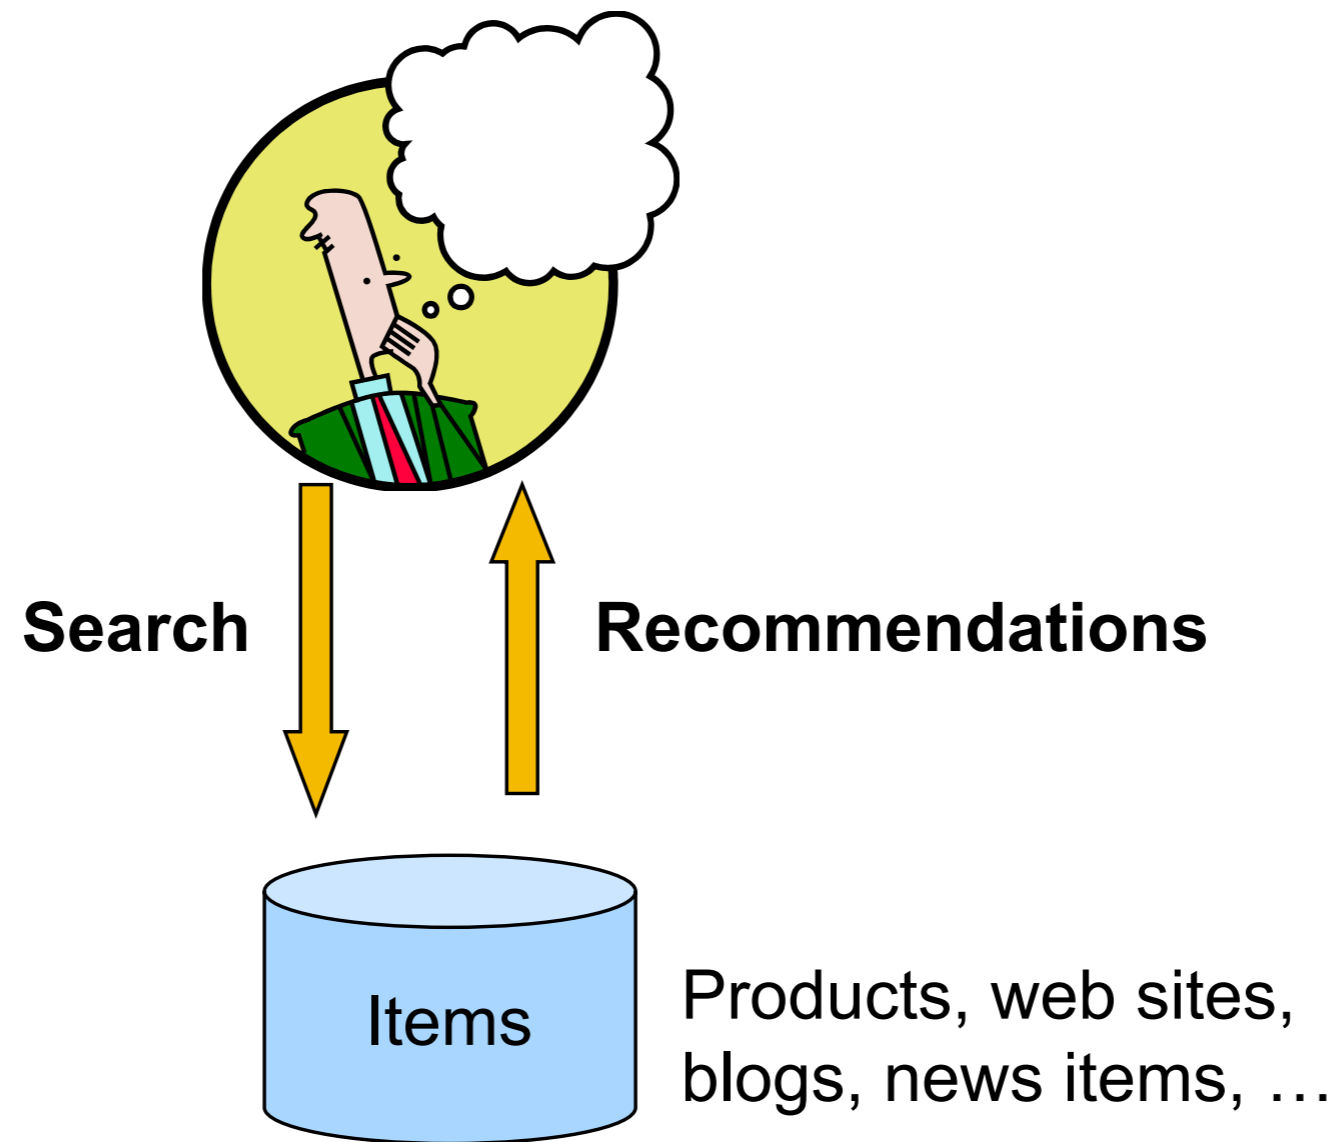


Recommendation Systems

CS 534: Machine Learning

Slides adapted from Alex Smola, Jure Leskovec, Anand Rajaraman, Jeff Ullman, Lester Mackey, Dietmar Jannach, and Gerhard Friedrich

Recommender Systems (RecSys)



RecSys is Everywhere



System that provides or suggests items to the end users

Long Tail Phenomenon




Physical vs Online Presence

THE BIT PLAYER ADVANTAGE

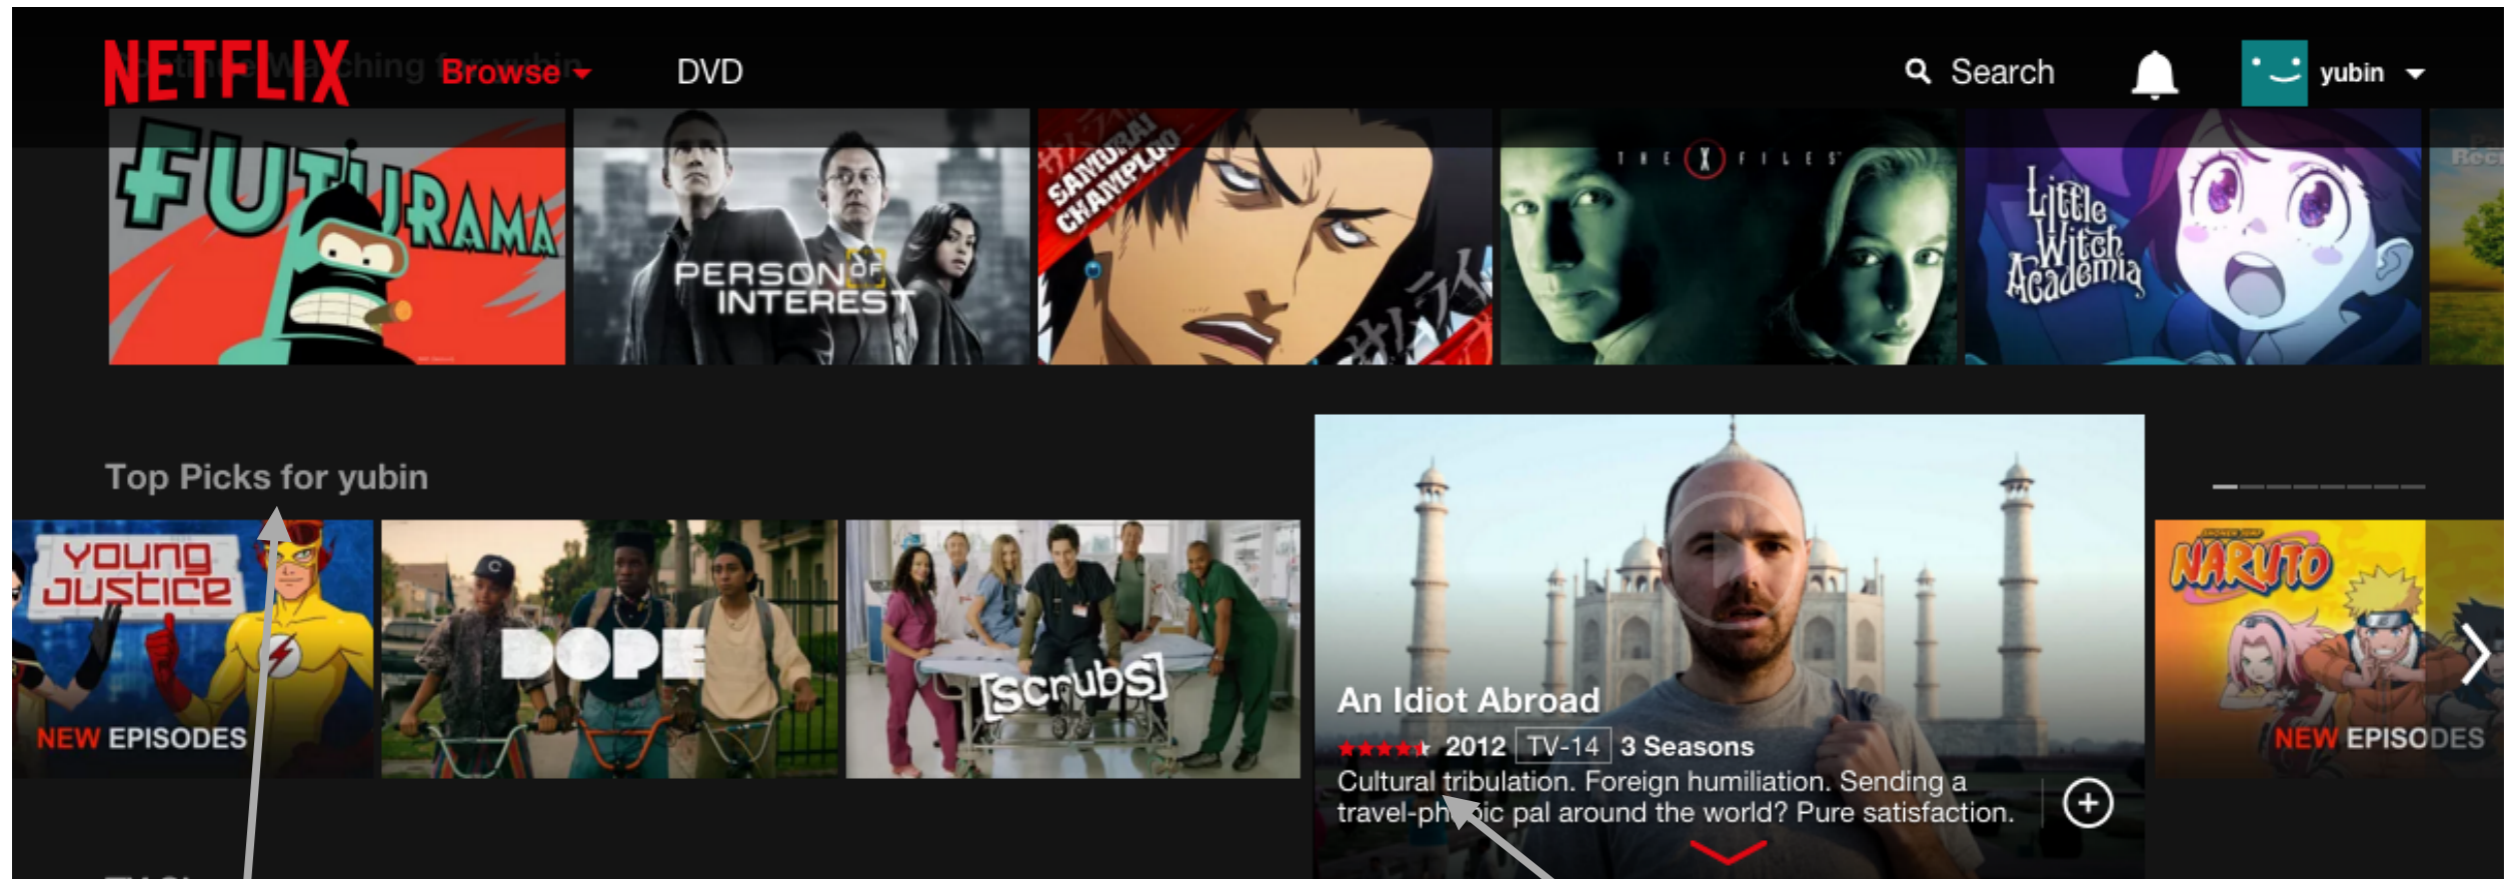
Beyond bricks and mortar there are two main retail models – one that gets halfway down the Long Tail and another that goes all the way. The first is the familiar hybrid model of Amazon and Netflix, companies that sell physical goods online. Digital catalogs allow them to offer unlimited selection along with search, reviews, and recommendations, while the cost savings of massive warehouses and no walk-in customers greatly expands the number of products they can sell profitably.

Pushing this even further are pure digital services, such as iTunes, which offer the additional savings of delivering their digital goods online at virtually no marginal cost. Since an extra database entry and a few megabytes of storage on a server cost effectively nothing, these retailers have no economic reason not to carry *everything* available.



