The science of retweets

October 6 2015

What's the best time to tweet, to ensure maximum audience engagement? Researchers at the University of Maryland have demonstrated that an algorithm that takes into account the past activity of each of your followers—and makes predictions about future tweeting—can lead to more "retweets" than other commonly used methods, such as posting at peak traffic times.

The internet is full of advice about when to tweet to gain maximum exposure, but the new study subjects marketing folk wisdom to scientific scrutiny.

William Rand, director of the Center for Complexity in Business in UMD's Robert H. Smith School of Business, with co-authors from the departments of scientific computation and physics, examined the retweeting patterns of 15,000 Twitter followers during two different five week intervals, in 2011 and 2012, from 6 a.m. to 10 p.m. Retweets are especially valuable to marketers because they help to spread a brand's message beyond core followers.

Most marketers are well aware there's a pattern to Twitter traffic. In the early morning, nothing much happens. Then people get into work and retweet intensely, as they do their morning surfing. The number of retweets drops as the day progresses, with a slight uptick at 5 p.m. Then it picks up again later "when people get back to their computers after dinner, or are out at a bar or restaurant using their phones," as Rand puts it. Monday through Friday follow roughly that pattern, but Saturday and Sunday show markedly different behavior, with much smaller morning spikes and less decline during the day.

A "seasonal" model of posting—the folk-wisdom model—would suggest posting whenever there are peaks in that recurring weekly pattern. (Which peaks you choose would depend how many tweets you expect to send.)

The authors compared that model to two others: The first added to the seasonal model a component that looked for unusual surges and declines (caused by, say, big news events) and adjusted posting patterns correspondingly. They built the final model from scratch: It took into account the individual tweeting behavior of each follower and predicted his or her likelihood of tweeting in the next 10 minutes.

The authors first had to write software that collected the tweets. For each five-week period studied, the authors used the first four weeks to build a model and the final week for testing it, by tweeting and watching what

happened.

All three models were reasonably effective, but the algorithm that the authors wrote, which took each individual's behavior into account, was the most successful at generating retweets. The paper serves as a demonstration that applying analytic methods to Twitter data can improve a brand's ability to spread its message. The authors made the open-source software developed for the study available online.

  More information: "Forecasting High Tide: Predicting Times of Elevated Activity in Online Social Media," *Proceedings of Advances in Social Networks Analysis and Mining* (*ASONAM*).

Provided by University of Maryland