Course Update

Out

Due

Current Plan (updated)

- HW 8 (online)
- Mini-project proposal
- HW 9 (online)
- HW 10 (written/prog)
- Midterm 2
- Mini-project

Sun	Mon	Tue	Wed	Thu	Fri	Sat
2	3	4	5 HW8 Proj	6	7	8
9 HW8	10 HW9 HW10	11	12	13	14	15
16	17 HW9	18	19 Prop	20	21	22 HW10
23	24	25	26 MT2	27	28	29
30	1	2	3	4	5 Proj	6

Plan

Last time

- Unsupervised Learning: Dimensionality Reduction

Today

- Recommender Systems
- Unsupervised Learning: Clustering
  - K-means

Next time

- Unsupervised Learning: Clustering
  - Gaussian mixture models and expectation maximization

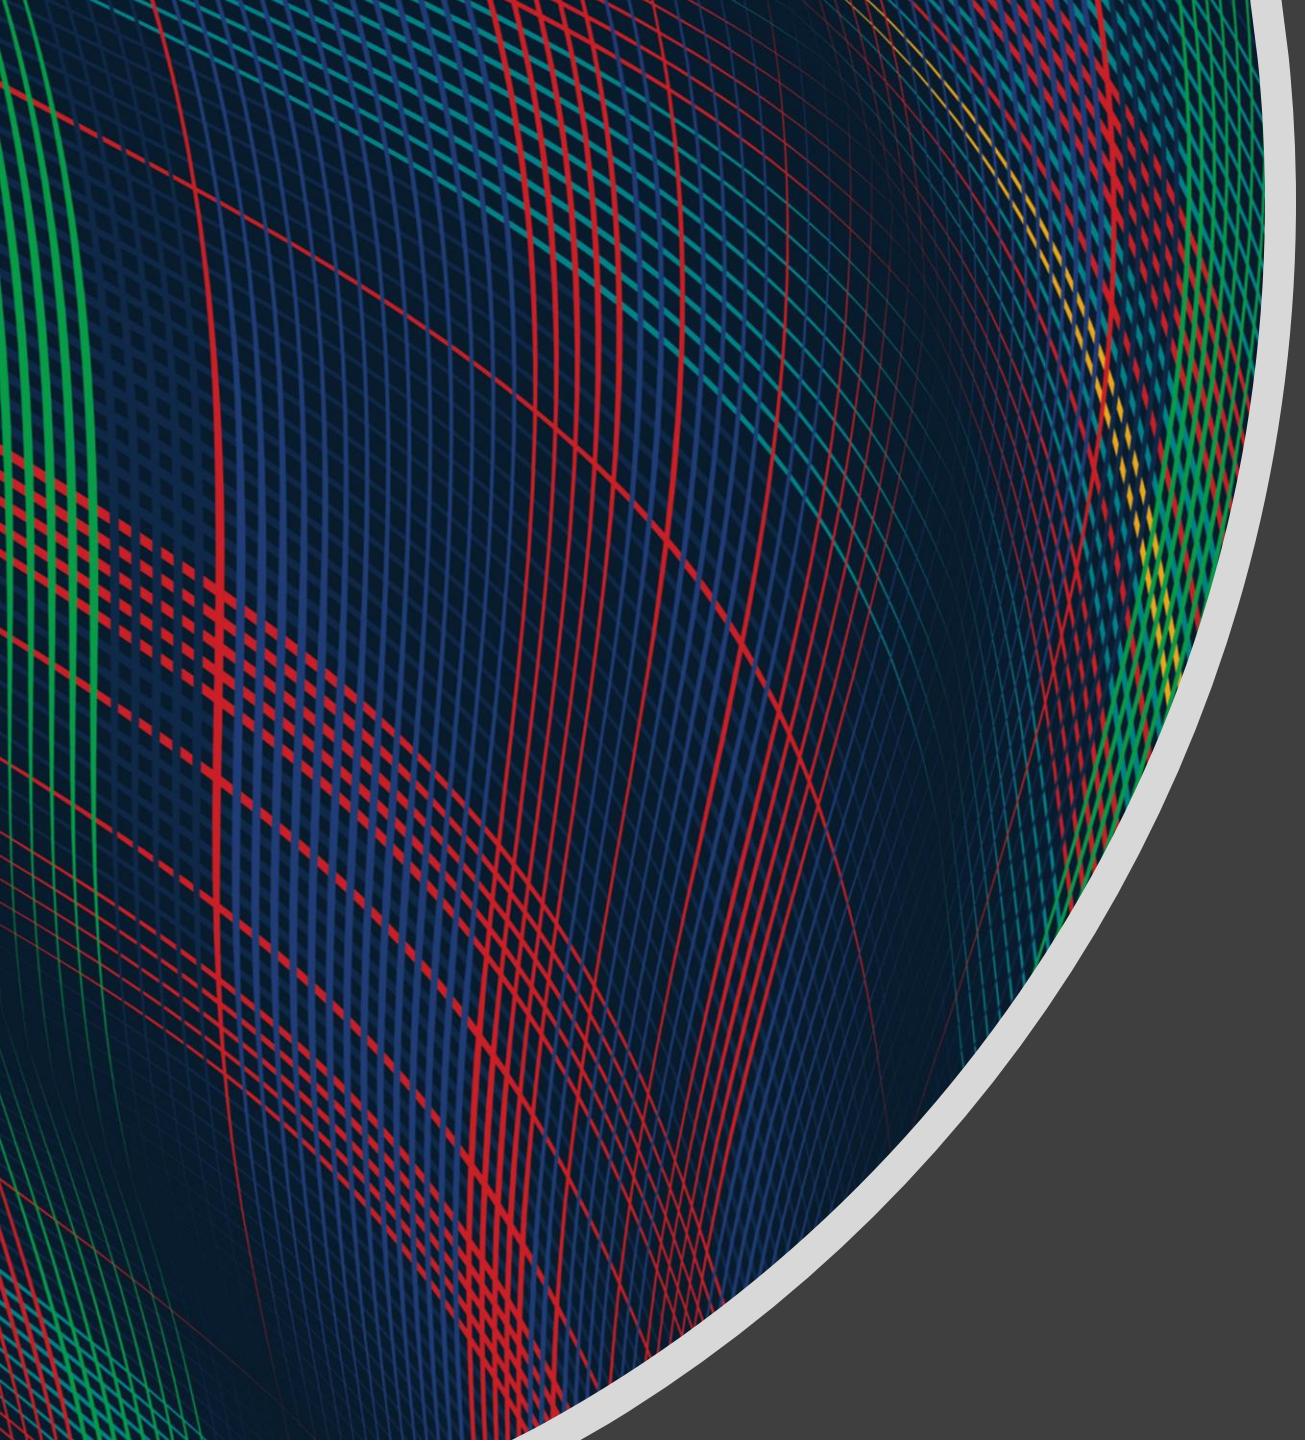

10-315 Introduction to ML

Recommender Systems

Instructor: Pat Virtue

Recommender Systems

A Common Challenge:

- Assume you're a company selling items of some sort:
  movies, songs, products, etc.
- Company collects millions of ratings from users of their items
- To maximize profit / user happiness, you want to
  recommend items that users are likely to want

