

Recommender Systems for Fun and Profit

Chris Volinsky
Director, Statistics Research
AT&T Research
Florham Park, NJ

Center for Data Analysis and Statistics
April 17, 2008



Recommender Systems

The screenshot shows the Amazon.com homepage for user Chris Volinsky. At the top, the Amazon logo is on the left, and a personalized greeting reads: "Hello, Chris Volinsky. We have [recommendations](#) for you. (Not [Chris?](#))". Below this is a navigation bar with "Chris's Amazon.com", "Today's Deals", "Gifts & Wish Lists", and "Gift Cards". A search bar contains "Amazon.com". A secondary navigation bar includes "Shop All Departments", "Chris's Amazon.com", "Your Browsing History", "Recommended For You", and "Rate These Items".

The main content area features a section titled "Today's Recommendations For You". Below the title, it says: "Here's a daily sample of items recommended for you. Click here to [see all recommendations](#)".

Three items are displayed in a carousel:

- Tribute (Audio CD)**: An audio CD cover with the word "Tribute" in large, stylized green letters on a green background.
- Java Concurrency in Practice (Paperback)**: A book cover by Brian Goetz, featuring the title "JAVA CONCURRENCY IN PRACTICE" and an image of a white boat on water.
- Car Talk (Audio CD)**: An audio CD cover featuring a man playing a guitar, with the text "CAR TALK BORNEOT DRUN" and a red "New Release" badge.

Below the items is a category filter bar with the following options: "Action & Adventure", "All Categories", "Applied", "Book Characters", "Classics", "Classics", "ECM Jazz & World", "Electronica", "Fiction", "Humorous", "Indie Rock", "Jazz", "Literature", "Rock", "Shakespeare", "Shakespeare, William", "Singer-Songwriters", "Software Development", and "Sports".

Recommender Systems

- Two camps on how to do it:
 - Collaborative Filtering
 - based solely on past user preferences
 - Ashley tends to like movies that Emily likes
 - Emily likes “Saw 3”
 - Recommend “Saw 3” to Ashley
 - Content based
 - uses preferences based on attributes
 - Ashley likes Will Smith and romantic comedies
 - Recommend “Hitch”
- CF is ‘cleaner’
- Content handles the ‘cold start’ problem

Netflix: Movies You'll Love

http://www.netflix.com/RecommendationsHome?Inkctr=mh2

Infolab Sports AT&TAdmin CTV Drudge MyYahoo POST weather Gmail KEXP Gcal CTVisME

Chris Volinsky | Your Account | Buy / Redeem

NETFLIX

Browse DVDs Browse Instant Your Queue Movies You'll Friends & Community DVD Sale \$5.99

Movies, actors, directors, genres

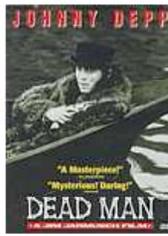
Suggestions (663) Suggestions by Genre Rate Movies Rate Genres Movies You've Rated (103)

Movies You'll Love

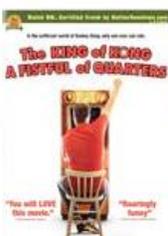
Suggestions based on your ratings

You have 6 Suggestions from 103 ratings

★★★★★ INDEPENDENT SUGGESTIONS (19) [See all 19 >](#)

 <p>Wristcutters: A Love Story</p> <p>Because you enjoyed: Lost in Translation Garden State Children of Men</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>	 <p>Dead Man</p> <p>Because you enjoyed: Taxi Driver Being John Malkovich Harold and Maude</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>	 <p>Trainspotting: Collector's Edition</p> <p>Because you enjoyed: Pulp Fiction Reservoir Dogs Taxi Driver</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>	 <p>Stranger Than Paradise</p> <p>Because you enjoyed: Annie Hall This Is Spinal Tap Taxi Driver</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>
---	--	--	--

★★★★☆ DOCUMENTARY SUGGESTIONS (107) [See all 107 >](#)

 <p>The King of Kong</p> <p>Because you enjoyed: This Is Spinal Tap Spellbound Children of Men</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>	 <p>The Business of Being Born</p> <p>Because you enjoyed: Life Is Beautiful Spellbound Super Size Me</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>	 <p>Jimmy Carter: Man from Plains</p> <p>Because you enjoyed: Annie Hall Being John Malkovich Lost in Translation</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>	 <p>Lake of Fire</p> <p>Because you enjoyed: Annie Hall Fargo The Graduate</p> <p>Add</p> <p>★★★★★</p> <p><input type="radio"/> Not Interested</p>
---	--	--	---

Find:

