

# A Geometric Approach for Learning Latent Mixed Membership Models

Venkatesh Saligrama

Boston University, Boston, MA

(joint work with Weicong Ding & Prakash Ishwar)



# Outline

- Latent Mixture Models
  - Text Documents, User Preferences, Community Networks, ...
  - Overall Goal/Objective: Algorithm with provable guarantees
- Topic Models & Estimation Problem
- Geometric Structure of Topic Models
- Algorithm & Guarantees: Exploiting Geometry
- Rank Aggregation Problem
- Real-World Expts.

Ding, Ishwar, Saligrama, ICML'13  
Ding, Ishwar, Saligrama, NIPS'14  
Ding, Ishwar, Saligrama, ITA'15  
Ding, Ishwar, Saligrama, AISTATS'14  
Ding, Ishwar, Saligrama, AISTATS'15

# Mixed membership latent variable model

- Text Docs  $\leftarrow$  (noisy) mixture of latent topics
- Connections in network  $\leftarrow$  mixture of latent communities
- User preferences  $\leftarrow$  mixtures of latent ranking factors

**Tweets**

**BostonUniversity ECE** @BU\_ece  
Participate in the 2nd Annual #imagineering Competition and have a chance to win money while making a difference! - bit.ly/USd98y  
Expand

**BostonUniversity ECE** @BU\_ece  
Scientists at @toshiba & @bu ensure internet security -  
View summary

**BostonUniversity ECE** @BU\_ece  
Biospired robots may steal danger (via @ScienceDaily)  
Expand

**BostonUniversity ECE** @BU\_ece  
Welcome back! We hope you had the last 3 weeks of the semester  
Expand

**Good News From Iraq**  
Pierre Goldschmidt, Toby Dalton  
Article, November 8, 2012

Good news from the Middle East is rare these days. But Iraq's ratification of its Additional Protocol safeguards agreement with the International Atomic Energy Agency is certainly something to celebrate.


**A Truly Credible Military Threat to Iran**  
David Rothkopf  
Foreign Policy, October 8, 2012

The Romney campaign has argued that Obama has not offered a credible military threat against the Iranians. The easiest way for the Obama team to defuse Romney's critique is to communicate better what options they are in fact considering.


**Avoiding the Iraq Experience in Syria**  
Katherine Wilkens  
National Interest, August 2, 2012

The U.S. experience in Iraq suggests that foreign military involvement could not have prevented the scenario we now see unfolding in Syria.


★★★★★ (291)  
4.4 out of 5 stars



★★★★★ (64)  
4.5 out of 5 stars



★★★★★ (149)  
3.9 out of 5 stars




No SIM 5:45 PM 58%

**Nearby Search**

Filter Check-In Offers Map

- Americas Florist** 0.1 miles  
1020 Ave of the Americas, Midtown West  
2 Reviews Florists  
Check-in Offer: 5% off One Dozen Roses
- Onyx & Jade Salon** 0.2 miles  
41 W 38 St, Midtown West  
8 Reviews Hair Salons  
Check-in Offer: 10% off Your Next Salon Service
- M&J Trimming** 0.1 miles  
1008 6th Ave, Midtown West  
54 Reviews Fabric Stores  
Check-in Offer: 1 free M&J Measuring Tape
- Pulse Karaoke** 0.2 miles  
135 W 41st St, Theater District

Home Nearby Search Bookmarks Check-Ins



# Mixed membership latent variable model

Text document:

**Seeking Life's Bare (Genetic) Necessities**

COLD SPRING HARBOR, NEW YORK— How many genes does an organism need to survive? Last week at the genome meeting here,\* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough.

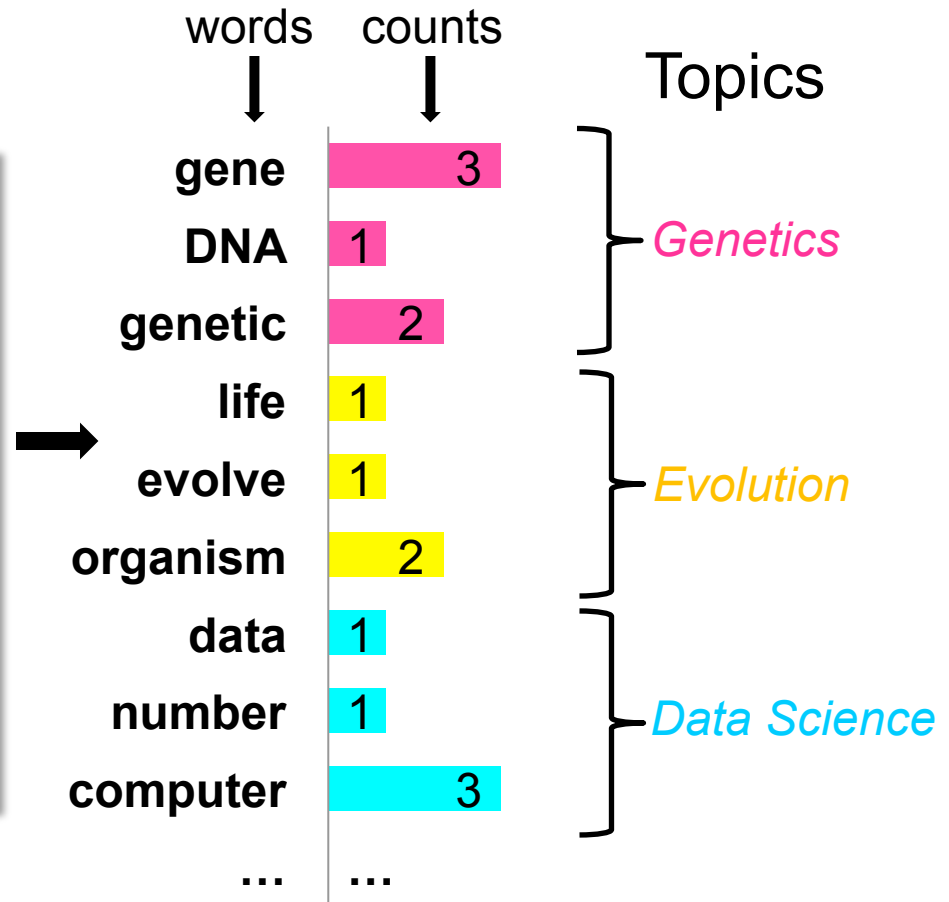
Although the numbers don't match precisely, those predictions

"are not all that far apart," especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. "It may be a way of organizing any newly sequenced genome," explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

**Stripping down. Computer analysis yields an estimate of the minimum modern and ancient genomes.**

\* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

ADAPTED FROM NCBI



**document = mixture of latent topics**

# Mixed membership latent variable model

User preferences

See All Formats >

House of Cards: Season 2 2014 | NR

DVD  
\$27.22 \$66.00 Prime  
Only 10 left in stock - order soon.  
More Buying Choices  
\$22.44 used & new (27 offers)

Blu-ray + UltraViolet  
\$31.49 \$86.00 Prime  
Get it by **Thursday, Mar 5**  
More Buying Choices  
\$26.49 used & new (37 offers)

See All Formats >

Seinfeld: The Complete Series 2013 | NR

DVD  
\$90.85 \$449.00 Prime  
Get it by **Wednesday, Mar 4**  
More Buying Choices  
\$74.99 used & new (20 offers)

