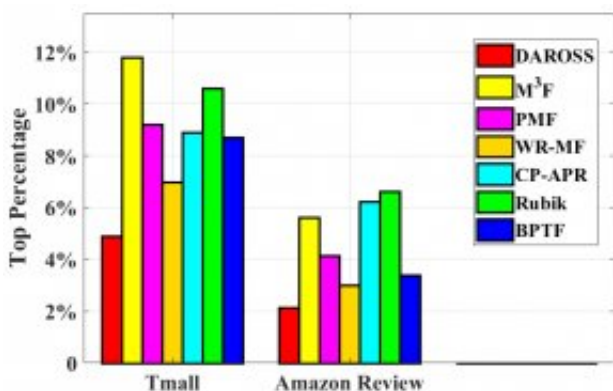


New AI algorithm recommends right products at the right time

December 8 2017



Prediction performance on real-world Tmall and Amazon Review datasets. With 1 being best, the IBM algorithm noted in red and referred to here as DAROSS, most accurately predicted categories. Credit: IBM

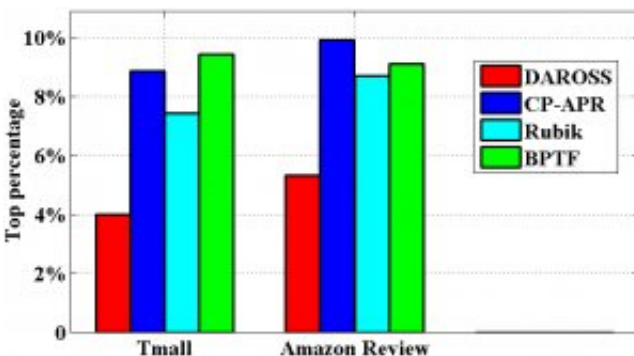
Consumers everywhere are exposed to AI algorithms that recommend products to them based on their past purchases and those of others. Of course, they don't always hit the mark. My IBM Research AI colleagues, our academic collaborators and I recently developed a new algorithm able to provide more accurate, timely product recommendations.

In a much-discussed, wide-ranging interview with IEEE Spectrum, UC Berkeley professor Michael Jordan summarized the situation this way: "If I buy a refrigerator, that doesn't show that I am interested in refrigerators in general. I've already bought my refrigerator, and I'm

probably not likely to still be interested in them. Whereas if I buy a song by Taylor Swift, I'm more likely to buy more songs by her. That has to do with the specific semantics of singers and products and items. To get that right across the wide spectrum of human interests requires a large amount of data and a large amount of engineering."

We have developed a recommendation algorithm that gets it right by going beyond just assuming that similar users like similar items (the common assumption in collaborative filtering-based recommenders). In our approach, we decompose the likelihood of a user purchasing an item into 'form utility' (its intrinsic appeal) and 'time utility' (whether it is the right time to purchase the item). Previous approaches tend to focus only on form utility. We detail the approach in a paper appearing at NIPS 2017 titled "Scalable Demand-Aware Recommendation." The work was led by Jinfeng Yi at IBM Research with contributions from Cho-Jui Hsieh (a professor at UC Davis who completed a summer internship at IBM Research) and his student Yao Li, as well as Lijun Zhang and me.

We quantify time utility via time between purchases within categories of items. Durable goods like refrigerators or televisions tend to have long inter-purchase durations, whereas non-durable goods like Taylor Swift songs or loafs of bread tend to have short or negligible inter-purchase durations.



With 0 being best, DAROSS most accurately predicted purchase time. Credit: IBM

Our AI algorithm infers the durations of all of the categories along with the typical form utilities. Moreover, the algorithm is able to work with purchase data (as opposed to ratings data common in media recommendation) in which the lack of a purchase does not necessarily indicate a dislike. In comparison to six other algorithms, our algorithm provides superior purchase predictions on two real-world datasets: the Amazon Product Data set and a subset of the Repeat Buyers Prediction Competition dataset from IJCAI 2015, in terms of category prediction accuracy and purchase time prediction accuracy. Our algorithm can solve problems with millions of users and millions of items efficiently.

The main technical challenges arise from the problem being posed as tensor completion (filling in missing values in a three-dimensional array) with an objective containing nested hinge losses (the composition of a specific non-smooth function with itself). We significantly reduce computational complexity by converting the tensor completion to matrix completion (two-dimensional array) by assuming that users' preferences for items do not change over time. (In other past work of mine, we focus precisely on the issue of time-varying user preferences.)

Even with this assumption, we still have a difficult problem with a non-smooth objective containing a huge number of terms (equal to the product of the number of users, number of items, and number of time steps). We address these challenges with an alternating minimization that takes advantage of the sparsity of [purchase](#) data. Specifically, among an entire large collection of user-item pairs, a very small fraction constitute purchases. The algorithm cleverly focuses mainly on the purchases, not on the non-purchases. Additionally, it does not attempt to infer both

form utility and time utility at once, but in separate steps.

Due to its accuracy and scalability, we can imagine that our [algorithm](#) would be useful to e-commerce providers.

Provided by IBM

Citation: New AI algorithm recommends right products at the right time (2017, December 8) retrieved 26 April 2024 from <https://phys.org/news/2017-12-ai-algorithm-products.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.