Social Networks 39 (2014) 62–70

Contents lists available at ScienceDirect

Social Networks

journal homepage: www.elsevier.com/locate/socnet

Anomaly detection in online social networks

David Savage[a,*], Xiuzhen Zhang[a], Xinghuo Yu[a], Pauline Chou[a,b], Qingmai Wang[a]

[a] School of CS&IT, RMIT University, GPO Box 2476, Melbourne, Victoria 3001, Australia
[b] Australian Transaction Reports and Analysis Centre, PO Box 13173, Law Courts, Melbourne, Victoria 8010, Australia

ARTICLE INFO

Keywords:
Anomaly detection
Link mining
Link analysis
Social network analysis
Online social networks

ABSTRACT

Anomalies in online social networks can signify irregular, and often illegal behaviour. Detection of such anomalies has been used to identify malicious individuals, including spammers, sexual predators, and online fraudsters. In this paper we survey existing computational techniques for detecting anomalies in online social networks. We characterise anomalies as being either static or dynamic, and as being labelled or unlabelled, and survey methods for detecting these different types of anomalies. We suggest that the detection of anomalies in online social networks is composed of two sub-processes; the selection and calculation of network features, and the classification of observations from this feature space. In addition, this paper provides an overview of the types of problems that anomaly detection can address and identifies key areas for future research.

1. Introduction

Anomalies arise in online social networks as a consequence of particular individuals, or groups of individuals, making sudden changes in their patterns of interaction or interacting in a manner that markedly differs from their peers. The impacts of this anomalous behaviour can be observed in the resulting network structure. For example, fraudulent individuals in an online auction system may collaborate to boost their reputation. Because these individuals have a heightened level of interaction, they tend to form highly interconnected subregions within the network (Pandit et al., 2007). In order to detect this type of behaviour, the structure of a network can be examined and compared to an assumed or derived model of normal, non-collaborative interaction. Regions of the network whose structure differs from that expected under the normal model can then be classified as anomalies (also known as outliers, exceptions, abnormalities, etc.).

In recent times, the rise of online social networks and the digitisation of many forms of communication has meant that online social networks have become an important part of social network analysis (SNA). This includes research into the detection of anomalies in social networks, and numerous methods have now been developed. This development has occurred over a wide range of problem domains, with anomaly detection being applied to the detection of important and influential network participants (e.g. Shetty and Adibi, 2005; Malm and Bichler, 2011; Cheng and Dickinson, 2013), clandestine organisational structures (e.g. Shetty and Adibi, 2005; Krebs, 2002; Reid et al., 2005), and fraudulent and predatory activity (e.g. Phua et al., 2010; Fire et al., 2012; Chau et al., 2006; Akoglu et al., 2013; Pandit et al., 2007).

Since anomaly detection is coming to play an increasingly important role in SNA, the purpose of this paper is to survey existing techniques, and to outline the types of challenges that can be addressed. To the best of our knowledge this survey represents the first attempt to examine anomaly detection with a specific focus on social networks. The contributions of this paper are as follows

- provide an overview of existing challenges in a range of problem domains associated with online social networks that can be addressed using anomaly detection
- provide an overview of existing techniques for anomaly detection, and the manner in which these have been applied to social network analysis
- explore future challenges for online social networks, and the role that anomaly detection can play
- outline key areas where future research can improve the use of anomaly detection techniques in SNA

In drafting this review we did not set out to consider particular problem domains. Rather, we aimed to identify tools specifically designed for detection of anomalies, regardless of the particular social networks they were designed to analyse. However, as we conducted our survey we found that relevant work was

* Corresponding author. Tel.: +61 3 9925 2774.
E-mail addresses: david.savage@rmit.edu.au, david.savage@mac.com (D. Savage).

http://dx.doi.org/10.1016/j.socnet.2014.05.002
0378-8733/

predominantly published in the area of computer science, and consequently, many of the applications of anomaly detection that we encountered were focused on anomalies in online systems. Therefore, unless specifically stated otherwise, the term social network will be used throughout this paper to mean an online social network.

Within the social sciences literature, we found a number of papers focusing on the concept of network change (see for example McCulloh and Carley, 2011; Arney and Arney, 2013; Tambayong, 2014), which attempts to characterise the evolution of social networks. We see anomaly detection as being a subset of change detection, as anomaly detection could be used to identify change points where an evolving social network undergoes a rapid change, however a network that evolves in a consistent fashion over an extended period of time is unlikely to be deemed anomalous. We have therefore elected to limit the scope of our review to those methods that deal specifically with anomaly detection.

2. Related work

Previous reviews of anomaly detection have provided an overview of the general, non-network based problem, describing the use of various algorithms and the particular types of problems to which these algorithms are most suited (Hodge and Austin, 2004; Chandola et al., 2009; Markou and Singh, 2003, 2003). A workshop on the detection of network based anomalies was also held at ACM 2013 (Akoglu and Faloutsos, 2013). The most recent review of general anomaly detection (Chandola et al., 2009), expands on previous works to define six categories of anomaly detection techniques; classification (supervised learning), clustering, nearest neighbour, statistical, information theoretic, and spectral analysis.

