A Social Network Analysis Tool for Uncovering Cybersecurity Threats

Vasileios Vlachos^{1,2}

¹Computer Technology Institute & Press — "DIOPHANTUS",
Nikou Kazantzaki str.,
University of Patras Campus, GR-256504,
Rio, Patras, Greece

²University of Thessaly, General Department,
Larissa – Trikala Ring Road, GR-41110
Larissa, Greece
vsvlachos@uth.gr

Yannis C. Stamatou^{1,2}

¹Computer Technology Institute & Press — "DIOPHANTUS",
Nikou Kazantzaki str.,
University of Patras Campus, GR-256504,
Rio, Patras, Greece

²Business Administration Department,
University of Patras, GR-26504,
Rio, Patras, Greece
stamatou@ceid.upatras.gr

Pantelis Tzamalīs^{1,2}

¹Computer Technology Institute & Press — "DIOPHANTUS",
Nikou Kazantzaki str.,
University of Patras Campus, GR-256504,
Rio, Patras, Greece

²Computer Engineering and Informatics Department
University of Patras GR-26504
Rio, Patras, Greece
tzamalīs@ceid.upatras.gr

Sotiris Nikolettseas^{1,2}

¹Computer Technology Institute & Press — "DIOPHANTUS",
Nikou Kazantzaki str.,
University of Patras Campus, GR-256504,
Rio, Patras, Greece

²Computer Engineering and Informatics Department
University of Patras GR-26504
Rio, Patras, Greece
nikolett@cti.gr

Kyriaki Chantzi

University of Thessaly,
Larissa – Trikala Ring Road, GR-41110
Larissa, Greece
kikichantzi@gmail.com

The paper presents the design and implementation of the Cybersecurity Social Network Analyser (CSNA) platform. The CSNA is a software system for real time automated analysis of cybersecurity-related information available on Social Networks, with a focus on Twitter. Its aim is to uncover cybersecurity breach related discussions and posts, identify emerging cybersecurity incidents and provide, in a timely fashion, early warning alerts. The CSNA platform presents the results from the analysis of the gathered information in a variety of graphical formats for easy visualisation of potential threats within user-specified time windows of interest. We complement our description with a demonstration of the real time information provided by the platform during its operation on the Twitter social network.

Cybercrime, Social Network Analysis, Computer Security

1. INTRODUCTION

Today, the Internet archives a vast amount of cybersecurity related data and information from a variety of sources. On the other hand, despite this exponentially increasing amount of information on cybercrime, cyberthreats appear to escalate, as recent studies and surveys demonstrate. More importantly, the most serious vulnerabilities, such as the ones that are unknown to the developer of a product, i.e. *0-day exploits*, are traded in various underground black hat hacker communities [3]. The fast growing number of existing software vulnerabilities and the corresponding exploits support a complex ecosystem of buyers and sellers of toolkits for compromising the security of widely used software applications, operating systems and hardware devices [15]. The transactions between the discoverers and the buyers of the vulnerabilities vary from the totally illegal black markets to the borderline (legal or semi-legal) grey markets. The buyers' profile, at least for the legal transactions, is usually that of law enforcement agencies, secret services and national security agencies. As expected, transactions information is rarely made public. Another approach which is considered both legal and ethical is to offer bug bounties to researchers for their work on discovering bugs and problems that could lead to security breaches [9].

As stated above, the amount of publicly available cybersecurity information is massive and exponentially growing. However, despite this wealth of information, the number of cybersecurity incidents is increasing at a fast pace. One plausible explanation for this, seemingly, contradictory fact is that a vast information volume does not automatically, by itself, constitute a solid, comprehensive, cybersecurity defence knowledge basis. Stated differently, it appears that this massive information openly available on the Internet exists in a, rather, fragmented and "noisy" form. It is fragmented because it comes from a tremendous number of heterogeneous sources having a non-standardised format with no semantic information or metadata that would aid processing and knowledge extraction. Also, the available information is noisy because the useful part is usually buried in lengthy texts that may address numerous other things, also in a wordy fashion, perhaps unrelated to cybersecurity or reflecting the writer's views in a form of free style discussion unsuitable for automated processing. However, these unstructured and full of redundancies information sources, are still our primary means of obtaining data, knowledge and alerts about cybercrime and its financial aspects which are also non-negligible [2] along with the technical ones. What is more interesting, in this context, are the direct financial consequences of software bugs to

the stock prices of a firm [24] or to the society in general [28].

