Probabilistic Topic Model

Jie Tang

Tsinghua University
October 2011

Why Topic?

- Math & Feature Space: SVD -> PCA
- Information Retrieval: Tfidf->lsi->plsi->lda
- More:
  - Lda->topic correlation model
  - Static->dynamic
  - Bow->word-order

Why Topic? (cont.)

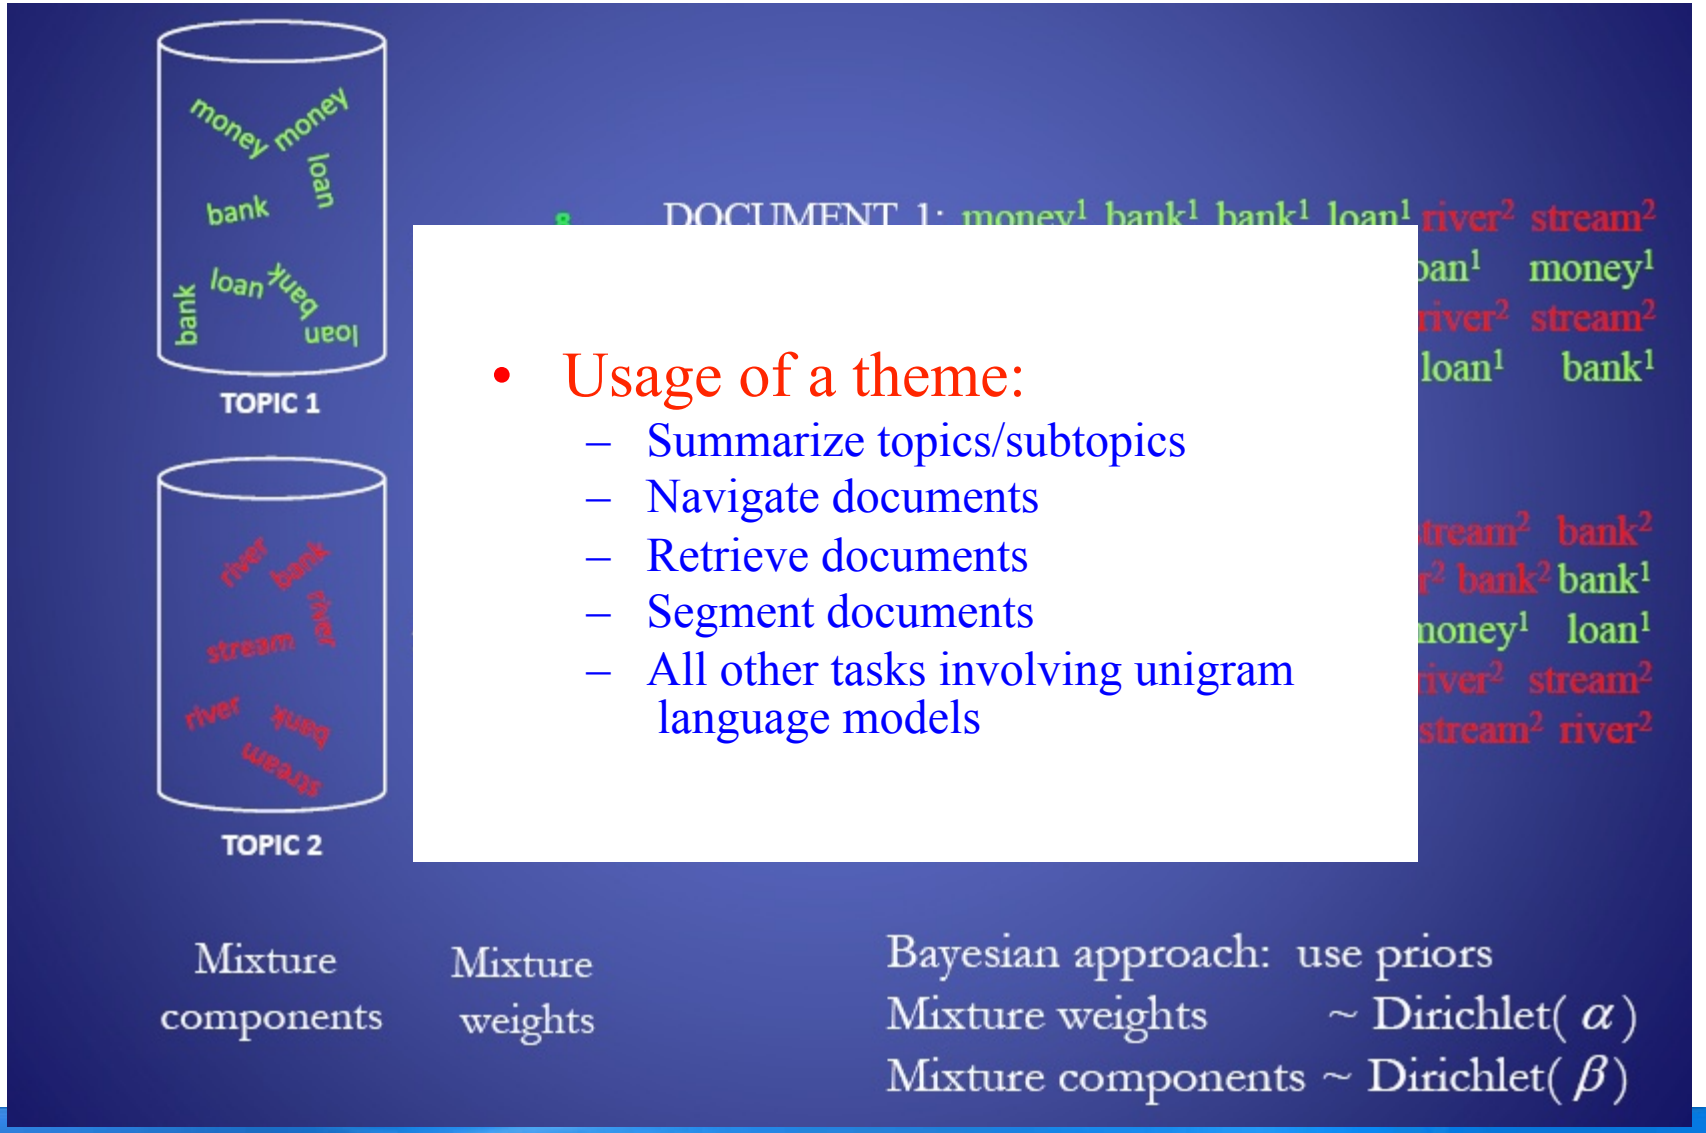
- IR – text to real number vector(Baeza-Yates and Ribeiro -Neto, 1999), tfidf (Salton and McGill, 1983)
  - Tfidf – shortcoming: (1) Lengthy and (2) Cannot model inter- and intra- document statistical structure
- LSI – dimension reduction (Deerwester et al., 1990)
  - Advantages: achieve significant compression in large collections and capture synonymy and polysemy.
- Generative probabilistic model – to study the ability of LSI (Papadimitriou et al., 1998)
  - Why LSI, we can model the data directly using maximum likelihood or Bayesian methods.

What is Topic?

"Arts"	"Budgets"	"Children"	"Education"
NEW	MILLION	CHILDREN	SCHOOL
FILM	TAX	WOMEN	STUDENTS
SHOW	PROGRAM	PEOPLE	SCHOOLS
MUSIC	BUDGET	CHILD	EDUCATION
MOVIE	BILLION	YEARS	TEACHERS
PLAY	FEDERAL	FAMILIES	HIGH
MUSICAL	YEAR	WORK	PUBLIC
BEST	SPENDING	PARENTS	TEACHER
ACTOR	NEW	SAYS	BENNETT
FIRST	STATE	FAMILY	MANIGAT
YORK	PLAN	WELFARE	NAMPHY
OPERA	MONEY	MEN	STATE
THEATER	PROGRAMS	PERCENT	PRESIDENT
ACTRESS	GOVERNMENT	CARE	ELEMENTARY
LOVE	CONGRESS	LIFE	HAITI

The William Randolph Hearst Foundation will give $1.25 million to Lincoln Center, Metropolitan Opera Co., New York Philharmonic and Juilliard School. "Our board felt that we had a real opportunity to make a mark on the future of the performing arts with these grants an act every bit as important as our traditional areas of support in health, medical research, education and the social services," Hearst Foundation President Randolph A. Hearst said Monday in announcing the grants. Lincoln Center's share will be $200,000 for its new building, which will house young artists and provide new public facilities. The Metropolitan Opera Co. and New York Philharmonic will receive $400,000 each. The Juilliard School, where music and the performing arts are taught, will get $250,000. The Hearst Foundation, a leading supporter of the Lincoln Center Consolidated Corporate Fund, will make its usual annual $100,000 donation, too.

What is Topic? (cont.)

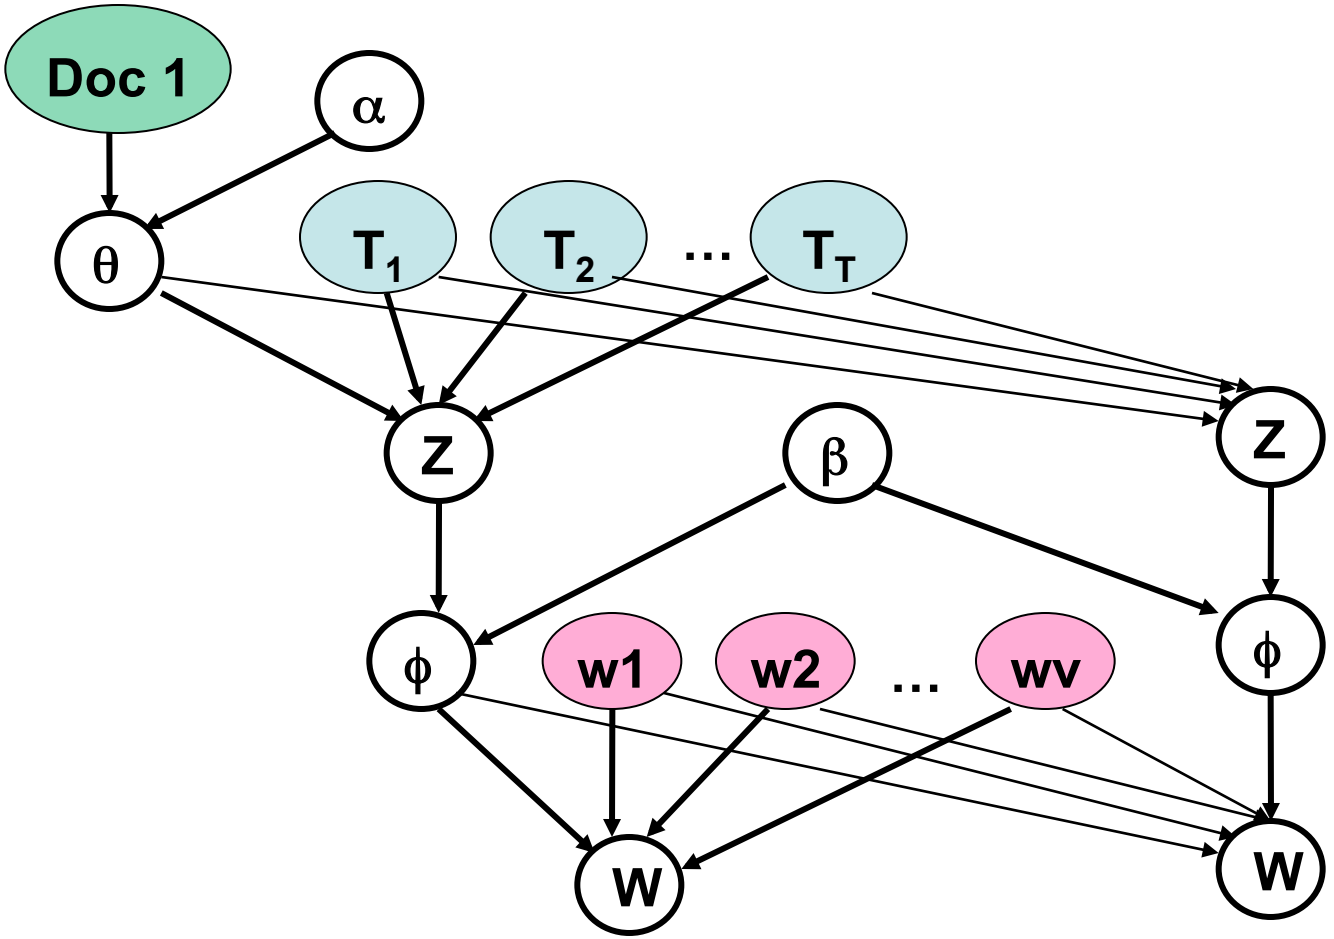


- Usage of a theme:
  - Summarize topics/subtopics
  - Navigate documents
  - Retrieve documents
  - Segment documents
  - All other tasks involving unigram language models

