Rothchild D. Finding Fraudulent Websites Using Twitter Streams. *Technology Science*. 2015092905. September 29, 2015.http://techscience.org/a/2015092905

Finding Fraudulent Websites Using Twitter Streams

Daniel Rothchild

Highlights

- I developed a monitoring program that searches Twitter in real time for tweets with potentially suspicious links
- The program found more than 70,000 suspicious tweets in 24 hours, with 56% of the tested links appearing fraudulent

Count	URL
12095	http://womanshealthlifestyle.com/PureGarciniaCambogia/
10328	http://muscleandhealth.info/
3556	http://muscleformen.com/Metaboosts
2033	http://healthyreport.co/nfl-wants-to-ban-supplement/index.html
1953	http://womenshealthmag.com-article.link/
324	http://www.forcefactor.com/h/
193	http://www.uniquegarcinia.com/
181	http://tmzf.itness.co/index.html

Most frequently occurring tweets in 24 hours that contain the words muscle, weight, diet, acai, cambogia, lose fast, or miracle pill.

Abstract

Social media offers new opportunities for Internet fraudsters to direct traffic to websites that make fraudulent offers. In an interesting twist, however, social media also enables consumer protection organizations and government agencies to monitor traffic for suspicious links to these websites. I developed an automated monitoring algorithm that searches Twitter in real time for suspicious links appearing in tweets.

Results summary: Searching for tweets containing a URL and at least one of the keywords muscle, weight, diet, acai, cambogia, lose fast, and miracle pill, my program downloaded more than 70,000 tweets during a 24-hour period. The most-tweeted URL was tweeted over

12,000 times. I visited the 50 most tweeted URLs and classified 28 of 50 (or 56 percent) of them as suspicious and 20 of 50 (or 40 percent) as unsuspicious. The 28 suspicious URLs redirected to only 8 distinct URLs. Of the 10 most frequently tweeted URLs, all 10 were suspicious.

Introduction

The Federal Trade Commission (FTC) released a statistical survey of online and offline fraud in the United States in 2011. The survey found that about 25.6 million American adults (10.8 percent of the adult population) were victims at least one fraud. About 5.1 million U.S. consumers were victims of fraudulent weight-loss products, and 2.4 million were victims of fraudulent prize offers. Unauthorized billing for buyers' club memberships and for Internet services tied for third place. Work-at-home programs also ranked high on the list [1]. The FBI's Internet Crime Complaint Center received about 300,000 complaints of Internet crime for 2014, and the true incidence of Internet crime is estimated to be more than 10 times this number [2].

In many cases, fraudsters lure victims to websites they have set up in order to collect their personal or financial information or to present them with advertisements that make deceptive offers. Given the enormous growth in popularity of social media in the past decade, it is not surprising that fraudsters use social media as a new tool to reach victims. However, what may be surprising is the extent to which it is possible for law enforcement agencies and consumer protection groups to turn the fraudsters' new tool against them in order to more effectively detect and prevent fraud. Does social media offer a new opportunity for fraud prevention? Might law enforcement agencies and consumer protection groups now have the opportunity to monitor social media in real time for posts from fraudsters that attempt to direct traffic to potentially fraudulent sites?

There are several advantages to proactively discovering and monitoring fraudulent websites. First, data collected in this way is unbiased by whether victims of different kinds of fraud choose to report or not. Some frauds are underreported, and many target specific demographics, so those who report frauds do not necessarily represent all those who are heavily targeted. Second, fraudsters frequently put up, take down or move their fraudulent websites on short timescales, perhaps to avoid detection. Continuously monitoring social media data might allow consumer protection groups to become aware of fraudulent websites much faster than they could otherwise. Finally, having a more complete record of suspicious activity on social media could prove invaluable to law enforcement when prosecuting fraudsters who try to conceal the extent of their fraud.

Background

Twitter serves as a convenient social medium to study. Unlike many other forms of social media, Twitter provides a live stream of every publicly posted tweet, which is easily accessible to anyone who digitally checks the stream.

Checking Twitter's stream can be automated. Twitter offers a streaming application program interface (API) that allows anyone to write a custom program to freely connect to Twitter's computers in order to monitor the public stream of tweets. The custom program specifies keywords to track, and Twitter's computers send to the custom program tweets that contain one or more of those keywords. In fact, Twitter's servers will continually send matching tweets as they are posted until the custom program ends the connection.

One important limitation of Twitter's streaming API is that, except for users with elevated privileges, the API will limit the number of tweets it sends to about 1 percent of the total number of tweets being posted at any given time [3]. This limitation did not affect my small-scale trial because I used few enough keywords that the volume of tweets I was requesting never approached 1 percent of all tweets. However, the 1 percent limitation might come into play were this script to be run with a larger number of keywords.