In view of this complex cybersecurity information landscape, our goal is to describe the design, implementation and functionalities of an online, real time system, that utilises Social Networks to identify, archive and analyse cybersecurity related discussions, posts and information exchanges. Through the information archiving and processing capabilities of the software, we can raise early warning alerts, develop information and knowledge bases in a structured and standardised, automated processing friendly, format.

The rest of the paper is organised as follows. Section 2 presents recent related work. Section 3 provides an overview of the CSNA architecture. Section 4 focuses on the implementation aspects of the CSNA as well as examples of its analysis and data presentation capabilities. Finally, section 5 discusses our conclusions from designing and building the CSNA platform and provides directions for future work.

2. RELATED WORK

Online data from web services, social networks, and Open Source Intelligence (OSINT) feeds can provide important input in virtually all scientific disciplines and research domains [11]. This information can play a major role in extracting valuable insights in the latest developments in many human activities and physical processes. For instance, one of the most well studied frameworks of forecasting based on social networks is related to health epidemics. Ritterman et al. [20] were among the first to use the Twitter to predict the evolution of the swine flu pandemic. In particular, they analysed fifty million tweets using Artificial Intelligence (AI) methods and Support Vector Machines (SVM) for Machine Learning (ML). Their encouraging results were also validated by other researchers. Successful monitoring algorithms for the influenza pandemic have been also documented in other research works [23]. Santillana et al. [21] further improved the accuracy of traditional approaches by combining social media and existing data sources in a single searching and processing framework. Nagar et al. [16] exploited geocoded Twitter data in order to improve the efficiency of the prediction of the New York city 2012 influenza epidemic. Other researchers [18] have used Wikipedia access logs to monitor and forecast six diseases. Their methods yielded promising results with high forecasting value. As our work also uses some basic epidemiological metrics in order to monitor the spread of malware and other cybersecurity threats, we also plan to

validate some of the aforementioned approaches using the CSNA framework.

Social networks analysis in support of forecasts in the stocks markets has attracted the attention of several researchers. Rao et al. [19] analysed stock market activity using Twitter *sentiment analysis*. The authors collected and processed four million tweets, identifying a strong correlation between stock prices and twitter sentiments. Mao et al. [14] demonstrated a correlation of S&P with Twitter data which, according to the authors, renders the predictions in the stock market domain feasible. Other researchers [17], focused on the information available in the StockTwits platform using sentiment analysis and posts volume indicators. The authors also highlight the volatility and complexity of stock markets, which poses limitations to such approaches. Interestingly, in previous works, researchers have tried to explore [13] the efficiency of forecasting models and techniques from the financial world to the ICT security landscape. A conceptually similar approach [22] evaluates the risks stemming from remotely launched cyberattacks.

With respect to political studies, Tumasjan et al. [25] described one of the first efforts to use Twitter to estimate the German Federal election results. Their research showed that the distribution of tweets for the main parties was similar to the corresponding distribution of votes in the election. Other researchers tried to use Social Networks to predict the election outcome in numerous other cases. Kegan et al. [12] worked on the Pakistani and Indian elections of 2013 and 2014 respectively using a diffusion estimation model. Ahmed and Skoric [1] emphasised on the nature of the Twitter campaign characteristics in these elections as a decisive factor for the final result. Choy et al. [7] presented a platform based on socio-demographics and sentiment analysis for modelling the Singapore elections. The authors, also, tried to predict the 2012 US presidential elections in [8] with partial, however, success. Bovet et al. [5] used a combination of statistical physics of complex systems and artificial intelligence to emphasise that social network analysis can provide more timely and less expensive estimates for the election results, as accurate as national polls. Likewise, Burnap et al. [6] used sentiment analysis of 13.899.073 tweets to predict the 2015 General Elections results in UK. A significant contribution of their work is that they identified the importance of the geolocation information of the tweets for a successful election results prediction. On the other hand, using Twitter to predict elections has a number of limitations and shortcomings which are summarised in Gayo-Avello's work [10].

