



# News

[Stories](#)[Media Highlights](#)[Media Resources](#)[Contact Us](#)

[News](#) › [Stories](#) › [Archives](#) › [2016](#) › [September](#) › CMU Algorithm Detects Online Fraudsters

September 08, 2016

## CMU Algorithm Detects Online Fraudsters

### Method Sees Through Camouflage To Reveal Fake Followers, Reviewers


By Byron Spice / 412-268-9068 /

[bspice@cs.cmu.edu](mailto:bspice@cs.cmu.edu)



An algorithm developed at Carnegie Mellon University makes it easier to determine if someone has faked an Amazon or Yelp review or if a politician with a suspiciously large number of Twitter followers might have bought and paid for that popularity.

The method, called FRAUDAR, marks the latest escalation in the cat-and-mouse game played by online fraudsters and the social media platforms that try to out them. In particular, the new algorithm makes it possible to see through camouflage that fraudsters use to make themselves look legitimate, said Christos Faloutsos, professor of [machine learning](#) and [computer science](#).



Professor Christos Faloutsos and his team won a Best Paper Award at the ACM Conference in San Francisco.

In real-world experiments using Twitter data for 41.7 million users and 1.47 billion followers, FRAUDAR fingered more than 4,000 accounts not previously identified as fraudulent, including many that used known follower-buying services such as TweepMe and TweeterGetter.

“We’re not identifying anything criminal here, but these sorts of frauds can undermine people’s faith in online reviews and behaviors,” Faloutsos said. He noted most social media platforms try to flush out such fakery, and FRAUDAR’s approach could be useful in keeping up with the latest practices of fraudsters.

The CMU algorithm is [available as open-source code](#). A research paper describing the algorithm

won the Best Paper Award last month at the Association for Computing Machinery's Conference on Knowledge Discovery and Data Mining (KDD2016) in San Francisco.

Faloutsos and his data analytics team specialize in graph mining, a method that looks for patterns in the data. In this case, social media interactions are plotted as a graph, with each user represented as a dot, or node, and transactions between users represented as lines, or edges.

The state-of-the-art for detecting fraudsters, with tools such as Faloutsos' [NetProbe](#), is to find a pattern known as a "bipartite core." These are groups of users who have many transactions with members of a second group, but no transactions with each other. This suggests a group of fraudsters, whose only purpose is to inflate the reputations of others by following them, by having fake interactions with them, or by posting flattering or unflattering reviews of products and businesses.

But fraudsters have learned to camouflage themselves, Faloutsos said. They link their fraudulent accounts with popular sites or celebrities, or they use legitimate user accounts they have hijacked. In either case, they try to look "normal."

FRAUDAR can prune away this camouflage. Essentially, the algorithm begins by finding accounts that it can confidently identify as legitimate — accounts that may follow a few random people, those that post only an occasional review and those that otherwise have normal behaviors. This pruning occurs repeatedly and rapidly. As these legitimate accounts are

eliminated, so is the camouflage the fraudsters rely upon. This makes bipartite cores easier to spot.

To test the algorithm, Faloutsos and his students used a massive Twitter database extracted from the social media platform in 2009 for research purposes. FRAUDAR found more than 4,000 accounts that appeared highly suspicious, though most of the tweets had not been removed and the accounts had not been suspended in the seven years since the data was collected.

The researchers randomly selected 125 followers and 125 followees from the suspicious group, along with two control groups of 100 users who had not been picked out by the algorithm. They examined each for links associated with malware or scams and for clear robot-like behavior, such as replying to large numbers of tweets with identical messages. They found 57 percent of the followers and 40 percent of the followees in the suspicious group were labeled as fraudulent, compared to 12 percent and 25 percent in the control groups.

Among the suspicious accounts, the researchers found 41 percent of the followers and 26 percent of the followees included advertising for follower-buying services — 62 percent and 42 percent, respectively, if deleted or suspended accounts are ignored. Few such mentions were found in the control groups.

“The algorithm is very fast and doesn’t require us to target anybody,” Faloutsos said. “We hope that by making this code available as open source, social media platforms can put it to good use.”

In addition to Faloutsos, the research team included Ph.D. students Bryan Hooi, Hyun Ah Song, Neil Shah and Kijung Shin, and Alex Beutel, who recently received his Ph.D. in computer science. The National Science Foundation supported this research.

[The Piper: Campus & Community News](#)

[Official Events Calendar](#)

Carnegie Mellon  
University  
5000 Forbes Avenue  
Pittsburgh, PA 15213  
412-268-2900

[Legal Info](#) |  
[www.cmu.edu](http://www.cmu.edu)  
© 2016 Carnegie Mellon  
University