

# Recommender Systems

Algorithms, Data and Security  
A.Y. 2024/25

**Valeria Cardellini**

Global Governance, 3rd year  
Science and Technology Major

## Recommendations

**Becoming**  
Michelle Obama (Author, Narrator), Random House Audio (Publisher)

**What other items do customers buy after viewing this item?**

- Educated: A Memoir** Audible Audiobook  
Tara Westover  
★★★★☆ 5,929  
\$0.00 Free with Audible trial
- Girl, Stop Apologizing (Audiobook Exclusive Edition): A Shame-Free Plan for Embracing and Achieving Your Goals** Audible Audiobook  
Rachel Hollis  
★★★★☆ 689  
\$0.00 Free with Audible trial
- Where the Crawdads Sing** Audible Audiobook  
Delia Owens  
★★★★☆ 7,377  
\$0.00 Free with Audible trial
- Girl, Wash Your Face: Stop Believing the Lies About Who You Are So You Can Become Who You Were Meant to Be** Audible Audiobook  
Rachel Hollis  
★★★★☆ 9,792  
\$0.00 Free with Audible trial

- how can i visit antarctica
- how can i visit the white house
- how can i visit north korea
- how can i visit canada
- how can i visit cuba

**Recommender Systems: The Textbook** 1st ed. 2016 Edition  
by Charu C. Aggarwal -- (Author)  
★★★★☆ 9 customer reviews

**Look inside**

ISBN-13: 978-3319296579  
ISBN-10: 3319296574  
Why is ISBN important?

Have one to sell? [Sell on Amazon](#)

[Add to List](#)

Share [Facebook](#) [Twitter](#) [LinkedIn](#) [Pinterest](#)

**Customers who bought this item also bought**


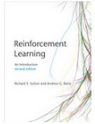
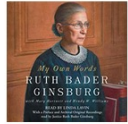
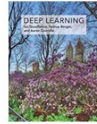




- CAUSALITY**  
Judea Pearl  
Causality: Models, Reasoning and Inference  
- Judea Pearl  
★★★★☆ 38  
Hardcover \$52.99
- HADOOP**  
Hadoop: The Definitive Guide: Storage and Analytics at Internet Scale  
- Tom White  
★★★★☆ 67  
Paperback \$36.06

*"We are leaving the age of information and entering the age of recommendation."* Chris Anderson

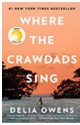
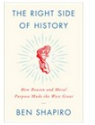



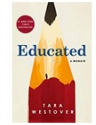

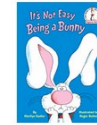
# Recommendations

Your recently viewed items and featured recommendations

Inspired by your browsing history

 <p><b>Statistical Methods for Recommender Systems</b> Deepak K. Agarwal ★★★★☆ 4 Hardcover \$59.29</p>	 <p><b>Reinforcement Learning: An Introduction...</b> Richard S. Sutton ★★★★☆ 13 Hardcover \$50.71</p>	 <p><b>My Own Words</b> Ruth Bader Ginsburg ★★★★☆ 234 Audiobook \$0.00 Free with Audible trial</p>	 <p><b>Deep Learning (Adaptive Computation and...)</b> Ian Goodfellow ★★★★☆ 193 Hardcover \$28.80</p>	 <p><b>BECOMING (German edition): Meine Geschichte</b> Michelle Obama ★★★★☆ 23 Audiobook \$0.00 Free with Audible trial</p>	 <p><b>Hands-On Machine Learning with Scikit-Learn &amp; TensorFlow</b> Aurélien Géron ★★★★☆ 265 Paperback \$56.99</p>	 <p><b>Born a Crime: Stories from a South African Childhood</b> Trevor Noah ★★★★☆ 6,109 Audiobook \$0.00 Free with Audible trial</p>	 <p><b>Machine Learning for Text</b> Charu C. Aggarwal ★★★★☆ 6 Hardcover \$48.61</p>
---	---	---	--	--	--	---	---

Best Sellers

 <p><b>Where the Crawdads Sing</b> Delia Owens ★★★★☆ 7,377 Hardcover \$15.60</p>	 <p><b>The Right Side of History: How Reason and Moral...</b> Ben Shapiro ★★★★☆ 202 Hardcover \$16.79</p>	 <p><b>Supermarket</b> Bobby Hall Paperback \$13.51</p>	 <p><b>Girl, Wash Your Face: Stop Believing the Lies About...</b> Rachel Hollis ★★★★☆ 9,792 Hardcover \$11.80</p>	 <p><b>Raise Your Hand</b> Alice Paul Tapper Hardcover \$11.04</p>	 <p><b>Educated: A Memoir</b> Tara Westover ★★★★☆ 5,929 Hardcover \$16.80</p>	 <p><b>Girl, Stop Apologizing: A Shame-Free Plan for...</b> Rachel Hollis ★★★★☆ 689 Hardcover \$14.77</p>	 <p><b>It's Not Easy Being a Bunny (Beginner Books(R))</b> Marilyn Sadler ★★★★☆ 266 Hardcover \$6.00</p>
---	--	--	--	---	---	--	---

Valeria Cardellini - ADS 2024/25

2

# Recommendations



A **recommender system** is a system able to provide or suggest items to end users

It helps to match users with items

Valeria Cardellini - ADS 2024/25

3

# From scarcity to abundance

---

- Shelf space is a scarce commodity for traditional retailers
  - Also: TV networks, movie theaters,...
- Web enables **near-zero-cost dissemination of information** about products
  - From scarcity to abundance
- More choice necessitates **better filters**
  - Recommender systems

# The power of recommender systems

---

- When does a recommender system do its job well?