Many researchers have tried to predict criminal activity based on social media, with encouraging results. Wang et al. [27] focused on the "hit-and-run" based crimes. They used automatic sentiment analysis and natural language processing for related Twitter posts. Bendler et al. [4] used geographically weighted regression analysis to improve the accuracy of delinquency prediction. A very interesting combination of techniques for improving crime predictions is presented in [26]. In this work, the author gathers data from Foursquare, Twitter and taxi trips to provide evidence for the importance and necessity of deploying hybrid approaches based on temporal and static analysis.

Although all crime types share some common patterns, our research focuses, exclusively, on cybercrime activity, as in [29]. In that work, the authors concentrate on analysing, understanding and predicting DDoS attacks with satisfactory accuracy. In contrast, in our work, we focus on evolving real-time cybersecurity threats using a dynamic intelligent threat detection approach.

3. THE SAINT SOCIAL NETWORK ANALYSER

The primary target of the Cybersecurity Social Network Analyser platform (CSNA) is Twitter. Twitter is a very useful and massive source for up-to-date and real time publicly available information generated continuously by users of any background and educational level, on virtually any subject. In our case, the focus is on cybersecurity related information.

However, simply obtaining raw data coming from the social medium's API (e.g. Twitter's API), does not provide much added value. The proposed software system adds value to the obtained data by easily processing, analysing, and preparing results to be archived in standard data formats and visualised in intuitive graphical representations. Using Python, which has excellent string processing capabilities and great ease-of-use, it is possible to facilitate the development of any conceivable data processing procedure on any social medium's data as long as one builds a suitable client application to interact with the target social networks's API.

The general architecture of the CSNA platform is shown in Figure 1. The data flow and data handling procedures of the CSNA are divided into six steps: (1) Setting up a search query: Starting with this step, a query containing search keywords, i.e. hashtags, is created and inserted into the crawler script as a parameter so that a search process is initiated for tweets that include the hashtags. (2) User Authentication (OAuth): At this point, an

Open Authorisation request is created by the App in the Twitter account where we retrieved the API keys, the API secrets, the Access Tokens and the Access Token Secrets as dictated by the OAuth policies and protocols. In this way, the Twitter Crawler can be granted access to the Twitter platform and especially to the Twitter Streaming API, which will be deployed for searching and gathering the tweets that are created in real time by Twitter users. (3) Data collection (Tweets): The Twitter Crawler searches and downloads the publicly available tweets that are related to the search hashtags queries into our server. The raw data items are stored into the NoSQL schema database, implemented in the MongoDB DBMS. (4) Data cleaning and pre-processing: Various data pre-processing techniques and methods are applied during this step looking for invalid data transformations and enforcing the desired, machine readable, formats for the subsequent analysis by the CSNA and its constituent modules. (5) Modelling and analysis: Natural Language Processing, Regular Expression Analysis, and other Data Analysis techniques (e.g. Statistical models), regarding the tweets that are collected to our server, take place during this step of the analysis. (6) Result presentation (Visualisation): Finally, data visualisation methods are used during this step for creating, in real time, bar graphs, time series line charts, and histograms for quick and easy presentation of data and analysis results to data analysts.

Figure 1: The Cybersecurity Social Network Analyser tool (CSNA) architecture.

The CSNA is one of the two primary subsystems of the SAINT platform. The Open Source Intelligence (OSINT) module is another complementary framework for monitoring ongoing threats. It is currently under development and therefore it will be discussed in more detail in the Future Work section of this

paper. The OSINT module collects related information from a number of publicly available web sites, services and datasets and stores them for further analysis in the SAINT database. One important aspect of this two-modular approach, based on the CSNA and the OSINT subsystems is the possibility to correlate and cross-validate both findings. During its last 12-month operation, the SAINT managed to identify imminent threats based on the correlation and concurrency of spikes in the received information on both subsystems. In Section 4.3 we will provide some relevant examples.

