551: Guest Lecture Low-rank Matrix Completion

Greg Ongie, Postdoctoral Fellow
EECS Department
University of Michigan
11/7/2017

Netflix Problem

Movies

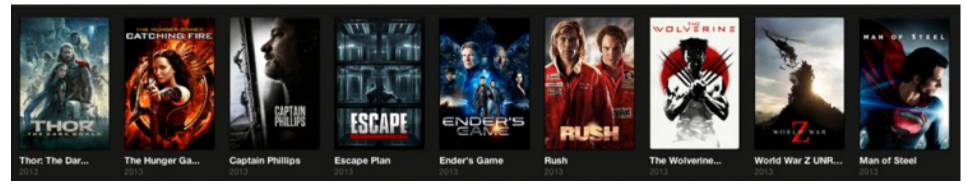

Users

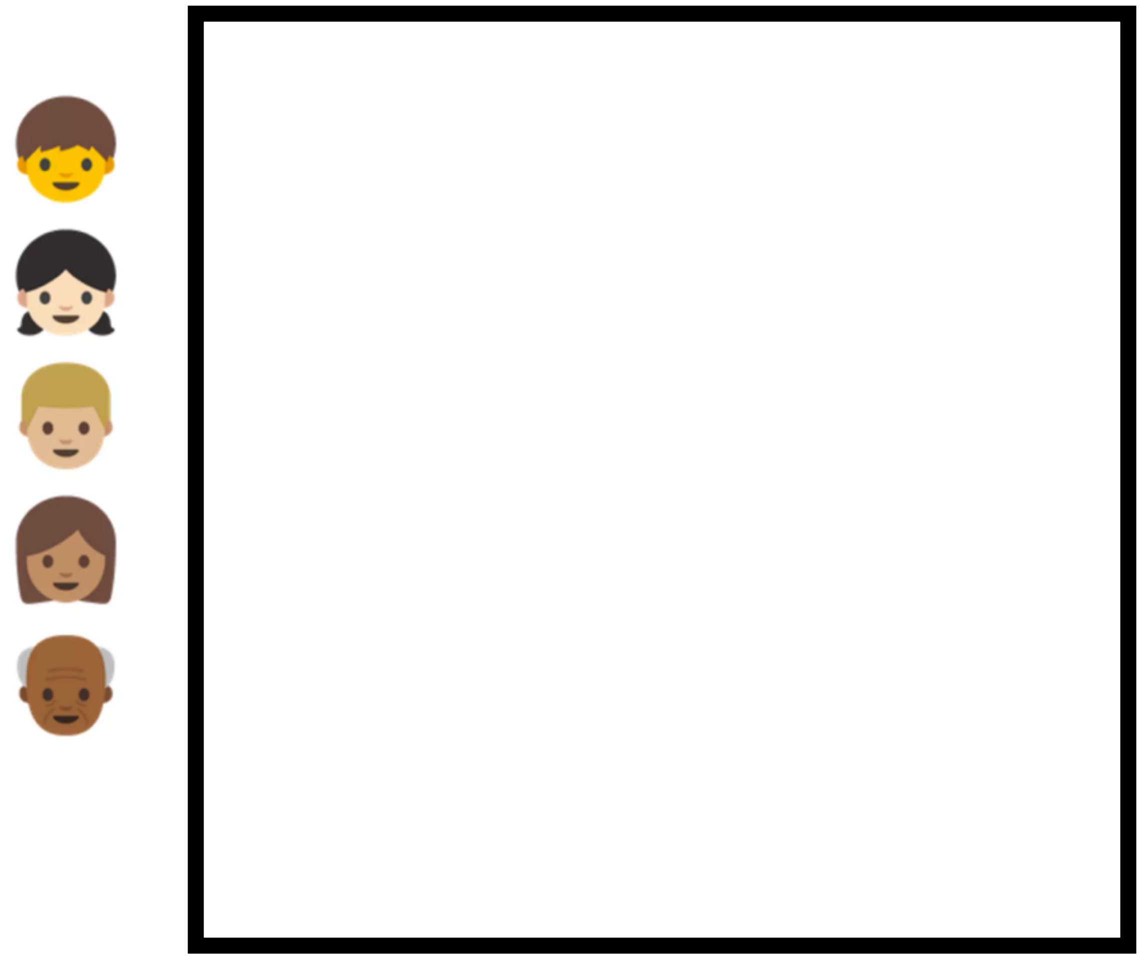

Netflix Problem

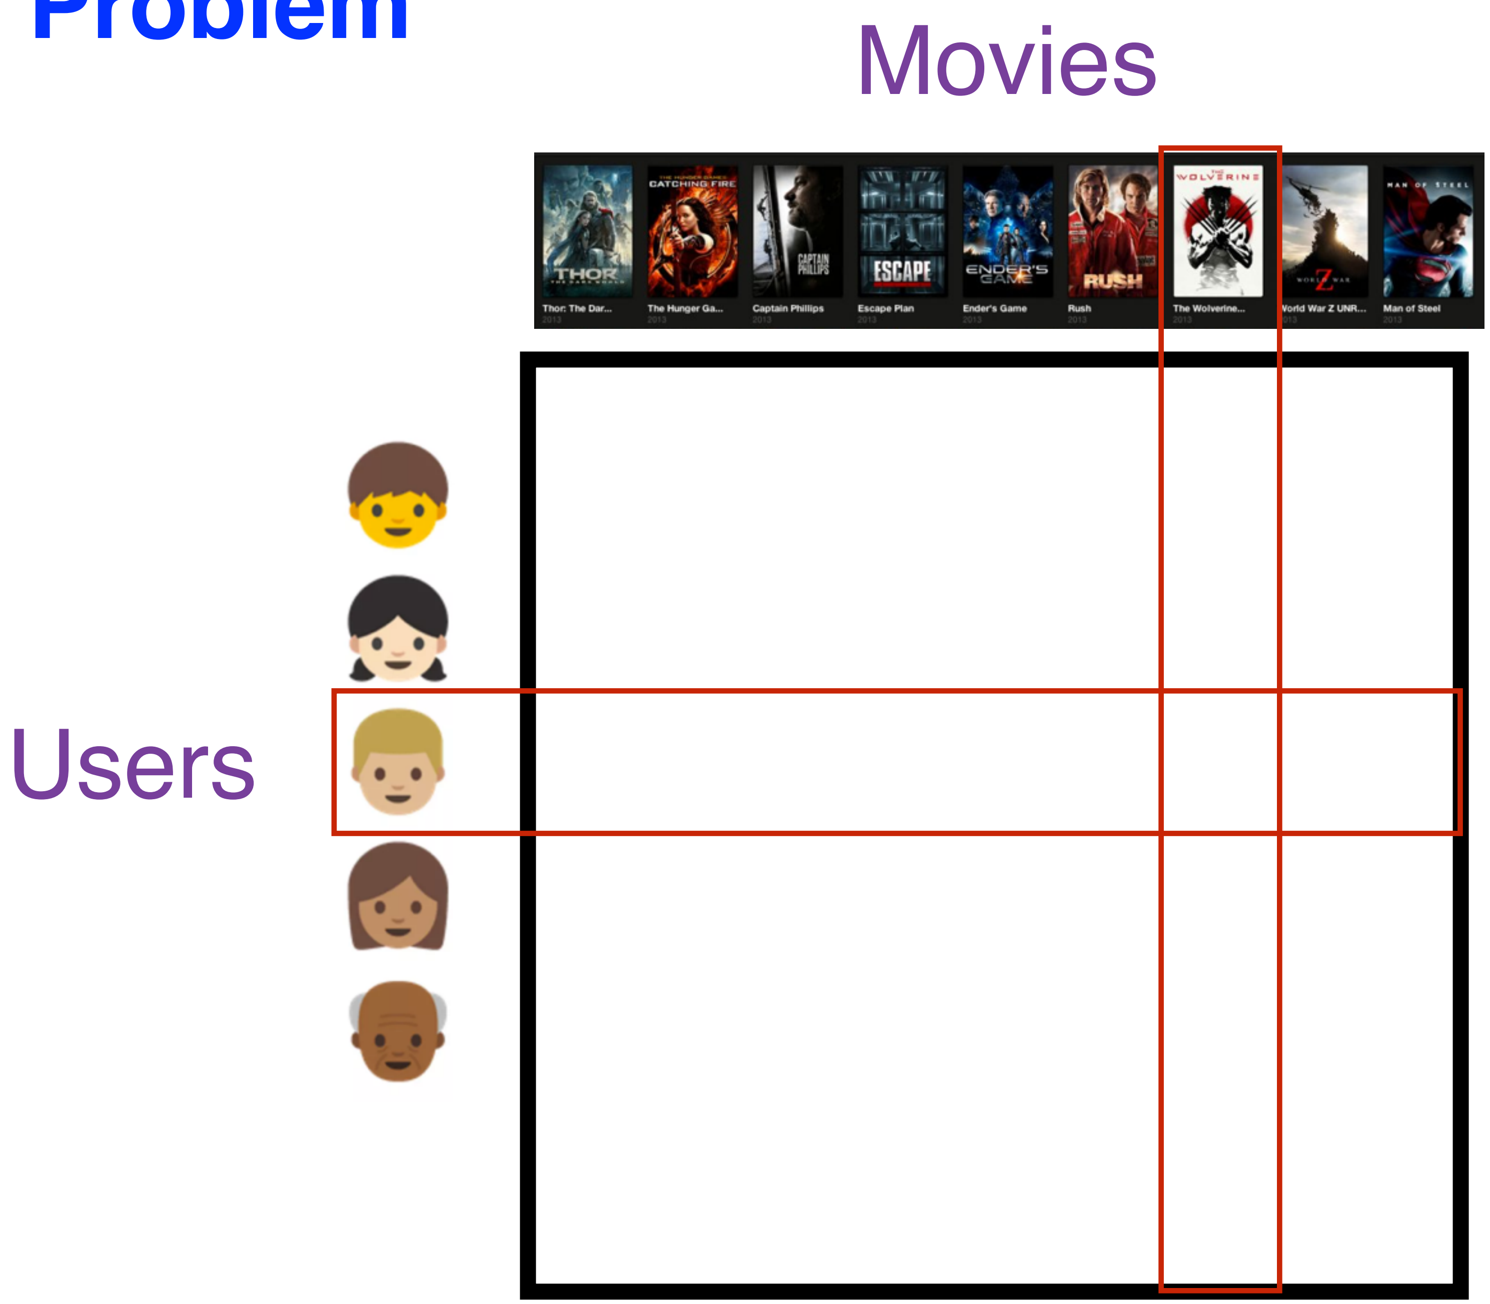

Netflix Problem

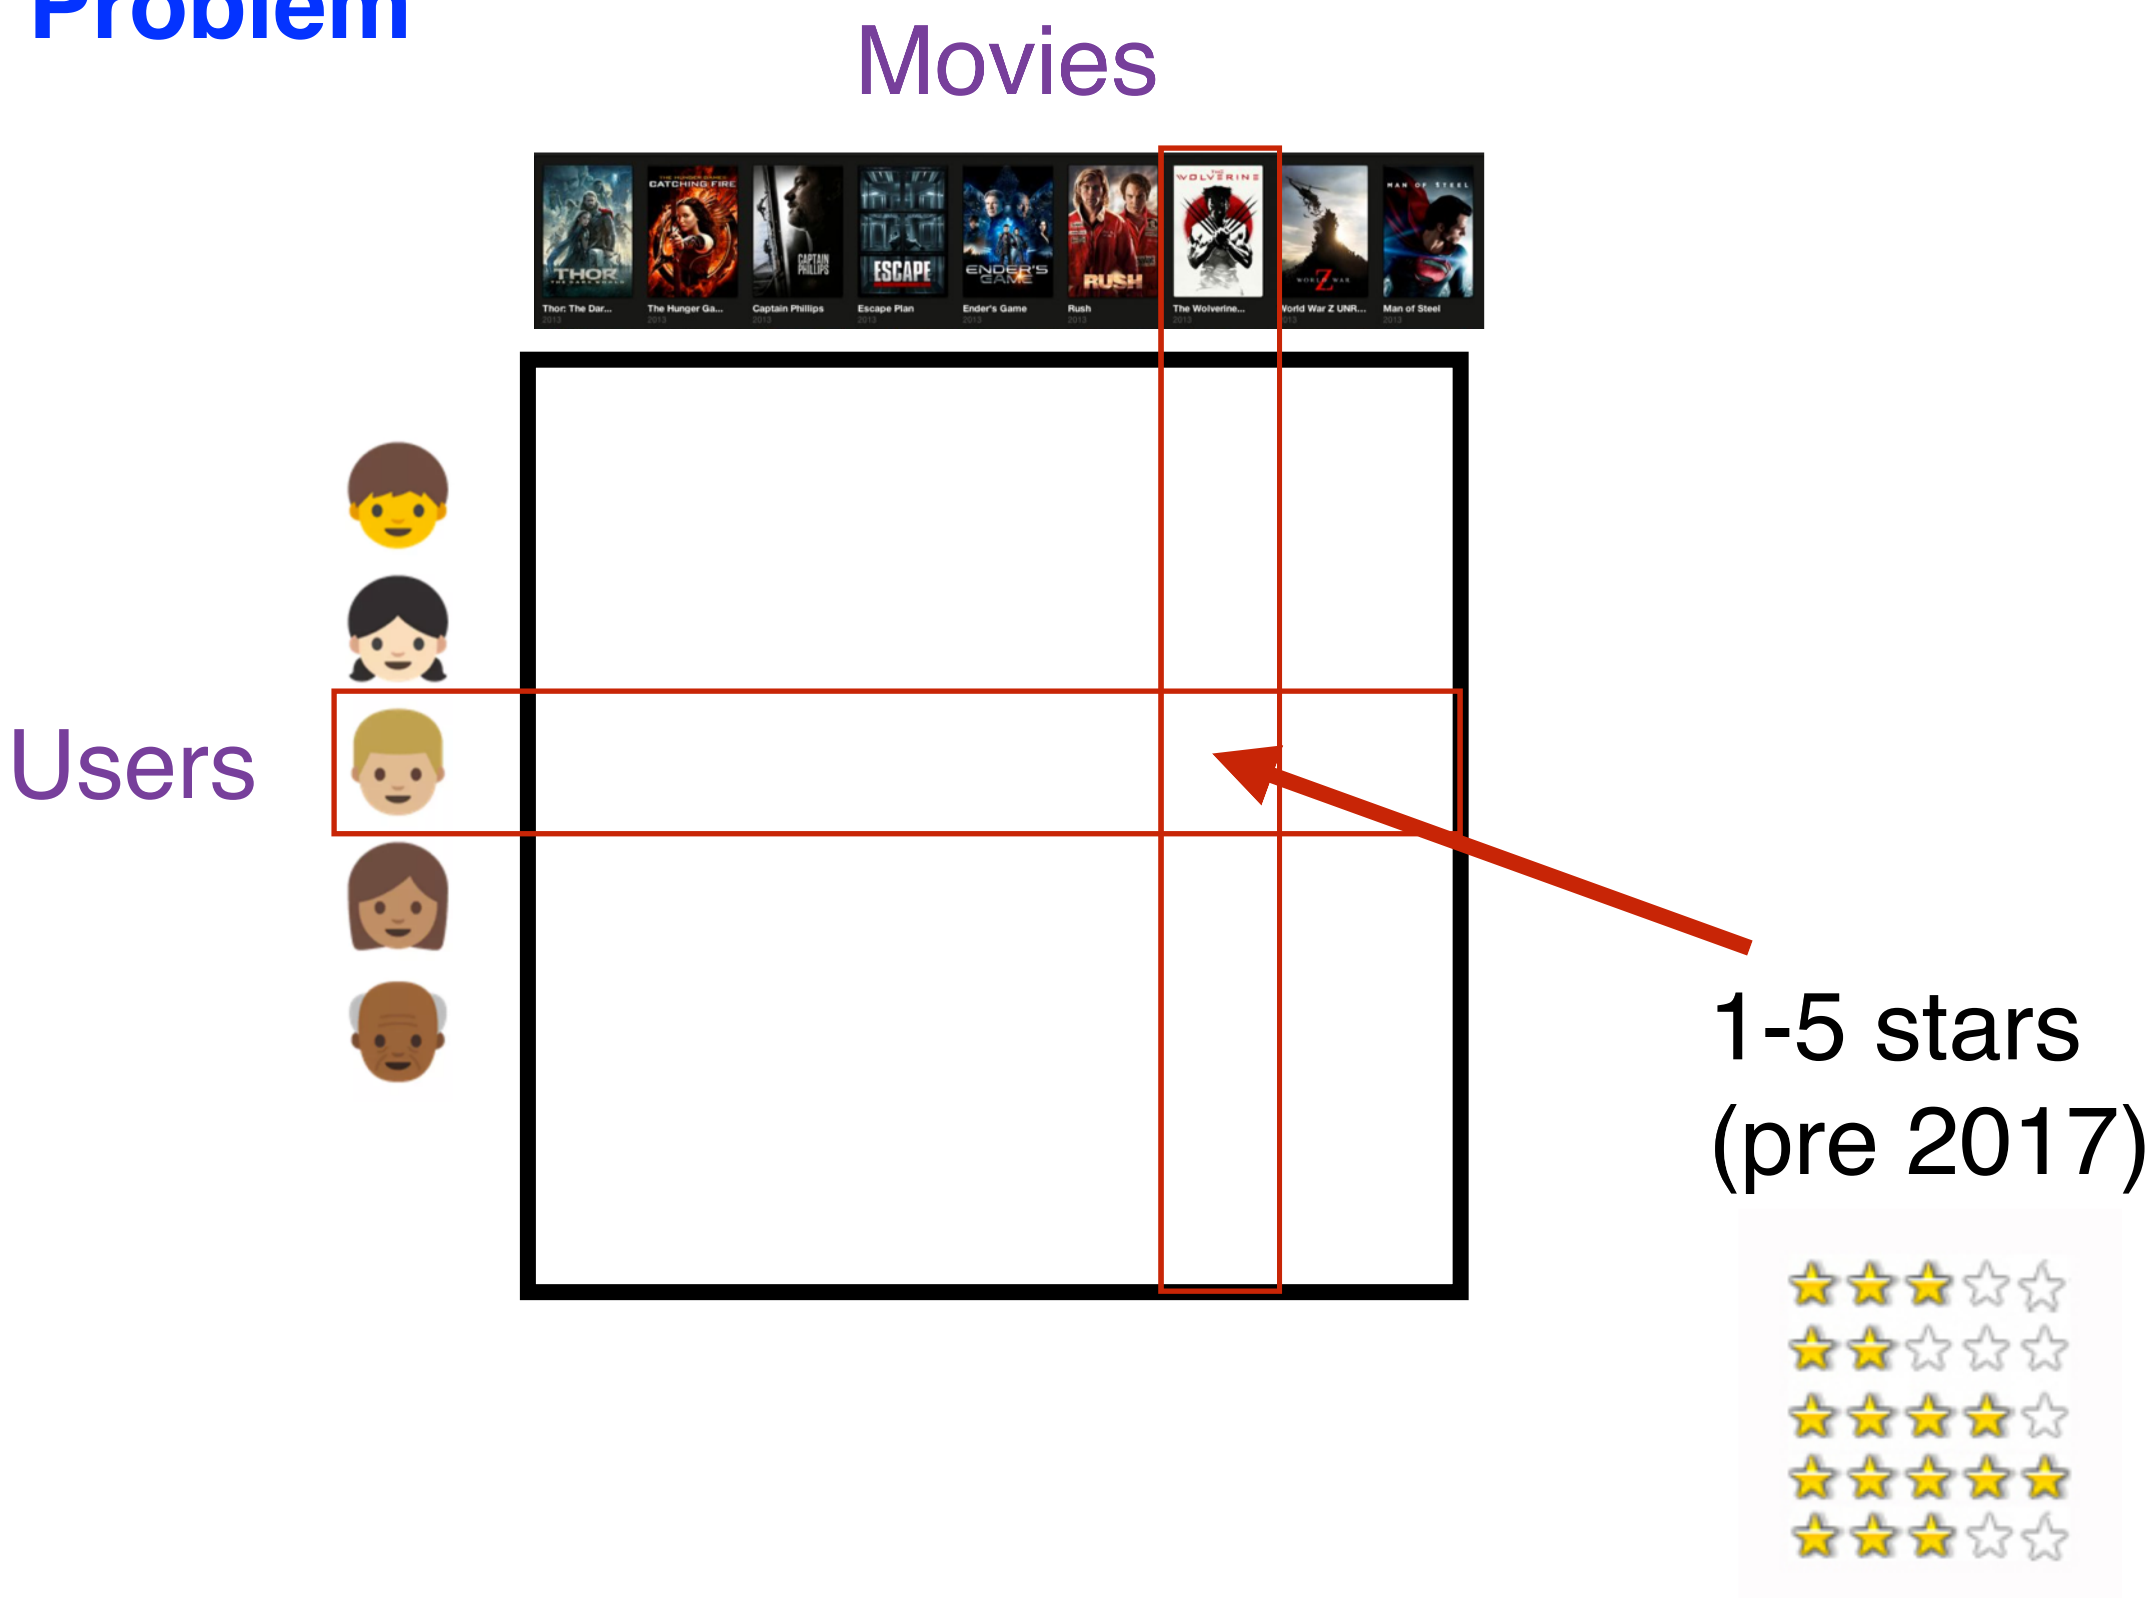

Netflix Problem

