

# Recommender Systems

Sargur N. Srihari  
srihari@cedar.buffalo.edu

# Topics in Recommender Systems

- Types of Recommender Systems
- Collaborative Filtering
- Word Embeddings to Item Embeddings
- Bilinear Prediction
- Content-based Recommender Systems
- Relationship to Reinforcement Learning
  - Contextual Bandits
  - Exploration vs. Exploitation

# Recommender Systems

- Major family of ML applications in the IT sector
- Ability to recommend items to potential users
  1. Media
    - Recommending news, video on demand, music
  2. eCommerce
    - Product recommendation
  3. Jobs Boards
    - Newsletters with job offers
  4. Travel and Real Estate:
    - Events, places, chatbot
  5. Education:
    - Personalizing education, recommend materials

# Commercial importance

- Internet is financed by online advertising
- Major parts of economy rely on online shopping
- Amazon, eBay use ML including deep learning for product recommendations
- Sometimes items are not products for sale, e.g.,
  - Posts to display on social network feeds
  - Recommending movies to watch
  - Recommending jokes
  - Recommending advise

# Recommenders have commonality

- Surprisingly, similar machine learning algorithms are used for recommending
  - News
  - Videos for media
  - Products
  - Personalization in travel and retail

# Recommender's prediction

- Both online advertising and item recommendations rely on
  1. Predicting probability of some action
    - User buys product or
    - Some *proxy* for this action
      - Clicks on ad, enters rating, spends time visiting page
  2. The expected gain
    - It may depend on value of product
      - If an ad is shown or
      - A recommendation is made regarding that product

# Interaction Matrix (Rating, Feedback)

Set of users  $U$

Set of Items  $I$

Recommendation Task: Items to be recommended to users

Task of ML: Learn a function based on past data that predicts utility of each item  $i \in I$  to each user  $u \in U$

Matrix is typically huge,  
very sparse and most  
values are missing

Items  $I$

				
John	5	1	3	5
Tom	?	?	?	2
Alice	4	?	3	?

Users  $U$

# Prediction as Supervised Learning

- Given item and user, predict proxy of interest
  - Clicks on ad, user enters a rating
  - Clicks on a “like” button
  - Buys product
  - Spends money on the product
  - Spends time visiting a page for the product
- This often ends up being either
  - Regression (predict a conditional expected value)
  - Classification (predict conditional probability of a discrete event)

# Two types of Recommenders

- Recommenders use
  1. Historical interactions
  2. Attributes of items/users

## 1. Content-based Methods

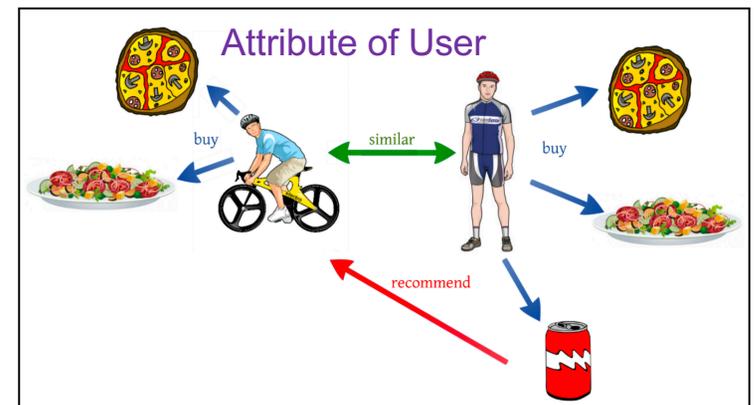
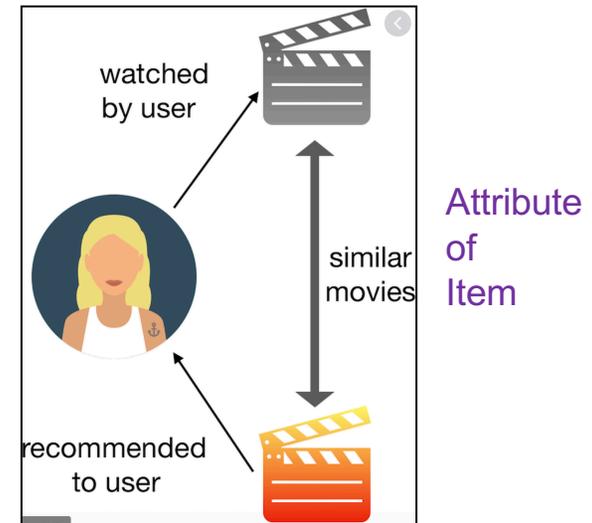
- Based on *similarity of item attributes*

