


Whose Likes Matter Most? Ask this Learning Algorithm

News, 2018



By Meredith Alexander Kunz, Adobe Research

Tracking the most influential voices online is critical to spreading information on the web. Whether it's for marketing a product, promoting a candidate, or sharing knowledge about health, many organizations are trying to locate the biggest influencers and leverage their online bullhorns.

Finding those individuals is not always straightforward, however, and needs study. Now, new work from Adobe Research offers a promising means to empower computers to find influencers on their own—using sophisticated methods in machine learning.

Zheng Wen and **Branislav Kveton**, both senior research scientists at Adobe Research, partnered with Michal Valko at France's INRIA and Sharan Vaswani at the University of British Columbia to develop this novel approach. They published their findings at December's Conference on Neural Information Processing Systems (NIPS) in *Online Influence Maximization under Independent Cascade Model with Semi-Bandit Feedback*.

The researchers began by creating an artificial social network based on real-world data, including friendships, likes, and followers, and they then deployed a new learning algorithm they had developed to figure out who influenced whom—starting from scratch.

Part of what makes this work so groundbreaking is that it doesn't let the past dictate the future. Instead of relying on old, possibly stale data to inform their algorithm, the researchers built an active learning algorithm that explores on its own and gathers new data as it interacts with the social network. That way, it improves over time. That fundamental shift has led to much more impressive results, says Kveton.

"In traditional machine learning, the problem was that past data is not representative of the future. Now, we have computers and systems that interact with people in the present, and that can actively learn," he explains. The algorithm is also able to adapt seamlessly from exploring the network for learning purposes to using that knowledge to spread influential information rapidly.

The system did well in identifying influencers on its own while gaining feedback from the network in real time. In an experiment on a Facebook subnetwork, the team's algorithm learned to locate and tap into social media influencers very effectively. Their approach influenced 50 percent more users than the best competing state-of-the-art baseline.

The researchers simplified their model to promote faster learning. But this approach is ultimately scalable to real-world social networks, they say.

This kind of active learning algorithm—where the system collects its own data and improves its model—could also have an important impact on many other efforts in marketing and communications. Viral marketing, for instance, is one of the most efficient ways to spread the word about a product or service. The influencers identified by this algorithm could be tapped to launch a trend virally. Someday, this experimental work could offer a major boost to those trying to get the word out online.

Contributors:

[Zheng Wen](#) and [Branislav Kveton](#), senior research scientists, Adobe Research

Michal Valko, INRIA

Sharan Vaswani, University of British Columbia