Slide credit: CMU MLD Matt Gormley

Recommender Systems

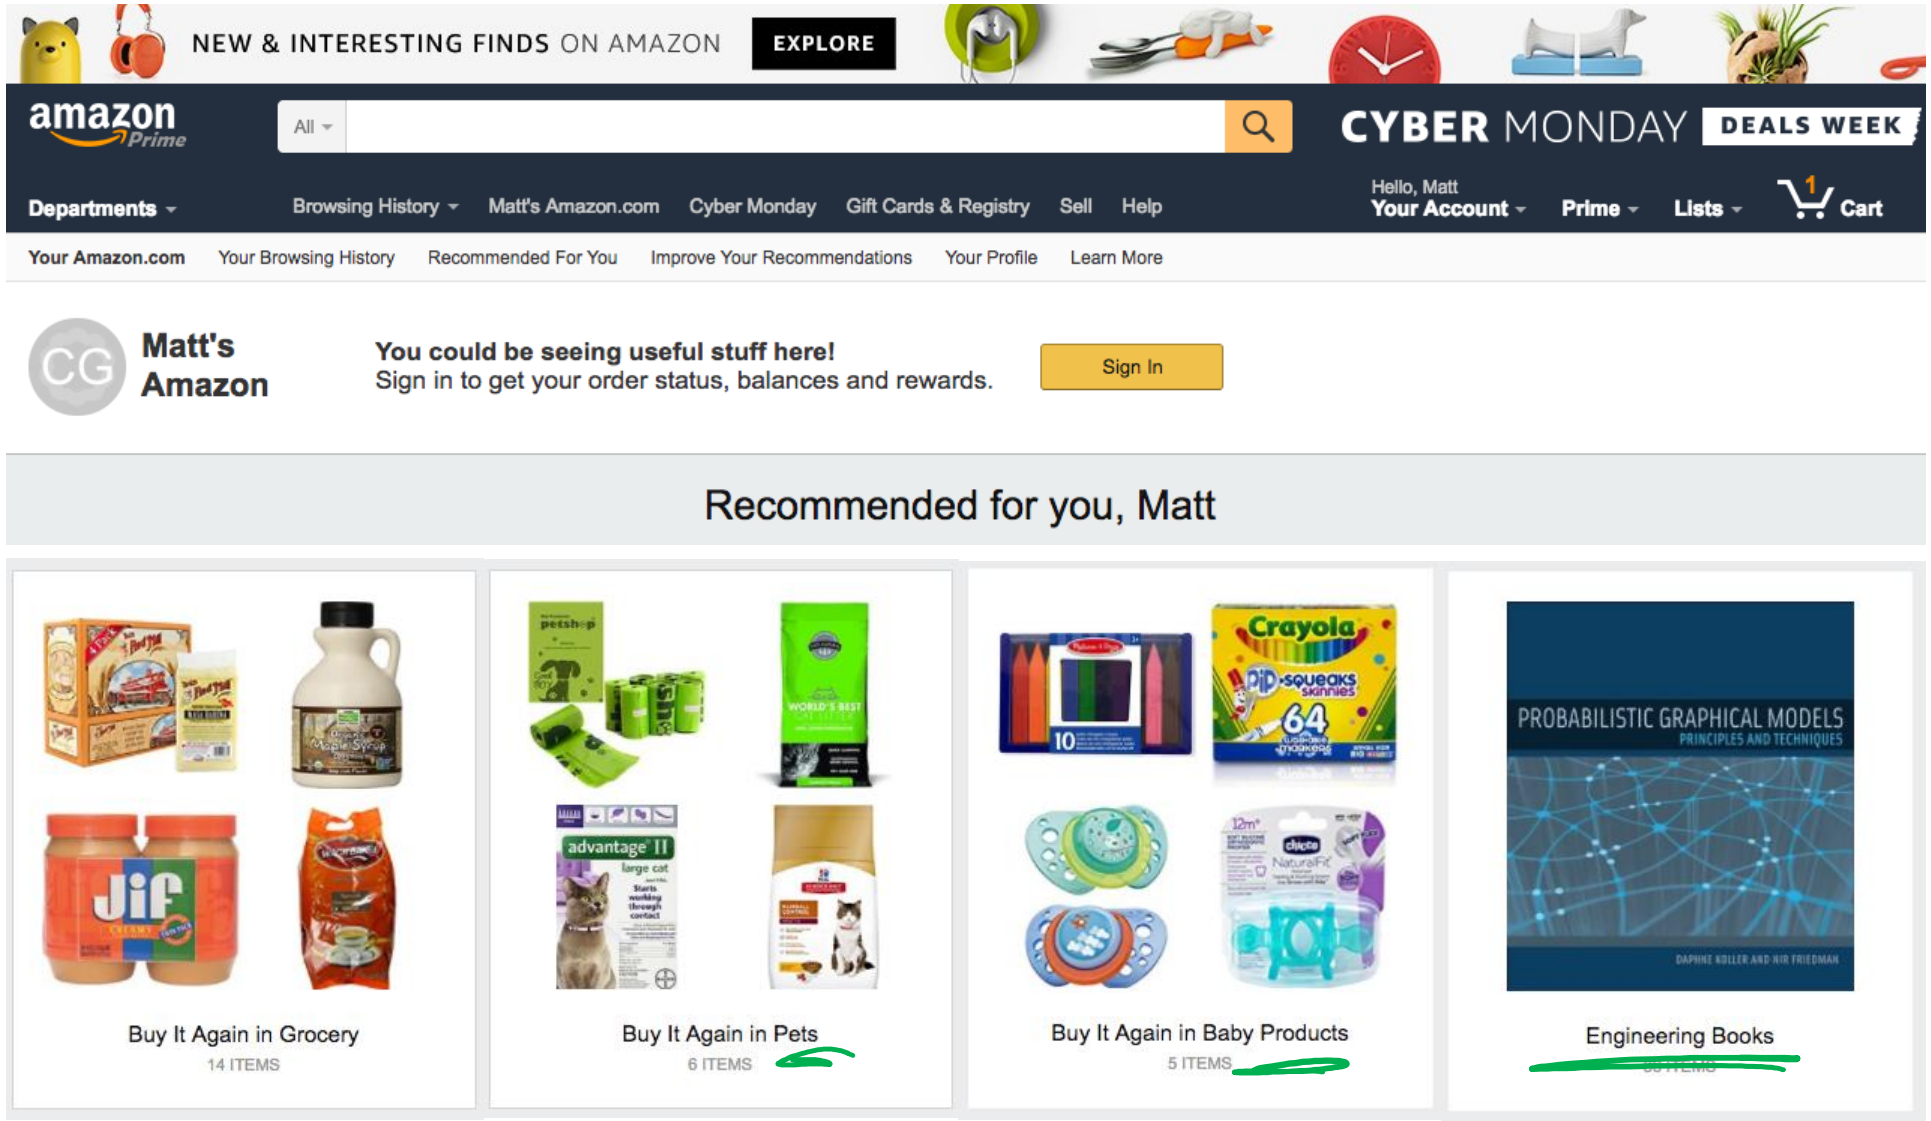

Slide credit: CMU MLD Matt Gormley

Recommender Systems

Slide credit: CMU MLD Matt Gormley

Recommender Systems

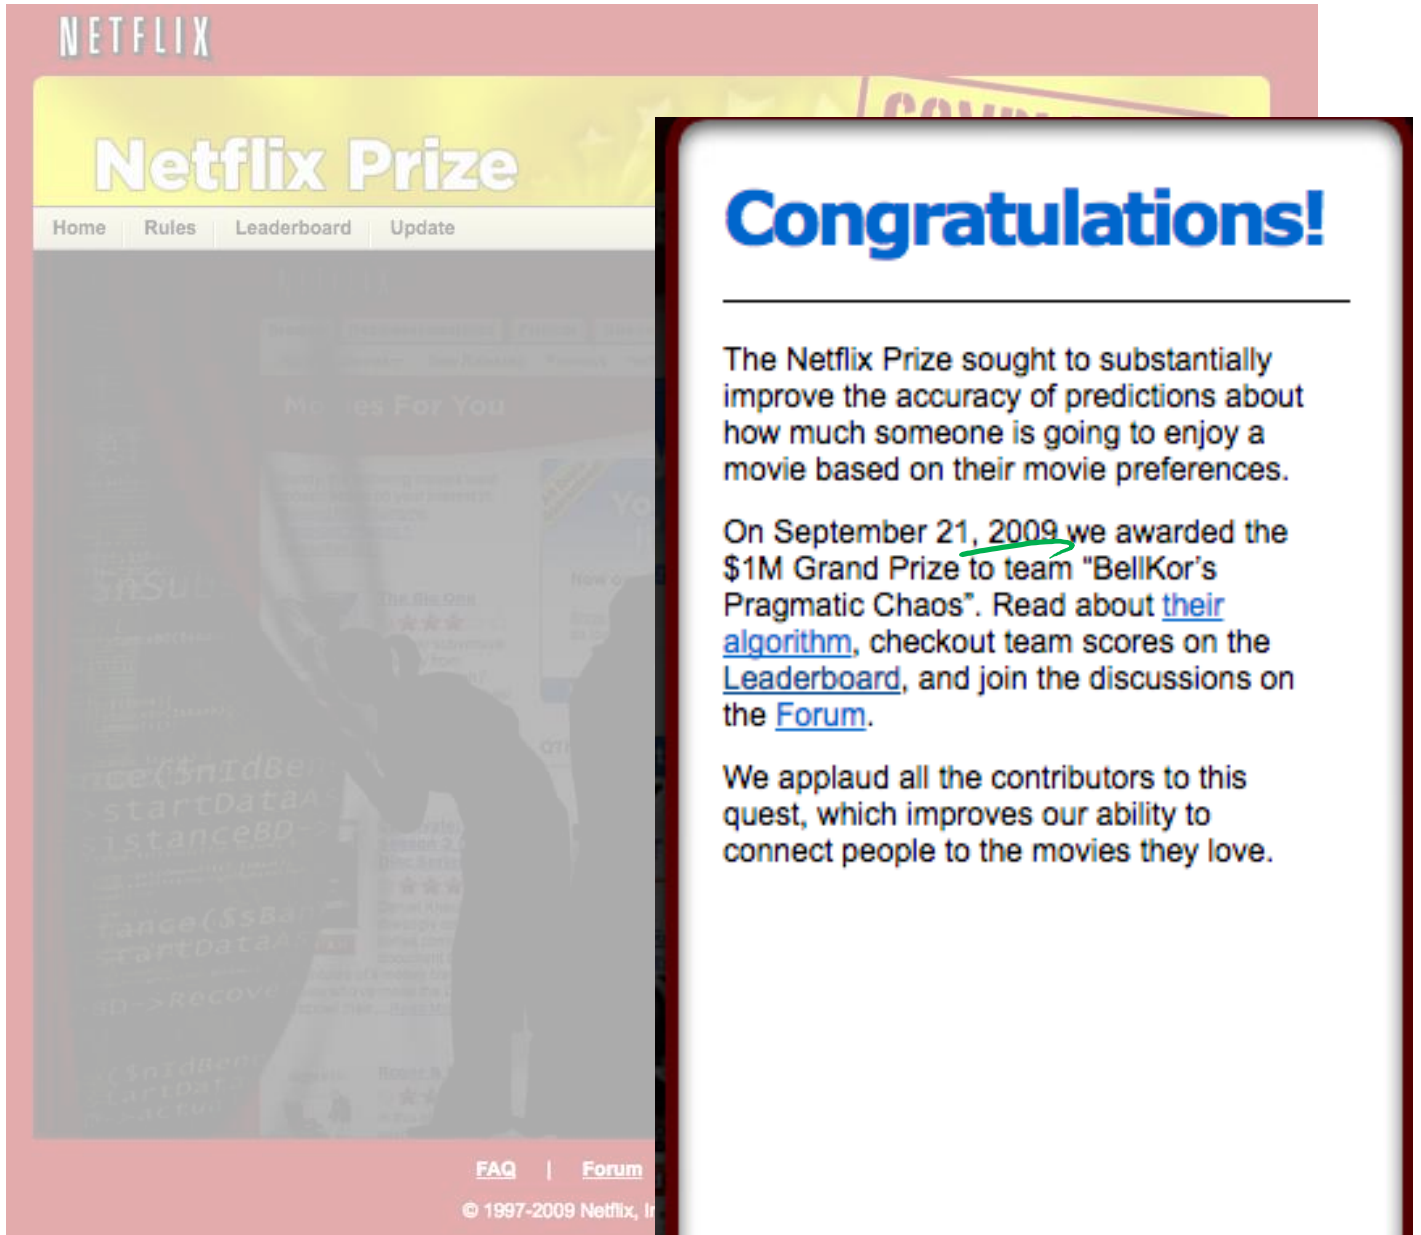

$1M

Slide credit: CMU MLD Matt Gormley

AI System Design: Movie Recommendation

Task

h(input) → output

user, item → rating 1-5

Experience

Data (training)

N (i, j, r)