Source: Amazon.com

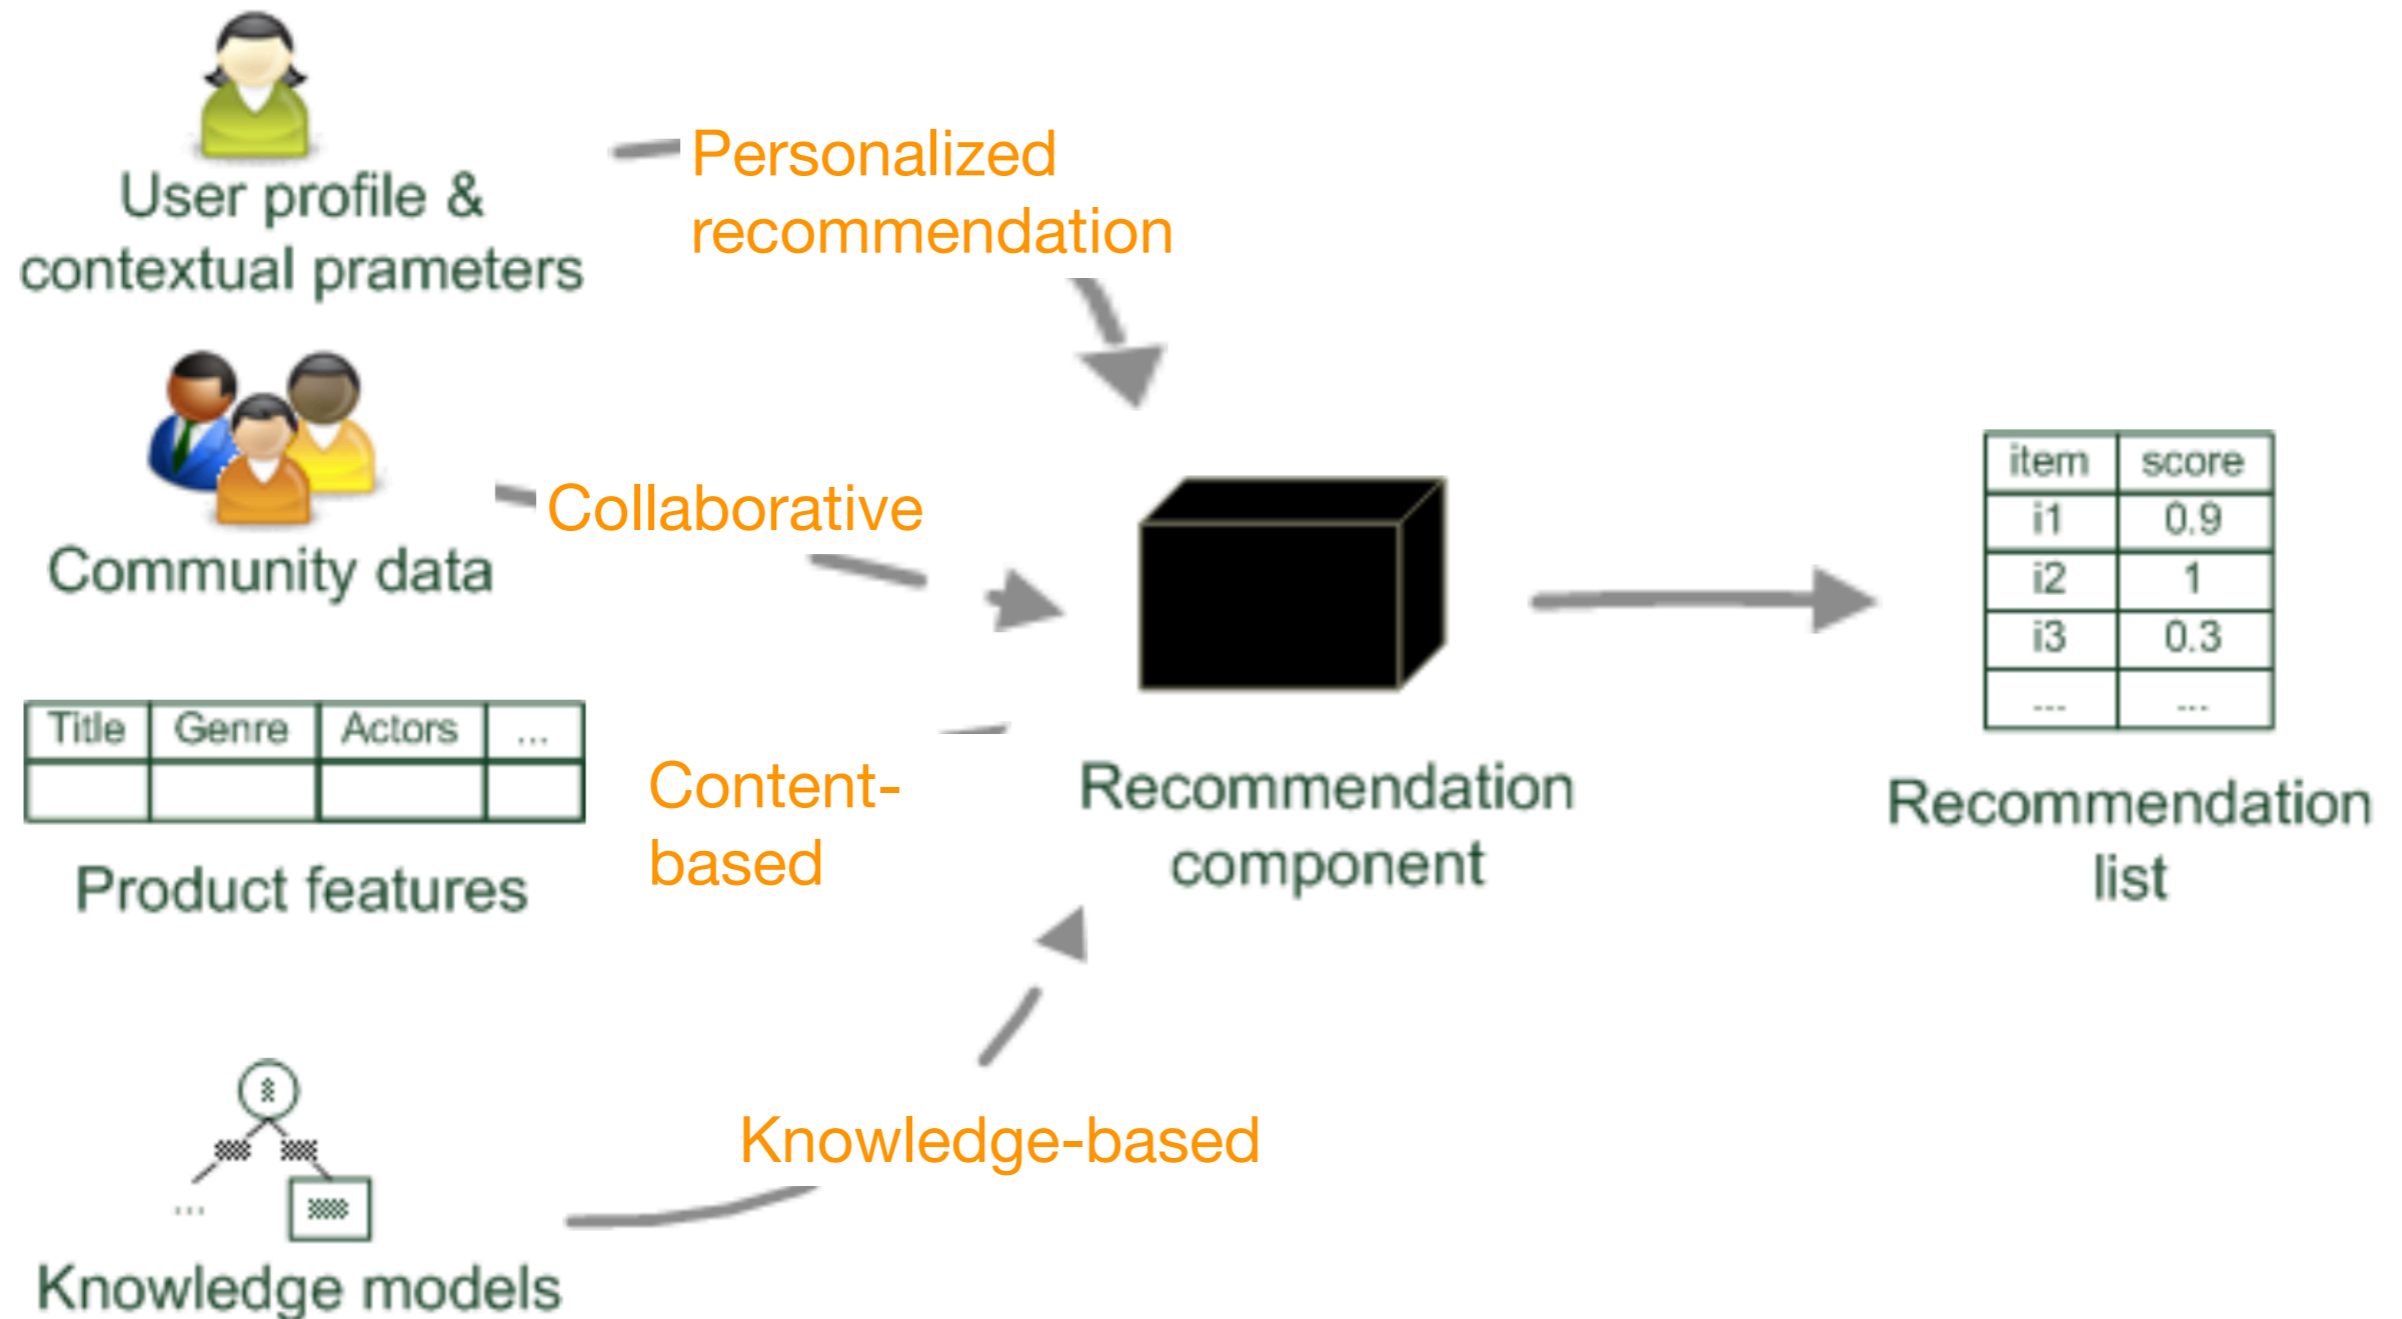
RecSys: Tasks



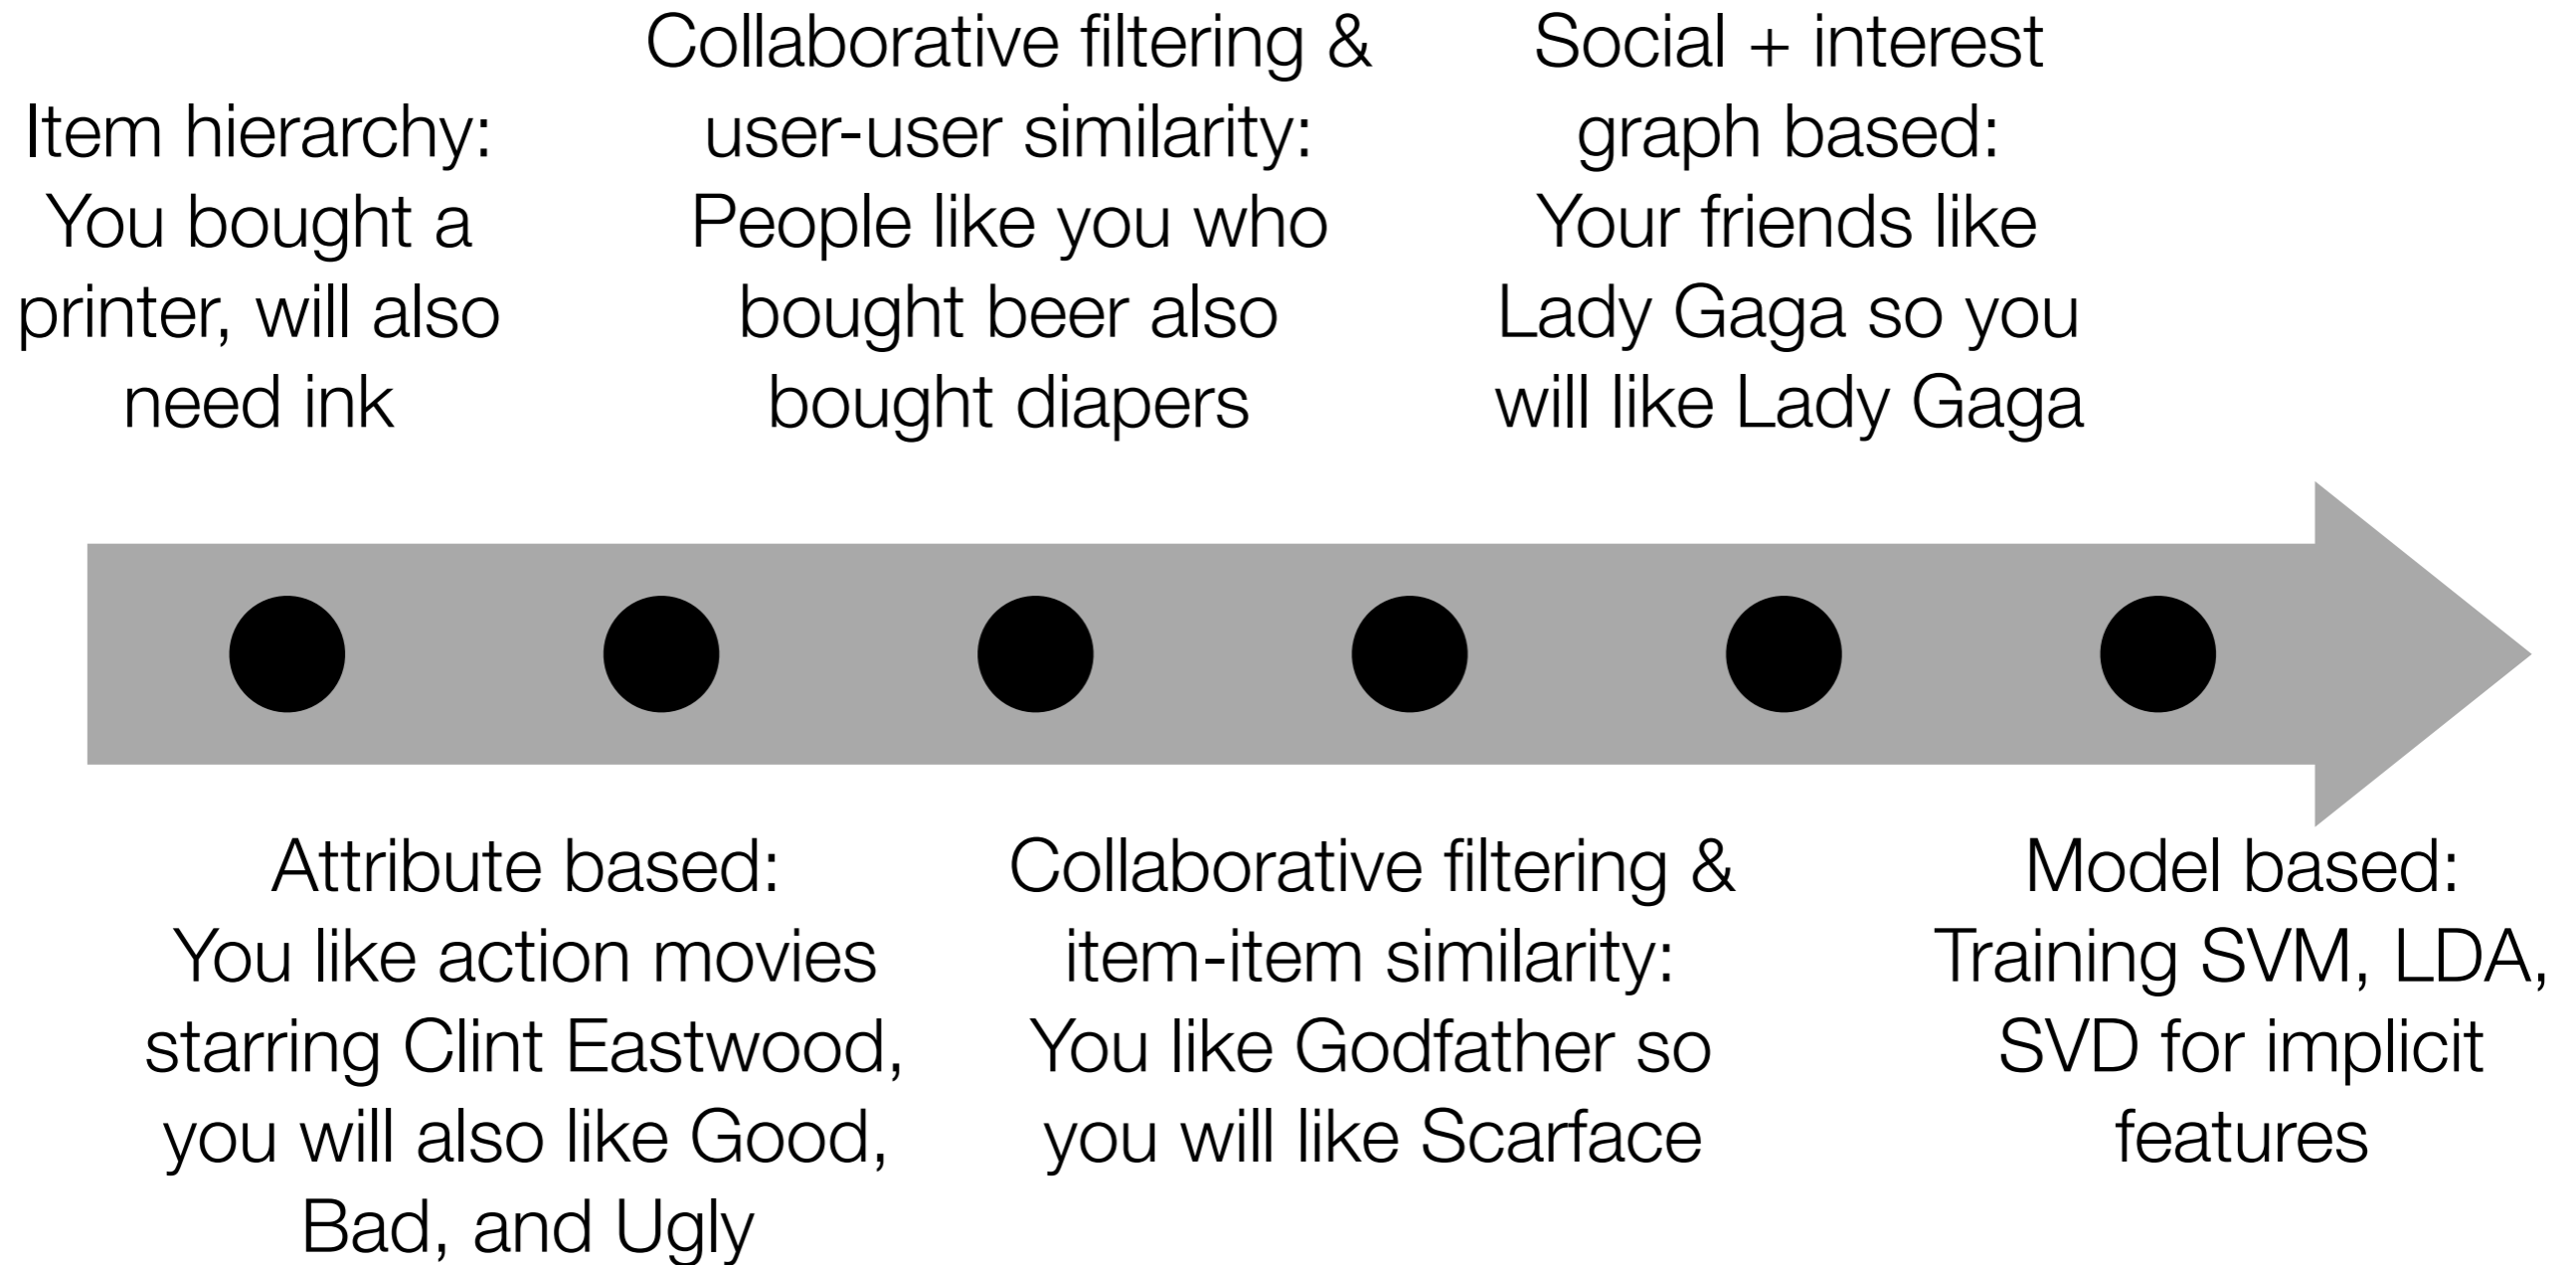
Task 1: Predict user rating

Task 2: Top N recommendation

RecSys: Paradigms



RecSys: Evolution



RecSys: Basic Techniques

	Pros	Cons
Collaborative	No knowledge- engineering effort, serendipity of results, learns market segments	Requires some form of rating feedback, cold start for new users and new items
Content-based	No community required, comparison between items possible	Content descriptions necessary, cold start for new users, no surprises
Knowledge-based	Deterministic recommendations, assured quality, no cold- start, can resemble sales dialogue	Knowledge engineering effort to bootstrap, basically static, does not react to short-term trends

RecSys: Challenges

- Scalability - millions of objects and users
- Cold start
 - Changing user base
 - Changing inventory (movies, stories, goods)
 - Attributes
- Imbalanced dataset - user activity / item reviews are power law distributed

Netflix Prize: \$1 M (2006-2009)



The image shows a screenshot of the Netflix Prize website. At the top, the Netflix logo is visible. Below it, a yellow banner reads "Netflix Prize". A navigation bar includes links for Home, Rules, Leaderboard, Register, Update, Submit, and Download. The main content area features a "Movies For You" section with recommendations like "The Big One" and "The Big Bang Theory". A "Welcome!" message box is overlaid on the right, explaining the prize and providing links to rules, registration, and frequently asked questions. Silhouettes of two people are shown in the foreground, looking at the screen. The background of the website is a red gradient with a green code-like pattern on the left side.

NETFLIX

Netflix Prize

Home Rules Leaderboard Register Update Submit Download

NETFLIX

Browse Recommendations Friends Queue Buy DVDs

Home Genres New Releases Previews Netflix Top 100 Crit