See All Formats >

Breaking Bad: The Complete Series 2013 | Unrated

DVD  
\$74.49 \$499.00 Prime  
Get it by **Wednesday, Mar 4**  
More Buying Choices  
\$70.49 used & new (31 offers)


Blu-ray  
\$105.99 \$399.00 Prime  
Get it by **Thursday, Mar 5**  
More Buying Choices  
\$84.99 used & new (52 offers)

See All Formats >

The Big Bang Theory: Season 7 2014 | NR

DVD  
\$19.99 \$44.00 Prime  
Get it by **Wednesday, Mar 4**  
More Buying Choices  
\$15.90 used & new (30 offers)

Blu-ray  
\$29.70 \$84.00 Prime  
Get it by **Wednesday, Mar 4**



>



words



counts



1

Influencing factors

“actor”




>




0

“music”



>



1

“special effect”

**document = mixture of latent influencing factors**

# Overall Goal

- Learn/Estimate Latent Factors from Observations(docs)
- Goal: develop algorithms with
  - **Provable** guarantees → **Model Fidelity**
    - How many Docs to estimate Latent Factors within a tolerance?
    - Computational Cost: How does Algorithm scale with #params?
  - Good **empirical** performance → **Web Scale applications**
    - Real-world datasets

# Outline

- Latent Mixture Models
  - Text Documents, User Preferences, Community Networks, ...
- Topic Models & Estimation Problem
  - Observation Model & Related Work
- Geometric Structure of Topic Models
- Algorithm & Guarantees: Exploiting Geometry
- Extensions to User Preference Model Estimation
- Empirical Results on Real-World Datasets

# “Bag of words” model: a text corpus example

One document in the collection:

### Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK—How many genes does an organism need to survive? Last week at the genome meeting here,\* two genome researchers with radically different approaches presented complementary views of the basic genes needed for life. One research team, using computer analyses to compare known genomes, concluded that today's organisms can be sustained with just 250 genes, and that the earliest life forms required a mere 128 genes. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough. Although the numbers don't match precisely, those predictions

“are not all that far apart,” especially in comparison to the 75,000 genes in the human genome, notes Siv Andersson of Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a genetic numbers game, particularly as more and more genomes are completely mapped and sequenced. “It may be a way of organizing any newly sequenced genome,” explains Arcady Mushegian, a computational molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

**Stripping down.** Computer analysis yields an estimate of the minimum modern and ancient genomes.

\* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.



words	counts
gene	3
DNA	1
genetic	2
life	1
evolve	1
organism	2
data	1
number	1
computer	3
...	...



$$\begin{array}{r} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{array} \begin{bmatrix} \beta_{11} & \beta_{1K} \\ \beta_{21} & \beta_{2K} \\ \vdots & \dots \\ \vdots & \vdots \\ \beta_{W1} & \beta_{WK} \end{bmatrix} \begin{array}{r} \\ \\ \\ \\ \text{topic 1} \\ \text{topic } K \end{array}$$

### Topic Matrix - $\beta$

- column = topic
- $W$  = vocabulary size
- $K$  = # topics

$$\begin{array}{c} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{array} \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1M} \\ A_{21} & A_{22} & \dots & A_{2M} \\ \vdots & \vdots & \dots & \vdots \\ \vdots & \vdots & \dots & \vdots \\ A_{W1} & A_{W2} & \dots & A_{WM} \end{bmatrix} = \begin{array}{c} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{array} \begin{bmatrix} \beta_{11} & \beta_{1K} \\ \beta_{21} & \beta_{2K} \\ \vdots & \vdots \\ \beta_{W1} & \beta_{WK} \end{bmatrix} \begin{bmatrix} \theta_{11} & \theta_{12} & \dots & \theta_{1M} \\ \vdots & \vdots & \dots & \vdots \\ \theta_{K1} & \theta_{K2} & \dots & \theta_{KM} \end{bmatrix}$$

doc. 1    doc. 2    ...    doc. M
topic 1    topic K
doc. 1    doc. 2    ...    doc. M

### Document Distribution matrix - $A$

- column = distb. of a doc.
- $M = \#$  docs.

### Weight matrix - $\theta$

- column = mixing weights
- $M = \#$  docs.

$$\begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{matrix} \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1M} \\ A_{21} & A_{22} & \dots & A_{2M} \\ \vdots & \vdots & \dots & \vdots \\ A_{W1} & A_{W2} & \dots & A_{WM} \end{bmatrix} = \begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{matrix} \begin{bmatrix} \beta_{11} & \beta_{1K} \\ \beta_{21} & \beta_{2K} \\ \vdots & \vdots \\ \beta_{W1} & \beta_{WK} \end{bmatrix} \begin{bmatrix} \theta_{11} & \theta_{12} & \dots & \theta_{1M} \\ \vdots & \vdots & \dots & \vdots \\ \theta_{K1} & \theta_{K2} & \dots & \theta_{KM} \end{bmatrix}$$

doc. 1    doc. 2    ...    doc. M
topic 1    topic K
doc. 1    doc. 2    ...    doc. M

$N$  iid samples + empirical words count

$$\begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{matrix} \begin{bmatrix} X_{11} & X_{12} & \dots & X_{1M} \\ X_{21} & X_{22} & \dots & X_{2M} \\ \vdots & \vdots & \dots & \vdots \\ X_{W1} & X_{W2} & \dots & X_{WM} \end{bmatrix}$$

**Observation matrix  $X$** 

- column = word-freq. of a doc.
- $N = \#$  word/doc.

$$\begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{matrix} \begin{bmatrix} A_{11} & A_{12} & \dots & A_{1M} \\ A_{21} & A_{22} & \dots & A_{2M} \\ \vdots & \vdots & \dots & \vdots \\ A_{W1} & A_{W2} & \dots & A_{WM} \end{bmatrix} = \begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{matrix} \begin{bmatrix} \beta_{11} & \beta_{1K} \\ \beta_{21} & \beta_{2K} \\ \vdots & \vdots \\ \beta_{W1} & \beta_{WK} \end{bmatrix} \begin{bmatrix} \theta_{11} & \theta_{12} & \dots & \theta_{1M} \\ \vdots & \vdots & \dots & \vdots \\ \theta_{K1} & \theta_{K2} & \dots & \theta_{KM} \end{bmatrix}$$

doc. 1    doc. 2    doc. M
topic 1    topic K
doc. 1    doc. 2    doc. M

↓

$$\begin{matrix} \text{word 1} \\ \text{word 2} \\ \vdots \\ \text{word } W \end{matrix} \begin{bmatrix} X_{11} & X_{12} & \dots & X_{1M} \\ X_{21} & X_{22} & \dots & X_{2M} \\ \vdots & \vdots & \dots & \vdots \\ X_{W1} & X_{W2} & \dots & X_{WM} \end{bmatrix}$$

**Problem**

- Given :  $X$  and  $K$
- Goal : estimate  $\beta$

$K$	# topics	~100
$W$	vocab. size	~10k
$N$	#word/doc.	~100
$M$	# doc.	~100k

# Related work

Topic matrix Weight matrix

Method	$\beta$	$\theta$	Approach	Issues
Nonnegative Matrix Factorization (NMF), e.g., [Cichocki et al., '09]	Deterministic	Deterministic	Regularized joint optimization	NP Hard (Arora'12) Non-convex. Need approximations and heuristics.
“Bayesian Methods” e.g., LDA, CTM [Blei et al., '03],	Deterministic or Prior	Prior	MAP or ML	Non-convex. Need approximations like MCMC.
Method of Moments [Anandkumar et al., '12, '13]	Deterministic and sparse	Prior	Tensor decomposition	No empirical performance reported
Topic-separability based [Several references]	Approximate Separability	Prior	Geometric	