# The power of recommender systems

---

- How *Into Thin Air* made *Touching the Void* a bestseller <http://www.wired.com/wired/archive/12.10/tail.html>
- In 1988, a British mountain climber named Joe Simpson wrote a book called *Touching the Void*, a harrowing account of near death in the Peruvian Andes. It got good reviews but, only a modest success, it was soon forgotten. Then, a decade later, a strange thing happened. Jon Krakauer wrote *Into Thin Air*, another book about a mountain-climbing tragedy, which became a publishing sensation. Suddenly *Touching the Void* started to sell again.
- Random House rushed out a new edition to keep up with demand. Booksellers began to promote it next to their *Into Thin Air* displays, and sales rose further. A revised paperback edition, which came out in January, spent 14 weeks on the *New York Times* bestseller list. That same month, IFC Films released a docudrama of the story to critical acclaim. Now *Touching the Void* outsells *Into Thin Air* more than two to one.
- What happened? In short, [Amazon.com recommendations](#). The online bookseller's software noted patterns in buying behavior and suggested that readers who liked *Into Thin Air* would also like *Touching the Void*. People took the suggestion, agreed wholeheartedly, wrote rhapsodic reviews. More sales, more algorithm-fueled recommendations, and the positive feedback loop kicked in.

## Why do we need recommender systems?

---

- Long tail and digital products/services
  - Reduce cognitive load on users
  - Enhance user experience
  - Increase loyalty and volume
  - Introduce quality
  - Help with inventory control
- ... and much more!

# Netflix

---

**NETFLIX**

- Media services provider
- Netflix's initial business included DVD sales and rental by mail
- Current primary business: subscription-based streaming OTT service which offers online streaming of a library of films and television programs, including those produced in-house
  - More than 220 million subscribers
  - Netflix video streaming generated 9.3% of global Internet traffic... (1 out of 10 bits)

[Netflix revenue and usage statistics](#)

## Netflix's prize (2006-2009)

---

- Goal: offer customers accurate movie recommendations so that they find content to watch and enjoy and Netflix maximizes customer satisfaction and retention
- On October 2006, Netflix offered a \$1,000,000 prize to the first developer of a movie-recommendation algorithm that could beat its existing algorithm Cinematch at predicting customer ratings for movies by more than 10% [en.wikipedia.org/wiki/Netflix\\_Prize](https://en.wikipedia.org/wiki/Netflix_Prize)
- Data made available: 100 Million ratings (1999-2005)
  - Ratings came from > 480,000 users and 17,770 movies
- Really difficult: big, sparse and skewed data
  - Many users rated only few movies, and very few users rated huge number of movies (one user rated over 17,000 movies)

## How to measure success?

---

- Which algorithm can find the predicted ratings most similar to the actual ones?
- From most relevant to most tractable metric:
  - Customer satisfaction
  - Prediction effectiveness: how well can we predict the ratings users give to movies they watched?
- **Prediction error**
  - Can be calculated for those  $(u,i)$  pairs (formed by user  $u$  and movie  $i$ ) for which we have both prediction and actual rating (let's say  $C$  such pairs)
  - $r_{ui}$ : actual rating given by user  $u$  to movie  $i$
  - $\hat{r}_{ui}$ : predicted rating for user  $u$  and movie  $i$

## Distance metrics

---

- How much are the prediction and actual rating distant? Use **distance metrics**
- Let's define two distance metrics
  - Hamming distance
  - Root Mean Square Error

## Distance metrics

---

- **Hamming distance**
  - Number of positions in which two binary strings (of equal length) differ, e.g.,
    - 100 and 011 differ in 3 positions:  $\text{dist}_{\text{Hamming}} = 3$
    - 01011 and 11010 differ in 2 positions:  $\text{dist}_{\text{Hamming}} = 2$
  - To apply this distance, we need to transform real values into binary codes

## Distance metrics

---

- **Root Mean Square Error (RMSE)**
  - Square root of the second sample moment of the differences between observed values  $r_{ui}$  and predicted values  $\hat{r}_{ui}$  or the quadratic mean of these differences

$$RMSE = \sqrt{\sum_{(u,i) \in C} \frac{(r_{ui} - \hat{r}_{ui})^2}{|C|}}$$

user  $u$   
movie  $i$

- RMSE is always non-negative, and a value of 0 indicates a perfect fit to data
- The smaller the RMSE, the better the recommender system
- RMSE is the metric used to evaluate solutions for Netflix's prize

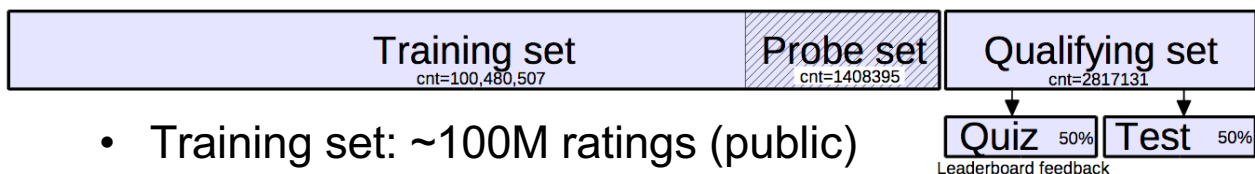
# Which recommendation to users?