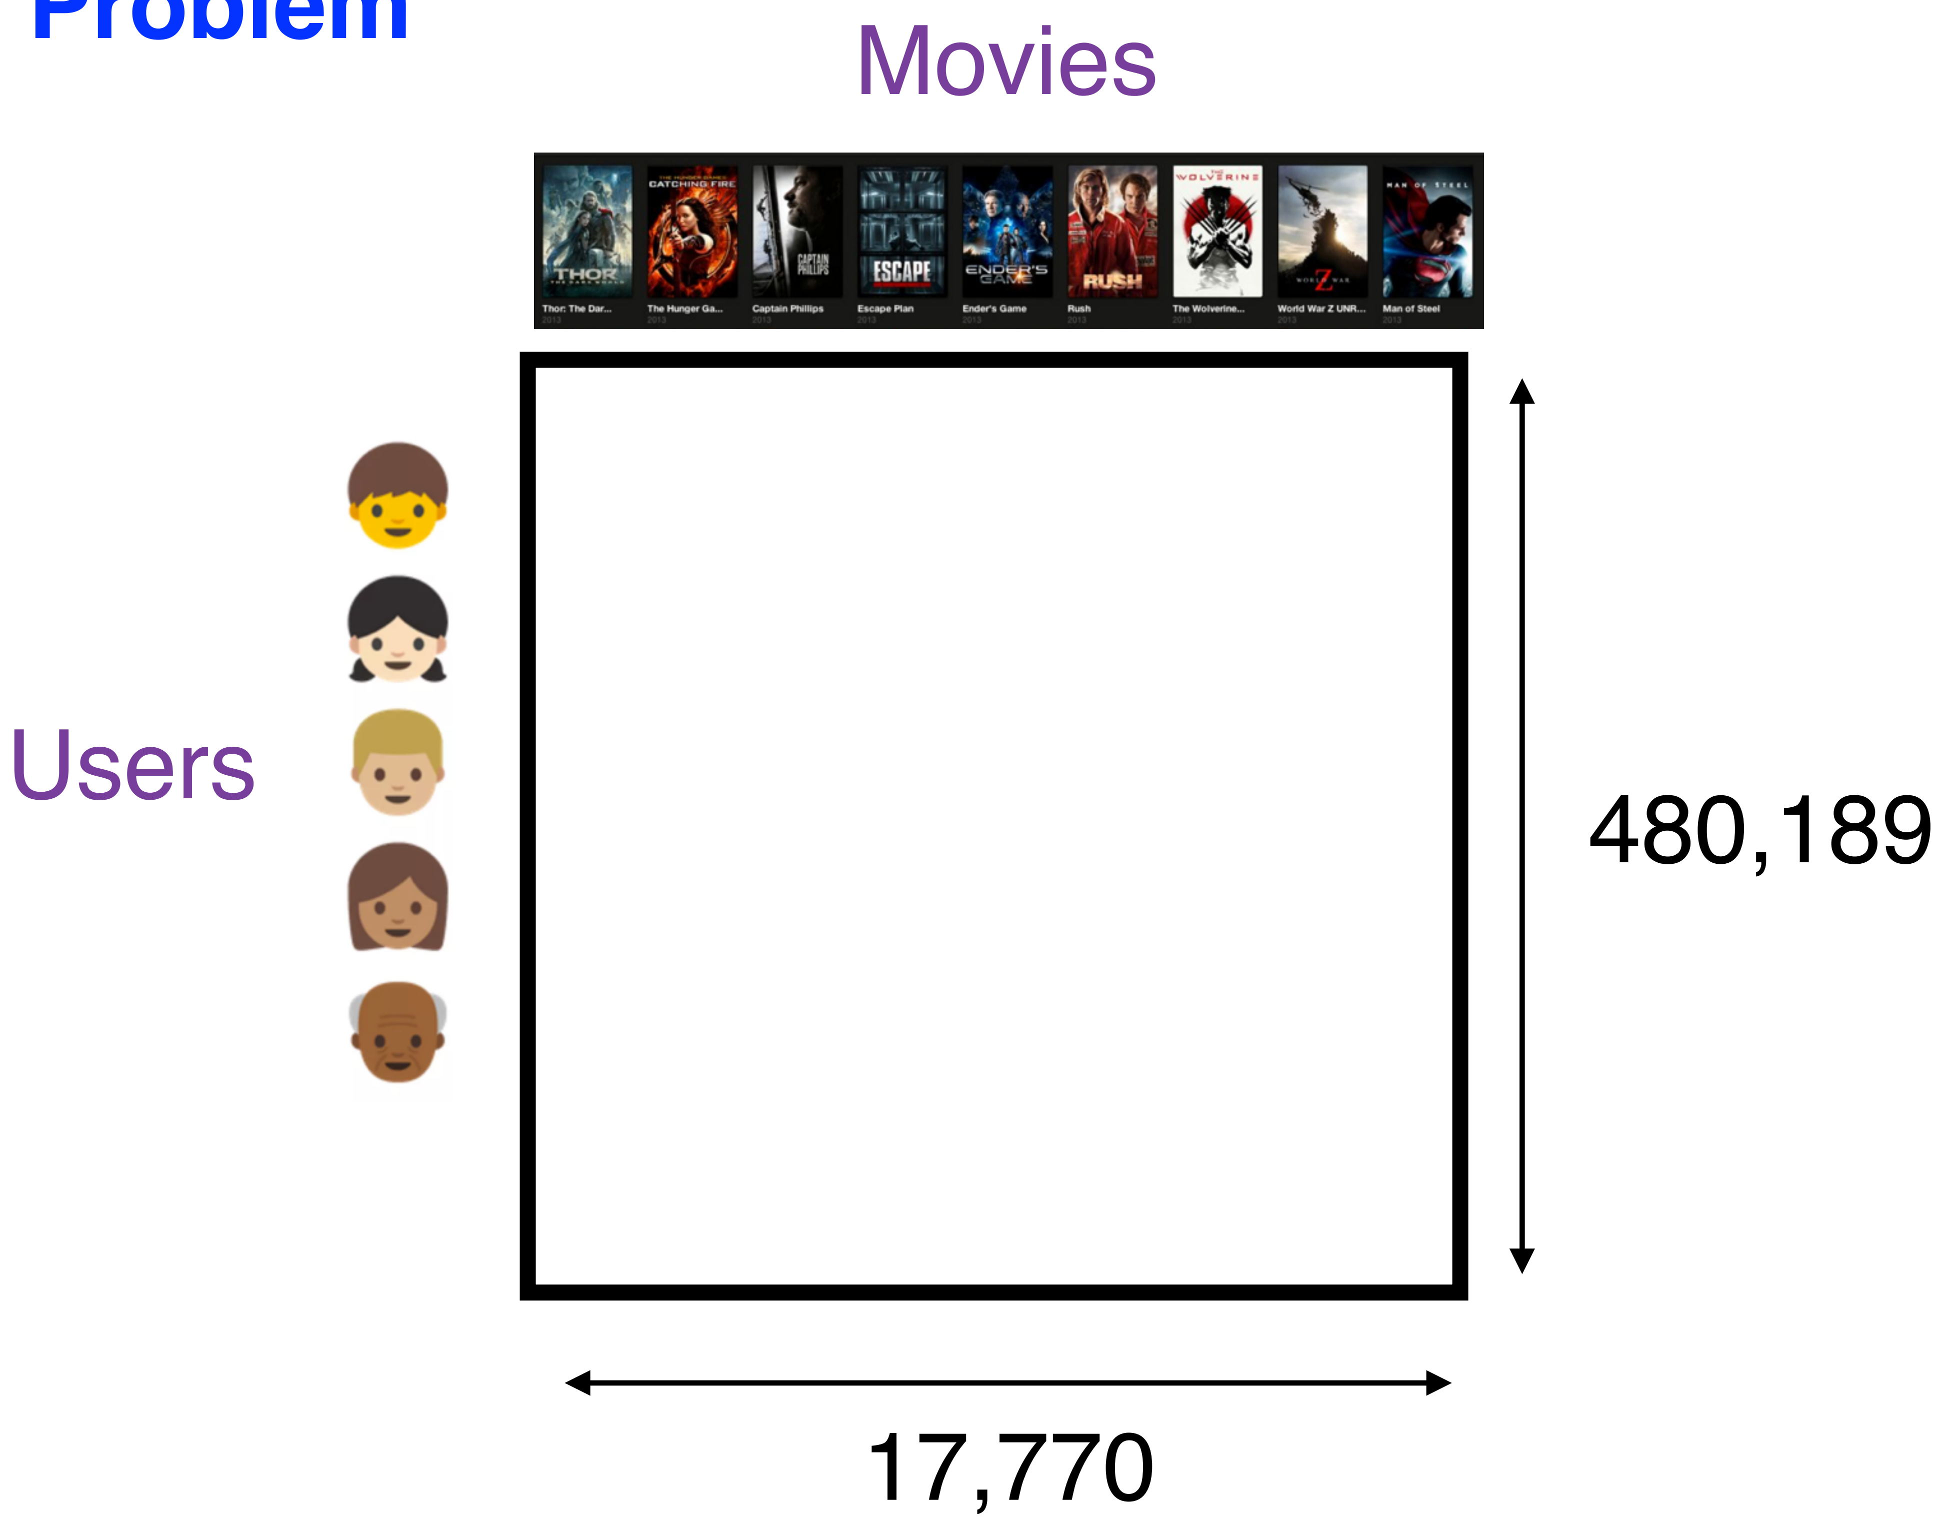

Netflix Problem

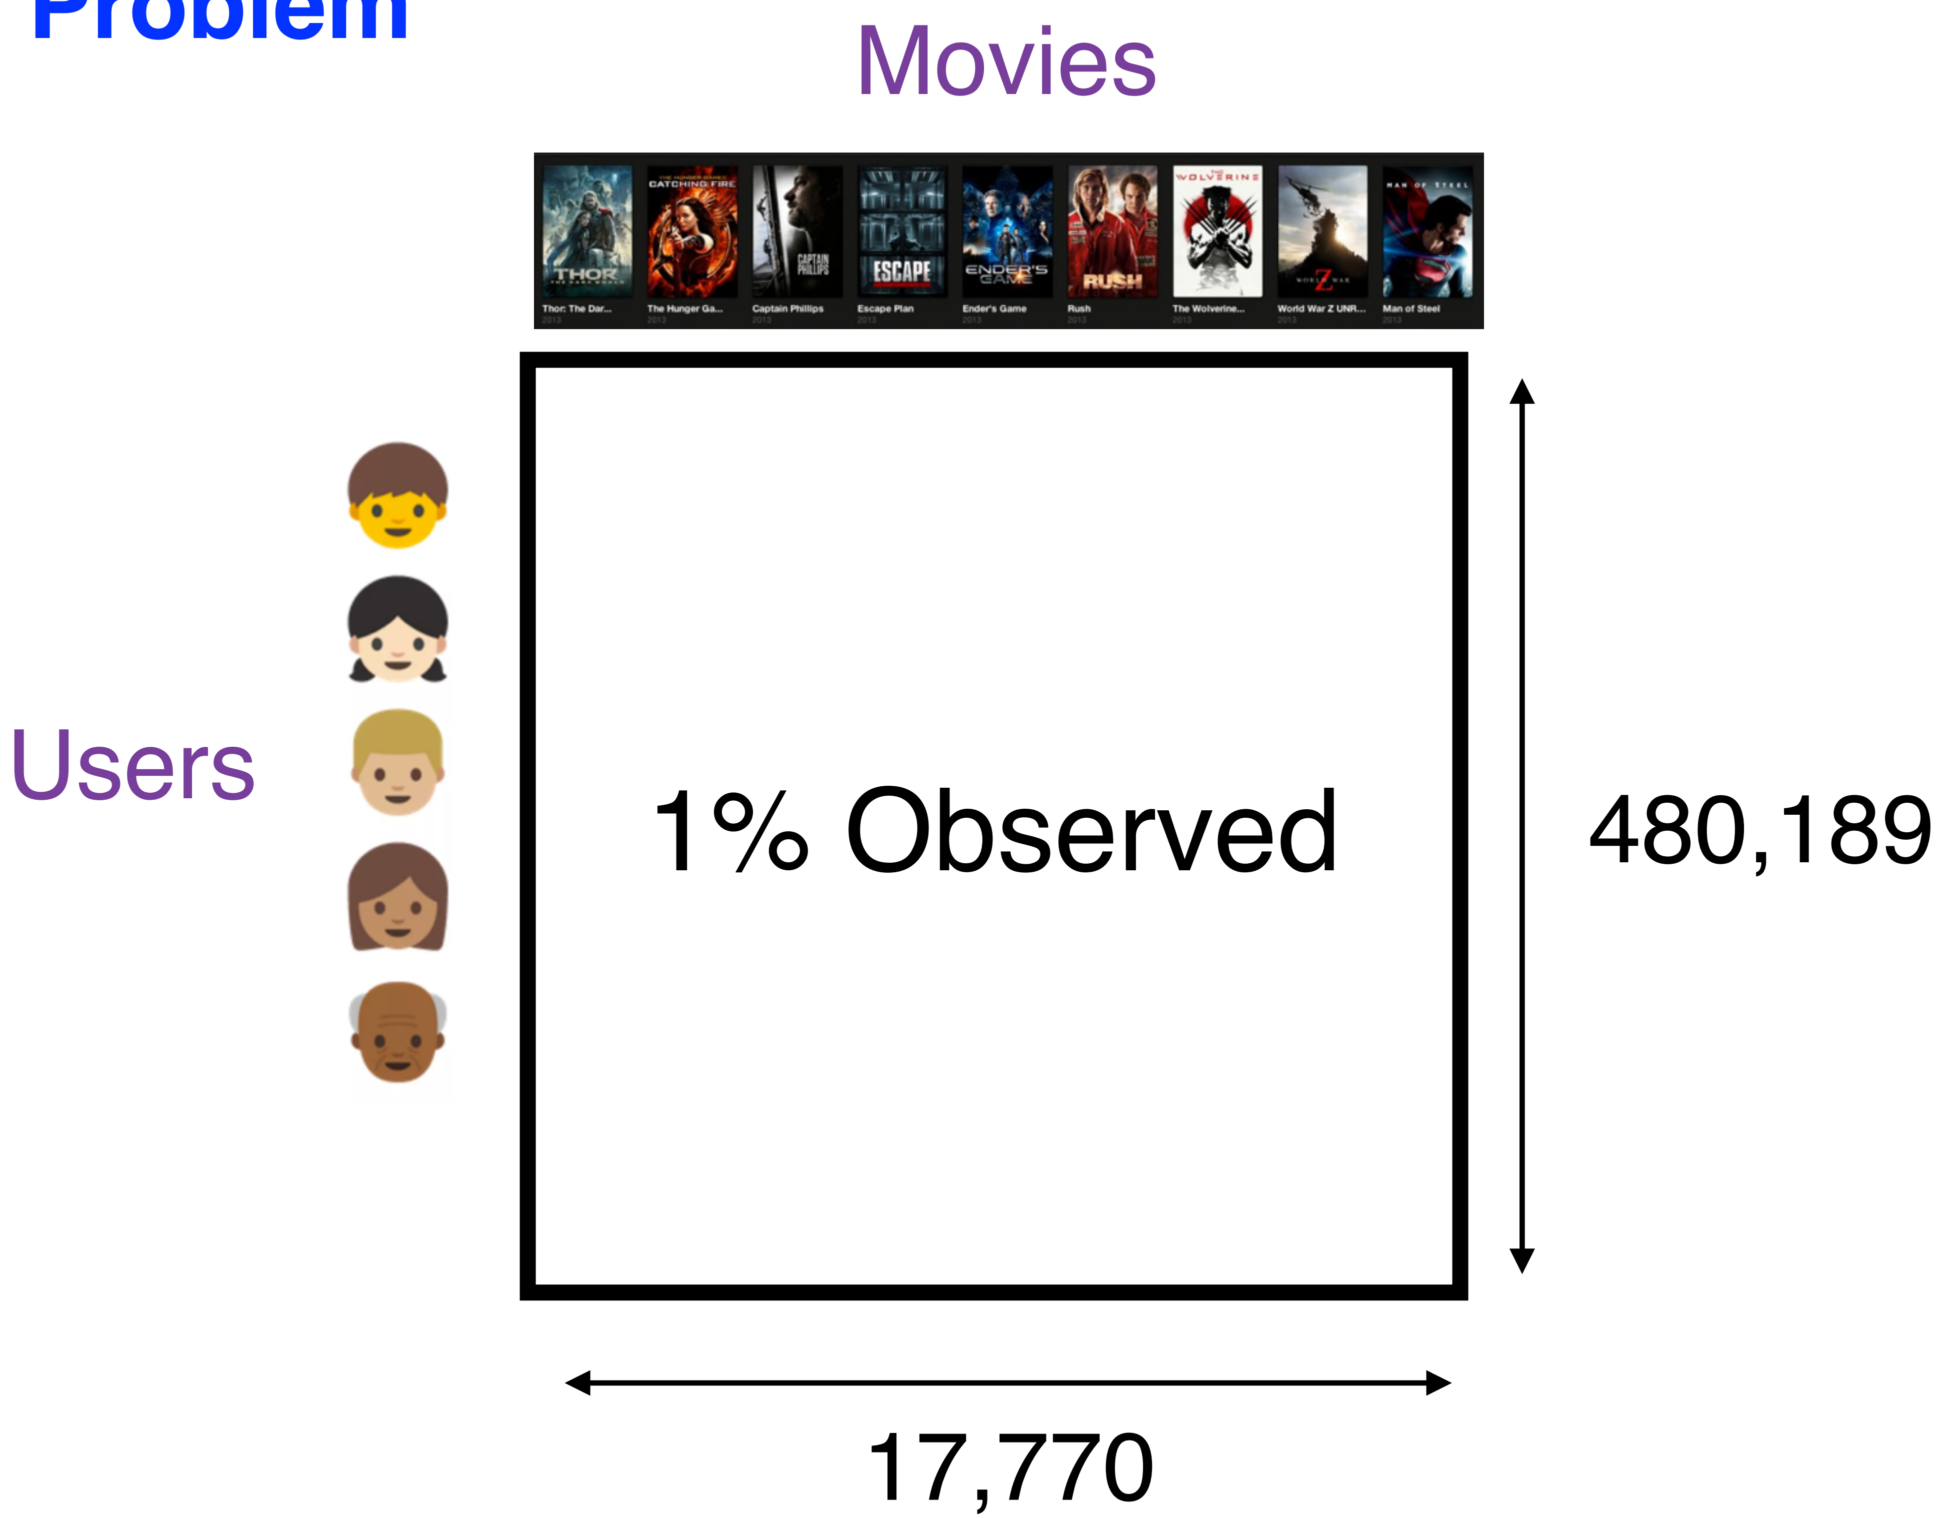

Netflix Problem

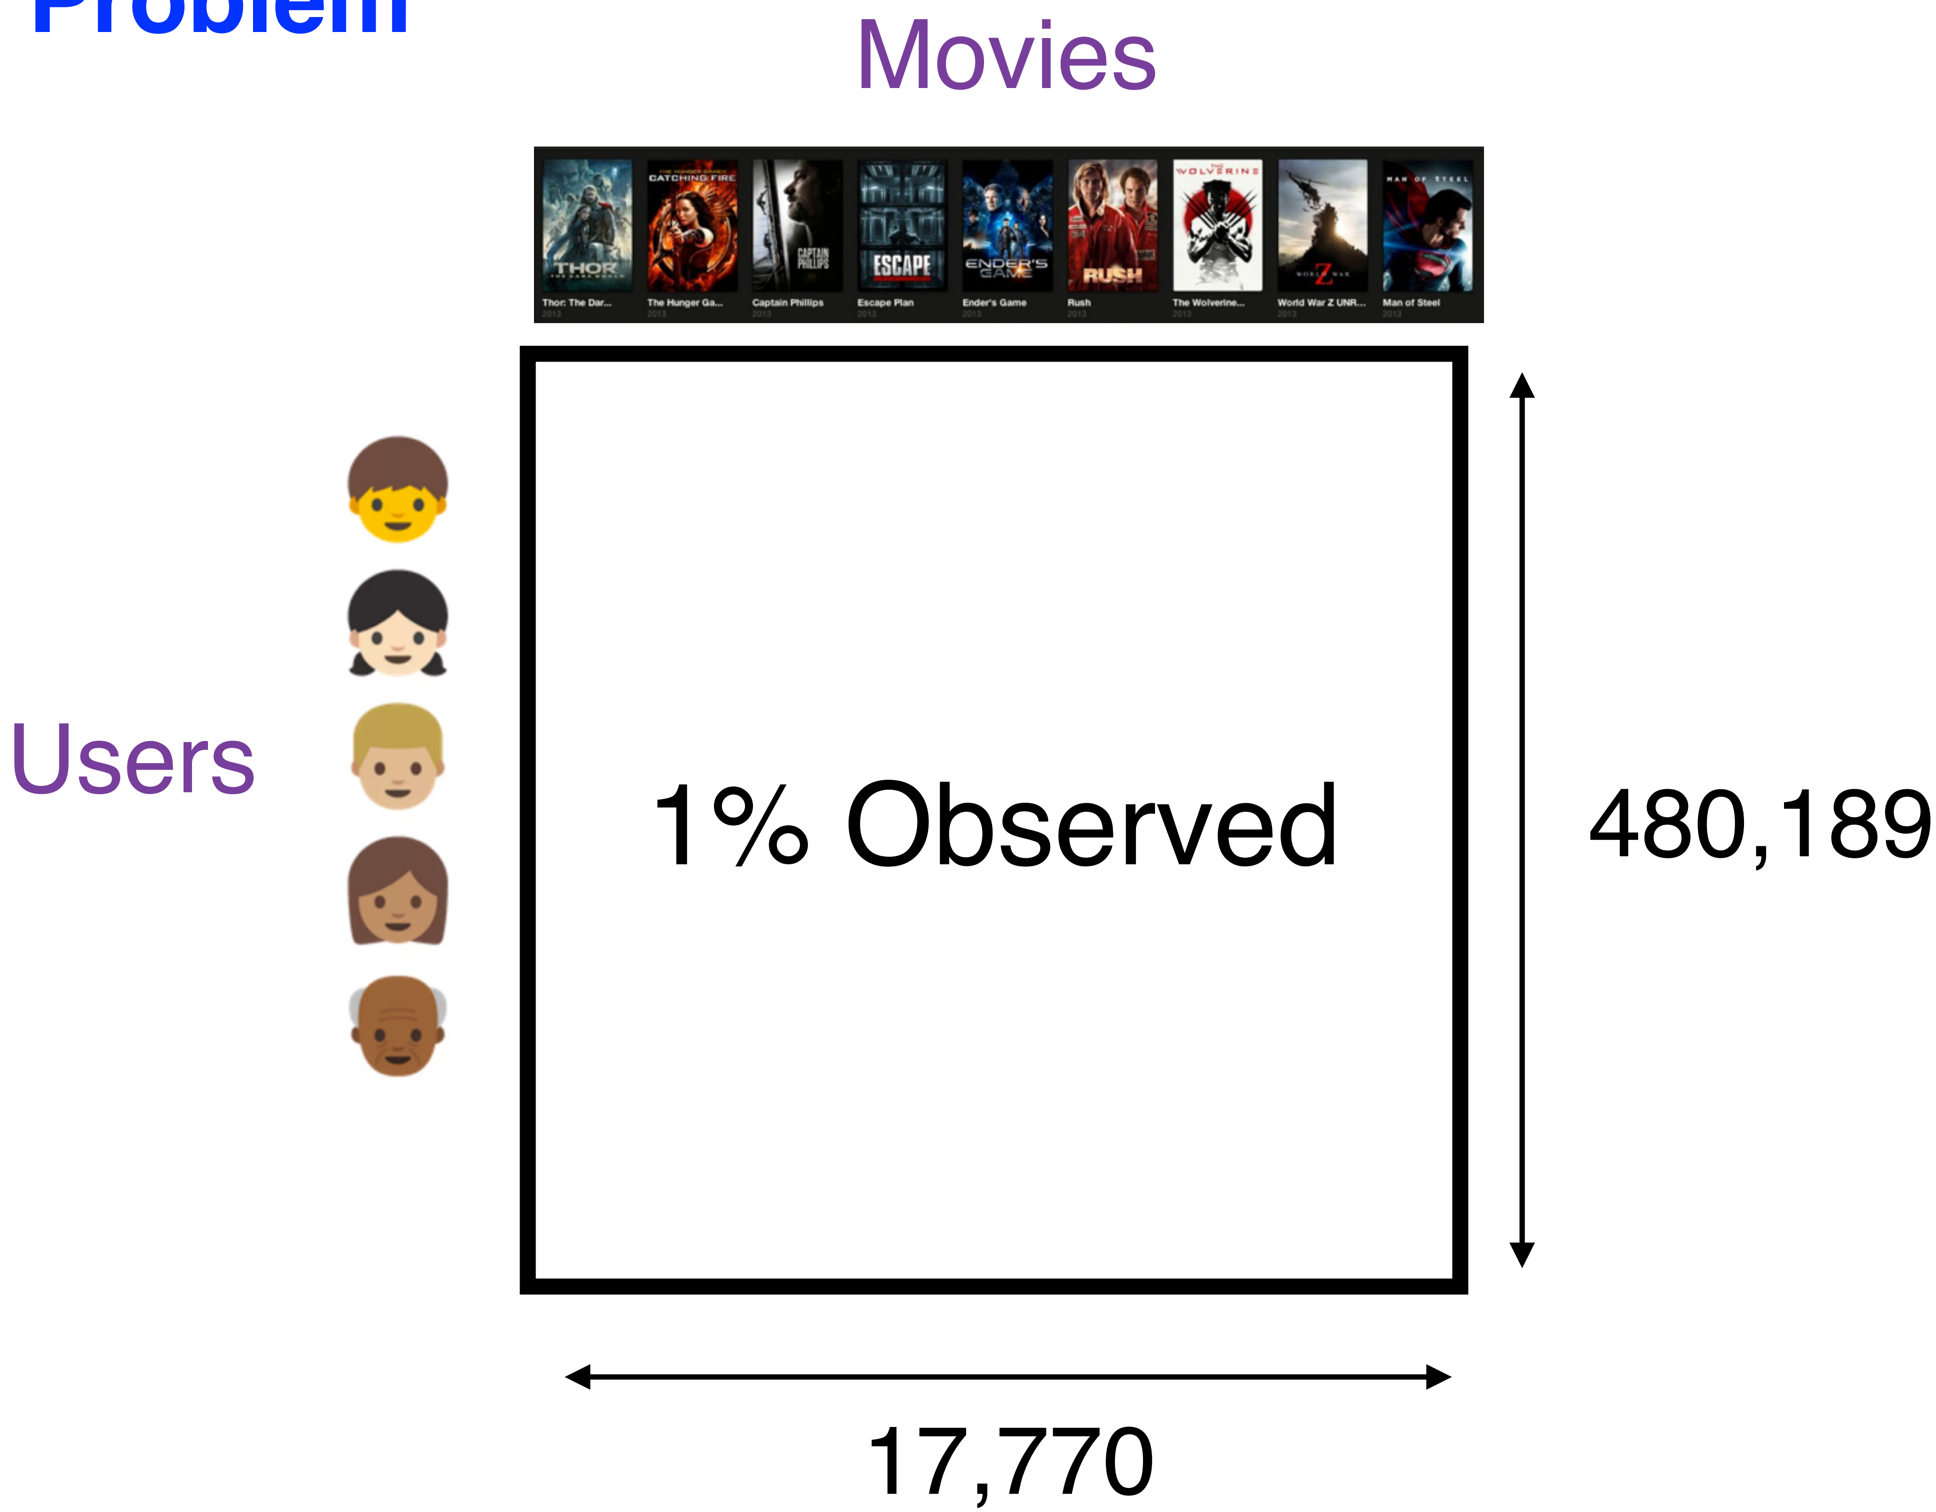

Can we (approximately) fill-in the missing entries?
(i.e., recommend movies to users)

Netflix Prize

Competition in 2006-2009

Netflix Prize

Home | Rules | Leaderboard | Update

Leaderboard

Showing Test Score. Click here to show quiz score

Rank	Team Name	Best Test Score	% Improvement	Best Submit Time
Grand Prize - RMSE = 0.8567 - Winning Team: BellKor's Pragmatic Chaos				
1	BellKor's Pragmatic Chaos	0.8567	10.06	2009-07-26 18:18:28
2	The Ensemble	0.8567	10.06	2009-07-26 18:38:22
3	Grand Prize Team	0.8582	9.90	2009-07-10 21:24:40
4	Opera Solutions and Vandelay United	0.8588	9.84	2009-07-10 01:12:31
5	Vandelay Industries !	0.8591	9.81	2009-07-10 00:32:20
6	PragmaticTheory	0.8594	9.77	2009-06-24 12:06:56
7	BellKor in BigChaos	0.8601	9.70	2009-05-13 08:14:09
8	Dace_	0.8612	9.59	2009-07-24 17:18:43
9	Feeds2	0.8622	9.48	2009-07-12 13:11:51
10	BigChaos	0.8623	9.47	2009-04-07 12:33:59
11	Opera Solutions	0.8623	9.47	2009-07-24 00:34:07
12	BellKor	0.8624	9.46	2009-07-26 17:19:11

