportunities to monitor algorithms or native cloud processes for change. For example, BERT neural networks have created a significant advance for natural language processing (Devlin et al. 2018). These methods are so powerful, algorithms can learn language pre-processing steps, making data and feature engineering more important than data pre-processing for ultimate performance. Perhaps network monitoring methods can be applied to neural networks to better detect linguistic drift, changes in discourse, or the presence of unintended bias over time.

Driscoll

Network monitoring applications have predominantly focused on network surveillance problems where the goal is to detect or prevent some adverse event. Examples include terrorist networks such as the Al-Qaeda network seen in Figure 1 of Woodall et al. (2017) and the Enron email communications depicted in Figure 2 of the same paper. Another relevant application area in this vein is cybersecurity, where interest lies in the development of tools and methods for the detection and subsequent prevention of cyber-attacks, a threat all too common in our society.

Future research should also consider applications in other areas where the goal is not to detect negative outcomes. Network monitoring can alert us to important or influential network participants in a wide variety of applications. In Motalebi, Stevens, and Steiner (2021), the researchers illustrate the use of hurdle blockmodels through two examples. One application uses the infamous Enron email data and the second example reveals the relevance of the method with a small research collaboration network from the University of Waterloo's Statistics & Actuarial Science Department. This thought-provoking example shows how we might create our own network datasets to illustrate our methods.

As we adapt and develop network monitoring techniques for use in additional application areas, we should also strive to produce more modern and diverse and publicly available datasets. The Enron email scandal resulted in the collapse of the company in 2001. This was 20 years ago, yet the email exchange data (Shetty and Adibi 2005) is still used extensively to illustrate and investigate monitoring methodologies. It would be beneficial to the field if additional datasets related to more current events were made available. However, given the proprietary nature of many network monitoring applications, sharing data publicly may prove to be difficult.

Perry

The predominant applications of statistical network monitoring strategies seem to emerge from computer networks, e.g., monitoring for computer network traffic anomalies (Neil et al. 2013). However, social science applications have also emerged as motivations for developing network monitoring strategies, e.g., see applications highlighted in Savage et al. (2014), Wilson, Stevens, and Woodall (2019), and Perry (2020). In the future, many emerging applications of statistical network monitoring will likely stem from Industry 4.0.

According to the National Institute of Standards and Technology (NIST): "Industry 4.0 refers to the fourth industrial revolution, which connects machines, people, and physical assets into an integrated digital ecosystem that seamlessly generates, analyzes, and communicates data and sometimes takes action based on that data without the need for human intervention." Given that the emphasis of Industry 4.0 is interconnectivity, automation, machine learning, and real-time data, it is my view that the future of statistical network monitoring research for Industry 4.0 applications is fruitful.

Sparks and Paris

We first note that by "network data", we refer primarily to social media data, in which people post messages, sometimes directed to someone else, sometimes simply to the public, or the whole social media community at large. We call each message a communication event, and these events form a network, either through time or in terms of the relationships among those who post these messages. Detecting changes within the network data in this context amounts to determining whether an unusual pattern of communication is occurring. We note that this work often requires a multidisciplinary approach; our team consists of statisticians, natural language processing experts, and domain experts to inform what type of change we should be investigating.

Social media networks have revolutionized the way users interact. In this context, sentiment analysis, opinion mining, or emotion detection are important tools for determining feeling or emotion through text (Ahmed, Tazi, and Hossny 2015; Kaur and Saini 2014; Ravi and Ravi 2015; Bakshi et al. 2016; Yadollahi, Shahraki, and Zaiane 2017). In these applications, natural language processing is used to define sentiments, emotions, and feelings directly from the text used in the message. The applications are broadly located in medical, social and industrial domains.

The interface between psychology, classical sociology and social networks is an expanding field for wellness monitoring (Healey and Logan 2005). The need to monitor the well-being of those within the social network requires the interface between computer science, classical sociology, natural language processing, statistics, social networks analysis and psychology (Assembling Teams of Experts in Bonchi et al. (2011)). Monitoring activities in this context aim to identify individuals who

are at risk of self-harm or suicide on the basis of the messages they send. In our own work, a team which includes Professor Maria Kangas (Macquarie University Psychology Department) is currently monitoring the network data from a social media platform dedicated to the expression and sharing of emotions (the Vent platform). Our aim is to identify at-risk individuals that belong to the Vent network (Malko et al. 2021). The vents (the posts on the platform) are self-annotated with a group of core emotions, and other event labels such as Halloween or Valentine's Day, group labels such as LGBT and Asian heritage, and interests labels, like Star Wars or gaming.

