

Recommendation Systems



Big Data Analytics, The Class

Goal: Generalizations
A model or summarization of the data.

Data Frameworks

Hadoop File System ✓
Streaming ✓
MapReduce ✓
Spark ✓
Tensorflow ✓

Algorithms and Analyses

Similarity Search ✓
Link Analysis
Large Scale Hyp. Testing
Recommendation Systems
Deep Learning

Big Data Analytics, The Class

Goal: Generalizations
A model or summarization of the data.

Data Frameworks

Hadoop File System ✓
Streaming ✓
MapReduce ✓
Spark ✓
Tensorflow ✓

Algorithms and Analyses


Similarity Search ✓
Link Analysis ✓
Large Scale Hyp. Testing
Recommendation Systems
Deep Learning

Recommendation Systems



- What other item will this **user** like?
(based on previously liked items)
- How much will user like item X?

Recommendation Systems



- What other item will this **user** like?
(based on previously liked items)

How much will user like item X?

?

Recommendation Systems




- What other item will this **user** like?
(based on previously liked items)

How much will user like item X?



Recommendation Systems



Recommendation Systems



Past User Ratings




Recommendation Systems

Why Big Data?

- Data with many potential features (and sometimes observations)
- An application of techniques for finding similar items
 - locality sensitive hashing
 - dimensionality reduction


Recommendation Systems: Example



- **Customer X**
 - Buys Metallica CD
 - Buys Megadeth CD



- **Customer Y**
 - Does search on Metallica
 - Recommender system suggests Megadeth from data collected about customer X



Examples:

amazon.com.



StumbleUpon



del.icio.us



movie lens

helping you find the *right* movies

last.fm
the social music revolution

Google
News

YouTube

XBOX
LIVE

Origins: Web Shopping

- Does Wal-Mart have everything you need?

Origins: Web Shopping

- Does Wal-Mart have everything you need?



(thelongtail.com)

Origins: Web Shopping


- Does Wal-Mart have everything you need?
- A lot of products are only of interest to a small population (i.e. “[long-tail products](#)”).
- However, most people buy many products that are from the long-tail.
- Web shopping enables more choices
 - Harder to search
 - Recommendation engines to the rescue



(thelongtail.com)

Origins: Web Shopping

- Does Wal-Mart have everything you need?
- A lot of products are only of interest to a small population (i.e. “[long-tail products](#)”).
- However, most people buy many products that are from
- Web shopping
 - Harder to
 - Recomm



Rec Systems Model

Given: *users, items, utility matrix*

Rec Systems Model

Given: *users*, *items*, *utility matrix*



<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
<i>A</i>	4	5	3		3
<i>B</i>	5			4	2
<i>C</i>			5	2	

Rec Systems Model

Given: *users*, *items*, *utility matrix*



<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
<i>A</i>	4	5	3		3
<i>B</i>	5			4	2
<i>C</i>	?	?	5	2	?

Rec Systems Model

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

Rec Systems Model

Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. Content-based
2. Collaborative
3. Latent Factor

Rec Systems Model

Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

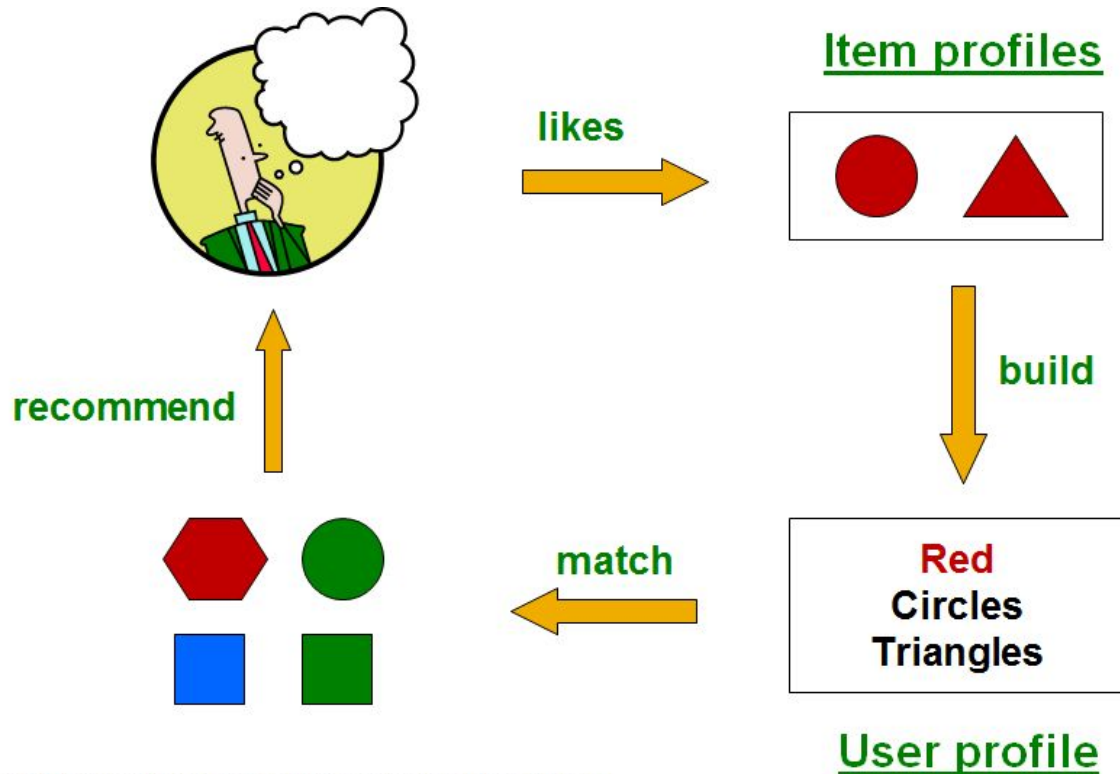
1. **Content-based**
2. Collaborative
3. Latent Factor

Content-Based Rec Systems

Based on similarity of items to past items that they have rated.

