Personalization and Recommendation

Here are some movie ratings.

	Movie 1	Movie 2	Movie 3	Movie 4
Alice	4	4		1
Bob		2	2	3
Carol	1	5	3	
Dennis	3		4	1
Emma	5	2	1	4
Flora	3	1		5

TECHNOLOGY

Google Knows You Better Than You Know Yourself

Predictive analysis combs through calendars and search histories—and gets in the way of routine self-deception.

JAMES CARMICHAEL AUGUST 19, 2014

Facebook Knows You Better than You Know Yourself

Erman Misirlisoy, PhD Oct 18, 2018 · 7 min read ★

The Internet Knows You Better Than You Know Yourself

When Amazon or eBay recommend us something we like but were not looking for, they effectively know us better than we know ourselves.

4

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

Kashmir Hill Former Staff

Welcome to The Not-So Private Parts where technology & privacy collide

5

https://www.youtube.com/embed/CxjZ7Ikjaqc?enablejsapi=1

What types of data do firms collect?

Data that firms collect

Search history: Almost every search engine collects data on your search histories. Whenever you search on Google, it knows what you are looking for, the items you viewed on Google, the webpages you clicked…

Data that firms collect

Geolocation and device information: A GPS and Wi-Fi chip are installed in every smartphone. The dozens of apps on our phones, most of them free, aren't just serving up information and entertainment. They are collecting and selling your data to digital marketers who will then offer you personalized ads.

Data that firms collect

Purchase histories: Online sellers frequently collect data on your purchase histories. They know what you have purchased and what you have not purchased. As we will figure out later, the purchase history data is the most important data for online sellers.

Data that firms collect

IP addresses: You have to get an IP address to search the internet. You can find your IP address here. Then, even if you are not using a mobile device, firms can still know your geolocation (e.g., your country and even neighborhood).

Google also collects...

Number of email exchanges you've had in Gmail; number of files in Drive; number of photos Google stores for you.

Your location or searches or browsing history: Google Maps keeps track of everywhere you go and when, alongside the photos taken that day and travel times down to the minute.

Your Google Account: your photo and birthdate.

https://www.youtube.com/embed/Oz3b0Qp2Ong?enablejsapi=1

How about traditional firms?

Even the traditional brick-and-mortar (offline) shops are also collecting your data.

- Your payment method (Credit? Mobile pay? Cash?)
- Loyalty program information (Are you using Yuu?)
- Personal profile (If you ever registered there…)

How about traditional firms?

With new technologies, brick-and-mortar stores can also get much more information than what they had before.

- As described in the video, if you use the free Wi-Fi they provide you, they will be able to collect data from your smartphone!
- Facial recognition and mobile payments help collect data from you.

Personalized Pricing

With personalized pricing, a seller offers each consumer an individualized price, and two persons can receive two different prices at the same time.

Note that personalized pricing is different from dynamic pricing. With dynamic pricing, the price is changing over time. For personalized pricing, the price is changing over consumers.

Example of dynamic pricing: Uber adjusts prices timely.

Price Discrimination

Broadly speaking, personalized pricing is a form of price discrimination. Let's review types of price discrimination:

- 1st degree: The firm sells a product at the maximum price that every consumer is willing to pay.
- 2nd degree: price varies according to quantity demanded.
- 3rd degree: charging a different price to different consumer groups.

Price Discrimination

Personalized pricing is close to first-degree price discrimination.

Firms can learn about your income (e.g., from your bank account), your geo-location (e.g., in the US or India), your neighborhood (a high-end one?), your device (iOS or Android), your purchase habits (bargain hunter?), your gender,…

Based on this information, firms can infer how much you are willing to pay for the product and offer you a personalized price.

THE WALL STREET JOURNAL.

English Edition ▼ | Print Edition | Video | Podcasts | Latest Headlines

Home   World   U.S.   Politics   Economy   Business   **Tech**   Markets   Opinion   Life & Arts   Real Estate   WSJ.M

On Orbitz, Mac Users Steered to Pricier Hotels

Are you using a Mac or PC?

On Orbitz, Mac users spend as much as 30% more a night on hotels that PC users do.