Movies For You

Randy, the following movies were chosen based on your interest in:

- Dancing Queen
- Carnival: Season 1
- Crash

You really liked it...

Now only for just \$5.99

The Big One

★ ★ ★ ☆ ☆

Original art

Welcome!

The Netflix Prize seeks to substantially improve the accuracy of predictions about how much someone is going to love a movie based on their movie preferences. Improve it enough and you win one (or more) Prizes. Winning the Netflix Prize improves our ability to connect people to the movies they love.

Read the [Rules](#) to see what is required to win the Prizes. If you are interested in joining the quest, you should [register a team](#).

You should also read the [frequently asked questions](#) about the Prize. And check out how various teams are doing on the [Leaderboard](#).

Good luck and thanks for helping!

FAQ | Forum | Netflix Home

© 1997-2008 Netflix, Inc. All rights reserved.

Netflix Movie Recommendation

- Training Data:
 - 480,000 users
 - 17,700 movies
 - 6 years of data:
2000-2005
- Test data: most recent ratings of each user

Training data

user	movie	date	score
1	21	5/7/02	1
1	213	8/2/04	5
2	345	3/6/01	4
2	123	5/1/05	4
2	768	7/15/02	3
3	76	1/22/01	5
4	45	8/3/00	4
5	568	9/10/05	1
5	342	3/5/03	2
5	234	12/28/00	2
6	76	8/11/02	5
6	56	6/15/03	4

Test data

user	movie	date	score
1	62	1/6/05	?
1	96	9/13/04	?
2	7	8/18/05	?
2	3	11/22/05	?
3	47	6/13/02	?
3	15	8/12/01	?
4	41	9/1/00	?
4	28	8/27/05	?
5	93	4/4/05	?
5	74	7/16/03	?
6	69	2/14/04	?
6	83	10/3/03	?

Evaluation Metrics

Error on unseen test set Q , not on training error

- Root Mean Square Error

$$\text{RMSE} = \sqrt{\frac{1}{|S|} \sum_{(i,u) \in S} (\hat{r}_{ui} - r_{ui})^2}$$

- Mean Absolute Error

$$\text{MAE} = \frac{1}{|S|} \sum_{(i,u) \in S} |\hat{r}_{ui} - r_{ui}|$$

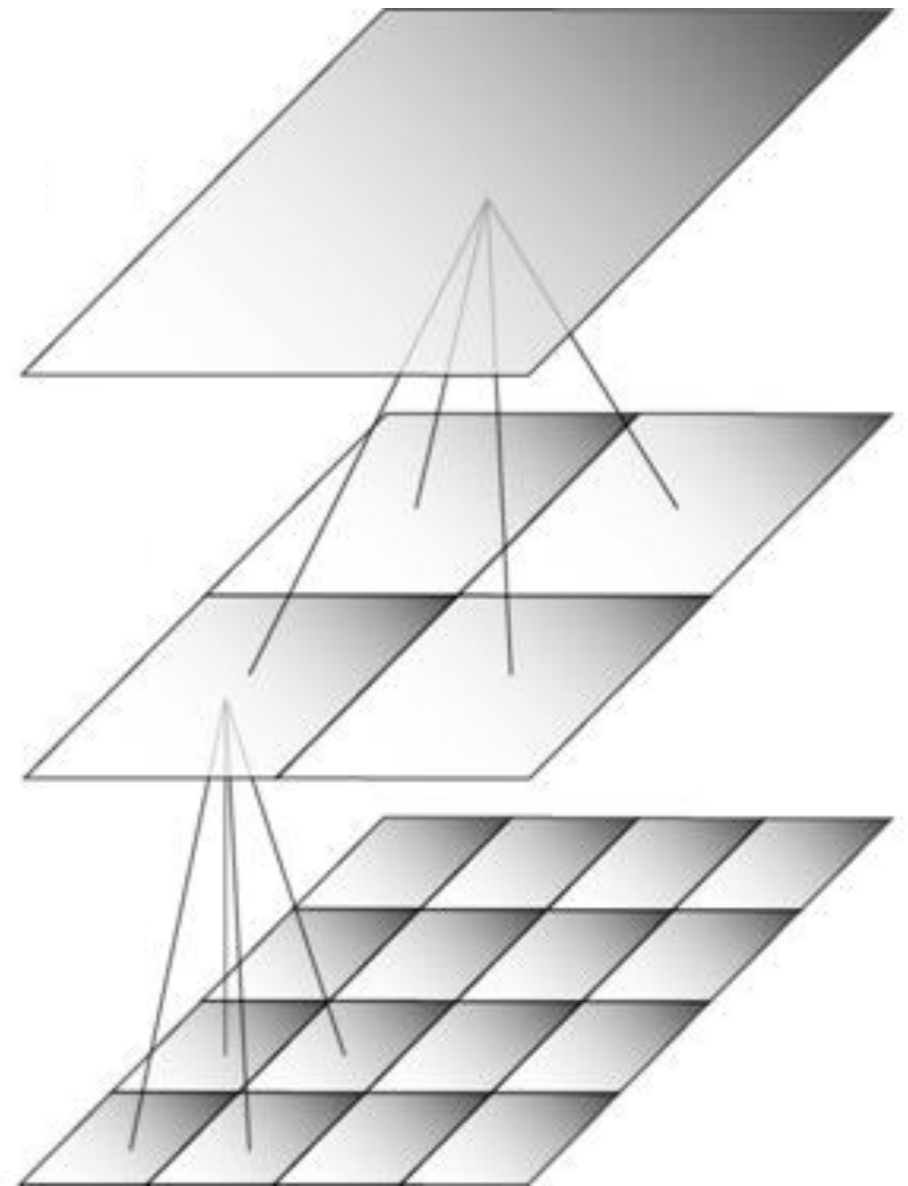
- Rank-based objectives (e.g., What fraction of true top-10 preferences are in predicted top 10?)