Performance measure

output

Y y

Recommender Systems

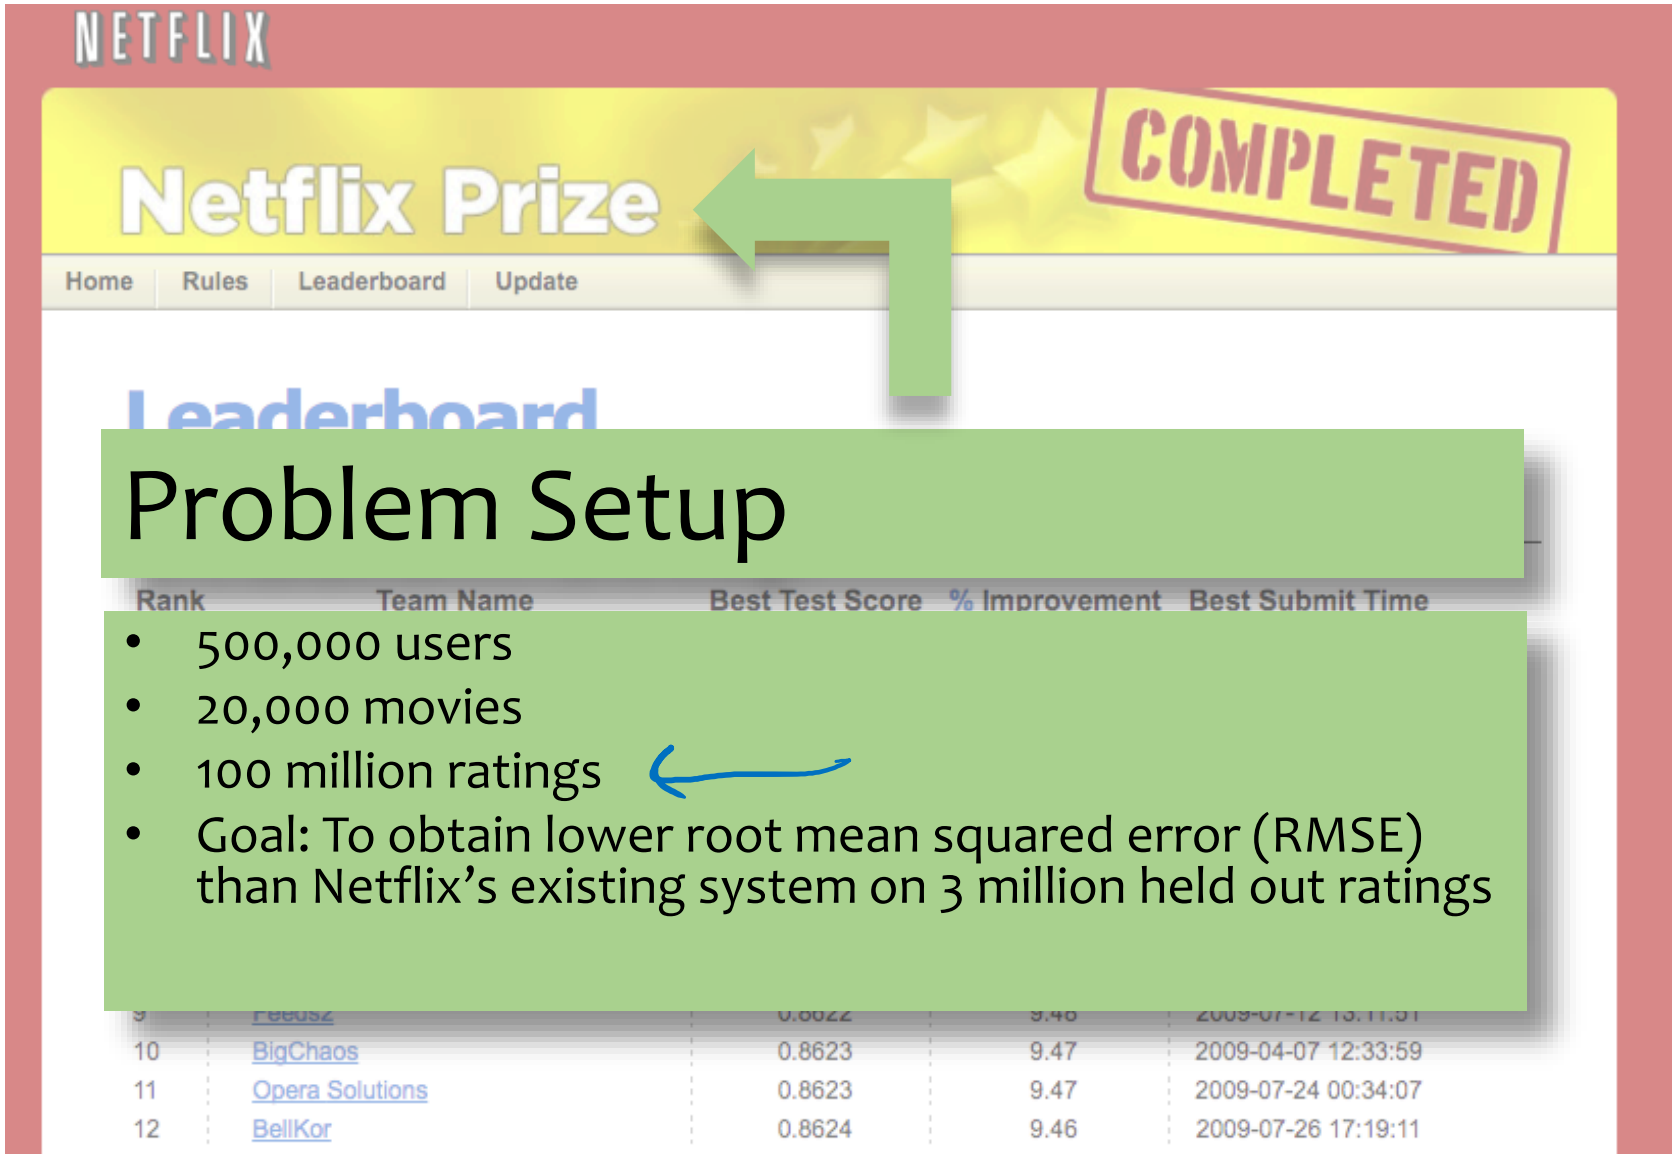

Slide credit: CMU MLD Matt Gormley

ML System Design: Movie Recommendation

Model

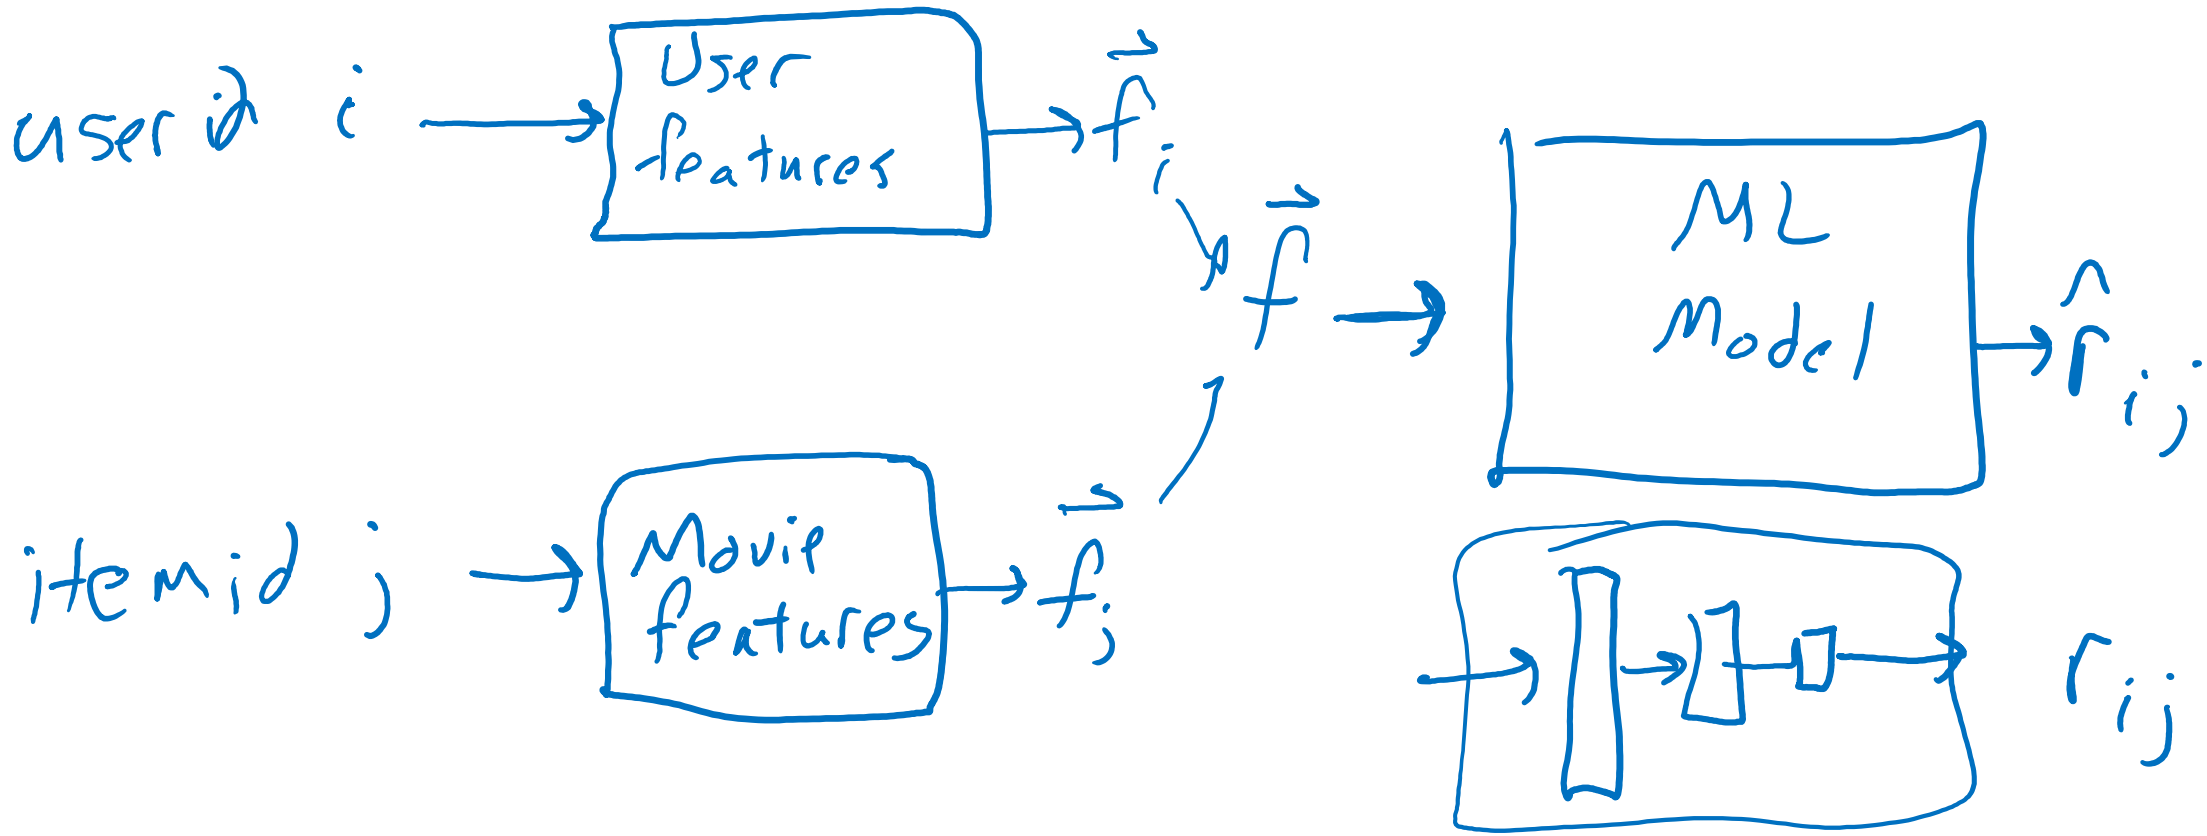

ML System Design: Movie Recommendation

Model

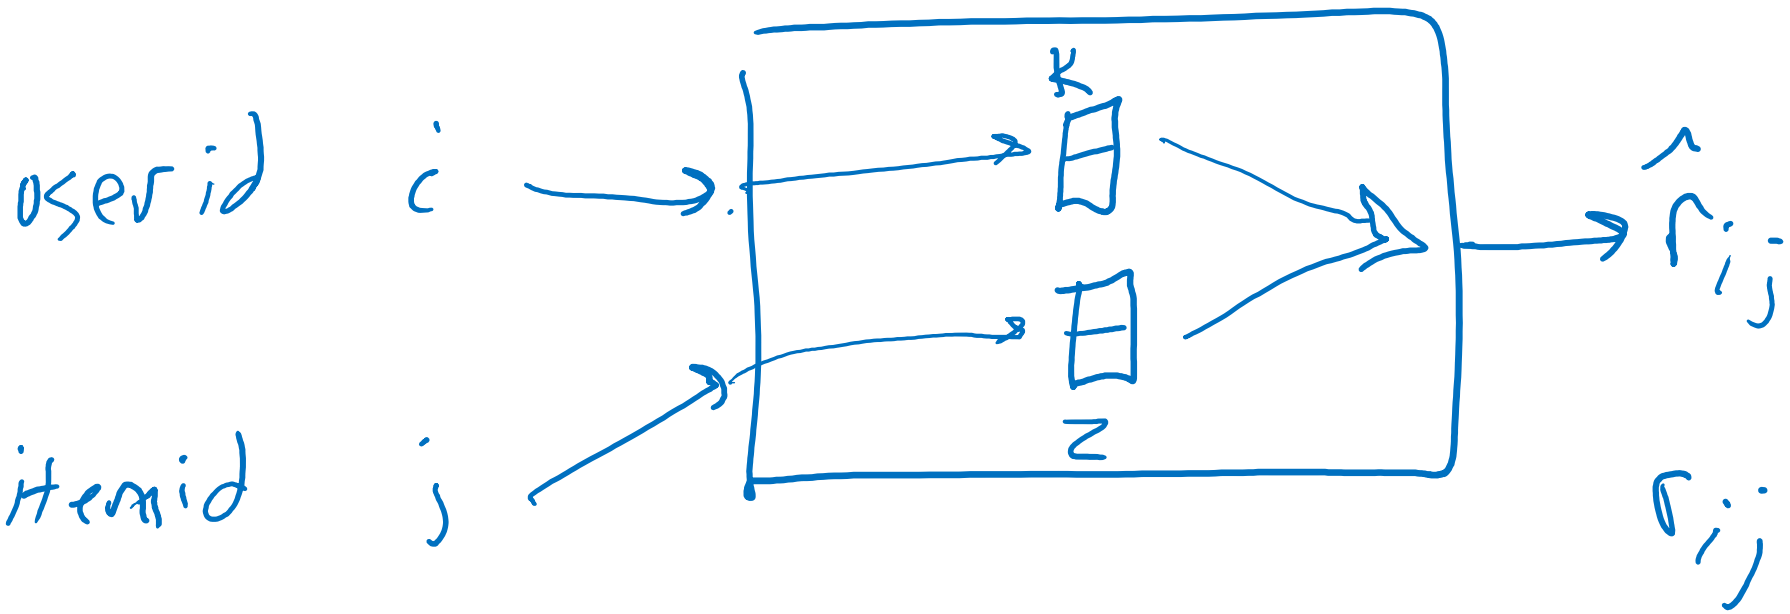

Recommender Systems

Setup:

- Items:
  movies, songs, products, etc.
  (often many thousands)
- Users:
  watchers, listeners, purchasers, etc.
  (often many millions)
- Feedback:
  5-star ratings, not-clicking ‘next’, purchases, etc.

Challenge:

- Users only rate a small number of items
  (the user/item rating data is sparse)

	Doctor Strange	Star Trek: Beyond	Zootopia
Alita	1		5
BB-8	3	4	
C-3P0	3	5	2
Pat		1	5

Slide credit: CMU MLD Matt Gormley

Different Approaches

Item-based (*Content filtering*)

- Features about each item
- Given an item, other “close” items have similar values
- e.g. Pandora.com, music genome project

Different Approaches

Item-based (*Content filtering*)

- Features about each item
- Given an item, other “close” items have similar values
- e.g. Pandora.com, music genome project

User-based

- Features about each user
- Given a user, other “close” users have similar preferences
- *Market segmentation*

Learning user-item relationship

- Can be done without features on either user or item
- Collaborative filtering techniques

Collaborative Filtering

Collaborative Filtering

Everyday Examples of Collaborative Filtering...