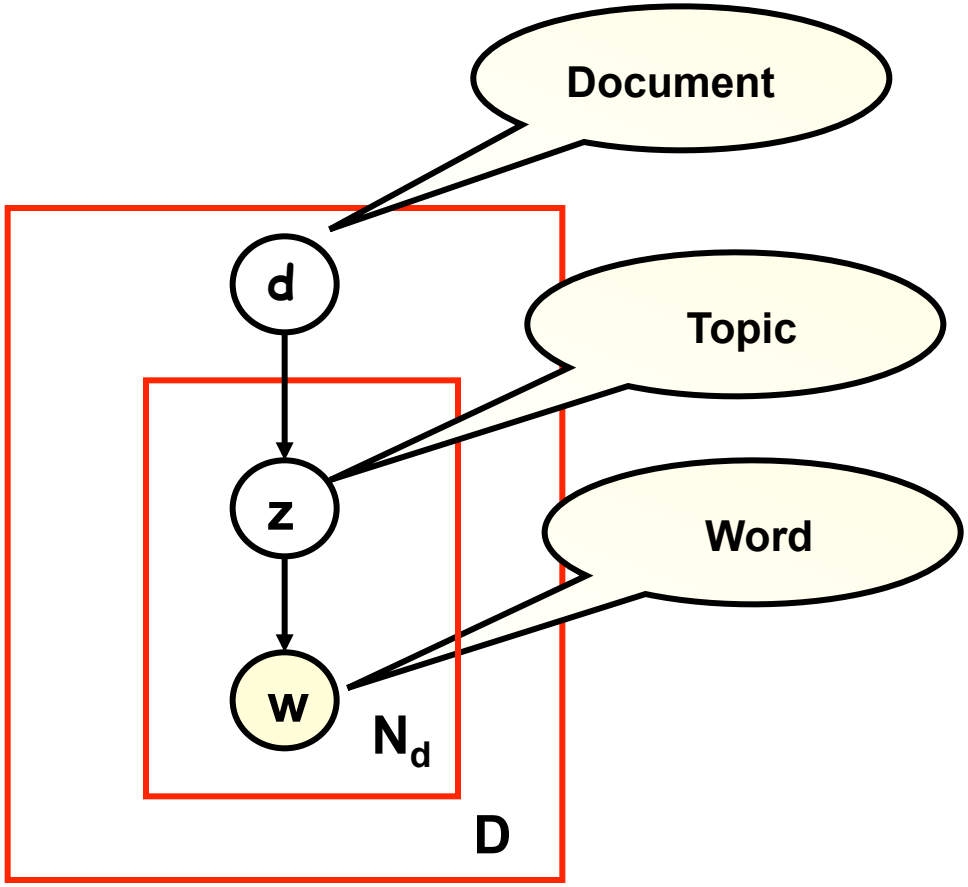
Topic Model Family Member

- Directed graphical model (Bayesian network model)
- Undirected graphical model (Markov random field, restricted boltzmann machines)
- We will only focus on directed graphical model in this talk.

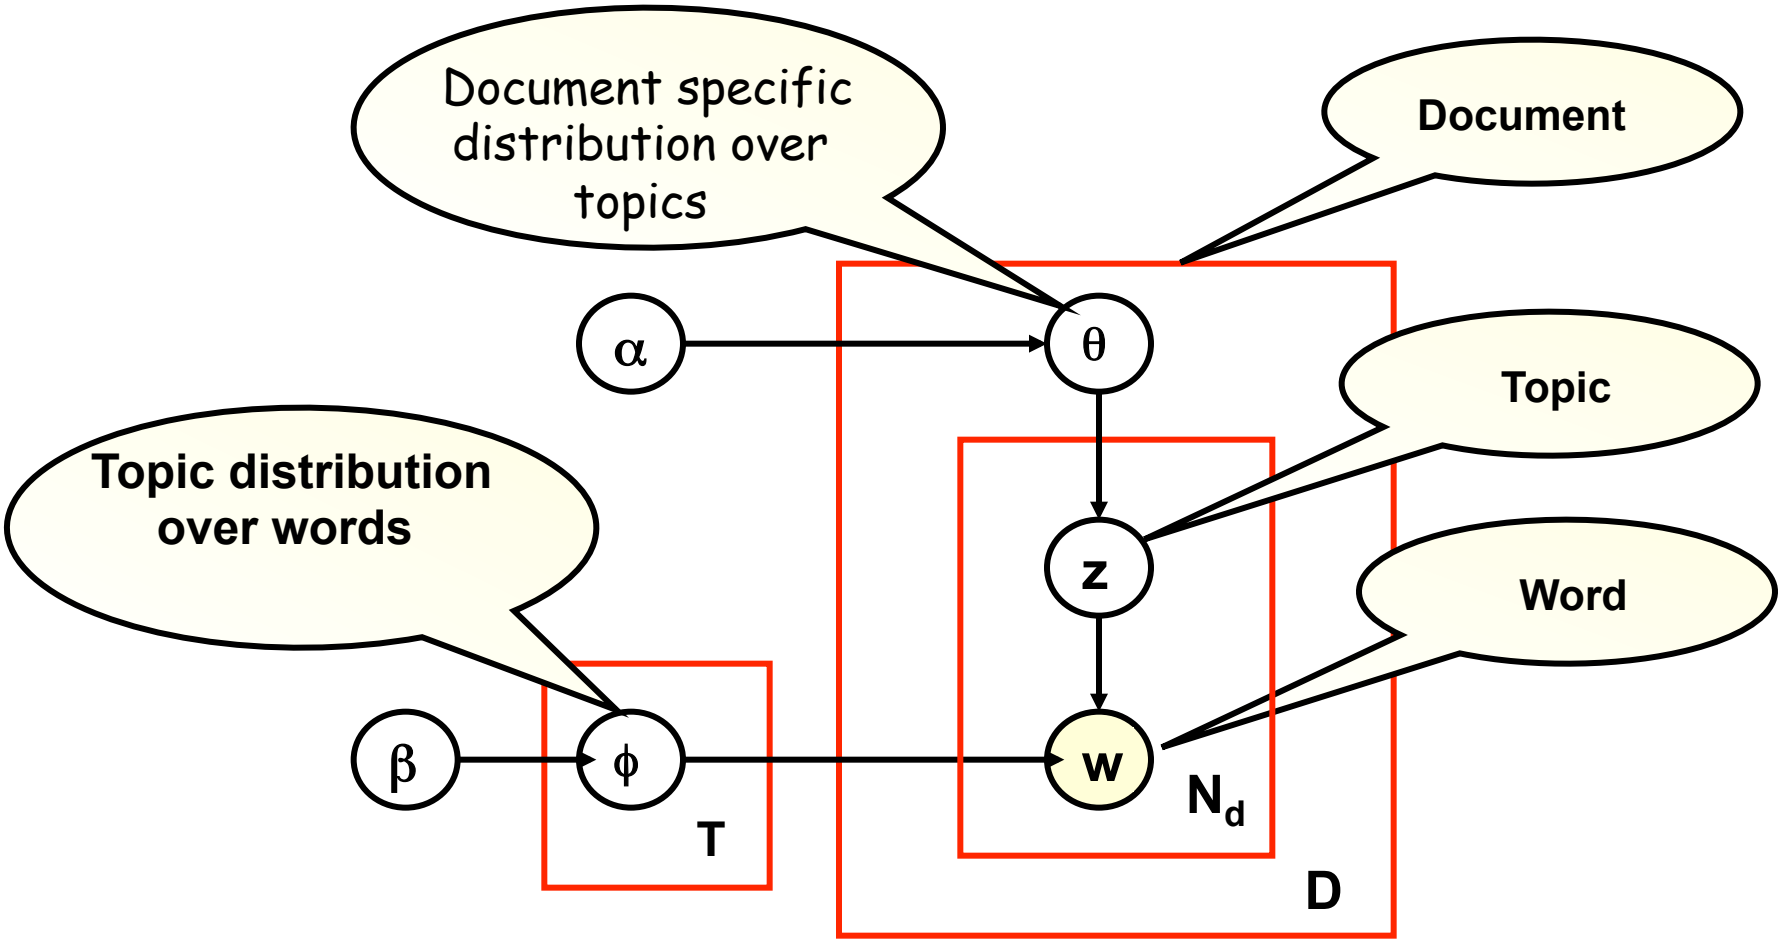
Bayesian Network for modeling document generation