## 2. Collaborative Filtering

- Recommend items by calculating *user similarity* from interactions

- Enable exploring diverse items

- Modern systems use both



# Collaborative Filtering

- Early work on recommender systems
  - Relied on minimal inputs for prediction
  - Rely on similarity between patterns of values of target variable for different users or different items
    - user 1 and user 2 both like items  $A$ ,  $B$  and  $C$
    - We may infer that user 1 and user 2 have similar tastes
    - If user 1 likes item  $D$  then this a strong cue that user 2 will also like  $D$
    - Algorithms based on this principle come under the name of *collaborative filtering*

# Feedback can be Explicit and Implicit

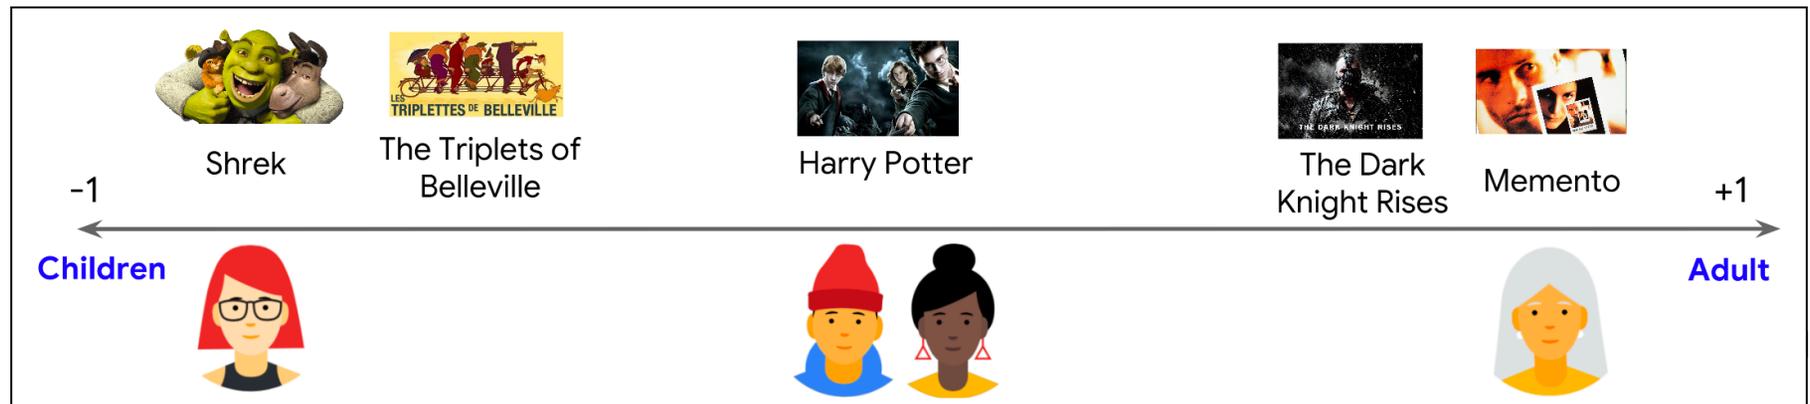
- Movie Recommendation Example
  - Row represents user, column represents movie
- Feedback falls into one of two categories:
  - **Explicit**— users provide a numerical rating
  - **Implicit**— if user watches movie, infer interested
- Binary feedback matrix: 1 indicates interest
- When user visits homepage, system should recommend movies based on both:
  - similarity to movies user has liked in the past
  - movies that similar users liked