- Bestseller lists
- Top 40 music lists
- The “recent returns” shelf at the library
- Unmarked but well-used paths thru the woods
- The printer room at work
- “Read any good books lately?”
- …

Common insight: personal tastes are correlated

- If Alita and BB-8 both like X and Alita likes Y then BB-8 is more likely to like Y
- especially (perhaps) if BB-8 knows Alita

Slide from William Cohen

Two Types of Collaborative Filtering

1. Neighborhood Methods

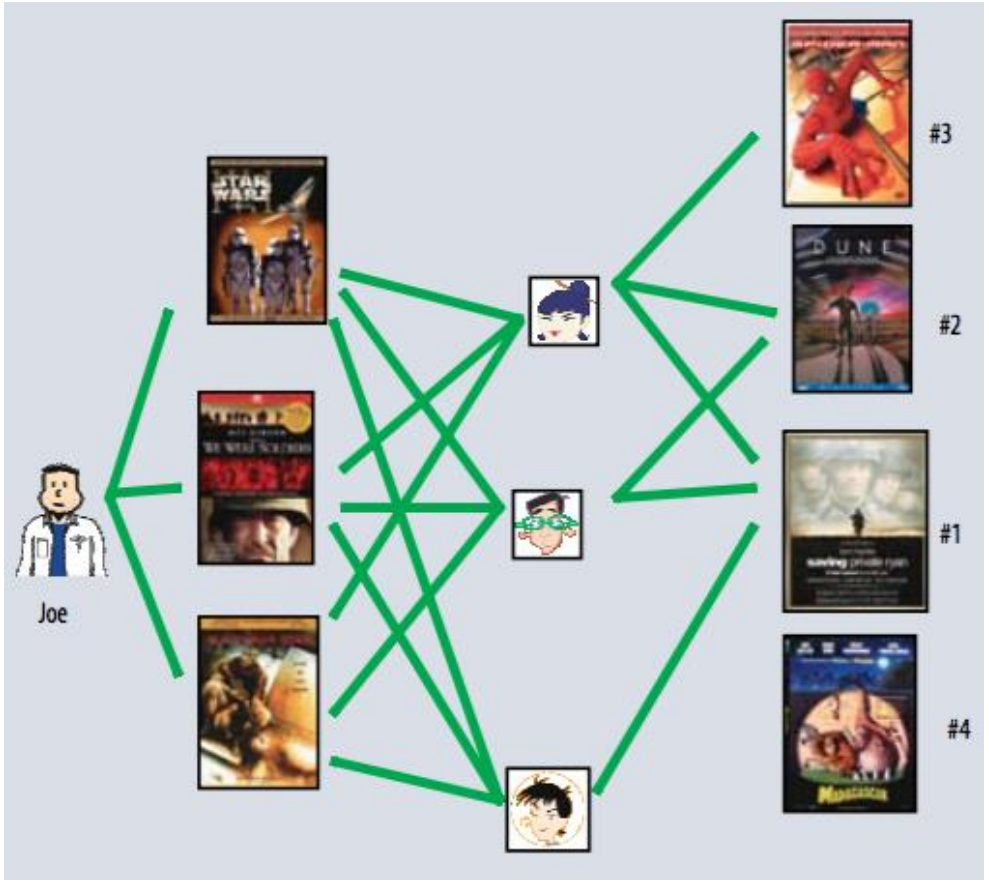

2. Latent Factor Methods

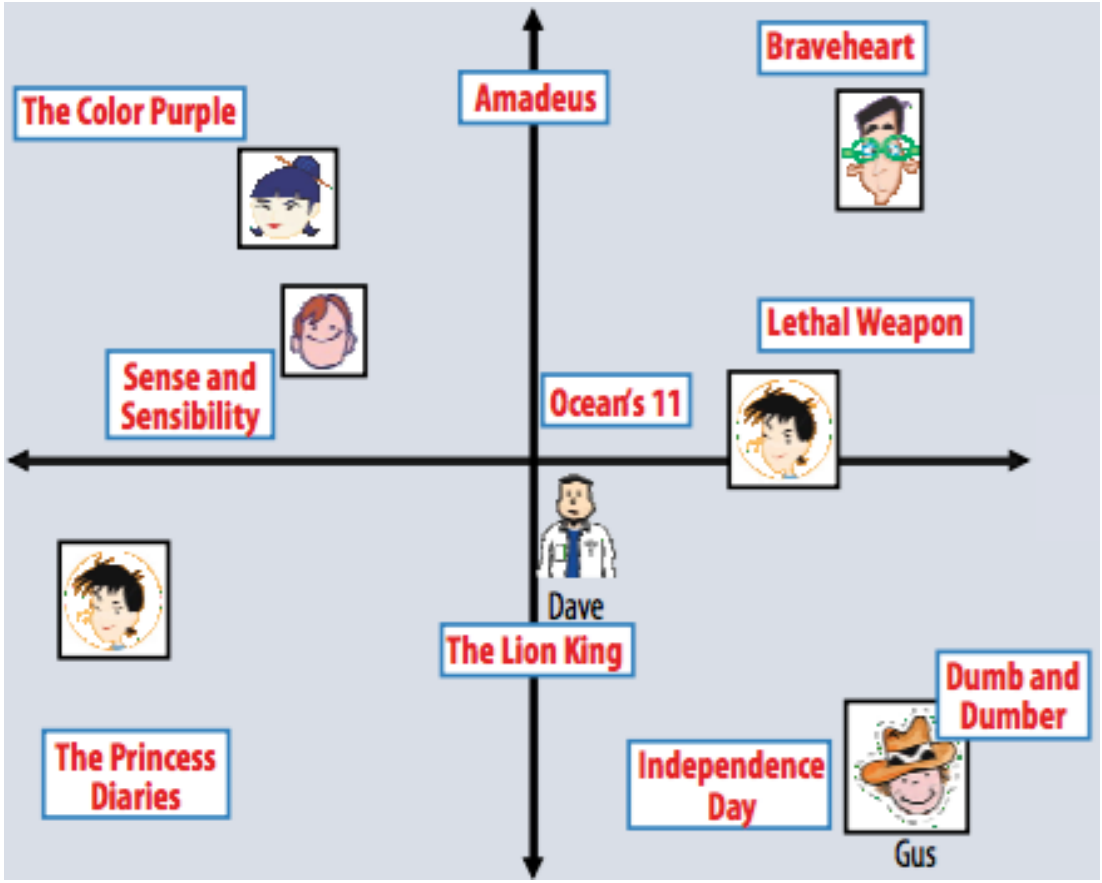

Figures from Koren et al. (2009)

Two Types of Collaborative Filtering

1. Neighborhood Methods

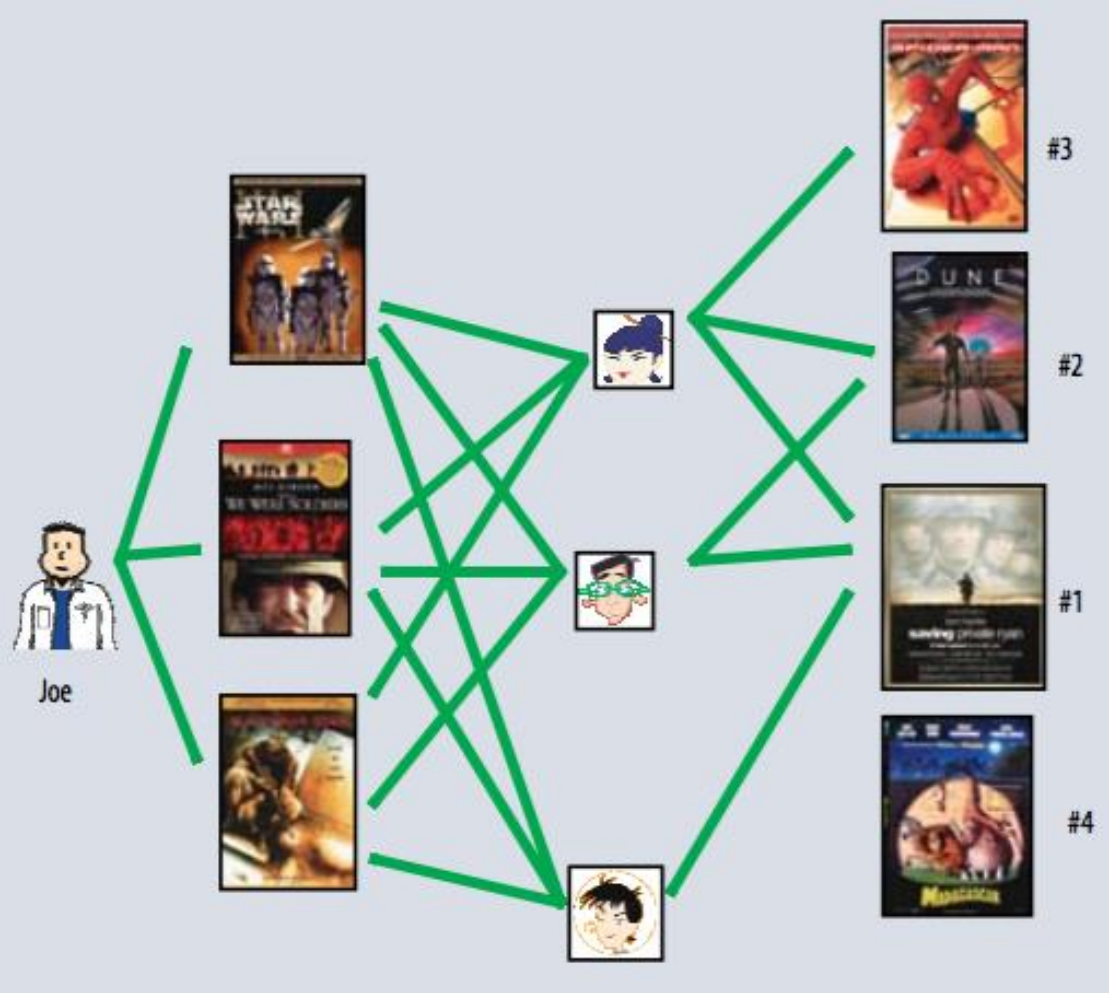

In the figure, assume that a green line indicates the movie was **watched**

Algorithm:

1. **Find neighbors** based on similarity of movie preferences
2. **Recommend** movies that those neighbors watched

Figures from Koren et al. (2009)

Two Types of Collaborative Filtering

- Assume that both movies and users can be mapped to the same **feature space**
- **Recommend** a movie based on its **proximity** to the user in the feature (latent) space
- **Example algorithm**: Matrix Factorization

2. Latent Factor Methods

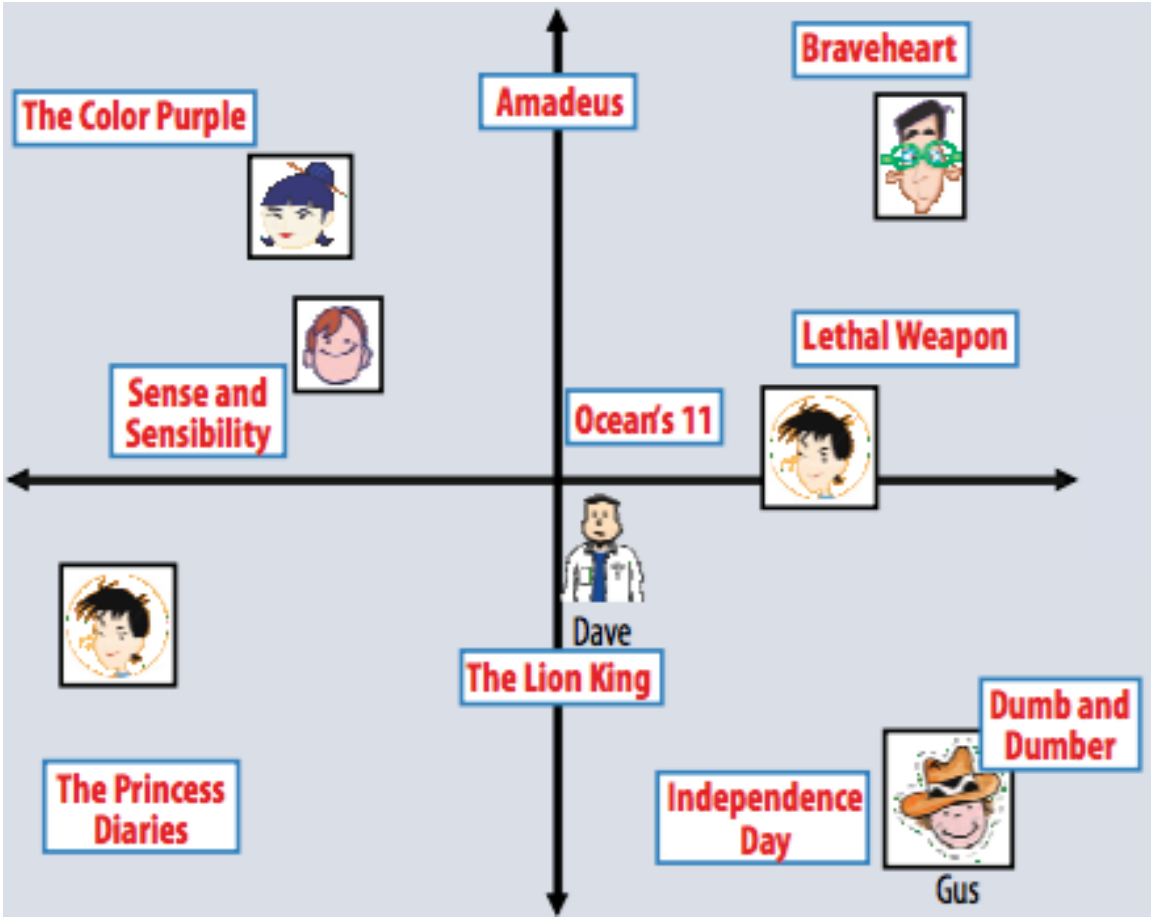

Figures from Koren et al. (2009)