http://www.netflixprize.com/leaderboard.html

Netflix Prize

Competition in 2006-2009

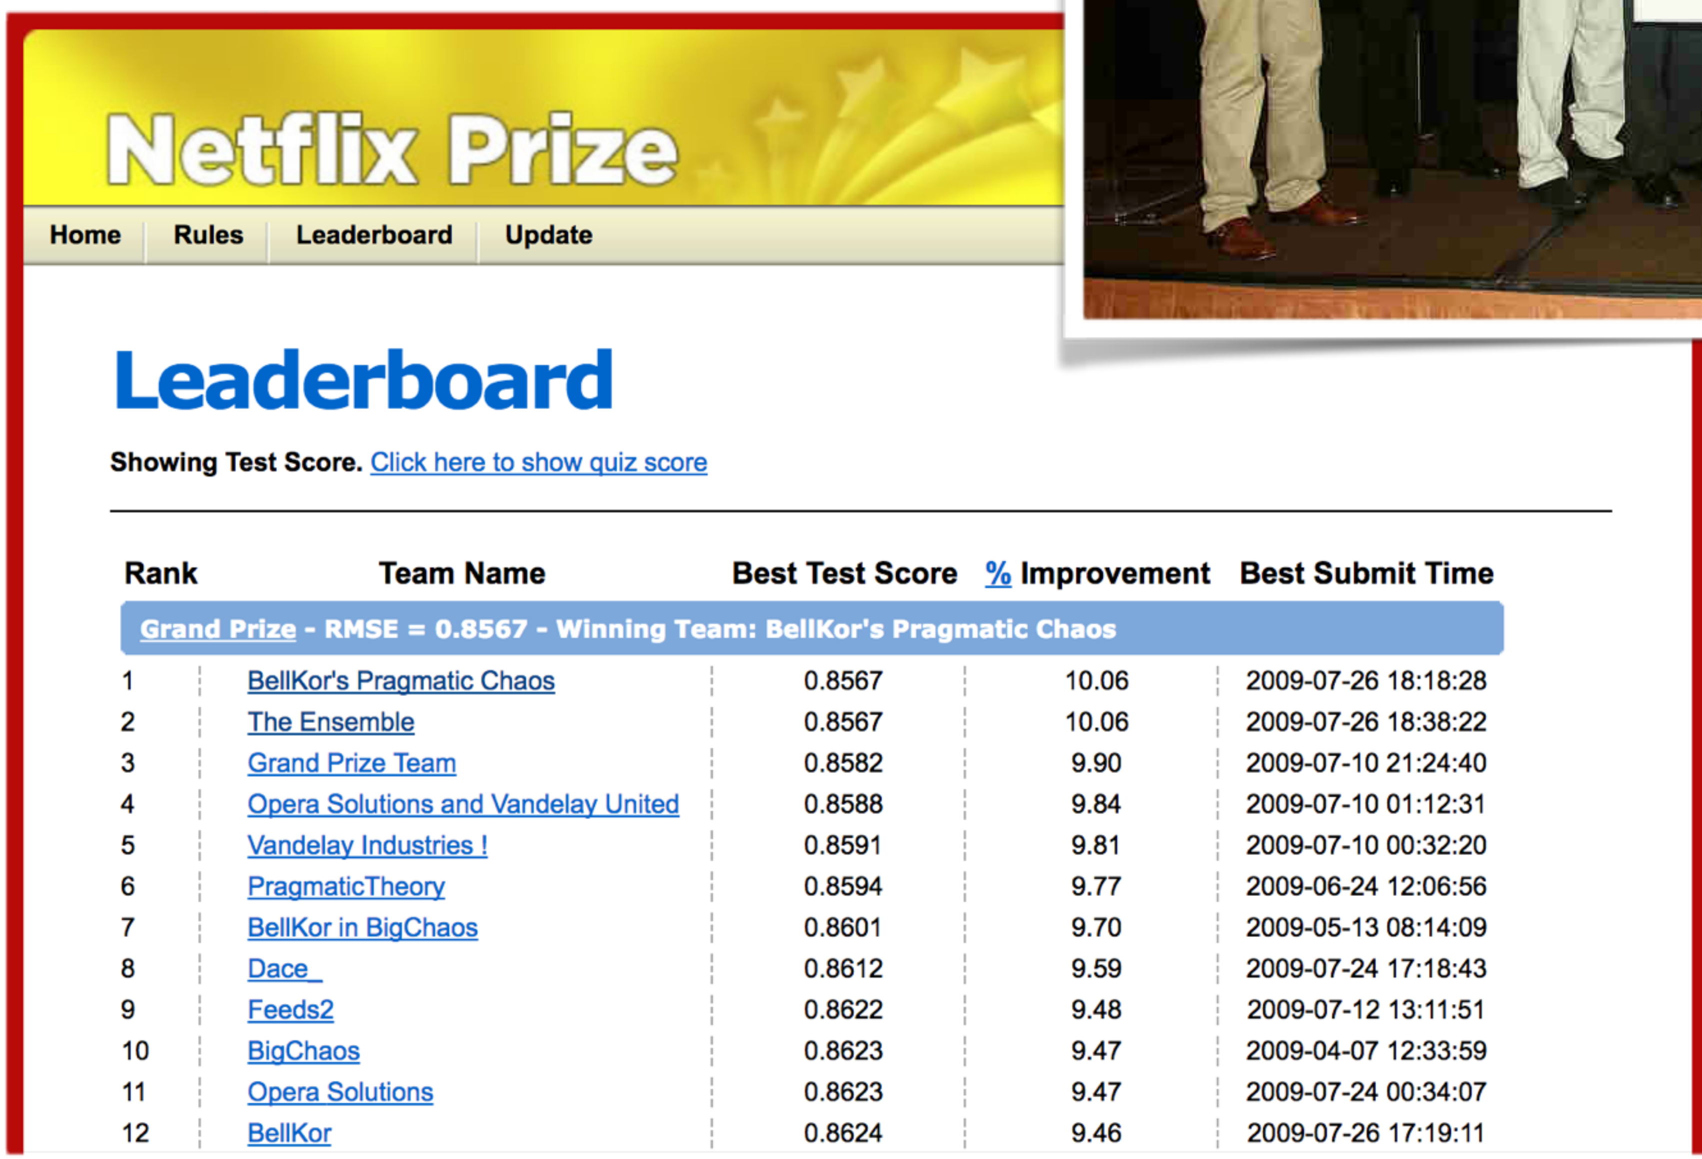

Leaderboard

Showing Test Score. Click here to show quiz score

Rank	Team Name	Best Test Score	% Improvement	Best Submit Time
Grand Prize - RMSE = 0.8567 - Winning Team: BellKor's Pragmatic Chaos				
1	BellKor's Pragmatic Chaos	0.8567	10.06	2009-07-26 18:18:28
2	The Ensemble	0.8567	10.06	2009-07-26 18:38:22
3	Grand Prize Team	0.8582	9.90	2009-07-10 21:24:40
4	Opera Solutions and Vandelay United	0.8588	9.84	2009-07-10 01:12:31
5	Vandelay Industries !	0.8591	9.81	2009-07-10 00:32:20
6	PragmaticTheory	0.8594	9.77	2009-06-24 12:06:56
7	BellKor in BigChaos	0.8601	9.70	2009-05-13 08:14:09
8	Dace_	0.8612	9.59	2009-07-24 17:18:43
9	Feeds2	0.8622	9.48	2009-07-12 13:11:51
10	BigChaos	0.8623	9.47	2009-04-07 12:33:59
11	Opera Solutions	0.8623	9.47	2009-07-24 00:34:07
12	BellKor	0.8624	9.46	2009-07-26 17:19:11

http://www.netflixprize.com/leaderboard.html

- Winning team used complex combination of methods
- Simple technique gets within 3% of top score in RMSE: **Low-rank matrix completion**

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$X \approx u_1 v_1' + u_2 v_2' + \cdots + u_r v_r'$$

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$\mathsf{X} \approx \mathsf{u}_1\mathsf{v}_1' + \mathsf{u}_2\mathsf{v}_2' + \cdots + \mathsf{u}_r\mathsf{v}_r'$$

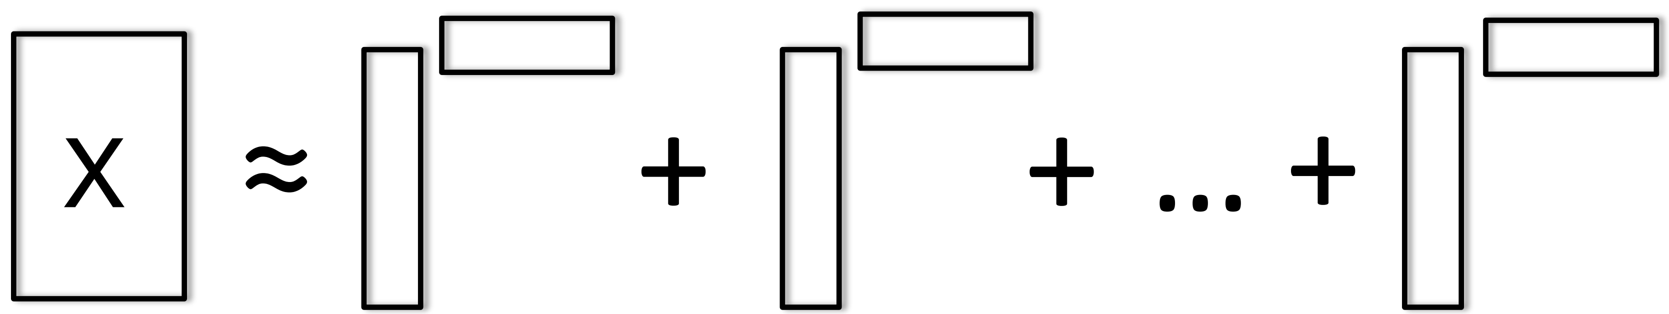

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$X \approx u_1 v_1' + u_2 v_2' + \cdots + u_r v_r'$$

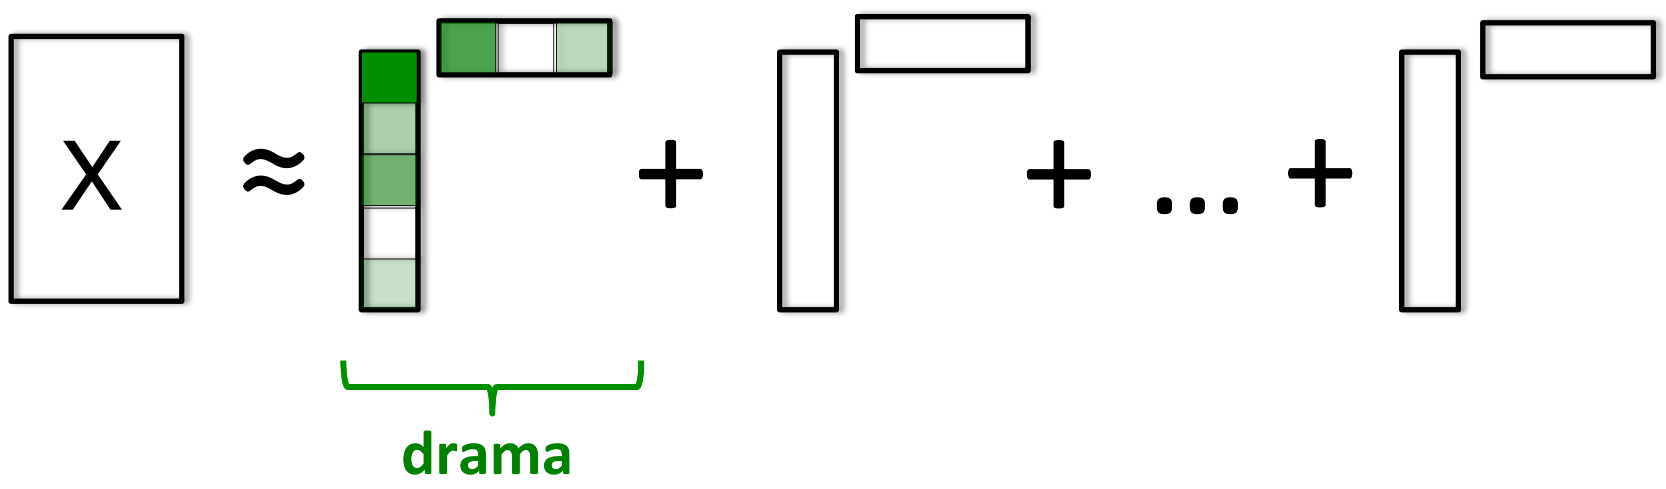

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$X \approx u_1 v_1' + u_2 v_2' + \cdots + u_r v_r'$$

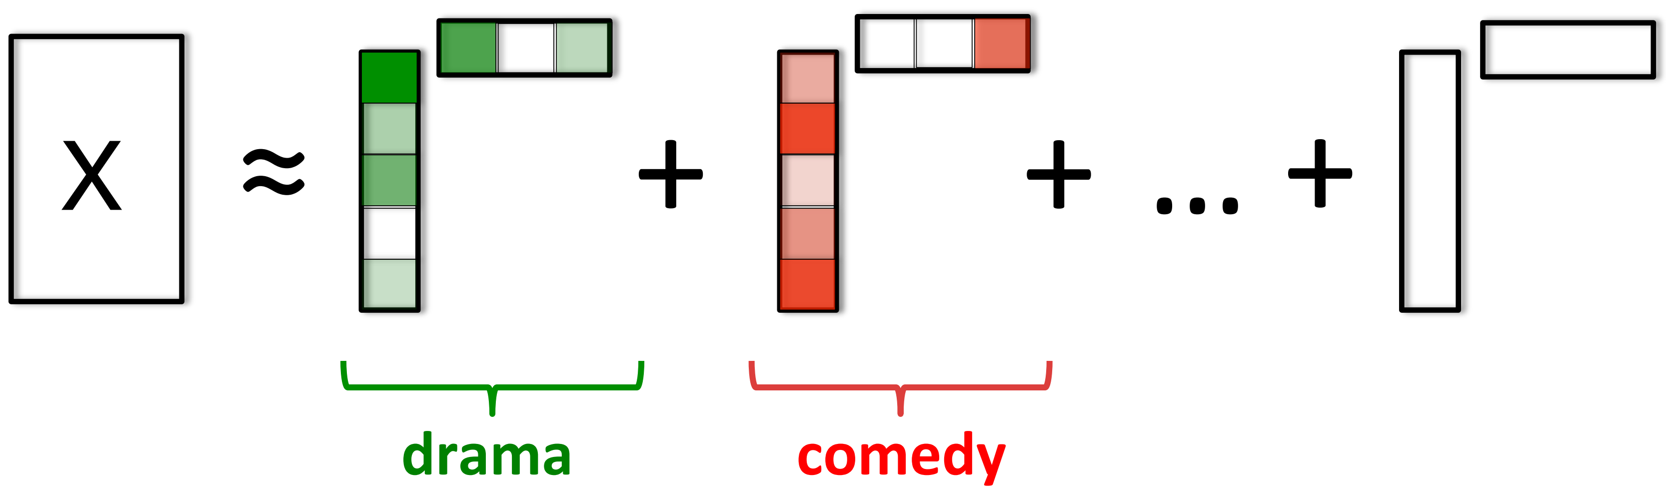

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$X \approx u_1 v_1' + u_2 v_2' + \cdots + u_r v_r'$$

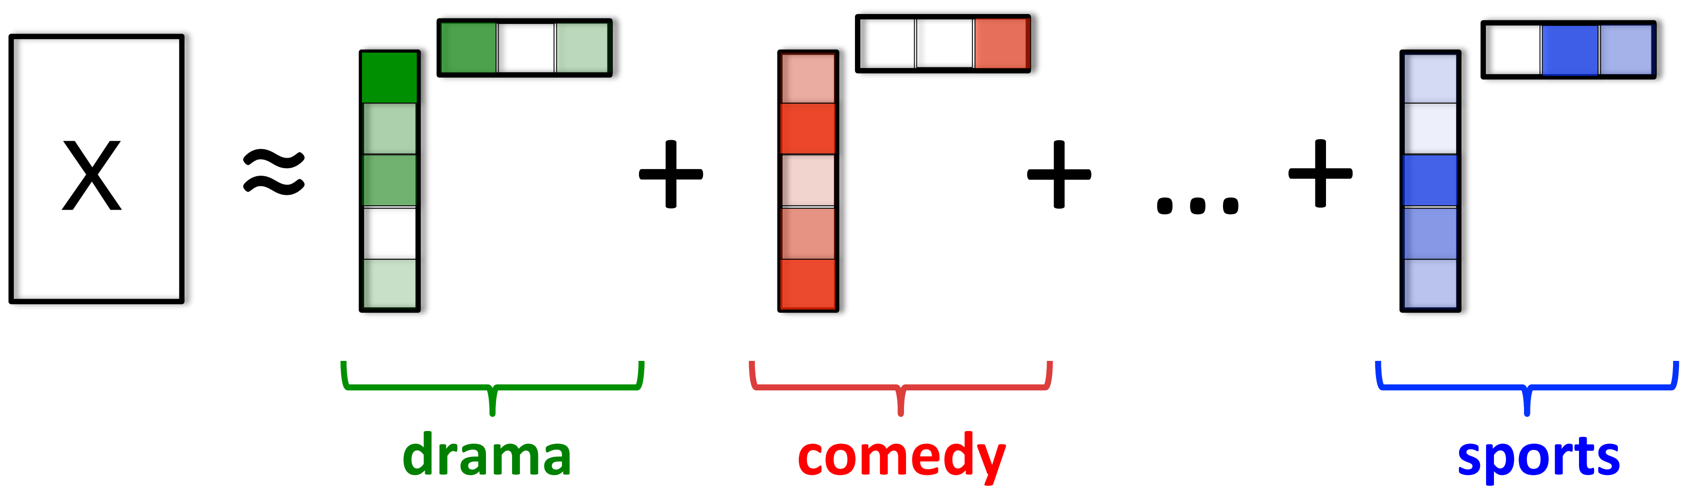

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$X \approx u_1 v_1' + u_2 v_2' + \cdots + u_r v_r'$$