Done Now: Sunny, 58° F Wed: 68° F Thu: 74° F



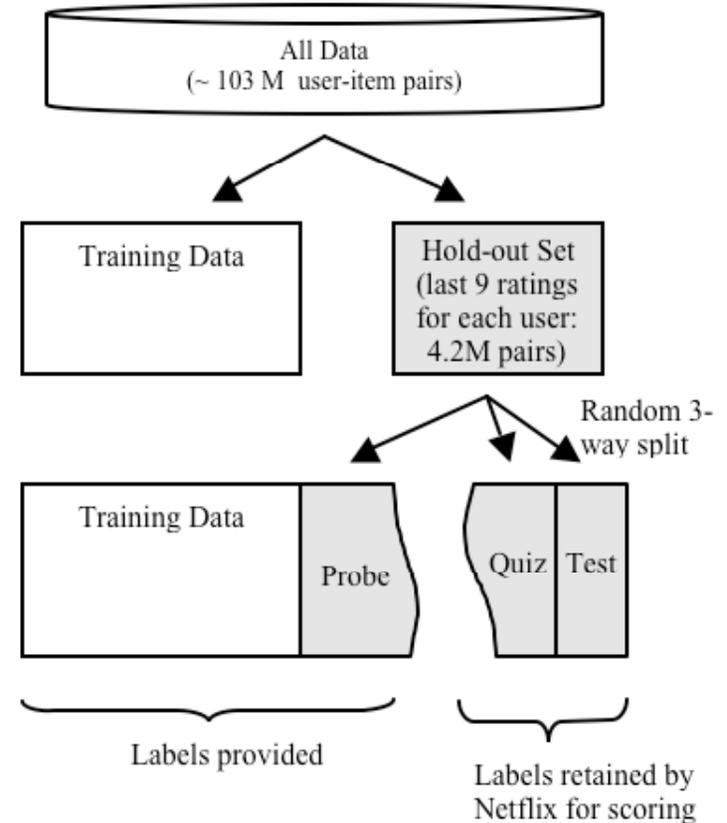
Netflix Prize

- **\$1,000,000** for an improved recommender algorithm
- Training data
 - 100 million ratings
 - 480,000 users
 - 17,770 movies
 - 6 years of data: 2000-2005
- Test data
 - Last few ratings of each user (2.8 million)
 - Evaluation criterion: root mean squared error (RMSE)
 - Netflix Cinematch RMSE: 0.9514
 - results submitted by email
- Competition
 - 3000 teams
 - \$1 million grand prize for 10% improvement on Cinematch result
 - If 10% not met, \$50,000 annual “Progress Prize” for best improvement

user	movie	score	date
1	21	?	2002-01-03
1	123	?	2002-04-04
2	125	?	2002-05-05
2	873	?	2002-05-05
2	8	?	2003-05-03
3	116	?	2003-10-10
4	245	?	2004-10-11
5	232	?	2004-10-11
5	998	?	2004-10-11
5	1012	?	2004-12-12
6	64	?	2005-01-02
6	156	?	2005-01-31

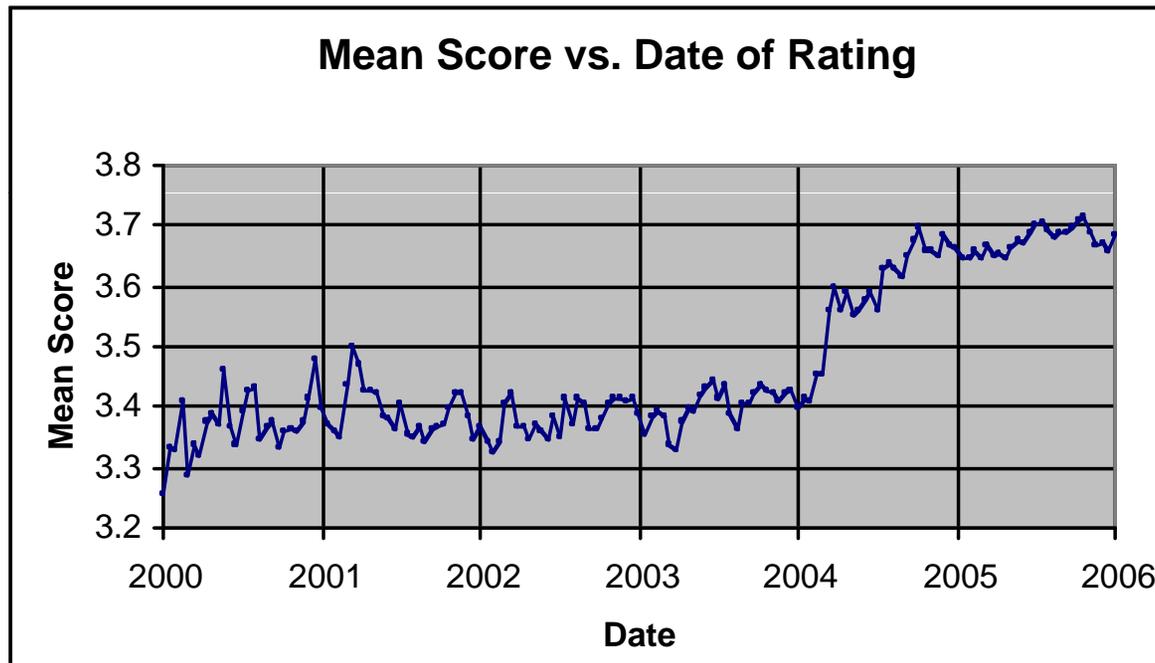
Netflix Prize

- Competition structure
 - results posted on public leaderboard
 - but you could be anonymous if you want!
 - active forums monitored by Netflix
 - prizemaster Jim Bennett a regular poster
 - KDD 2006 Workshop on Netflix prize
 - KDD 2006 cup on a related task
 - winner keeps IP but must announce methodology to the public
- Inspired design!



Data

- date effect



Data

- Most Loved Movies

Title	Avg rating	Count
The Shawshank Redemption	4.593	137812
Lord of the Rings: The Return of the King	4.545	133597
The Green Mile	4.306	180883
Lord of the Rings: The Two Towers	4.460	150676
Finding Nemo	4.415	139050
Raiders of the Lost Ark	4.504	117456