Background: Learn a Feature Space

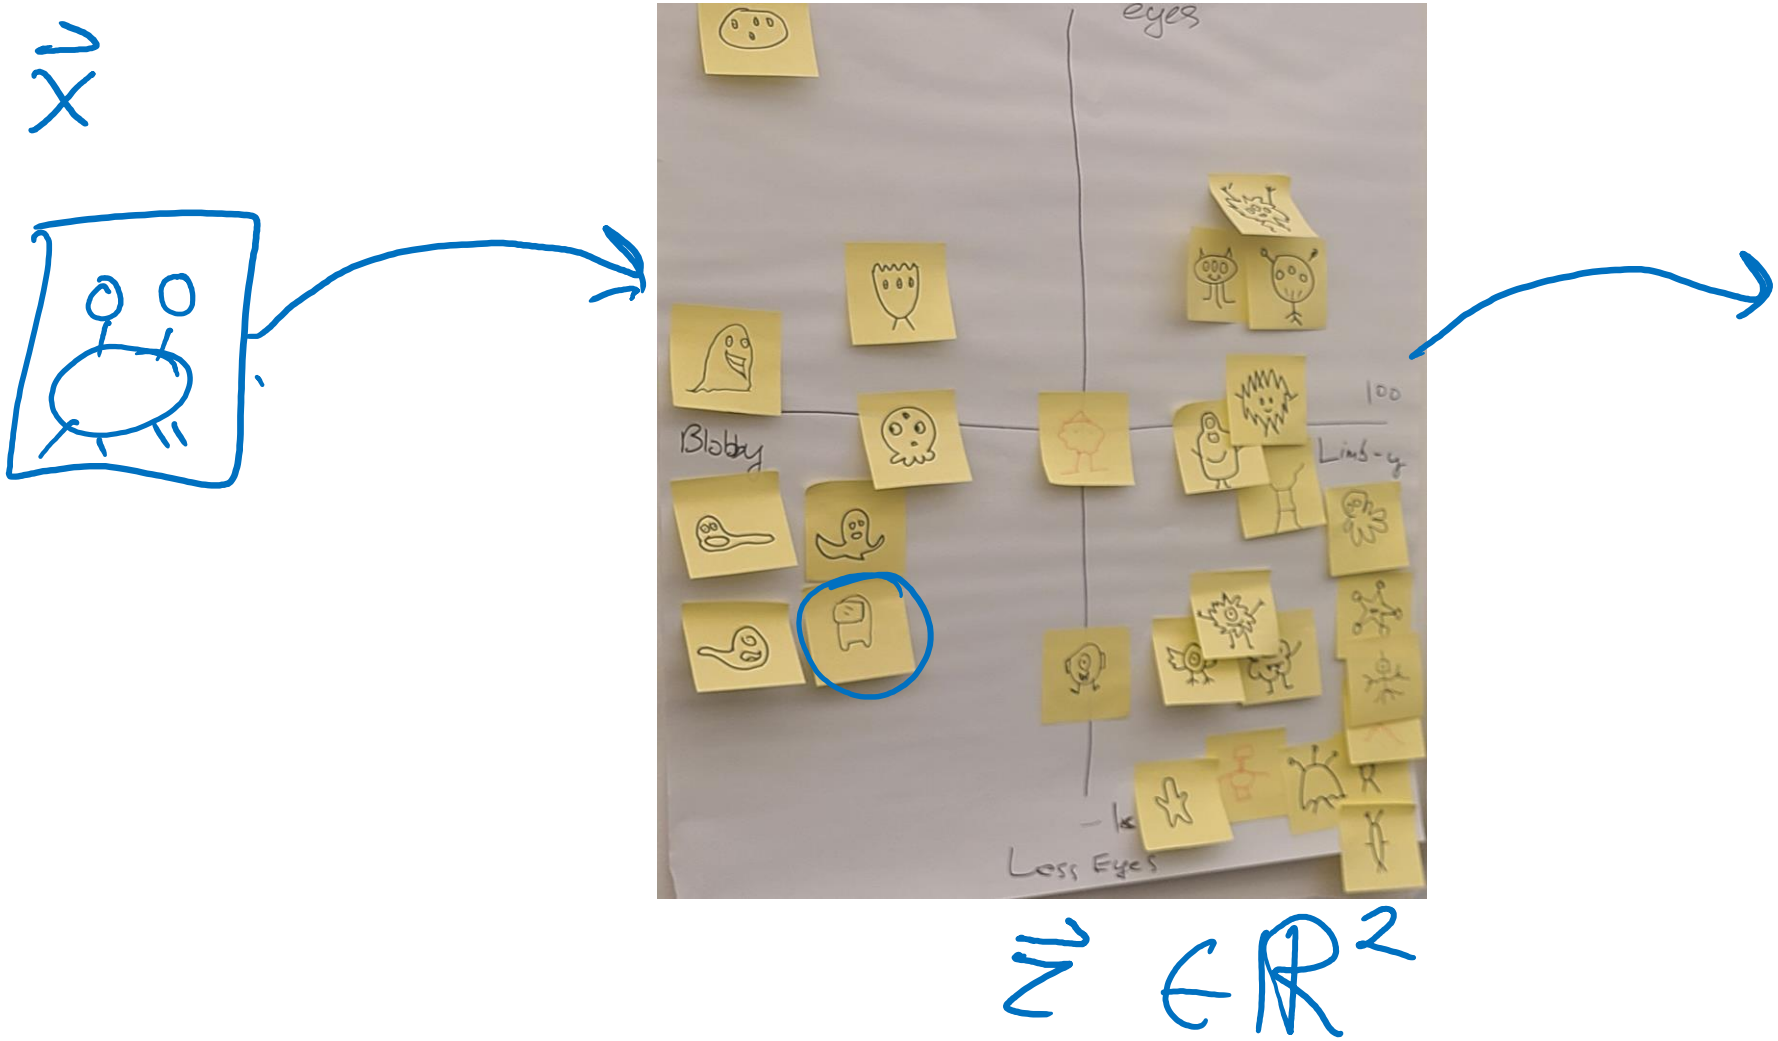

Background: Learn a Feature Space

Word2vec:

Feature space for words

CLIP:

Common feature space for text and images

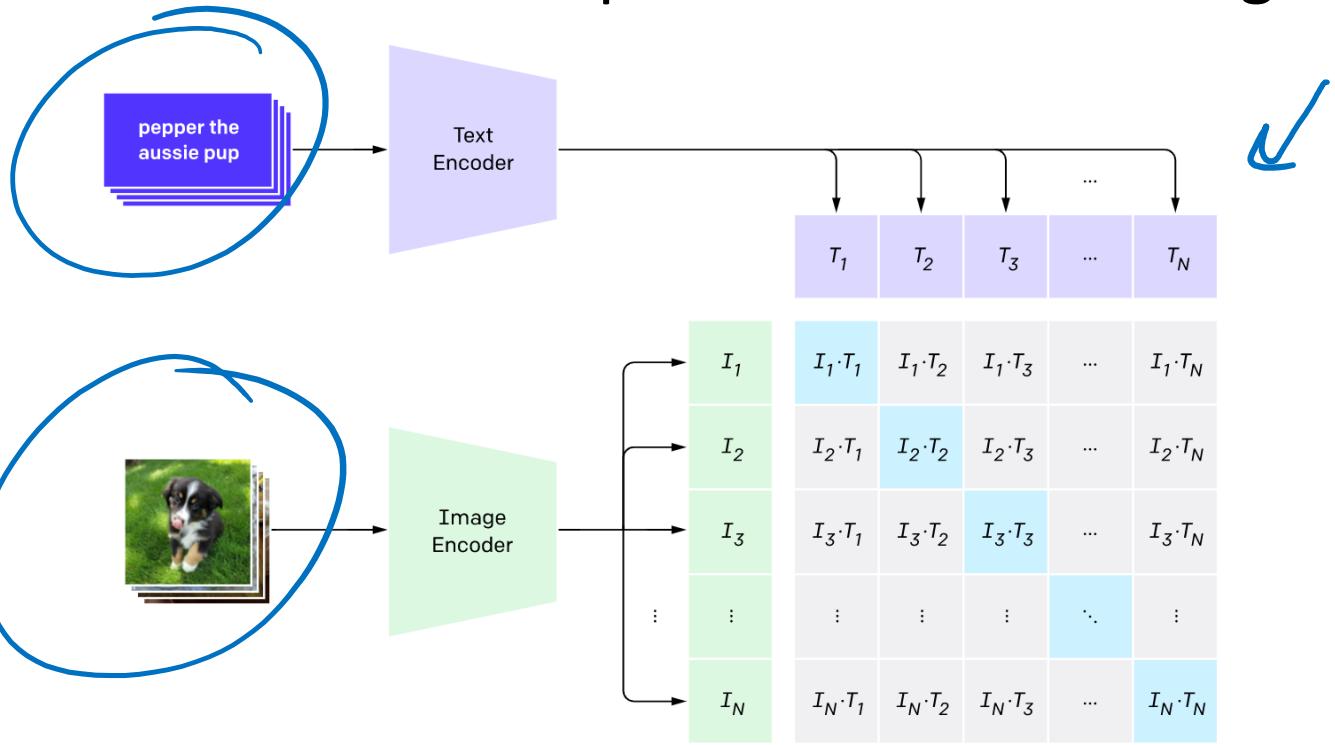

Background: Learn a Feature Space

Why might low dimensional embeddings be useful?

- Example: MNIST digit classification with nearest neighbor

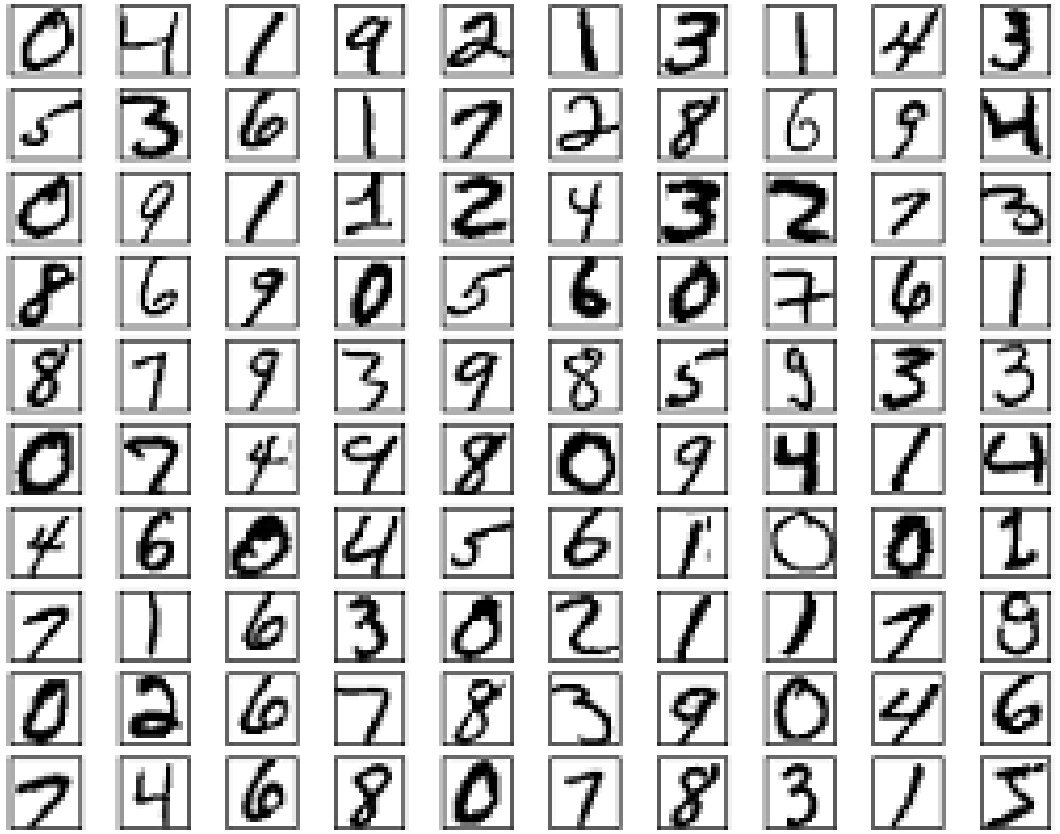

Background: Learn a Feature Space

Why might low dimensional embeddings be useful?