---

- Only top few predictions in rank will be recommended to each user
  - E.g., top 5 movies with the highest predicted rating
- Ultimate test is whether user decides to watch recommended movies, and whether she likes them or not

## Netflix's prize dataset

---



- Training set: ~100M ratings (public)
- Probe set: ~1.4M ratings (public)
  - Similar statistical properties to Test and Quiz sets
  - Competitors could use Probe set to test their algorithms
- Quiz set: ~1.4M ratings (hidden)
  - Competitors could submit an algorithm that would run on quiz test, but not more than once a day
- RMSE scores updated on Netflix prize's website were based on performance on quiz set
- Final decision was based on comparison of RMSE on test set
  - Cinematch's RMSE is 0.9514



## Netflix's prize: the contest

---

- Within a week, Cinematch was beaten
- By June 2007, over 20,000 teams had registered from over 150 countries
- By Sep. 2007, team BellKor made 8.26% improvement
- First place changed hands a few times, until after 1 year BellKor got 8.43%
- Teams ranked in the first positions merged and in June 2009, BellKor's Pragmatic Chaos were first to achieve  $> 10\%$
- On test set both BellKor's Pragmatic Chaos and The Ensemble got RMSE equal to 0.8567!

Full story at <https://www.wired.com/2009/09/bellkors-pragmatic-chaos-wins-1-million-netflix-prize/>

## Netflix: Predicting the best user ratings

---

- Netflix was willing to pay over \$1M for the best algorithm, which shows how critical the recommender system was to their business
  - Even though recommendation systems existed before this prize, new research began
- What data can be used to predict Netflix user ratings?
  - Every movie has the rating from users who have ranked it
  - We also know attributes about the movie itself: actors, director, genre classification, year released, etc.

## More in general: input data

---

- What input data could be used to predict user ratings? How input data can be gathered?
- Explicit data: users intentionally provide indications about their preferences, e.g.,
  - Customer ratings
  - Feedback
  - Demographics
- Implicit data: stems from monitoring user behavior, e.g.,
  - Purchase history
  - Click or browse history
- Product information
  - Product taxonomy, attributes, description

## Netflix's prize: Recommendation algorithm


---

- What algorithm used in winning solution?
- Netflix's prize winner is a **cocktail of many methods combined**, with hundreds of ingredient algorithms blended together and thousands of model parameters re-tuned specifically to training set provided by Netflix
  - Overview of used techniques at <https://blog.echen.me/2011/10/24/winning-the-netflix-prize-a-summary/>
- This variety of approaches is a general principle of analytics
  - We will focus on main approaches, big ideas, and few key methodologies

