available at www.sciencedirect.com

ScienceDirect

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

Digital Investigation

Using Author Topic to detect insider threats from email traffic☆

James S. Okolica, Gilbert L. Peterson, Robert F. Mills*

Air Force Institute of Technology, AFIT/ENG, Building 641, Room 220, 2950 Hobson Way, Wright Patterson AFB, OH 45433-7765, USA

ARTICLE INFO

Article history:
Received 13 February 2007
Revised 8 October 2007
Accepted 31 October 2007

Keywords:
Author Topic (AT)
Insider threat
Datamining
Social networks
Large data set

ABSTRACT

One means of preventing insider theft is by stopping potential insiders from becoming actual thieves. This article discusses an approach to assist managers in identifying potential insider threats. By using the Author Topic [Rosen-Zvi Michal, Griffiths Thomas, Steyvers Mark, Smyth Padhraic. The author-topic model for authors and documents. In: Proceedings of the 20th conference on uncertainty in artificial intelligence; 2004. p. 487–94.] clustering algorithm, we discern employees' interests from their daily emails. These interests then provide a means to create an implicit and an explicit social network graph. This approach locates potential insiders by finding individuals who either (1) feel alienated from the organization (a key warning sign of a possible disgruntled worker) or (2) have a hidden interest in a sensitive (e.g. proprietary or classified) topic. In both cases, this is revealed when someone demonstrates an interest in a topic but does not share that interest with anyone in the organization. By applying this technique to the Enron email corpus, we produce coherent, sensible topics and reveal Sherron Watkins, the famous Enron whistleblower, as a potential insider threat from the viewpoint of the individuals behind the Enron scandal.

Published by Elsevier Ltd.

1. Introduction

The best time to address the insider threat is before it occurs. Mostly, individuals do not enter an organization with the intent to do harm (Shaw et al.). Instead, something changes while they work at the organization. An automated way is needed to detect when individuals begin feeling alienated from an organization. Managers can then use the results to allocate scarce resources to determine if active intervention is required.

One of the best indicators of a person's interests in today's organizations is their email traffic. Through datamining, topics of interest are extracted and people categorized by those topics they are most interested in. The topics of interest are then used to detect warning signs such as alienation from the organization as well as inappropriate interest in sensitive or classified topics. Especially likely are people who have shown an interest in a sensitive topic but never communicated that interest with anyone within the organization. These people either have a secret interest in the topic or generally feel alienated from the organization and so communicate their interest only outside of it.

In this paper, Author Topic (Rosen-Zvi et al., 2004) is used on the Enron email corpus to test its applicability on generating insider threat investigative leads. The results produce 48 clear topics and reveal Sherron Watkins, Enron's well-known whistleblower, as one of three individuals who has both a clandestine interest in the off-book partnerships' topic and a feeling of alienation from Enron.

This paper is organized as follows. In Section 2, the nature of the insider threat is discussed as well as a short overview of

☆ The views expressed in this article are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

* *Corresponding author*. Tel.: +1 937 255 3636; fax: +1 937 656 7061.
E-mail addresses: james.okolica@afit.edu (J.S. Okolica), gilbert.peterson@afit.edu (G.L. Peterson), robert.mills@afit.edu (R.F. Mills).

1742-2876/$ – see front matter Published by Elsevier Ltd.
doi:10.1016/j.diin.2007.10.002

the Author Topic Model. This is followed by the methodology and results of the experiment in Sections 3 and 4. The article concludes in Section 5 with conclusions and suggestions for future investigation.

2. Background

Espionage is the practice of spying or using spies to gather information about a competitor. An insider is someone in an organization who has a legitimate right to (some of) the organization's information but uses it for an illegitimate purpose. Herbig and Wiskoff (2002) published a report detailing all of the publicly known cases of espionage against the United States between 1947 and 2001. They found that, although in most cases the individual responsible was an insider, they did not enter the organization intending to commit espionage. In many cases, an organizational change, such as restructuring, or a personal crisis, such as the end of a relationship or a severe financial problem, contributed to the individual choosing to commit espionage. Managers can no longer spend sufficient time to detect these warning signs. Instead, they must rely on some tool to direct their attention on a few individuals.