Content-Based Rec Systems

Based on similarity of items to past items that they have rated.



Content-Based Rec Systems

Based on similarity of items to past items that they have rated.

1. Build profiles of items (set of features); examples:

shows: producer, actors, theme, review


people: friends, posts

pick words with tf-idf



Content-Based Rec Systems

Based on similarity of items to past items that they have rated.

1. Build profiles of items (set of features); examples:
 - shows*: producer, actors, theme, review
 - people*: friends, posts

pick words with tf-idf
2. Construct user profile from item profiles; approach:
 - average all item profiles of items they've purchased
 - variation: weight by difference from their average

Content-Based Rec Systems

Based on similarity of items to past items that they have rated.

1. Build profiles of items (set of features); examples:

shows: producer, actors, theme, review

people: friends, posts

pick words with tf-idf


2. Construct user profile from item profiles; approach:

average all item profiles of items they've purchased

variation: weight by difference from their average ratings

3. Predict ratings for new items; approach:

find similarity between user and items



Content-Based Rec Systems

Based on similarity of items to past items that they have rated.

1. Build profiles of items (set of features); examples:

shows: producer, actors, theme, review

people: friends, posts

pick words with tf-idf

2. Construct user profile from item profiles; approach:


average all item profiles of items they've purchased

variation: weight by difference from their average ratings

3. Predict ratings for new items; approach:

find similarity between user and items

$$utility(user, i) = \cos(x, i) = \frac{x \cdot i}{||x|| \cdot ||i||}$$



Distance Metrics (for Similarity)


finding *near-neighbors* in *high-dimensional space*

Typical properties of a distance metric, d :

$$d(a, a) = 0$$

$$d(a, b) = d(b, a)$$

$$d(a, b) \leq d(a, c) + d(c, b)$$



Distance Metrics (for Similarity)


finding *near-neighbors* in *high-dimensional space*

There are other metrics of similarity. e.g:

- Euclidean Distance
- Cosine Distance
- ...
- Edit Distance
- Hamming Distance

$$distance(X, Y) = \sqrt{\sum_i^n (x_i - y_i)^2} \quad (\text{"L2 Norm"})$$

$$distance(X, Y) = 1 - \frac{\sum x_i y_i}{\sqrt{\sum x_i^2} \sqrt{\sum y_i^2}}$$



Distance Metrics (for Similarity)

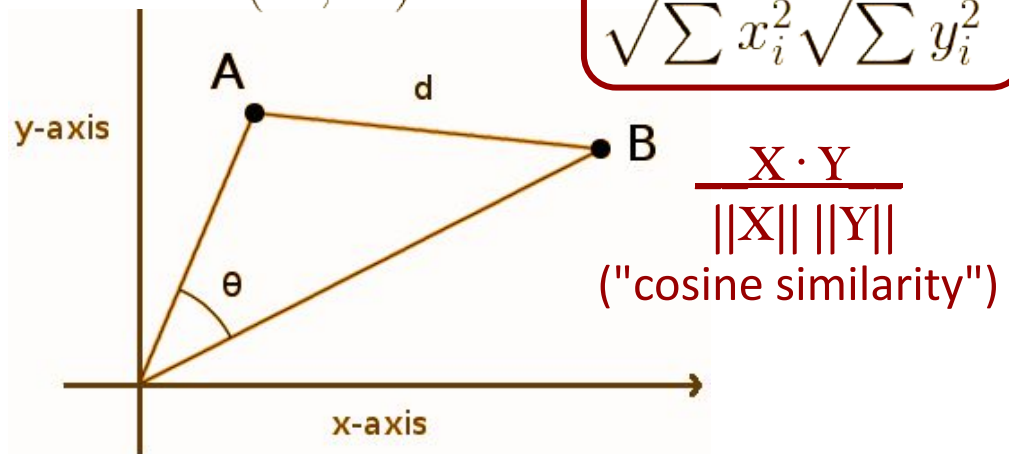
finding *near-neighbors* in *high-dimensional space*

There are other metrics of similarity. e.g:

- Euclidean Distance
- Cosine Distance
- ...
- Edit Distance
- Hamming Distance

$$\text{distance}(X, Y) = \sqrt{\sum_i^n (x_i - y_i)^2} \quad (\text{"L2 Norm"})$$

$$\text{distance}(X, Y) = 1 - \frac{\sum x_i y_i}{\sqrt{\sum x_i^2} \sqrt{\sum y_i^2}}$$



Content-Based Rec Systems

- Only need users history
- Captures unique tastes
- Can recommend new items
- Can provide explanations

Content-Based Rec Systems


- Only need users history
- Captures unique tastes
- Can recommend new items
- Can provide explanations
- Need good features
- New users don't have history
- Doesn't venture "outside the box"
(Overspecialized)

Content-Based Rec Systems

- Only need users history
- Captures unique tastes
- Can recommend new items
- Can provide explanations
- Need good features
- New users don't have history
- Doesn't venture "outside the box"
(Overspecialized)

(not exploiting other users judgments)

Collaborative Filtering



(not exploiting other users judgments)

Rec Systems

Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. **Content-based**
2. Collaborative
3. Latent Factor