# Features to explain feedback

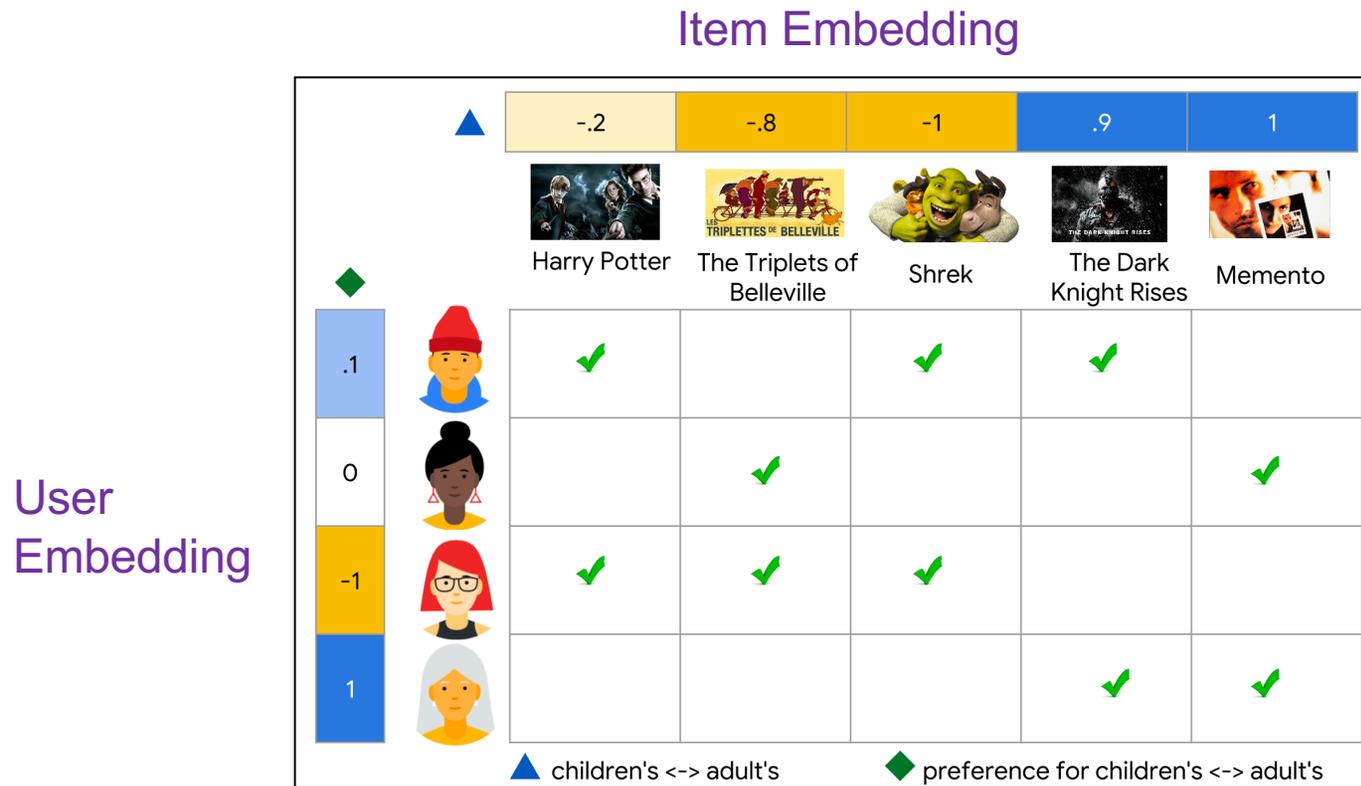
- Movies:** Feature is a scalar in  $[-1, 1]$  that describes whether the movie is for children (negative values) or adults (positive values).

Movie	Rating	Description
<a href="#">The Dark Knight Rises</a>	PG-13	Batman endeavors to save Gotham City from nuclear annihilation in this sequel to <a href="#">The Dark Knight</a> , set in the DC Comics universe.
<a href="#">Harry Potter and the Sorcerer's Stone</a>	PG	A orphaned boy discovers he is a wizard and enrolls in Hogwarts School of Witchcraft and Wizardry, where he wages his first battle against the evil Lord Voldemort.
<a href="#">Shrek</a>	PG	A lovable ogre and his donkey sidekick set off on a mission to rescue Princess Fiona, who is imprisoned in her castle by a dragon.
<a href="#">The Triplets of Belleville</a>	PG-13	When professional cyclist Champion is kidnapped during the Tour de France, his grandmother and overweight dog journey overseas to rescue him, with the help of a trio of elderly jazz singers.
<a href="#">Memento</a>	R	An amnesiac desperately seeks to solve his wife's murder by tattooing clues onto his body.