4. DEVELOPMENT & RESULTS

4.1. Details on the Implementation and Operation of the System

Python has excellent string processing capabilities, great ease-of-use, and it is possible to facilitate the development of a variety of data processing procedures on any social medium's content as long as a suitable client application is built to interact with the target social network's API. In our case, one of the CSNA's main development applications is **Tweepy**, which is an open source library that allows Python to communicate with the Twitter platform using its APIs. Additionally, several efficient Python libraries that are relevant to data analysis and NLP are used on the back-end development, while HTML and JavaScript (including visualisation libraries, such as **Highcharts/Highstocks**) are the main programming languages used for the implementation of the user interface and the front-end part.

CSNA is composed of the following data analysis modules: (1) Python's **Pandas**, **Numpy**, and **SciPy** for data processing. (2) Python's **NLTK** and **RE** (Regular Expressions) for Natural Language Processing. (3) JavaScript **Highcharts/Highstocks** libraries for data visualisation. (4) **Bootstrap**, an open source toolkit for developing with HTML, CSS, and JS. (5) Python's **Wordcloud** module for making the word clouds data visualisation.

Our experiments demonstrated that interesting and relevant results can be obtained by the construction of the appropriate search queries. Two examples follow:

twitterQuery1: *search tweets including hashtags related to Bug Bounty programs and security contests: '#BugBounty' OR '#vulnerability #market' OR '#security #contest' OR '#bugsbounty' OR '#securitycontest' OR '#securitycontest #vulnerability' OR '#vulnerability #bugsbounty' OR '#bugbounty #securitycontest' OR '#securitycontest #bugsbounty'*

twitterQuery2: search tweets including hashtags related to cyberattacks: '#malware' OR '#apt' OR '#ransomware' OR '#spyware' OR '#xss' OR '#lfi' OR '#rfi' OR '#websecurity' OR '#xee' OR '#webappsec' OR '#botnet' OR '#DoS' OR '#DDoS' OR '#botnet' OR '#bot' OR '#botmaster' OR '#botnets' OR '#phishing' OR '#phish' OR '#pharming' OR '#spam' OR '#idtheft' OR '#trojan'

To converge to the most targeted and information-promising Twitter search queries, we experimented with numerous search keywords, hashtags, terms and word combinations to see which ones give the best results on cybersecurity related information. Our experiments demonstrated that interesting results can be obtained by the construction of appropriate search queries divided into two big classes. The first one contains keywords related to cybersecurity markets and contests keyword and the second one contains keywords relevant to cyberthreat attacks are included. Thus, two large, concurrent, instances of sub-crawlers are created, corresponding to each keywords class.

Each tweet instance that is obtained by the CSNA is collected into the **MongoDB** database, stored as a JSON-style format. The use of MongoDB as a data repository, renders the process of the data manipulation and analysis more convenient. We use Python's **pymongo** library to connect each of our Twitter sub-crawlers to the MongoDB database system. Finally, each sub-crawler writes new records (tweets) to the corresponding Twitter collection (based on cybersecurity markets or cyberattacks related data) into the general database, that is named '**twitter**'.

4.2. Entity Analysis

We performed frequency analysis using the data collected by the Twitter sub-crawlers into the two **twitter** database collections. Slicing and dicing the data entities allows us to produce interesting statistics. The results of such processing are the tweet entities analysis on the contents of the collections.

We were interested in extracting the 10 **most frequently appearing user mentions, hashtags, and word terms** over the last 7 days that were written in the tweets collected by the relevant sub-crawler. We searched over all the tweets that are stored in the categorised collections (markets-based, threats-based). The results from this analysis can be seen in the Figures 2, 3, and 4 respectively.