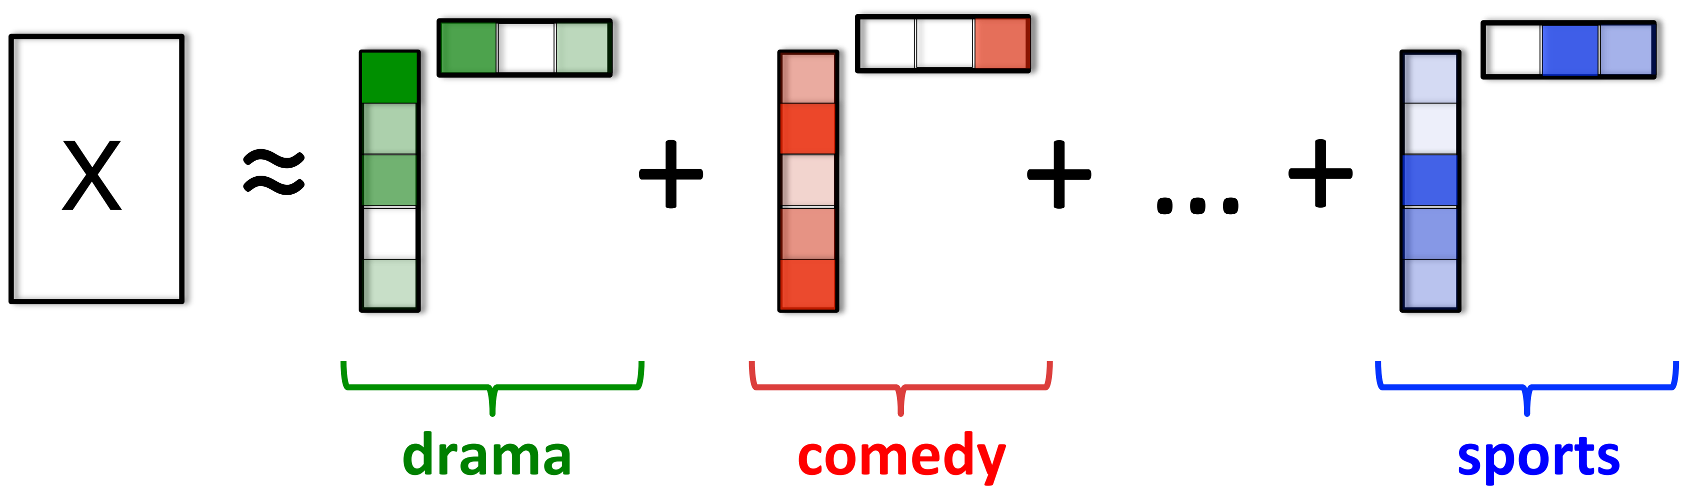

Model: We are given random observations

$X_{i,j}$ for all $(i,j) \in \Omega$ (observation set)

Low-Rank Matrix Completion (LRMC)

Let X be an approximately low-rank matrix

$$X \approx u_1 v_1' + u_2 v_2' + \cdots + u_r v_r'$$

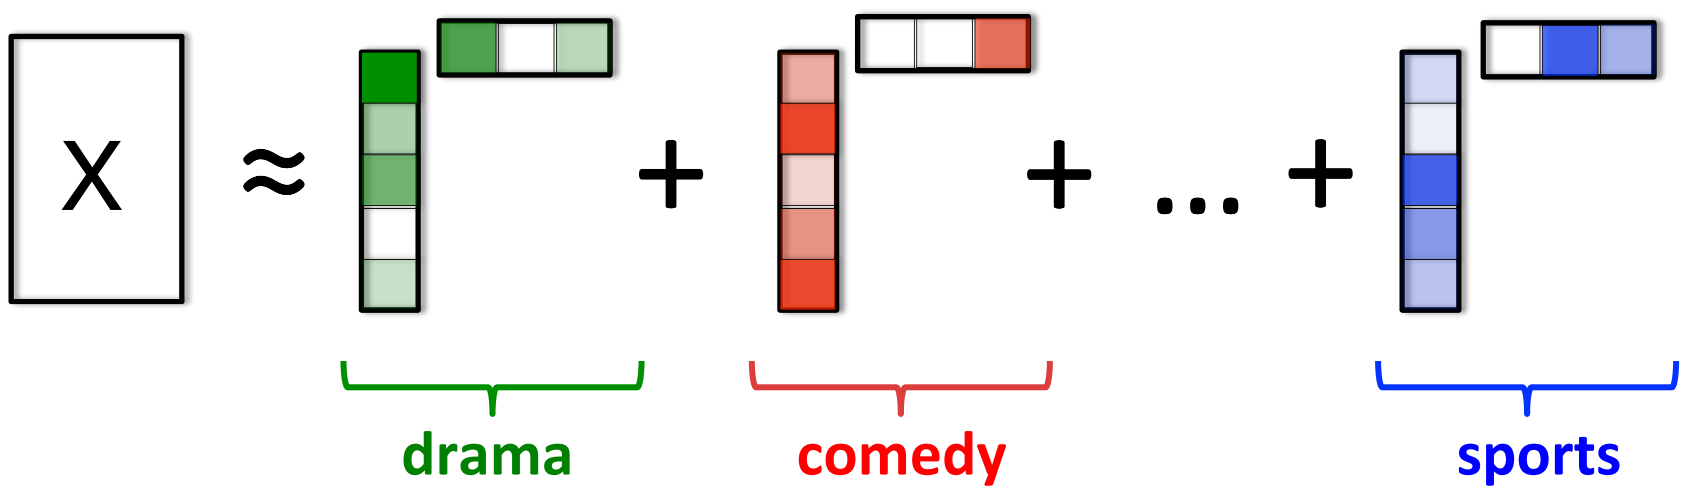

Model: We are given random observations

$X_{i,j}$ for all $(i,j) \in \Omega$ (observation set)

Is there an efficient algorithm to recover X?

Successful Applications of LRMC

- Recommender systems ("Netflix prize")
- Imaging: denoising, reconstruction in medical, hyperspectral imaging.
- Anomaly detection in network flows
- Source localization and target tracking in radar and sonar
- Computer vision: background subtraction, object tracking, and to represent a single scene under varying illuminations
- Environmental monitoring of soil and crop conditions, water contamination, and air pollution, also sensor calibration
- Seismological activity and modal estimation in materials and manmade structures

...and so on

LRMC Sampling Complexity

Degrees of freedom (DoF) of an $n \times s$ rank r matrix:

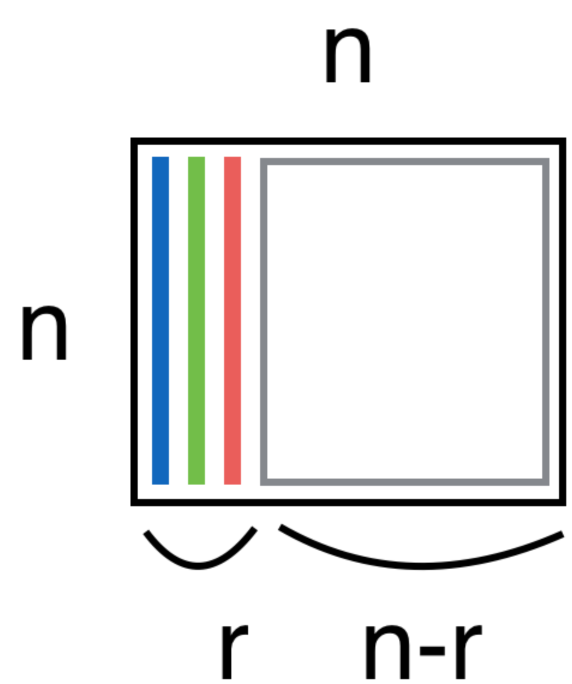

$$\text{DoF} = nr + r(n-r)$$
$$\approx 2nr$$

LRMC Sampling Complexity

Degrees of freedom (DoF) of an $n \times s$ rank r matrix:

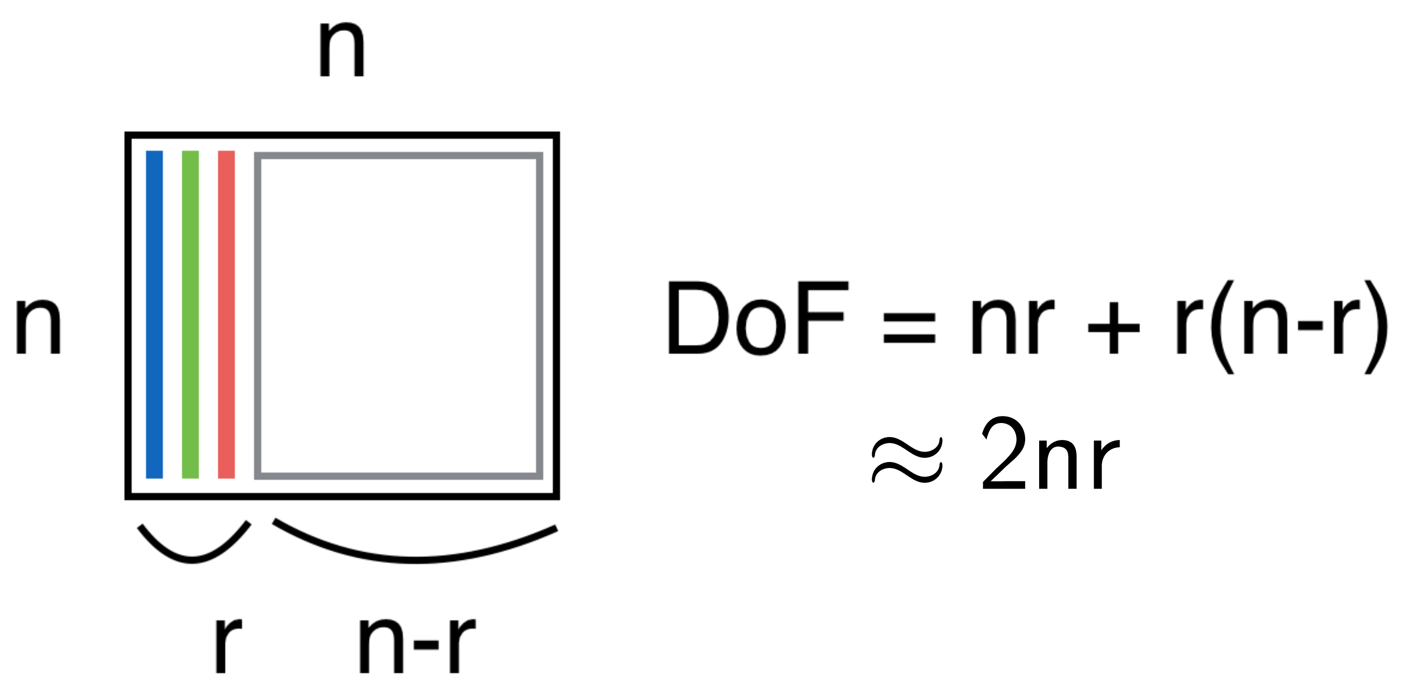

Implies we need $O(nr)$ samples for LRMC to even be possible.

State-of-the-art algorithms *provably* complete low-rank matrices from $O(nr\,\text{polylog}(n))$ random samples.

LRMC Sampling Complexity

Degrees of freedom (DoF) of an $n \times s$ rank r matrix:

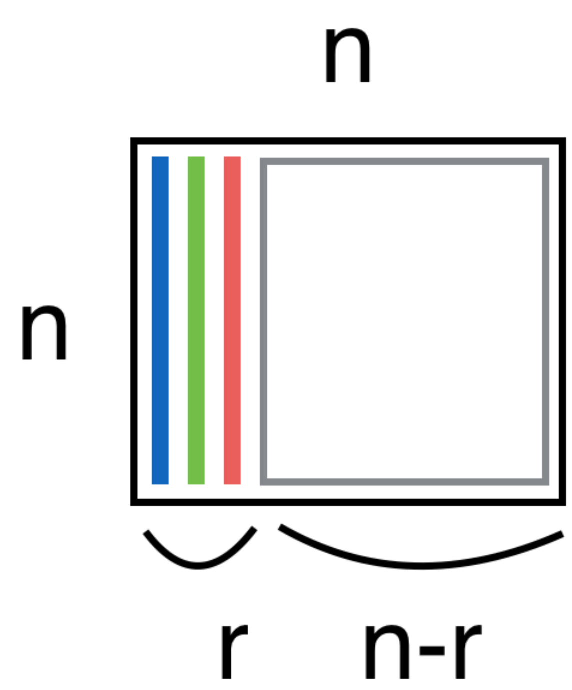

$$\text{DoF} = nr + r(n-r) \approx 2nr$$

Implies we need $O(nr)$ samples for LRMC to even be possible.

State-of-the-art algorithms *provably* complete low-rank matrices from $O(nr\,\text{polylog}(n))$ random samples.

Today will we talk about one algorithm:
Nuclear norm minimization

Clicker Questions

Approximately how many observations per column are necessary to recover a low-rank matrix?