## 1-D Embedding



# Feedback matrix with one feature



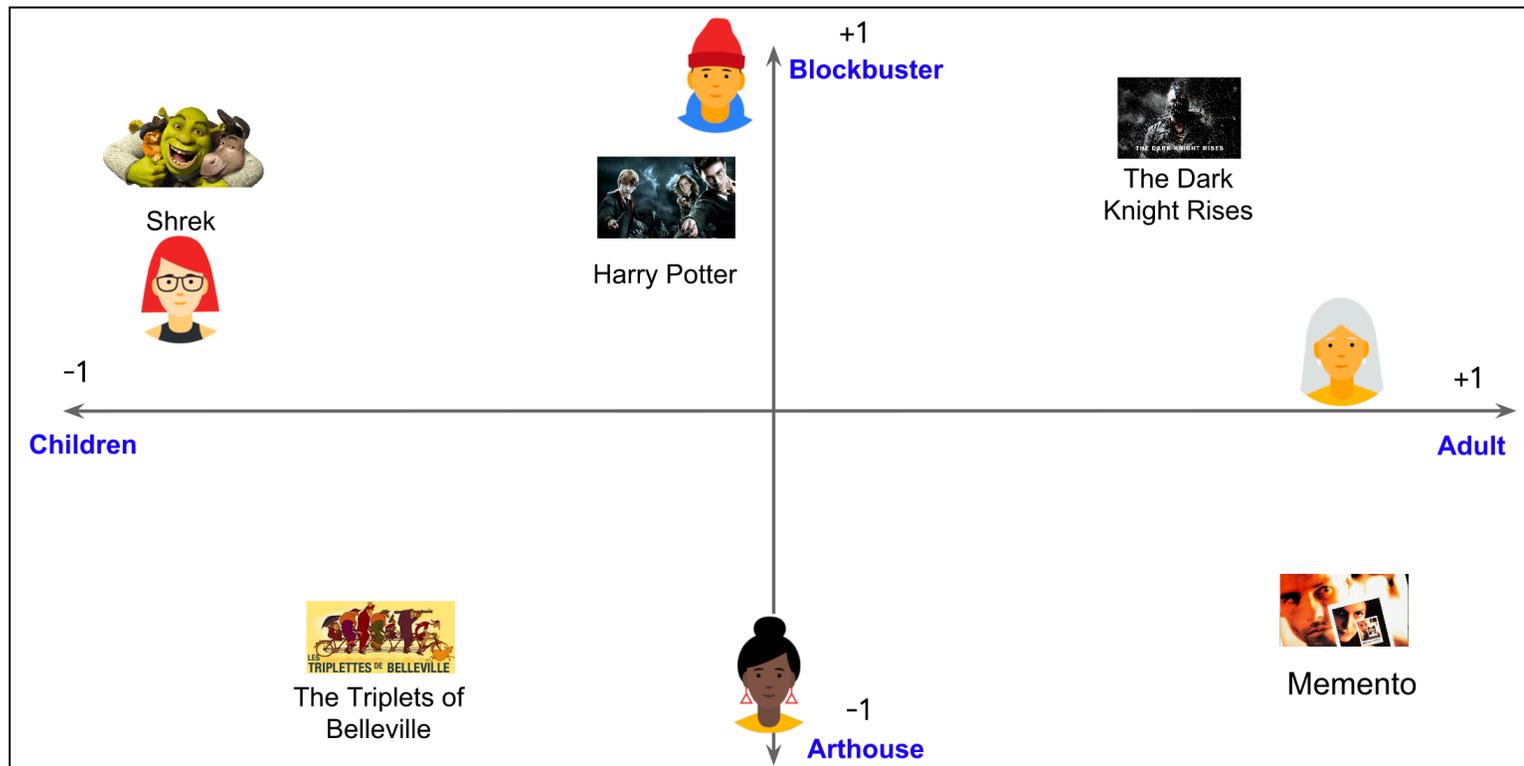
Each checkmark identifies a movie that a user watched.

Third and fourth users have preferences that are well explained by this feature—the third user prefers movies for children and the fourth user prefers movies for adults.