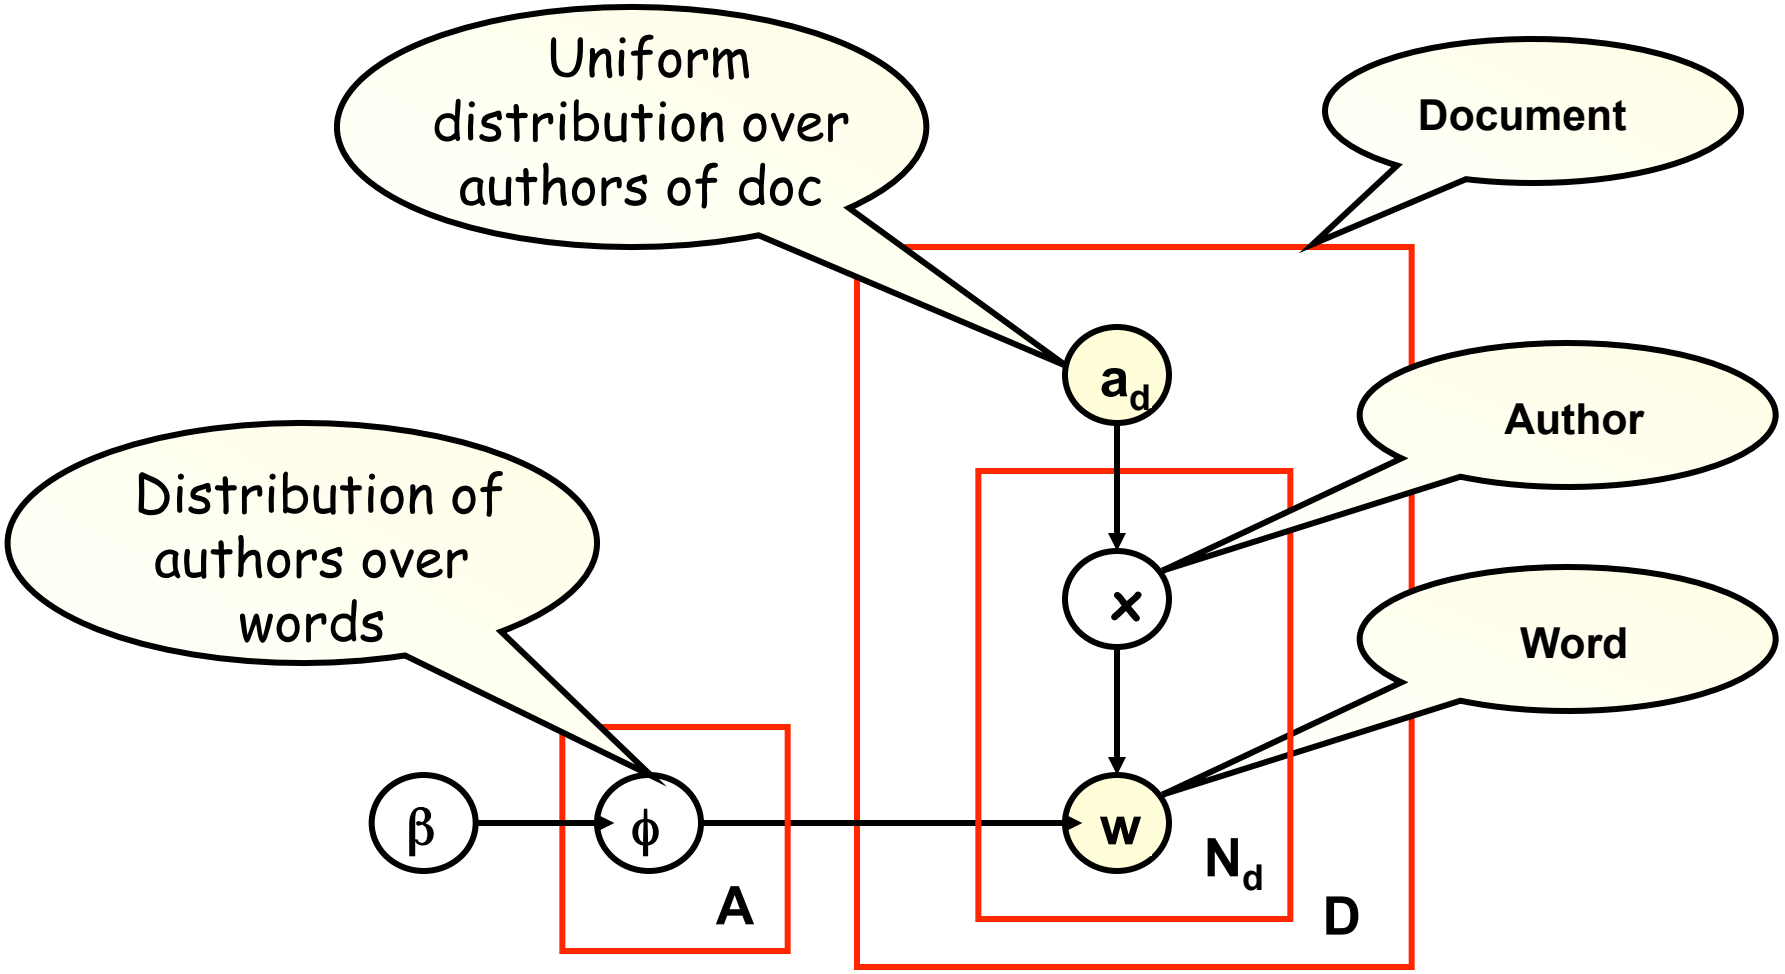
PLSI(SIGIR'99)



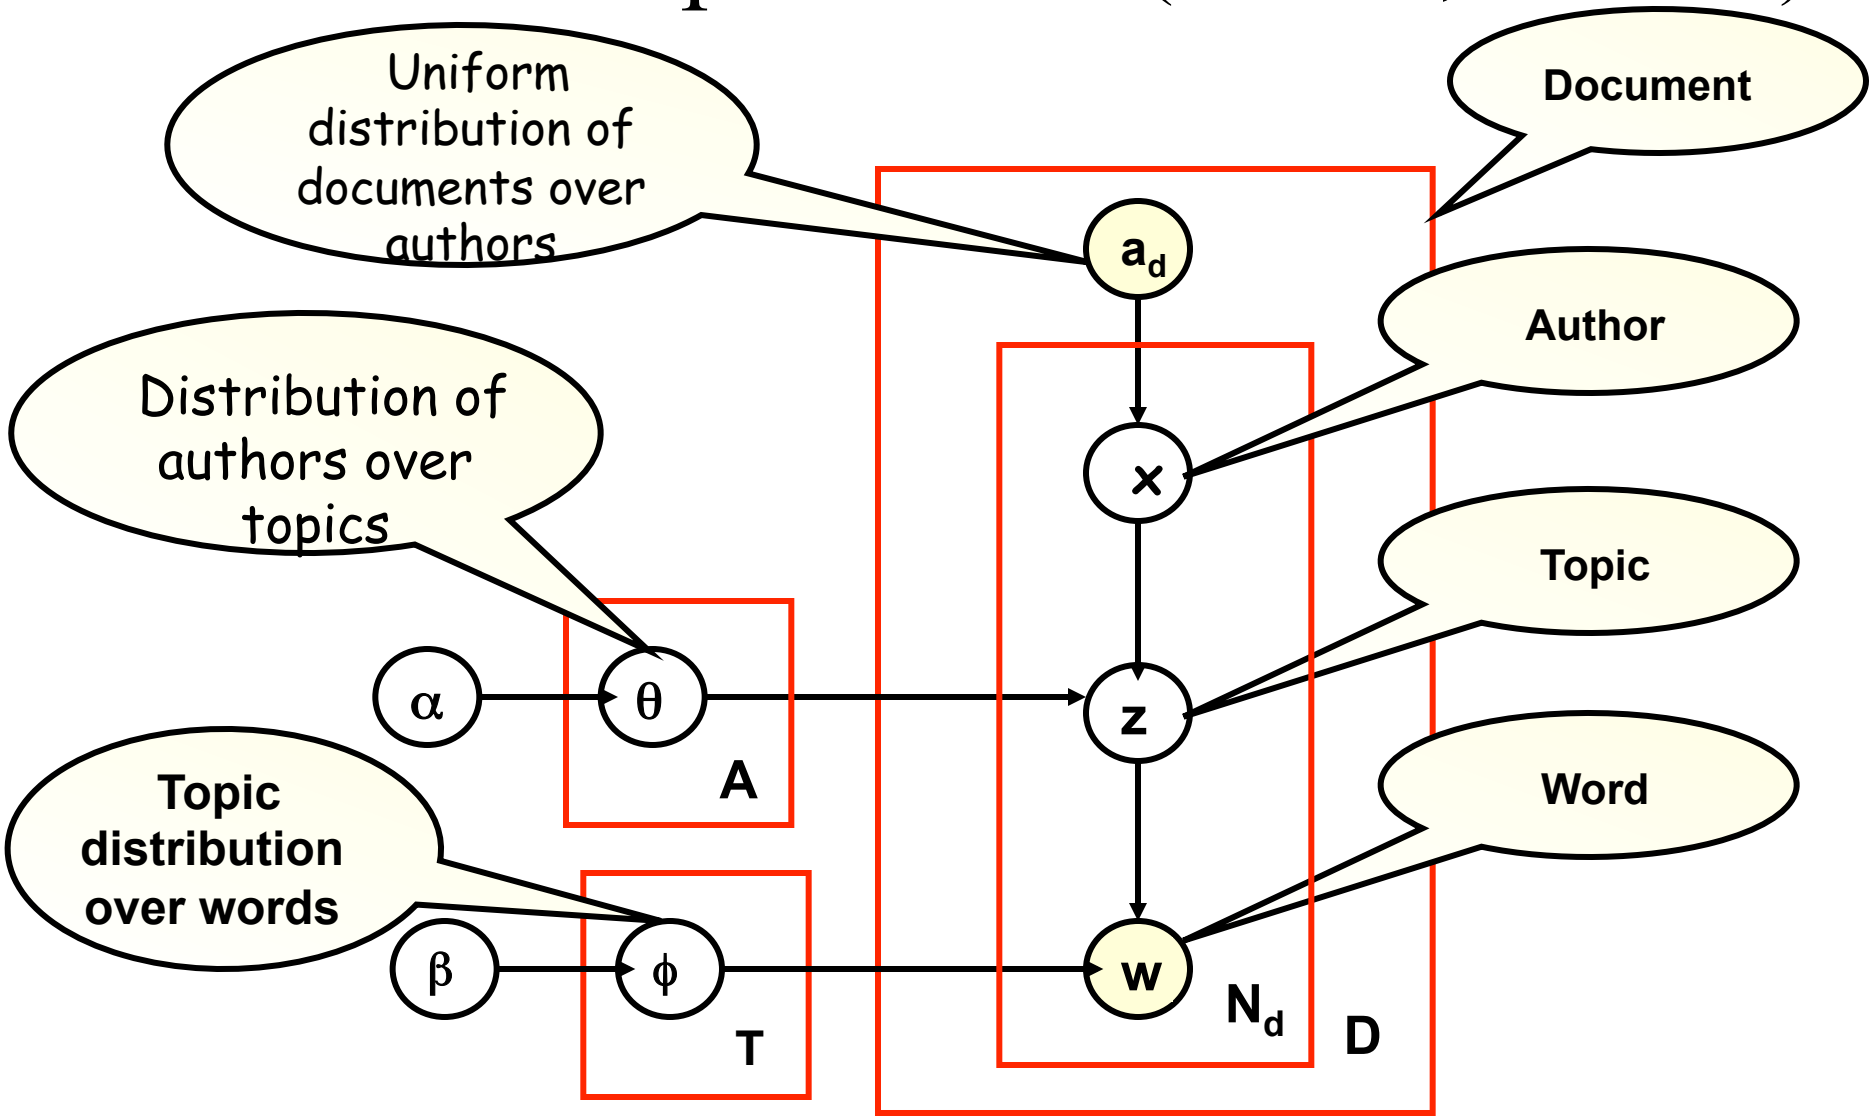
LDA (JMLR'03)


Document specific distribution over topics
Document
α
θ
Topic
Topic distribution over words
z
Word
β
ϕ
w
T
N_d
D

Modeling Authors with Words
(AAAI'09 workshop)



Author-Topic Model (UAI'04, KDD'05)

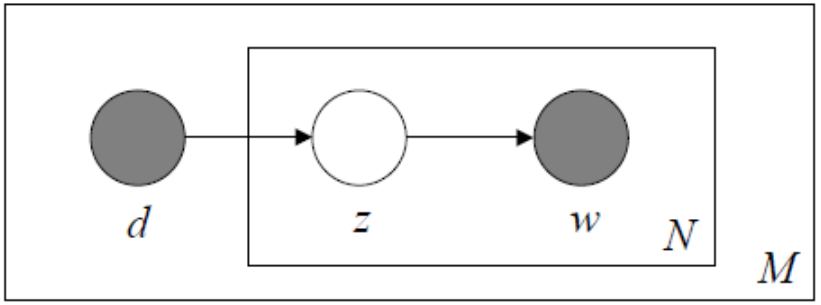


Two state-of-the-art topic models:

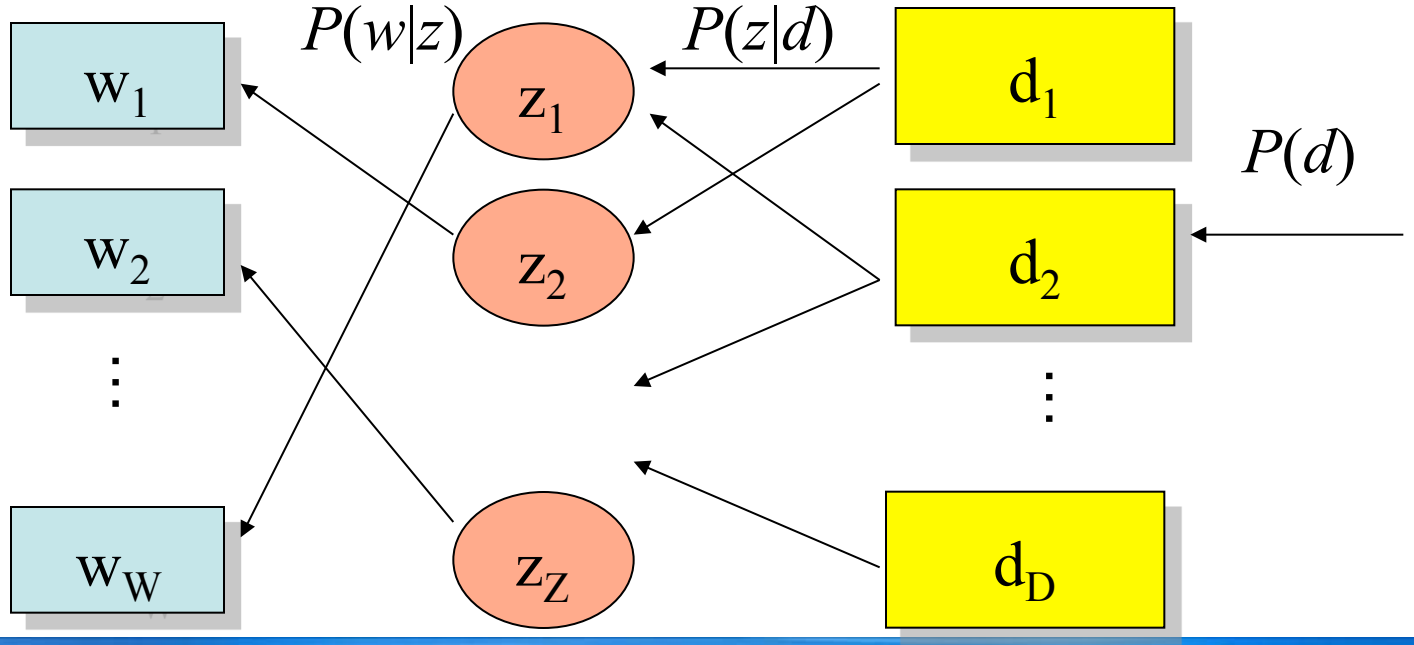
- PLSI
- LDA

Probabilistic Latent Semantic Indexing (PLSI)

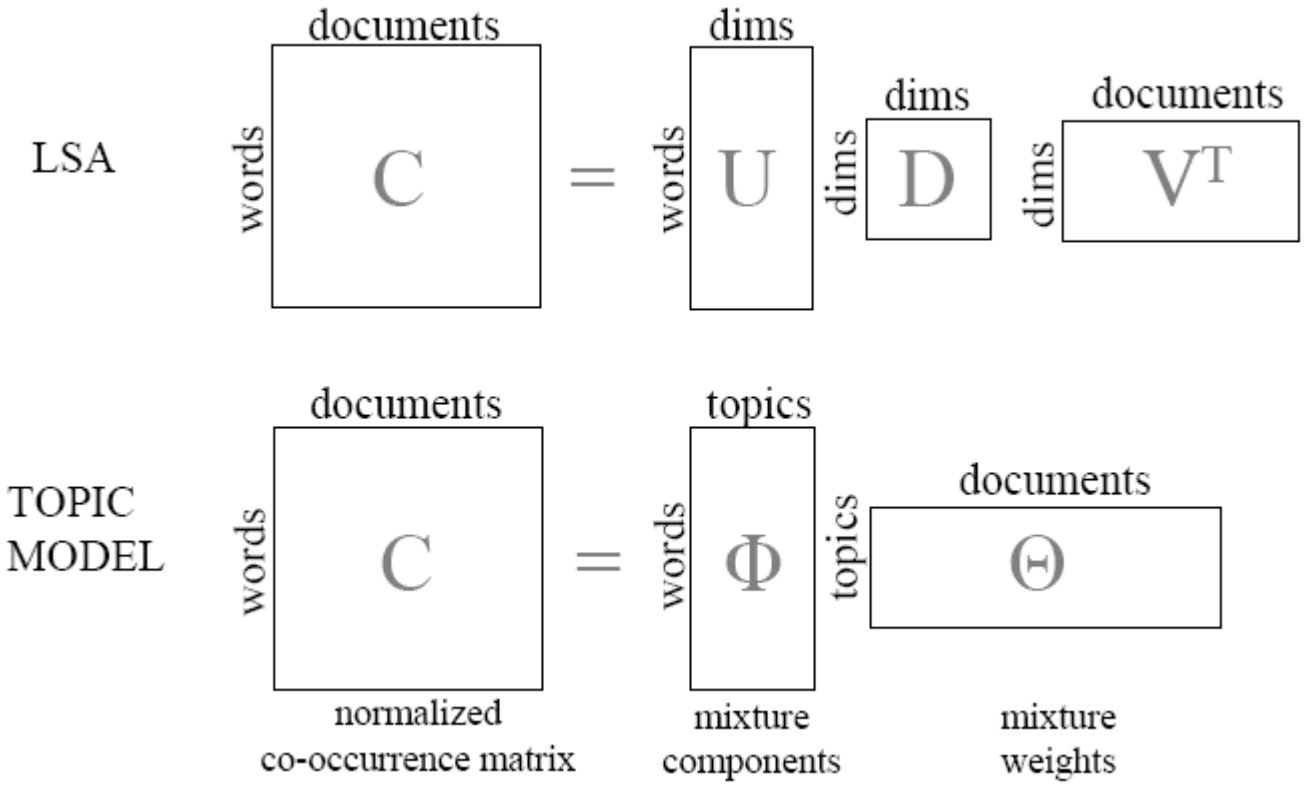
Model



- A generative model for generating the co-occurrence of documents $d \in D=\{d_1,\ldots,d_D\}$ and terms $w \in W=\{w_1,\ldots,w_W\}$, which associates latent variable $z \in Z=\{z_1,\ldots,z_Z\}$.
- The generative processing is:

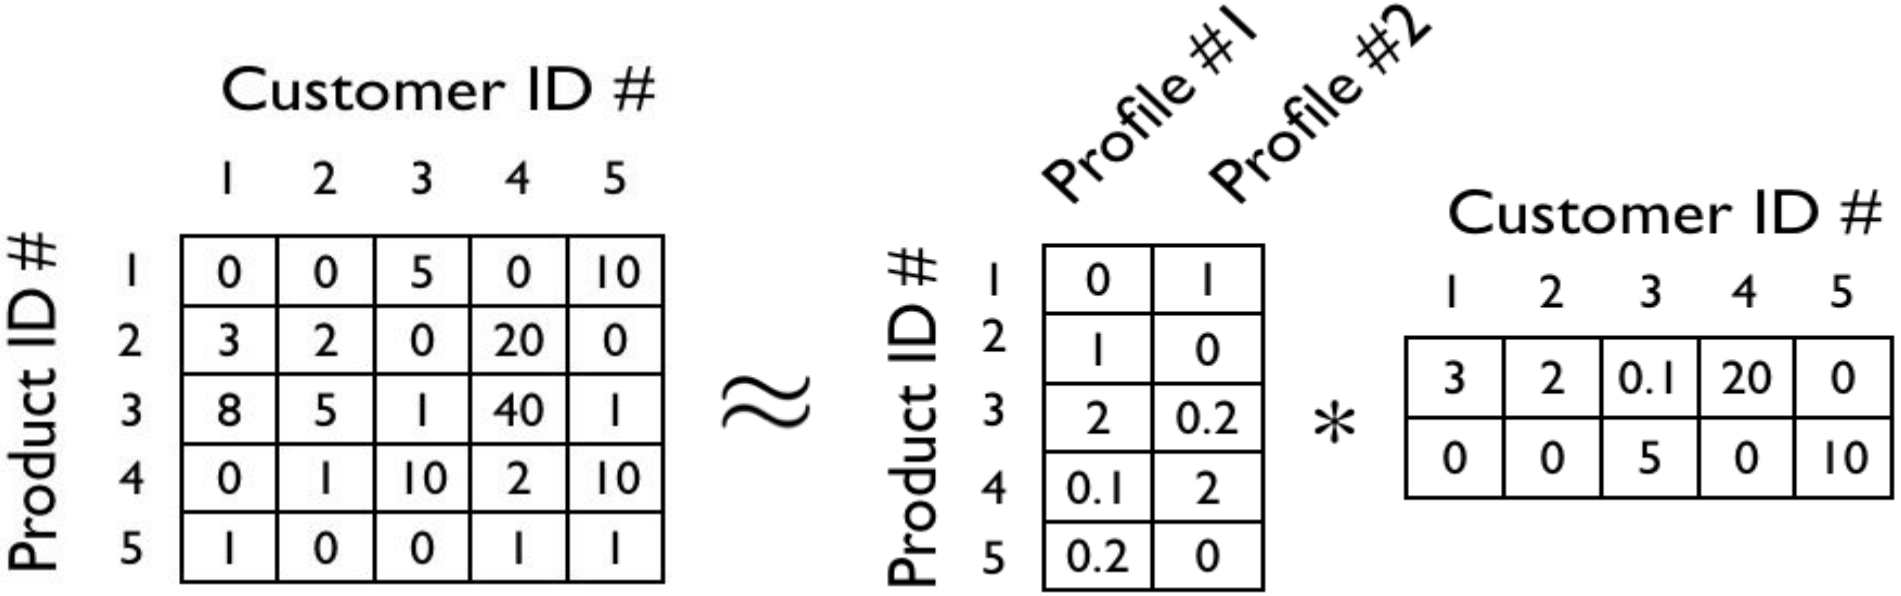


Matrix Interpretation



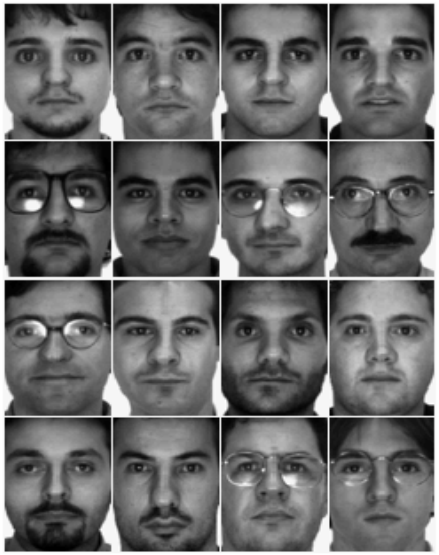
- In the topic model, the word-document co-occurrence matrix is split into two parts: a topic matrix Φ and a document matrix Θ. Note that the diagonal matrix D in LSA can be absorbed in the matrix U or V, making the similarity between the two representations even clearer.

