

Why AI Isn't Providing Better Product Recommendations

Originally published at <https://www.unite.ai/why-ai-isnt-providing-better-product-recommendations/>



If you're interested in obscure things, there are two reasons why your searches for items and products are likely to be less related to your interests than those of your 'mainstream' peers; either you're a monetization 'edge case' whose interests will only be catered to if you're also in the upper categories of economic purchasing power (for example, products and services related to '*wealth management*'); or the search algorithms that you're using are leveraging [collaborative filtering](#) (CF), which favors the interests of the majority.

Since collaborative filtering is cheaper and more established than other potentially more capable algorithms and frameworks, it's possible that both these cases apply.

CF-based search results will prioritize items that are perceived to be popular among 'people like you', as best the host framework can understand what kind of a consumer you are.

If you're wary of providing data profiling information to the host system – for instance, not inclined to press the 'Like' buttons in Netflix and other video content services – you're likely to be classified quite generically in your earliest interactions with the system, and the recommendations you receive will reflect the most popular trends.


On a streaming platform, that could mean being recommended whatever shows and movies are currently 'hot', such as reality TV and forensic murder documentaries, irrespective of your interest in these. Likewise for book recommendation platforms, which will tend to proffer current and recent best-sellers, apparently arbitrarily.

In theory, even data-circumspect users should eventually get better results from such systems based on the way that they use them and the things that they search for, since most search frameworks give users limited ability to edit their usage history.

Any You Like, so Long as It's Black

However, according to a new study from Austria, the ascendancy of collaborative filtering over *content-based filtering* (which seeks to define relationships between products instead of just taking aggregate popularity into account), and other alternative approaches, inclines search systems towards long term *popularity bias*, where obviously popular results are pushed towards end users that are unlikely to be enthused by them.

The paper finds that users who are uninterested in popular items receive 'significantly worse' recommendations than users with medium or high interest in popularity, and (perhaps tautologically) that popular items are recommended more frequently than unpopular items. The researchers also conclude that users with low interest in popular items tend to have larger user profiles that could potentially improve recommender systems – if only the systems could kick their addiction to 'herd' metrics.



Comparing popularity to complexity of user profiles shows that ‘marginal’ users disinterested in mainstream content actually have more potential content for recommender systems to mine; but since such users don’t conform to trends, it seems to be a lost opportunity. Source: <https://arxiv.org/pdf/2203.00376.pdf>

The [paper](#) is titled *Popularity Bias in Collaborative Filtering-Based Multimedia Recommender Systems*, and comes from researchers at now-Center GmbH in Graz, and the Graz University of Technology.

Domains Covered

Building on prior works that studied individual sectors (such as book recommendations), the new paper examines four domains: digital books (via the BookCrossing dataset); movies (via [MovieLens](#)); music (via [Last.fm](#)); and animes (via [MyAnimeList](#)).

The study applied four popular multimedia recommender systems (MMRS) collaborative filtering algorithms against datasets [split](#) into three user groups, according to their inclination to be receptive to ‘popular’ results: *LowPop*, *MedPop*, and *HighPop*. The user groups were filtered down to 1000 equal size groups, based on least, average, and most likely to favor ‘popular’ results.

Commenting on the results, the authors state:


‘[We] find that the probability of a multimedia item to be recommended strongly correlates with this items’ popularity [and] that users with less inclination to popularity (LowPop) receive statistically significantly worse multimedia recommendations than users with medium (MedPop) and high (HighPop) inclination to popular items...’

‘Our results demonstrate that although users with little interest into popular items tend to have the largest user profiles, they receive the lowest recommendation accuracy. Hence, future research is needed to mitigate popularity bias in MMRS, both on the item and the user level.’

Among the algorithms evaluated were two [K-Nearest Neighbors](#) (KNN) variants, [UserKNN](#) and UserKNNAvg. The first of these does not generate an average rating for the target user and item. A [non-negative matrix factorization](#) variant (NMF) was also tested, along with a CoClustering algorithm.

The evaluation protocol considered the recommendation task as a prediction challenge, measured by the researchers in terms of mean absolute error (MAE), against a five-fold cross validation protocol that exceeds the usual 80/20 split between trained and test data.

The results indicate a near-guarantee of popularity bias under collaborative filtering. The question, arguably, is whether this is perceived as a problem by the multi-billion dollar companies currently incorporating CF into their search algorithms.




Across all four datasets studied over four popular Collaborative Filtering recommendations, each result indicates that popular media items are more likely to be recommended than unpopular offerings.

The 'Easy' Way Out

Though collaborative filtering is increasingly used as only one plank of a broader search algorithm strategy, it has a strong stake in the search sector, and its logic and potential profitability is attractively easy to understand.

In itself, CF essentially offloads the task of evaluating content value to end users, and uses their uptake of the content as an index of its value and potential attractiveness to other customers. By analogy, it's essentially a map of 'water cooler buzz'.

Content-based filtering (CBF) is more difficult, but could potentially provide more relevant results. In the computer vision sector, an increasing amount of research is currently being expended on categorizing video content and [attempting to derive](#) domains, features, and high level concepts through analysis of audio and video in movie and TV output.



One of many research projects over the last five years that attempts to derive semantic features from the content of movies, in order to generate more intelligent 'adjacent' recommendations. Source: <https://arxiv.org/pdf/1701.00199.pdf>

However, this is a relatively nascent pursuit, and bound up in the current, more general struggle to quantify, isolate and exploit high level concepts and features in domain knowledge.

Who Uses Collaborative Filtering?

At the time of writing, Netflix's [oft-criticized](#) recommendation engine remains fixated on various collaborative filtering approaches, applying a variety of [adjunct technologies](#) in ongoing attempts to generate more user-relevant recommendations.

Amazon's search engine [evolved](#) from its early adoption of user-based collaborative filtering to an item-item collaborative filtering method, which places greater emphasis on the customer's purchase history. Naturally, this can lead to different types of inaccuracy, such as [filter bubbles](#), or over-emphasis on sparse data. In the latter case, if an infrequent Amazon customer makes an 'unusual' purchase, such as a set of operettas for an opera-loving friend, there may not be adequate alternative purchases that reflect the customer's own preferences to stop this purchase from becoming an influence on their own recommendations.

Collaborative filtering is also extensively used by Facebook, [in concert with other approaches](#), and also by [LinkedIn](#), YouTube, and Twitter.

First published 2nd March 2022.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



[Discover More](#)