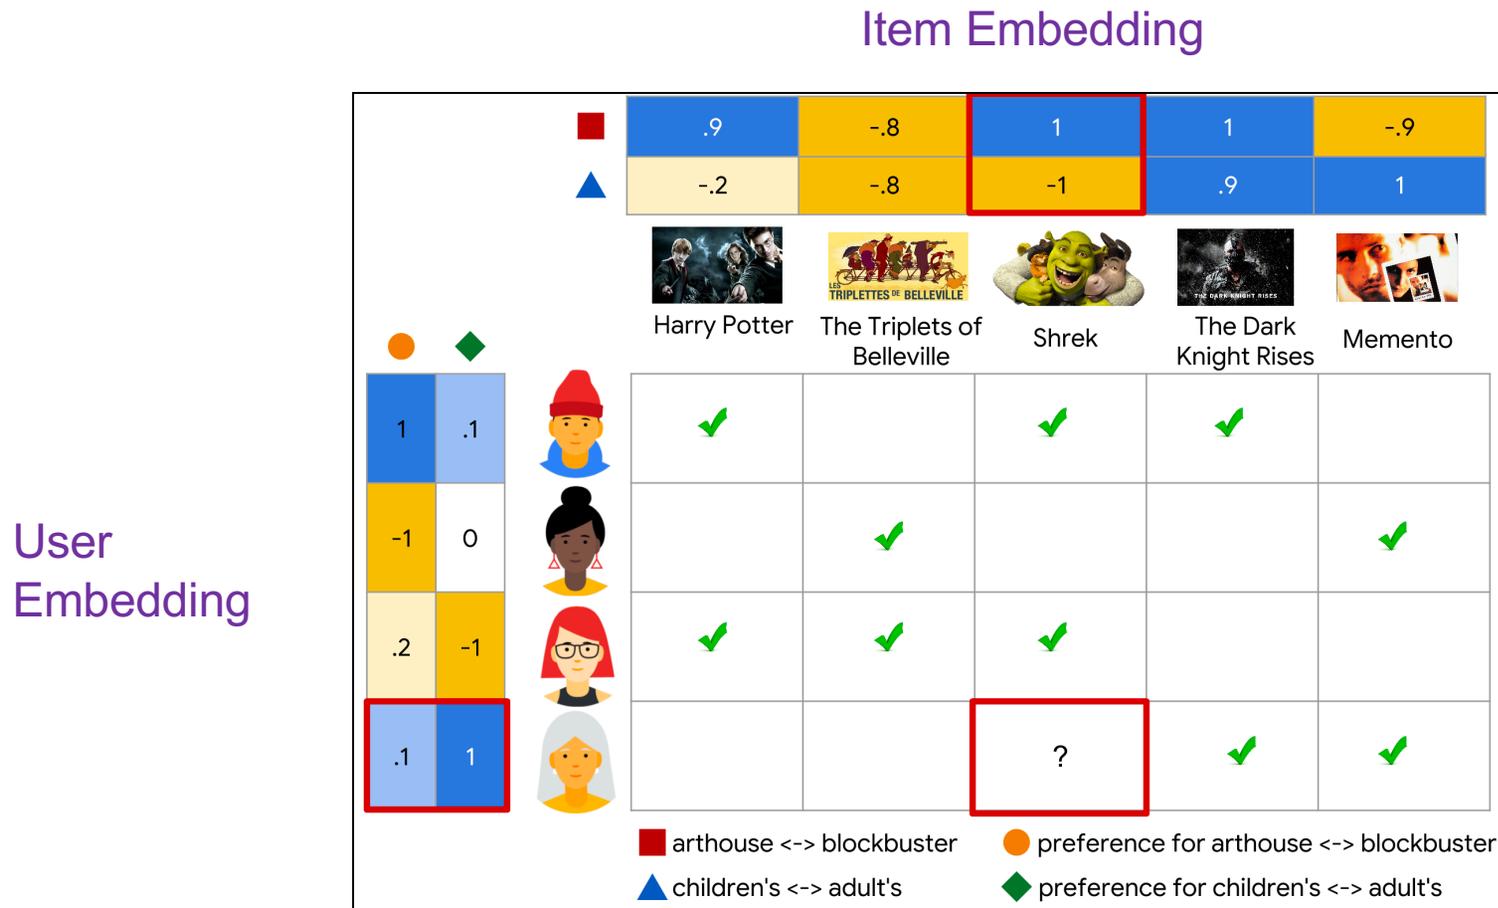
However, the first and second users' preferences are not well explained by this single feature.

# 2-D embedding

One feature was not enough to explain preferences of all users.  
To overcome this problem, let's add a second feature:  
degree to which movie is a blockbuster or arthouse



# Feedback matrix with 2 features



We again place our users in the same embedding space to best explain the feedback matrix: for each (user, item) pair, we would like the dot product of the user embedding and the item embedding to be close to 1 when the user watched the movie, and to 0 otherwise.

# Note on Embedding space

- We represented both items and users in the same embedding space.
  - This may seem surprising.
  - After all, users and items are two different entities.
- However, you can think of the embedding space as an abstract representation common to both items and users
  - in which we can measure similarity or relevance using a similarity metric.

# Collaboration

- Collaboration is apparent when model learns embeddings
- If embeddings for movies are fixed
  - Can learn embeddings for users to explain preferences
    - User embeddings with similar preferences will be close
- If embeddings for users are fixed
  - Can learn movie embeddings to best explain feedback matrix
    - Embeddings of movies liked by similar users will be close in the embedding space.