# Outline

- Latent Mixture Models
  - Text Documents, User Preferences, Community Networks, ...
- Topic Models & Estimation Problem
  - Related Work
- Geometric Structure of Topic Models
  - Inevitability of Separability in high-dimensions
- Algorithm & Guarantees: Exploiting Geometry
- Extensions to User Preference Model Estimation
- Empirical Results on Real-World Datasets

# Approximately Separable Topic Matrix ([Ding-Ishwar-S'14])

		Genetics			Data		
		topic 1	topic 2	topic 3	topic 1	topic 2	topic 3
DNA	w1	$\beta_1$	0	0	0	0	0
	w2	$\beta_2$	0	0	0	0	0
Computer	w3	0	$\beta_3$	0	0	0	0
	w4	0	$\beta_4$	0	0	0	0
	w5	0	0	$\beta_5$	0	0	0
	w6	0	0	$\beta_6$	0	0	0
		$\beta_{71}$	$\beta_{72}$	$\beta_{73}$	...		

**Separable Topic Matrix**

$\lambda$ -approximately separable if one word for each topic is predominantly unique

$\lambda = 0$  Case: Novel Word(s)

**unique** to each topic

[Boardman'93, Donoho'04],

[Arora'13, Ding et. al.'13]

		topic 1	topic 2	topic 3
w1	$\beta_1$	$\leq \lambda \beta_1$	$\leq \lambda \beta_1$	$\leq \lambda \beta_1$
w2	$\beta_2$	$\beta_{21}$	$\lambda \beta_2$	
w3	$\leq \lambda \beta_3$	$\beta_3$	$\leq \lambda \beta_3$	
w4	$\lambda \beta_4$	$\beta_4$	$\lambda \beta_4$	
w5	$\lambda \beta_5$	$\leq \lambda \beta_5$	$\beta_5$	
w6	$\lambda \beta_6$	$\lambda \beta_6$	$\beta_6$	
		$\beta_{71}$	$\beta_{72}$	$\beta_{73}$
		...		


**Approximately Separable**

# Is Approximately Separable Fundamental?

- In real-world problems
  - Size of vocab.  $W \gg \#$ . Topics  $K$

Dataset	$W$	$K$
Wikipedia	109,611	50
Twitter	122,035	50
New York Times	102,660	100
PubMed	141,043	150

- Main result:** Separability is an inevitable consequence of high-dimensionality!
  - Satisfied in estimates produced by NMF, LDA, and other algorithms (Bayesian Models)






# Why is Separability inevitable for $W \gg K$ ?

- Theorem:** Suppose  $W \geq tK e^{\beta_0 K \log(K+1/\lambda)}$

$$\text{Prob}\{\beta \text{ not } \lambda\text{-sep}\} = O(W^{-t})$$

Dataset	W	K
Wikipedia	109,611	50
Twitter	122,035	50
New York Times	102,660	100
PubMed	141,043	150




# Separability in Practice

Dataset	Vocab. size $W$	# Topics $K$	Prob. 0.01-separable
NIPS	12,419	50	100±0%
Wikipedia	109,611	50	99.9±0.3%
Twitter	122,035	50	100±0%
New York Times	102,660	100	99.6±0.6%
PubMed	141,043	150	99.9±0.3%

$\beta_0 = 0.01$ ,  
1000 MC runs

## Generative Model



- $\beta_0$  moderately small positive value in practice.
  - $\beta_0 = 0.01$  for  $K \in [50, 200]$
- Some packages suggest  $\beta_0 = c/W$  to get satisfactory empirical results.
 
$$W \geq tK e^{\beta_0 K} \log(K+1/\lambda)$$
- Analysis explains reasoning for this choice!

# Outline

- Latent Mixture Models
  - Text Documents, User Preferences, Community Networks, ...
- Topic Models & Estimation Problem
  - Related Work
- Geometric Structure of Topic Models
  - Inevitability of Separability in high-dimensions
- Algorithm & Guarantees: Exploiting Geometry
  - Efficiently Identifying Extreme Points
  - Empirical Results on Real-World Datasets
- Extensions to User Preference Model Estimation

# Key Idea ( $\lambda = 0$ case)

$$\beta_{W \times K} \theta_{K \times M} = A_{W \times M}$$

	topic 1	topic 2	topic 3
w1	$\beta_1$	0	0
w2	$\beta_2$	0	0
w3	0	$\beta_3$	0
w4	0	$\beta_4$	0
w5	0	0	$\beta_5$
w6	0	0	$\beta_6$
	$\beta_{71}$	$\beta_{72}$	$\beta_{73}$
		...	

doc. 1	...	doc. M
← $\theta_1$ →		
← $\theta_2$ →		
← $\theta_3$ →		


**Weight Matrix**



doc. 1	...	doc. M
← $\beta_1 \theta_1$ →		
← $\beta_2 \theta_1$ →		
← $\beta_3 \theta_2$ →		
← $\beta_4 \theta_2$ →		
← $\beta_5 \theta_3$ →		
← $\beta_6 \theta_3$ →		
$\beta_{71} \theta_1 + \beta_{72} \theta_2 + \beta_{73} \theta_3$		
...		

**Distribution Matrix A**

**Separable Topic Matrix  $\beta$**



# Key Idea ( $\lambda = 0$ case)

$$\beta_{W \times K} \theta_{K \times M} = A_{W \times M}$$

	topic 1	topic 2	topic 3
w1	$\beta_1$	0	0
w2	$\beta_2$	0	0
w3	0	$\beta_3$	0
w4	0	$\beta_4$	0
w5	0	0	$\beta_5$
w6	0	0	$\beta_6$
	$\beta_{71}$	$\beta_{72}$	$\beta_{73}$
		...	

doc. 1	...	doc. M
← $\theta_1$ →		
← $\theta_2$ →		
← $\theta_3$ →		


**Weight Matrix**



doc. 1	...	doc. M
← $\beta_1 \theta_1$ →		
← $\beta_2 \theta_1$ →		
← $\beta_3 \theta_2$ →		
← $\beta_4 \theta_2$ →		
← $\beta_5 \theta_3$ →		
← $\beta_6 \theta_3$ →		
$\beta_{71} \theta_1 + \beta_{72} \theta_2 + \beta_{73} \theta_3$		
...		