# Approaches to recommender systems


---

Recommender systems reduce information overload by estimating relevance




# Approaches to recommender systems

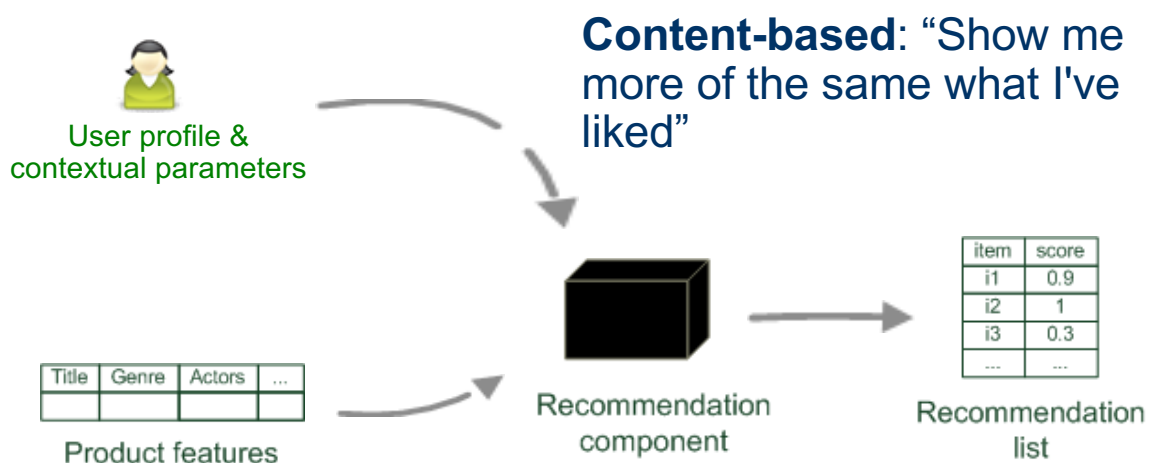
---



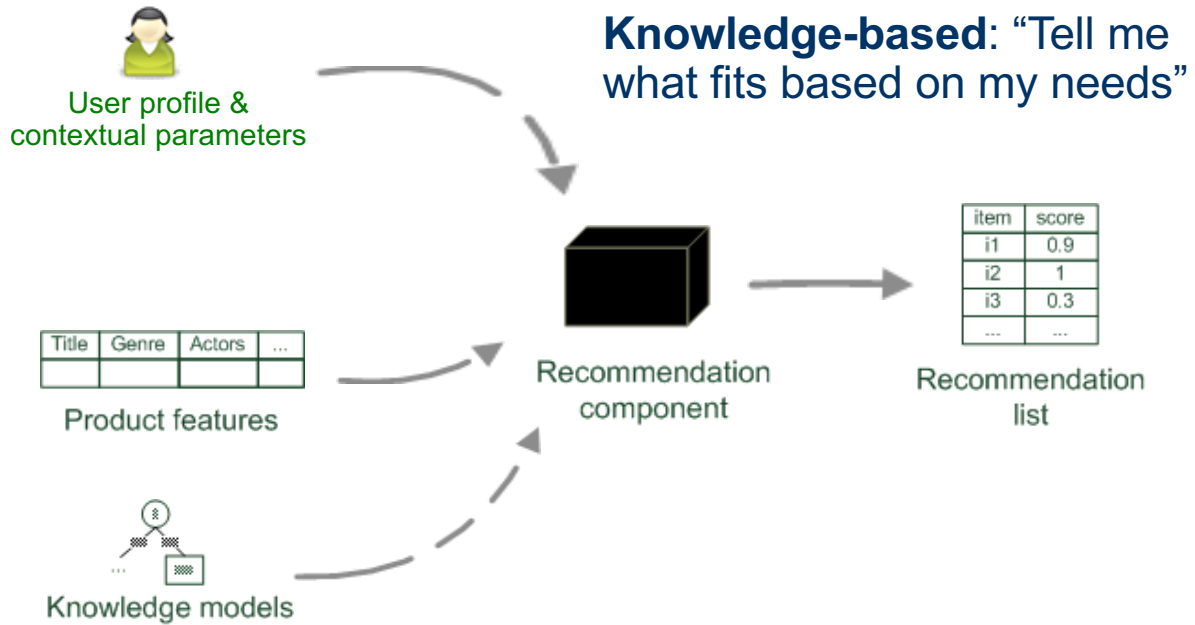
# Approaches to recommender systems



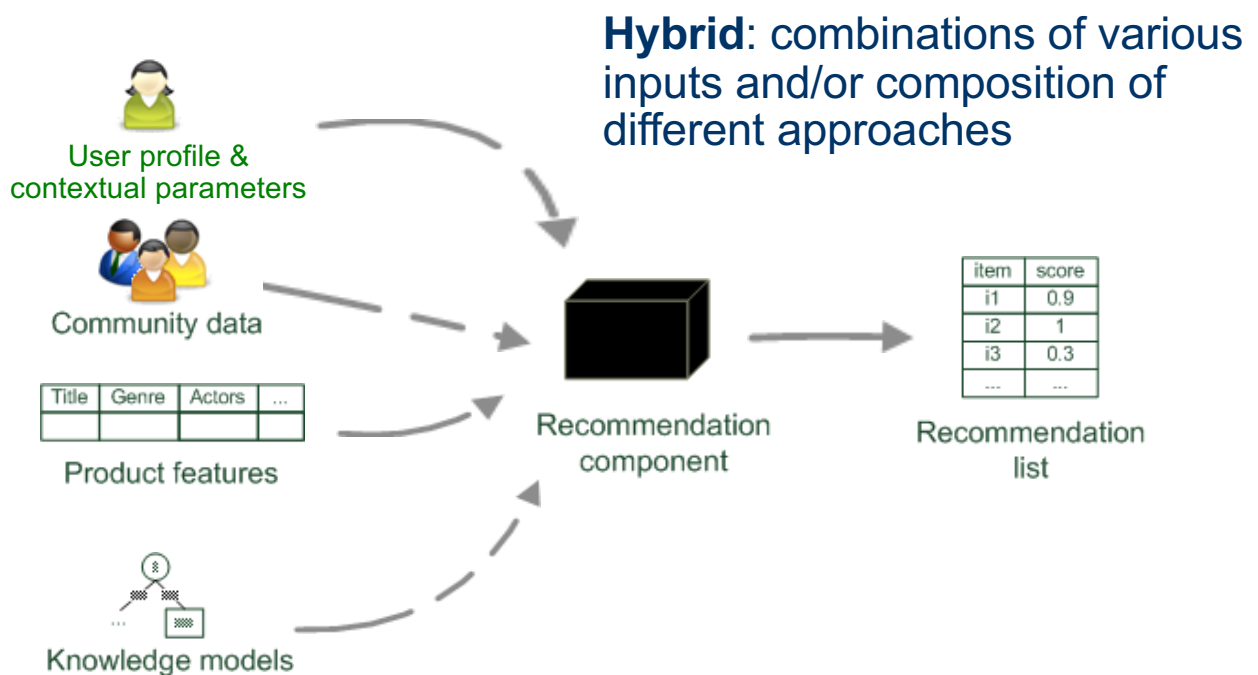
# Approaches to recommender systems



# Approaches to recommender systems



# Approaches to recommender systems



# Paradigms of recommender systems

---

- We analyze:
- Collaborative filtering
  - Make **automatic predictions** (filter) about the interests of a user by collecting **preferences information from many users** (collaborative)
  - Assumption: similar users will have similar ratings
- Content-based filtering
  - Recommend items based on a **comparison** between **content of items** and **user profile**

## Collaborative filtering

---

- Use other users' rankings to make predictions about user's rating for an item, which the user hasn't rated yet
  - These predictions are built upon the known ratings of other users, who have similar ratings with the user
  - E.g., predict Eva's rating for Inception

**Ratings matrix**

	Forrest Gump	Godfather	Inception	Jaws
Amy	5		4	3
Bob	3	5	2	5
Carl		3	5	4
Dan	4	5	4	
Eva	4	4	?	3

Unknown rating we want to predict →

## Collaborative filtering

- Let's calculate mean values for ratings

	Forrest Gump	Godfather	Inception	Jaws	Mean rating
Amy	5		4	3	4.00
Bob	3	5	2	5	3.75
Carl		3	5	4	4.00
Dan	4	5	4		4.33
Eva	4	4		3	3.67