# Collaborative Filtering Algorithms

## 1. Non-parametric method

- Nearest neighbor based on similarity between patterns of preferences

## 2. Parametric method

- Rely on learning a distributed representation called an *embedding* for each user and each item
- Bilinear prediction of the target variable

# Non-parametric: User-based knn

- Compute similarity of users
- Find  $k$  most similar users to  $a$ 
  - Alice most similar to John
- Recommend movies not seen by user  $a$ 
  - Items 2 and 4 of Alice get values 1 and 5

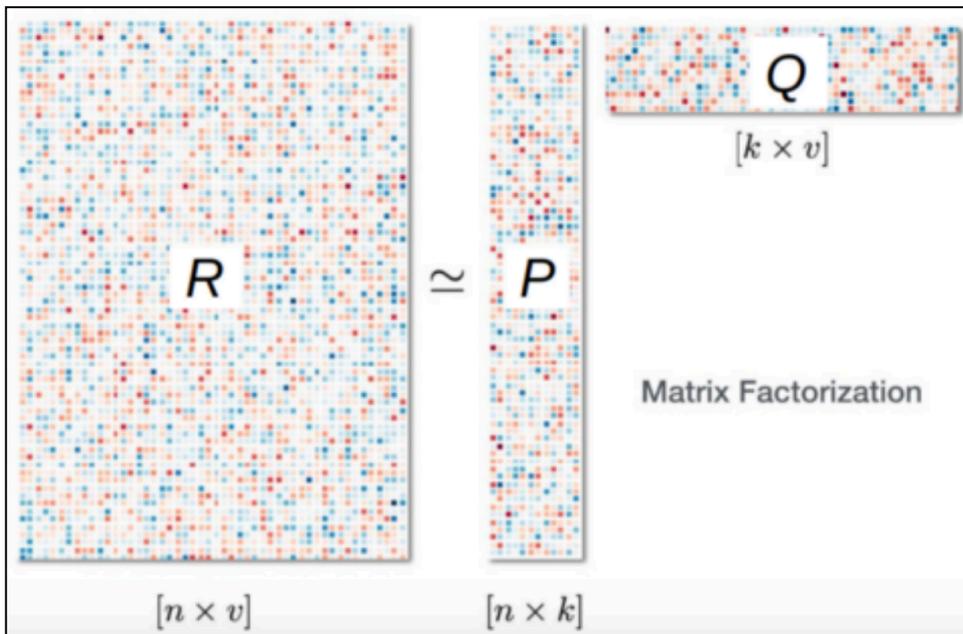
Cosine similarity

$$\text{sim}(a, b) = \frac{a \cdot b}{\|a\| \cdot \|b\|}$$

	ARCA	The Matrix	The Godfather	The Matrix Reloaded
John	5	1	3	5
Tom	?	?	?	2
Alice	4	?	3	?

# Matrix Factorization

- Attempt to reduce dimensionality of the interaction matrix and approximate it by two or more small matrices with  $k$  latent components



$$L = \|R - P \times Q^T\|_2 + \lambda (\|P\|_2 + \|Q\|_2)$$

$$\hat{r}_{ij} = p_i^T q_j = \sum_k p_{ik} q_{kj}$$

$$\operatorname{argmin}_{q,p} \sum_{i,j} (r_{ij} - p_i^T q_j)^2$$

By multiplying corresponding row and column you predict rating of item by user. Training error can be obtained by comparing non empty ratings to predicted ratings. One can also regularize training loss by adding a penalty term keeping values of latent vectors low.

# Parametric Method: Item Embeddings

## Sentence-word data

s0: "Man walked his dog"  
s1: "Man took his dog to park"  
s2: "Dog went to park"

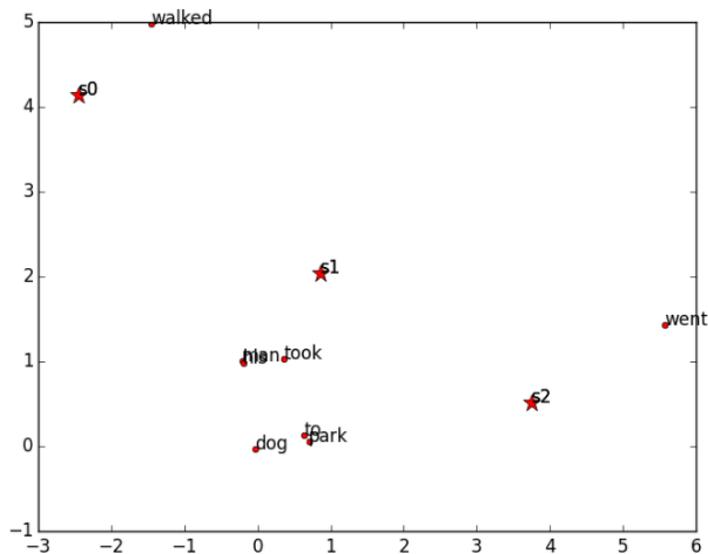
Vocabulary:  
{man, walked, his, dog, took, to, park, went}

Vector representation:  
s0: [1,1,1,1,0,0,0,0]  
s1: [1,0,1,1,1,1,1,0]  
s2: [0,0,0,1,0,1,1,1]