The US retailer Office Depots uses customers' browsing history and location data to vary prices

These Brands Have Some of the Best Abandoned Cart Email Strategies

Aug 28, 2019 5:03:58 PM

When you abandon an item from your online shopping cart, e-tailers may issue you a discount to lure you to make a purchase.

Behavior Based Pricing

The more common approach is pricing with consumers' purchase history, a practice known as "behavior-based pricing."

The idea is very simple: The price you receive depends on whether or not you have purchased the products before. In other words, we offer new and existing consumers different prices.

Recommendation is everywhere!

The Importance of Recommendation

- Netflix: 2/3 of the movies watched are recommended.
- Google News: recommendations generate 38% more click-throughs.
- Amazon: 35% sales from recommendations.
- ChoiceStream: 28% of the people would buy more music if they found what they liked.

How to recommend?

A recommendation system must have three inputs:

- **Items** to be recommended: songs, movies, products, restaurants etc. (often many thousands)
- **Users** of the items: watchers, listeners, purchasers, shoppers etc. (often many millions)
- **Feedback** of users on items: 5-star ratings, upvotes/downvotes, clicking "next" or "skipping the ad", purchases or clicks.

Collaboratives Filtering

Collaborative filtering is not something new. We have done it in many places in the past. Here are a few examples:

- Bestseller list for books
- Top 50 music list
- The "recent returns" shelf at libraries

The intuition behind: People's tastes are correlated.

Collaboratives Filtering

However, in the above examples, recommendations are not personalized, i.e., everybody receives the same recommendation. How to make recommendations personalized?

The intuition: If Alice and Bob both like X and Alice also likes Y, then Bob is more likely to like Y, especially when Alice and Bob know each other.

Suppose that you want to recommend a movie to Emma, which movie will you recommend?

	1	2	3	4	5	6
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

User-Based Collaborative Filtering:
The Neighbourhood Method

Step 1: Find all the movies rated by Emma before, we get movies 3 and 6

	1	**2**	**3**	**4**	**5**	**6**
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Step 2: Identify other users that have rated the same movie, we get Bob, Carol, and Dennis

	1	**2**	**3**	**4**	**5**	**6**
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Step 3: Compare the similarity between Emma and her "neighbors" to see who are close to Emma.

	1	**2**	**3**	**4**	**5**	**6**
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Step 4: Select the top k most similar neighbors and use their average ratings to predict Emma's rating.

	1	2	3	4	5	6
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Item-Based Collaborative Filtering

Item-based collaborative filtering

Suppose that we are predicting the who will like movie 5.

Step 1: Who have rated movie 5 before? We get Alice and Carol.

	1	**2**	**3**	**4**	**5**	**6**
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Step 2: Identify other movies that have rated the same users, we get movies 1 and 3.

	1	**2**	**3**	**4**	**5**	**6**
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Step 3: Compare the similarity between movie 5 and its "neighbors" to see which movie is close to movie 5.

	1	2	3	4	5	6
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Step 4: Select the top k most similar neighbors and use their average ratings to predict movie 5's rating.

	1	2	3	4	5	6
Alice	2			4	5	
Bob	5		4			1
Carol			5		2	
Dennis		1		5		4
Emma			4			2
Flora	4	5		1		

Model-based Collaborative Filtering

What did Netflix do to make recommendations?

Issues with this approach:

- It is takes users much effort to complete the questionnaire, and users do not always bother to answer the questions carefully and seriously.
- Sometimes users themselves cannot clearly state their preferences.
- Maybe there are some hidden factors that are not captured by the questionnaire, e.g., the duration of the movie.

The Matrix Factorization Method

Matrix Factorization

Here, we assume that each movie has a number of "latent" or "hidden" factors that affect user preferences. Examples of the factors include the length of the movie, the amount of actions in the movie, the seriousness of the movie, the orientation of the movie for children etc.

The factors are "latent" or "hidden," implying that we do not know which these factors are, and we do not need to know either.

Matrix Factorization

Each user also has his or her own preference for each factor. For instance, some users prefer long movies over short movies, and some users prefer to have more actions in their movie. If we know the preferences of a user and a movie's attribute values, we can match the movie with the user to see whether the user will like the movie.

Matrix Factorization

For instance, suppose that a movie is long and contains a lot of actions. We also know that

- Alice likes short movies and hates action movies,
- Bob prefers long movies and enjoys action movies.