**Distribution Matrix A**

**Separable Topic Matrix  $\beta$**




$\mathcal{R}_+^M$

# Key Idea

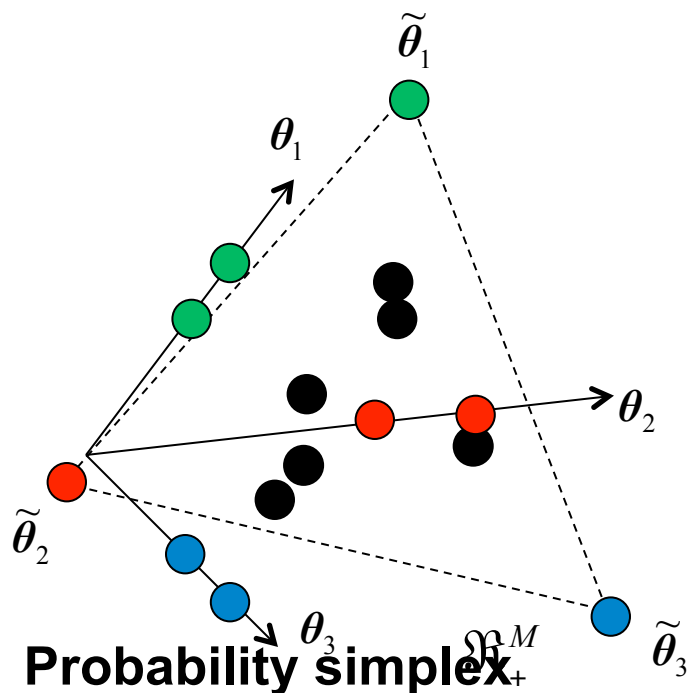
$$\tilde{\beta}_{W \times K} \tilde{\theta}_{K \times M} = \tilde{A}_{W \times M}$$

Novel Word = Extreme Point


Novel word detection + Topic matrix estimation



Distribution Matrix  $A$



Probability simplex  $S_+^M$




Row Normalized Distribution Matrix  $\tilde{A}$

# Finite words/doc.

$$\begin{array}{c}
 \text{doc.} \quad \dots \quad \text{doc.} \\
 1 \quad \quad \quad M \\
 \left[ \begin{array}{ccc}
 \leftarrow & \beta_1 \boldsymbol{\theta}_1 & \rightarrow \\
 \leftarrow & \beta_2 \boldsymbol{\theta}_1 & \rightarrow \\
 \leftarrow & \beta_3 \boldsymbol{\theta}_2 & \rightarrow \\
 \leftarrow & \beta_4 \boldsymbol{\theta}_2 & \rightarrow \\
 \leftarrow & \beta_5 \boldsymbol{\theta}_3 & \rightarrow \\
 \leftarrow & \beta_6 \boldsymbol{\theta}_3 & \rightarrow \\
 \beta_{71} \boldsymbol{\theta}_1 + \beta_{72} \boldsymbol{\theta}_2 + \beta_{73} \boldsymbol{\theta}_3 & & \\
 \dots & & 
 \end{array} \right]
 \end{array}$$

**Distribution Matrix  $A$**

row normalization




$$\begin{array}{c}
 \text{doc.} \quad \dots \quad \text{doc.} \\
 1 \quad \quad \quad M \\
 \left[ \begin{array}{ccc}
 \leftarrow & \tilde{\boldsymbol{\theta}}_1 & \rightarrow \\
 \leftarrow & \tilde{\boldsymbol{\theta}}_1 & \rightarrow \\
 \leftarrow & \tilde{\boldsymbol{\theta}}_2 & \rightarrow \\
 \leftarrow & \tilde{\boldsymbol{\theta}}_2 & \rightarrow \\
 \leftarrow & \tilde{\boldsymbol{\theta}}_3 & \rightarrow \\
 \leftarrow & \tilde{\boldsymbol{\theta}}_3 & \rightarrow \\
 \tilde{\beta}_{71} \tilde{\boldsymbol{\theta}}_1 + \tilde{\beta}_{72} \tilde{\boldsymbol{\theta}}_2 + \tilde{\beta}_{73} \tilde{\boldsymbol{\theta}}_3 & & \\
 \dots & & 
 \end{array} \right]
 \end{array}$$

**Row Normalized Distribution Matrix  $A$**

# Finite words/doc.


$$\tilde{\beta}_{W \times K} \tilde{\theta}_{K \times M} = \tilde{A}_{W \times M} \approx \tilde{X}_{W \times M}$$

**Key issue:**  
 $N$  fixed  $\rightarrow$  perturbation  
 does not vanish




**Row Normalized  
 Distribution  
 Matrix  $\tilde{A}$**

$K$	# topics	$\sim 100$
$W$	vocab. size	$\sim 10k$
$N$	#word/doc.	$\sim 100$
$M$	# doc.	$\sim 100k$



**Probability simplex**



**Row Normalized  
 Observation  
 Matrix:  $\tilde{X}$**



**Nonnegative**

$$\tilde{\beta}_{W \times K} \tilde{\theta}_{K \times M} = \tilde{A}_{W \times M}$$



**Separable**  
**Nonnegative**

$K$	# topics	~100
$W$	vocab. size	~10k
$N$	#word/doc.	~100
$M$	# doc.	~100k

**Nonnegative**

$$M \tilde{X} \tilde{X}^T \xrightarrow[M \rightarrow \infty]{a.s.} \tilde{\beta} (M \tilde{\theta} \tilde{\theta}^T) \tilde{\beta}^T = \mathbf{E}_{W \times W}$$


**Separable**  
**Nonnegative**



Novel word detection + Topic matrix estimation


# Detect Novel Words via Projections

- Max/Min of projection  
→ extreme points of convex hull
- Which directions to use  
→ Generate  $P$  iid  
Isotropy directions




# “Robustness” of extreme points

- Max/Min of projection  
→ extreme points of convex hull
- Which directions to use  
→ Generate a few iid Isotropy directions
- Freq. of maximum  $\approx$   
Solid Angle of an  
Extreme Point



# Extension to Approximately Separable ( $\lambda > 0$ )

- Approx. novel words  $\Leftrightarrow$  larger solid angles
- Solution
  - Sort solid angles
  - Take the top – K extreme points



# Main result [Ding et al, '13,'14]

- Computational complexity :

$$O(MNK + WK + WK^3)$$

$K$	# topics	~100
$W$	vocab. size	~10k
$N$	#word/doc.	~100
$M$	# doc.	~100k

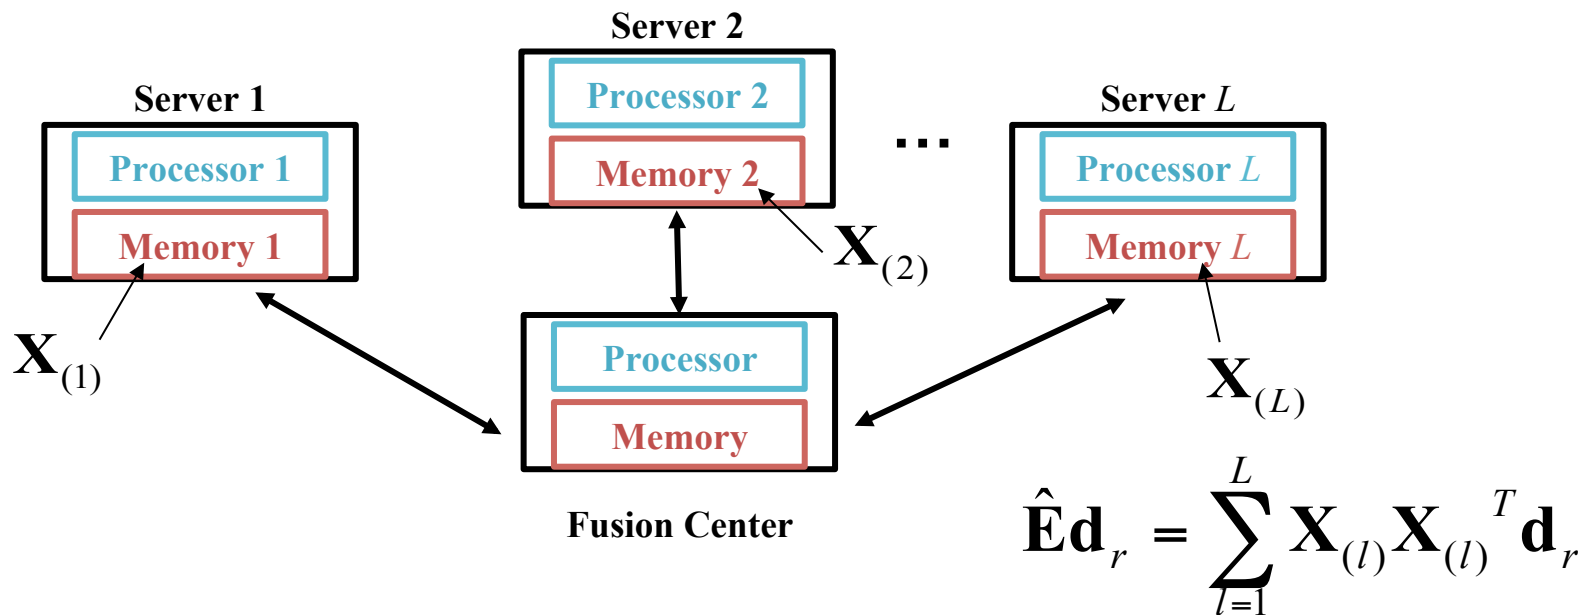
- Sample complexity :