Rec Systems

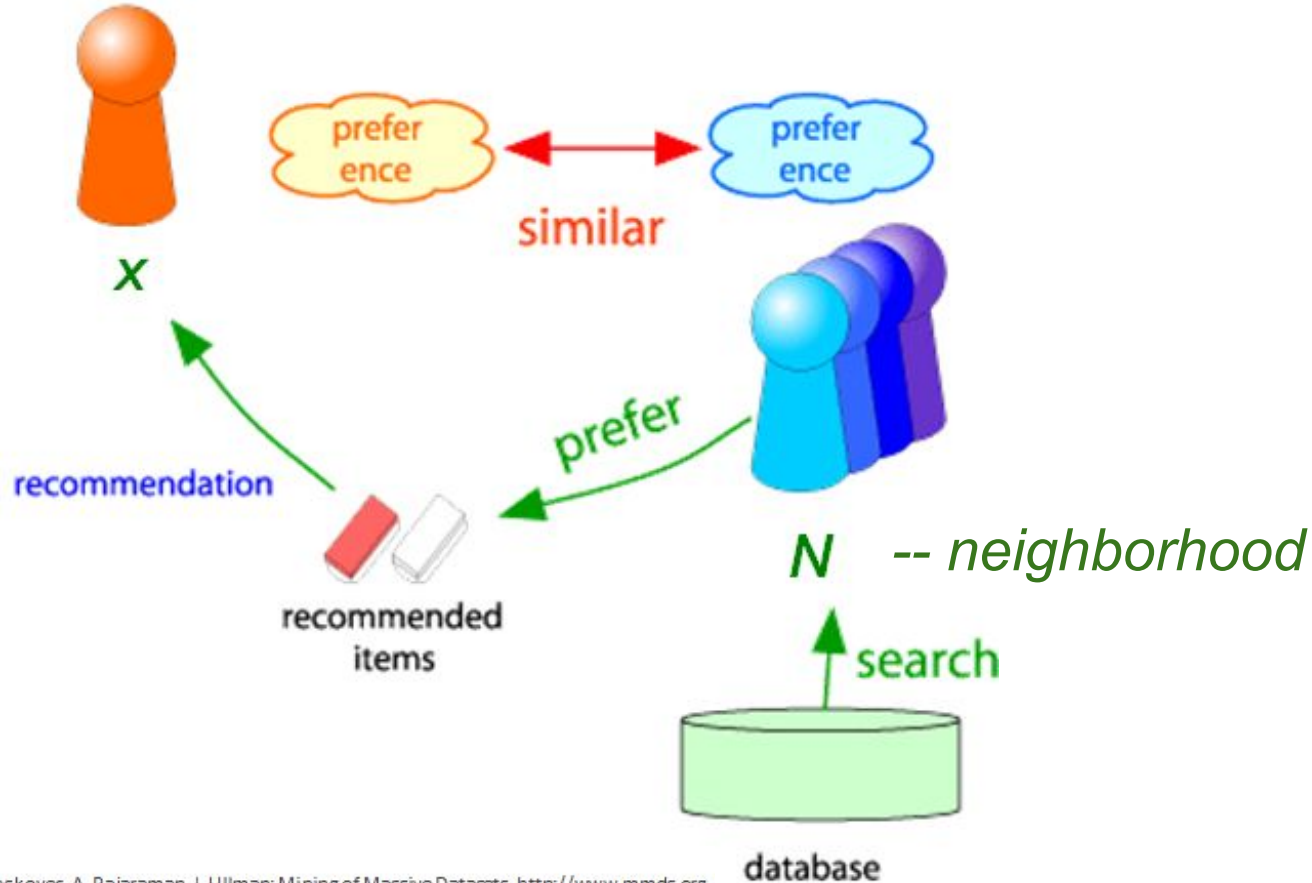
Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. Content-based
2. **Collaborative**
3. Latent Factor

Collaborative Filtering



Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
<i>A</i>	4	5	2		3
<i>B</i>	5			4	2
<i>C</i>			5	2	

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
<i>A</i>	4	5	2		3
<i>B</i>	5			4	2
<i>C</i>			5	2	

General Idea:

- 1) Find similar users = "neighborhood"*
- 2) Infer rating based on how similar users rated*

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
<i>A</i>	4	5	2		3
<i>B</i>	5			4	2
<i>C</i>			5	2	

Given: *user, x; item, i; utility matrix, u*

1. Find neighborhood, N # set of k users most similar to x who have also rated i

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
<i>A</i>	4	5	2		3
<i>B</i>	5			4	2
<i>C</i>			5	2	

Given: *user, x; item, i; utility matrix, u*

1. Find neighborhood, N # set of k users most similar to x who have also rated i

Two Challenges: (1) user bias, (2) missing values

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
A	4 => 0.5	5 => 1.5	2 => -1.5	=> 0	3 => -0.5
B	5			4	2
C			5	2	

Given: *user, x; item, i; utility matrix, u*

1. Find neighborhood, N # set of k users most similar to x who have also rated i

Two Challenges: (1) user bias, (2) missing values

Solution: subtract user's mean, add zeros for missing

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
A	4 => 0.5	5 => 1.5	2 => -1.5	=> 0	3 => -0.5
B	5			4	2
C			5	2	

- Given: *user, x; item, i; utility matrix, u*
0. Update *u*: mean center, missing to 0
 1. Find neighborhood, *N* # set of *k* users most similar to *x* who have also rated *i*
 - $\text{sim}(x, \text{other}) = \text{cosine_sim}(u[x], u[\text{other}])$
 - threshold to top *k* (e.g. *k* = 30)

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
A	4 => 0.5	5 => 1.5	2 => -1.5	=> 0	3 => -0.5
B	5			4	2
C			5	2	