Other applications of social media monitoring include monitoring criminal activities via social networks; see Sparks and Wilson (2019), and Basu and Sen (2021), for examples. The organizational criminal activities of outlaw motorcycle gangs have been explored by Bright and Deegan (2021).

People have also studied social influence in an attempt to quantify the amount of influence some people have on others within a social network, thereby identifying highly influential individuals. Such analyses are crucial to accommodate the needs of social network applications; identifying opinion leaders has widespread applicability (Oueslati et al. 2021). There are many new problems and challenges in this context, however. We believe that the problems are dynamic in nature and depend on the topic of discussion because of the diversity of social networks (Peng et al. 2018; Samanta, Dubey, and Sarkar 2021). Distinguishing between positive influence and negative influence is also a challenge. Difficulty measuring influence of individuals within a large social network with big data is another challenge.

Social media monitoring has also been shown to enhance natural disaster management by providing real-time data on disasters such as tsunamis, typhoons and floods (Kim and Hastak 2018; Xue et al. 2021). In this context, a substantial amount of research has been done by various researchers to analyze communication events during disasters. The aim is to develop effective ways of analyzing and extracting critical information from such communications so as to detect missing people during and after an earthquake or tsunami, for example.

Enterprise applications of social media monitoring are also common. For example, a company can "listen" to social networks to learn what their customers say about the company, their competitors, and the market in general. Taking advantage of social media in this way provides important social business intelligence that may improve a company's competitive edge and their overall success; such insights enable the company to better understand their customer-base, and perceptions of their brand, facilitating tailored marketing strategies and customer-focused decision making. There are a broad range of monitoring tools available for that purpose, e.g., Awario; Audiense; Brand24 and Brandwatch Consumer Research to name a few.

Editors' comments

The panelists identify a number of existing and emerging network monitoring applications, ranging from cybersecurity in computer networks to disease diagnosis in brain networks. In terms of existing applications, network monitoring has traditionally been used for purposes of threat detection in online social and communication networks. Network monitoring methods have also been used to detect behavioral change in organizational networks, as well as malicious and adversarial attacks on physical, computer networks. While these applications are still very relevant, the modern emphasis on interconnected and networked systems is giving rise to a variety of novel contexts for which network monitoring is a useful tool. These emerging application areas include institutional networks in finance, protein-protein interaction networks in biology, functional brain networks in neuroscience, and physical infrastructure networks such as power networks, transportation networks, and supply chain networks, to name a few. As the world becomes evermore interconnected, and as advances in technology facilitate the observation and measurement of this interconnectedness, new networks and hence network monitoring applications will arise. Like Dr. Perry, we anticipate that Industry 4.0 will provide many such opportunities.

In the context of social media monitoring, Drs. Sparks and Paris identify a number of interesting and important applications including monitoring disease spread, mental well-being and criminal activities. Social media monitoring may also be used to good effect in the contexts of natural disaster management, and business insights. As social media becomes increasingly enmeshed in our daily lives, such monitoring applications are bound to increase in prominence and importance.

ORCID

Nathaniel T. Stevens http://orcid.org/0000-0001-6149-5797
James D. Wilson http://orcid.org/0000-0002-2354-935X
Ian McCulloh http://orcid.org/0000-0003-2916-3914
George Michailidis http://orcid.org/0000-0002-3676-1739

Kamran Paynabar http://orcid.org/0000-0002-6906-3611
Srijan Sengupta http://orcid.org/0000-0001-6889-8599
Ross Sparks http://orcid.org/0000-0001-5852-5334

References

Ahmed, K., N. E. Tazi, and A. H. Hossny. 2015. Sentiment analysis over social networks: An overview. In 2015 IEEE international conference on Systems, Man, and Cybernetics, pp. 2174–2179. IEEE.

Akoglu, L., M. McGlohon, and C. Faloutsos. 2010. Oddball: Spotting anomalies in weighted graphs. Pacific-Asia Conference on Knowledge Discovery and Data Mining, pp. 410–421. Springer.