Under the **Simplicial Condition** on  $R'$ , with  $u \sim N(\mathbf{0}, I_W)$ , the proposed Random Projection algorithm can detect all novel words of  $K$  topics with probability  $1-\delta$  if

$$M \geq \text{Poly}\left(W, \log\left(\frac{1}{\delta}\right), K, \frac{1}{N}\right), P \geq \text{Poly}\left(\log W, \log\left(\frac{1}{\delta}\right), K\right)$$

Moreover, if  $R$  is **full-rank**, can recover  $\beta$  with  $\varepsilon$  element-wise error with probability  $1-\delta$ .

# Distributed Implementation



- Modern web-scale database are distributed
- $M$  document archived on  $L$  servers
- Goal: **low communication cost  $O(WK)$**

# Experimental Results (semi-synthetic data)

Real-world corpus  
New York Times articles

Topic matrix learnt  
by Gibbs sampling

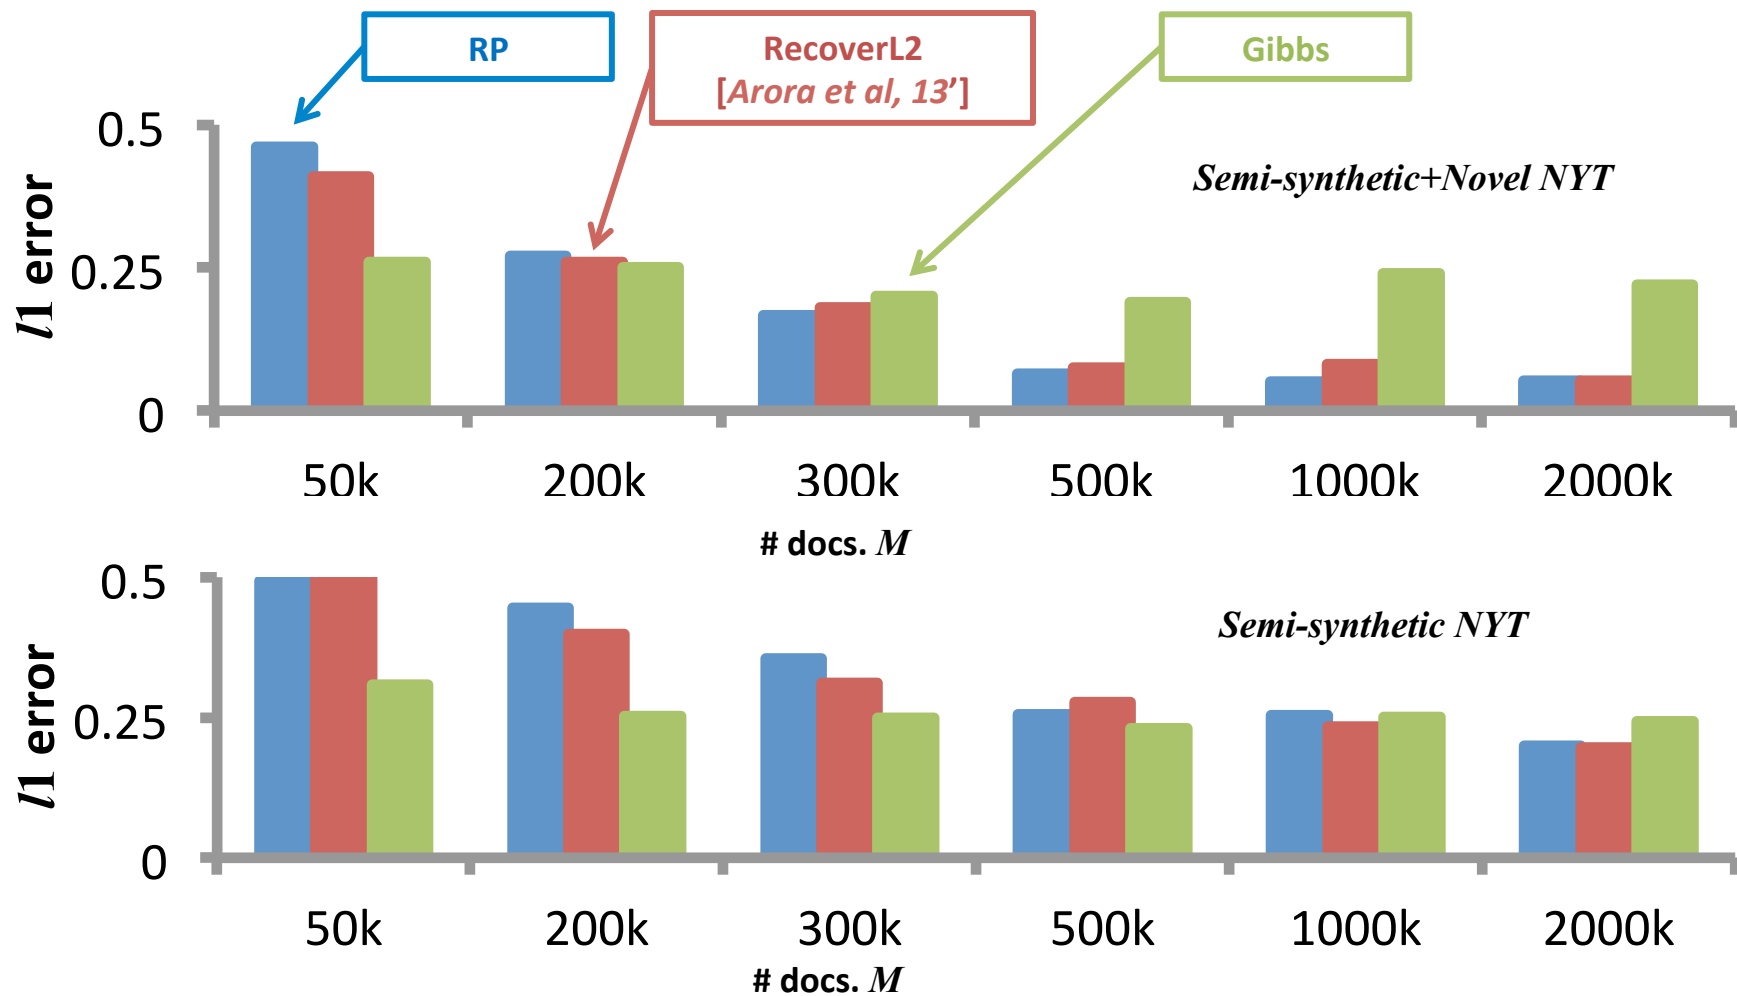
Generate synthetic  
docs. with Dirichlet  
prior

Add artificial novel  
words;  
Generate synthetic  
docs.