Most rated movies

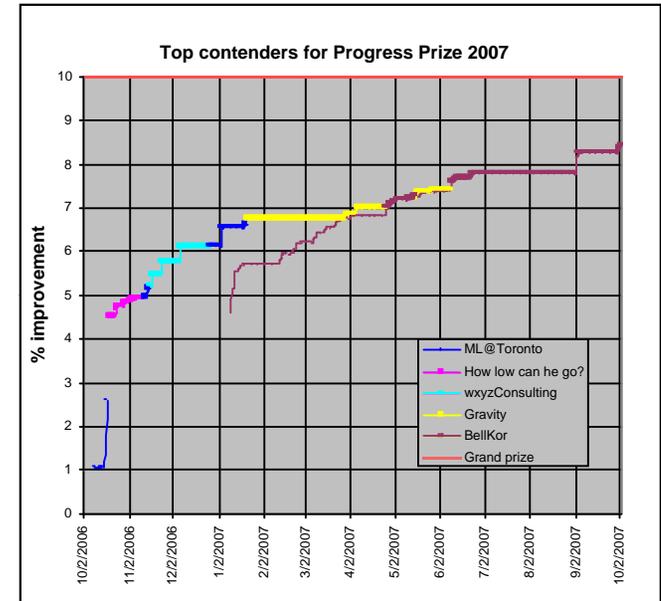
Title
Miss Congeniality
Independence Day
The Patriot
The Day After Tomorrow
Pretty Woman
Pirates of the Caribbean

Highest variance

Title
The Royal Tenenbaums
Lost In Translation
Pearl Harbor
Miss Congeniality
Napolean Dynamite
Fahrenheit 9/11

As the Netflix Prize Turns...

- Cinematch beaten in two weeks
- Halfway to 10% in 6 weeks
- Impact of Simon Funk
- With two days to go, we had a comfortable lead...and hadn't submitted our best yet



Leaderboard

05:00 pm
Sept 30

Rank	Team Name	Best Score	% Improvement
..	No Grand Prize candidates yet
Grand Prize - RMSE <= 0.8563			
1	BellKor	0.8728	8.26
2	Gravity	0.8750	8.03
3	Dinosaur Planet	0.8753	8.00
4	ML@UTorontoA	0.8787	7.64
5	Arek Paterek	0.8789	7.62
6	bashe	0.8805	7.45
7	NIPS Reject	0.8808	7.42
8	Ensemble Experts	0.8841	7.07
Progress Prize 2007 - RMSE: 0.9419			



Leaderboard

06:00 pm
Sept 30

Rank	Team Name	Best Score	% Improvement
--	No Grand Prize candidates yet	--	--
Grand Prize - RMSE <= 0.8563			
1	BellKer	0.8728	8.26
2	bashe	0.8746	8.07
3	Gravity	0.8750	8.03
4	Dinosaur Planet	0.8753	8.00
5	ML@UTeronto A	0.8787	7.64
6	Arek Paterek	0.8789	7.62
7	NIPS Reject	0.8808	7.42
8	Ensemble Experts	0.8841	7.07

Progress Prize 2007 - RMSE: 0.9419

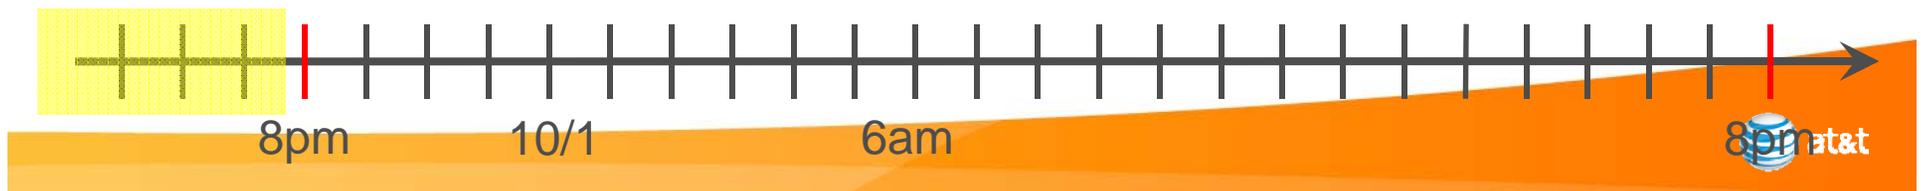
wanna split 50/50?

- so now we submit our best solution, with 25 hours to go...

Rank	Team Name	Best Score	% Improvement	Last Submit Time
--	No Grand Prize candidates yet	--	--	--
Grand Prize: -- RMSE <= 0.8563				
1	When Gravity and Dinosaurs Unite	0.8717	8.38	2007-09-30 23:39:36
2	BellKor	0.8717	8.38	2007-09-30 23:40:52
3	bashe	0.8746	8.07	2007-09-30 21:48:09
4	Gravity	0.8750	8.03	2007-09-30 23:14:39
5	Dinosaur Planet	0.8753	8.00	2007-09-30 23:17:14
6	ML@UTorontoA	0.8787	7.64	2007-09-30 20:41:54
7	Arek Paterek	0.8789	7.62	2007-09-30 11:35:42
8	NIPS Reject	0.8808	7.42	2007-09-13 21:02:32
9	Ensemble Experts	0.8841	7.07	2007-09-30 04:37:18

Progress Prize 2007 - RMSE: 0.9419

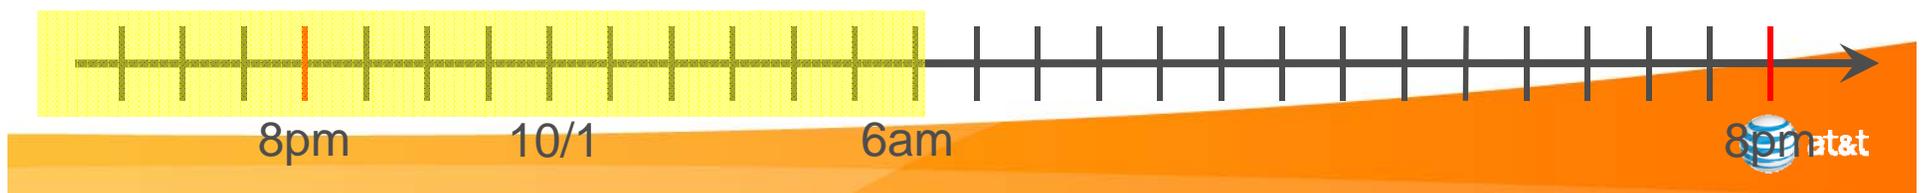
- ARRRRGH! We have one more chance....



