Journal of Network and Computer Applications 65 (2016) 131–143

Contents lists available at ScienceDirect

Journal of Network and Computer Applications

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

Economic metric to improve spam detectors

Fida Gillani [a,*], Ehab Al-Shaer [a], Basil AsSadhan [b]

[a] University of North Carolina Charlotte, United States
[b] College of Engineering King Saud University, Saudi Arabia

ARTICLE INFO

Article history:
Received 7 July 2015
Received in revised form
9 November 2015
Accepted 29 February 2016
Available online 5 March 2016

Keywords:
Spam economics
Spam detection
Anomaly detection
Email spam
Consumer economics theory

ABSTRACT

Economic lifting has made email spam a scathing threat to the society due to its related exploits. Many spam detection schemes have been proposed employing the tendency of spam to alter the normal statistical behavior of mail traffic. Threshold tuning of these detectors is still a challenging task. Since, shooting down benign emails as spam (false positive), in pursuit of higher detection rates, can be detrimental. In this paper, we introduce a novel economic metric, based on the underlying spam economic system, to assist detectors in keeping their false positives at bay by associating detection accuracy to the spammer's cost. Hence, the sensitivity of a detector does not need to be tuned all the way up to maximize detection, but enough to make spamming cost unbearable to the spammer. Since, spam is all about making money ultimately. We also show that the statistical features used in our spam detectors can easily differentiate spam from benign and we also show that these features are hard to evade by the spammer. Our evaluation shows the effectiveness of this approach in considerably reducing the false positives for these detectors.

1. Introduction

Economic lifting has made email spam a scathing threat (Hann et al., 2006) to the society due to its related exploits (like malware propagation) (Kanich et al., 2008; Isacenkova et al., 2013; Allodi et al., 2013). Existing spam detectors (Hao et al., 2013; Ramachandran and Feamster, 2006; Clayton; Xie et al., 2008; Smith et al., 2009) employ the logic that the spam activity always changes the normal mail behavior. Different detectors use different statistical features to segregate spam from benign (normal). Nevertheless, threshold tuning is still a challenging task as, shooting down benign emails as spam, in pursuit of higher detection rates, can be detrimental. Since, emails are no longer used to share just fun stories rather they have gained a mission critical status. Therefore, there exists an inherent tradeoff between the accuracy and the efficiency of such behavioral detection.

Such tradeoff and the monetary benefits of spam inspired researchers to study and model the spam economy (Stone-Gross et al., 2011; Ford and Gordon, 2006; Li and Liao, 2009) and to quantify the spammer's earning (Kanich et al., 2011, 2008). Although, the fundamental goal of these approaches is to understand spam economy and halt its progression. But, none has developed a metric for the spam detectors to improve their efficiency. In our original work (Gillani and Al-Shaer, 2014), we introduced a novel economic metric, based on the underlying spam economic system, to assist detectors reduce their false positives by associating detection accuracy to the spammer's cost. Hence, the sensitivity of a detector does not need to be tuned all the way up to maximize detection, but enough to make spamming cost unbearable to the spammer. As shown in Fig. 1, we want to assist detectors finding this sweet spot to defeat spam. A sweet spot represents a threshold where detector is causing enough increase in the spamming cost of the spammer (or enough decrease in the profit of the spammer) with least false positives such that the spamming activity becomes useless to the spammer.

The first contribution of the original paper (Gillani and Al-Shaer, 2014) was in identifying 4 effective statistical mail traffic features that can distinguish spam from benign. In this extended version, we have considered 6 more features to rigorously test all available features presented in the existing literature (Smith et al., 2009; Hao et al., 2013; Ramachandran and Feamster, 2006; Clayton; Xie et al., 2008). Furthermore, we want to assure that the final features must be hard to evade by the spammer. Therefore, we have also added evasion analysis of all these features in this extended version. We want to use best features available in our economic modeling. We use K directed divergence (Cover and Thomas, 1991) measure to analyze the discerning capacity of these features. We perform this analysis on our own dataset that is

* Corresponding author.
E-mail addresses: sgillan4@uncc.edu (F. Gillani), ealshaer@uncc.edu (E. Al-Shaer), bsadhan@ksu.edu.sa (B. AsSadhan).

http://dx.doi.org/10.1016/j.jnca.2016.02.019
1084-8045/

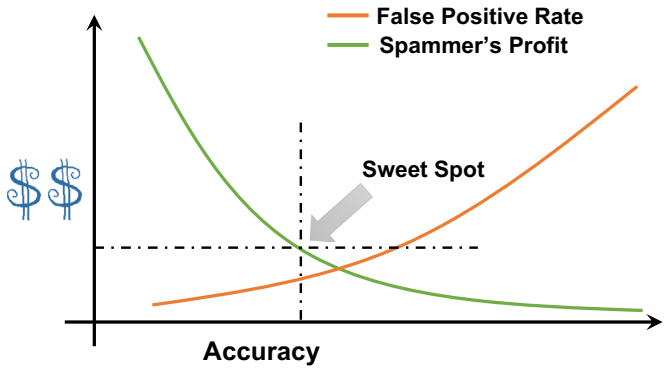

Fig. 1. Economic metric.

comprised of around 75,000 benign emails and around 3 million spam emails.

