


Blocking 'fake engagement' to keep the count honest

May 20 2016, by Bill Steele



When you see that a YouTube video has "16,685 views," take that with a grain of salt. Not all of those views may have been by human beings.

There are services that will, for a fee, spam a social media site with computer-generated views, likes, comments and other actions to boost a posting's apparent popularity and draw more attention. Videos with a lot of views, for example, will be featured on YouTube's opening page.

"Bad actors have been trying to game the system," said Yixuan Li, a graduate student working with John Hopcroft, the IBM Professor of Engineering and Applied Mathematics in the Department of Computer Science. The problem is not limited to YouTube, the researchers pointed out, noting "Twitter followers, Amazon reviews and Facebook likes are all buyable by the thousand."

In "a world that counts," the researchers said, the count should reflect genuine interest.

The good news is that Li, Hopcroft and colleagues at Google have developed a way to recognize and block this "fake [social engagement](#)." Li began the project while interning at Google, and the system is now coming into use on their sites, he said. Li described the system, called "LEAS" (Local Expansion At Scale) in a paper presented at the 25th International World Wide Web Conference held April 11 to 15 in Montreal.

A tipoff, Li explained, is that the accounts posting the fake hits are in "lockstep," posting to the same video targets around the same times. LEAS creates a map – officially known as an "engagement relationship graph" – of accounts and links between them, and behavioral similarity over time. It learns by looking at known spamming accounts (called "seeds"), then searches on the engagement graph for sets of accounts similar to the seeds performing orchestrated actions that have very low likelihood of happening spontaneously. It works best, the researchers said, to focus on small "local" sections of the graph.

To evaluate the system, humans manually reviewed postings from accounts LEAS had identified as spammers on YouTube. Some of those accounts had been created very recently but had run up a long list of postings. The comments they posted often amounted to just "good video," "Yeah," "Cool" and other all-purpose bits of text, and identical comments had been posted to several videos. Some of the comments included malicious web links and advertisements.

Provided by Cornell University

Citation: Blocking 'fake engagement' to keep the count honest (2016, May 20) retrieved 8 July 2024 from <https://phys.org/news/2016-05-blocking-fake-engagement-honest.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.