

Rating, ranking and recommending—three R's for the internet age

December 20 2016, by Christopher Brinton And Mung Chiang



Credit: AI-generated image (disclaimer)

This holiday season, when we Google for the most trending gifts, compare different items on Amazon or take a break to watch a holiday movie on Netflix, we are making use of what might be called "the three R's" of the Internet Age: rating, ranking and recommending.



Much like the traditional "three R's" of education – "<u>reading, 'riting and</u> <u>'rithmetic</u>" – no modern education is complete without understanding how websites' algorithms combine, process and synthesize information before presenting it to us.

As we explore in our new book, <u>"The Power of Networks: Six Principles</u> <u>that Connect Our Lives</u>," the three tasks of rating, ranking and recommending are interdependent, though it may not be initially obvious. Before we can rank a set of items, we need some measure by which they can be ordered. This is really a rating of each item's quality according to some criterion.

With ranked lists in hand, we may turn around and make recommendations about specific items to people who may be interested in purchasing them. This interrelationship highlights the importance of how the quality and attractiveness of an item is quantified into a rating in the first place.

Ranking

What consumers and internet users often call "rating," tech companies may call "scoring." This is key to, for example, how Google's search engine returns high-quality links at the top of its <u>search results</u>, with the most relevant information usually contained in the first page of responses. When a person enters a search query, Google assigns two main scores to each page in its <u>database of trillions</u>, and uses these to generate the order for its results.

The first of these scores is a "relevance score," a combination of <u>dozens</u> <u>of factors</u> that measure how closely related the page and its content are to the query. For example, it takes into account how prominently placed search keywords are on the result page. The second is an "importance score," which captures the way the network of webpages are connected



to one another via hyperlinks to quantify how important each page is.

The combination of these two scores, along with other information, gives a rating for each page, quantifying how useful it might be to the end user. Higher ratings will be placed toward the top of the search results. These are the pages Google is implicitly recommending that the user visit.

Rating

The three Rs also pervade online retail. Amazon and other e-commerce sites <u>allow customers to enter reviews</u> for products they have purchased. The star ratings contained in these reviews are usually aggregated into a single number representing customers' overall opinion. The principle behind this is called "<u>the wisdom of crowds</u>," the assumption that combining many independent opinions will be more reflective of reality than any single individual's evaluation.

Key to the wisdom of crowds is that the reviews accurately reflect customers' experiences, and are not biased or influenced by, say, the manufacturer adding a series of positive assessments to its own items. Amazon has mechanisms in place to screen out these sorts of reviews – for example, by requiring a purchase to have been made from a given account before it can <u>submit a review</u>. Amazon then averages the star ratings for the reviews that remain.

Averaging ratings is fairly straightforward. But it's more complicated to figure out how to effectively rank products based on those ratings. For example, is an item that has 4.0 stars based on 200 reviews better than one that has 4.5 stars but only from 20 reviews? Both the average rating and sample size need to be accounted for in the ranking score.

There are even more factors that may be taken into consideration, such



as reviewer reputation (ratings based on reviewers with higher reputations may be trusted more) and rating disparity (products with widely varying ratings may be demoted in the ordering). Amazon may also present products to different users in <u>varying orders</u> based on their browsing history and records of previous purchases on the site.

Recommending

The prime example of recommendation systems is <u>Netflix's method for</u> <u>determining which movies a user will enjoy</u>. Algorithms predict how each specific user would rate different movies she has not yet seen by looking at the past history of her own ratings and comparing them with those of similar users. The movies with the highest predictions are those that will then make the final cut for a particular user.

The quality of these recommendations depends heavily on the algorithm's accuracy and its use of machine learning, data mining and the data itself. The more ratings we start with for each user and each movie, the better we can expect the predictions to be.

A simple rating predictor might assign one parameter to each user that captures how lenient or harsh a critic she tends to be. Another parameter might be assigned to each movie, capturing how well-received the movie is relative to others. More sophisticated models will <u>identify similarities</u> among users and movies – so if people who like the kinds of movies you like have given a high rating to a movie you haven't seen, the system might suggest you'll like it too.

This can involve hidden dimensions that underlie user preferences and movie characteristics. It can also involve measuring how the ratings for any given movie have changed over time. If a previously unknown film becomes a cult classic, it might start appearing more in people's recommendation lists. A key aspect of dealing with several models is



combining and tuning them effectively: The algorithm that won the <u>Netflix Prize competition</u> of predicting movie ratings in 2009, for example, was a blend of hundreds of individual algorithms.

This combination of rating, ranking and recommendation algorithms has transformed our daily online activities, far beyond shopping, searching and entertainment. Their interconnection brings us clearer – and sometimes unexpected – insights into what we want and how we get it.

This article was originally published on <u>The Conversation</u>. Read the <u>original article</u>.

Provided by The Conversation

Citation: Rating, ranking and recommending—three R's for the internet age (2016, December 20) retrieved 1 May 2024 from <u>https://phys.org/news/2016-12-recommendingthree-internet-age.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.