$M$	300,000
$N$	300
$W$	14943
$K$	100
$L$	200

- Semi-synthetic data can resemble the dimensionality and sparsity of real-world data

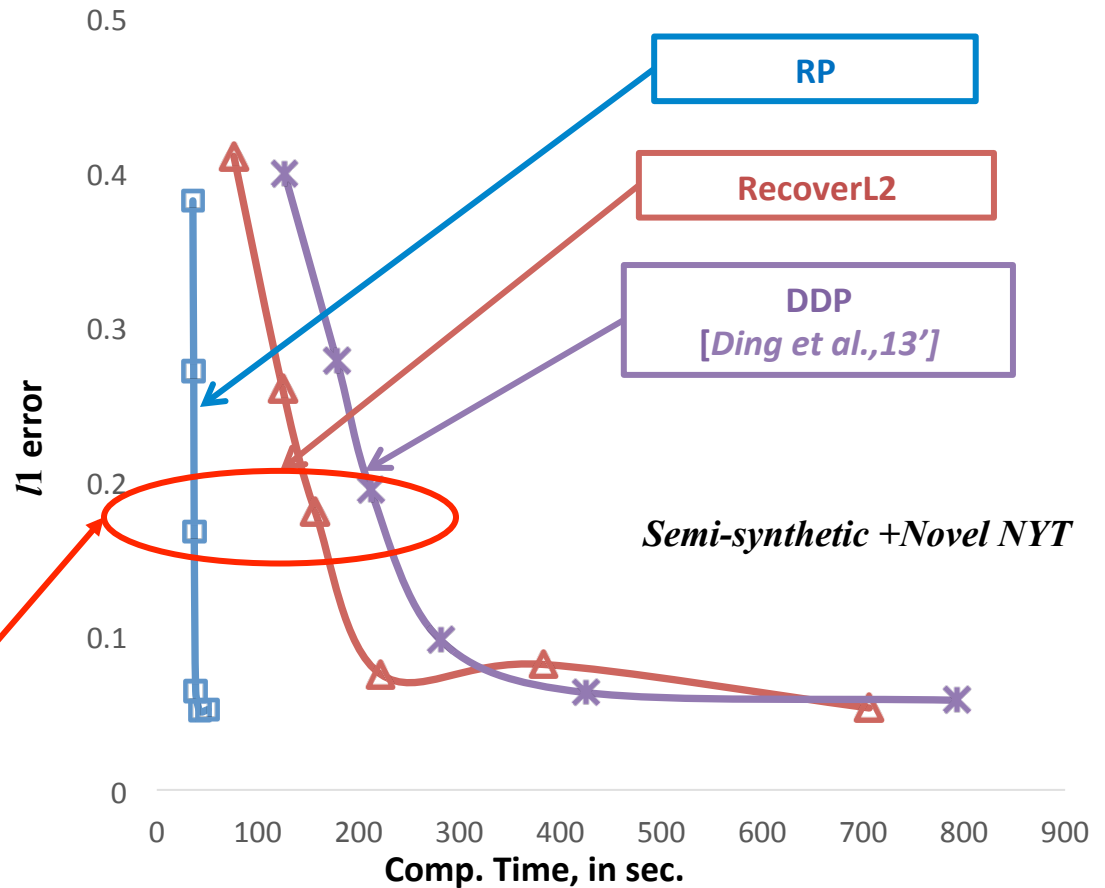
# Experimental Results (semi-synthetic data)





# Experimental Results (semi-synthetic data)

	<i>Semi-syn +novel NYT</i>	<i>Semi-syn NYT</i>
$M$	Variable	Variable
$N$	300	300
$W$	15043 words, 100 novel	14943
$K$	100	100
$L$	200	200



\*  $T$  (Gibbs) ~ 6918 sec

# *Experimental Results* (Real-World Text Corpus)

New York Times dataset

$M$	300,000 = 240k train + 60k test
$N$	300 words/doc. (avg.)
$W$	14,943
$K$	50/100/150 topics

Decreasing Prob. ↓

“Weather”	“Emotion”	“Politics”	“Football”
Weather	Feeling	Election	Yard
Wind	Sense	<i>Florida</i>	Game
Air	Love	Ballot	Team
Storm	Character	Vote	Season
Rain	Heart	<i>Al_gore</i>	Play
Cold	Emotion	Recount	<i>NFL</i>

(See [Ding et al., '13](#) for more example topics)

# Outline

- Latent Mixture Models
  - Text Documents, User Preferences, Community Networks, ...
- Topic Models & Estimation Problem
  - Related Work
- Geometric Structure of Topic Models
  - Inevitability of Separability in high-dimensions
- Algorithm & Guarantees: Exploiting Geometry
  - Efficiently Identifying Extreme Points
  - Empirical Results on Real-World Datasets
- Rank Aggregation Problem
  - Heterogenous Population ~ Mixture of Mallows Model
  - Reduction to Topic Modeling Problem
  - Empirical Results on Real-World Datasets

# Mixed membership latent variable model

User preferences

See All Formats >

House of Cards: Season 2 2014 | NR

DVD  
\$27.22 \$66.00 Prime  
Only 10 left in stock - order soon.  
More Buying Choices  
\$22.44 used & new (27 offers)

Blu-ray + UltraViolet  
\$31.49 \$86.00 Prime  
Get it by **Thursday, Mar 5**  
More Buying Choices  
\$26.49 used & new (37 offers)

See All Formats >

Seinfeld: The Complete Series 2013 | NR

DVD  
\$90.85 \$449.00 Prime  
Get it by **Wednesday, Mar 4**  
More Buying Choices  
\$74.99 used & new (20 offers)

See All Formats >

Breaking Bad: The Complete Series 2013 | Unrated

DVD  
\$74.49 \$499.00 Prime  
Get it by **Wednesday, Mar 4**  
More Buying Choices  
\$70.49 used & new (31 offers)


Blu-ray  
\$105.99 \$399.00 Prime  
Get it by **Thursday, Mar 5**  
More Buying Choices  
\$84.99 used & new (52 offers)

See All Formats >

The Big Bang Theory: Season 7 2014 | NR

DVD  
\$19.99 \$44.00 Prime  
Get it by **Wednesday, Mar 4**  
More Buying Choices  
\$15.90 used & new (30 offers)