Offers come pouring in....

Rank	Team Name	Best Score	% Improvement	Last Submit Time
--	No Grand Prize candidates yet	--	--	--
Grand Prize: -- RMSE: <= 0.8563				
1	When Gravity and Dinosaurs Unite	0.8717	8.38	2007-09-30 23:39:36
2	BellKor	0.8717	8.38	2007-09-30 23:40:52
3	bashe	0.8746	8.07	2007-09-30 21:48:09
4	Gravity	0.8750	8.03	2007-09-30 23:14:39
5	Dinosaur Planet	0.8753	8.00	2007-09-30 23:17:14
6	ML@UTorontoA	0.8787	7.64	2007-09-30 20:41:54
7	Arak Paterek	0.8789	7.62	2007-09-30 11:35:42
8	NIPS Reject	0.8808	7.42	2007-09-13 21:02:32
9	Ensemble Experts	0.8841	7.07	2007-09-30 04:37:18

Progress Prize: 2007 - RMSE: 0.9419



Our final submission(s)...

Rank	Team Name	Best Score	% Improvement
--	No Grand Prize candidates yet	--	--
Grand Prize - RMSE <= 0.8563			
1	KerBell	0.8712	8.43
2	BellKer	0.8714	8.41
3	When Gravity and Dinosaurs Unite	0.8717	8.38
4	bashe	0.8746	8.07
5	Gravity	0.8750	8.03
6	Dinosaur Planet	0.8753	8.00
7	ML@UTorontoA	0.8787	7.64
8	Arek Paterek	0.8789	7.62
9	NIPS Reject	0.8808	7.42
10	Ensemble Experts	0.8841	7.07
Progress Prize 2007 - RMSE: 0.9419			





Who is currently in the lead?

NETFLIX

Netflix Prize

Home Rules Leaderboard Register Update Submit Download

Leaderboard

Display top leaders.

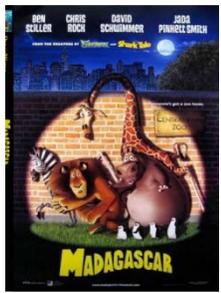
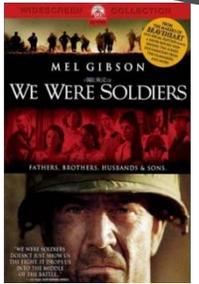
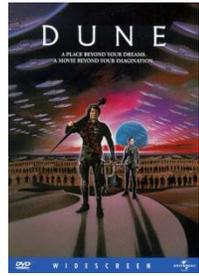
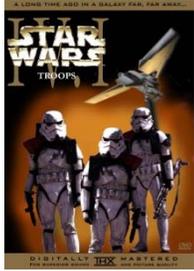
Rank	Team Name	Best Score	% Improvement	Last Submit Time
--	No Grand Prize candidates yet	--	--	--
Grand Prize - RMSE <= 0.8563				
--	No Progress Prize candidates yet	--	--	--
Progress Prize - RMSE <= 0.8625				
1	BellKor	0.8650	9.08	2008-04-01 17:02:59
2	When Gravity and Dinosaurs Unite	0.8675	8.82	2008-04-06 19:34:42
3	BigChaos	0.8677	8.80	2008-04-15 10:59:00
4	Gravity	0.8693	8.63	2008-04-09 10:18:08
Progress Prize 2007 - RMSE = 0.8712 - Winning Team: KorBell				
5	KorBell	0.8712	8.43	2007-10-01 23:25:23
6	acmehill	0.8718	8.37	2008-03-26 02:07:17
7	basho	0.8722	8.32	2008-03-03 06:18:08
8	Dan Tillberg	0.8723	8.31	2008-03-28 20:58:05
9	Just a guy in a garage	0.8724	8.30	2008-04-08 21:22:16
10	Reel Ingenuity	0.8747	8.06	2008-04-02 21:59:23
11	Dinosaur Planet	0.8753	8.00	2007-10-04 04:56:45
12	Ces	0.8759	7.94	2008-04-01 06:57:28
13	Three Blind Mice	0.8778	7.74	2008-02-16 20:47:39
14	Efratko	0.8778	7.74	2008-03-31 02:02:10
15	ML@UToronto A	0.8787	7.64	2007-09-30 20:41:54
16	Arek Paterek	0.8789	7.62	2007-09-30 11:35:42

Find: Next Previous Highlight all Match case Phrase not found

Done Adblock Now: Sunny, 58° F Wed: 68° F Thu: 74° F

BellKor solution...

- Solution is a mixture of 107 (!) different models
 - mostly variants of Nearest Neighbors and Latent Factor models.
 - helps to look at the data from different perspectives
 - lots of parameters with heavy regularization

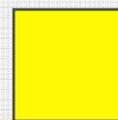


Nearest Neighbors

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



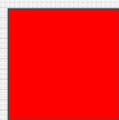
- unknown rating



- rating between 1 to 5

Nearest Neighbors

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



- estimate rating of movie 1 by user 5

Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		?	5			5		4	
	2			5	4			4			2	1	3
	<u>3</u>	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	<u>6</u>	1		3		3			2			4	

Neighbor selection:
Identify movies similar to 1, rated by user 5

Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		?	5			5		4	
	2			5	4			4			2	1	3
	<u>3</u>	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	<u>6</u>	1		3		3			2			4	

Compute similarity weights:

$$s_{13}=0.2, s_{16}=0.3$$

Nearest Neighbors

		users											
		1	2	3	4	5	6	7	8	9	10	11	12
movies	1	1		3		2.6	5			5		4	
	2			5	4			4			2	1	3
	<u>3</u>	2	4		1	2		3		4	3	5	
	4		2	4		5			4			2	
	5			4	3	4	2					2	5
	<u>6</u>	1		3		3			2			4	