User:  $u$

Movie:  $i$

Rating: 1-5

$n_u$ : number of ratings by user  $u$

Mean rating  $\mu_u = \frac{1}{n_u} \sum_i r_{ui}$

## Collaborative filtering

- Let's measure the similarity between each pair of users
  - We will see later how
  - For now: similarity values range between -1 and 1, where -1 is perfectly dissimilar and 1 is perfectly similar

	Amy	Bob	Carl	Dan	Eva
Amy	1	-0.54	0.00	-0.29	0.87
Bob	-0.54	1	-0.82	0.78	-0.32
Carl	0.00	-0.82	1	-0.87	-0.29
Dan	-0.29	0.78	-0.87	1	0.17
Eva	0.87	-0.32	-0.29	0.17	1

# Collaborative filtering

	Forrest Gump	Godfather	Inception	Jaws	Mean rating
Amy	5		4	3	4.00
Bob	3	5	2	5	3.75
Carl		3	5	4	4.00
Dan	4	5	4		4.33
Eva	4	4		3	3.67

	Amy	Bob	Carl	Dan	Eva
Amy	1	-0.54	0.00	-0.29	0.87
Bob	-0.54	1	-0.82	0.78	-0.32
Carl	0.00	-0.82	1	-0.87	-0.29
Dan	-0.29	0.78	-0.87	1	0.17
Eva	0.87	-0.32	-0.29	0.17	1

- Can predict Eva's rating for Inception:  

$$3.67 + [0.87(4-4) - 0.32(2-3.75) - 0.29(5-4) + 0.17(4-4.33)] / (0.87 - 0.32 - 0.29 + 0.17) = 4.17$$
- Eva will like Inception more than average

## Exploit similarities between users

	Forrest Gump	Godfather	Inception	Jaws	Mean rating
Amy	5		4	3	4.00
Bob	3	5	2	5	3.75
Carl		3	5	4	4.00
Dan	4	5	4		4.33
Eva	4	4		3	3.67


	Amy	Bob	Carl	Dan	Eva
Amy	1	-0.54	0.00	-0.29	0.87
Bob	-0.54	1	-0.82	0.78	-0.32
Carl	0.00	-0.82	1	-0.87	-0.29
Dan	-0.29	0.78	-0.87	1	0.17
Eva	0.87	-0.32	-0.29	0.17	1

- Consider suggesting to Eva "Inception", since Amy rated it more than her average, and Eva and Amy seem to have similar preferences
- This technique is called **user-based collaborative filtering**



# Collaborative filtering: measuring similarity

- How to measure similarity of users?
- We consider **cosine similarity**
  - Cosine of angle  $\theta$  between two user vectors A and B in an  $n$ -dimensional space (figure:  $n=2$ )
  - It measures the similarity in the directions of the two vectors



Valeria Cardellini - ADS 2024/25

32

## Cosine similarity

- We consider **cosine similarity**
  - Cosine of angle  $\theta$  between two vectors A and B in an  $n$ -dimensional space
  - Similarity score ranges between -1 and 1
    - $\cos(0^\circ) = 1$ : perfectly similar
    - $\cos(90^\circ) = 0$ : orthogonal (or uncorrelated)
    - $\cos(180^\circ) = -1$ : perfectly dissimilar
- For a pair of users: the larger the score, the closer their taste

$$\text{cosine similarity} = \cos(\theta) = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}}$$

- $A_i$  and  $B_i$  are components of vector A and B, respectively

Valeria Cardellini - ADS 2024/25


33

## Cosine similarity: Example

- To deal with biases in the ratings (e.g., movie critic who always gives out low ratings), let's **normalize ratings** by subtracting from each rating the average rating of that user

	Forrest Gump	Godfather	Inception	Jaws	Mean rating
Amy	5		4	3	4.00
Bob	3	5	2	5	3.75
Carl		3	5	4	4.00
Dan	4	5	4		4.33
Eva	4	4		3	3.67

	Forrest Gump	Godfather	Inception	Jaws
Amy	5-4=1		4-4=0	3-4=-1
Bob	3-3.75=-0.75	5-3.75=1.25	2-3.75=-1.75	5-3.75=1.25
Carl		3-4=-1	5-4=1	4-4=0
Dan	4-4.33=-0.33	5-4.33=0.67	4-4.33=-0.33	
Eva	4-3.67=0.33	4-3.67=0.33		3-3.67=-0.67



## Cosine similarity: Example

### Normalized ratings matrix

	Forrest Gump	Godfather	Inception	Jaws
Amy	1		0	-1
Bob	-0.75	1.25	-1.75	1.25
Carl		-1	1	0
Dan	-0.33	0.67	-0.33	
Eva	0.33	0.33		-0.67

Note: after normalization, a negative number means below average rating and a positive number means above average ratings given by the same user

- Now compute the cosine similarity for each pair (A, B) of users

$$\text{cosine similarity} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}}$$

- E.g., cosine similarity between Amy and Bob:

$$\frac{1 * (-0.75) + 0 * (-1.75) + (-1) * 1.25}{\sqrt{1^2 + 0^2 + (-1)^2} \sqrt{(-0.75)^2 + 1.25^2 + (-1.75)^2 + 1.25^2}} = \frac{-2}{3.6742} = -0.5443$$

- Note: lack of ratings is considered as 0 ratings

## Cosine similarity: Example

- Now compute the cosine similarity for each pair (A, B) of user

$$\text{cosine similarity} = \frac{\sum_{i=1}^n A_i B_i}{\sqrt{\sum_{i=1}^n A_i^2} \sqrt{\sum_{i=1}^n B_i^2}}$$