As well as categorising anomaly detection techniques, previous reviews describe a number of challenges for anomaly detection, mainly associated with the problem of defining normal behaviour, particularly in the face of evolving systems, or systems where anomalies result from malicious activities (Chandola et al., 2009; Hodge and Austin, 2004). In particular, Chandola et al. (2009) note that the development of general solutions to anomaly detection remains a significant challenge and that novel methods are often developed to solve specific problems, accommodating the specific requirements of these problems and the specific representation of the underlying systems. As discussed in Section 7, this has also been the case for some methods focused on anomaly detection in social networks.

In addition to the major reviews of anomaly detection described above, other works have considered anomaly detection as part of methodological surveys for particular problem domains. For example, methods for performing anomaly detection have been discussed as part of more general reviews in areas of fraud detection (Bolton and Hand, 2002; Phua et al., 2010), network intrusion (Jyothsna et al., 2011; Gogoi et al., 2011; Patcha and Park, 2007), and the deployment of wireless sensor networks (Zhang et al., 2010; Janakiram et al., 2006). While significant overlap exists between the analysis of computer and sensor networks and social networks, there are also a number of differences that must be taken into account. In particular, social networks are typically composed of many inter-connected communities, which has important consequences for the distribution of node degree, and the transitivity of the network (Newman and Park, 2003). Moreover, anomaly detection in both sensor and computer networks is typically required to occur online in (soft) real-time, and while this constraint may also apply in some SNA scenarios, it is not typically required. In addition, anomaly detection in sensor networks generally requires algorithms that reduce network traffic and have a low computational complexity (Zhang et al., 2010).

3. Problem domains for the application of anomaly detection in social networks

Anomalies in social networks are often representative of illegal and unwanted behaviour. The recent explosion of social media and online social systems, means that many social networks have become key targets for malicious individuals attempting to illegally profit from, or otherwise cause harm to, the users of these systems.

Many users of online social systems such as Facebook, Google+, and Twitter are regularly subjected to a barrage of spam and otherwise offensive material (Shrivastava et al., 2008; Fire et al., 2012; Akoglu et al., 2010; Hassanzadeh et al., 2012). Moreover, the relative anonymity and the unsupervised nature of interaction in many online systems provides a means for sexual predators to engage with young, vulnerable individuals (Fire et al., 2012). Since the perpetrators of these behaviours often display patterns of interaction that are quite different from regular users, they can be identified through the application of anomaly detection techniques. For example, sexual predators often interact with a set of individuals who are otherwise unconnected, leading to the formation of star like structures (Fire et al., 2012). These types of structures can be identified by examining a range of network features (Akoglu et al., 2010; Shrivastava et al., 2008; Hassanzadeh et al., 2012), or through the use of trained classifiers (Fire et al., 2012).

Online retailers and online auctions have also become a key target for malicious individuals. By subverting the reputation systems of online auction systems, fraudsters are able to masquerade as honest users, fooling buyers into paying for expensive goods that are never delivered. This process is facilitated by the use of Sybil attacks (the use of multiple fake accounts) and through collaboration between fraudulent individuals to artificially boost reputation to a point where honest buyers are willing to participate in large transactions (Chau et al., 2006; Pandit et al., 2007). In many online stores, opinion spam, in the form of fake product reviews, is used in an attempt to distort consumers' perceptions of product quality and to influence buyer behaviour (Akoglu et al., 2013). Again, the malicious individuals who engage in these types of behaviour often form anomalous structures within the network, as their patterns of interaction can be quite different from regular users.

In addition to the social networks supported by dedicated online systems, mining of the social networks induced by mobile phone communications, financial transactions, etc. can also be used to identify illegal activities. Detection of anomalies in these types of networks have previously been used to identify organised criminal behaviour, including insurance fraud (˜Subelj et al., 2011), and terrorist activities (Reid et al., 2005; Krebs, 2002). Given the highly detrimental impact of these types of behaviour, anomaly detection in social networks can be seen as an extremely important component in the growing tool-box for performing social network analysis (SNA).

Outside of criminal or malicious behaviour, anomaly detection has also been used to detect important and influential individuals (Shetty and Adibi, 2005), individuals fulfilling particular roles within a community (Welser et al., 2011), levels of community participation (Bird et al., 2008), and unusual patterns in email traffic (Eberle and Holder, 2007).

4. Definitions

Anomalies are typically defined in terms of deviation from some expected behaviour. A recent review of general, non-network based anomaly detection defined anomalies as "patterns in data that do not conform to a well defined notion of normal behaviour" (Chandola et al., 2009). Another recent review defines anomalies as "an observation (or subset of observations) which appears to