Recommendation systems¹

Jake Hofman

Yahoo! Research

April 4, 2011

¹For large data sets

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 1 / 31

▲ロト ▲掃 ト ▲ 臣 ト ▲ 臣 ト 一 臣 - つへぐ

Personalized recommendations



@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 2 / 31

▲□▶ ▲圖▶ ▲ 臣▶ ▲ 臣▶ ― 臣 … のへで

Personalized recommendations



S Not Interested Our best guess for Jake: 5 stars Average of 4,2/5,920 ratings: 3.8 stars

The Big Lebowski

1998 R 117 minutes

Stacker Jeff "The Dude" Lebowski (Jeff Bridges) gets involved in a gargantuan mess of events when he's mistaker for another man named Lebowski, whose wife has been kidnapped and is being held for ransom. All the while, Dude's friend, Walter (Jack Bridger and State State) and an antiter and direct this cuit cornedy classic that also stars Steve Buscemi, Philip Seymon Hoffman, Julianen Movies and Jack Bridger and State Stars Steve Buscemi, Philip Seymon Hoffman, Julianen Movies and Jack Brance Steve Buscemi, Philip Seymon Hoffman, Julianen Movies and Jack Brance Brance Steve Buscemi, Philip Seymon Hoffman (Steve Buscemi, Philip Seymon Hoffman) and the Steve Buscemi, Philip Seymon Hoffman (Steve Buscemi, Philip Seymon (Steve Buscemi, Philip Seymon Hoffman) and Steve Buscemi, Philip Seymon (Steve Buscemi, Steve Buscemi, Philip Seymon (Steve Buscemi, Steve Bu

Cast: Jeff Bridges, John Goodman, Philip Seymour Hoffman, Stave Busceni, Juliane Moore, Tara Reid, Peter Stormare, David Huddleston, Philip Moon, Mark Pellegrino, Flea, Torsten Voges, Jimmie Dale Gilmore, Jack Kehler, John Turturno, James G. Hoosler, Richard Gant, Christian Clemenson, David Thewils, Peter Siragusa, Sam Elliott, Ben Gazzara, Jon Polito, Asia Carrera, Paris Themmen

Director: Joel Coen

- Genres: Comedy, Cult Comedies, Universal Studios Home Entertainment, Blu-ray
- This movie is: Quirky, Witty
 - Format: DVD and streaming (Blu-ray availability date unknown) (HD available)



Recommended based on your interest in: Fargo, O Brother, Where Art Thou? and No Country for Old Men

▲□▶ ▲圖▶ ▲ 臣▶ ▲ 臣▶ ― 臣 … のへで

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 3 / 31

http://netflixprize.com



@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 4 / 31

http://netflixprize.com/rules

Terms and Conditions in a Nutshell

- Contest begins October 2, 2006 and continues through at least October 2, 2011.
- · Contest is open to anyone, anywhere (except certain countries listed below).
- · You have to register to enter.
- Once you register and agree to these Rules, you'll have access to the Contest training data and qualifying test sets.
- To qualify for the \$1,000,000 Grand Prize, the accuracy of your submitted predictions on the qualifying set must be at least 10% better than the accuracy Cinematch can achieve on the same training data set at the start of the Contest.
- To qualify for a year's \$50,000 Progress Prize the accuracy of any of your submitted predictions that year must be less than or equal to the accuracy value established by the judges the preceding year.
- To win and take home either prize, your qualifying submissions must have the largest accuracy improvement verified by the Contest judges, you must share your method with (and non-exclusively license it to) Netflix, and you must describe to the world how you did it and why it works.

http://netflixprize.com/faq



How does Cinematch do it?

Straightforward statistical linear models with a lot of data conditioning. But a real-world system is much more than an algorithm, and Cinemath does a lot more than just optimize for RMSE. After all, we have a website to support. In production we have to worry about system scaling and performance, and we have additional sources to data we can use to guide our recommendations. But, as mentioned in the <u>Rules</u> and just to be perfectly clear, for the purposes of the Prize the RMSE values we report here do not use any of this extra data.

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 6 / 31

イロト イポト イヨト イヨト

Netflix prize: results

The final standing of the Leaderboard at that time showed that two teams met the minimum requirements for the Grand Prize. "The Ensemble" with 10.10% mprovement over Cinematch on the Qualifying set (an RMSE of 0.8553), and "BellKor's Pragmatic Chaos" with 10.09% mprovement over Cinematch on the Qualifying set (an RMSE of 0.8554).^[21] The Grand Prize winner was to be the one with the better performance on the Test set.