- E.g., cosine similarity between Amy and Eva:

$$\frac{1 * (0.33) + (-1) * (-0.67)}{\sqrt{1^2 + 0^2 + (-1)^2} \sqrt{0.33^2 + 0.33^2 + (-0.67)^2}} = \frac{1}{1.1547} = 0.8660$$

- We obtain the following similarity values:

	Amy	Bob	Carl	Dan	Eva
Amy	1	-0.54	0.00	-0.29	0.87
Bob	-0.54	1	-0.82	0.78	-0.32
Carl	0.00	-0.82	1	-0.87	-0.29
Dan	-0.29	0.78	-0.87	1	0.17
Eva	0.87	-0.32	-0.29	0.17	1

## How to predict the rating?

- Given the similarity table, the predicted rating  $\hat{r}_{ui}$  of item  $i$  for user  $u$  is given by:

$$\hat{r}_{ui} = \mu_u + \frac{\sum_{v \neq u} s_{uv} (r_{vi} - \mu_v)}{\sum_{v \neq u} s_{uv}}$$

$s_{uv}$ : similarity between  $u$  and  $v$   
 $\mu_u$ : mean rating for  $u$

	Forrest Gump	Godfather	Inception	Jaws	Mean rating
Amy	5		4	3	4.00
Bob	3	5	2	5	3.75
Carl		3	5	4	4.00
Dan	4	5	4		4.33
Eva	4	4		3	3.67

	Amy	Bob	Carl	Dan	Eva
Amy	1	-0.54	0.00	-0.29	0.87
Bob	-0.54	1	-0.82	0.78	-0.32
Carl	0.00	-0.82	1	-0.87	-0.29
Dan	-0.29	0.78	-0.87	1	0.17
Eva	0.87	-0.32	-0.29	0.17	1

- Eva's rating for Inception:

$$\hat{r}_{Eva, Inception} = 3.67 + \frac{[0.87(4-4) - 0.32(2-3.75) - 0.29(5-4) + 0.17(4-4.33)]}{(0.87 - 0.32 - 0.29 + 0.17)} = 4.1674$$

## More on the predicted rating formula

---

- Given the similarity table, the predicted rating  $\hat{r}_{ui}$  of item  $i$  for user  $u$  is given by:

$$\hat{r}_{ui} = \mu_u + \frac{\sum_{v \neq u} s_{uv}(r_{vi} - \mu_v)}{\sum_{v \neq u} s_{uv}} \quad \begin{array}{l} s_{uv}: \text{similarity between } u \text{ and } v \\ \mu_u: \text{mean rating for } u \end{array}$$

- What does this formula mean?
  - $\mu_u$  is a baseline predictor for  $\hat{r}_{ui}$
  - Weighted average approach: we multiply each normalized rating by a similarity factor (which tells how similar the users are)
  - We add weights to the ratings: the heavier the weight, the more the rating would matter

## Collaborative filtering: all the steps so far

---

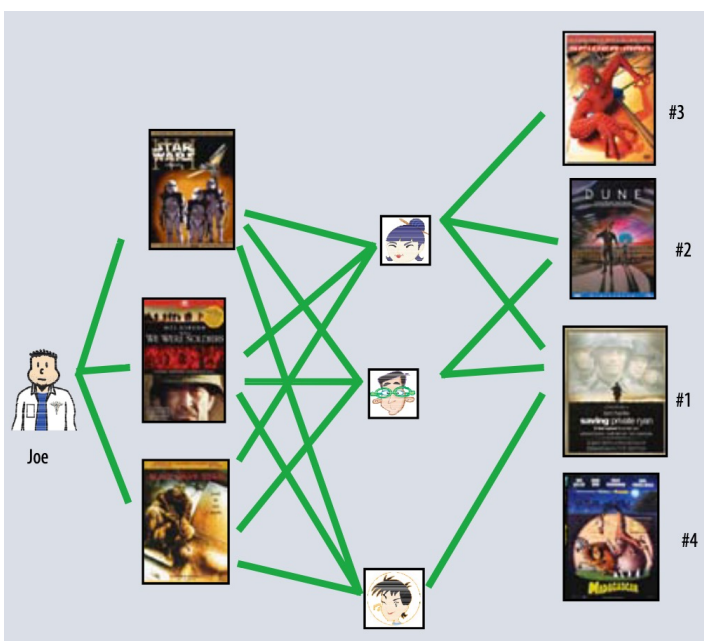
1. Calculate the mean rating for each user
2. Normalize the ratings matrix
3. Compute the cosine similarity for each pair of users
4. Compute the predicted rating(s)

## Collaborative filtering: neighborhood-based

- Real problems are much larger and sparser than our example
- When dealing with large dataset, we do not consider all the users (too expensive!) but only like-minded users that is, users with high similarity (**neighbors** set K) to the user in question
  - E.g., select top-K similar users and use their ratings to recommend movies: weighted average of the ratings given only by **similar users**

$$\hat{r}_{ui} = \mu_u + \frac{\sum_{v \in K} s_{uv}(r_{vi} - \mu_v)}{\sum_{v \in K} s_{uv}}$$

## Collaborative filtering: neighborhood-based



Joe likes the 3 movies on the left.

To make a prediction for him, the recommender system finds similar users who also liked those movies, and then determines which other movies they liked.

In this case, all three liked Saving Private Ryan, so that is the first recommendation. Two of them liked Dune, so that is next, and so on.