- Example: MNIST digit classification with nearest neighbor

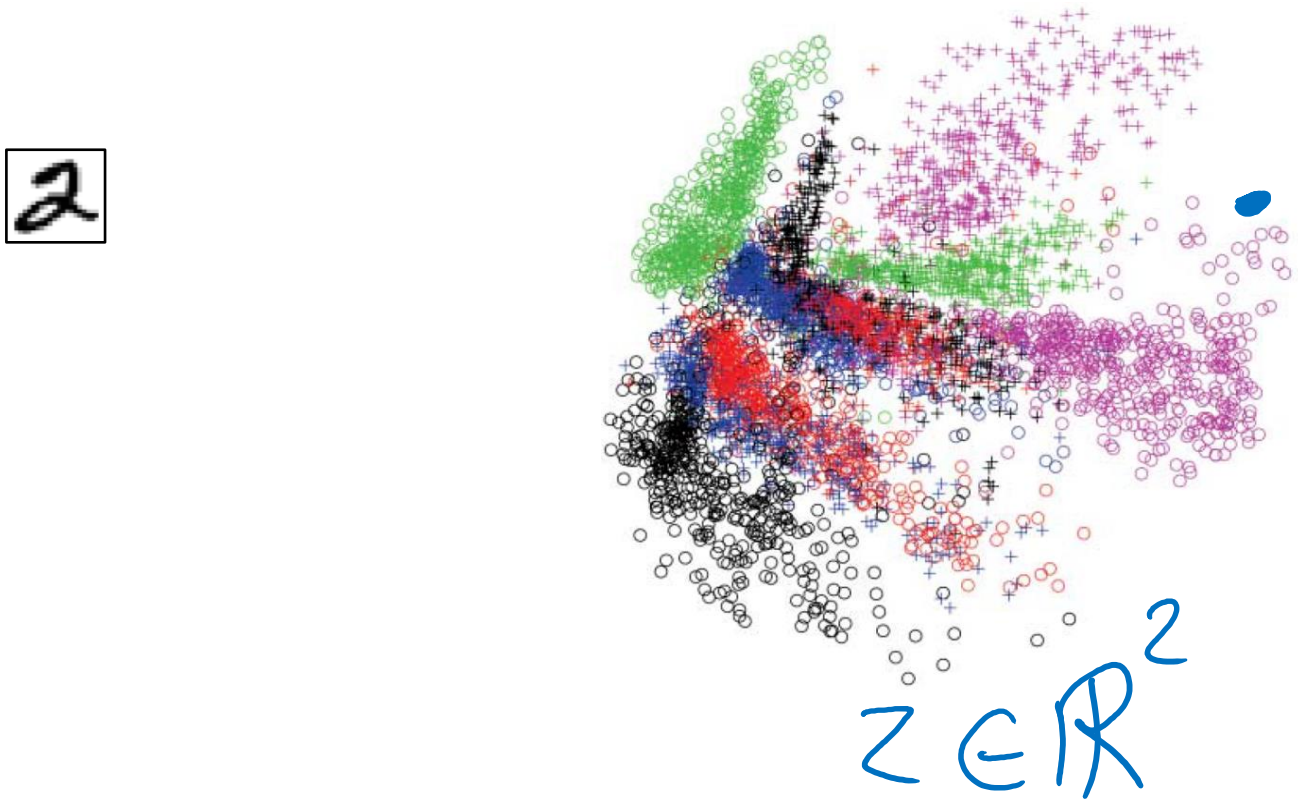

Image: Hinton & Salakhutdinov. *Science* 313.5786 (2006): 504-507.

Background: Measure of Similarity

We've been using Euclidean distance

- $d(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\|_2$

Cosine similarity

- To vectors are similar if the angle between them is small
- $d(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T\mathbf{v}$

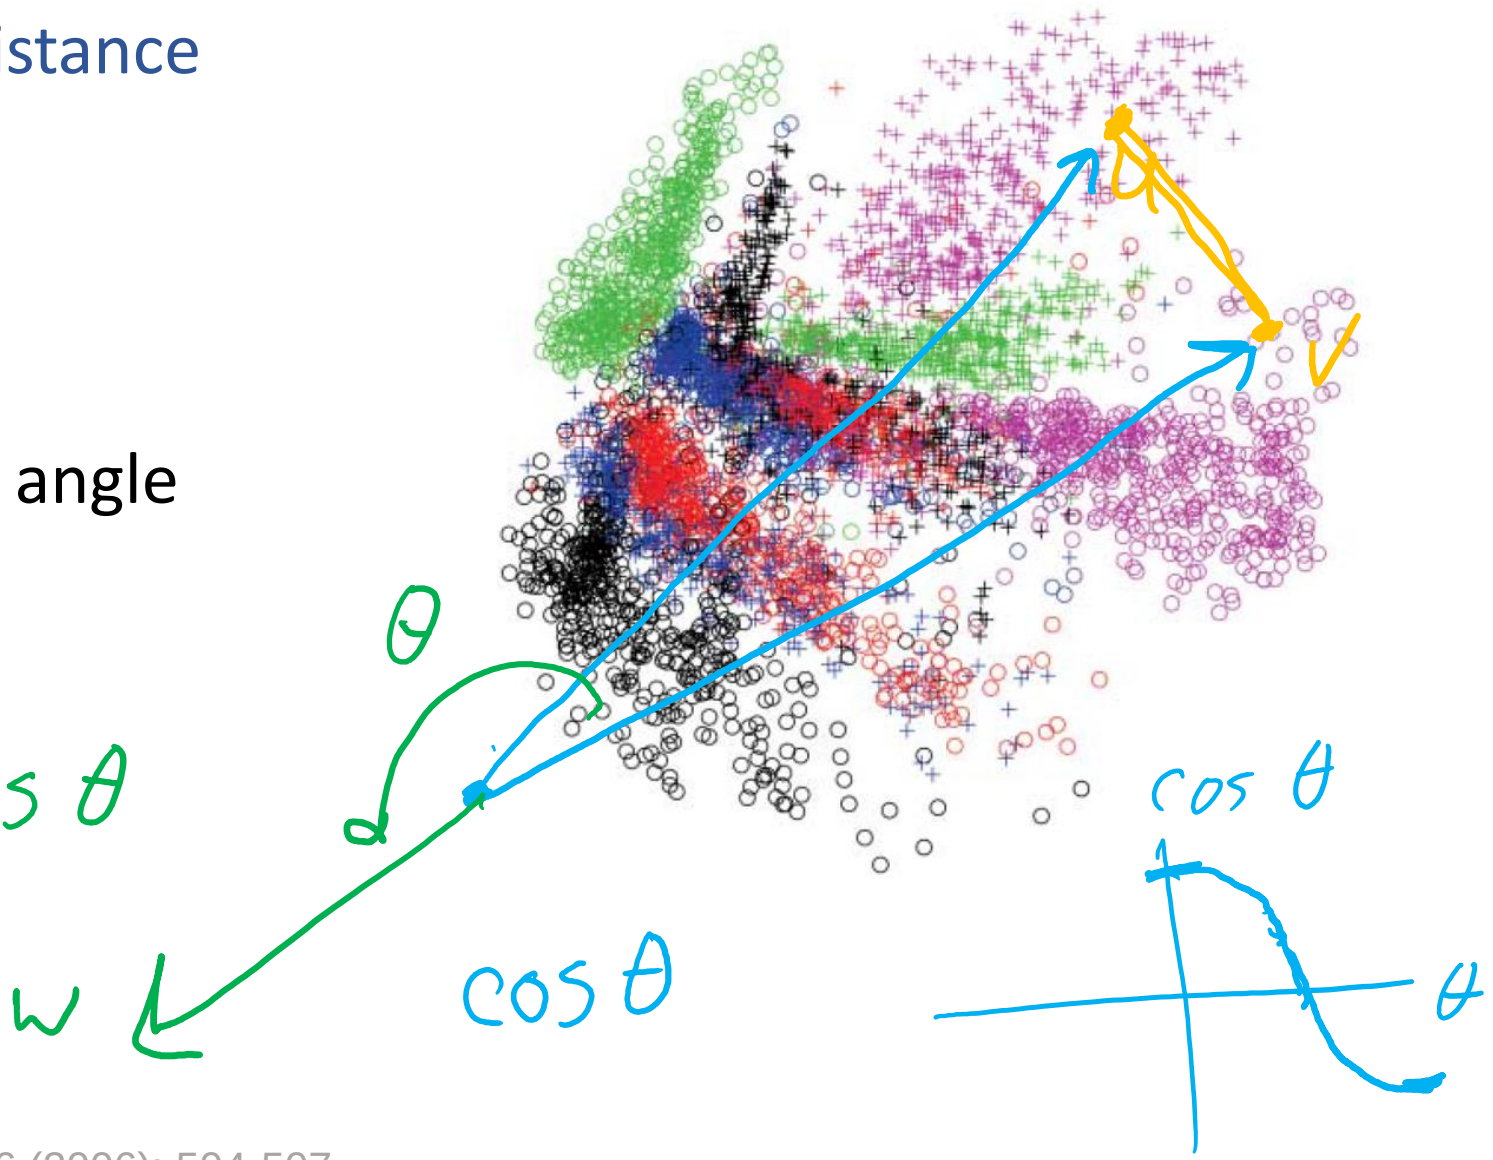

Image: Hinton & Salakhutdinov. *Science* 313.5786 (2006): 504-507.