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.

Bakshi, R. K., N. Kaur, R. Kaur, and G. Kaur. 2016. Opinion mining and sentiment analysis. 2016 3rd international conference on computing for sustainable global development (INDIACom), pp. 452–455. IEEE.

Barabási, A.-L., and R. Albert. 1999. Emergence of scaling in random networks. *Science (New York, N.Y.)* 286 (5439):509–12. doi:10.1126/science.286.5439.509.

Barnett, I., and J.-P. Onnela. 2016. Change point detection in correlation networks. *Scientific Reports* 6 (1):1–11. doi:10.1038/srep18893.

Bassett, D. S., and E. Bullmore. 2006. Small-world brain networks. *The Neuroscientist : A Review Journal Bringing Neurobiology, Neurology and Psychiatry* 12 (6):512–23. doi:10.1177/1073858406293182.

Bassett, D. S., N. F. Wymbs, M. A. Porter, P. J. Mucha, J. M. Carlson, and S. T. Grafton. 2011. Dynamic reconfiguration of human brain networks during learning. *Proceedings of the National Academy of Sciences of the United States of America* 108 (18):7641–6. doi:10.1073/pnas.1018985108.

Basu, K., and A. Sen. 2021. Identifying individuals associated with organized criminal networks: A social network analysis. *Social Networks* 64:42–54. doi:10.1016/j.socnet.2020.07.009.

Bhattacharjee, M., M. Banerjee, and G. Michailidis. 2018. Change point estimation in a dynamic stochastic block model. *arXiv preprint arXiv:1812.03090*.

Bonchi, F., C. Castillo, A. Gionis, and A. Jaimes. 2011. Social network analysis and mining for business applications. *ACM Transactions on Intelligent Systems and Technology* 2 (3):1–37. doi:10.1145/1961189.1961194.

Bright, D., and S. J. Deegan. 2021. The organisational structure, social networks and criminal activities of outlaw motorcycle gangs: Literature review. In *Trends and Issues in Crime and Criminal Justice (621)*, 1–16. Canberra: Australian Institute of Criminology.

Brunetti, C., J. H. Harris, S. Mankad, and G. Michailidis. 2019. Interconnectedness in the interbank market. *Journal of Financial Economics* 133 (2):520–38. doi:10.1016/j.jfineco.2019.02.006.

Declich, S., and A. O. Carter. 1994. Public health surveillance: Historical origins, methods and evaluation. *Bulletin of the World Health Organization* 72 (2):285–304.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dey, S., and J. Stori. 2005. A bayesian network approach to root cause diagnosis of process variations. *International Journal of Machine Tools and Manufacture* 45 (1):75–91. doi:10.1016/j.ijmachtools.2004.06.018.

Dong, H., N. Chen, and K. Wang. 2020. Modeling and change detection for count-weighted multilayer networks. *Technometrics* 62 (2):184–95. doi:10.1080/00401706.2019.1625812.

Eberle, W., and L. Holder. 2006. Detecting anomalies in cargo using graph properties. International Conference on Intelligence and Security Informatics, pp. 728–730. Springer.

Ebrahimi, S., M. Reisi-Gahrooei, K. Paynabar, and S. Mankad. 2021. Monitoring sparse and attributed networks with online hurdle models. *IISE Transactions*: 1–14. doi:10.1080/24725854.2020.1861390.

Gahrooei, M. R., and K. Paynabar. 2018. Change detection in a dynamic stream of attributed networks. *Journal of Quality Technology* 50 (4):418–30. doi:10.1080/00224065.2018.1507558.

Healey, J., and B. Logan. 2005. Wearable wellness monitoring using ECG and accelerometer data. Ninth IEEE International Symposium on Wearable Computers (ISWC'05), pp. 220–221. IEEE.

Ilbeigi, M. 2019. Statistical process control for analyzing resilience of transportation networks. *International Journal of Disaster Risk Reduction* 33:155–61. doi:10.1016/j.ijdrr.2018.10.002.

Jeske, D. R., N. T. Stevens, A. G. Tartakovsky, and J. D. Wilson. 2018a. Statistical methods for network surveillance. *Applied Stochastic Models in Business and Industry* 34 (4):425–45. doi:10.1002/asmb.2326.