Predict by taking weighted average:

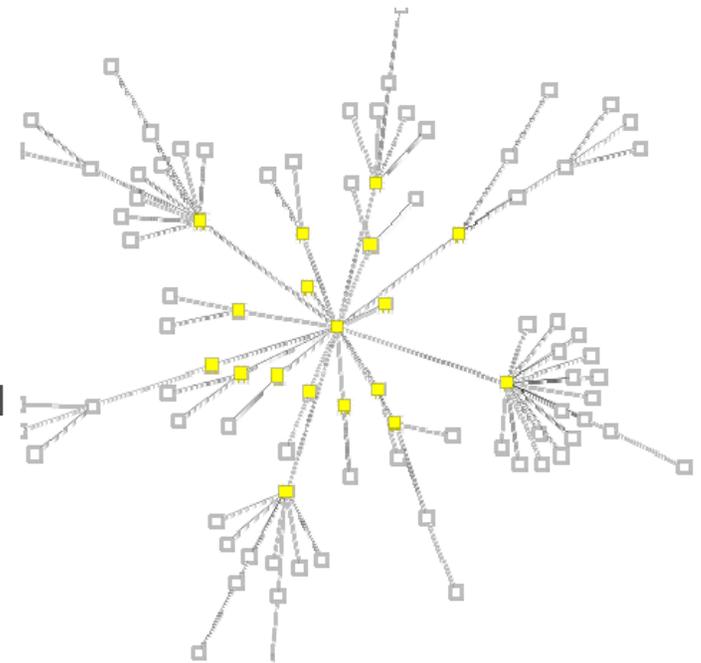
$$(0.2*2+0.3*3)/(0.2+0.3)=2.6$$

Properties of NN

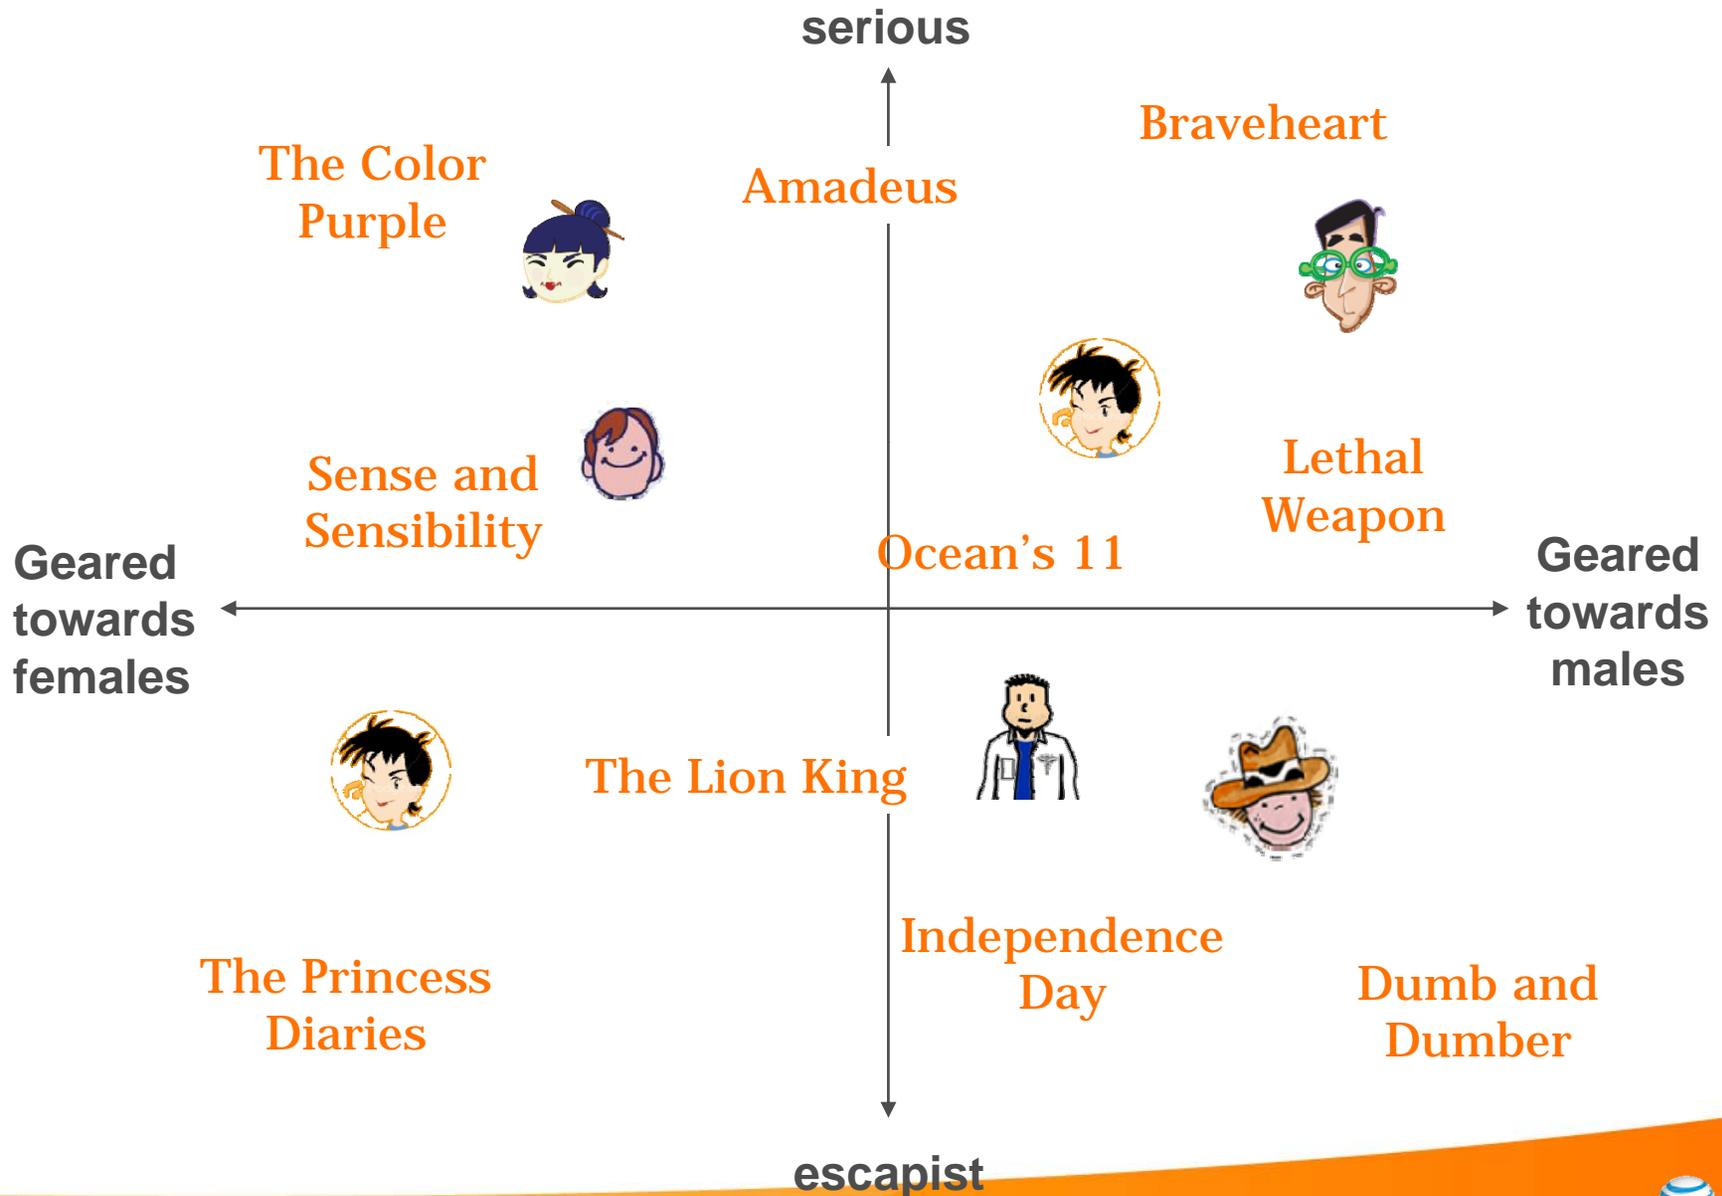
more specifically:

$$\hat{r}_{ui} = b_{ui} + \frac{\sum_{j \in N(i,u)} s_{ij}(r_{uj} - b_{uj})}{\sum_{j \in N(i,u)} s_{ij}}$$

- Intuitive
- No substantial preprocessing is required
- Easy to explain reasoning behind a recommendation
 - every user and every movie has a neighborhood
- Accurate



Latent factor models



Latent Factors - SVD

users

	1		3			5			5			4	
items			5	4	?		4			2	1	3	
	2	4		1	2		3		4	3	5		
		2	4		5			4			2		
			4	3	4	2					2	5	
	1		3		3			2			4		

~

users

	.1	-.4	.2
items	-.5	.6	.5
	-.2	.3	.5
	1.1	2.1	.3
	-.7	2.1	-2
	-1	.7	.3

•

	1.1	-.2	.3	.5	-2	-.5	.8	-.4	.3	1.4	2.4	-.9
	-.8	.7	.5	1.4	.3	-1	1.4	2.9	-.7	1.2	-.1	1.3
	2.1	-.4	.6	1.7	2.4	.9	-.3	.4	.8	.7	-.6	.1

A rank-3 SVD approximation

Factorization-based modeling

1		3			5			5		4	
		5	4			4			2	1	3
2	4		1	2		3		4	3	5	
	2	4		5			4			2	
		4	3	4	2					2	5
1		3		3			2			4	

.1	-.4	.2
-.5	.6	.5
-.2	.3	.5
1.1	2.1	.3
-.7	2.1	-2
-1	.7	.3

1.1	-.2	.3	.5	-2	-.5	.8	-.4	.3	1.4	2.4	-.9
-.8	.7	.5	1.4	.3	-1	1.4	2.9	-.7	1.2	-.1	1.3
2.1	-.4	.6	1.7	2.4	.9	-.3	.4	.8	.7	-.6	.1

$$\sum_{u,i} (r_{ui} - p_u^T q_i)^2 + \lambda(\|p_u\|^2 + \|q_i\|^2)$$

:

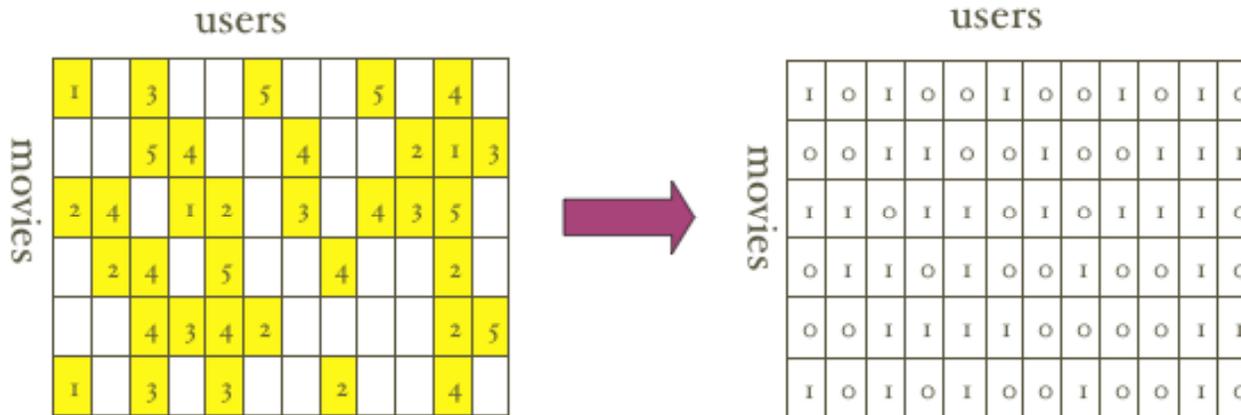
- can easily overfit, sensitive to regularization
- Probably most popular model among contestants
 - 12/11/2006: Simon Funk describes an SVD based method