- Given: *user, x; item, i; utility matrix, u*
0. Update *u*: mean center, missing to 0
 1. Find neighborhood, *N* # set of *k* users most similar to *x* who have also rated *i*
 - $\text{sim}(x, \text{other}) = \text{cosine_sim}(u[x], u[\text{other}])$
 - threshold to top *k* (e.g. *k* = 30)
 2. Predict utility (rating) of *i* based on *N*

Collaborative Filtering

<i>user</i>	Game of Thrones	Fargo	Brooklyn Nine-Nine	Silicon Valley	Walking Dead
A	4 => 0.5	5 => 1.5	2 => -1.5	=> 0	3 => -0.5
B	5			4	2
C			5	2	

Given: *user*, *x*; *item*, *i*; *utility matrix*, *u*

0. Update *u*: mean center, missing to 0

1. Find neighborhood, *N* # set of *k* users most similar to *x* who have also rated *i*

-- $\text{sim}(x, \text{other}) = \text{cosine_sim}(u[x], u[\text{other}])$

-- threshold to top *k* (e.g. *k* = 30)

2. Predict utility (rating) of *i* based on *N*

-- average, weighted by sim
$$\text{utility}(x, i) = \frac{\sum_{y \in N} \text{Sim}(x, y) \cdot \text{utility}(y, i)}{\sum_{y \in N} \text{Sim}(x, y)}$$

Collaborative Filtering

“User-User collaborative filtering”



- Given: *user, x; item, i; utility matrix, u*
0. Update *u*: mean center, missing to \emptyset
 1. Find neighborhood, *N* # set of *k* users most similar to *x* who have also rated *i*
 - $\text{sim}(x, \text{other}) = \text{cosine_sim}(u[x], u[\text{other}])$
 - threshold to top *k* (e.g. $k = 30$)
 2. Predict utility (rating) of *i* based on *N*
 - average, weighted by sim
$$\text{utility}(x, i) = \frac{\sum_{y \in N} \text{Sim}(x, y) \cdot \text{utility}(y, i)}{\sum_{y \in N} \text{Sim}(x, y)}$$

Collaborative Filtering

“User-User collaborative filtering”

Item-Item:

Flip rows/columns of utility matrix and use same methods.
(i.e. estimate rating of item i , by finding similar items, j)

Given: *user*, x ; *item*, i ; *utility matrix*, u

0. Update u : mean center, missing to \emptyset
1. Find neighborhood, N # set of k users most similar to x who have also rated i
 - $\text{sim}(x, \text{other}) = \text{cosine_sim}(u[x], u[\text{other}])$
 - threshold to top k (e.g. $k = 30$)
2. Predict utility (rating) of i based on N
 - average, weighted by sim
$$\text{utility}(x, i) = \frac{\sum_{y \in N} \text{Sim}(x, y) \cdot \text{utility}(y, i)}{\sum_{y \in N} \text{Sim}(x, y)}$$

Collaborative Filtering

“User-User collaborative filtering”

Item-Item:

Flip rows/columns of utility matrix and use same methods.
(i.e. estimate rating of item i , by finding similar items, j)

Given: *user*, x ; *item*, i ; *utility matrix*, u

0. Update u : mean center, missing to \emptyset

1. Find neighborhood, N # set of k **items** most similar to i also rated by x

-- $\text{sim}(i, \text{other}) = \text{cosine_sim}(u[i], u[\text{other}])$

-- threshold to top k (e.g. $k = 30$)

2. Predict utility (rating) by x based on N

-- average, weighted by sim
$$\text{utility}(x, i) = \frac{\sum_{j \in N} \text{Sim}(i, j) \cdot \text{utility}(x, j)}{\sum_{j \in N} \text{Sim}(i, j)}$$

item-item vs user-user

Item-item often works better than user-user. Why?

Users tend to be more different from each other than items are from other items.

*e.g. Mary likes jazz + rock, Coleman likes classical + rock,
but Mary may still have same rock preferences as Bob*

item-item vs user-user

Item-item often works better than user-user. Why?

Users tend to be more different from each other than items are from other items.

*e.g. Mary likes jazz + rock, Coleman likes classical + rock,
but Mary may still have same rock preferences as Bob*

In other words, users span genres but items usually do not.


Item-Item: Example

movies

	1	2	3	4	5	6	7	8	9	10	11	12
1	1		3			5			5		4	
2			5	4			4			2	1	3
3	2	4		1	2		3		4	3	5	
4		2	4		5			4			2	
5			4	3	4	2					2	5
6	1		3		3			2			4	



- unknown rating




- rating between 1 to 5

Item-Item: Example

	1	2	3	4	5	6	7	8	9	10	11	12
1	1		3		?	5			5		4	
2			5	4			4			2	1	3
3	2	4		1	2		3		4	3	5	
4		2	4		5			4			2	
5			4	3	4	2					2	5
6	1		3		3			2			4	

movies



- estimate rating of movie **1** by user **5**

Item-Item: Example

	1	2	3	4	5	6	7	8	9	10	11	12	$\text{sim}(1,m)$
1	1		3		?	5			5		4		1.00
2			5	4			4			2	1	3	-0.18
<u>3</u>	2	4		1	2		3		4	3	5		<u>0.41</u>
4		2	4		5			4			2		-0.10
5			4	3	4	2					2	5	-0.31
<u>6</u>	1		3		3			2			4		<u>0.59</u>

movies

Same as cosine sim when subtracting the mean

Neighbor selection:

Identify movies similar to movie 1, rated by user 5

Here we use Pearson correlation as similarity:

1) Subtract mean rating m_i from each movie i

$$m_1 = (1+3+5+5+4)/5 = 3.6$$

row 1: [-2.6, 0, -0.6, 0, 0, 1.4, 0, 0, 1.4, 0, 0.4, 0]

2) Compute cosine similarities between rows

Item-Item: Example

	1	2	3	4	5	6	7	8	9	10	11	12	$\text{sim}(1,m)$
1	1		3		?	5			5		4		1.00
2			5	4			4			2	1	3	-0.18
<u>3</u>	2	4		1	2		3		4	3	5		<u>0.41</u>
4		2	4		5			4			2		-0.10
5			4	3	4	2					2	5	-0.31
<u>6</u>	1		3		3			2			4		<u>0.59</u>

Compute similarity weights:

$$s_{1,3}=0.41, s_{1,6}=0.59$$

Item-Item: Example

	1	2	3	4	5	6	7	8	9	10	11	12	$\text{sim}(1,m)$
1	1		3		?	5			5		4		1.00
2			5	4			4			2	1	3	-0.18
<u>3</u>	2	4		1	2		3		4	3	5		<u>0.41</u>
4		2	4		5			4			2		-0.10
5			4	3	4	2					2	5	-0.31
<u>6</u>	1		3		3			2			4		<u>0.59</u>

$$\text{utility}(1, 5) = (0.41 \cdot 2 + 0.59 \cdot 3) / (0.41 + 0.59)$$

$$\text{utility}(x, i) = \frac{\sum_{j \in N} \text{Sim}(i, j) \cdot \text{utility}(x, j)}{\sum_{j \in N} \text{Sim}(i, j)}$$

Rec Systems

Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. Content-based
2. Collaborative
3. Latent Factor

Rec Systems

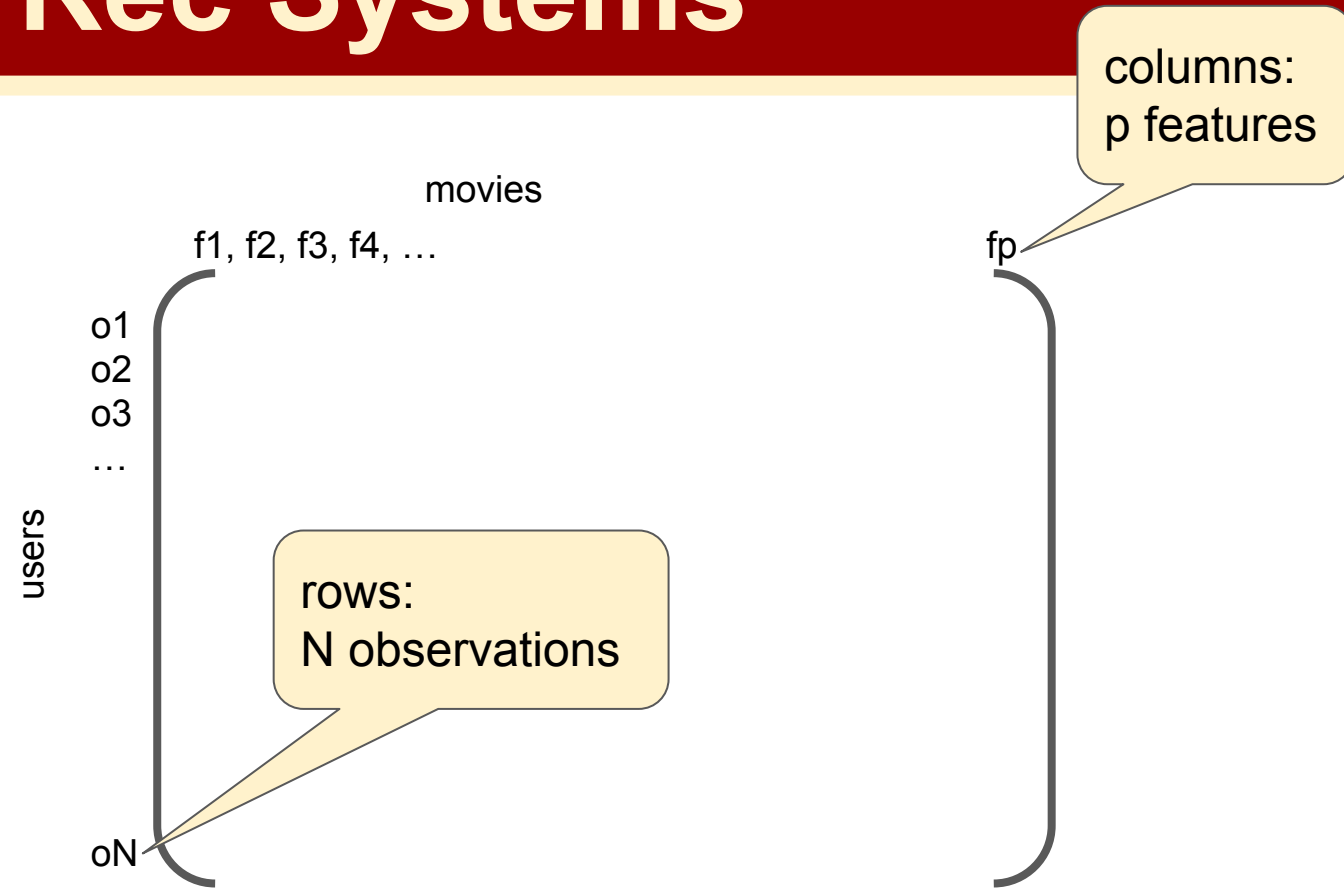
Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. Content-based
2. Collaborative
3. **Latent Factor**