On September 18, 2009, Netflix announced team "BellKor's Pragmatic Chaos" as the prize winner, and the prize was awarded to the team in a ceremony on September 21, 2009.^[22] "The Ensemble" team had in fact succeeded to match BellKor's result, but since BellKor submitted their results 20 minutes earlier; the rules award the prize to them.^[23]

http://en.wikipedia.org/wiki/Netflix_Prize

@jakehofman (Yahoo! Research)

オロト 本理 ト オヨト オヨト ニヨー ろんの

Netflix prize: results

Add an asymmetric frequency feature $\mathbf{y}_{j,f_{ut}}^{(3)}$: **SBRAMF-UTB-UTF-MTF-ATF-MFF-AFF**

$$\widehat{r_{uit}} = \mu_i + \mu_u + \mu_{u,t} + \mu_{i,\text{bin}(t)} + \left(\mathbf{p}_i^{(1)} + \mathbf{p}_{i,\text{bin}(t)}^{(2)} + \mathbf{p}_{i,f_{ut}}^{(3)}\right)^T \left(\mathbf{q}_u^{(1)} + \mathbf{q}_{u,t}^{(2)} + \frac{1}{\sqrt{|N(u)|}} \sum_{j \in N(u)} \left(\mathbf{y}_j^{(1)} + \mathbf{y}_{j,\text{bin}(t)}^{(2)} + \mathbf{y}_{j,f_{ut}}^{(3)}\right)\right)$$
(34)

Model extension (+)	epoch time	#epochs	probeRMSE, $k = 50$ features
SBRMF - SVD with biases	17[s]	69	0.9054
SBRAMF - asymmetric part	50[s]	30	0.8974
+UTB - user time bias	61[s]	50	0.8919
+UTF - user time feature	62[s]	38	0.8911
+MTF - movie time feature	74[s]	37	0.8908
+ATF - asymmetric time feature	74[s]	44	0.8905
+MFF - movie frequency feature	149[s]	46	0.8900
+AFF - asymmetric frequency feature	206[s]	45	$0.8886 \ (0.8846 \text{ with } k = 1000)$

See [TJB09] and [Kor09] for more gory details.

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 8 / 31

◆□▶ ◆□▶ ◆三▶ ◆三▶ ◆□ ◆ ◇◇◇

Recommendation systems

High-level approaches:

- Content-based methods (e.g., $w_{\text{genre: thrillers}} = +2.3$, $w_{\text{director: coen brothers}} = +1.7$)
- Collaborative methods (e.g., "Users who liked this also liked")

Netflix prize: data

(userid, movieid, rating, date)

15124	5	9	3	2004
4506	3	9	3	2004
5069	3	9	3	2004
5924	5	9	3	2004
7040	3	9	3	2004
2095	4	9	3	2004
8301	3	9	3	2004
17295	5	9	3	2004
17195	3	9	3	2004
13007	4	9	3	2004
16912	3	9	3	2004
14240	5	9	3	2004
705	3	9	3	2004
	15124 4506 5069 5924 7040 2095 8301 17295 17195 13007 16912 14240 705	$\begin{array}{ccccc} 15124 & 5 \\ 4506 & 3 \\ 5924 & 5 \\ 7040 & 3 \\ 2095 & 4 \\ 8301 & 3 \\ 17295 & 5 \\ 17195 & 3 \\ 13007 & 4 \\ 16912 & 3 \\ 14240 & 5 \\ 705 & 3 \end{array}$	$\begin{array}{cccccccc} 15124 & 5 & 9 \\ 4506 & 3 & 9 \\ 5069 & 3 & 9 \\ 5924 & 5 & 9 \\ 7040 & 3 & 9 \\ 2095 & 4 & 9 \\ 8301 & 3 & 9 \\ 17295 & 5 & 9 \\ 17195 & 3 & 9 \\ 17195 & 3 & 9 \\ 13007 & 4 & 9 \\ 16912 & 3 & 9 \\ 14240 & 5 & 9 \\ 705 & 3 & 9 \end{array}$	$\begin{array}{cccccccccccccccccccccccccccccccccccc$

@jakehofman (Yahoo! Research)

April 4, 2011 10 / 31

▲□▶ ▲圖▶ ▲ 臣▶ ▲ 臣▶ ― 臣 … のへで

Netflix prize: data

(movieid, year, title)

1	60034316	2003	Dinosaur Planet
2	70014695	2004	Isle of Man TT 2004 Review
3	60010225	1997	Character
4	60033124	1994	Paula Abdul's Get Up & Dance
5	70017841	2004	The Rise and Fall of ECW
6	60031092	1997	Sick
7	19826698	1992	8 Man
8	70011191	2004	What the #\$*! Do We Know!?
9	70002378	1991	Class of Nuke 'Em High 2
10	60024682	2001	Fighter

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 10 / 31

Recommendation systems

High-level approaches:

- Content-based methods (e.g., $w_{\text{genre: thrillers}} = +2.3$, $w_{\text{director: coen brothers}} = +1.7$)
- Collaborative methods (e.g., "Users who liked this also liked")