Methods

I wrote a Python script to search for keywords in Twitter's stream. The keywords are terms I loosely associated with common types of fraud, namely: muscle, weight, diet, acai, cambogia, lose fast, and miracle pill. The FTC report [1] mentioned above identifies several terms that may be associated with fraudulent offers, including weight, diet, exercise, weight-loss, prize, sweepstakes, lottery, and winner. The exact choice of keywords is unimportant because it is possible to use any keywords that are suspected of being associated with fraud. Regardless of the specific words, many innocent tweets and legitimate offers may also include those terms. Therefore, the critical question is whether searching for simple keywords in tweets that include URLs will be sufficient to identify suspicious offers.

Whenever a new tweet is sent over, the script detects whether the tweet contains a URL; if so, it stores the text of the tweet and the twitter account that posted the tweet in a database. This script can be left running indefinitely, and it is possible to query the database at any time to see which URLs have been found. The algorithm ranks how suspicious each URL is by finding the total number of times it was tweeted during the time the script was running. Ordering the tweets by URL frequency and examining the 50 most frequently captured URLs should be sufficient to determine whether the approach finds suspicious websites.

I used visual inspection to determine whether a website was suspicious. My criteria were apparent false claims about a product, inclusion of comments endorsing the product that appear to be posted by fake users, or being exceedingly misleading about the price the consumer would be charged by only listing the true price in the fine print of the website.

Results

I ran the script on 5/1/2015 for approximately 24 hours. During that time, the Twitter computers sent over 71,067 tweets that contained URLs. Of these, 19,567 contained the term muscle, 32,150 contained weight, 24,098 contained diet, 196 contained acai, 517 contained cambogia, 0 contained lose fast, and 4 contained miracle pill.

Table 1 lists the top 50 sites found by my script that accounted for a total of 36,522 tweets (51 percent of all tweets surveyed). Several of them are variations of the same websites. In the top 50 URLs my script identified, there were only 29 distinct websites once the URL redirections were followed. Table 2 shows the 8 of these 29 distinct websites that I found to be suspicious upon visual inspection. The most-tweeted site, once resolved, was tweeted over 12,000 times.

Count	URL
7184	http://skinnypills.co/1264
4195	http://muscle.co/1nGqHH
3618	http://skinnypills.co/1228
3095	http://muscle.co/SXxlgv
2156	http://musclepills.co/1722
1400	http://musclepills.co/1094
1188	http://skinnypills.co/1722
1100	http://m.uscle.co/SH1Y2
925	http://muscle.co/8xpT6R
845	http://m.uscle.co/SH1vS
743	https://cards.twitter.com/cards/18ce53y62oj/into
731	http://flt.bz/2JiW7N
623	http://muscle.co/NrSWAK
587	http://goo.gl/ux9SZb
549	http://muscle.co/YcSDUg
485	http://flt.bz/YMDwv9
438	http://bit.ly/1DtO38n
412	http://muscle.co/14fy6b
409	http://muscle.co/MGSylE
368	http://flt.bz/xFBQGm
350	http://bit.ly/1QPTlCO
348	http://bit.ly/1gCtNnf
344	https://vine.co/v/e7DBL6LM6qw
325	http://bit.ly/1Ne8uyv
324	http://bit.ly/1dvDJVK
295	http://my-extreme-weight-loss.com
294	http://bit.ly/1DNRD17
243	http://goo.gl/ogXEVb
233	http://bit.ly/1GuEDLZ
202	http://Express.co.uk

Count	URL
201	http://bit.ly/1DwZgsC
193	http://loseweightburnfat.info
182	http://bit.ly/1Dx66yz
181	http://f.itness.co/SH1is
149	http://bit.ly/1fQJVb5
142	http://flt.bz/X3CrA2
134	http://ift.tt/1JPMq7W
124	http://ift.tt/18ASnJq
122	http://flt.bz/N4SXko
120	http://muscie.co/TPr69K
114	http://Examiner.com
107	http://M.Diet
105	http://skinnypills.co/1324
105	http://flt.bz/XbQHhO
104	http://ift.tt/1fceOQL
90	http://es.pn/1Q8X11d
88	http://m.uscle.co/SH37k
86	http://goo.gl/ktydxK
83	http://goo.gl/PZKjcd
83	http://bit.ly/1zAHqUm

Table 1. The 50 most frequently tweeted URLs in tweets containing muscle, weight, diet, acai, cambogia, lose fast, or miracle pill in a 24-hour period. Count is the number of times the URL appeared in a tweet. A total of 71,067 tweets containing URLs were surveyed.

Count	URL
12095	http://womanshealthlifestyle.com/PureGarciniaCambogia/
10328	http://muscleandhealth.info/
3556	http://muscleformen.com/Metaboosts
2033	http://healthyreport.co/nfl-wants-to-ban-supplement/index.html
1953	http://womenshealthmag.com-article.link/
324	http://www.forcefactor.com/h/
193	http://www.uniquegarcinia.com/
181	http://tmzf.itness.co/index.html

Table 2. The 8 distinct URLs found by resolving redirections of the 28 URLs in Table 1 that appeared to be suspicious upon visual inspection. Count is the number of times a URL which redirected to the listed URL appeared in a tweet. A total of 71,067 tweets were surveyed.