Background: Learn a Feature Space

CLIP: Common feature space for text and images

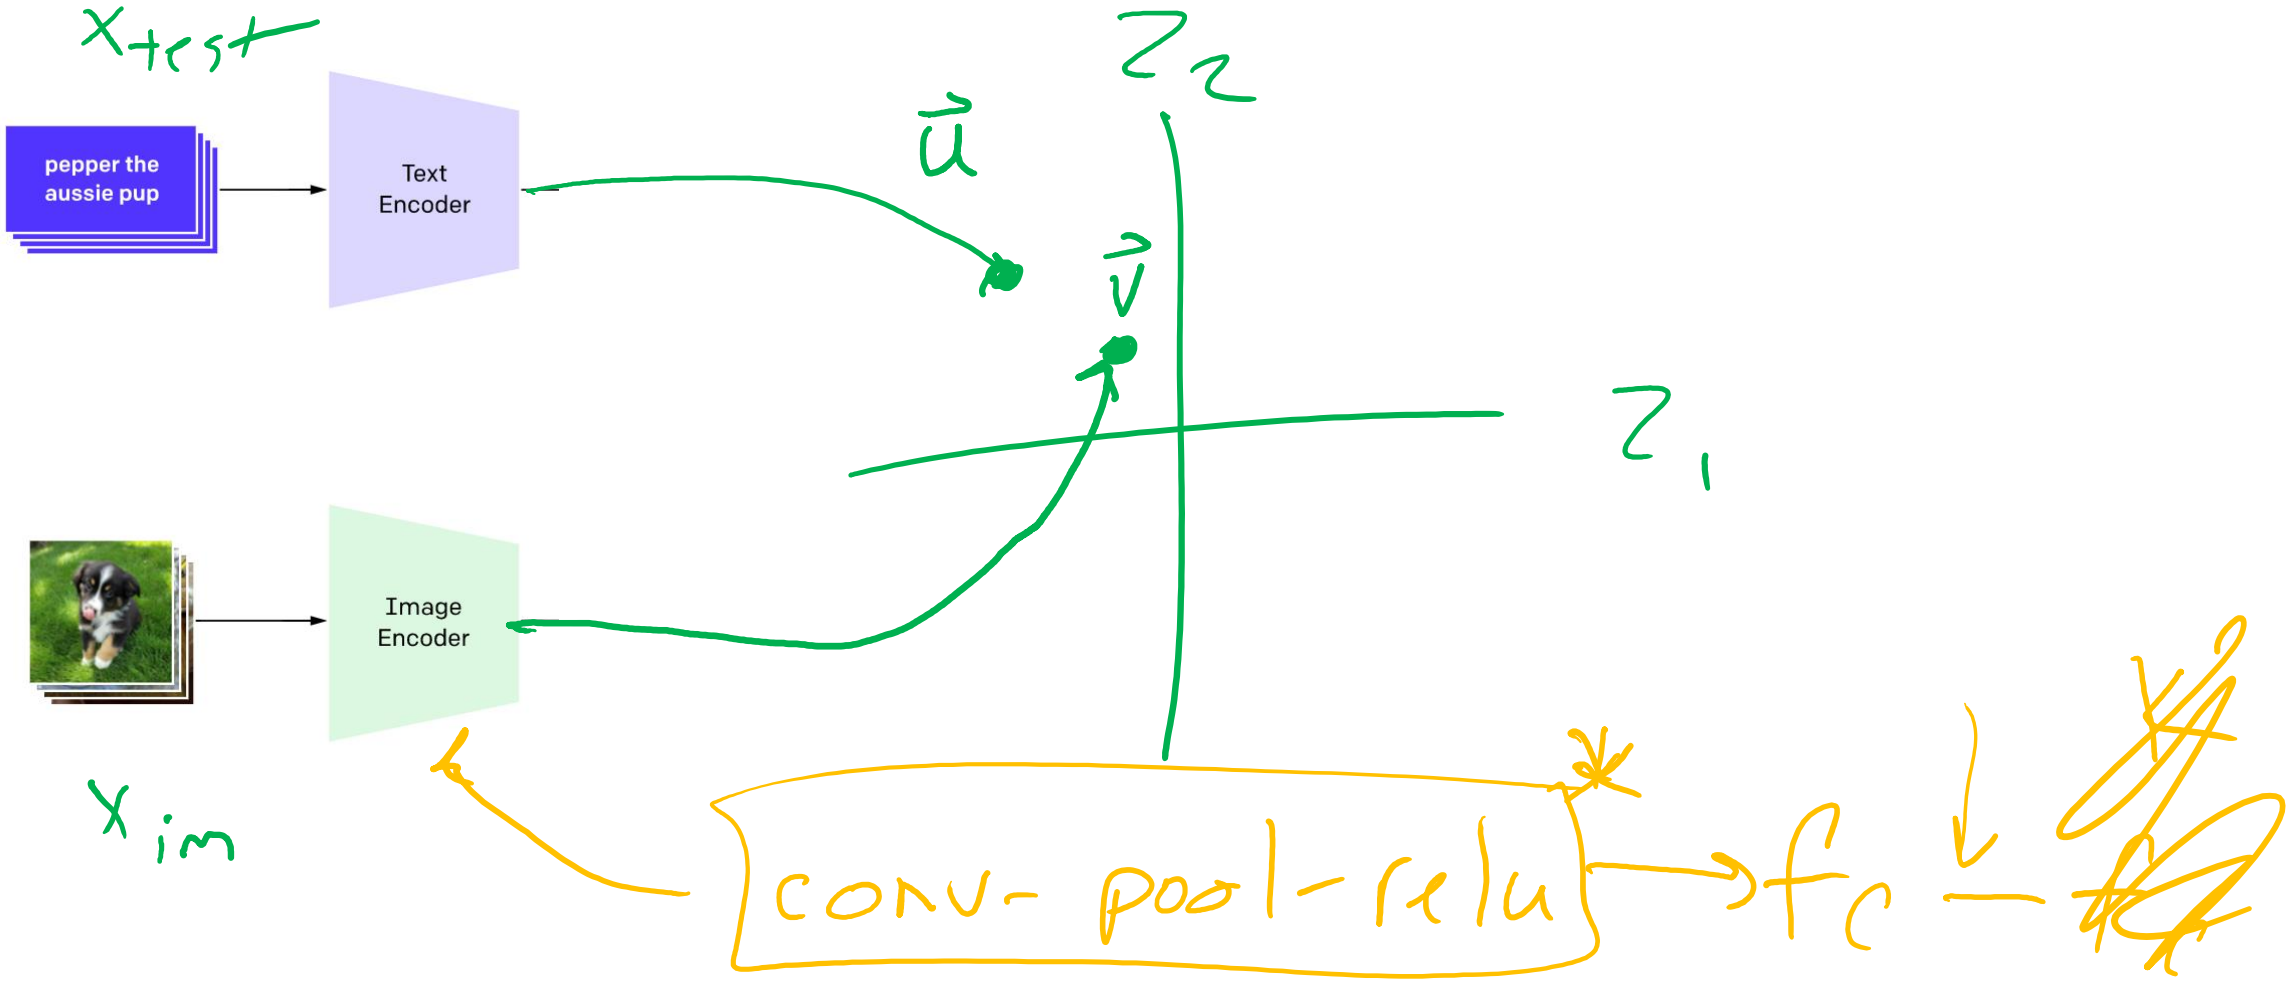

Recommender System: Matrix Factorization

Learning to map items and users to the same feature space

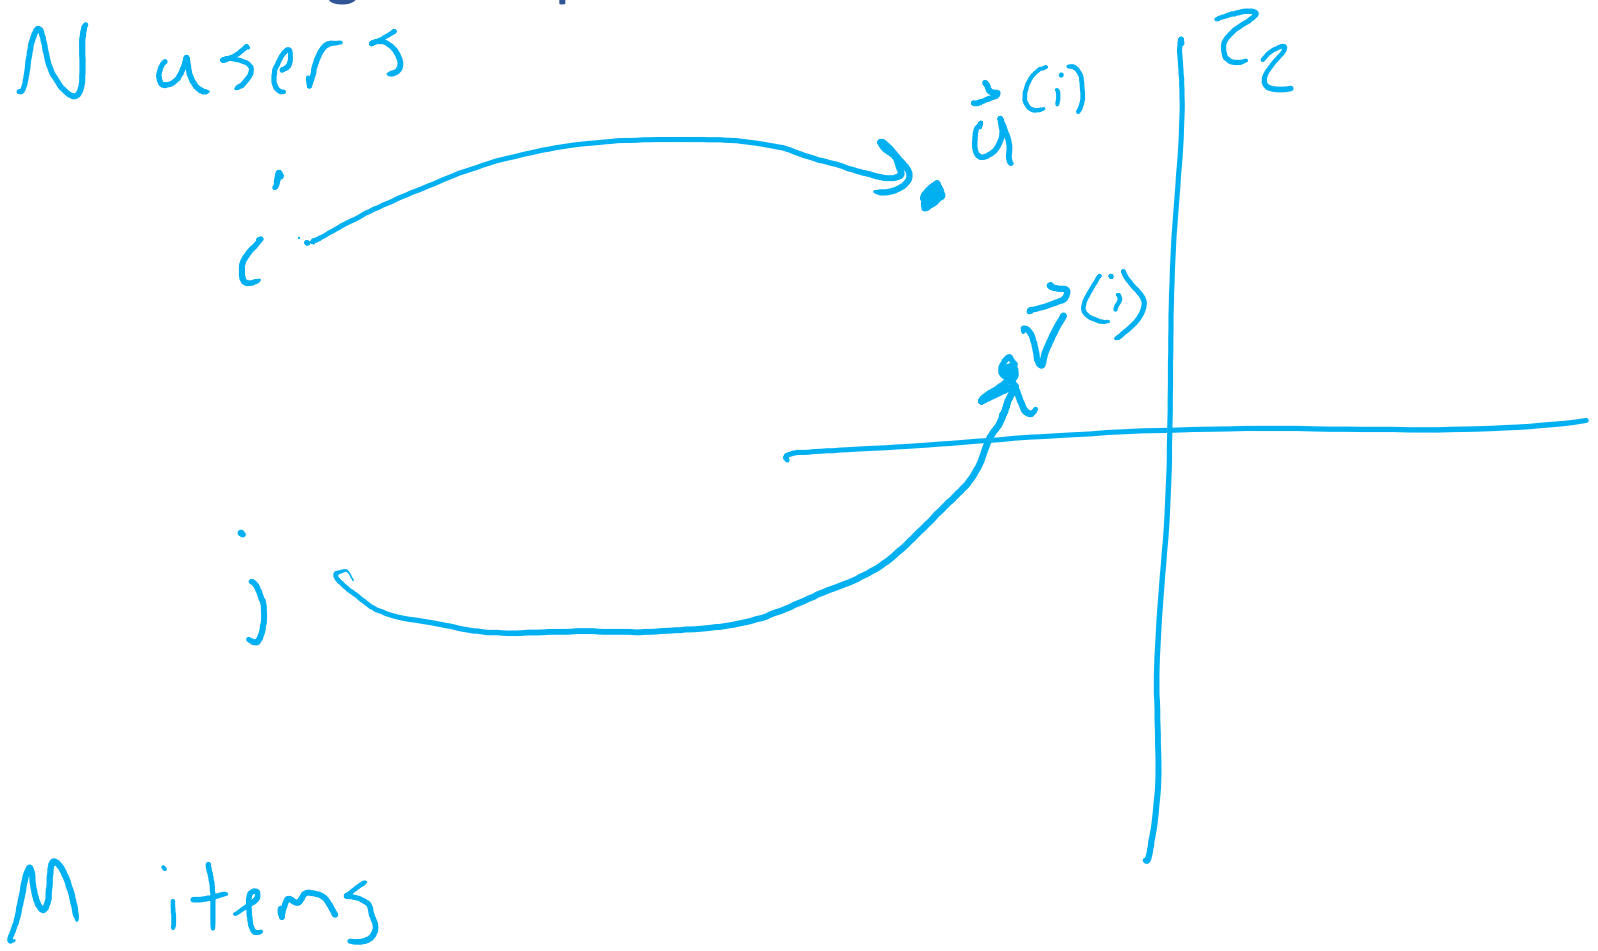

$$\vec{u}^{(i)T}\vec{v}^{(j)} = \hat{r}_{ij}$$

$$\min_{(i,j,r)} \sum \left(r_{ij} - \hat{r}_{ij}\right)^2$$

Recommender System: Matrix Factorization

Learning to map items and users to the same feature space

Recommender System: Matrix Factorization

Learning to map items and users to the same lower dimensional space

N user

M item

$U \in \mathbb{R}^{N \times K}$

$V \in \mathbb{R}^{M \times K}$

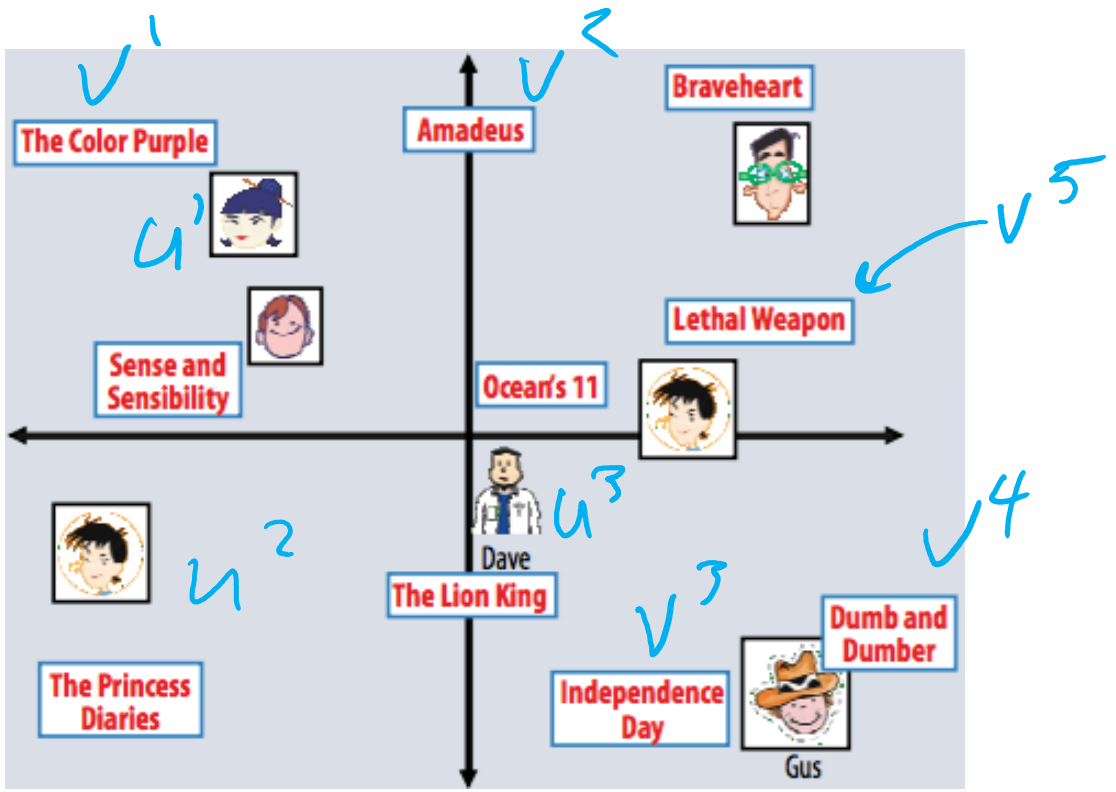

$K = 2$

Figures from Koren et al. (2009)

Recommender System: Matrix Factorization

Optimization: Objective function using only the labels we have

$$\underset{\theta, V}{\operatorname{argmin}} \sum_{(i,j),r_{ij}} \left(r_{ij} - \theta^{(i)T} v^{(j)} \right)^2$$

Poll 1

Is the following optimization a quadratic optimization?

$$\min_{U,V} \sum_{i,j \in S} \left(r_{ij} - \mathbf{u}^{(i)^T} \mathbf{v}^{(j)} \right)^2$$

