On the hunt for Facebook's army of fakes

May 12 2016, by Emiliano De Cristofaro

Credit: AI-generated image (disclaimer)

Ever wonder why so many people like certain Facebook pages, no matter how boring or badly updated they are? They could well be the subject of "like farming", the process of artificially inflating the number of Facebook page likes.

Researchers like myself have developed computer algorithms that can tell genuine likes from artificial ones generated by farm-controlled

accounts. But it turns out that more sophisticated farms are evading detection tools, including those deployed by Facebook itself. So we've now developed an experimental way of looking for farmed accounts, including those that are run by real human users.

Facebook pages allow their owners to publicise products and events, communicate with customers and fans and promote themselves using targeted ads. More than 40m small businesses reportedly have active pages, and almost 2m of them use Facebook's advertising platform possibly to broaden their audience and engage with more customers.

If someone wants to quickly increase their page's number of likes, they can also purchase them from farmers for between around $10 (£7) and $100 (£70) per 100 likes, depending on whether they want to target specific regions. For example, likes from US-based accounts are usually more expensive. You can even buy entire pre-liked pages with large numbers of followers that you can then adapt to promote your own organisation. While these paid-for likes may not come from engaged customers, they can make the page or its owner appear more popular, in turn increasing its appeal to potential customers or followers.

There are several ways that farms can generate fake likes, and the method they use significantly affects both their cost and how hard it is to detect them. One obvious way is to create fake accounts, although this is somewhat cumbersome because Facebook has checks in place, such as having to input a code displayed on screen or sent to a mobile phone, to prevent this being done automatically by computer "bots". Another strategy is to take control of real accounts whose passwords have been leaked or captured using software that spies on people's computers.

But, importantly, there are also networks of real users who will like pages on request in return for other services or small payments. And you can lure users to like a page by promising them access to lotteries,

discounts or exclusive content.

Different farms also use different strategies to avoid detection. Some deliver likes in bursts and employ accounts that are not really connected to the rest of the social network, making them easier to spot. Others use a stealthier approach, mimicking regular users' behaviour such as liking genuinely popular pages and paid adverts. Each account only likes a small number of pages and relies on many accounts, each connected with many different friends, to gradually deliver likes.

This strategy of using fake accounts to like genuinely popular pages can cause embarrassment if exposed. For example, Hillary Clinton was criticised when her Facebook account suddenly received thousands of likes from Thailand and Myanmar overnight. But it can also harm legitimate Facebook users running advertising campaigns, who pay for clicks from real users but receive them from fake ones.

In an attempt to counter farming, Facebook, in collaboration with university researchers, has developed and deployed several tools to detect spam and fake likes. One, called CopyCatch, detects groups of fraudsters acting together, generally liking the same pages at around the same time. Another method, called SynchroTrap, relies on the fact that malicious accounts usually perform similar actions around the same time. So the algorithm can detect these fakes when it spots a cluster of them acting together over a sustained period of time.

The problem is that these methods are unlikely to spot the stealthier (and more expensive) farms that rely on the accounts of real people rather than fake or compromised profiles. This is because focusing on activity patterns of pages and users fails to capture important characteristics of these "real" accounts used by the farms. These profiles are often created mainly as a money-making tool and so their activity is different from a typical account used for social networking.

Not so "real" users

In our recent study, my colleagues and I set out to address this gap by looking at how and what users post on Facebook, in order to improve the accuracy of detection mechanisms. We found that posts made by these "real" farm accounts had fewer words, a more limited vocabulary, and lower readability than normal users' posts. Their posts were also highly focused on some specific topics, generate significantly more comments and likes, and a large fraction of their activity was simply sharing content such as articles, videos and posts made by other users.

We then trained machine-learning algorithms to use these patterns to analyse a set of accounts we knew included farmed likes. We found that the algorithms were nearly perfectly accurate at detecting farm accounts, including the more stealthy "real" ones.

We've yet to see if the same techniques could be used to accurately detect farmed likes across Facebook's 1.2 billion users and many billions more posts. What we may find is that as these techniques become better at spotting farmed accounts, those accounts find new ways of changing their posting behaviour to become even better at mimicking "innocent" users, in an economic game of cat and mouse. The question is how much this will cost them and whether creating even more realistic farmed accounts will be worth it.

This article was originally published on The Conversation. *Read the* original article.

Source: The Conversation