Then, we can predict that Alice will hate the movie and Bob will like the movie.

Mathematically, our model is as follows:

Your rating = Your preference for length $\times$ Movie's length
+Your preference for action $\times$ Movie's amount of action

Suppose that the movie's length is 1 and amount of action is 2. Alice's preference for length is 0.5, for action is 0; Bob's preference for length is 1, for action is 1.5. We can predict:

- Alice's rating: $0.5 \times 1 + 0 \times 2 = 0.5$.
- Bob's rating: $1 \times 1 + 1.5 \times 2 = 4$.

Let's generalize the above discussion. Suppose that Alice, Bob, Carol's preferences for length and action are as follows:

	Length	Action
Alice	0.5	0
Bob	1	1.5
Carol	1.5	0.5

There are two movies, whose length and action values are

	Length	Action
1	1	2
2	0	3

We can multiply the two matrix to get user-movie ratings:

$$\begin{bmatrix} 0.5 & 0 \\ 1 & 1.5 \\ 1.5 & 0.5 \end{bmatrix} \times \begin{bmatrix} 1 & 0 \\ 2 & 3 \end{bmatrix} = \begin{bmatrix} 0.5 & 0 \\ 4 & 4.5 \\ 2.5 & 1.5 \end{bmatrix}$$

In other words, our prediction is as follows.

	Movie 1	**Movie 2**
Alice	0.5	0.0
Bob	4.0	4.5
Carol	2.5	1.5

Overall, if we know the user matrix and the movie matrix, we can multiply the two to get the user-movie rating matrix. The issue is: We do not yet know the user matrix and the movie matrix.

But how to get the user matrix and the movie matrix?

The short answer is, we can guess. We guess different user matrices and movie matrices, and see if the predicted rating matrix is close to the actual rating given by users. When the two are close enough, we can use the two matrices to construct the new rating matrix. This is known as matrix factorization.

We can do better than guessing. There are some advanced statistical methods for estimating the user matrix and movie matrix, but this is beyond the scope of our class. If you are interested, you can search for "stochastic gradient descent."

The matrix factorization algorithm

```r
1  matrix_factorization <- function(R, P, Q, K, steps=5000, alpha=0.0002,
   beta=0.02) {
2    Q <- t(Q)
3    for (step in 1:steps) {
4      for (i in 1:nrow(R)) {
5        for (j in 1:ncol(R)) {
6          if (R[i, j] > 0) {
7            eij <- R[i, j] - sum(P[i,] * Q[,j])
8            for (k in 1:K) {
9              P[i, k] <- P[i, k]+alpha*(2*eij*Q[k, j] - beta*P[i, k])
10             Q[k, j] <- Q[k, j]+alpha*(2*eij*P[i, k] - beta*Q[k, j])
11           }}}}
12   eR <- P %*% Q
13   e <- 0
14   for (i in 1:nrow(R)) {
15     for (j in 1:ncol(R)) {
16       if (R[i, j] > 0) {
17         e <- e + (R[i, j] - sum(P[i,] * Q[,j]))^2
18           for (k in 1:K) {
19           e <- e + (beta/2) * (P[i, k]^2 + Q[k, j]^2)
20           }}}}
21     if (e < 0.001){break}}
22   return(list(P = P, Q = t(Q)))
23 }
```

```r
set.seed(123)
R <- matrix(c(5, 3, 0, 1,
              4, 0, 0, 1,
              1, 1, 0, 5,
              1, 0, 0, 4,
              0, 1, 5, 4,
              2, 1, 3, 0), nrow = 6, ncol = 4, byrow = TRUE)
# N: num of User
N <- nrow(R)
# M: num of Movie
M <- ncol(R)
# Num of Features
K <- 2
P <- matrix(runif(N * K), nrow = N, ncol = K)
Q <- matrix(runif(M * K), nrow = M, ncol = K)
result <- matrix_factorization(R, P, Q, K)
nP <- result$P
nQ <- result$Q
nR <- nP %*% t(nQ)
print(nP)
print(nQ)
print(nR)
```

	Movie 1	Movie 2	Movie 3	Movie 4
Alice	4	4		1
Bob		2	2	3
Carol	1	5	3	
Dennis	3		4	1
Emma	5	2	1	4
Flora	3	1		5

The Cold Start Problem