Rec Systems



Rec Systems

Goal: Complete Matrix

movies

f1, f2, f3, f4, ...


fp

o1
o2
o3
..
oN

2	0	3	7	5	0	8	4	2	8	6	4	5	2	3	6	7	3	4	2	9	2	0	9	8	7				
2	5	3	8	9	2	3	1	0	5	4	2	7	8	2	6	9	7	0	4	9	5	3	7	3	9	4	9		
6	7	9	8	3	7	4	7	2	4	1	5	0	4	8	3	6	7	6	6	2	0	4	1	5	4	9	0		
9	5	2	2	0	8	4	9	0	2	6	8	8	1	1	9	4	9	4	5	4	2	4	5	4	8	5	1	7	
6	6	9	4	9	6	0	0	4	4	1	8	8	5	8	2	5	8	3	2	3	5	4	9	6	3	8	1		
4	8	5	0	1	1	9	7	4	3	8	5	4	3	3	5	6	3	5	5	3	3	6	4	2	9	4	3	2	
5	5	8	0	9	3	3	4	1	0	4	9	8	9	6	3	7	6	1	6	4	7	4	0	0	7	9	7	1	
8	7	2	0	3	8	0	2	4	2	7	6	2	6	7	6	3	9	6	8	1	0	7	2	4	0	1	6	1	
7	9	1	4	9	7	1	1	1	2	9	3	0	6	8	2	9	8	2	5	3	1	9	4	0	2	5	5	4	
6	5	6	0	7	4	5	1	7	0	3	3	4	0	2	4	3	7	1	7	4	2	6	8	7	7	1	6	8	
2	2	3	4	4	1	9	4	3	0	4	0	4	7	5	6	3	2	6	9	0	1	9	4	9	6	8	1	2	
7	0	8	8	7	7	6	3	8	9	0	5	4	6	6	9	1	4	3	7	5	8	6	0	4	1	1	4	7	
9	5	2	9	2	8	8	5	2	9	9	8	4	2	5	5	7	9	7	2	2	6	7	2	0	8	6	5	0	9
0	0	5	8	9	8	6	7	6	1	5	3	8	9	4	6	9	5	0	9	6	0	5	4	2	2	3	0	7	
1	7	0	6	9	8	5	9	7	6	6	9	7	2	0	6	6	9	6	3	7	7	1	8	8	1	6	9	9	
1	8	1	1	4	8	9	5	7	8	1	1	5	1	6	3	8	5	4	7	8	1	5	0	9	5	8	6	9	1
0	5	2	3	7	2	3	1	9	2	0	5	6	5	2	4	8	1	8	3	0	1	5	4	6	9	5	9	5	4