Collaborative filtering

Memory-based (e.g., *k*-nearest neighbors)



Model-based (e.g., matrix factorization)



characterizes both users and movies using two axes and serious versus escapist.

http://research.yahoo.com/pub/2859

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 12 / 31

Problem statement

- Given a set of past ratings R_{ui} that user u gave item i
 - Users may explicitly assign ratings, e.g., $R_{ui} \in [1, 5]$ is number of stars for movie rating
 - Or we may infer implicit ratings from user actions, e.g. $R_{ui} = 1$ if *u* purchased *i*; otherwise $R_{ui} = ?$

Problem statement

- Given a set of past ratings R_{ui} that user u gave item i
 - Users may explicitly assign ratings, e.g., $R_{ui} \in [1, 5]$ is number of stars for movie rating
 - Or we may infer implicit ratings from user actions, e.g. $R_{ui} = 1$ if *u* purchased *i*; otherwise $R_{ui} = ?$
- Make recommendations of several forms
 - Predict unseen item ratings for a particular user
 - Suggest items for a particular user
 - Suggest items similar to a particular item
 - . . .

Problem statement

- Given a set of past ratings R_{ui} that user u gave item i
 - Users may explicitly assign ratings, e.g., $R_{ui} \in [1, 5]$ is number of stars for movie rating
 - Or we may infer implicit ratings from user actions, e.g. $R_{ui} = 1$ if *u* purchased *i*; otherwise $R_{ui} = ?$
- Make recommendations of several forms
 - Predict unseen item ratings for a particular user
 - Suggest items for a particular user
 - Suggest items similar to a particular item
 - ...
- Compare to natural baselines
 - Guess global average for item ratings
 - Suggest globally popular items

@jakehofman (Yahoo! Research)

April 4, 2011 13 / 31

k-nearest neighbors

Key intuition:

Take a local popularity vote amongst "similar" users

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 14 / 31

▲□▶ ▲圖▶ ▲ 臣▶ ▲ 臣▶ ― 臣 … のへで

k-nearest neighbors _{User similarity}

Quantify similarity as a function of users' past ratings, e.g.

• Fraction of items *u* and *v* have in common

$$S_{uv} = \frac{|R_u \cap R_v|}{|R_u \cup R_v|} = \frac{\sum_i R_{ui} R_{vi}}{\sum_i (R_{ui} + R_{vi} - R_{ui} R_{vi})}$$
(1)

Retain top-k most similar neighbors v for each user u

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 15 / 31

k-nearest neighbors _{User similarity}

Quantify similarity as a function of users' past ratings, e.g.

Angle between rating vectors

$$S_{uv} = \frac{R_u \cdot R_v}{|R_u| \, |R_v|} = \frac{\sum_i R_{ui} R_{vi}}{\sqrt{\sum_i R_{ui}^2 \sum_j R_{vj}^2}}$$
(1)

Retain top-k most similar neighbors v for each user u

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 15 / 31

▲□▶ ▲□▶ ▲目▶ ▲目▶ 三日 - つくつ

k-nearest neighbors Predicted ratings

Predict unseen ratings \hat{R}_{ui} as a weighted vote over *u*'s neighbors' ratings for item *i*

$$\hat{R}_{ui} = \frac{\sum_{v} R_{vi} S_{uv}}{\sum_{v} S_{uv}}$$
(2)

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 16 / 31

k-nearest neighbors Practical notes

We expect most users have nothing in common, so calculate similarities as:

for each item *i*: for all pairs of users u, v that have rated *i*: calculate S_{uv} (if not already calculated)

k-nearest neighbors Practical notes

Alternatively, we can make recommendations using an item-based approach [LSY03]:

- Compute similarities S_{ij} between all pairs of items
- Predict ratings with a weighted vote $\hat{R}_{ui} = \sum_{i} R_{uj} S_{ij} / \sum_{i} S_{ij}$

k-nearest neighbors Practical notes

Several (relatively) simple ways to scale:

- Sample a subset of ratings for each user (by, e.g., recency)
- Use MinHash to cluster users [DDGR07]
- Distribute calculations with MapReduce

Key intuition:

Model item attributes as belonging to a set of unobserved "topics and user preferences across these "topics"

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 18 / 31

イロト 不得下 イヨト イヨト 二日

Start with a simple linear model:

 $\hat{R}_{ui} = \underbrace{b_0} + \underbrace{b_u} + \underbrace{b_i}$ global average user bias item bias

Recommendation systems

April 4, 2011 19 / 31

▲ロト ▲圖ト ▲画ト ▲画ト 三直 - のへで

(3)

