

Model for evaluating productrecommendation algorithms suggests that trial and error get it right

November 14 2014, by Larry Hardesty



Credit: AI-generated image (disclaimer)

Devavrat Shah's group at MIT's Laboratory for Information and Decision Systems (LIDS) specializes in analyzing how social networks process information. In 2012, the group demonstrated algorithms that could predict what topics would trend on Twitter up to five hours in



advance; this year, they used the same framework to predict fluctuations in the prices of the online currency known as Bitcoin.

Next month, at the Conference on Neural Information Processing Systems, they'll present a paper that applies their model to the recommendation engines that are familiar from websites like Amazon and Netflix—with surprising results.

"Our interest was, we have a nice model for understanding dataprocessing from social data," says Shah, the Jamieson Associate Professor of Electrical Engineering and Computer Science. "It makes sense in terms of how people make decisions, exhibit preferences, or take actions. So let's go and exploit it and design a better, simple, basic recommendation algorithm, and it will be something very different. But it turns out that under that model, the standard recommendation algorithm is the right thing to do."

The standard algorithm is known as "collaborative filtering." To get a sense of how it works, imagine a movie-streaming service that lets users rate movies they've seen. To generate recommendations specific to you, the algorithm would first assign the other users similarity scores based on the degree to which their ratings overlap with yours. Then, to predict your response to a particular movie, it would aggregate the ratings the movie received from other users, weighted according to similarity scores.

To simplify their analysis, Shah and his collaborators—Guy Bresler, a postdoc in LIDS, and George Chen, a graduate student in MIT's Department of Electrical Engineering and Computer Science (EECS) who is co-advised by Shah and EECS associate professor Polina Golland—assumed that the ratings system had two values, thumbs-up or thumbs-down. The taste of every user could thus be described, with perfect accuracy, by a string of ones and zeroes, where the position in



the string corresponds to a particular movie and the number at that location indicates the rating.

Birds of a feather

The MIT researchers' model assumes that large groups of such strings can be clustered together, and that those clusters can be described probabilistically. Rather than ones and zeroes at each location in the string, a probabilistic cluster model would feature probabilities: an 80 percent chance that the members of the cluster will like movie "A," a 20 percent chance that they'll like movie "B," and so on.

The question is how many such clusters are required to characterize a population. If half the people who like "Die Hard" also like "Shakespeare in Love," but the other half hate it, then ideally, you'd like to split "Die Hard" fans into two clusters. Otherwise, you'd lose correlations between their preferences that could be predictively useful. On the other hand, the more clusters you have, the more ratings you need to determine which of them a given user belongs to. Reliable prediction from limited data becomes impossible.

In their new paper, the MIT researchers show that so long as the number of clusters required to describe the variation in a population is low, collaborative filtering yields nearly optimal predictions. But in practice, how low is that number?

To answer that question, the researchers examined data on 10 million users of a movie-streaming site and identified 200 who had rated the same 500 movies. They found that, in fact, just five clusters—five probabilistic models—were enough to account for most of the variation in the population.



Missing links

While the researchers' model corroborates the effectiveness of collaborative filtering, it also suggests ways to improve it. In general, the more information a collaborative-filtering algorithm has about users' preferences, the more accurate its predictions will be. But not all additional information is created equal. If a user likes "The Godfather," the information that he also likes "The Godfather: Part II" will probably have less predictive power than the information that he also likes "The Notebook."

Using their analytic framework, the LIDS researchers show how to select a small number of products that carry a disproportionate amount of information about users' tastes. If the service provider recommended those products to all its customers, then, based on the resulting ratings, it could much more efficiently sort them into probability clusters, which should improve the quality of its recommendations.

Sujay Sanghavi, an associate professor of electrical and computer engineering at the University of Texas at Austin, considers this the most interesting aspect of the research. "If you do some kind of collaborative filtering, two things are happening," he says. "I'm getting value from it as a user, but other people are getting value, too. Potentially, there is a tradeoff between these things. If there's a popular movie, you can easily show that I'll like it, but it won't improve the recommendations for other people."

That trade-off, Sanghavi says, "has been looked at in an empirical context, but there's been nothing that's principled. To me, what is appealing about this paper is that they have a principled look at this issue, which no other work has done. They've found a new kind of problem. They are looking at a new issue."



This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Model for evaluating product-recommendation algorithms suggests that trial and error get it right (2014, November 14) retrieved 2 May 2024 from <u>https://phys.org/news/2014-11-product-recommendation-algorithms-trial-error.html</u>

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.