Netflix Prize

- Evaluation criterion: RMSE
- Cinematch (Netflix) system RMSE: 0.9514
- Competition
 - 2700+ teams
 - \$1 M prize for 10% improvement on Netflix

Netflix Winner: BellKor

- Multi-scale modeling of the data:
 - Global: Overall deviations of users & movies
 - Factorization: “Regional” effects
 - Collaborative filtering: Extract local patterns



Normalization / Global Bias

- Mean movie rating across all movies
- Some users tend to give higher ratings than others
- Some movies tend to receive higher rating than others

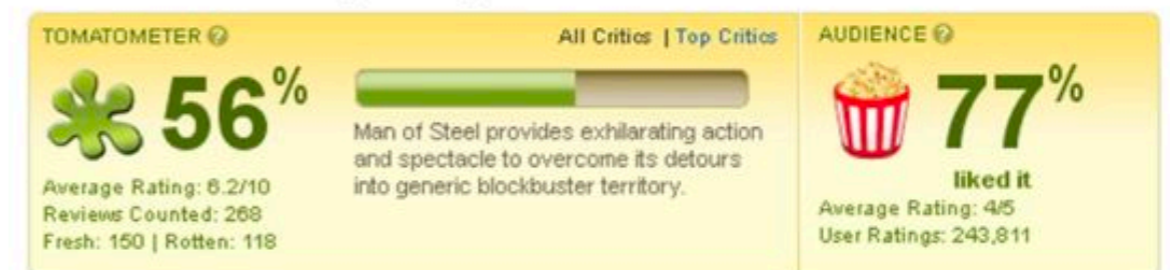
The Wolverine (2013)



Iron Man 3 (2013)




Man of Steel (2013)



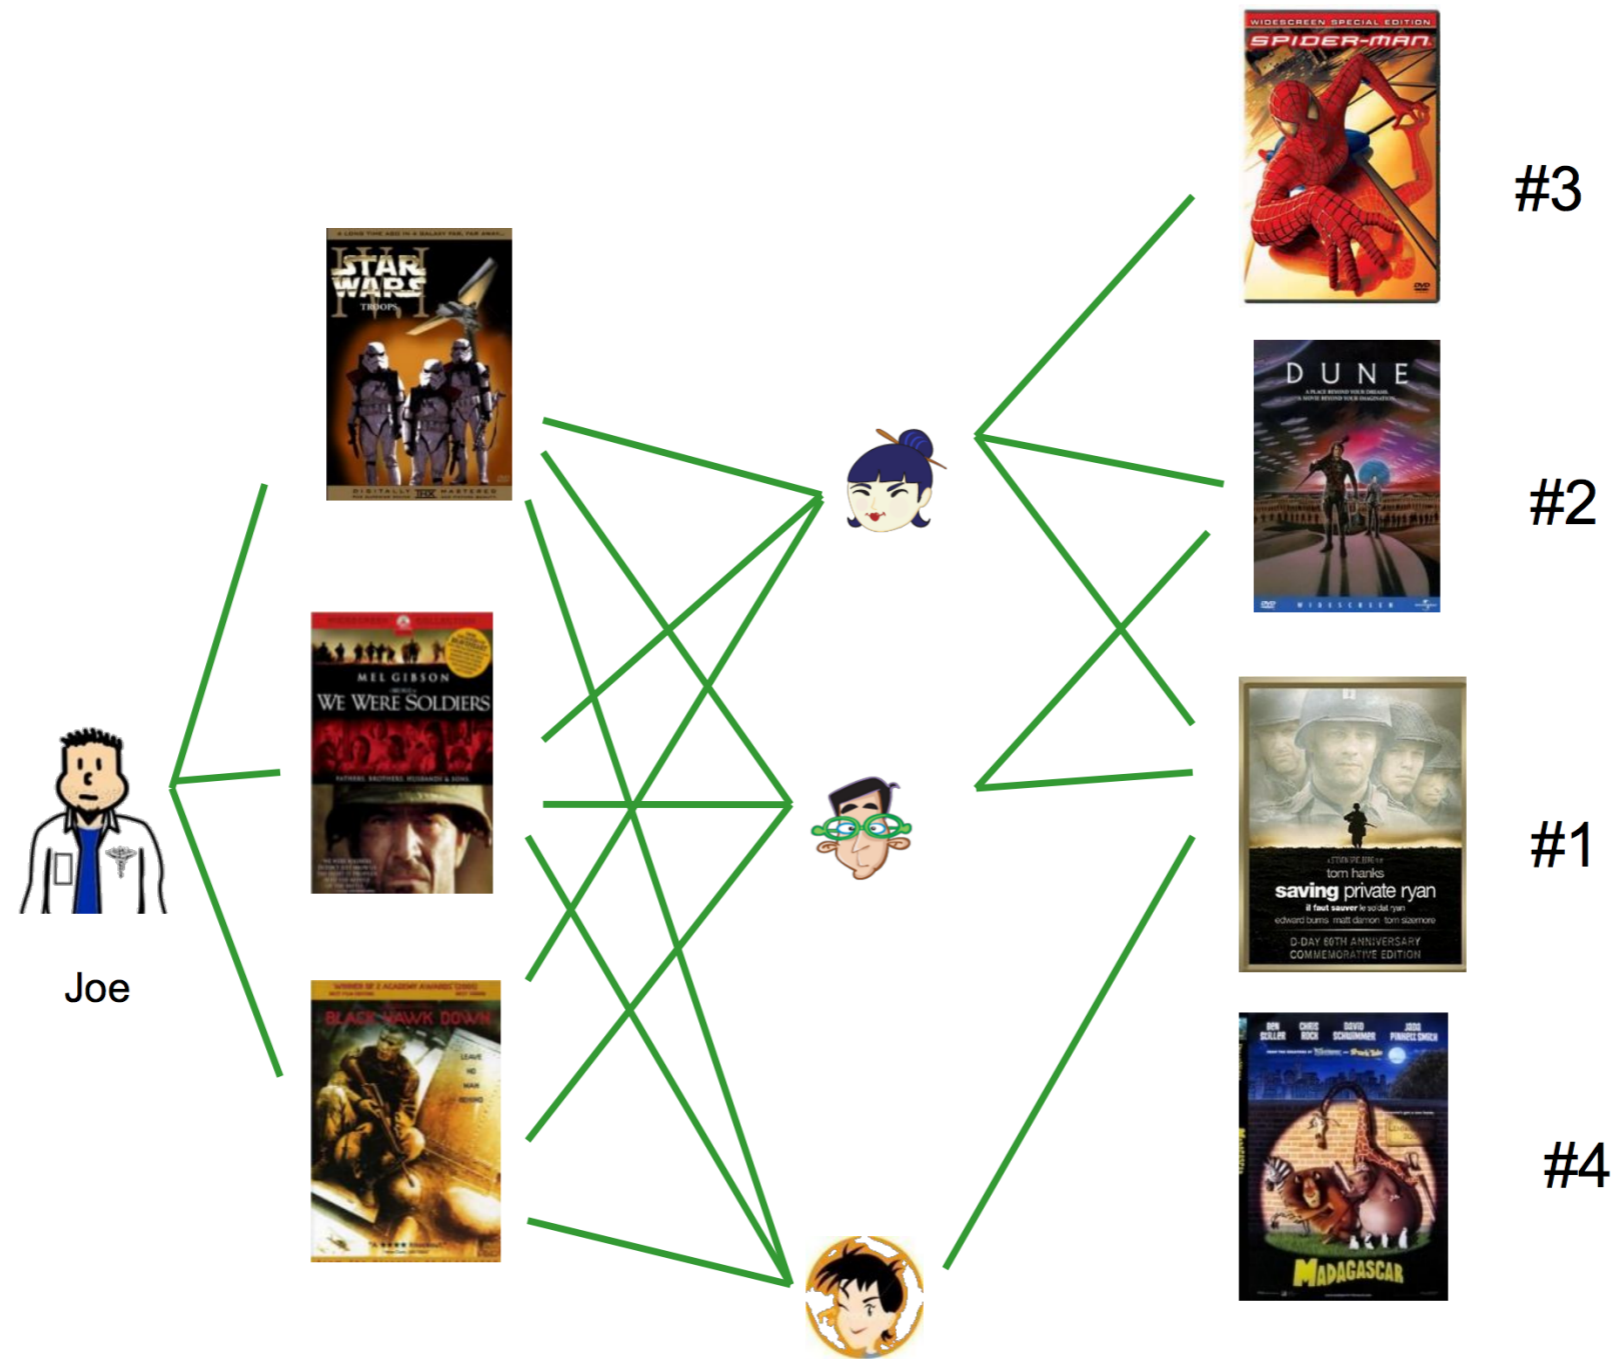
Example: Global & Local Effects

- Global effect
 - Mean movie rating: 3.7 stars
 - The Sixth Sense is 0.5 stars above average
 - Joe rates 0.2 stars below average
 - Local effect
 - Joe doesn't like related movie Signs
- Baseline estimate: 4 stars
- Final estimate: 3.8 stars

Netflix Performance




Neighborhood Methods: Basic Idea



Review: k-NN

- Examine the k -“closest” training data points to new point x
- Closest depends on distance metric used
- Assign the object the most frequently occurring class (majority vote) or the average value (regression)



k-NN: User-based

- Intuition: Similar users will rate the item similarly
- Represent each user as incomplete vector of item ratings
- Find set of N users who are 'similar' to Joe's ratings
- Estimate Joe's ratings based on ratings of users in set N

k-NN: User-based

- What is the right distance metric then?

- Jaccard similarity: $D(\mathbf{x}, \mathbf{y}) = \frac{|\mathbf{x} \cap \mathbf{y}|}{|\mathbf{x} \cup \mathbf{y}|}$ ignores value of the rating

- Cosine similarity: $D(\mathbf{x}, \mathbf{y}) = \frac{\mathbf{x} \cdot \mathbf{y}}{\|\mathbf{x}\|_2 \|\mathbf{y}\|_2}$ missing ratings are “negative”

- Pearson correlation coefficient

$$D(\mathbf{x}, \mathbf{y}) = \frac{\sum_{s \in S_{xy}} (\mathbf{x}_s - \bar{\mathbf{x}})(\mathbf{y}_s - \bar{\mathbf{y}})^{\top}}{\sqrt{\sum_{s \in S_{xy}} (\mathbf{x}_s - \bar{\mathbf{x}})^2} \sqrt{\sum_{s \in S_{xy}} (\mathbf{y}_s - \bar{\mathbf{y}})^2}}$$

k-NN: Item-based

- Intuition: Users rate similar items similarly
- Represent each item as incomplete vector of user ratings
- Find other similar items
- Estimate rating for item based on ratings for similar items

k-NN: Item-based

users

	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

k-NN: Item-based

	users												
	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
3	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
6	1		3		3			2			4		<u>0.59</u>