We employ various aspects of text analytics and NLP methods, such as text pre-processing, encoding, normalisation and complex tokenisation, for dealing

(a)

(b)

Figure 2: [a] Top 10 User Mentions over the last 7 days in Cybersecurity Markets and [b] Cyberattacks collections.

with phrases, symbols, etc., as well as statistical analysis approaches for the tweets and various regular expressions methods applied to the raw text. Tokenisation is, also, a vital step of the process in which a stream of text is broken down into individual units called tokens. In their simplest form, these units are words, but we also work on more complex tokenisations that deal with phrases, symbols etc. Another useful pre-processing operation is the **stop word removal**. Stop words are words that do not bear any content when taken in isolation, such as articles, propositions, adverbs, and so on. Various **regular expression (regex)** methods are applied to avoid duplicate counts, recognise emoticons, HTML tags, URLs, numbers and other characters as well as targeted words. To complete the stop word list, we finally included the 'rt', 'via', and '...' tokens (a single-character Unicode symbol for horizontal ellipses).

Applying the NLP techniques that were described above to the **threats-based** and **markets-based** collections of tweets, a categorical analysis step is, subsequently, performed and two bar chart plots are created for each kind of entity analysis, in real-time.

(a)

(a)

(b)

(b)

Figure 3: [a] Top 10 frequently used Hashtags over the last 7 days in Cybersecurity Markets and [b] Cyberattacks collections.

Figure 4: [a] Top 10 frequently used Terms over the last 7 days in Cybersecurity Markets and [b] Cyberattacks collections.

All the figures given in this paper are screen-shots from the outputs generated from our online CSNA platform.

Figure 2 shows the most frequently mentioned users in the collected raw tweets that are stored into the two kinds of collections. From Figure 3, we can see that some of the most common hashtags that were mentioned in the markets-based tweets (up bar-plot) were: security, bugbounty, and Hackerone. This agrees with our intuition since Hackerone is one of the most well-known organisation specializing in organising bug-bounties on behalf of companies and, in general, discovering the most serious security issues before they can be exploited by criminals. Additionally, in the same Figure, in the down bar-plot, malicious software hashtag mentions such as malware, ransomware, phishing and others are among the top mentioned threats. We can observe that some of the top attack methods that cybercriminals use at the time, such as malware and

ransomware viruses, stir interest and ignite users' posts about IoT security, cyberattacks, hacking, artificial intelligence, cybercrime and cybersecurity in general. From the last category of bar-plots exports (Figure 4), there are a few terms on the left-hand side of the plots with very high frequency of appearance. For example, the most frequently appearing term is twice more frequent than the ones after index position 10. As we move towards the right-hand side of the plots, the curve becomes less steep, implying that the terms on the right-hand portion of the plots have similar frequency of appearance.

4.3. Time Series Descriptive Analysis

In this section, we discuss another aspect of our analysis, that of studying the distribution of tweets over time. A time series is a sequence of data points that consists of successive observations over a given time interval. As Twitter provides a 'created_at' field with a precise timestamp/datetime of the tweet, we can reorder tweets into "time buckets" so that we

can examine how users react to the appearance of events in real time, i.e. while an event is on its onset or it is still ongoing.

Time series descriptive analysis can provide information on how users may react to trending events or a data related to the **online reputation management**. The dynamic and massive nature of Twitter appears to be a perfect feature that enables to identify the public's reactions to events and breaking news, with a focus on cybersecurity. In this case, our goal is to observe how users react to various cyberattacks (**threats-based collection**) and bug bounty related news (**markets-based collection**) as time evolves and to derive valuable information about potential cybersecurity risks and related data breach costs. We exploited **Panda's** capabilities for handling time series, and we used **Highcharts** for a visual interpretation of the series.

The processing by the CSNA uses the **Highcharts** JavaScript library to create informative visualisations of the time series data. **Highcharts** is, in general, more verbose than other data visualisation libraries, but its outputs are not complicated and they are easy to interpret. The charts are exported into the online CSNA platform, including the corresponding function and query that were used for the Time Series Descriptive Analysis of each category. The generated CSV and JSON files are also available for download. Figure 5 shows a Time Series Descriptive analysis outcome that come of the corresponding cybersecurity markets collection.