Example 1: products

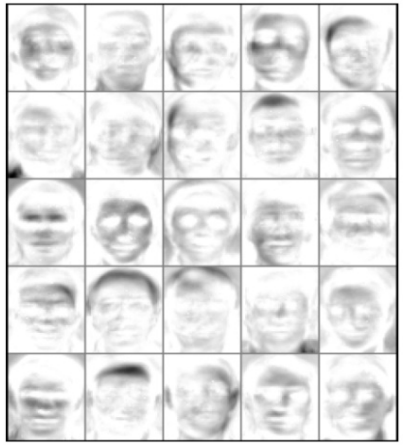


- Benefits:
  - Discovering "latent profile" in input matrix
  - Recommending new products to customers
    e.g., P5→C2(0.4); P4→C1 (0.3)
  - More

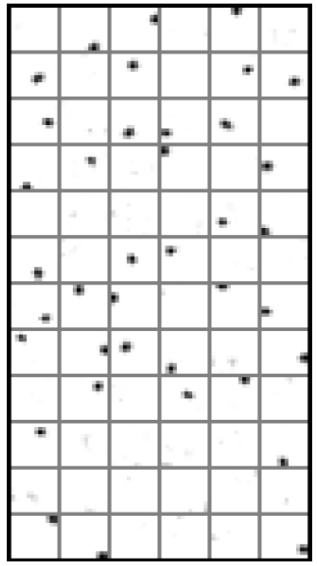
Example 2: images



$\approx$



$\times$



$\mathbf{X}$

input images

$\mathbf{U}$

face patterns

$\mathbf{V}^{\mathrm{T}}$

images represented with patterns

Example 3: documents

Encyclopedia entry:
'Constitution of the United States'

president (148)
congress (124)
power (120)
united (104)
constitution (81)
amendment (71)
government (57)
law (49)

$\approx$

court government council culture supreme constitutional rights justice	president served governor secretary senate congress presidential elected
flowers leaves plant perennial flower plants growing annual	disease behaviour glands contact symptoms skin pain infection

$\times$



X
input documents

U
latent topics

$\mathbf{V}^{\mathrm{T}}$
docs represented with topics

Example 4: videos

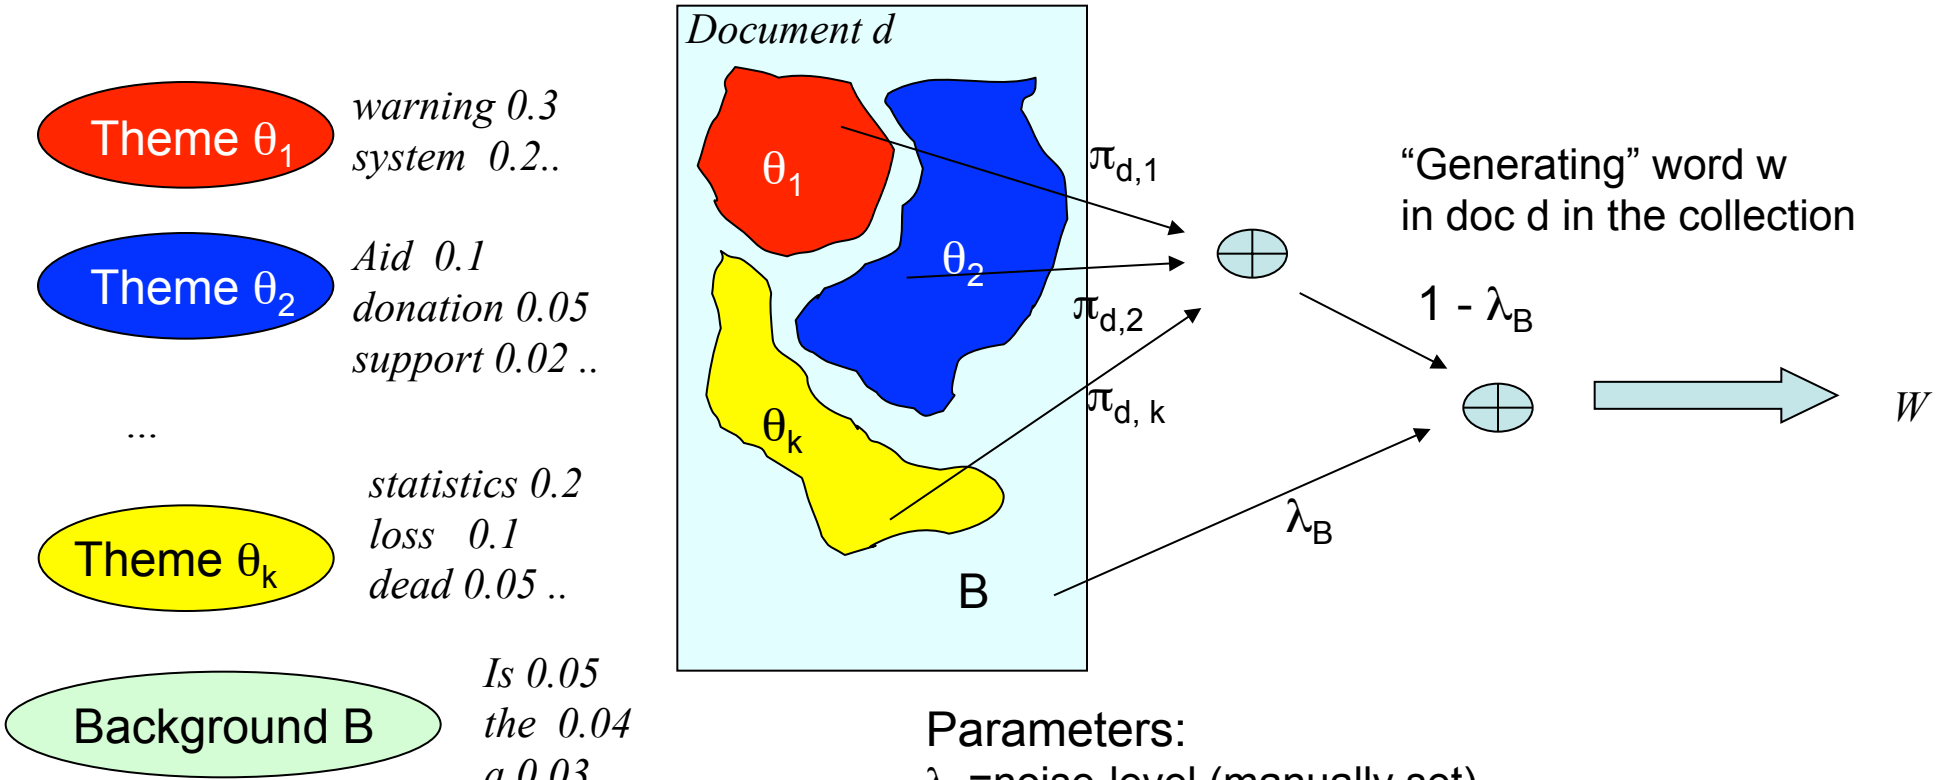
Video D = Low-rank appx. A + Sparse error E



PLSI as a Mixture Model

$$p_d(w) = \lambda_B p(w \mid \theta_B) + (1-\lambda)\sum_{j=1}^{k} \pi_{d,j} p(w \mid \theta_j)$$

$$\log p(d) = \sum_{w \in V} c(w,d) \log[\lambda_B p(w \mid \theta_B) + (1-\lambda)\sum_{j=1}^{k} \pi_{d,j} p(w \mid \theta_j)]$$



Parameters:

λ_B=noise-level (manually set)

θ's and π's are estimated with Maximum Likelihood

But how can we obtain the Eqs?

Maximum-likelihood

- Definition
  - We have a density function $P(x|\Theta)$ that is govened by the set of parameters Θ, e.g., P might be a set of Gaussians and Θ could be the means and covariances
  - We also have a data set $X=\{x_1,\ldots,x_N\}$, supposedly drawn from this distribution P, and *a*ssume these data vectors are i.i.d. with P.
  - Then the log-likehihood function is:

$$L(\Theta \mid X) = \log p(X \mid \Theta) = \log \prod_i p(x_i \mid \Theta) = \sum_i \log p(x_i \mid \Theta)$$

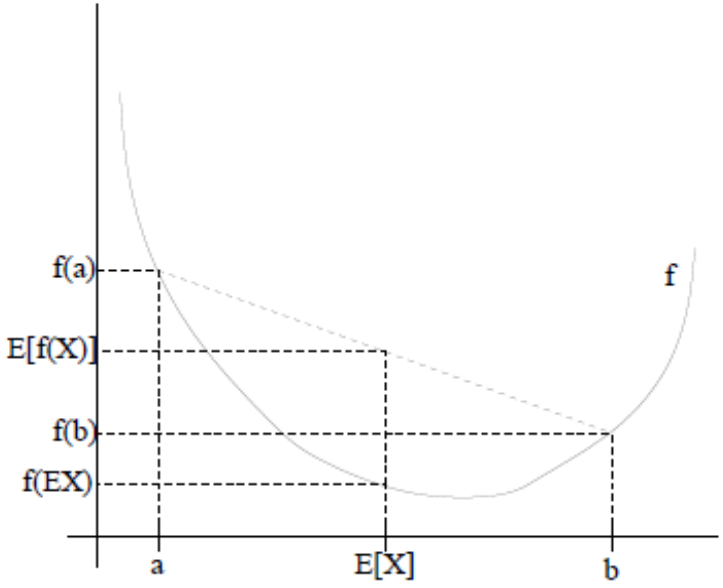
  - The log-likelihood is thought of as a function of the parameters Θ where the data X is fixed. Our goal is to find the Θ that maximizes L. That is

$$\Theta^* = \arg\max_{\Theta} L(\Theta \mid X)$$

Jensen's Inequality

- Recall that f is a convex function if f''(x)≥0, and f is strictly convex function if if f''(x)>0
- Let f be a convex function, and let X be a random variable, then:

$$E[f(X)] \geq f(EX)$$



- Moreover, if f is strictly convex, then [illegible] X=E[x] with probability 1 (i.e., if X is a constant)

Basic EM Algorithm

- However, Optimizing the likelihood function is analytically intractable but when the likelihood function can be simplified by assuming the existence of and values for additional but missing (or hidden) parameters:

$$L(\Theta \mid X) = \sum_i \log p(x_i \mid \Theta) = \sum_i \log \sum_z p(x_i, z \mid \Theta)$$

- Maximizing L(Θ) explicitly might be difficult, and the strategy is to instead repeatedly construct a lower-bound on L(E-step), and then optimize that lower bound (M-step).
  - For each i, let Q_i be some distribution over z ($\sum_z Q_i(z)=1$, $Q_i(z)\geq 0$), then

$$\sum_i \log \sum_{z^{(i)}} p(x^{(i)}, z^{(i)}; \Theta) = \sum_i \log \sum_{z^{(i)}} Q_i(z^{(i)}) \frac{p(x^{(i)}, z^{(i)}; \Theta)}{Q_i(z^{(i)})} \geq \sum_i \sum_{z^{(i)}} Q_i(z^{(i)}) \log \frac{p(x^{(i)}, z^{(i)}; \Theta)}{Q_i(z^{(i)})}$$

  - The above derivation used Jensen's inequality. Specifically, f(x) = logx is a concave function, since f''(x)=$-1/x^2<0$

Basic EM Algorithm (cont')

- In the summation is just an expectation of the quantity $[p(x^{(i)},z^{(i)}|\Theta)/Q_i(z^{(i)})]$ with respect to $z^{(i)}$ drawn according to the distribution given by Q_i. By Jensen's inequality, we have

$$f(E_{z^{(i)}\sim Q_i}[\frac{p(x^{(i)},z^{(i)};\Theta)}{Q_i(z^{(i)})}]) \geq E_{z^{(i)}\sim Q_i}[f(\frac{p(x^{(i)},z^{(i)};\Theta)}{Q_i(z^{(i)})})]$$

  - Now, for any set of distributions Q_i, we gives a lower-bound on $L(\Theta)$.
- There are many possible choices for the Q_i. Which should we choose?
  - Have current guess Θ, try to make the lower-bound tight at the value of Θ.
  - To make the bound tight for a particular value of Θ, we need Jensen's inequality to hold with equality. we require that X is a constant, i.e.,

$$\frac{p(x^{(i)},z^{(i)};\Theta)}{Q_i(z^{(i)})} = c$$

  - Since we know that $\sum_z Q_i(z)=1$, this further tells us that

$$Q_i(z^{(i)}) = \frac{p(x^{(i)},z^{(i)};\Theta)}{\sum_z p(x^{(i)},z;\Theta)} = \frac{p(x^{(i)},z^{(i)};\Theta)}{p(x^{(i)};\Theta)} = p(z^{(i)}\,|\,x^{(i)};\Theta)$$