use weighted average to predict

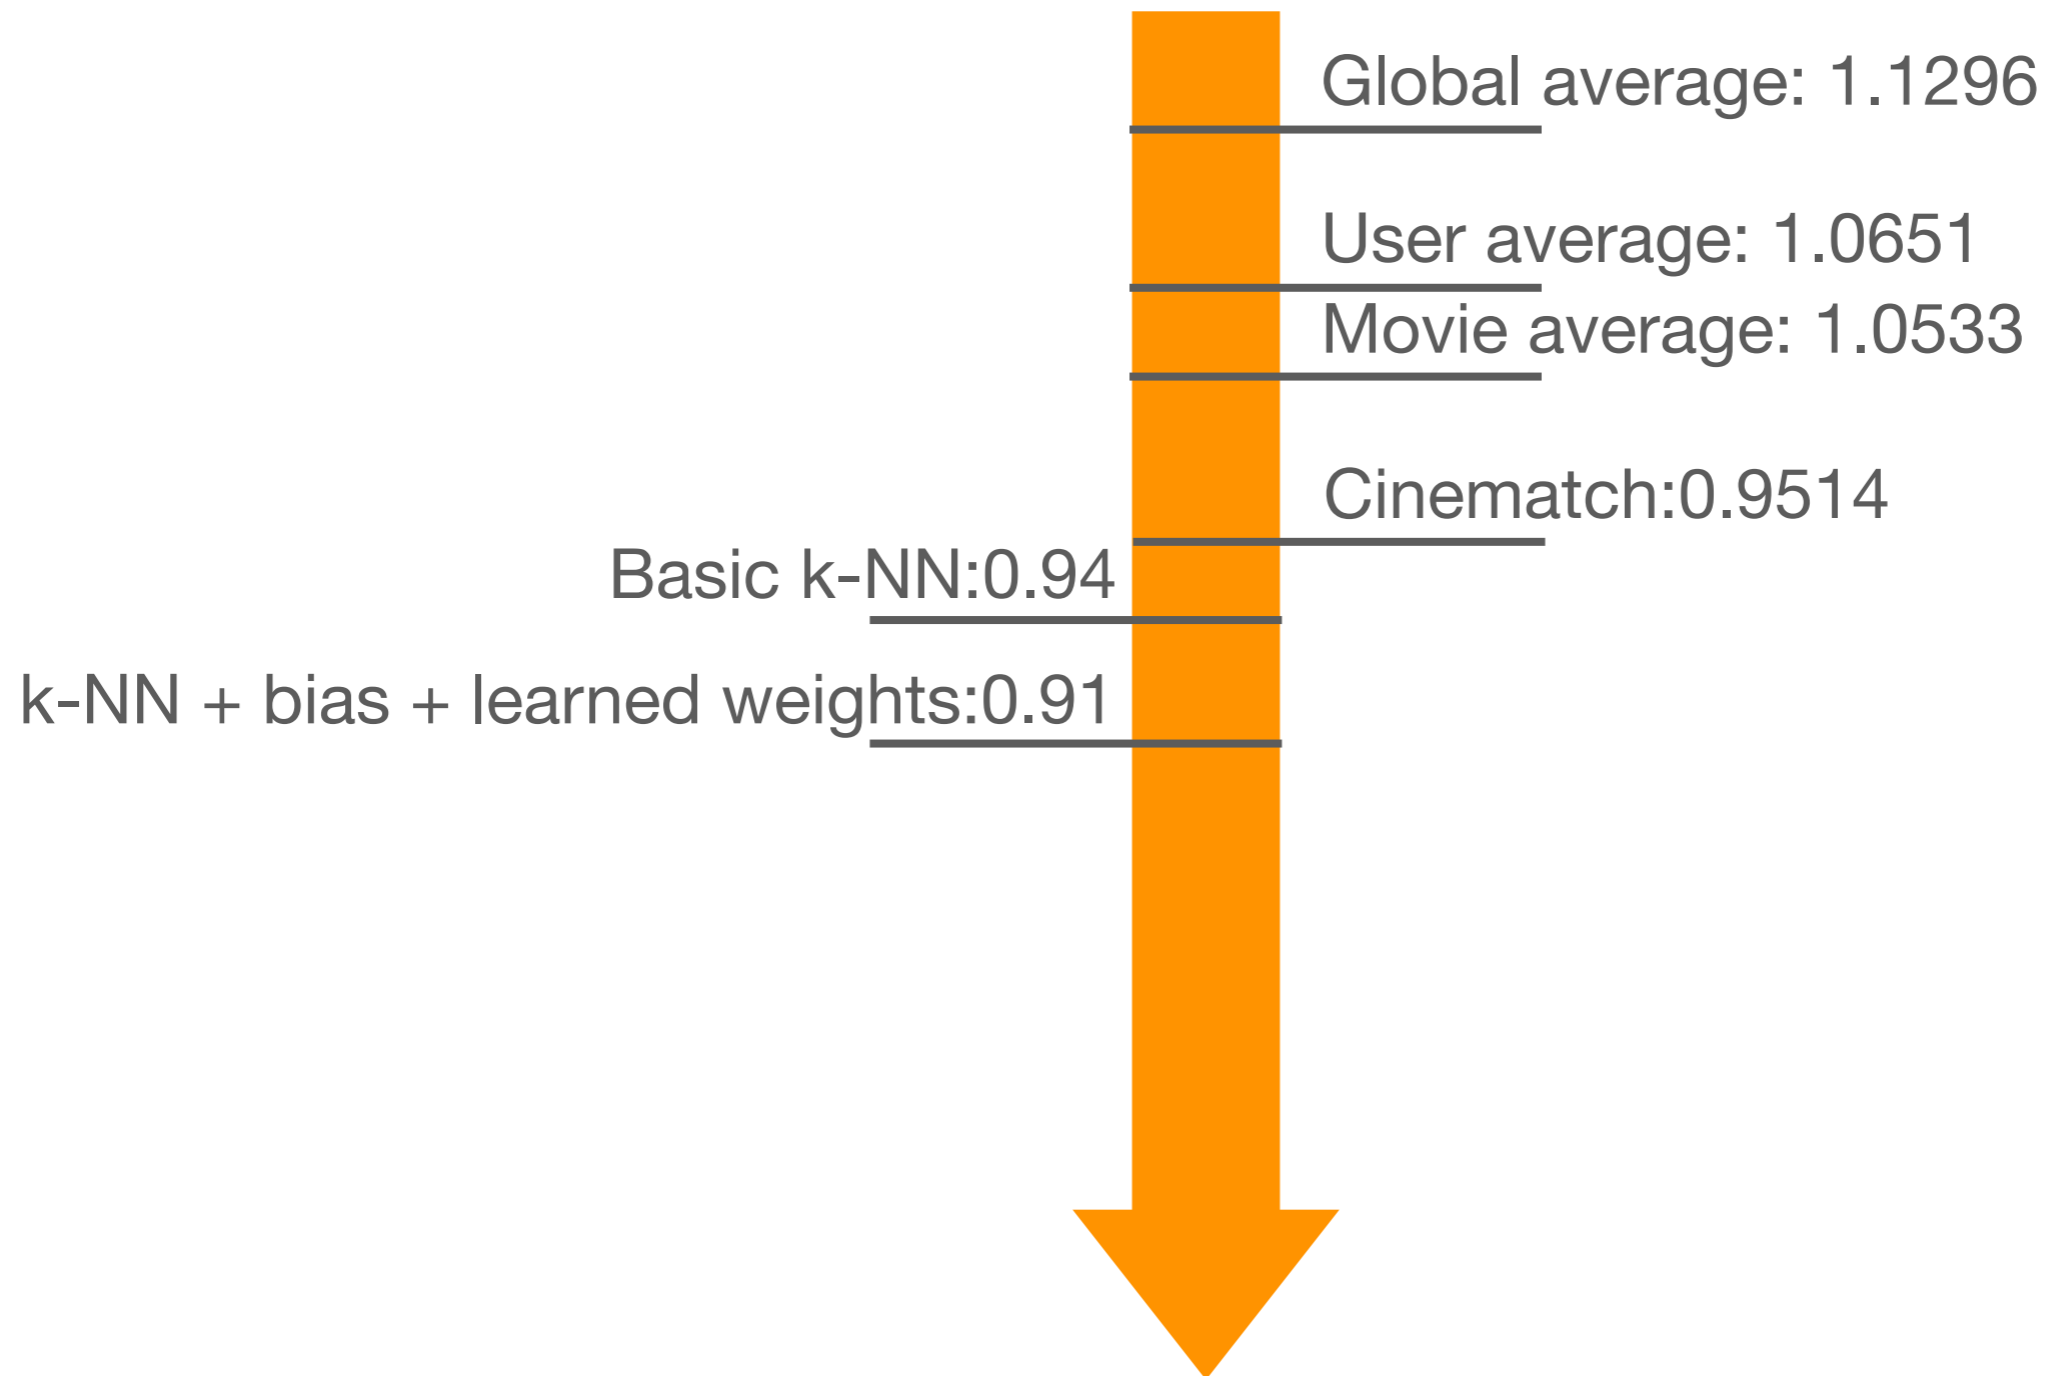
k-NN: Advantages

- Intuitive interpretation: you will like what your neighbors like
- Easy to implement and zero training time
- No feature selection needed — works for any kind of item

k-NN: Disadvantages

- Cold start
 - Need enough users in the system to find a match
 - New items and esoteric items may not have any ratings
- Sparse, high-dimensional similarity search is not easy
- Tends to recommend popular items
- Need to store all items or user vectors in memory