Blu-ray  
\$29.70 \$64.00 Prime  
Get it by **Wednesday, Mar 4**



words



counts



1


Influencing factors

*“actor”*



0

*“music”*



1

*“special effect”*

**document = mixture of latent influence factors**

Prob. Prefer movie 1  
over 2 in 1<sup>st</sup> latent  
factor

$$\begin{array}{c}
 \text{actor} \\
 \text{musical} \\
 \begin{array}{c}
 \xrightarrow{(1,2)} \\
 \xrightarrow{(1,3)} \\
 \vdots \\
 \vdots \\
 \xrightarrow{(Q-1,Q)}
 \end{array}
 \left[ \begin{array}{cc}
 \beta_{11} & \beta_{1K} \\
 \beta_{21} & \beta_{2K} \\
 \vdots & \vdots \\
 \vdots & \vdots \\
 \beta_{W1} & \beta_{WK}
 \end{array} \right] \\
 \begin{array}{cc}
 \text{topic 1} & \text{topic } K
 \end{array}
 \end{array}$$

Ranking matrix -  $\beta$

- column = "topic"
- $W = \#$  ordered pairs  
=  $Q(Q - 1)$
- $K = \#$  topics

- Generative Model for Latent factor

- Mallows model

- Baseline permutation  $\sigma_0$


- Prob. of permutation  $\sigma$ :  $Prob\{\sigma \mid \sigma_0\} \propto \phi^{dist(\sigma, \sigma_0)} \Rightarrow \beta$

- Heterogeneous population

- Dispersion factor  $\phi$

# Mixed Membership Rank Aggregation Problem

“Plate” representation



$w$ : comparisons  $(i, j)$

Key parameters

$Q$	# items
$W=Q(Q-1)$	# ordered pairs
$K$	# latent rankings
$M$	# users
$N$	# comps./user

→ Words

→ Topics

→ Documents

# Related Work

Category	Models	Issues	Heterogeneity	User-inconsistency
Single ranking models	Mallows [ <i>Mallows, '57</i> ], BTL [ <i>Bradley &amp; Terry, '52</i> ], etc.	Only one global ranking	Only via deviation from base ranking scheme	Via noisy observation model
Mixture of ranking models	Mixture of Mallows model [ <i>Lu &amp; Boutilier ICML11, Awasthi et al. NIPS14</i> ], Mixture of single rankings [ <i>Farias et al. NIPS09</i> ], Mixture of BTLs [ <i>Oh &amp; Shah NIPS14</i> ]	Each user is dominated by one type	Multiple mixture components	Via noisy observation model
“Topic” ranking models	[ <i>Ding et al. NIPS14</i> ], Mixed membership Mallows, etc.		Multiple <b>shared</b> rankings	Via probabilistic mixture + noisy observation

# Approximate Separability

- Most ranking matrix are  $\lambda$ -Approximate separable, # items  $Q \gg$  # factors  $K$

$$\Pr(\sigma \text{ is } \lambda\text{-separable}) \geq 1 - K \exp(-QL(\lambda; \phi)^{-2K+1})$$

$\phi$	Prob. of 0.05-separable
0.0	93.3%
0.1	87.0%
0.2	79.3%
0.5	42.6%

	Rank 1	Rank 2	Rank 3
Pair 1	0.98	0.01	0.01
Pair 2	0.01	0.99	0.01
Pair 3	0.01	0.01	0.90
Pair 4	0.98	0.90	0.10
Pair 5	0.10	0.09	0.90
		...	
		$\beta$	

Approximately  
Separable  
ranking matrix

$$Q = 100$$

$$K = 10$$

1000 Monte Carlo runs



# Main result

- Computational complexity :

$$O(MNP + Q^2 K^3)$$

- Sample complexity :

- If
  - Correlation matrix of weight prior has **full-rank** and
  - The ranking matrix  $\sigma$  is  **$\lambda$ -separable** and
  - $\lambda \leq cK^{-2}$
- Then
  - Proposed Random Projection algorithm can estimate the ranking matrix correctly with probability at least  $1-\delta$  for all

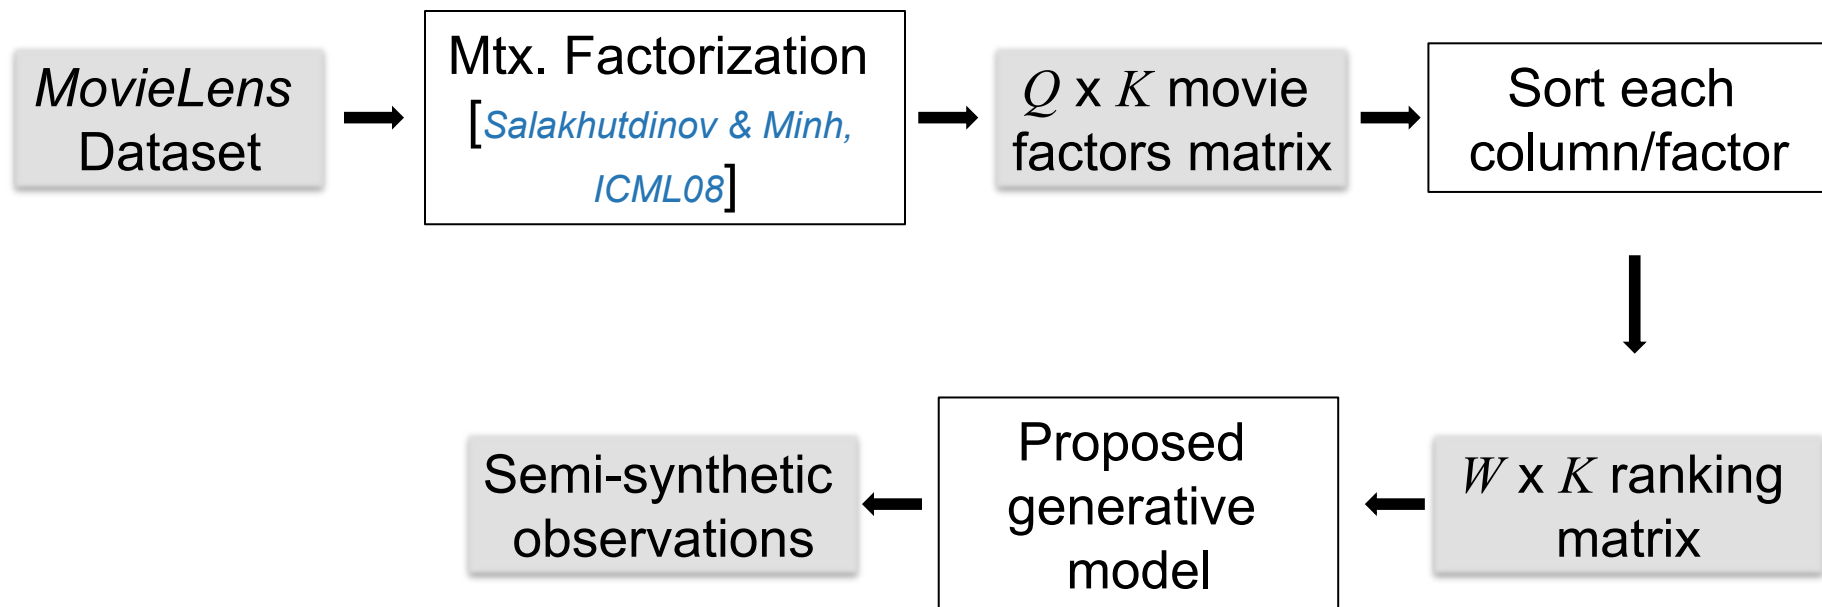
$$M \geq \text{Poly}\left(W, K, \frac{1}{N}, \log\left(\frac{1}{\delta}\right)\right), P \geq \text{Poly}\left(\log W, \log\left(\frac{1}{\delta}\right), K\right)$$