## Collaborative filtering: neighborhood-based

---

- Summing up the general approach for neighborhood-based collaborative filtering:
  1. Find the cohort of other similar users who have rated the item in question
  2. Predict the rating from the ratings of similar users

$$\hat{r}_{ui} = \mu_u + \frac{\sum_{v \in K} s_{uv}(r_{vi} - \mu_v)}{\sum_{v \in K} s_{uv}}$$

## Exercise: Collaborative filtering

---

- Applying user-based collaborative filtering, predict Eric's rating for Titanic

	The Matrix	Titanic	Die Hard	Forrest Gump	Wall•E
John	5	1		2	2
Lucy	1	5	2	5	5
Eric	2	?	3	5	4
Diane	4	3	5	3	

# Netflix's prize solution

- Netflix's prize solution uses **matrix factorization**, a **more complex collaborative filtering algorithm** than the user-based one
- Idea: represent users and items in a lower dimensional latent space, exploiting relationships among users and items (that could also be hidden that is, *latent*) to predict ratings
  - For a gentle introduction: watch <https://www.youtube.com/watch?v=ZspR5PZemcs> and read "Recommendation Systems: Collaborative Filtering using Matrix Factorization - Simplified" <https://medium.com/sfu-csmp/recommendation-systems-collaborative-filtering-using-matrix-factorization-simplified-2118f4ef2cd3>
  - For more details: read "Matrix factorization techniques for recommender systems" <https://www.inf.unibz.it/~ricci/ISR/papers/ieeecomputer.pdf>

Valeria Cardellini - ADS 2024/25

44

## Matrix factorization: dependencies

- Let's analyze some example of row and column dependencies
  - Relationship is easy: two very similar movies

	M1	M2	M3	M4	M5
U1	3			3	
U2	1			1	
U3	3			3	
U4	4			4	



M1 = M4



Mall Cop

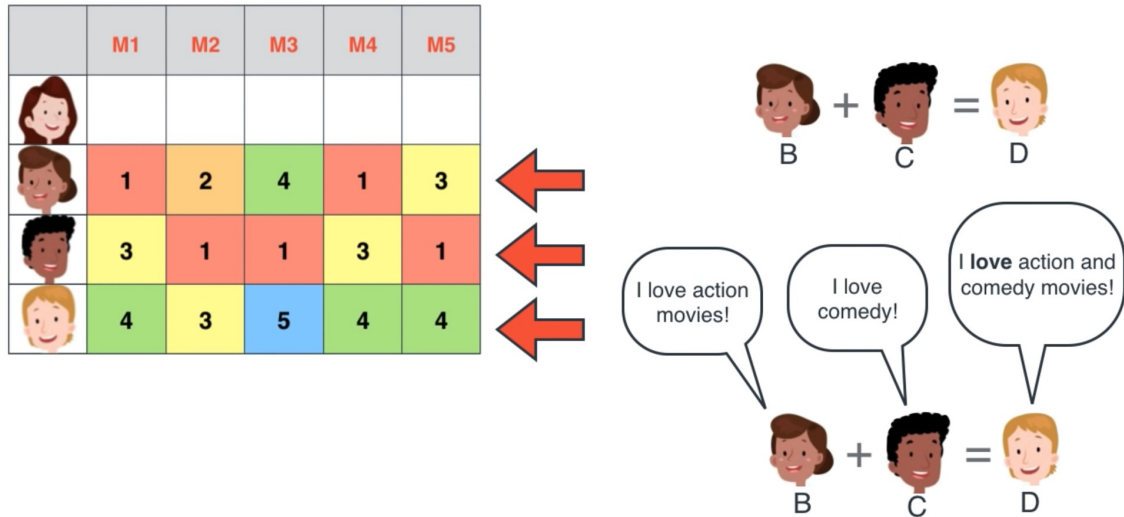


Observe and Report



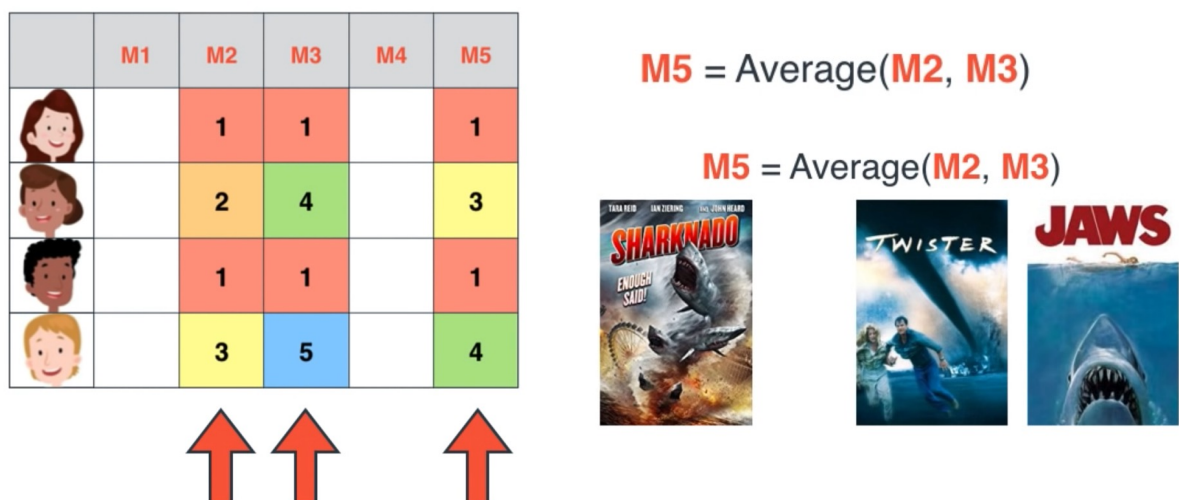
## Matrix factorization: dependencies

- Let's analyze some example of row and column dependencies
  - Relationship is less easy