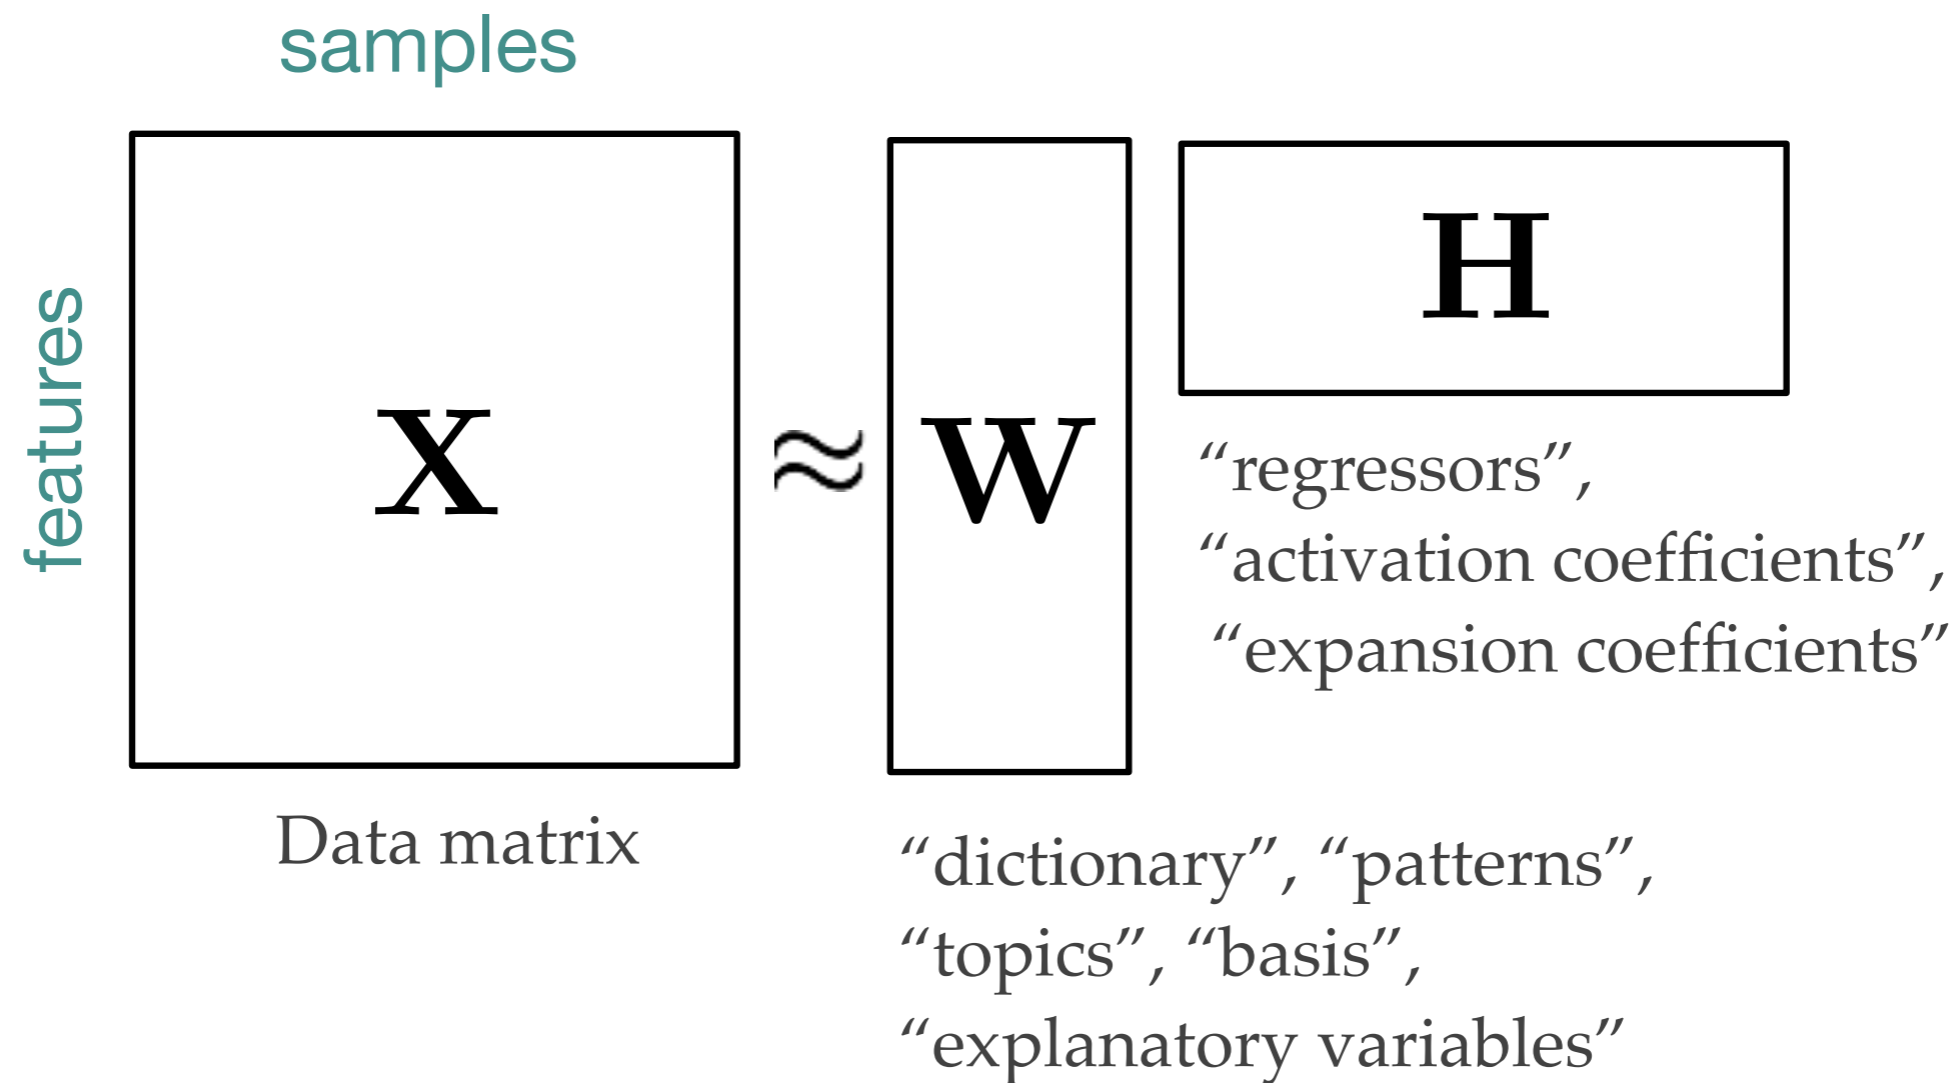
Netflix Performance



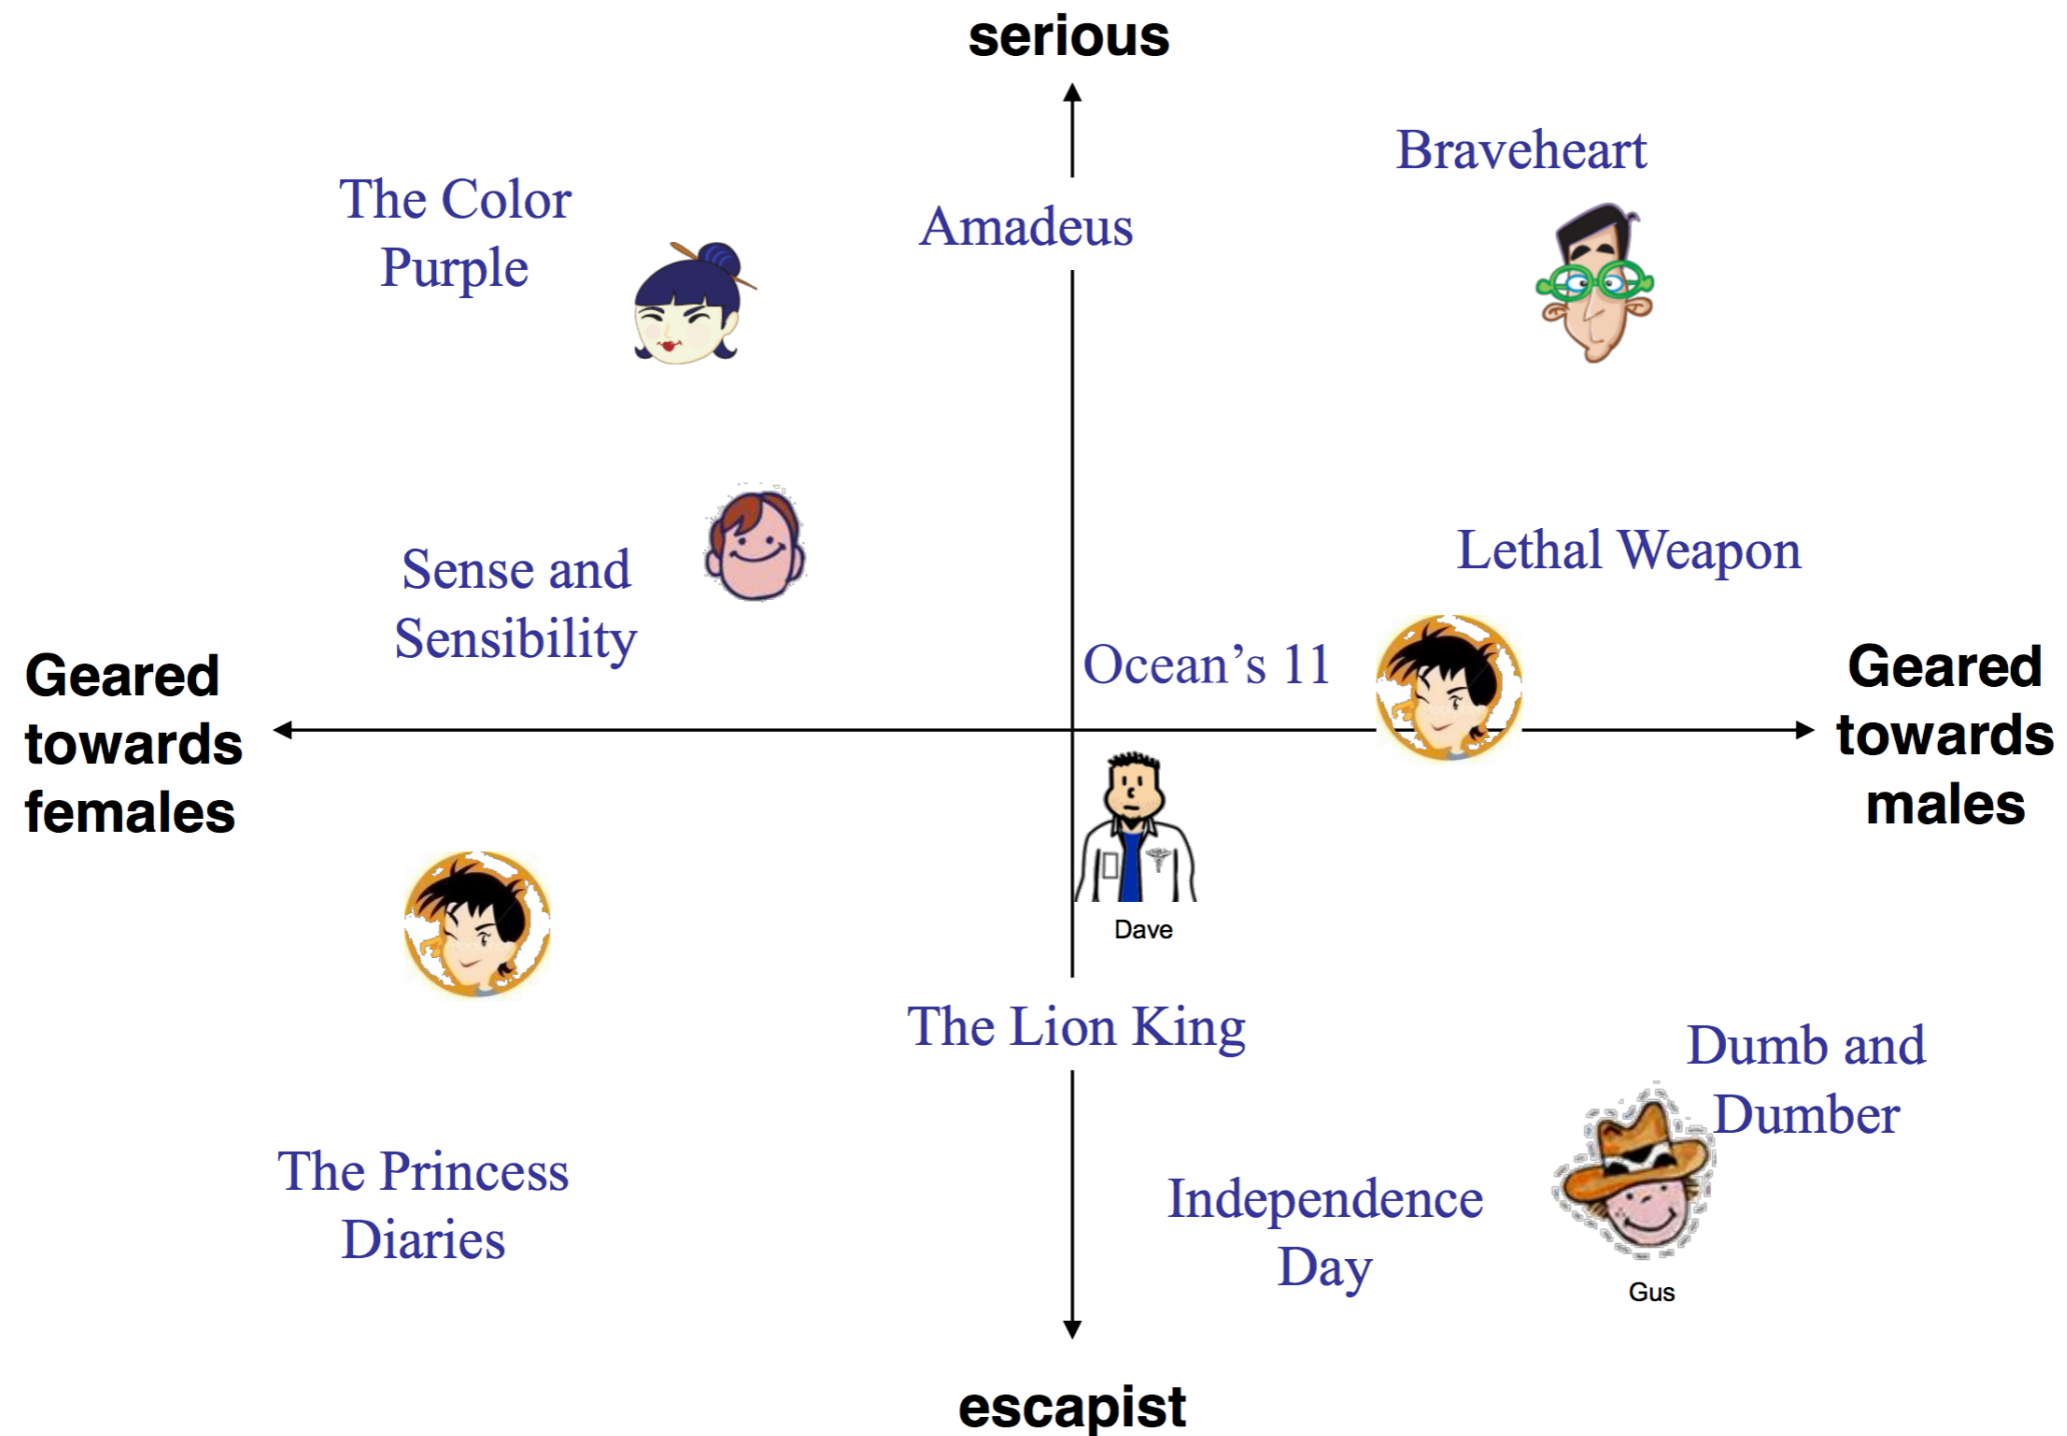
Review: Dimensionality Reduction

- Generate a low-dimensional encoding of a high-dimensional space
- Purposes:
 - Data compression / visualization
 - Robustness to noise and uncertainty
 - Potentially easier to interpret

Review: Matrix Factorization



Dimensionality Reduction



Review: SVD

- Each matrix can be decomposed using singular value decomposition (SVD):

$$\underbrace{\mathbf{X}}_{n \times p} = \underbrace{\mathbf{U}}_{n \times p} \underbrace{\mathbf{D}}_{p \times p} \underbrace{\mathbf{V}}_{p \times p}^T$$


orthonormal columns
which are principal
components

orthonormal columns
which are normalized
PC scores

diagonal matrix which if each
diagonal element is squared
and divided by n gives
variance explained

SVD to MF

- Create two new matrices (user and item matrices) where the square root of the singular values are distributed to each matrix U and V



- Interpretation:
 - p_u indicates how much user likes each latent factor f
 - q_i means the contribution of item to each of the latent factors f

RecSys: SVD

- SVD is great as it minimizes SSE which is monotonically related to RMSE
- Conventional SVD is undefined for missing entries
 - No rating can be interpreted as zero rating — is that right?