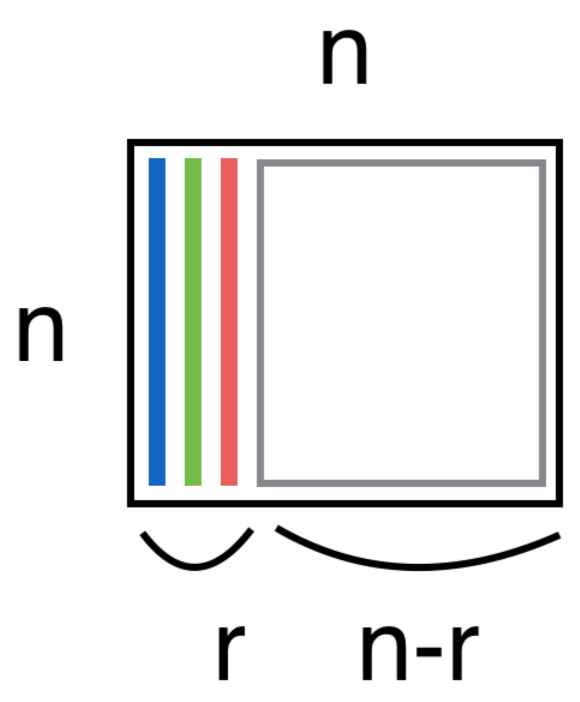

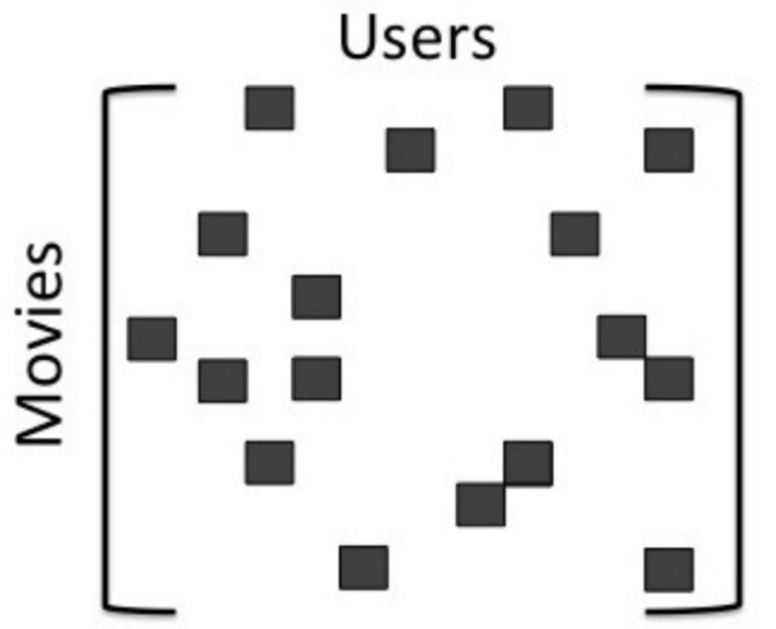

Assume n x n rank r matrix is determined by 2nr parameters.

A) n
B) 2r
C) 2n
D) nr
E) r^2

Approximately how many observations per column are necessary to recover a low-rank matrix?

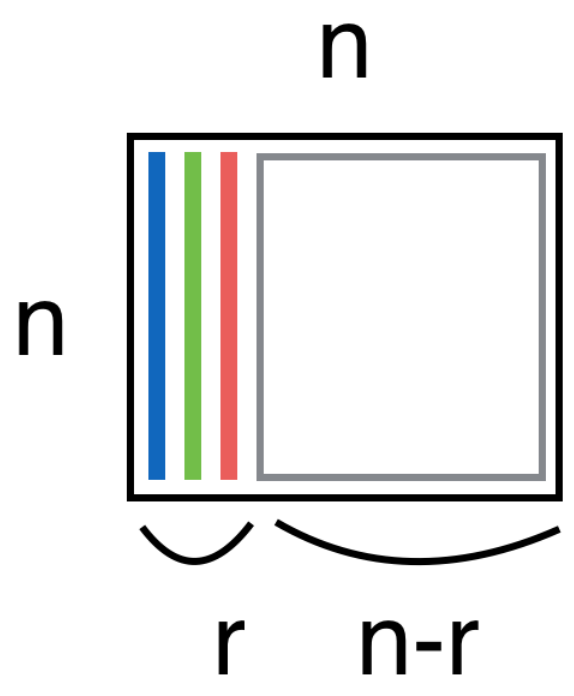

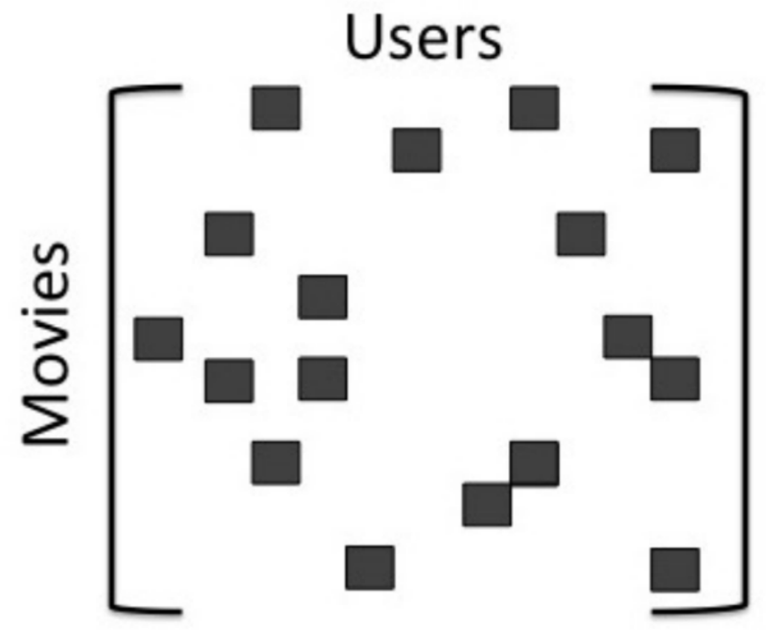

Assume n x n rank r matrix is determined by 2nr parameters.

A) n
B) 2r
C) 2n
D) nr
E) r^2

Approximately how many observations **per column** are necessary to recover a low-rank matrix?

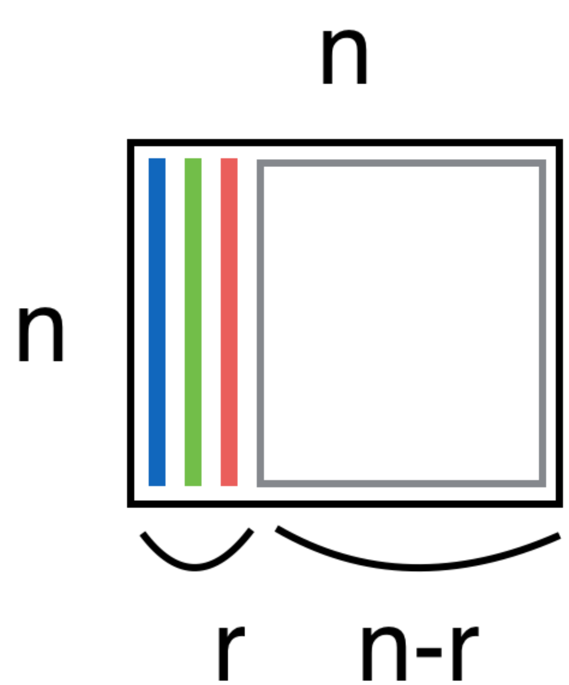

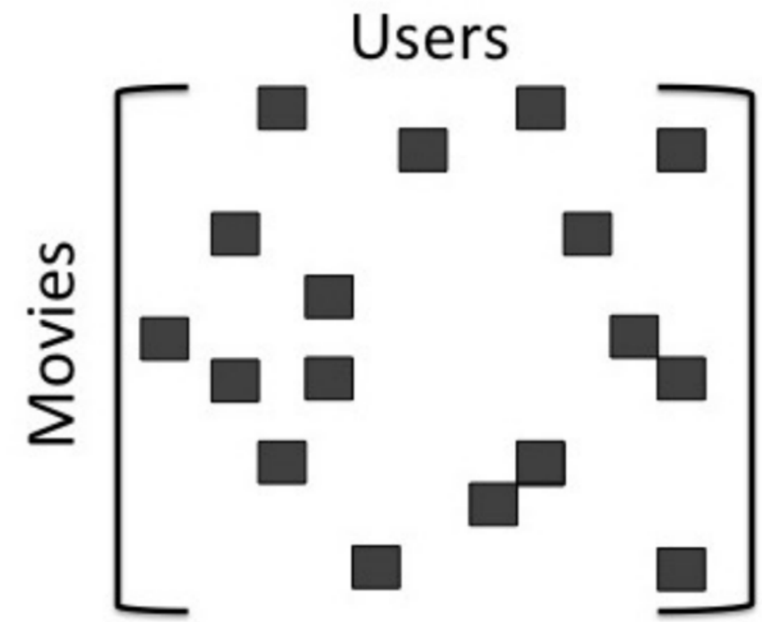

Assume n x n rank r matrix is determined by 2nr parameters.

A) n
B) 2r
C) 2n
D) nr
E) r^2

if low-rank:

$$2r \ll n$$

What is the nuclear norm of a 1x1 matrix x?

$$\left|\left|\left| x \right|\right|\right|_* = \ ?$$

scalar / 1x1 matrix

A) x
B) Heaviside(x)
C) |x|
D) max(0,x)
E) x^2

What is the nuclear norm of a 1x1 matrix x?

$$|||x|||_* = \ ?$$

scalar / 1x1 matrix

A) x
B) Heaviside(x)
C) |x|
D) max(0,x)
E) x^2

What is the nuclear norm of a 1x1 matrix x?

$$|||x|||_* = \ ?$$

scalar / 1x1 matrix

$$|||x|||_* = |x|$$

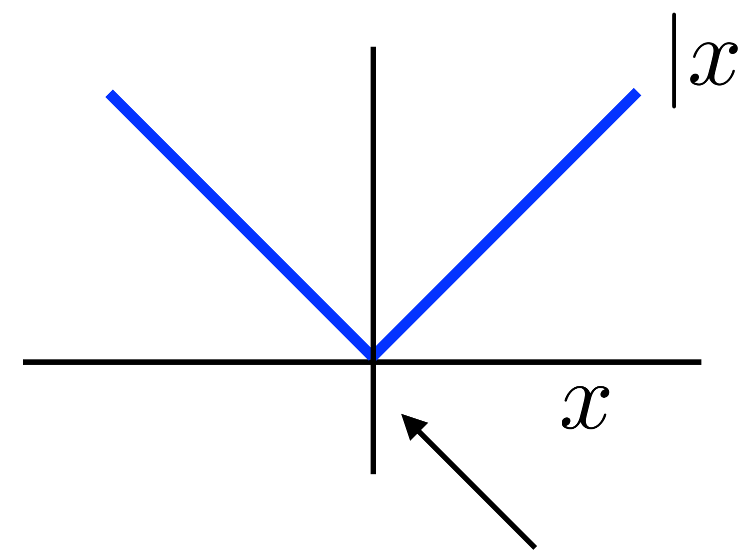

A) x
B) Heaviside(x)
C) |x|
D) max(0,x)
E) x^2

What went wrong with the ISTA algorithm?