Figure 5: #BugBounty Time Series Descriptive Analysis.

Currently the SAINT platform intends to act as a comprehensive analytical framework to assist researches, scientist and cybersecurity experts to identify imminent threats. The collection, processing and analysis of the data is fully automated and can be adjusted and parameterised to target additional types of cyberattacks. Although much effort is put into making the detection of serious cybersecurity incidents also automated as we will discuss later, the final decision relies on the platform operator based on the visual inspection of the data and the graph charts. Our effort is to make the system self-intelligent by using outlier detection techniques to identify suspicious events. To make that decision less complicated the SAINT system utilises a

comparative framework. The operator, before taking any action, can compare the data from the CSNA and the OSINT module. Generally, as it is observed in Figure 4, when a spike related to an issue is shown in an OSINT indicator's Time Series descriptive results, after a few hours or the next day, a similar spike is also shown to the relevant statistical result of the CSNA, indicating that Twitter users are involved with that issue too. Thus, the two subsystems can cross-validate each other. Empirical results suggest that when a trend is observable in both subsystems or in other words in social networks and the open source intelligence feeds, there is a very high probability of an ongoing or imminent incident. Next steps are the automatization of this procedure as well.

(a)

(b)

Figure 6: Phishing attacks' descriptive results in [a] SNA and [b] OSINT Crawler.

4.4. Word Clouds

A **Word Cloud** is a visual representation of text data, typically used to display keyword metadata (tags) on websites or to visualise free form text. Tags are usually single words depicted in the word cloud using different font sizes and/or colors depending on their importance.

The word cloud format is useful for quickly identifying the most important terms in the word cloud picture comparing, at the same time, the terms according to their importance. Moreover, significant textual data items can be highlighted. Word clouds are also widely used for analysing data from social networks.

There are three main reasons behind our decision to use word clouds to present textual data:

1. They convey information with simplicity and clarity.
2. They are powerful communication tools that are easy to understand and interpret.
3. They are more visually engaging than a simple table of data.

In our case, we use the word cloud visualisation method to present the most important terms (tags) for each **twitter** database collection over the past 30 days. The data analysis process is similar to the one described in Section 4.2. The Categorical Analysis part of the CSNA analysis framework is used to identify the 10 most frequently used, by Twitter users, terms. This data visualisation method is useful for understanding, in depth, the trends in a topic (in our case, cybersecurity markets and cyberattacks) as it facilitates the application of human intelligence methods to help stakeholders make more informed decisions.

Two Python scripts were constructed for the CSNA, each one related to the corresponding **twitter** collection. These scripts include all the processes and functions that were presented in Section 4.2. Initially, we perform the textual analysis of the tweets that were archived into each collection during the past 30 days. Thereafter a collection of the twitter database is created in .png image format. An example is presented in Figure 7 for the results of **twitterQuery1**, which was given in Section 4).

information is stored in a database using a standardised electronic format for further processing and visualisation.

We focused on Twitter, since Social Networks can provide significant insights in various trends related to ongoing cybersecurity activities and cyber security breach incidents. For instance, the number and frequency of cybersecurity relevant discussions in social media can be highly indicative for the current malware epidemics and other types of cybercriminal activities. The identification of trends is achieved by the platform through a variety of approaches. Hashtags act as a filter to select tweets of interest. The NLP methods that the software deploys are particularly effective in extracting relevant text from the main body of the raw tweet data. Twitter metadata information items are, also, particularly effective in correlating structured and unstructured tweet information.

The visualisation of the information in a time-series form provides valuable timing data related to the development and onset of these trends. The graphical display of the rates of appearance of specific topics in tweets highlights their significance at the time they were posted. Empirical results correlate the social network analysis outcomes with specific incidents such as ongoing phishing campaigns and eminent ransomware attacks. Our findings are consistent with the work of other researchers regarding the deployment of social networks in monitoring and forecasting various phenomena in, virtually, any aspect of human activity. To increase the accuracy of our detection mechanisms we also deploy some basic mathematical epidemiological models such as the **epidemic curve** to identify points of time where an escalation is taking place.