RecSys: SVD

- One idea: Expectation maximization as form of imputation
 - Fill in unknown entries with best guess
 - Apply SVD
 - Repeat
- Can be expensive and inaccurate imputation can distort data

SVD w/ Missing Values

- New idea: Model only the observed entries and avoid overfitting via regularization

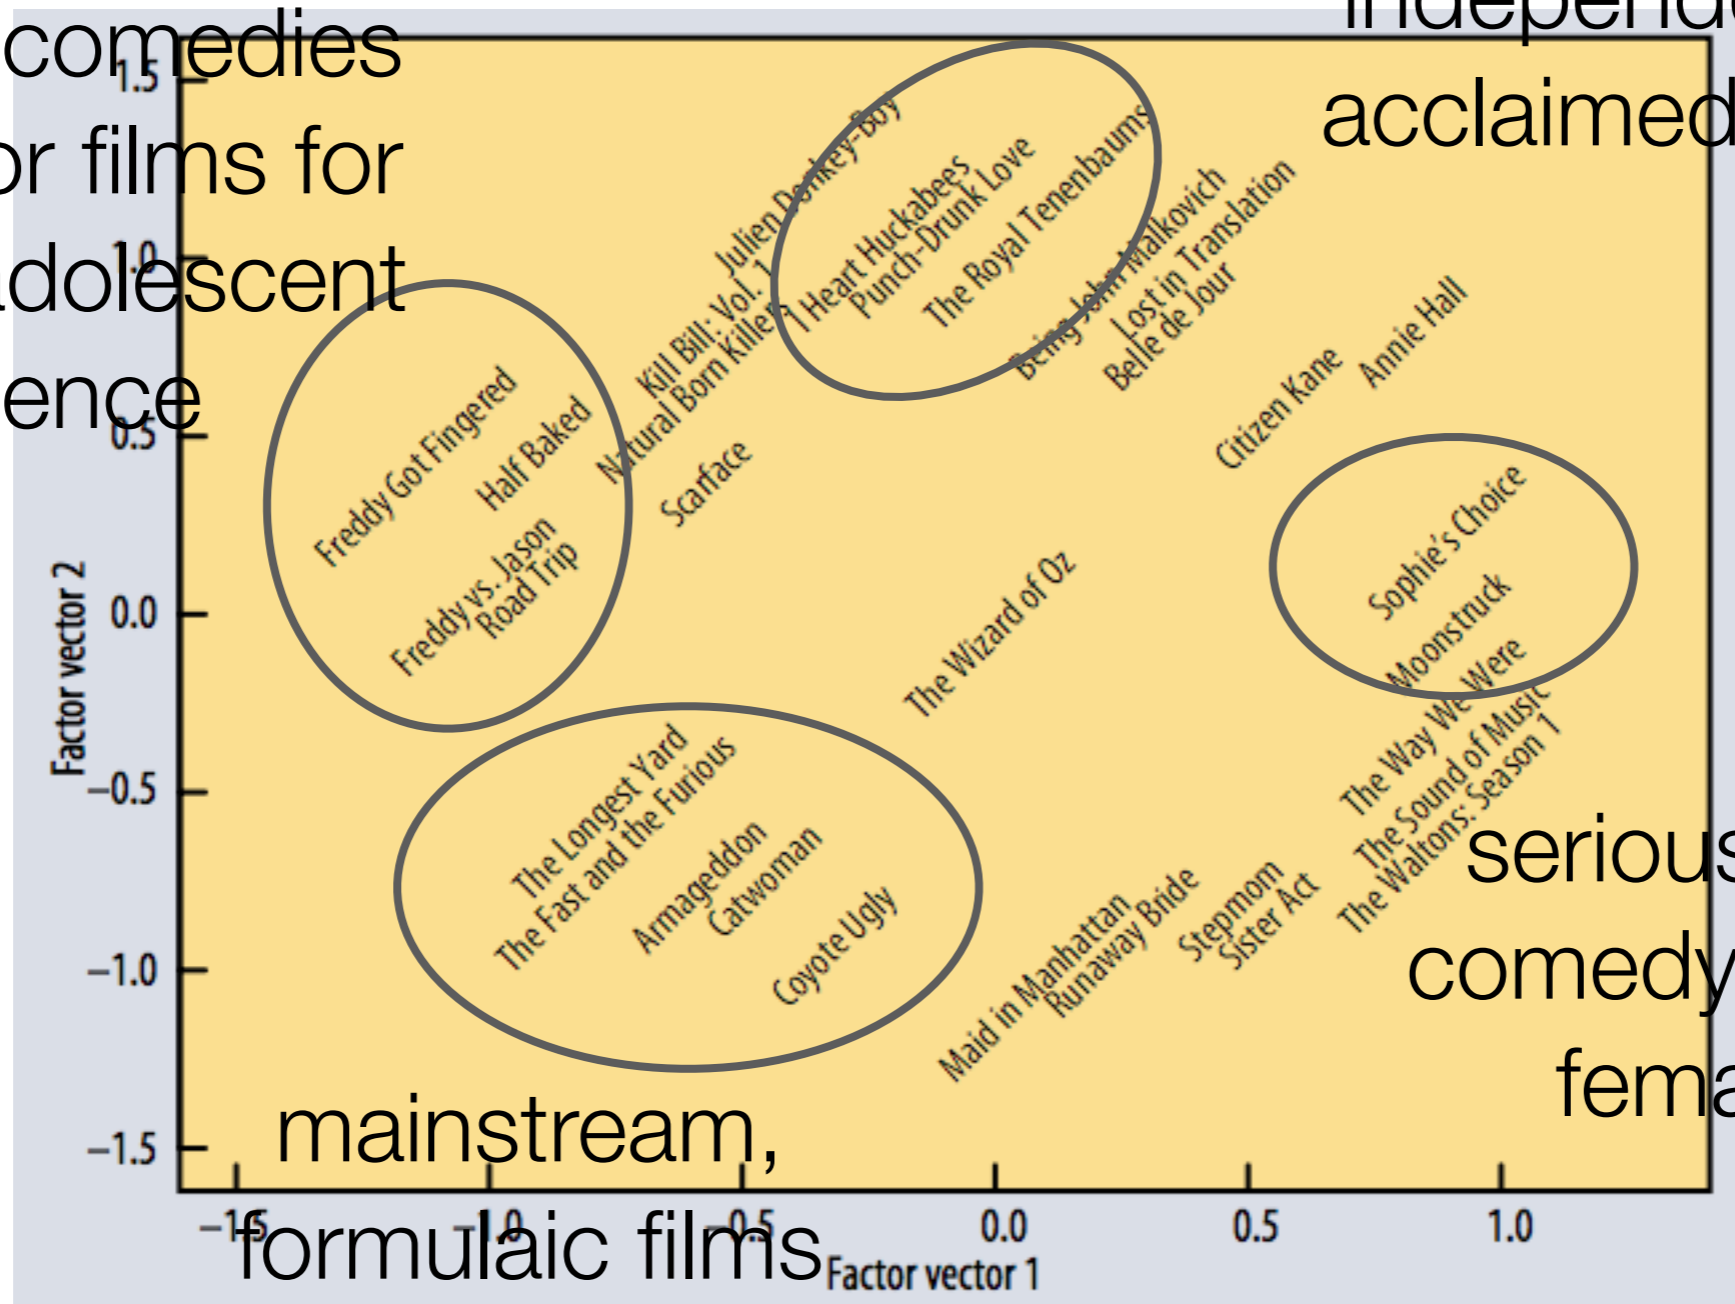
$$\min_{q^*, p^*} \sum_{(u,i) \in \mathcal{K}} (r_{ui} - q_i^T p_u)^2 + \lambda(\|q_i\|^2 + \|p_u\|^2)$$

- Two methods for solving the new model
 - Stochastic gradient descent
 - Alternating least squares - easier to parallelize as each q_i is independent and more memory efficient

Netflix Results: Latent Factors

lowbrow comedies
and horror films for
male or adolescent
audience


independent, critically
acclaimed, quirky films



mainstream,
formulaic films

serious drama or
comedy with strong
female leads

SVD with Bias



Baseline predictor


- Separates users and movies
- Benefits from insights into user's behavior
- Among the main practical contributions of the competition

User-Movie interaction

- Characterizes the matching between users and movies
- Attracts most research in the field
- Benefits from algorithmic and mathematical innovations

$$\begin{aligned} \text{minimize}_{p,q} \quad & \sum_{(u,i) \in S} (r_{ui} - (\mu + b_u + b_i + \langle p_u, q_i \rangle))^2 + \\ & \lambda \left[\|p\|_{\text{Frob}}^2 + \|q\|_{\text{Frob}}^2 + \|b_{\text{users}}\|^2 + \|b_{\text{items}}\|^2 \right] \end{aligned}$$