  - Thus, we simply set the Q_i's to be the posterior distribution of the $z^{(i)}$'s given $x^{(i)}$ and the setting of the parameters Θ

Basic EM Algorithm (cont')

Repeat until convergence {

(E-step) For each i, set

$$Q_i(z^{(i)}) = p(z^{(i)} \mid x^{(i)}; \Theta)$$

(M-step) Set

$$\Theta := \arg\max_{\Theta} \sum_i \sum_{z^{(i)}} Q_i(z^{(i)}) \log \frac{p(x^{(i)}, z^{(i)}; \Theta)}{Q_i(z^{(i)})}$$

}

Proof of Convergence

- Suppose $\Theta^{(t)}$ and $\Theta^{(t+1)}$ are the parameters from two successive iterations of EM. We will now prove that $L(\Theta^{(t)}) \leq L(\Theta^{(t+1)})$, which shows EM always monotonically improves the log-likelihood.

$$L(\Theta^{(t+1)}) \geq \sum_i \sum_{z^{(i)}} Q_i(z^{(i)}) \log \frac{p(x^{(i)}, z^{(i)}; \Theta^{(t+1)})}{Q_i(z^{(i)})} \qquad (1)$$

$$\geq \sum_i \sum_{z^{(i)}} Q_i(z^{(i)}) \log \frac{p(x^{(i)}, z^{(i)}; \Theta^{(t)})}{Q_i(z^{(i)})} \qquad (2)$$

$$= L(\Theta^{(t)})$$

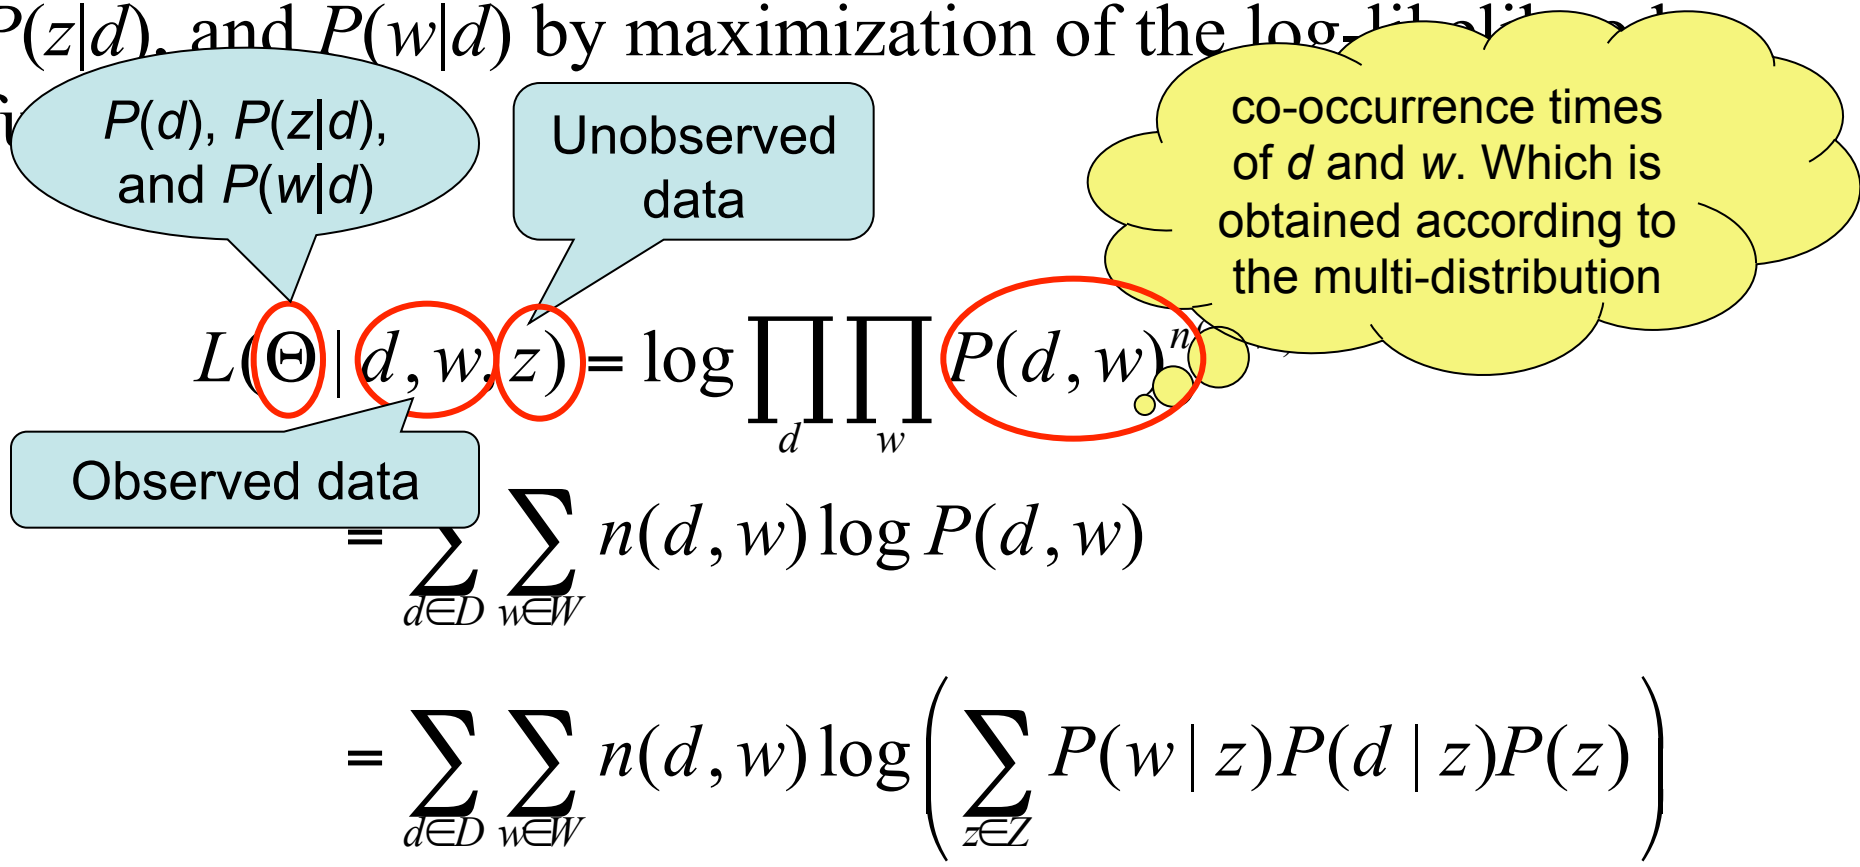
  - Inequality (1) comes from the fact that

$$L(\Theta) \geq \sum_i \sum_{z^{(i)}} Q_i(z^{(i)}) \log \frac{p(x^{(i)}, z^{(i)}; \Theta)}{Q_i(z^{(i)})}$$

    - Holds for any values of Q_i and Θ, and in particular holds for $Q_i = Q_i^{(t)}$, $\Theta = \Theta^{(t+1)}$.
  - Inequality (2) used the fact of M-step that $\Theta^{(t+1)}$ is chosen to maximize the lower bound.

PLSI Model

- Following the likelihood principle, we determines $P(d)$, $P(z|d)$, and $P(w|d)$ by maximization of the log-li
fu



$$L(\Theta \mid d, w, z) = \log \prod_{d} \prod_{w} P(d,w)^{n(d,w)}$$

$$= \sum_{d \in D} \sum_{w \in W} n(d,w) \log P(d,w)$$

$$= \sum_{d \in D} \sum_{w \in W} n(d,w) \log \left(\sum_{z \in Z} P(w \mid z) P(d \mid z) P(z) \right)$$

Parameter Estimation-Using EM

- According to Basic EM:

$$Q_i(z^{(i)}) = p(z^{(i)} \mid x^{(i)}; \Theta)$$

- Then we define

$$Q_i(z^{(i)}) = p(z \mid d, w)$$

- Thus according to Jensen's inequality

$$L(\Theta) = \sum_{d \in D} \sum_{w \in W} n(d,w) \log \sum_{z \in Z} p(z \mid d, w) \frac{p(w \mid z) p(d \mid z) p(z)}{p(z \mid d, w)}$$

$$\geq \sum_{d \in D} \sum_{w \in W} n(d,w) \sum_{z \in Z} p(z \mid d, w) \log \frac{p(w \mid z) p(d \mid z) p(z)}{p(z \mid d, w)}$$

(1)Solve $P(w|z)$

- We introduce Lagrange multiplier λwith the constraint that $\sum_w P(w|z)=1$, and solve the following equation:

$$\frac{\partial}{\partial P(w|z)}\left\{\sum_{d\in D}\sum_{w\in W} n(d,w)\sum_{z\in Z} p(z|d,w)\log\frac{p(w|z)p(d|z)p(z)}{p(z|d,w)}+\lambda\left[\sum_z P(w|z)-1\right]\right\}=0$$

$$\therefore \frac{\sum_{d\in D} n(d,w)P(z|d,w)}{P(w|z)}+\lambda=0,$$

$$\therefore P(w|z)=-\frac{\sum_{d\in D} n(d,w)P(z|d,w)}{\lambda},$$

$$\sum_w P(w|z)=1,$$

$$\therefore \lambda=-\sum_{w\in W}\sum_{d\in D} n(d,w)P(z|d,w),$$

$$\therefore P(w|z)=\frac{\sum_{d\in D} n(d,w)P(z|d,w)}{\sum_{w\in W}\sum_{d\in D} n(d,w)P(z|d,w)}$$

(2)Solve $P(d|z)$

- We introduce Lagrange multiplier λ with the constraint that $\sum_d P(d|z)=1$, and get the following result:

$$\therefore P(d \mid z) = \frac{\sum_{w \in W} n(d,w) P(z \mid d,w)}{\sum_{d \in D} \sum_{w \in W} n(d,w) P(z \mid d,w)}$$

(3)Solve $P(z)$

- We introduce Lagrange multiplier λ with the constraint that $\sum_z P(z)=1$, and solve the following equation:

$$\frac{\partial}{\partial P(z)}\left\{\sum_{d\in D}\sum_{w\in W} n(d,w)\sum_{z\in Z} p(z\mid d,w)\log\frac{p(w\mid z)p(d\mid z)p(z)}{p(z\mid d,w)}+\lambda\left[\sum_z P(z)-1\right]\right\}=0$$

$$\therefore \frac{\sum_{d\in D}\sum_{w\in W} n(d,w)P(z\mid d,w)}{P(z)}+\lambda=0,$$

$$\therefore P(z)=-\frac{\sum_{d\in D}\sum_{w\in W} n(d,w)P(z\mid d,w)}{\lambda},$$

$$\sum_z P(z)=1,$$

$$\therefore \lambda=-\sum_{d\in D}\sum_{w\in W} n(d,w)\sum_z P(z\mid d,w)=-\sum_{d\in D}\sum_{w\in W} n(d,w),$$

$$\therefore P(z)=\frac{\sum_{d\in D}\sum_{w\in W} n(d,w)P(z\mid d,w)}{\sum_{w\in W}\sum_{d\in D} n(d,w)}$$

(4)Solve $P(z|d,w)$ – M step

$$P(z \mid d,w) = \frac{P(d,w,z)}{P(d,w)}$$

$$= \frac{P(w,d \mid z)P(z)}{P(d,w)}$$

$$= \frac{P(w \mid z)P(d \mid z)P(z)}{\sum_{z \in Z} P(w \mid z)P(d \mid z)P(z)}$$

The final update Equations

- E-step:

$$P(z \mid d, w) = \frac{P(w \mid z)P(d \mid z)P(z)}{\sum_{z \in Z} P(w \mid z)P(d \mid z)P(z)}$$