For example, we might predict that a harsh critic would score a popular movie as



April 4, 2011 19 / 31

Add an interaction term:



where $W_{ui} = P_u \cdot Q_i = \sum_k P_{uk} Q_{ik}$

- P_{uk} is user *u*'s preference for topic *k*
- Q_{ik} is item *i*'s association with topic k

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 20 / 31

Measure quality of model fit with squared-loss:

$$\mathcal{L} = \sum_{(u,i)} \left(\hat{R}_{ui} - R_{ui} \right)^2$$
(6)
=
$$\sum_{(u,i)} \left([PQ]_{ui} - R_{ui} \right)^2$$
(7)

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 21 / 31

The loss is non-convex in (P, Q), so no global minimum exists

Instead we can optimize \mathcal{L} iteratively, e.g.:

- Alternating least squares: update entire matrix *P*, holding *Q* fixed, and vice-versa
- Stochastic gradient descent: update individual rows P_u and Q_i for each observed R_{ui}

Matrix factorization Alternating least squares

 \mathcal{L} is convex in P (Q fixed), so take partials and solve linear system for updates:

$$0 = \frac{\partial \mathcal{L}}{\partial P} \rightarrow \hat{P} = RQ \left[Q^T Q \right]^{-1}$$
(8)

$$0 = \frac{\partial \mathcal{L}}{\partial Q} \quad \rightarrow \quad \hat{Q} = R^{T} P \left[P^{T} P \right]^{-1} \tag{9}$$

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 23 / 31

▲□▶ ▲圖▶ ▲ 臣▶ ▲ 臣▶ ― 臣 … のへで

Matrix factorization Stochastic gradient descent

Alternatively, we can avoid inverting matrices by taking steps in the direction of the negative gradient for each observed rating:

$$P_u \leftarrow P_u - \eta \frac{\partial \mathcal{L}}{\partial P_u} = P_u + \left(R_{ui} - \hat{R}_{ui}\right) Q_i$$
 (10)

$$Q_i \leftarrow Q_i - \eta \frac{\partial \mathcal{L}}{\partial Q_i} = Q_i + \left(R_{ui} - \hat{R}_{ui}\right) P_u$$
 (11)

for some step-size η

April 4, 2011 24 / 31

Matrix factorization Practical notes

Several ways to scale:

- Distribute matrix operations with MapReduce [GHNS11]
- Parallelize stochastic gradient descent [ZWSL10]
- Expectation-maximization for pLSI with MapReduce [DDGR07]

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 25 / 31

Datasets

- Movielens http://www.grouplens.org/node/12
- Reddit http://bit.ly/redditdata
- CU "million songs" http://labrosa.ee.columbia.edu/millionsong/
- Yahoo Music KDDcup http://kddcup.yahoo.com/
- AudioScrobbler http://bit.ly/audioscrobblerdata
- Delicious http://bit.ly/deliciousdata

• . . .

Photo recommendations



http://koala.sandbox.yahoo.com

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 27 / 31

▲ロト ▲掃 ト ▲ 臣 ト ▲ 臣 ト 一 臣 - つへぐ

Thanks. Questions?²

 $^{2}hofman@yahoo-inc.com$

@jakehofman (Yahoo! Research)

Recommendation systems

April 4, 2011 28 / 31

◆□ > ◆□ > ◆臣 > ◆臣 > 三臣 - のへで

References I

- RM Bell, Y Koren, and C Volinsky. The bellkor 2008 solution to the netflix prize. Statistics Research Department at AT&T Research, 2008.
- 🔋 AS Das, M Datar, A Garg, and S Rajaram.

Google news personalization: scalable online collaborative filtering.

page 280, 2007.

- R Gemulla, PJ Haas, E Nijkamp, and Y Sismanis. Large-scale matrix factorization with distributed stochastic gradient descent. 2011.
- JL Herlocker, JA Konstan, LG Terveen, and JT Riedl.
 Evaluating collaborative filtering recommender systems.
 ACM Transactions on Information Systems, 22(1):5–53, 2004.

イロト 不得 トイヨト イヨト 二日

References II



T Hofmann.

Latent semantic models for collaborative filtering. ACM Transactions on Information Systems, 22(1):89–115, 2004.

Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. <u>Computer</u>, 42(8):30–37, 2009.

Yehuda Koren.

The bellkor solution to the netflix grand prize. pages 1–10, Aug 2009.

🔋 G Linden, B Smith, and J York.

Amazon. com recommendations: Item-to-item collaborative filtering.

```
IEEE Internet computing, 7(1):76–80, 2003.
```

References III

A Toscher, M Jahrer, and RM Bell. The bigchaos solution to the netflix grand prize. 2009.

M. Zinkevich, M. Weimer, A. Smola, and L. Li. Parallelized stochastic gradient descent.

In <u>Neural Information Processing Systems (NIPS)</u>, 2010.

イロト 不得 とうき とうとう ほう