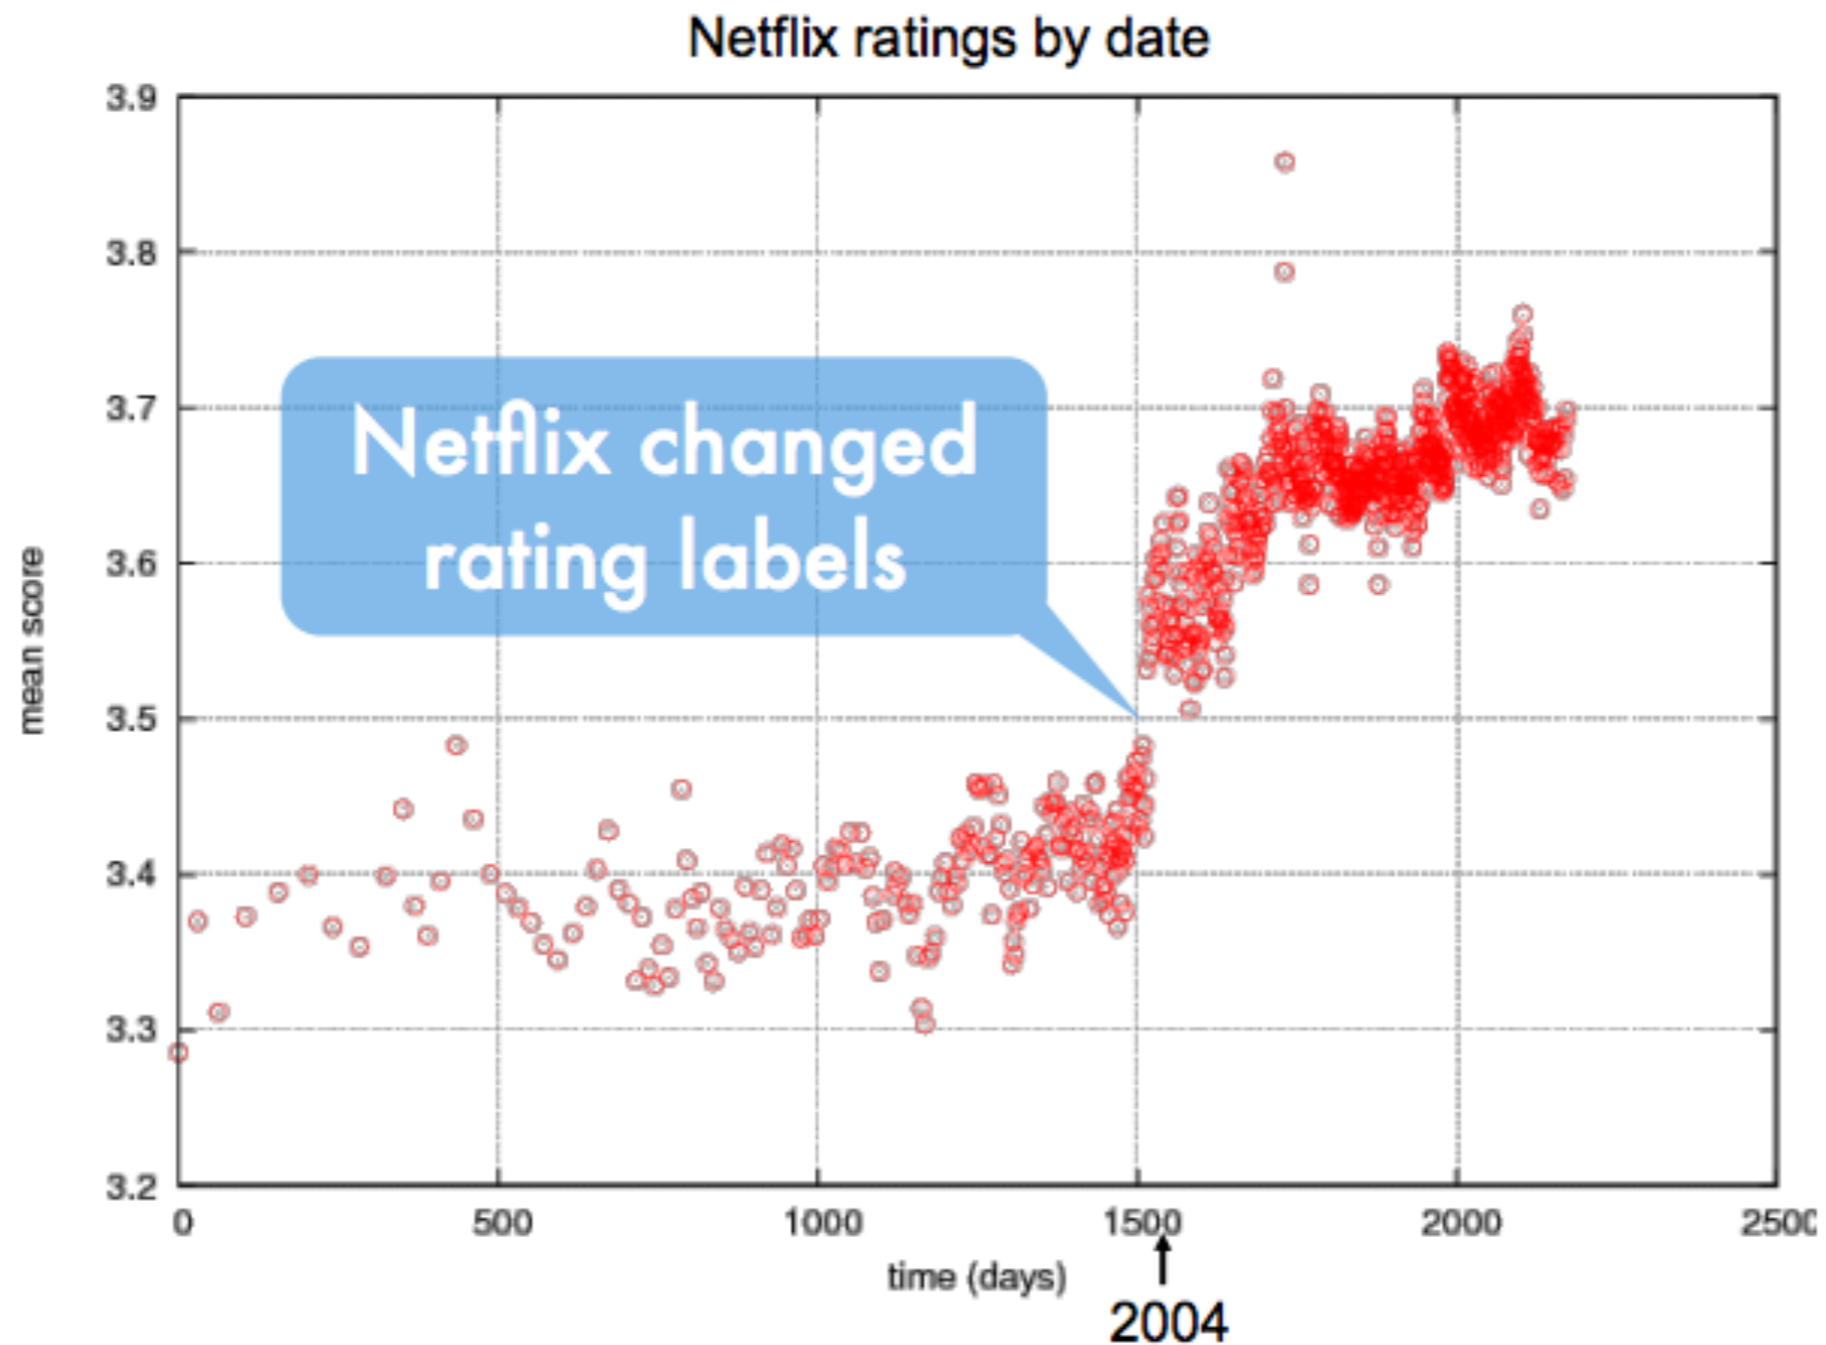
Netflix Performance



Implicit Feedback

- May have access to binary information reflecting implicit user preferences
 - Is a movie in a user's queue?
- Test source can be source — we know that user u rated item i , just don't know the rating
 - Data is not “missing at random”
 - Fact that user rated item provides information

Netflix: Temporal Bias



Netflix: Temporal Bias

- Items
 - Seasonal effects
 - Public perception (Oscars, SAG, etc.)
 - Grow and fade in popularity
 - ...
- Users
 - Changed review labels
 - Anchoring (relative to previous movie)
 - Selection bias for time of viewing
 - ...

Temporal SVD


linear modeling of
user biases

movie-related
temporal effects


$$\min \sum_{(u,i,t) \in \mathcal{K}} (r_{ui}(t) - \mu - b_u - \alpha_u \text{dev}_u(t) - b_{u,t} - b_i - b_{i, \text{Bin}(t)})^2 + \lambda(b_u^2 + \alpha_u^2 + b_{u,t}^2 + b_i^2 + b_{i, \text{Bin}(t)}^2)$$

single day effect —
sudden “spikes”

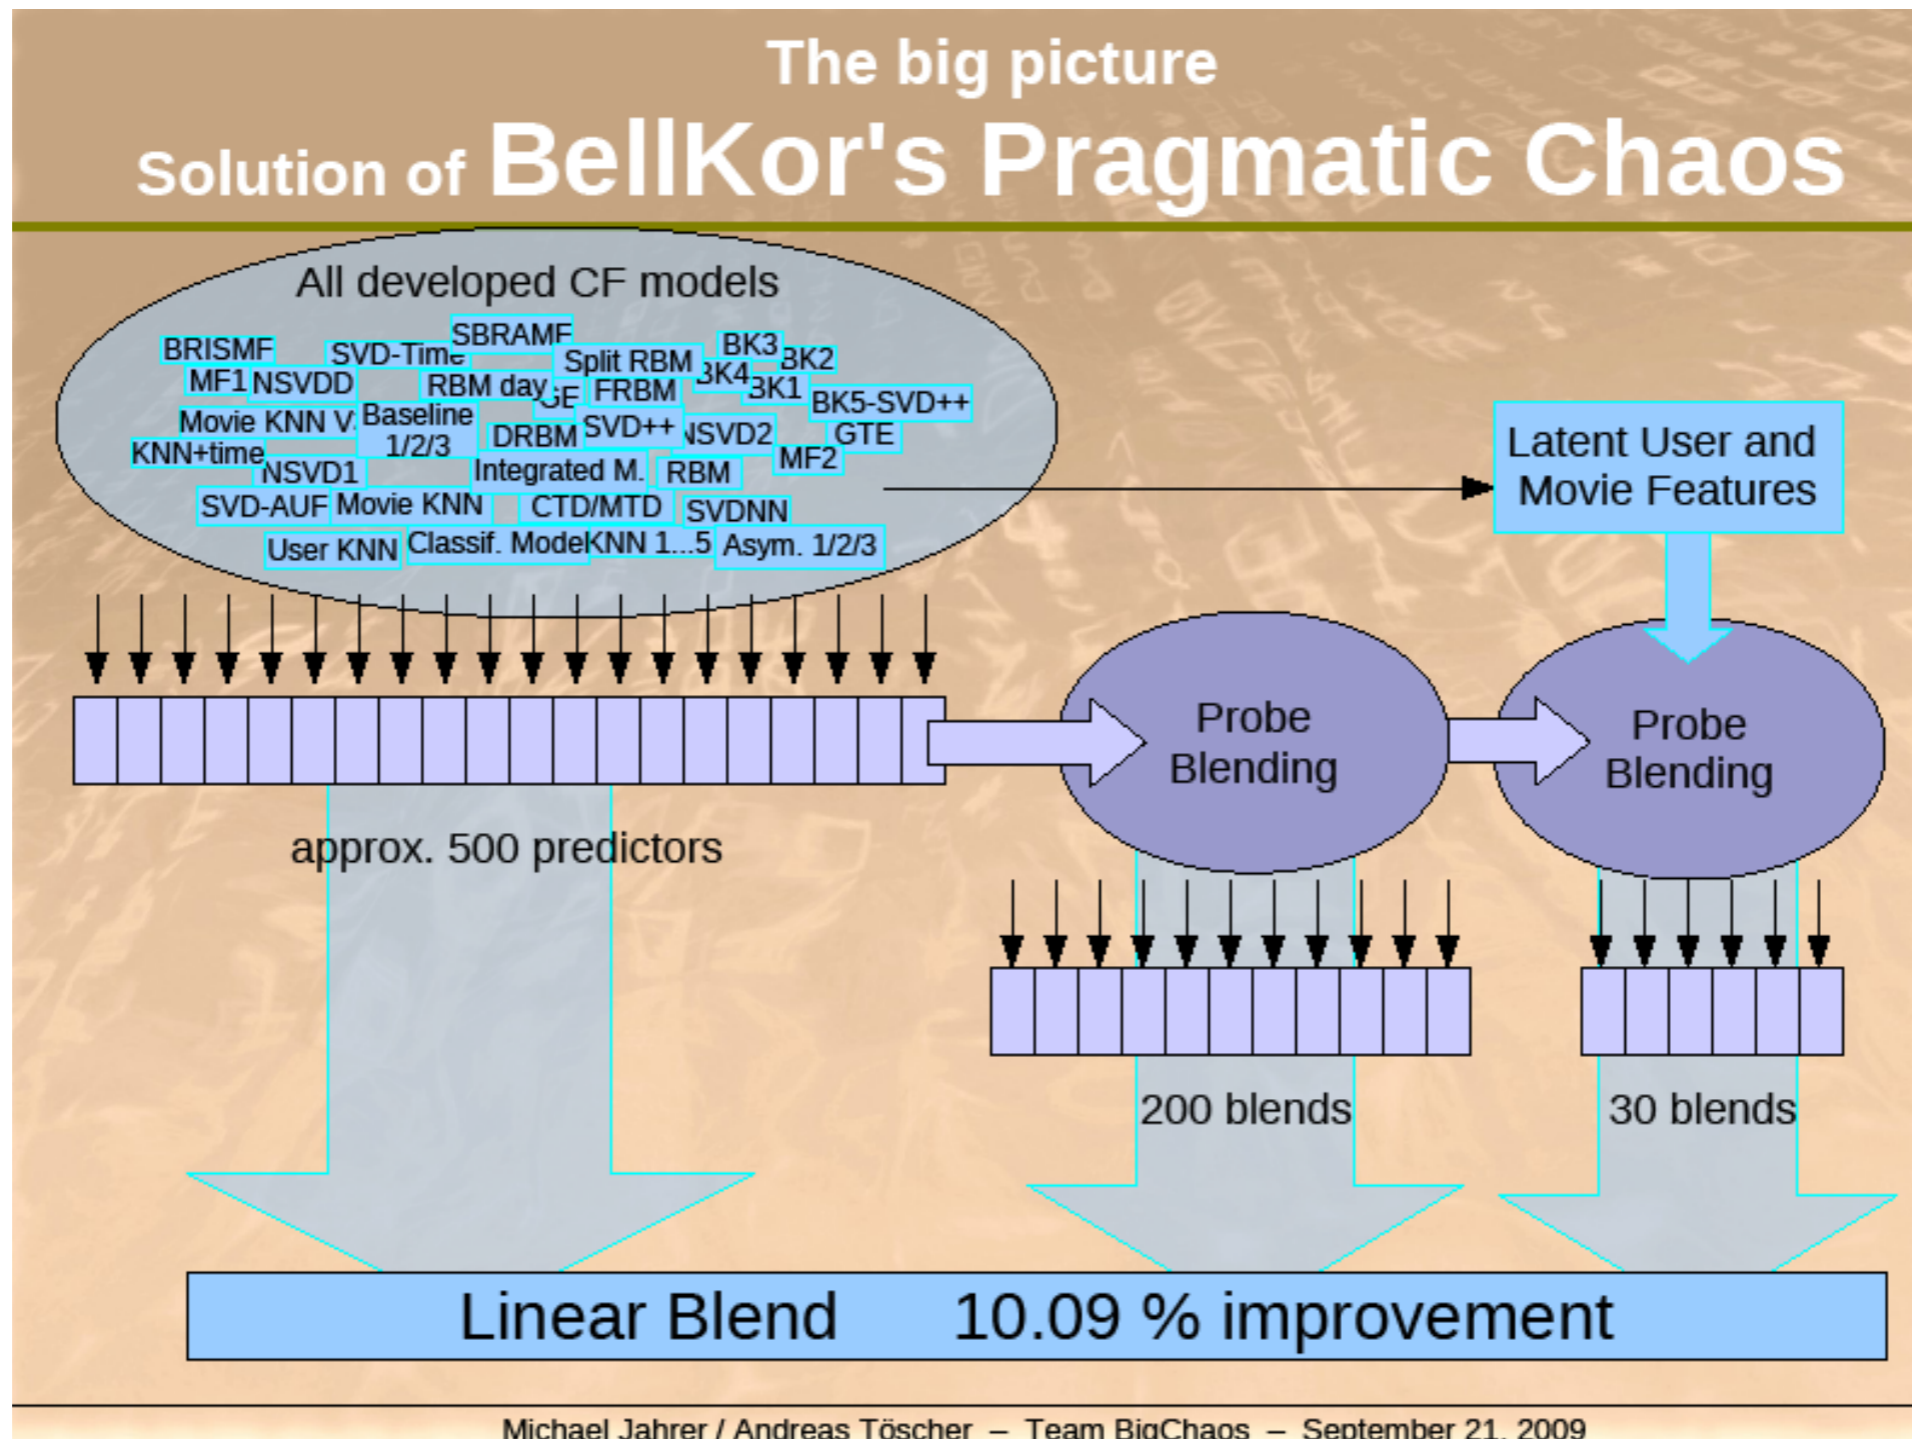
Netflix Performance



Netflix Results: RMSE



Final Solution: Kitchen Sink Approach



Winning Solution

- Beat Netflix by 10.06% in RMSE
- Tied with another team but won because submitted 20 minutes earlier
- Computationally intensive and impractical



Many More Ideas

- Cold start (new users)
- Different regularization for different parameter groups and differs users
- Sharing of statistical strength between users
- Hierarchical matrix co-clustering / factorization
- Incorporate social network, user profiles, item profiles
- ...

RecSys: Challenges

- Relevant objectives
 - Predicting actual rating may be useless!
 - May care more about ranking of items
- Missing at random assumption
 - How can our models capture information in choices of our ratings?
- Handling users and items with few ratings

RecSys: Challenges

- Multiple individuals using the same account — individual preference
- Preference versus intention
 - Distinguish between liking and interested in seeing / purchasing
 - Worthless to recommend an item a user already has
- Scalability
 - Simple and parallelizable algorithms are preferred