## Matrix factorization: dependencies

- Let's analyze some example of row and column dependencies
  - Another example of hidden relationship






# Matrix factorization: the idea

- We need a mathematical tool to figure out all the dependencies
- Specifically, we conjecture that the rating matrix  $R$  is actually the **product of two matrices**
- Why “matrix factorization”? We decompose  $R$  into the products of two matrices  $P$  and  $Q$

this  $\times$  that =



	M1	M2	M3	M4	M5
	3	1	1	3	1
	1	2	4	1	3
	3	1	1	3	1
	4	3	5	4	4

- Formally,  $R \approx P \times Q = \hat{R}$







## Matrix factorization





- We need a mathematical tool to figure out all the dependencies

Movie features  $Q$

	M1	M2	M3	M4	M5
 Comedy	3	1	1	3	1
 Action	1	2	4	1	3

User features  $P$

	 Comedy	 Action
 A	✓	✗
 B	✗	✓
 C	✓	✗
 D	✓	✓


	M1	M2	M3	M4	M5
	3	1	1	3	1
	1	2	4	1	3
	3	1	1	3	1
	4	3	5	4	4

Ratings  $R$

## Matrix factorization: dot product

- How to get each value of matrix  $\hat{R}$ ? We use the **dot product**

predicted rating  $\rightarrow \hat{r}_{ij} = \mathbf{p}_i \cdot \mathbf{q}_j = \sum_{k=1}^d p_{ik} q_{kj}$



## Matrix factorization: dot product

- We can use the entries of the product  $P \times Q$  to estimate the corresponding unknown entries (i.e., the predicted ratings) in the matrix  $\hat{R}$

Movie features  $Q$

	M1	M2	M3	M4	M5
F1	3	1	1	3	1
F2	1	2	4	1	3

User features  $P$

	F1	F2
A	1	0
B	0	1
C	1	0
D	1	1

	M1	M2	M3	M4	M5
A	3	?	1		1
B	1		4	1	
C	3	1		3	1
D		3		4	4

## Matrix factorization: dot product

- Let's use the dot product to predict the unknown ratings
- Predict user A rating for movie M2 using the dot product

Movie features Q

F1	3	1	1	3	1
F2	1	2	4	1	3

$$\hat{r}_{12} = (1, 0) \cdot (1, 2) = 1 \times 1 + 0 \times 2 = 1$$

User features P

	F1	F2
A	1	0
B	0	1
C	1	0
D	1	1

	M1	M2	M3	M4	M5
A	3	1	1	3	1
B	1	2	4	1	3
C	3	1	1	3	1
D	4	3	5	4	4

## Matrix factorization

- How do we find P and Q? Let's use Machine Learning
- Idea: the system learns the model to find latent factors by fitting the known ratings
- First initialize the two matrices P and Q with some values, calculate how different their product is to the known ratings and then try to minimize this difference (i.e., error) iteratively
- Such a method is called **gradient descent**, aiming at finding a local minimum of the error

## Matrix factorization

---

- First initialize the two matrices P and Q with some values, calculate how different their product is to the known ratings

$$e_{ij}^2 = (r_{ij} - \hat{r}_{ij})^2 = \left( r_{ij} - \sum_{k=1}^d p_{ik} q_{kj} \right)^2$$

- A good measure of how close the product PQ is to the given rating matrix is the RMSE
- Then try to minimize this difference iteratively by adjusting some element of P or Q

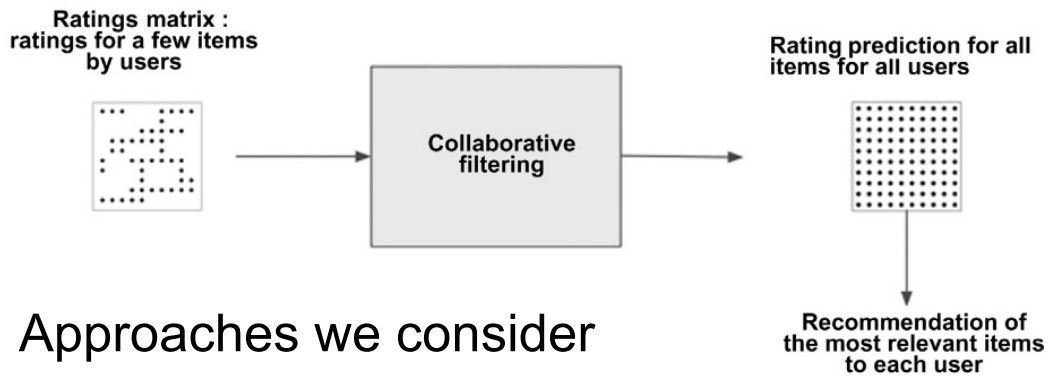
## User-based vs. item-based collaborative filtering

---

- So far, we have considered **user-based collaborative filtering**
- Collaborative filtering can also be **item-based**
  - Based on the assumption that users prefer items that are similar to previously preferred items
  - Compute item similarities based on user ratings
  - How to? Transpose the normalized ratings matrix (swapping rows and columns) and then compute the cosine similarity between each pair of items

# Collaborative filtering: summing up

- Overall picture



- Approaches we consider