Other methods

- Restricted Boltzmann Machines - UToronto
 - stochastic artificial neural network
- GOLR!
 - with lots of shrinkage

Combining Models

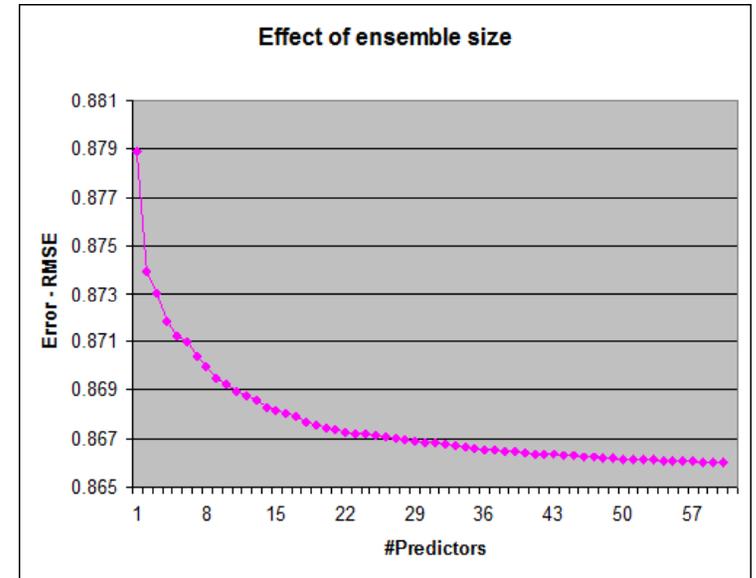
- Nearest neighbors captures “low level” effects
 - Latent factors capture overall structure
 - they are good to combine!
 - plus hybrids of these
-
- Also, look at data from another perspective:



Lessons learned

- It took 107 models to win, but:
- combination of models with different perspectives is powerful. Think different!
- regularization is extremely important to prevent overfitting
 - varying support for certain parameters
- content does not seem to help!

- power of open competition and collaboration
 - contestants aren't motivated by \$\$!



So why does AT&T let me work on this???

- Collaborative filtering can be used in any domain where data is collected on user preferences or purchases:
 - Web browsing
 - Ad placement (online or otherwise)
 - Market basket
 - customer care
 - ????? Lets get creative!
- Also, we have applied this to AT&T's U-Verse (digital cable) data to build a TV recommender system

MyVerse: TV Recommendations

- For AT&T's U-Verse customers, we know
 - What they watched
 - When they change channels
 - What they record
 - What movies they order (on demand)

- We can use this data for
 - Recommendations
 - Customer understanding
 - 3-screen advertising
 - ???



Uverse Mining Tool

[Overview](#) [Channels](#) [Users/STBs](#) [Programs](#)

[Viewing Profile](#) [Top Programs](#) [Recommendations: SVD kNN](#)

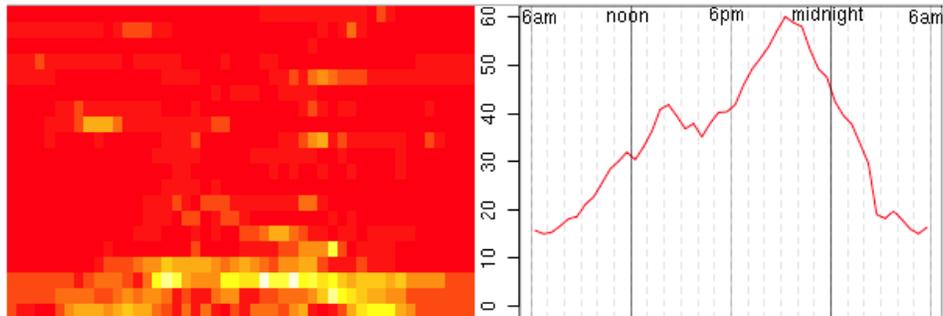
STB UUID: `jecc324d-c961-4400-933a-ffbef462ca68` [GO](#) ([Show test STB list](#))

Use only period from to [SET](#) [clear](#)

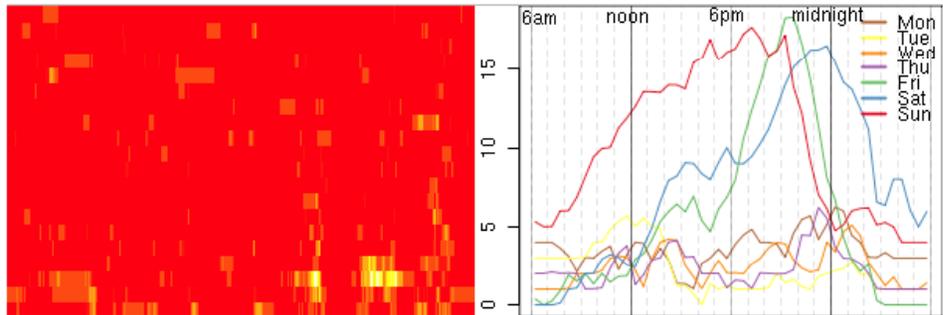
Usage: 2007-08-28 23:02:25 CDT..2008-02-26 17:30:29 CST(181 days)