In today's Information Age, one of the best sources of personal information available at work is an individual's email activity. Probabilistic clustering takes a collection of email and lumps it into coherent ''topics''. In addition it discovers the principle topics an individual is interested in. Prior to betraying an organization, individuals generally feel alienated from it. Whether such an alienation is initiated by the individual or the organization is irrelevant. What is required is for this alienation to exist (Shaw et al.). To test for this, social networks can be used. If individuals have an interest in a topic but do not send or receive email about that topic with others within the organization, they can be considered to have a clandestine interest in that topic. Depending on the number and type of such topics, these clandestine interests can be considered an indicator that an individual feels alienated from the organization and has insider threat potential.

2.1. *Enron as the data source*

As part of their investigation into Enron, the Federal Energy Regulatory Commission (FERC) seized Enron's email and made a portion of it publicly available (FindLaw). While it only includes the email folders of 151 employees, it still contains over 250,000 email messages. Furthermore, due to the number of individuals the emails were sent to, the resulting corpus has sufficient data on over 34,000 Enron employees to make probabilistic clustering effective. Prior to the public disclosure of Enron's questionable accounting, Ms. Watkins, a vice president in Enron's corporate accounting division, sent a letter to Ken Lay, Enron's chairman, detailing the dubious accounting practices and their likely impact on Enron's future. Her activities were considered insider theft by her boss, Andy Fastow, Enron's Chief Financial Officer, who demanded that she be fired immediately (McLean and Elkind, 2003). This is an example of the importance of perspective on distinguishing whistleblowers from insider threats.

2.2. *Current social network and insider threat research*

A social network is generally constructed based on perceived common interests. For example, ReferralWeb (Kautz et al., 1997) uses the co-occurrence of names in close proximity in documents publicly available on the WWW (e.g. journal articles, newsgroups, chatrooms, etc.) to denote a close relationship. Adamic and Adar have done similar research by using mailing lists and the homepages of students at Stanford and MIT. Since, when people create homepages, they link to their friends' homepages (and ask their friends to link to theirs), she postulated that using homepages would result in an appropriate social network. She also used the text present on the web pages to further predict relationships (i.e. common interests) between people. While she was able to show that the text provided strong indications of friendships between people, it is unclear if this would generalize beyond the rather closed community of a university. Culotta et al. (2004) approached this problem differently but from a more general population. They began by extracting names from email messages. Then they used the WWW to find the person's ''web presence'' (generally his or her homepage) and used that to describe the person and to find friends of that person. After the network was created, they used graph partitioning algorithms to find highly connected components. While their data set was small (53 email correspondents), their results were promising. However, one of the biggest drawbacks was a lack of web presence for many of the correspondents (31 of 53).

Since September 11, 2001 there has been increased research in uncovering potential threats through the use of social networks. However, despite several organization, such as Rand Corporation and Mitre, making proposals for using social networks to detect insider threats (RAND, 2004 Mitre, 2004), little public research has been done. Symonenko et al. (2004) have generated social networks of intelligence analysts and then used semantic analysis to detect when individuals are showing interests in areas outside of their group. While the results have been promising, the technique requires a large number of interviews with experts to provide the semantic analysis. In addition, this expert knowledge is then only applicable to the specific group and needs to be repeated each time the application is moved to another organization. Yee et al. (2005) have also performed some initial research into generating social networks from email headers for later analysis by social network analysts.

Other insider threat research has examined the problem from a formal security model and management perspective. By mapping employee performance metrics to insider threat traits a manager can determine an individual's potential for insider threat behavior (Butts et al., 2006). Additionally, Bayes network models of behavior have also been shown to be potentially helpful in detecting insider threats (AlGhamdi et al., 2005).

2.3. *Author Topic*

This work examines the application of Author Topic (Rosen-Zvi et al., 2004), a probabilistic clustering technique, on the task of revealing insiders from the Enron email corpus. Author Topic is a probabilistic clustering technique that takes the