- M-step:

$$P(w \mid z) = \frac{\sum_{d \in D} n(d, w)P(z \mid d, w)}{\sum_{w \in W} \sum_{d \in D} n(d, w)P(z \mid d, w)}$$

$$P(d \mid z) = \frac{\sum_{w \in W} n(d, w)P(z \mid d, w)}{\sum_{d \in D} \sum_{w \in W} n(d, w)P(z \mid d, w)}$$

$$P(z) = \frac{\sum_{d \in D} \sum_{w \in W} n(d, w)P(z \mid d, w)}{\sum_{w \in W} \sum_{d \in D} n(d, w)}$$

Coding Design

- Variables:
  - **double**[][] *p_dz_n* // *p(d|z)*, $|D|*|Z|$
  - **double**[][] *p_wz_n* // *p(w|z)*, $|W|*|Z|$
  - **double**[] *p_z_n* // *p(z)*, $|Z|$
- Running Processing:
  1. Read dataset from file

     ArrayList<DocWordPair> doc; // all the docs

     DocWordPair – (word_id, word_frequency_in_doc)
  2. Parameter Initialization

     Assign each elements of *p_dz_n, p_wz_n and p_z_n* with a random double value, satisfying $\sum_d p_dz_n=1$, $\sum_d p_wz_n=1$, and $\sum_d p_z_n=1$
  3. Estimation (Iterative processing)
     1. Update *p_dz_n, p_wz_n and p_z_n*
     2. Calculate Log-likelihood function to see where (|Log-likelihood – old_Log-likelihood| < threshold)
  4. Output *p_dz_n, p_wz_n and p_z_n*

Coding Design

- **Update *p_dz_n***

```
For each doc d{
   For each word w included in d   {
       denominator = 0;
       nominator = new double[Z];
       For each topic z   {
          nominator[z] = p_dz_n[d][z]* p_wz_n[w][z]* p_z_n[z]
           denominator  +=nominator[z];
       } // end for each topic z
       For each topic z    {
          P_z_condition_d_w = nominator[j]/denominator;
          nominator_p_dz_n[d][z] += tf_wd*P_z_condition_d_w;
          denominator_p_dz_n[z] += tf_wd*P_z_condition_d_w;
       } // end for each topic z
    }// end for each word w included in d
}// end for each doc d

For each doc d {
     For each topic z                {
        p_dz_n_new[d][z] = nominator_p_dz_n[d][z]/ denominator_p_dz_n[z];
     } // end for each topic z
}// end for each doc d
```

Coding Design

- **Update *p_wz_n***

```
For each doc d{
   For each word w included in d   {
       denominator = 0;
       nominator = new double[Z];
       For each topic z   {
          nominator[z] = p_dz_n[d][z]* p_wz_n[w][z]* p_z_n[z]
           denominator  +=nominator[z];
        } // end for each topic z
        For each topic z    {
          P_z_condition_d_w = nominator[j]/denominator;
          nominator_p_wz_n[w][z] += tf_wd*P_z_condition_d_w;
           denominator_p_wz_n[z] += tf_wd*P_z_condition_d_w;
        } // end for each topic z
    }// end for each word w included in d
}// end for each doc d

For each w {
     For each topic z                {
        p_wz_n_new[w][z] = nominator_p_wz_n[w][z]/ denominator_p_wz_n[z];
     } // end for each topic z
}// end for each doc d
```

Coding Design

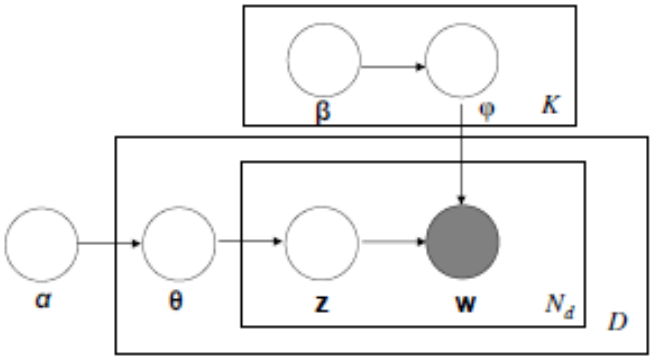
- **Update *p_z_n***

```
For each doc d{
   For each word w included in d   {
      denominator = 0;
      nominator = new double[Z];
      For each topic z   {
         nominator[z] = p_dz_n[d][z]* p_wz_n[w][z]* p_z_n[z]
          denominator  +=nominator[z];
      } // end for each topic z
      For each topic z    {
         P_z_condition_d_w = nominator[j]/denominator;
         nominator_p_z_n[z] += tf_wd*P_z_condition_d_w;
      } // end for each topic z
      denominator_p_z_n[z] += tf_wd;
   }// end for each word w included in d
}// end for each doc d

For each topic z{
       p_dz_n_new[d][j] = nominator_p_z_n[z]/ denominator_p_z_n;
} // end for each topic z
```

Latent Dirichlet Allocation (LDA)

Model



Generative Process:

For each document d in the corpus, the LDA model

1. picks a multinomial distribution $\theta_d = [\theta_{d1} \ldots \theta_{dK}]^T$ from the dirichlet distribution $\alpha = [\alpha_1 \ldots \alpha_K]^T$
2. assigns a topic $z_{di} = k$ to the i^{th} word in the document d according to the multinomial distribution θ d .
3. given a topic zdi=k , the model then pick a word w_{di} from the vocabulary of V words according to the multinomial distribution $[\Phi_{k1} \ldots \Phi_{kV}]^T$ which is generated from the Dirichlet distribution $[\beta_{k1} \ldots \beta_{kV}]^T$

Joint Distribution

- Given the parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, the joint distribution of a topic mixture $\boldsymbol{\theta}$, a set of N topics $\mathbf{z}$, and a set of N words $\mathbf{w}$ is given by:

$$p(\boldsymbol{\theta}, \mathbf{z}, \mathbf{w} \mid \boldsymbol{\alpha}) = p(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) \prod_{n=1}^{N} p(z_n \mid \boldsymbol{\theta}) p(w_n \mid z_n, \boldsymbol{\beta})$$

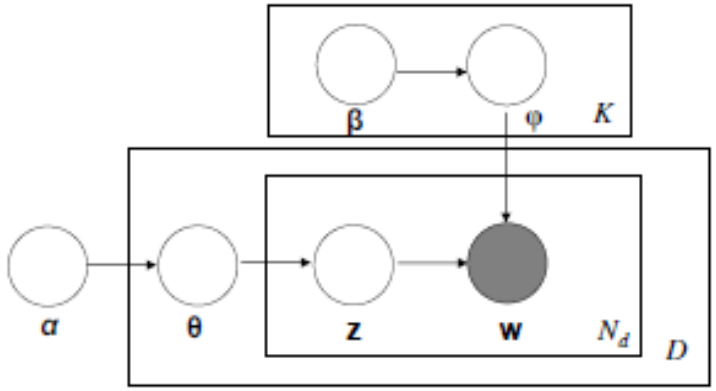
- where $p(z_n|\boldsymbol{\theta})$ is simply $\boldsymbol{\theta}_i$ for the unique i such that $z_n^i=1$. Integrating over $\boldsymbol{\theta}$ and summing over $\mathbf{z}$ we obtain the marginal distribution of a document

$$p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \int p(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) \left(\prod_{n=1}^{N} \sum_{z_n} p(z_n \mid \boldsymbol{\theta}) p(w_n \mid z_n, \boldsymbol{\beta}) \right) d\boldsymbol{\theta}$$

- Taking the products of the marginal probabilities of single documents, we obtain the probability of a corpus:

$$p(D \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \prod_{d=1}^{M} \int p(\boldsymbol{\theta}_d \mid \boldsymbol{\alpha}) \left(\prod_{n=1}^{N_d} \sum_{z_{nd}} p(z_{dn} \mid \boldsymbol{\theta}_d) p(w_{dn} \mid z_{dn}, \boldsymbol{\beta}) \right) d\boldsymbol{\theta}_d$$

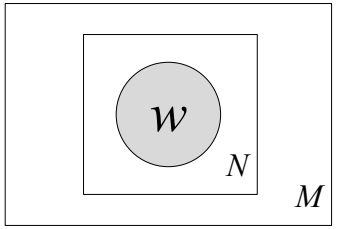
Three-level Representation



Similar model can be referred to hierarchical models (Gelman et al., 1995), or more precisely as conditionally independent hierarchical models (Kass and Steffey, 1989).

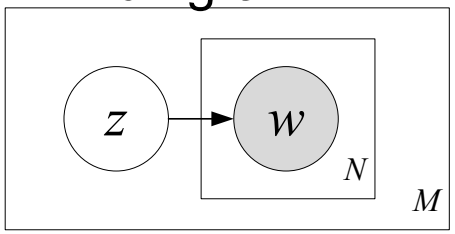
- **α ,β,** Φ are corpus-level parameters
  - Sampled once when generating a corpus
- **θ** $_d$ are document-level variables
  - Sampled once per document
- z_{dn},w_{dn} are word-level variables
  - Sampled once for each word in each document

Relationships with Other Latent Models



unigram

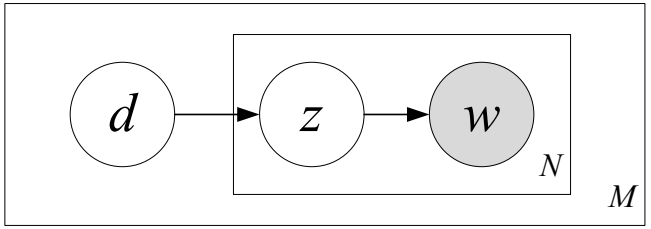
$$p(\mathbf{w}) = \prod_{n=1}^{N} p(w_n)$$



Mixture of unigrams

$$p(\mathbf{w}) = \sum_{z} p(z) \prod_{n=1}^{N} p(w_n \mid z)$$

Problems: Each document exhibits exactly one topic. This assumption is often too limiting to effectively model a large collection of documents.



PLSI/aspect model

$$p(d, w_n) = p(d) \sum_{z} p(w_n \mid z) p(z \mid d)$$

Problems: (1) $p(z|d)$ only model the documents in the training data set and cannot for the unseen document; and (2) the parameter $kV+kM$ grows linearly in M and thus overfitting.

$$\text{In } p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \int p(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) \left(\prod_{n=1}^{N} \sum_{z_n} p(z_n \mid \boldsymbol{\theta}) p(w_n \mid z_n, \boldsymbol{\beta}) \right) d\boldsymbol{\theta} \text{ on}$$

- The key inferential problem is that of computing the posterior distribution of the hidden variables given a document:

$$p(\boldsymbol{\theta}, \mathbf{z} \mid \mathbf{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \frac{p(\boldsymbol{\theta}, \mathbf{z}, \mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta})}{p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta})} \qquad \text{(a)}$$

- However, such distribution is intractable to compute in general. For normalization in the above distribution, we have to marginalize over the hidden variables and write the Equation (a) in terms of the model parameters:

$$p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \frac{\Gamma(\sum_i \boldsymbol{\alpha}_i)}{\prod_i \Gamma(\boldsymbol{\alpha}_i)} \int \left(\prod_{i=1}^{k} \boldsymbol{\theta}_i^{\boldsymbol{\alpha}_i - 1} \right) \left(\prod_{n=1}^{N} \sum_{i=1}^{k} \prod_{j=1}^{V} (\boldsymbol{\theta}_i \boldsymbol{\beta}_{ij})^{w_n^j} \right) d\boldsymbol{\theta}$$

Inference and Parameter Estimation (cont.)

$$p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \frac{\Gamma(\sum_i \boldsymbol{\alpha}_i)}{\prod_i \Gamma(\boldsymbol{\alpha}_i)} \int \left(\prod_{i=1}^{k} \boldsymbol{\theta}_i^{\boldsymbol{\alpha}_i - 1} \right) \left(\prod_{n=1}^{N} \sum_{i=1}^{k} \prod_{j=1}^{V} (\boldsymbol{\theta}_i \boldsymbol{\beta}_{ij})^{w_n^j} \right) d\boldsymbol{\theta}$$

- This function is intractable due to the coupling between $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ in the summation over latent topics (Dickey, 1983).
- Rather than the intractable exact inference, we can use some other approximate inference algorithms, e.g., Laplace approximation, variational approximation, and Markov chain Monte Carlo (Jordan, 1999).