Jeske, D. R., N. T. Stevens, J. D. Wilson, and A. G. Tartakovsky. 2018b. Statistical network surveillance. *Wiley StatsRef: Statistics Reference Online*:1–12.

Kan, C., and H. Yang. 2017. Dynamic network monitoring and control of in situ image profiles from ultraprecision machining and biomanufacturing processes. *Quality and Reliability Engineering International* 33 (8):2003–22. doi:10.1002/qre.2163.

Kaur, J., and J. R. Saini. 2014. Emotion detection and sentiment analysis in text corpus: A differential study with informal and formal writing styles. *International Journal of Computer Applications* 101 (9):1–9. doi:10.5120/17712-8078.

Kim, J., and M. Hastak. 2018. Social network analysis: Characteristics of online social networks after a disaster. *International Journal of Information Management* 38 (1):86–96. doi:10.1016/j.ijinfomgt.2017.08.003.

Kodali, L., S. Sengupta, L. House, and W. H. Woodall. 2020. The value of summary statistics for anomaly detection in temporally evolving networks: A performance evaluation study. *Applied Stochastic Models in Business and Industry* 36 (6):980–1013. doi:10.1002/asmb.2548.

Kulldorff, M., L. Huang, L. Pickle, and L. Duczmal. 2006. An elliptic spatial scan statistic. *Statistics in Medicine* 25 (22):3929–43. doi:10.1002/sim.2490.

Li, D., N. Gebraeel, K. Paynabar, and A. Meliopoulos. 2021. An online approach to cyberattack detection and localization in smart grid. *arXiv preprint arXiv:2102.11401*.

Lynall, M.-E., D. S. Bassett, R. Kerwin, P. J. McKenna, M. Kitzbichler, U. Muller, and E. Bullmore. 2010. Functional connectivity and brain networks in schizophrenia. *The Journal of Neuroscience* 30 (28):9477–87. doi:10.1523/JNEUROSCI.0333-10.2010.

Malko, A., C. Paris, A. Duenser, M. Kangas, D. Mollá, R. Sparks, and S. Wan. 2021. Demonstrating the reliability of self-annotated emotion data. Proceedings of the Seventh Workshop on Computational Linguistics and Clinical Psychology: Improving Access, pp. 45–54. doi:10.18653/v1/2021.clpsych-1.5.

McCulloh, I., K. Carley, and M. Webb. 2007. Social network monitoring of al-qaeda. *Network Science* 1 (1):25–30.

McCulloh, I., and K. M. Carley. 2011. Detecting change in longitudinal social networks. Technical report, Military Academy West Point NY Network Science Center (NSC).

Moreno, J. L. 1934. *Who shall survive?: A new approach to the problem of human interrelations*. Washington DC: Nervous and Mental Disease Publishing Co.

Motalebi, N., N. T. Stevens, and S. H. Steiner. 2021. Hurdle blockmodels for sparse network modeling. *The American Statistician*:1–11. doi:10.1080/00031305.2020.1865199.

Neil, J., C. Hash, A. Brugh, M. Fisk, and C. B. Storlie. 2013. Scan statistics for the online detection of locally anomalous subgraphs. *Technometrics* 55 (4):403–14. doi:10.1080/00401706.2013.822830.

Newman, M. 2018. *Networks*. Oxford: Oxford University Press.

Noble, C. C., and D. J. Cook. 2003. Graph-based anomaly detection. In Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 631–636. doi:10.1145/956750.956831.

Oueslati, W., S. Arrami, Z. Dhouioui, and M. Massaabi. 2021. Opinion leaders' detection in dynamic social networks. *Concurrency and Computation: Practice and Experience* 33 (1):e5692. doi:10.1002/cpe.5692.

Patil, G., and C. Taillie. 2003. Geographic and network surveillance via scan statistics for critical area detection. *Statistical Science* 18 (4):457–65. doi:10.1214/ss/1081443229.

Peng, S., Y. Zhou, L. Cao, S. Yu, J. Niu, and W. Jia. 2018. Influence analysis in social networks: A survey. *Journal of Network and Computer Applications* 106:17–32. doi:10.1016/j.jnca.2018.01.005.