Rec Systems

Problem: Given Incomplete Matrix



Rec Systems

Complete Matrix using Latent Factors

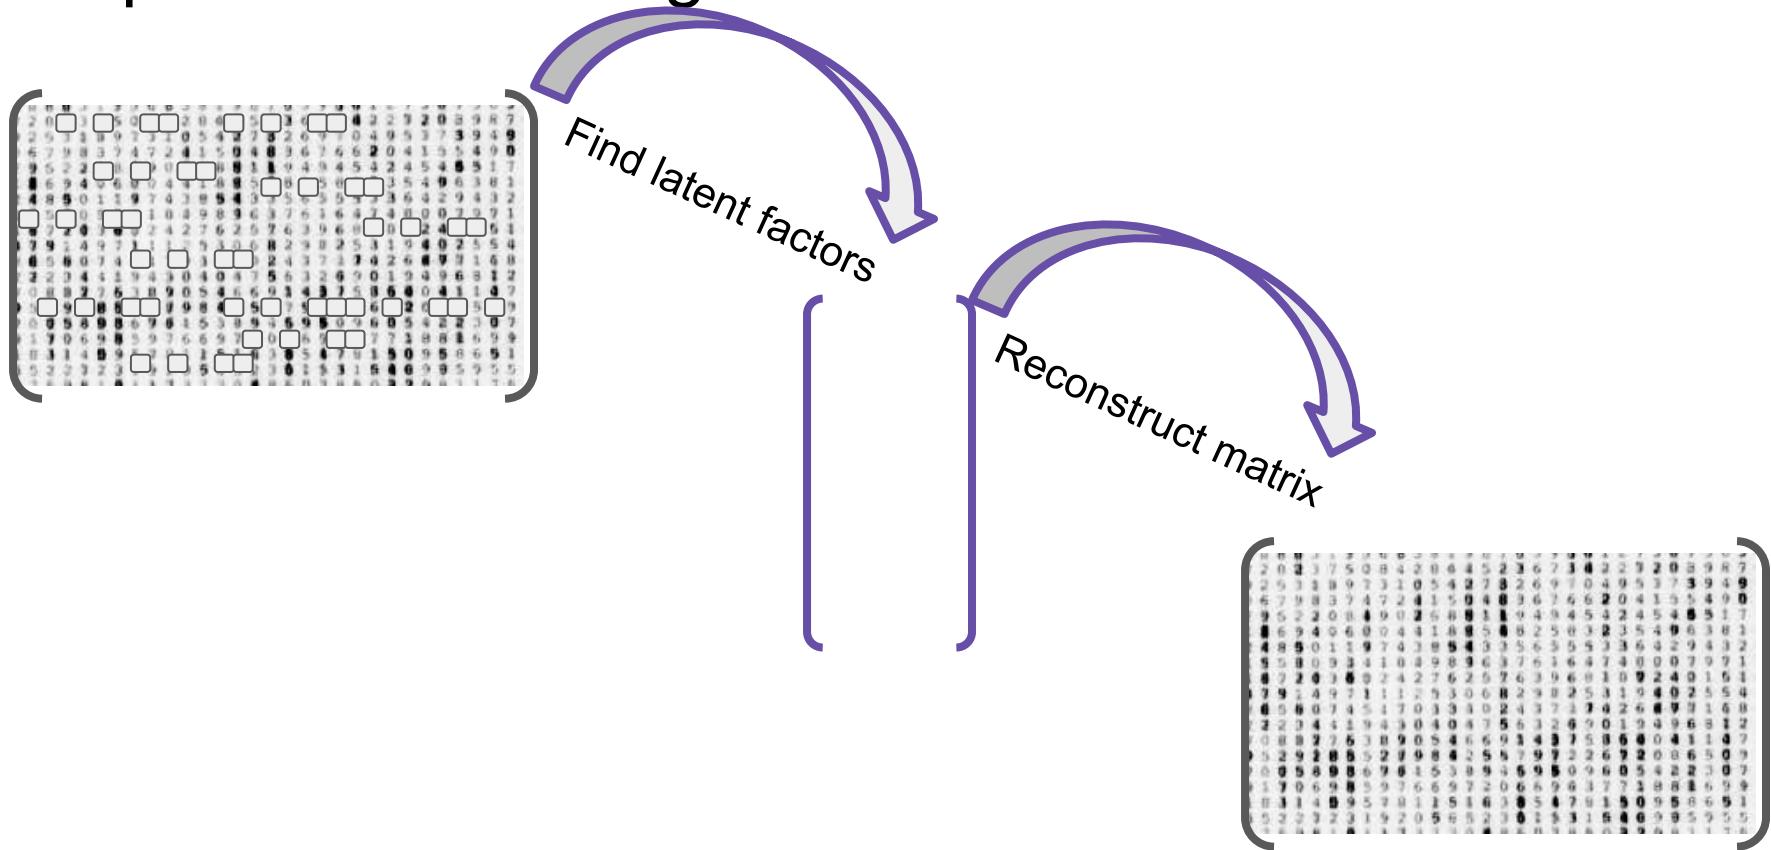


Dimensionality reduction

Try to best represent but with oN p' columns.

Rec Systems

Complete Matrix using Latent Factors



Dimensionality Reduction PCA

Linear approximates of data in r dimensions.

Found via *Singular Value Decomposition*:

$$X_{[n \times p]} = U_{[n \times r]} D_{[r \times r]} V_{[p \times r]}^T$$

X: original matrix,

U: “left singular vectors”,

D: “singular values” (diagonal),

V: “right singular vectors”

Projection (dimensionality reduced space) in 3 dimensions:

$$(U_{[n \times 3]} D_{[3 \times 3]} V_{[p \times 3]}^T)$$

To reduce features in new dataset:

$$X_{\text{new}} V = X_{\text{new_small}}$$

Dimensionality Reduction PCA

Linear approximates of data in r dimensions.

Found via *Singular Value Decomposition*:


$$X_{[n \times p]} = U_{[n \times r]} D_{[r \times r]} V_{[p \times r]}^T$$

X: original matrix,

D: “singular values” (diagonal),

U: “left singular vectors”,

V: “right singular vectors”