Our analysis reveals four features to stand out among all: (1) inter-departure time (IDT), which is the time between two consecutive emails, (2) emails per recipients (EPR), which is the number of emails sent to a recipient, (3) email size (ES), which gives the average email size, and (4) distribution of new recipients (DNR), which provides the frequency of new recipients appearing in a time window of email inspection. Afterwards, we benchmark the performance of these features using ROC (Fawcett, 2006) curves to establish a baseline to later test the impact of our economic metric.

The second contribution of the original paper was in developing a spam economic model to quantify the spammer's utility associated to a spam activity. We use the classical consumer theory of economics to model spam economy, where spammer acts as a consumer looking to buy a product (commodity) that could maximize his/her utility. To define commodity, we have used all the parameters that the spammer would look for in renting a botnet. In this extended version, we describe the intuition of these parameters in detail. In our economic model, we assume a rational spammer behavior. According to which, a spammer will choose a commodity that would yield maximum utility. This intuitive assumption is largely used in the existing literature (Ford and Gordon, 2006; Li and Liao, 2009).

The cost of generating a spam activity is calculated from the price quotes available in the current botnet market (Goncharov; Namestnikov). In actual, detectors force spammer to invest more by either increasing the duration of the spam activity or using more resources (bots) without detectors knowing that. We want to develop this insight into the detectors to look for both accuracy and spamming cost. For this purpose, we use the statistical features to constrain our economic model. Now, detectors adjust these statistical features to gain accuracy which in-turn constrain the spam economic model (utility). This reduced utility forces the spammer to add cost (resources or time) which we calculate and provide to the detectors. We repeat our spam detector benchmarks to observe this increase of spamming cost. In our evaluation, we map these accuracy and cost results together to show the final improvement in the entire detection process.

We structure the remainder of this paper as follows: in Section 2 we establish the novelty of our approach through literature review. We explain the feature selection mechanism and evasion analysis of all features in Section 3, followed by the performance benchmarks of the selected features in Section 4. The discussion of the economic model is presented in Section 5. Section 7 provides the performance benchmark of the detector with the spammer's cost. In the end Section 8 provides the conclusion and future directions of our work.

2. Related work

The purpose of this related work is to show that all spam detection techniques focus only on the behavioral divergence aspect of the problem disregarding the economic aspect completely. On the other extreme, all spam economic based studies try to understand and quantify spam economy without introducing any metric for detectors to exploit these findings. As per our understanding this is the first study that is proposing an economic metric based on spam economic model to assist detectors in tuning their thresholds.

In spam detection, some studies suggested intervention from the service providers to stop dissemination of large volumes of spam. For example, the study (Hao et al., 2013) proposed proactive blacklisting of spammer's domain through registrar monitoring and domain registration frequency to dampen spam and another study (van Eeten et al., 2010) proposed ISPs to monitor the involvement of different IPs in spamming to filter their traffic. These studies do not diminish the end point remedies to cope spam problem. An end point based spam detection technique (Smith et al., 2009) used entropy to measure effectiveness of different statistical features of email traffic to differentiate spam from benign without considering spam economics. In Xie et al. (2008), a framework called AutoRE was presented to filter out any legitimate URLs and focused on the URL that the spammer wants his victims to click on to buy his merchandise or download his malware. Using their signature method, they were able to identify botnet membership and determine which bots were used in the various spam campaigns. The work in Ramachandran and Feamster (2006) focused on the network properties of spam and showed that network-level characteristics of spam are sufficiently different than those of legitimate emails. Then another work in Clayton detected spam from email server logs by measuring the change in the mail behavior of a source over time. All of these studies were very effective bot/spam detectors but without any concern to the underlying economic model.

On the economic front, the economic study (Li and Liao, 2009) proposed an abstract economic model of botnet usage for DDoS attack from both bot botmaster and spammer's perspective. They introduced the concept of honeypots (fake bots) to increase the probability of failure for the attacker. However, the authors do not associate their model to any parameters used by their filters. Some other very prominent studies (Kanich et al., 2011, 2008; Böhme and Holz, 2006; Stone-Gross et al., 2011) rigorously analyzed the spam economics using empirical measurements based approach. They have tried to quantify the spam revenue by analyzing the spam market for months, but they did not establish any connection with the amount of spam activity required for such revenue. They provided an estimate that it requires almost 10 million spam emails to get a positive response, though. A similar study (Garg et al., 2013) provided microeconomic analyses of ecrime to develop a set of hypotheses to predict potential participating crowd. Another study (Herley, 2012) provided an abstract model that described the impact of reducing target density for the spammer but it did not provide any metric for the detectors to actually reduce this target density.

3. Feature selection for spam detection

We want to identify the statistical features from the existing literature that could effectively throttle the spam activity. The effectiveness of any feature depends upon its capability to differentiate spam from the benign emails and its ability to make evasion harder for the spammer. We have used our own collected dataset to test the feasibility of different mail traffic statistical