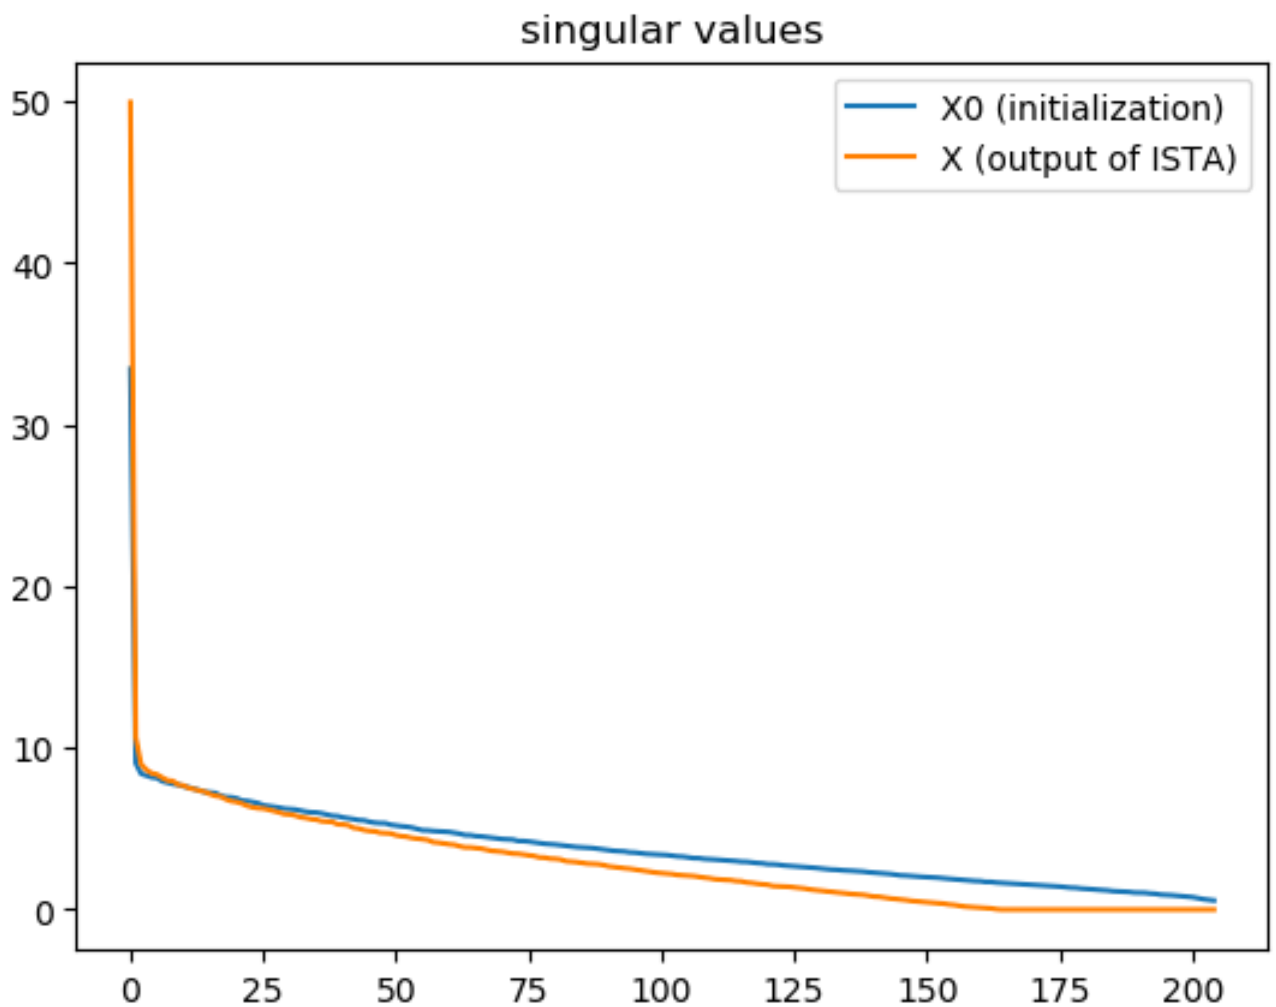

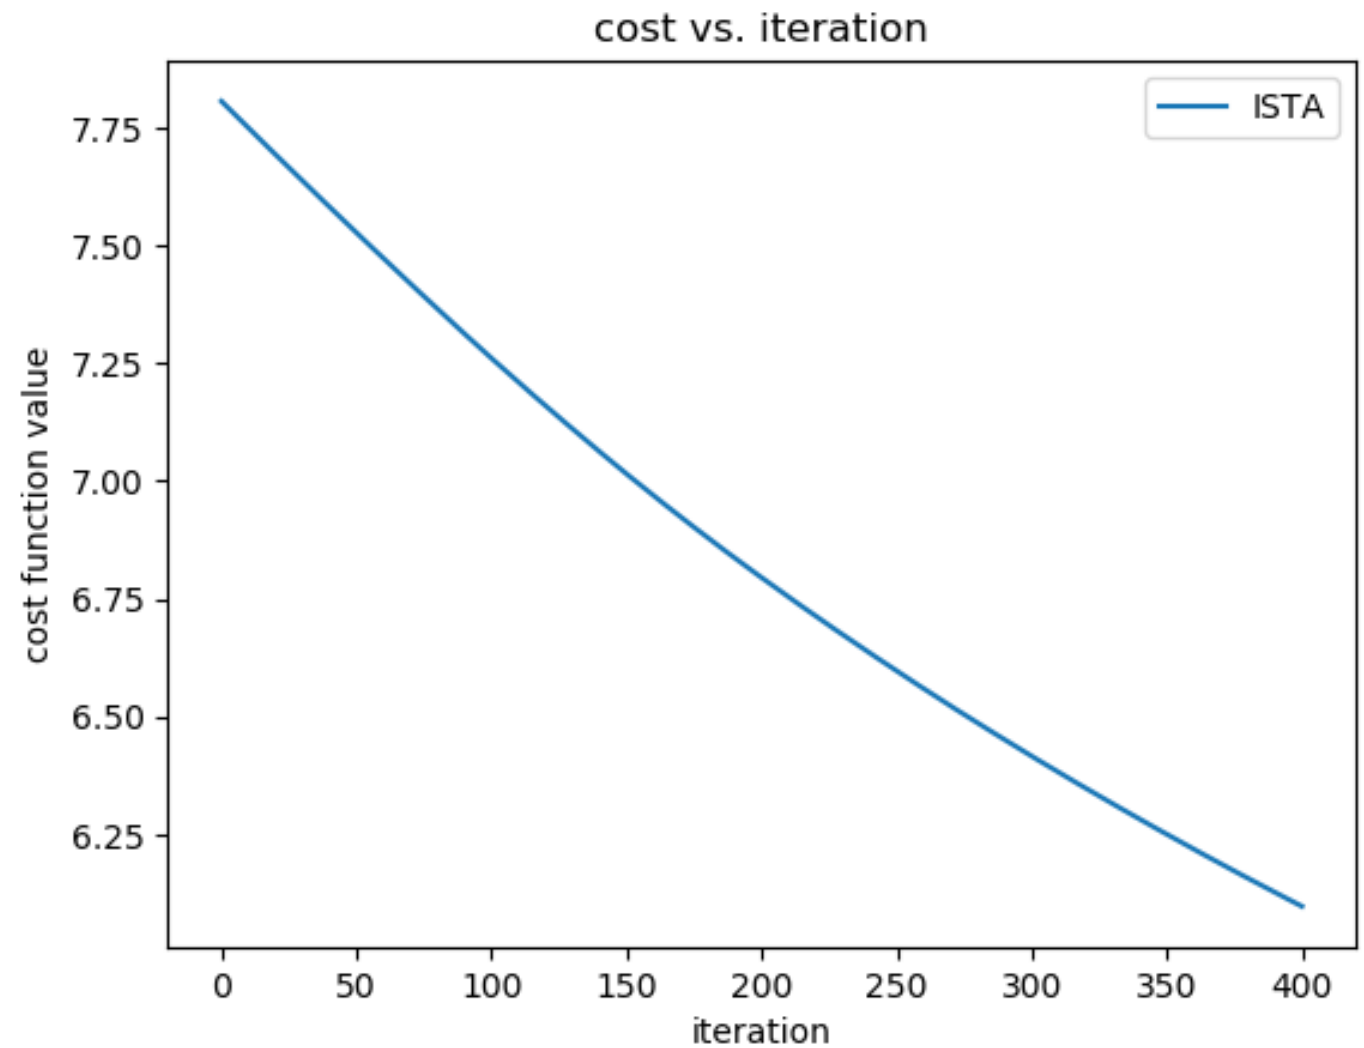

A) Nothing. The algorithm found a low-rank solution.
B) Bug in the implementation of SVT.
C) Bad initialization.
D) Algorithm has not converged—run more iterations.
E) Ran too many iterations.

What went wrong with the ISTA algorithm?

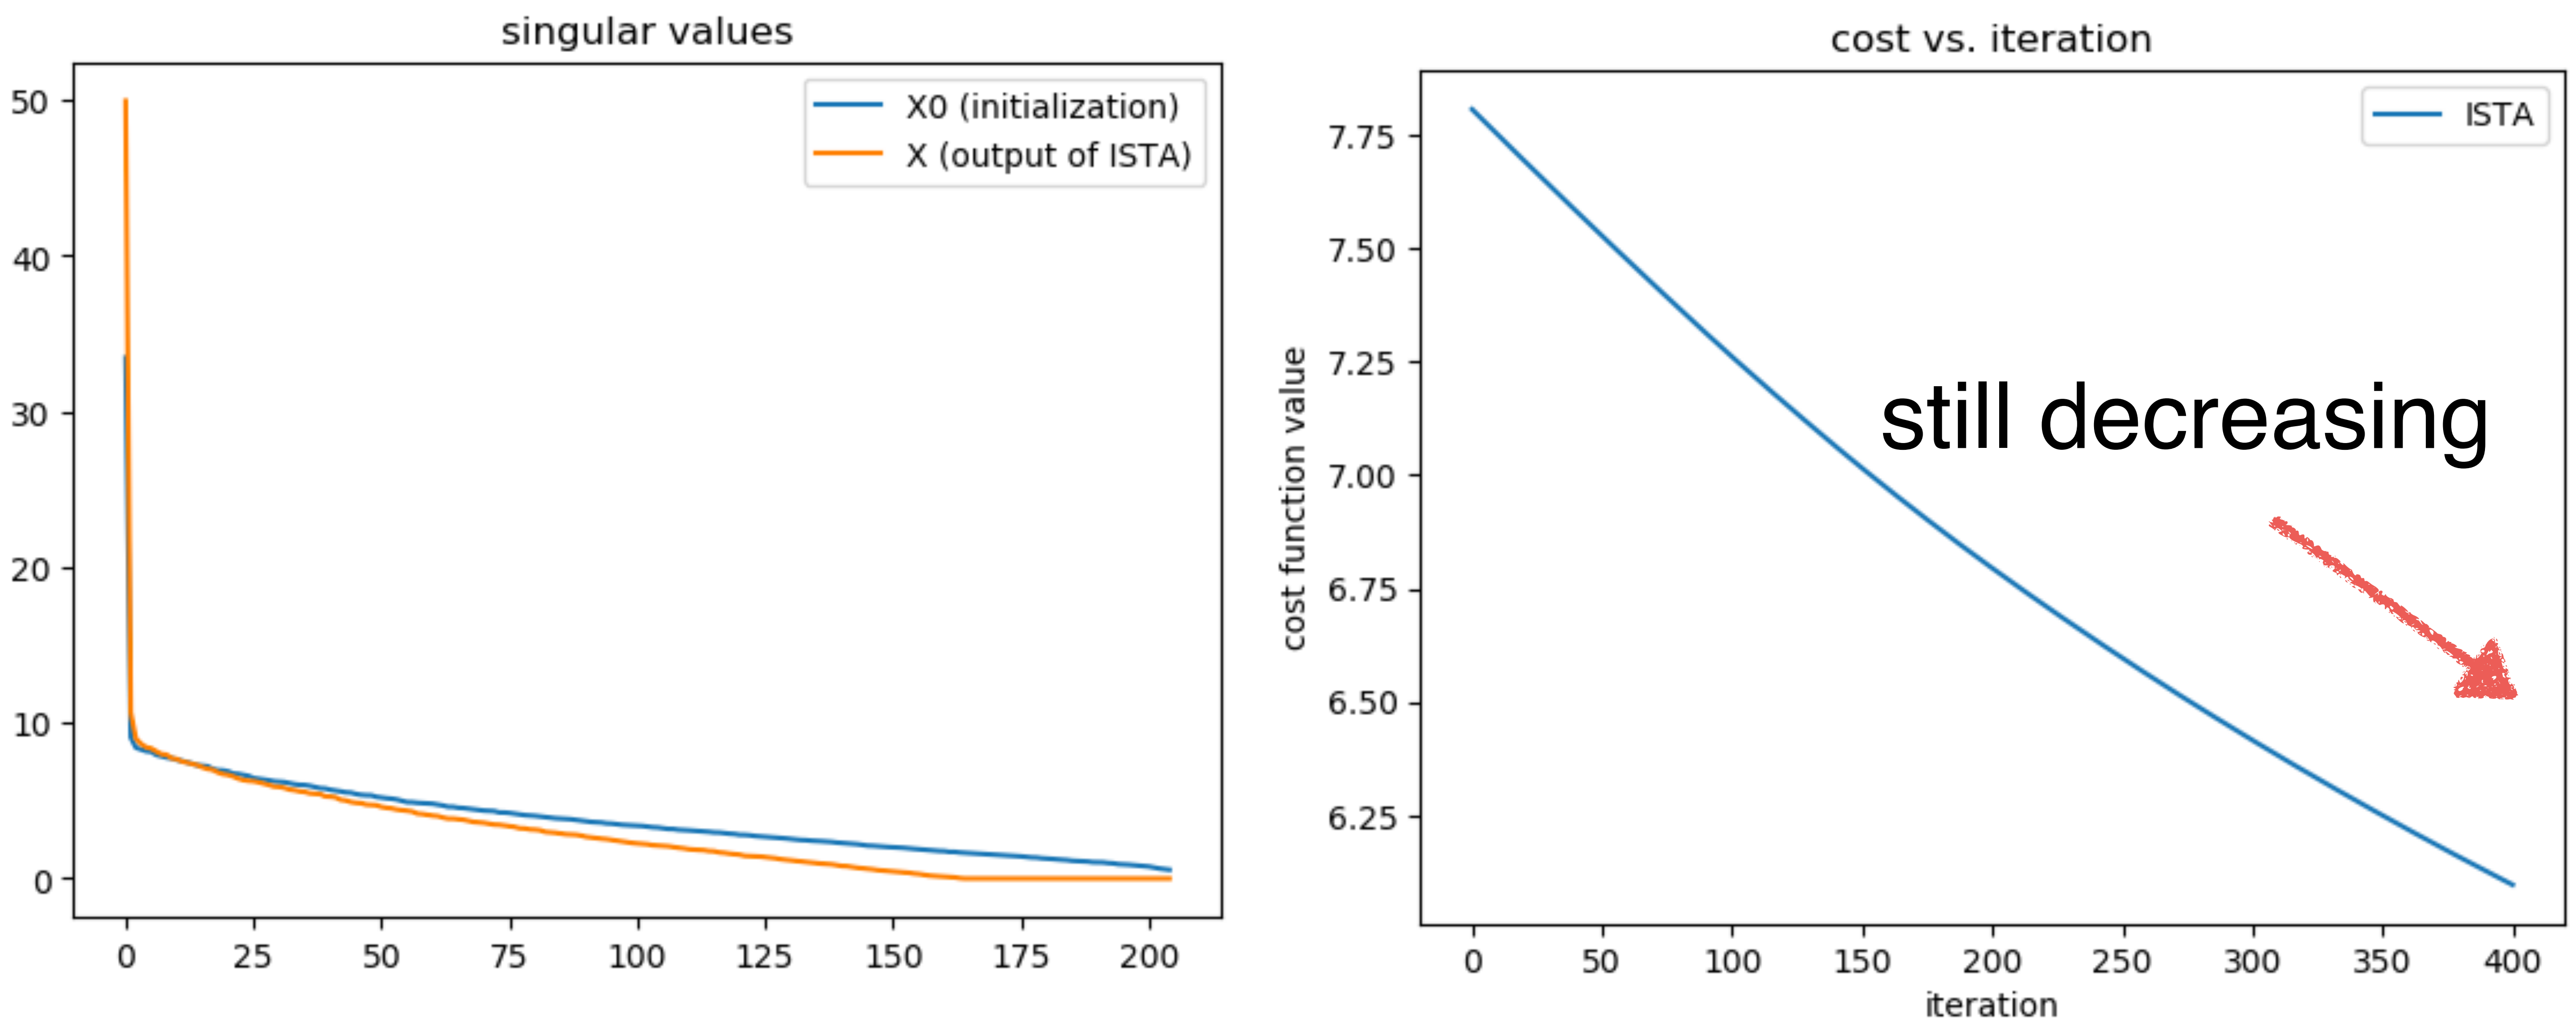

A) Nothing. The algorithm found a low-rank solution.
B) Bug in the implementation of SVT.
C) Bad initialization.
D) Algorithm has not converged—run more iterations.
E) Ran too many iterations.

CSP Seminar:
"Non-Linear Models for Matrix Completion"
Greg Ongie

Date: Thursday, November 9, 2017
Time: 4:00pm
Location: 1005 EECS

Can we complete a partially observed matrix ***X*** assuming its columns lie on an algebraic variety *V*?

5	9		4	5			6	10		1	6	6	7		4	10		9	8
	2		16	7		4	7	5			8	12		7	12	5	9		
13		5	2		4	8			11	6	12		2	10		7		13	10
	16		8	13			10	4	8	2		3		11	10		7	5	4
1		12		8	6	6		5	9		10		5	2		7	6		
5	8	13		1	9	8	2		7	12		7	12		10			4	9