Dimensionality Reduction PCA

$$X_{[n \times p]} = U_{[n \times r]} D_{[r \times r]} V_{[p \times r]}^T$$

<div style="display: flex; flex-direction: column; align-items: center; gap: 10px;"> <div style="display: flex; align-items: center; gap: 5px;"> ↑ SciFi </div> <div style="display: flex; align-items: center; gap: 5px;"> ↓ Romnce </div> </div>	$=$	<table border="1" style="border-collapse: collapse; text-align: center; margin-bottom: 10px;"> <thead> <tr> <th style="color: green; writing-mode: vertical-rl; transform: rotate(180deg);">Matrix</th> <th style="color: green; writing-mode: vertical-rl; transform: rotate(180deg);">Alien</th> <th style="color: green; writing-mode: vertical-rl; transform: rotate(180deg);">Serenity</th> <th style="color: green; writing-mode: vertical-rl; transform: rotate(180deg);">Casablanca</th> <th style="color: green; writing-mode: vertical-rl; transform: rotate(180deg);">Amelie</th> </tr> </thead> <tbody> <tr><td>1</td><td>1</td><td>1</td><td>0</td><td>0</td></tr> <tr><td>3</td><td>3</td><td>3</td><td>0</td><td>0</td></tr> <tr><td>4</td><td>4</td><td>4</td><td>0</td><td>0</td></tr> <tr><td>5</td><td>5</td><td>5</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>2</td><td>0</td><td>4</td><td>4</td></tr> <tr><td>0</td><td>0</td><td>0</td><td>5</td><td>5</td></tr> <tr><td>0</td><td>1</td><td>0</td><td>2</td><td>2</td></tr> </tbody> </table>	Matrix	Alien	Serenity	Casablanca	Amelie	1	1	1	0	0	3	3	3	0	0	4	4	4	0	0	5	5	5	0	0	0	2	0	4	4	0	0	0	5	5	0	1	0	2	2	\times	<table border="1" style="border-collapse: collapse; text-align: center; margin-bottom: 10px;"> <tbody> <tr><td>0.13</td><td>0.02</td><td>-0.01</td></tr> <tr><td>0.41</td><td>0.07</td><td>-0.03</td></tr> <tr><td>0.55</td><td>0.09</td><td>-0.04</td></tr> <tr><td>0.68</td><td>0.11</td><td>-0.05</td></tr> <tr><td>0.15</td><td>-0.59</td><td>0.65</td></tr> <tr><td>0.07</td><td>-0.73</td><td>-0.67</td></tr> <tr><td>0.07</td><td>-0.29</td><td>0.32</td></tr> </tbody> </table>	0.13	0.02	-0.01	0.41	0.07	-0.03	0.55	0.09	-0.04	0.68	0.11	-0.05	0.15	-0.59	0.65	0.07	-0.73	-0.67	0.07	-0.29	0.32	\times	<table border="1" style="border-collapse: collapse; text-align: center; margin-bottom: 10px;"> <tbody> <tr><td>12.4</td><td>0</td><td>0</td></tr> <tr><td>0</td><td>9.5</td><td>0</td></tr> <tr><td>0</td><td>0</td><td>1.3</td></tr> </tbody> </table>	12.4	0	0	0	9.5	0	0	0	1.3	\times	<table border="1" style="border-collapse: collapse; text-align: center;"> <tbody> <tr><td>0.56</td><td>0.59</td><td>0.56</td><td>0.09</td><td>0.09</td></tr> <tr><td>0.12</td><td>-0.02</td><td>0.12</td><td>-0.69</td><td>-0.69</td></tr> <tr><td>0.40</td><td>-0.80</td><td>0.40</td><td>0.09</td><td>0.09</td></tr> </tbody> </table>	0.56	0.59	0.56	0.09	0.09	0.12	-0.02	0.12	-0.69	-0.69	0.40	-0.80	0.40	0.09	0.09
Matrix	Alien	Serenity	Casablanca	Amelie																																																																																									
1	1	1	0	0																																																																																									
3	3	3	0	0																																																																																									
4	4	4	0	0																																																																																									
5	5	5	0	0																																																																																									
0	2	0	4	4																																																																																									
0	0	0	5	5																																																																																									
0	1	0	2	2																																																																																									
0.13	0.02	-0.01																																																																																											
0.41	0.07	-0.03																																																																																											
0.55	0.09	-0.04																																																																																											
0.68	0.11	-0.05																																																																																											
0.15	-0.59	0.65																																																																																											
0.07	-0.73	-0.67																																																																																											
0.07	-0.29	0.32																																																																																											
12.4	0	0																																																																																											
0	9.5	0																																																																																											
0	0	1.3																																																																																											
0.56	0.59	0.56	0.09	0.09																																																																																									
0.12	-0.02	0.12	-0.69	-0.69																																																																																									
0.40	-0.80	0.40	0.09	0.09																																																																																									

Users to movies matrix

Dimensionality Reduction PCA

Linear approximates of data in r dimensions.

Found via *Singular Value Decomposition*:


$$X_{[n \times p]} = U_{[n \times r]} D_{[r \times r]} V_{[p \times r]}^T$$

X: original matrix,

D: “singular values” (diagonal),

U: “left singular vectors”,

V: “right singular vectors”



Dimensionality Reduction PCA

- Goal: Minimize the sum of reconstruction errors:

$$\sum_{i=1}^N \sum_{j=1}^D \|x_{ij} - z_{ij}\|^2$$

X: original matrix
D: “singular values”


- where x_{ij} are the “old” and z_{ij} are the “new” coordinates
- “singular vectors”,
“singular vectors”

To check how well the original matrix can be reproduced:

$$Z_{[n \times p]} = U D V^T, \text{ How does } Z \text{ compare to original } X?$$

Dimensionality Reduction PCA

$$X_{[n \times p]} = U_{[n \times r]} D_{[r \times r]} V_{[p \times r]}^T$$



PCA - Parallelized

1. Approximate solutions to PCA (very large speedups with little drawback!):
 - a. **Stochastic Sampling** (also sometimes called "randomized" which is ambiguous): Only using a sample rows (i.e. only some users for recommendation systems)
 - b. **Truncated SVD**: Only optimizing for minimizing reconstruction error based on up to r dimensions (full SVD solves for up to $\min(n, p)$ dimensions and then you just truncate the result for the lower rank version). One you do this, by the way, using a smaller sample becomes much less of a problem.
 - c. **Limiting power iterations to a few iterations**: Power iterations from pagerank solves for the first principle component. This can be extended to multiple components.
(more [here](#).)

PCA - Parallelized

1. Approximate solutions to PCA (very large speedups with little drawback!):
 - a. **Stochastic Sampling** (also sometimes called "randomized" which is ambiguous): Only using a sample rows (i.e. users for recommendation systems)
 - b. **Truncated SVD**: Only optimizing for minimizing reconstruction error based on up to r dimensions (full SVD solves for up to $\min(n, p)$ dimensions and then you just truncate the result for the lower rank version). One you do this, by the way, using a smaller sample becomes much less of a problem.
 - c. **Limiting power iterations to a few iterations**: Power iterations from pagerank solves for the first principle component. This can be extended to multiple components.
(more [here](#).)
2. Distribute the matrix operations. Complex; not as flexible (usually done across processors within node)
3. Data Parallelism: As in other instances stochastic or mini-batch gradient descent.

Rec Systems

Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. Content-based
2. Collaborative
3. Latent Factor

Rec Systems

Common Approaches

Problems to tackle:

1. Gathering ratings
2. Extrapolate unknown ratings
 - a. Explicit: based on user ratings and reviews
(problem: only a few users engage in such tasks)
 - b. Implicit: Learn from actions (e.g. purchases, clicks)
(problem: hard to learn low ratings)
3. Evaluation

1. Content-based
2. Collaborative
3. Latent Factor