A. ✗ A ∨

B. Yes

C. No

D. ✗ A ∨

Matrix Factorization

Method of *alternating minimization*

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j \in S} \left(r_{ij} - \mathbf{u}^{(i)T} \mathbf{v}^{(j)} \right)^2$$

do

1. $\min_U J(U,V)$
2. $\min_V J(U,V)$

$$U^{(t+1)} \leftarrow U^{(t)} - \alpha \nabla_U J(U^t, V^t)$$

$$V^{(t+1)} \leftarrow V^{(t)} - \alpha \nabla_V J(U^{t+1}, V^t)$$

Matrix Factorization

Method of *alternating minimization*

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j\in\mathcal{S}} \left(r_{ij} - \mathbf{u}_i^T \mathbf{v}_j\right)^2$$

1) $\operatorname{argmin}_{U} J(U,V)$
2) $\operatorname{argmin}_{V} J(U,V)$

Matrix Factorization

Method of *alternating minimization*

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j \in S} \left(r_{ij} - \mathbf{u}_i^T \mathbf{v}_j\right)^2$$

1) $\operatorname{argmin}_{U} J(U,V)$

2) $\operatorname{argmin}_{V} J(U,V)$

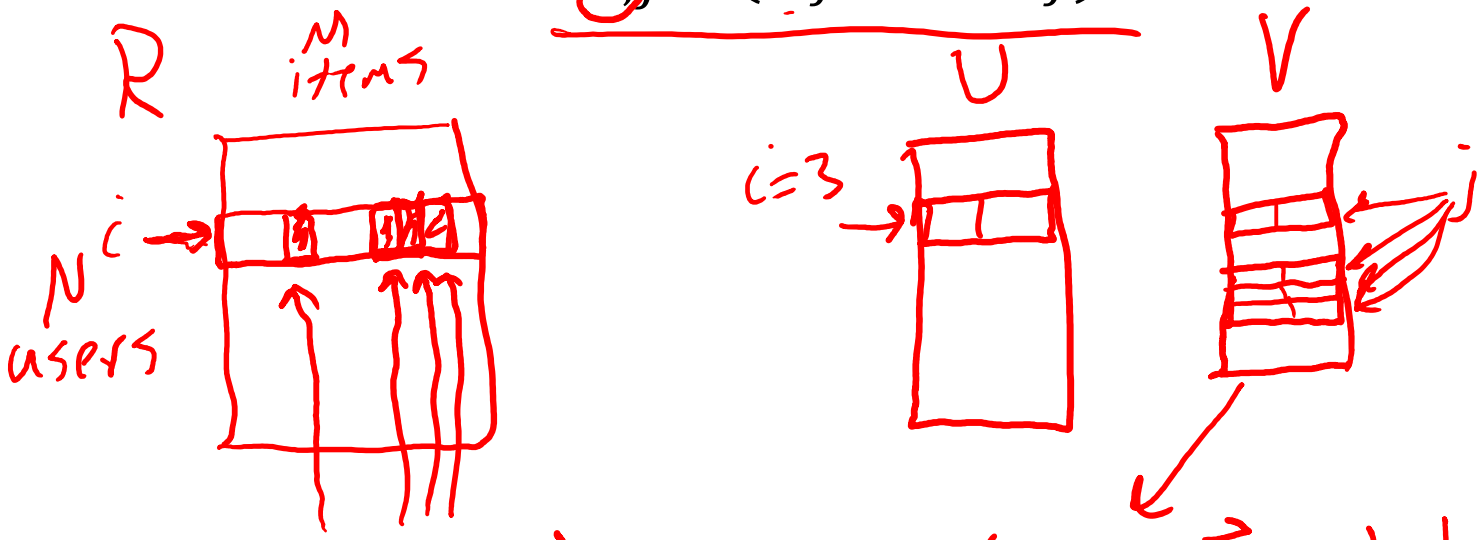

1) For all i, $\operatorname{argmin}_{\mathbf{u}_i} J(U,V)$

2) For all j, $\operatorname{argmin}_{\mathbf{v}_j} J(U,V)$

$$\left\| \vec{r}_{sub} - V_{sub} @ \vec{u}_i \right\|_2^2$$

$$y - X\theta$$

Matrix Factorization

Method of *alternating minimization*

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j\in S}\left(r_{ij} - \mathbf{u}_i^T\mathbf{v}_j\right)^2$$

$$\left\| \vec{r}_{sub} - V_{sub}\,\vec{u}_i \right\|_2^2$$

b A

$$(x^T x)^{-1} x^T y \rightarrow \hat{\theta}$$

1) For all i, $\underset{\mathbf{u}_i}{\operatorname{argmin}} J(U,V)$ $\qquad \mathbf{u}_i^{(t+1)} = (A^T A)^{-1} A^T \boldsymbol{b}$

2) For all j, $\underset{\mathbf{v}_j}{\operatorname{argmin}} J(U,V)$ $\qquad \mathbf{v}_j^{(t+1)} = (C^T C)^{-1} C^T \mathbf{d}$

Matrix Factorization

Method of *alternating minimization*

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j\in\mathcal{S}} \left(r_{ij} - \mathbf{u}_i^T \mathbf{v}_j\right)^2$$

1) For all i, $\underset{\mathbf{u}_i}{\operatorname{argmin}} J(U,V)$ $\qquad \mathbf{u}_i^{(t+1)} = (A^T A)^{-1} A^T \boldsymbol{b}$

2) For all j, $\underset{\mathbf{v}_j}{\operatorname{argmin}} J(U,V)$ $\qquad \mathbf{v}_j^{(t+1)} = (C^T C)^{-1} C^T \mathbf{d}$

Matrix Factorization

First solve:

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j \in S} \left(r_{ij} - \mathbf{u}^{(i)T} \mathbf{v}^{(j)} \right)^2$$

U V

What then?

predict

$h(i,j) \rightarrow \hat{r}_{ij}$

$u^{(i)T} v^{(j)} \rightarrow \hat{r}_{ij}$

	Doctor Strange	Star Trek: Beyond	Zootopia
Alita	1	6.4	5
BB-8	3	4	2.5
C-3Po	3	5	2

Matrix Factorization

First solve:

$$\min_{U,V} J(U,V) \qquad J(U,V) = \sum_{i,j \in S} \left(r_{ij} - \mathbf{u}^{(i)T} \mathbf{v}^{(j)} \right)^2$$

items (j)

users (i)

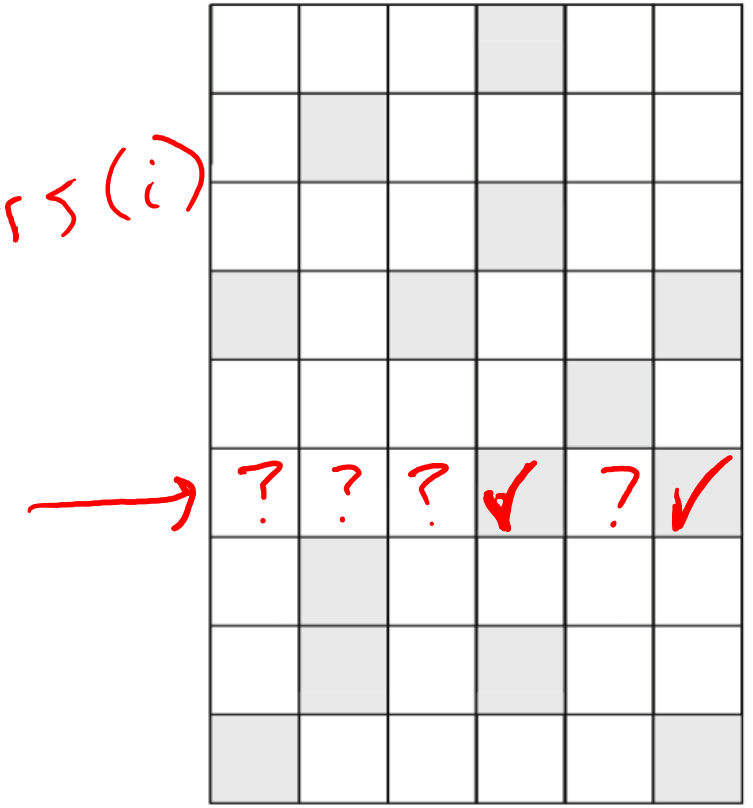

What then?

Recommender System: Matrix Factorization

Why is it called matrix factorization?

$$\min_{U,V} \sum_{i,j \in S} \left(r_{ij} - \mathbf{u}^{(i)T} \mathbf{v}^{(j)} \right)^2$$

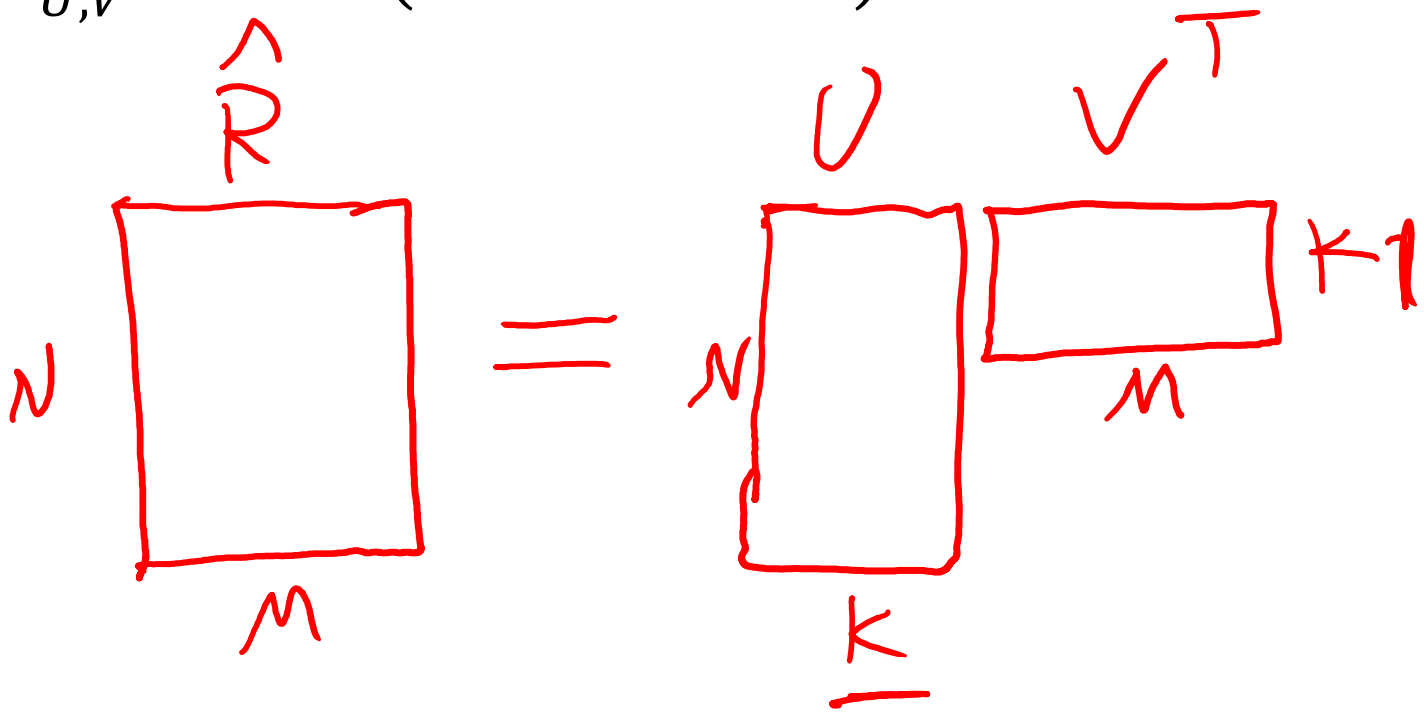

SVD

$X = U S V^T$

$R = U_1 V^T$

matrix factorization

Recommender System: Matrix Factorization

Why is it called matrix factorization?

$$\min_{\boldsymbol{U},\boldsymbol{V}} \|\boldsymbol{R} - \boldsymbol{U}\boldsymbol{V}^T\|_F^2$$

Sparse labels ☹

	Doctor Strange	Star Trek: Beyond	Zootopia
Alita	1		5
BB-8	3	4	
C-3Po	3	5	2

Matrix Factorization

Add regularization to avoid overfitting

$$\min_{U,V} J(U,V) + \lambda_1 \|U\|_F^2 + \lambda_2 \|V\|_F^2$$

$$J(U,V) = \sum_{i,j \in \mathcal{S}} \left(r_{ij} - \mathbf{u}^{(i)^T} \mathbf{v}^{(j)} \right)^2$$

Bias in Recommender Systems

What high-level problems can occur from recommender systems?

Summary

Recommender systems solve many **real-world** (*large-scale) **problems**

Collaborative filtering by matrix factorization (MF) is an **efficient** and **effective** approach

(Sparse matrix makes MF more challenging)

MF is just another example of a **common recipe**:

1. define a model
2. define an objective function
3. optimize

Optimization

- Use alternating minimization
- Add regularization to avoid overfitting