I visited each of the 50 websites listed in Table 1. Figure 1 shows a screenshot of the most tweeted website that my script found, with 12,095 total tweets. The page shown is the final

checkout page, where the user is asked for his or her credit card information. The top half of the page (which on a regular-sized screen may be all that the user sees without scrolling down) indicates that the user will only be charged $5 for shipping and handling. However, as the difficult-to-read fine print that I have highlighted reads: “By placing an order, you will pay S&H to receive a 30 day supply. You will also be automatically enrolled in our membership program. The program will charge you $86.94 on the 14th day of your order date for a monthly supply and every 30 days thereafter until you cancel. You can cancel at any time by calling 1-855-511-1315. If you cancel before the 14th day of your order date, you pay the S&H of your 30-day supply. If you cancel after the 14th day of your order date, you shall pay for the 30 day supply plus any future supplies without refund.”

Figure 1: A suspicious site found by my script. It ranked 1 out of 50 and had a count of 12,095 tweet occurrences. This image was captured from on 5/3/2015 but by the time of this writing had been taken down.

I classified 28 of 50 (or 56 percent) of URLs in Table 1 as suspicious and 20 of 50 (or 40 percent) as unsuspicious. The other 2 (or 4 percent) were pornographic (and I did not investigate them for possible fraud). The 28 suspicious URLs redirected to 8 distinct suspicious URLs. Of the top 10 URLs, all 10 (or 100 percent) of them were suspicious, and all of them redirected to one of the 8 suspicious URLs. Table 3 summarizes these results.

	Suspicious	Not Suspicious	Other
Top 10 URLs	10	0	0
Top 50 URLs	28	20	2
Distinct Top 50 URLs	8	20	1

Table 3. Number of URLs found suspicious after visual inspection of its website. The 10 most frequently tweeted URLs were all suspicious. Overall, 28 of the 50 URLs were suspicious, and 20 were not. The 28 suspicious URLs redirected to 8 distinct URLs.

Discussion

My experiment demonstrates that simple keyword searches of tweets can help identify potentially fraudulent offers. This could inspire a new tool for law enforcement and could also potentially be used in other ways. For example, it might be possible to create a browser extension that would warn users when they visit a site that has been automatically identified by the tool as suspicious. This would also protect users from newly created fraudulent websites, and would operate without requiring the user to take any action besides installing the extension. This functionality could even potentially be integrated within the browser.

There are many opportunities for further work that would improve my approach. First, the method by which I ranked sites as being suspicious is simple: links were counted as suspicious solely based on how many times they were tweeted. A more sophisticated method might be to detect whether a link that was tweeted many times has such a high count because a small number of users tweeted it many times (which would likely be the case if bots were doing the posting) or whether many different users tweeted the same link (which might be true in the case of a popular, non-fraudulent web post). Another improvement might be to investigate the tweet contents. For example, bots tend to post the same link many times with the same or similar text in the body of the tweet. Real users, on the other hand, are less likely to tweet the same text repeatedly. Lastly, I chose the keywords in an ad hoc manner. A future study could determine which terms are best to use. Employees of law enforcement agencies or consumer protection groups could help researchers by providing text from websites found to be fraudulent.

References

1. Federal Trade Commission. Consumer Fraud in the United States, 2011. The Third FTC Survey. April 2013.

https://www.ftc.gov/sites/default/files/documents/reports/consumer-fraud-united-states-2011-third-ftc-survey/130419fraudsurvey_0.pdf

2. Federal Bureau of Investigation. 2014 Internet Crime Report. Inter- net Crime Complaint Center. Accessed September 15, 2015. https://www.fbi.gov/news/news blog/2014-ic3-annual- report/

3. Twitter Developers. Is There a Limit to the Amount of Data the Streaming API Will Send Out? Twitter, March 30, 2012. https://twittercommunity.com/t/is-there-a-limit-to-the-amount-of-data-the-streaming-api-will-send-out/8482

Authors

Daniel Rothchild is a junior at Harvard College majoring in Physics. He has internship experience analyzing datasets such as employment record data and biological data, and last summer performed research at Harvard as a PRISE fellow under computational linguistics professor Stuart Shieber. He is a board member and cellist of the River Charles Ensemble, and also plays in a piano trio as part of Harvard's chamber music program. Daniel was a Presidential Scholar, a winner of the Detur Book Prize, and a John Harvard Scholar.

Editor: Latanya Sweeney

Citation

Rothchild D. Finding Fraudulent Websites Using Twitter Streams. *Technology Science*. 2015092905. September 29, 2015.http://techscience.org/a/2015092905

Data

Under review for data sharing classification. Data release available October 19.