*Word2Vec* produces word embeddings in a low-dimensional continuous space and carry semantic and syntactic information of words

Continuous vector representations

Word2Vec



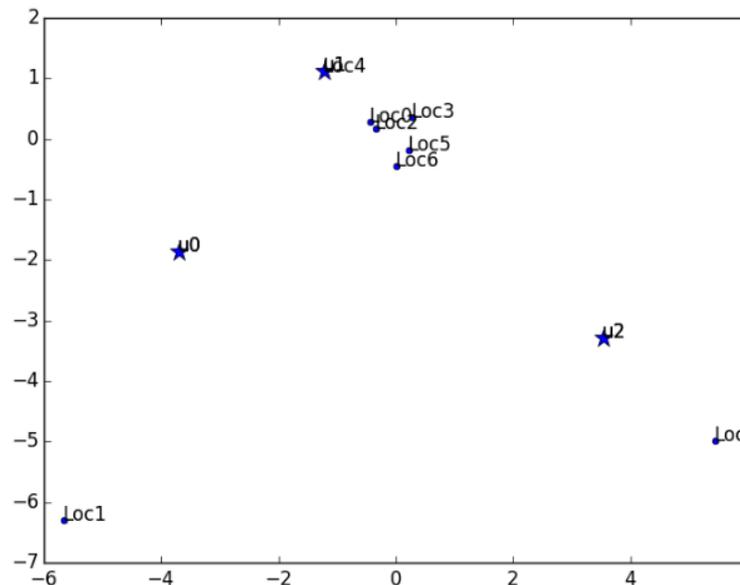
## User-item data

user0: Loc0, Loc1, Loc2, Loc3  
user1: Loc0, Loc4, Loc2, Loc3, Loc5, Loc6  
user2: Loc3, Loc7, Loc5, Loc6

Vocabulary:  
{Loc0, Loc1, Loc2, Loc3, Loc4, Loc5, Loc6, Loc7}

Vector representation:  
user0: [1,1,1,1,0,0,0,0]  
user1: [1,0,1,1,1,1,1,0]  
user2: [0,0,0,1,0,1,1,1]

Item2Vec



Similarly we get user2Vec

# Prediction of target variable

- Prediction of target variable (such as a rating)
  - Bilinear prediction is a simple parametric method
    - Highly successful
    - Found as a component in state-of-the-art systems
- Prediction is obtained by a dot product between the user embedding and the item embedding
  - Possibly corrected by by constants that depend only on either the user ID or item ID

# Bilinear Prediction

- Definitions

- Let  $\hat{R}$  be the matrix containing our predictions
- $A$  a matrix with user embeddings in its rows
- $B$  a matrix with item embeddings in its columns
- Let  $b$  and  $c$  be vectors that contain respectively
  - a kind of bias for each user
    - Representing how grumpy or positive that user is
  - For each item
    - Representing its general popularity

- The bilinear prediction is  $\hat{R}_{u,i} = b_u + c_i + \sum_j A_{u,j} B_{j,i}$

- Typical goal: minimize squared error between predicted ratings  $\hat{R}_{u,i}$  and actual ratings  $R_{u,i}$

# Use of embeddings

- User embeddings and item embeddings can then be conveniently visualized when they are first reduced to a low dimension (two or three)
- Or they can be used to compare users or items against each other, just like word embeddings

# Obtaining embeddings

- One way to obtain these embeddings is by performing singular value decomposition (SVD) of the matrix  $\mathbf{R}$  of actual targets (such as ratings)
- This corresponds to factorizing  $\mathbf{R}=\mathbf{UDV}'$  (or a normalized variant) into the product of two factors, the lower rank matrices  $\mathbf{A}=\mathbf{UD}$  and  $\mathbf{B}=\mathbf{V}'$

# Missing Entry problem with SVD

- One problem with SVD is that it treats missing entries in an arbitrary way, as if they corresponded to a target value of 0
- Instead we would like to avoid paying any cost for the predictions made on missing entries
- Fortunately the sum-of-squared-errors on the observed ratings can also be easily minimized by gradient-based optimization

# Netflix Competition

- Both SVD and bilinear prediction performed very well
- Competition was to predict ratings for films, based on previous ratings by a large set of anonymous users
- Even though it did not win by itself, the simple bilinear prediction of SVD was a component of the ensemble models presented by most competitors including the winners

$$\hat{R}_{u,i} = b_u + c_i + \sum_j A_{u,j} B_{j,i}$$

# Limitation of Collaborative Filtering

- When a new item or a new user is introduced, its lack of rating history means that there is no way to evaluate its similarity with other items or users (respectively),
- Or the degree of association between that user and existing items
- This is the problem of *cold-start recommendations*