Our future work, partially implemented, focuses on the improvement of the CSNA platform using more intelligent methods for locating and extracting information. This is work in progress and aims to create a dynamical semi-autonomous system that will be in the position to automatically refine and improve its search methods and goals. Using the new version of the SAINT tools, we expect to be able to identify the correlation between the data diffused in Social Networks and Open Source Intelligence (OSINT) feeds. We are already in the process of extracting data from numerous information sources. In particular, we are able to collect public data from more than 10 different datasets to cover threats related with malware attacks, ransomware operations, botnets activity, spamming and phishing campaigns. The datasets are updated periodically on an hourly or on a daily basis and therefore we are able to process and correlate them with CSNA's

Figure 7: Cybersecurity Markets Word Cloud.

5. CONCLUSIONS AND DIRECTIONS FOR FUTURE WORK

Our work presented the design, implementation, and operation of a Cybersecurity Social Network Analyser (CSNA). We discussed the applicability of the CSNA in discovering, extracting, and analysing information contained in tweets, based on appropriate keywords or keyword combinations. The obtained

output and identify common patterns. In the next version, a unified processing and analysis framework is expected to provide automatic detection of suspicious patterns in both subsystems. Finally, an experimental Deep Web Crawler DWC is also under heavy development. The goal is to make the DWC to follow leads from the CSNA and the OSINT subsystem. As cyberattacks become more sophisticated, innovative and multidisciplinary approaches are worth exploring.

ACKNOWLEDGEMENT

This research has received funding from the European Union's Horizon 2020 research and innovation program through the SAINT research project under grant agreement No 740829.

REFERENCES

- [1] Ahmed, S. and M. Skoric (2015). *Twitter and 2013 Pakistan general election: The case of David 2.0 against goliaths*, pp. 139–161. Springer International Publishing.
- [2] Anderson, R., C. Barton, R. Böhme, R. Clayton, M. J. G. van Eeten, M. Levi, T. Moore, and S. Savage (2013). *Measuring the Cost of Cybercrime*, pp. 265–300. Berlin, Heidelberg: Springer Berlin Heidelberg.
- [3] Armin, J., P. Foti, and M. Cremonini (2015). 0-day vulnerabilities and cybercrime. In *Proceedings of the 2015 10th International Conference on Availability, Reliability and Security, ARES '15*, Washington, DC, USA, pp. 711–718. IEEE Computer Society.
- [4] Bendler, J., A. Ratku, and D. Neumann (2014). Crime mapping through geo-spatial social media activity. See [4].
- [5] Bovet, A., F. Morone, and H. A. Makse (2016). Predicting election trends with twitter: Hillary clinton versus donald trump. *CoRR*.
- [6] Burnap, P., R. Gibson, L. Sloan, R. Southern, and M. L. Williams (2015). 140 characters to victory?: Using twitter to predict the UK 2015 general election. *CoRR abs/1505.01511*.
- [7] Choy, M., M. L. F. Cheong, N. L. Ma, and K. P. Shung (2011). A sentiment analysis of singapore presidential election 2011 using twitter data with census correction. *CoRR*.
- [8] Choy, M., M. L. F. Cheong, N. L. Ma, and K. P. Shung (2012). Us presidential election 2012 prediction using census corrected twitter model. *CoRR*.
- [9] Egelman, S., C. Herley, and P. C. van Oorschot (2013). Markets for zero-day exploits: Ethics and implications. In *Proceedings of the 2013 New Security Paradigms Workshop, NSPW '13*, New York, NY, USA, pp. 41–46. ACM.
- [10] Gayo-Avello, D. (2012). "i wanted to predict elections with twitter and all i got was this lousy paper" - a balanced survey on election prediction using twitter data. *CoRR abs/1204.6441*.
- [11] Jungherr, A. and P. Jürgens (2013). Forecasting the pulse: How deviations from regular patterns in online data can identify offline phenomena. *Internet Research* 23(5), 589–607.
- [12] Kagan, V., A. Stevens, and V. S. Subrahmanian (2015). Using twitter sentiment to forecast the 2013 pakistani election and the 2014 indian election. *IEEE Intelligent Systems* 30(1), 2–5.
- [13] Kollias, S. K., V. Vlachos, A. Papanikolaou, and V. Assimakopoulos (2011). Adapting econometric models, technical analysis and correlation data to computer security data. In *2011 First SysSec Workshop*, pp. 59–62.
- [14] Mao, Y., W. Wei, B. Wang, and B. Liu (2012). Correlating s&p 500 stocks with twitter data. In *Proceedings of the First ACM International Workshop on Hot Topics on Interdisciplinary Social Networks Research, HotSocial '12*, New York, NY, USA, pp. 69–72. ACM.
- [15] Miller, C. (2007). The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales. In *In Sixth Workshop on the Economics of Information Security*.
- [16] Nagar, R., Q. Yuan, C. C. Freifeld, M. Santillana, A. Nojima, R. Chunara, and J. S. Brownstein (2014, Oct). A case study of the new york city 2012-2013 influenza season with daily geocoded twitter data from temporal and spatiotemporal perspectives. *J Med Internet Res* 16(10), e236.
- [17] Oliveira, N., P. Cortez, and N. Areal (2013). On the predictability of stock market behavior using stocktwits sentiment and posting volume. In *EPIA*.
- [18] Priedhorsky, R., D. Osthus, A. R. Daughton, K. R. Moran, N. Generous, G. Fairchild, A. Deshpande, and S. Y. Del Valle (2017). Measuring global disease with wikipedia: Success, failure, and a research agenda. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing, CSCW '17*, New York, NY, USA, pp. 1812–1834. ACM.
- [19] Rao, T. and S. Srivastava (2012). Analyzing stock market movements using twitter sentiment analysis. In *Proceedings of the 2012 International*

Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012), ASONAM '12, Washington, DC, USA, pp. 119–123. IEEE Computer Society.

- [20] Ritterman, J., M. Osborne, and E. Klein (2009). Using prediction markets and twitter to predict a swine flu pandemic. In *In Proceedings of the 1st International Workshop on Mining Social*.
- [21] Santillana, M., A. T. Nguyen, M. Dredze, M. J. Paul, E. O. Nsoesie, and J. S. Brownstein (2015). Combining search, social media, and traditional data sources to improve influenza surveillance. In *PLoS Computational Biology*.
- [22] Schechter, S. (2005). Towards econometric models of the security risk from remote attacks. *IEEE Security & Privacy* 3(1), 40–44.
- [23] Signorini, A., A. M. Segre, and P. M. Polgreen (2011). The Use of Twitter to Track Levels of Disease Activity and Public Concern in the U.S. during the Influenza A H1N1 Pandemic. *PLoS ONE* 6, e19467.
- [24] Telang, R. and S. Wattal (2007). An empirical analysis of the impact of software vulnerability announcements on firm stock price. *IEEE Transactions on Software Engineering* 33, 544–557.
- [25] Tumasjan, A., T. Sprenger, P. Sandner, and I. Welp (2010). Predicting elections with twitter: What 140 characters reveal about political sentiment. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, pp. 178–185.
- [26] Vomfell, L., W. K. Härdle, and S. Lessmann (2018). Improving crime count forecasts using twitter and taxi data. *Decision Support Systems* 113, 73–85.
- [27] Wang, X., M. S. Gerber, and D. E. Brown (2012). Automatic crime prediction using events extracted from twitter posts. In S. J. Yang, A. M. Greenberg, and M. Endsley (Eds.), *Social Computing, Behavioral - Cultural Modeling and Prediction*, Berlin, Heidelberg, pp. 231–238. Springer Berlin Heidelberg.
- [28] Zhivich, M. and R. K. Cunningham (2009). The real cost of software errors. *IEEE Security & Privacy* 7(2), 87–90.
- [29] Zhongqing, W. and Z. Yue (2017). Ddos event forecasting using twitter data. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI-17)*, Melbourne, Australia, pp. 4151–4157.