Watching Profile

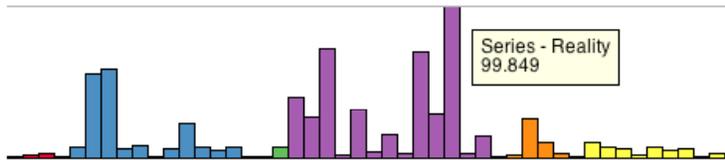
Time of Day



Day of Week



Genre



Uverse Mining Tool

[Overview](#) [Channels](#) [Users/STBs](#) [Programs](#)

[Viewing Profile](#) [Top Programs](#) [Recommendations: SVD](#) [kNN](#)

STB UUID: [GO](#) ([Show test STB list](#))

Use only period from to [SET](#) [clear](#)

Note: Recommendations are based on the period Jan 20-Feb 3 2008 ([set](#)), testing the week after

include previously watched shows in recommendations

Top programs

America's Next Top Model	6.689
Pros vs. Joes	4.677
The Wayans Bros.	4.004
George Lopez	2.590
MTV Jams	2.331
Snoop Dogg's Father Hood	1.979
Off Air	1.947
icons.	1.596
48 Hours Mystery	1.565
Fresh Prince of Bel-Air	1.437

Recommendations

America's Next Top Model	0.062	(watched)
George Lopez	0.049	(watched)
The Wayans Bros.	0.049	(watched)
Pros vs. Joes	0.049	(watched)
Fresh Prince of Bel-Air	0.048	(watched)
Snoop Dogg's Father Hood	0.043	(watched)
MTV Jams	0.040	(watched)
Off Air	0.033	(watched)
The Fairly OddParents	0.033	(watched)
Scott Baio is 46 ... and Pregnant	0.032	(watched)

Uverse Mining Tool

[Overview](#) [Channels](#) [Users/STBs](#) [Programs](#)

[Viewing Profile](#) [Top Programs](#) [Recommendations: SVD kNN](#)

STB UUID: [GO](#) ([Show test STB list](#))

Use only period from to [SET](#) [clear](#)

Note: Recommendations are based on the period Jan 20-Feb 3 2008 ([set](#)), testing the week after

include previously watched shows in recommendations

Top programs

America's Next Top Model	6.689
Pros vs. Joes	4.677
The Wayans Bros.	4.004
George Lopez	2.590
MTV Jams	2.331
Snoop Dogg's Father Hood	1.979
Off Air	1.947
icons.	1.596
48 Hours Mystery	1.565
Fresh Prince of Bel-Air	1.437

Recommendations

SpongeBob SquarePants	0.032
Drake & Josh	0.025
NFL Football	0.021
Rock of Love	0.020
iCarly	0.020
The Adventures of Jimmy Neutron: Boy Genius	0.020
Hitlist	0.018
Hannah Montana	0.018
The Suite Life of Zack & Cody	0.017
SportsCenter	0.017

Using Text Mining

Show blurbs can help solve the so-called “cold start” problem:

292984080|American Idol|American Idol|**With the semi-finalists narrowed down to 20, the 10 men perform for judges Paula, Randy and Simon before the viewers across the country call in their votes to decide who will continue to compete for the Idol crown and who will go home.**|The 10 male semi-finalists perform for judges Paula, Randy and Simon before viewers across the country vote to decide who stays and who goes.|en-US|Series|Series:Music,Series:Reality|292959|718|60|||02/26/2008||TV-G||||

amateur, annual, answer, award, casino,

challeng, challenge, chanc, choos, chosen,

compet, competit, contest, danc, dancer,

elimin, enter, **final, game, gameshow,**

group, **head, hope,** judg, member,

open, opportun, order, particip, **plai,**

poker, prize, question, remain, round,

show, stake, talent, **team, test,**

titt, top, tournament, try, **vega,**

vegas, vote, **Win,** winner, won,

access, analysi, athlet, basebal, baseball,

basketbal, bowl, **coach,** colleg, college,

confer, fan, featur, field, **footbal,**

football, game, head, highlight, hockei,

hockey, insid, interview, latest, leagu,

look, match, nfl, **parentprogram,** plai,

player, post, press, preview, pro,

profession, profil, recap, replai, review,

score, season, **sport, sports,** spotlight,

team, univers, up, upcom, weekli,

caus, **china,** chines, clock, describ,

destruct, devast, **disast,** effect, effort,

emerg, erupt, explos, **fire,** flood,

forc, frozen, **global,** hand, huge,

hurrican, **ic,** impact, major, matter,

mega, **minut,** natur, neil, occur,

prevent, produc, **reach, rush, send,**

storm, strike, strong, **surviv,** survivor,

time, tornado, tower, volcano, warm,

wast, **water,** wave, weather, wind,

Users will have a distribution among the clusters, as will shows

evil, forc, **power,** save, super,

team, turn, **friend, help, parti,**

plan, take, tri, want, dog,

friend, learn, plai, thing, walk,

kid, kids, **seri, series,** boi,

children, fantasi, fantasy, magic, young,

decid, **find,** help, learn, run,

start, take, tell, thing, think,

tri, troubl, try, turn, visit,

want, girl, **high, school,** student,

teen,

investig, killer, murder, david, mark,

series, with, accus, commit, crime,

jackson, kill, law, **murder,** million,

monei, owner, place, plan, rich,

boi, child, father, **mother,** date,

friend, help, parti, plan, take,

tri, want, wed, find, tri,

try, turn, want, daughter, husband,

wife, **woman,** at, fox, local,

new,news,**parentprogram,**weather,

In this way, we can map users to new shows that no one has watched

Thank you!