Perry, M. B. 2020. An EWMA control chart for categorical processes with applications to social network monitoring. *Journal of Quality Technology* 52 (2):182–97. doi:10.1080/00224065.2019.1571343.

Ranshous, S., S. Shen, D. Koutra, S. Harenberg, C. Faloutsos, and N. F. Samatova. 2015. Anomaly detection in dynamic networks: A survey. *Wiley Interdisciplinary Reviews: Computational Statistics* 7 (3):223–47. doi:10.1002/wics.1347.

Ravi, K., and V. Ravi. 2015. A survey on opinion mining and sentiment analysis: Tasks, approaches and applications. *Knowledge-Based Systems* 89:14–46. doi:10.1016/j.knosys.2015.06.015.

Samanta, S., V. K. Dubey, and B. Sarkar. 2021. Measure of influences in social networks. *Applied Soft Computing* 99:106858. doi:10.1016/j.asoc.2020.106858.

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.

Sengupta, S., and W. H. Woodall. 2018. Discussion of statistical methods for network surveillance. *Applied Stochastic Models in Business and Industry* 34 (4):446–8. doi:10.1002/asmb.2354.

Shetty, J., and J. Adibi. 2005. Discovering important nodes through graph entropy the case of enron email database. In *Proceedings of the 3rd International Workshop on Link Discovery*, pp. 74–81. doi:10.1145/1134271.1134282.

Sparks, R., and J. D. Wilson. 2019. Monitoring communication outbreaks among an unknown team of actors in dynamic networks. *Journal of Quality Technology* 51 (4):353–74. doi:10.1080/00224065.2018.1507557.

Stevens, M. C., K. A. Kiehl, G. D. Pearlson, and V. D. Calhoun. 2009. Brain network dynamics during error commission. *Human Brain Mapping* 30 (1):24–37. doi:10.1002/hbm.20478.

Tambayong, L. 2014. Change detection in dynamic political networks: The case of sudan. In *Theories and simulations of complex social systems*, eds. V. Dabbaghian and V. K. Mago, 43–59. Berlin: Springer.

Teutsch, S. M., and R. E. Churchill, eds. 2000. *Principles and practice of public health surveillance*. New York: Oxford University Press.

Wang, D., Y. Yu, and A. Rinaldo. 2021. Optimal change point detection and localization in sparse dynamic networks. *The Annals of Statistics* 49 (1):203–32. doi:10.1214/20-AOS1953.

Wilson, J. D., N. T. Stevens, and W. H. Woodall. 2019. Modeling and detecting change in temporal networks via the degree corrected stochastic block model. *Quality and Reliability Engineering International* 35 (5):1363–78. doi:10.1002/qre.2520.

Woodall, W. H., and D. C. Montgomery. 2014. Some current directions in the theory and application of statistical process monitoring. *Journal of Quality Technology* 46 (1):78–94. doi:10.1080/00224065.2014.11917955.

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.

Xue, K., S. Guo, Y. Liu, S. Liu, and D. Xu. 2021. Social networks, trust, and disaster-risk perceptions of rural residents in a multi-disaster environment: Evidence from Sichuan, China. *International Journal of Environmental Research and Public Health* 18 (4):2106. doi:10.3390/ijerph18042106.

Yadollahi, A., A. G. Shahraki, and O. R. Zaiane. 2017. Current state of text sentiment analysis from opinion to emotion mining. *ACM Computing Surveys* 50 (2):1–33. doi:10.1145/3057270.

Zhao, M. J., A. R. Driscoll, S. Sengupta, R. D. Fricker, Jr, D. J. Spitzner, and W. H. Woodall. 2018a. Performance evaluation of social network anomaly detection using a moving window–based scan method. *Quality and Reliability Engineering International* 34 (8):1699–716. doi:10.1002/qre.2364.

Zhao, M. J., A. R. Driscoll, S. Sengupta, N. T. Stevens, R. D. Fricker, Jr, and W. H. Woodall. 2018b. The effect of temporal aggregation level in social network monitoring. *PloS One* 13 (12):e0209075. doi:10.1371/journal.pone.0209075.

Zhou, S., Y. Chen, and J. Shi. 2004. Statistical estimation and testing for variation root-cause identification of multistage manufacturing processes. *IEEE Transactions on Automation Science and Engineering* 1 (1):73–83. doi:10.1109/TASE.2004.829427.