$Q$	# items
$W$	# ordered pairs
$K$	# latent ranking
$M$	# users
$N$	# comp. / user
$P$	# projections


# Semi-synthetic data

$Q$	100 most rated movies
$K$	10 latent rankings
$M$	5940 users
	~200K ratings

- To resemble the dimensionality and characteristics of real-world data



# Semi-synthetic data



$Q$	100 most rated movies
$K$	10 latent rankings
$N$	300 comparisons/user
$M$	# user, variable

- Dirichlet Prior for  $\theta$
- Uniform distribution for  $\mu$
- $\phi = \phi_1 = \dots = \phi_K$

# Movielens dataset – new comparison

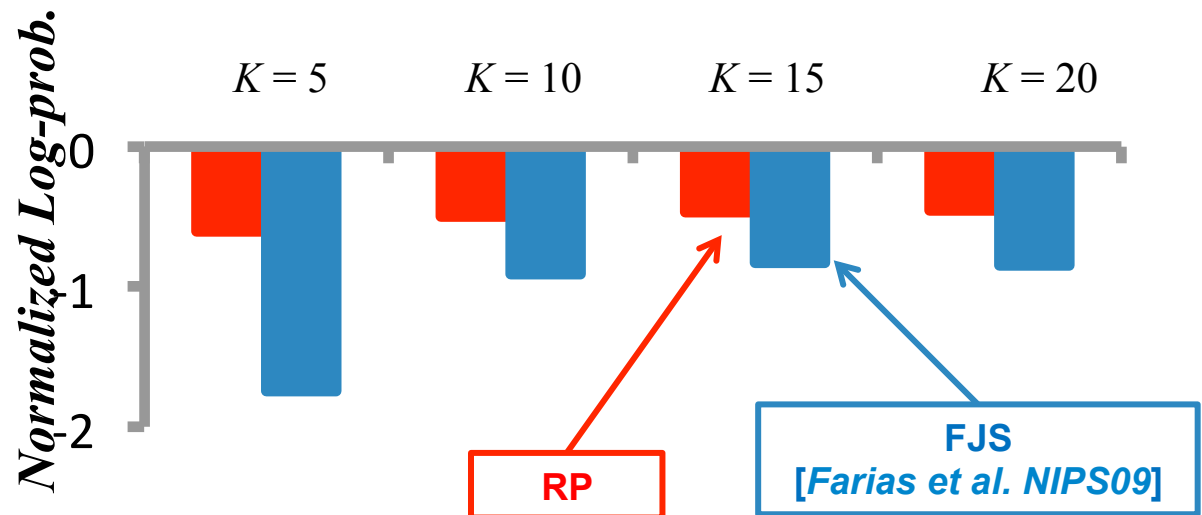
- **Data:** generate comparisons from ratings

User 1: Movie A, 4 star Movie C, 2 star  $\rightarrow w = (A, C)$  for user 1

- **Task:** predict new comparisons for users in the training set
- **Measure:** predictive log-probability [*Wallach et al. ICML09*]

$$\frac{1}{\text{Total \# test pairs}} \sum_{i,m} \log\{p(i \text{ th test pair} \mid \hat{\sigma}, \text{Training pairs of user } m)\}$$

$Q$	100 most rated movies
$K$	variable
$M$	5940 users



# Movielens dataset – new user

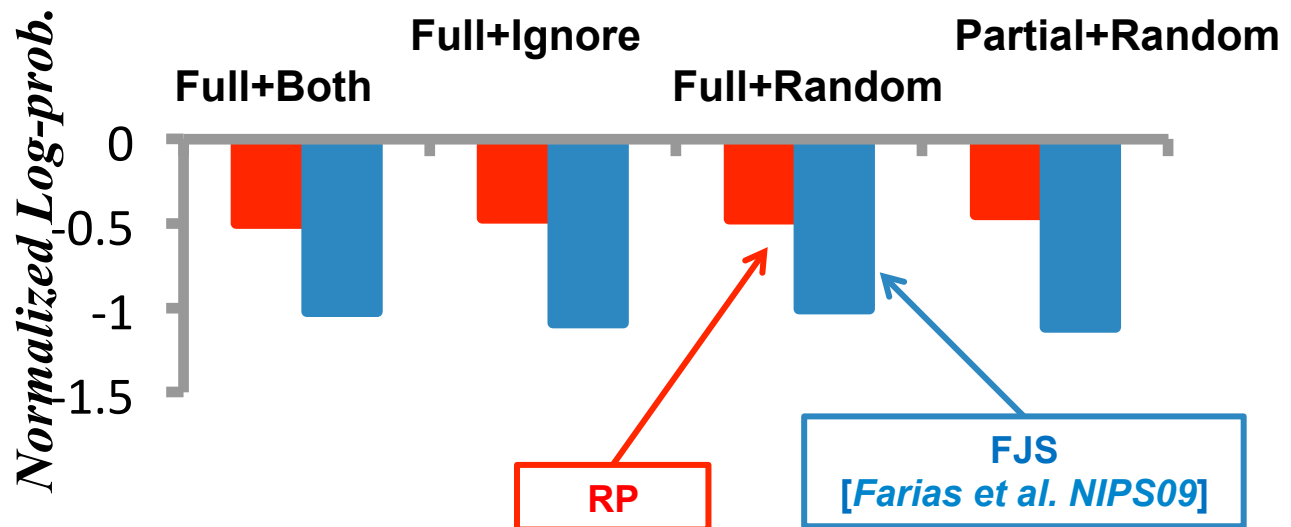
- **Data:** generate comparisons from ratings

User 1: Movie A, 4 star Movie C, 2 star  $\rightarrow w = (A, C)$  for user 1

- **Task:** predict comparison of new users
- **Measure:** predictive log-probability [*Wallach et al. ICML09*]

$$\frac{1}{\text{Total \# test pairs}} \sum_{i,m} \log \{p(i \text{ th test pair of user } m \mid \hat{\sigma})\}$$

$Q$	100 most rated movies
$K$	10 latent rankings
$M$	4000 <i>training</i> user
$M$	2040 <i>testing</i> user




# *Movielens* dataset – predicting stars

- Predict star ratings via ranking models
  - Generate comparisons from training ratings
  - Learn mixed membership Mallows model with Dirichlet prior
  - For each testing movie review:
- Measure: RMSE of estimated star ratings

K	PMF	BPMF	BPMF-int	TM(Ding et al.,14)	MMMM
10	1.0491	0.8254	0.8723	0.8840	0.8509
15	0.9127	0.8236	0.8734	0.8780	0.8296
20	0.9250	0.8213	0.8678	0.8721	0.8241

Rating based models

# Summary



High-D Latent Factor Models  
Geometry  $\sim$  Approx Sep.

Simple geometric picture

Efficient randomized algorithm

Consistency, efficiency, state-of-the-art performance

	Rank 1	Rank 2	Rank 3
Pair 1	0.98	0.01	0.01
Pair 2	0.01	0.99	0.01
Pair 3	0.01	0.01	0.90
Pair 4	0.98	0.90	0.10
Pair 5	0.10	0.09	0.90
		...	
		$\beta$	

Approximately  
Separable  
ranking matrix