# Solution to cold-start recommendation

- Introduce extra information about individual users or items
- Extra information could be user profile information or features of each item
- Systems that use such information are called *content-based recommender systems*
- The mapping from a rich set of user features or item features to an embedding can be learned through a deep learning architecture

# Content-based recommender systems

- Deep learning architectures such as CNNs learn to extract from rich content, e.g., musical audio tracks, for music recommendation
  - CNN takes acoustic features as input and computes an embedding for the associated song
  - Dot product between this song embedding and the embedding for a user is then used to predict whether a user will listen to the song

# Relationship to Reinforcement Learning

- A recommendation issue goes beyond supervised learning and into reinforcement learning
- Most recommendation problems are accurately described theoretically as *contextual bandits*
  - When recommendation system collects data, we get a biased and incomplete view of user preferences
    - We only see responses of users to items they were recommended and not to other items
    - In some cases no information on users for whom no recommendation has been made
      - E.g., with ad auctions, price proposed was below minimum, or

# Need for additional information

- System gets no information on what outcome would result from recommending any other item
  - It would be like training a classifier by picking one class  $\hat{y}$  for each training example  $\mathbf{x}$  (typically class with highest probability) and then getting feedback whether this was correct or not
    - Each example conveys less information than in the supervised case where the true label  $y$  is directly observable, so more examples are necessary
      - We may keep on picking the wrong model output to show
    - The correct decision may have a low probability
      - Until learner picks the correct decision it does not learn about the correct decision

# Reinforcement Learning and Bandits

- In reinforcement learning only the reward for the selected action is observed
- In general, reinforcement learning can involve a sequence of many actions and many rewards
- Bandits scenario is a special case of reinforcement learning, in which the learner takes only a single action and receives a single reward
  - Bandit problem is easier in that the learner knows which reward is associated with which action

# Multi-armed Bandit Problem

1. Which machines to play
2. How many times to play each machine
3. In which order to play
  - Each machine has
    - a probability distribution,  $B = \{R_1, \dots, R_K\}$ 
      - Mean values  $\mu_1, \dots, \mu_K$  associated with rewards
  - Objective:
    - Maximize reward through sequence of lever pulls
      - Minimize regret  $\rho$ , expected difference between optimal strategy and collected rewards  $r_t$ , after  $T$  rounds

One-armed  
bandit



Multi-armed  
bandit



$$\rho = T\mu^* - \sum_{t=1}^T \hat{r}_t$$

# Contextual Bandits

- In general reinforcement learning, a high or low reward might have been caused by a recent action or by an action in the distant past
- Contextual bandits refers to where the action is taken in context of an input variable that can inform the decision
  - E.g., we at least want to know user identity, and we want to pick an item
- Mapping from context to action is called policy
  - Feedback loop between learner and data distribution (which depends on actions of learner)
    - a central research issue in reinforcement learning

# Exploration vs Exploitation

- Reinforcement learning requires a tradeoff between exploration and exploitation
- Exploitation comes from taking actions that come from the current best version of the learned policy
  - Actions that we know will achieve a high reward
- Exploitation refers to taking actions specifically in order to obtain more training data

# Exploration example

- We know: in context  $\mathbf{x}$ , action  $a$  has reward=1
  - We don't know whether it's the best possible reward
  - We may want to exploit our current policy and continue taking action  $a$  in order to be relatively sure of obtaining a reward of 1
- However we may also want to explore action  $a'$ 
  - We do not know what will happen if we try action  $a'$
  - We hope for reward=2 but we may get reward=0
  - Either way we gain some knowledge

# Implementing Exploration-Exploitation

- Exploration can be implemented in many ways
  1. Random actions to cover all possible actions
  2. Model-based approaches
    - compute a choice of action based on its expected reward and the model's amount of uncertainty about that reward
- Factors determining exploration or exploitation
  - One prominent factor: Time scale
    - Agent has short time to accrue reward: exploitation
    - Agent has more time: begin with exploration so that future actions can be planned more effectively
      - As time progresses we move towards more exploitation

# Supervised Learning has no tradeoff

- No tradeoff between exploration-exploitation
- Supervision signal always specifies which output is correct for each input
- There is no need to try out different outputs to determine if one is better than the model's current output
  - We always know that the label is the best output

# Evaluating policies

- Another difficulty arising in the context of reinforcement learning
  - Besides exploration-exploitation tradeoff
- Difficulty of evaluating and comparing different policies
  - Reinforcement learning involves interaction between learner and environment
    - It is not straightforward to evaluate the learner's performance using a fixed set of test input values
    - The policy itself determines which inputs will be seen