Gibbs Sampling Derivation for LDA

Notation Definition

- D: #documents in the whole corpus
- K: #topics
- V: #vocabulary

- **w, z** : bold font means words or topics in the whole corpus

- d: the index of document, d = 1 to D
- N_d: #words in d^{th} document
- i: the index of word in a document, i=1 to N_d
  - w_{di} : i^{th} word of d^{th} document
  - z_{di} : i^{th} word of d^{th} document is denoted by topic z_{di}
  - d_{di} : i^{th} word of d^{th} document if from document d_{di}
  - Note that for different i, j, w_{di} and w_{dj} may indicate both the same word and different words, the same as z_{di} and d_{di}

- For the whole corpus, there are totally $D*N_d$ words and $D*N_d$ topics. Actually, the whole corpus contains V words and T topic without repeat, so
  - z : the index of topic without repeat, z=1 to K
  - v: the index of word without repeat, v=1 to V

- $n_{z,v}$: the cooccurrence of z^{th} topic and v^{th} word (i.e., the number of times word v is generated by topic z)
- $n_{d,z}$: the cooccurrence of d^{th} document and z^{th} topic (i.e., the number of times topic z occurs in document d)

- Φ_{zdi} : the topic-words distribution corresponding to topic z_{di}
- Φ_{z} : the topic-words distribution corresponding to topic z
- θ_d: the document-topic distribution corresponding to document d
- β: K*V matrix, $\Phi \sim Dirichlet(\beta)$
- α: K vector, $\theta \sim Dirichlet(\alpha)$

Gibbs Sampling

- The key idea of the Gibbs sampling is to sequentially update each variable of the distribution according the conditional probability given all the other variables.
- Take a distribution of the three variable (X, Y, Z) for example, the Gibbs sampler makes the Markov chain transit from (x_i, y_i, z_i) to $(x_{i+1}, y_{i+1}, z_{i+1})$ by the following steps:
  - Draw x_{i+1} from $p(X|Y=y_i,Z=z_i)$
  - Draw y_{i+1} from $p(Y|X=x_{i+1},Z=z_i)$
  - Draw z_{i+1} from $p(Z|X=x_{i+1},Y=y_{i+1})$

Gibbs Sampling Derivation for LDA

- To apply Gibbs Sampling, we need the full conditional distribution $p(z_{di}|\mathbf{w},\mathbf{z_{-di}})$. Then w_{di} and z_{di} can be sampled based on the probability. When the sampling completes, each w_{di} has assigned a topic, and each document also has assigned topics. So the document-topics distribution θ_d and topic-words distribution Φ_z can be easily estimated.
- According to $p(z_{di}, \mathbf{z_{-di}}) = p(\mathbf{z})$ and bayesian formula, we can get:

$$p(z_{di} \mid \mathbf{w}, \mathbf{z}_{-di}) = \frac{p(\mathbf{z}, \mathbf{w})}{p(\mathbf{z}_{-di}, \mathbf{w})} \propto \frac{p(\mathbf{z}, \mathbf{w})}{p(\mathbf{z}_{-di}, \mathbf{w}_{-di})}$$

w_{di} is generated by z_{di}, so we can ignore it if z_{di} is not considered.

- The key point is to calculate the joint distribution:

$$p(\mathbf{z}, \mathbf{w}) = p(\mathbf{w} \mid \mathbf{z}, \beta) p(\mathbf{z} \mid \alpha)$$

Then we calculate $p(\mathbf{w}|\mathbf{z},\beta)$ and $p(\mathbf{z}|\alpha)$ respectively.

$p(\mathbf{z}|\alpha)$

Import Θ:

$$p(\mathbf{z}\,|\,\alpha)=\int p(\mathbf{z}\,|\,\Theta)p(\Theta\,|\,\alpha)d\Theta \tag{1}$$

Expand all the bold and capital character:

$$p(\mathbf{z}\,|\,\alpha)=\int\prod_{d=1}^{D}\prod_{i=1}^{N_d}p(z_{di}\,|\,\theta_d)p(\theta_d\,|\,\alpha)d\theta_d \tag{2}$$

Because z_n~Multinomial(θ) and θ~Dir(α):

$$p(\mathbf{z}\,|\,\alpha)=\int\prod_{d=1}^{D}\prod_{z=1}^{K}\theta_{d,z}^{n_{d,z}}\prod_{d=1}^{D}\left(\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}\prod_{z=1}^{K}\theta_{d,z}^{\alpha_z-1}\right)d\theta_d \tag{3}$$

Because $\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}$ is unrelated with θ:

$$p(\mathbf{z}\,|\,\alpha)=\left(\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}\right)^{D}\int\prod_{d=1}^{D}\prod_{z=1}^{K}\theta_{d,z}^{n_{d,z}+\alpha_z-1}d\theta_d \tag{4}$$

For different document, its document-topics distribution θ_d is independent, so:

$$p(\mathbf{z}\,|\,\alpha)=\left(\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}\right)^{D}\prod_{d=1}^{D}\int\prod_{z=1}^{K}\theta_{d,z}^{n_{d,z}+\alpha_z-1}d\theta_d \tag{5}$$

According to the beta function:

$$B(\alpha)=\frac{\prod_{z=1}^{K}\Gamma(\alpha_z)}{\Gamma(\sum_{z=1}^{K}\alpha_z)}=\int_{\sum_{z=1}^{K}\theta_z=1}\prod_{z=1}^{K}\theta_z^{\alpha_z-1}d\theta$$

We get

$$p(\mathbf{z}\,|\,\alpha)=\left(\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}\right)^{D}\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))} \tag{6}$$

$p(\mathbf{z}|\alpha)\ /p(\mathbf{z}_{-di}|\alpha)$

$$\frac{p(\mathbf{z}\mid\alpha)}{p(\mathbf{z}_{-di}\mid\alpha)}=\frac{\left(\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}\right)^{D}\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}}{\left[\left(\frac{\Gamma(\sum_{z=1}^{K}\alpha_z)}{\prod_{z=1}^{K}\Gamma(\alpha_z)}\right)^{D}\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}\right]_{-di}}=\frac{\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}}{\left[\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}\right]_{-di}} \tag{7}$$

Ignoring di means ignoring z_{di} and w_{di}. That is to say, when considering the times of co occurrence of topic z and word v ($n_{z,v}$), ignore this cooccurrence of z_{di} and w_{di}. This makes $n_{zdi,wdi}$-1, $n_{ddi,zdi}$-1, however, do not impact $n_{z,\,wdi}$and $n_{d,zdi}$

$$\frac{\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}}{\left[\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}\right]_{-di}}=\frac{\Gamma(n_{d_{di},z_{di}}+\alpha_{z_{di}})\times\left[\prod_{d=1}^{D}\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)\right]_{-d_{di},z_{di}}}{\Gamma(\sum_{z=1}^{K}(n_{d_{di},z}+\alpha_z))\times\prod_{d\neq d_{di}}^{D}\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}$$
$$\frac{\Gamma(n_{d_{di},z_{di}}+\alpha_{z_{di}}-1)\times\left[\prod_{d=1}^{D}\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)\right]_{-d_{di},z_{di}}}{\Gamma(\sum_{z=1}^{K}(n_{d_{di},z}+\alpha_z)-1)\times\prod_{d\neq d_{di}}^{D}\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))} \tag{8}$$

Because : $\Gamma(x+1)=x\Gamma(x)$

We get:

$$\frac{\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}}{\left[\prod_{d=1}^{D}\frac{\prod_{z=1}^{K}\Gamma(n_{d,z}+\alpha_z)}{\Gamma(\sum_{z=1}^{K}(n_{d,z}+\alpha_z))}\right]_{-di}}=\frac{n_{d_{di},z_{di}}+\alpha_{z_{di}}-1}{\sum_{z=1}^{K}(n_{d_{di},z}+\alpha_z)-1} \tag{9}$$

The Full Conditional Probability

The same as (1)-(9), we can get:

$$\frac{p(\mathbf{w}\mid\mathbf{z},\beta)}{p(\mathbf{w}_{-di}\mid\mathbf{z}_{-di},\beta)}=\frac{n_{z_{di},w_{di}}+\beta_{w_{di}}-1}{\sum_{v=1}^{V}(n_{z_{di},v}+\beta_v)-1} \tag{10}$$

Finally, the full conditional probability is :

$$\begin{aligned} p(z_{di}\mid\mathbf{w},\mathbf{z}_{-di}) \propto \frac{p(\mathbf{z},\mathbf{w})}{p(\mathbf{z}_{-di},\mathbf{w}_{-di})} &= \frac{p(\mathbf{w}\mid\mathbf{z},\beta)p(\mathbf{z}\mid\alpha)}{p(\mathbf{w}_{-di}\mid\mathbf{z}_{-di},\beta)p(\mathbf{z}_{-di}\mid\alpha)} \\ &= \frac{n_{z_{di},w_{di}}+\beta_{w_{di}}-1}{\sum_{v=1}^{V}(n_{z_{di},v}+\beta_v)-1}\cdot\frac{n_{d_{di},z_{di}}+\alpha_{z_{di}}-1}{\sum_{z=1}^{K}(n_{d_{di},z}+\alpha_z)-1} \end{aligned} \tag{11}$$

Then we can use the full conditional probability to sample z_{di}

Finally

Finally, we need to calculate the topic of a new observation (word *w*), hence we can evaluate the distribution of a new topic-word pair (z, w) by

$$p(z = k, w = t \mid \mathbf{w}, \mathbf{z}) = \frac{p(z = k, w = t, \mathbf{w}, \mathbf{z})}{p(\mathbf{w}, \mathbf{z})}$$

$$= \frac{\frac{\Gamma(n_{z_{di},w_{di}} + 1 + \beta_{w_{di}})}{\Gamma(\sum_{v=1}^{V}(n_{z_{di},v} + 1 + \beta_v))} \cdot \frac{\Gamma(n_{d_{di},z_{di}} + 1 + \alpha_{z_{di}})}{\Gamma(\sum_{z=1}^{K}(n_{d_{di},z} + 1 + \alpha_z))}}{\frac{\Gamma(n_{z_{di},w_{di}} + \beta_{w_{di}})}{\Gamma(\sum_{v=1}^{V}(n_{z_{di},v} + \beta_v))} \cdot \frac{\Gamma(n_{d_{di},z_{di}} + \alpha_{z_{di}})}{\Gamma(\sum_{z=1}^{K}(n_{d_{di},z} + \alpha_z))}}$$

$$= \frac{n_{z_{di},w_{di}} + \beta_{w_{di}}}{\sum_{v=1}^{V}(n_{z_{di},v} + \beta_v)} \cdot \frac{n_{d_{di},z_{di}} + \alpha_{z_{di}}}{\sum_{z=1}^{K}(n_{d_{di},z} + \alpha_z)}$$

Multinomial Parameters

During parameter estimation, the algorithm keeps track of a *D*x*K* (document by topic) count matrix, a *K*x*V* (topic by word) count matrix. Given these matrices, we can easily estimate the probability θ_d^z of a topic given a document and the probability Φ_z^v of a word given a topic by:

$$\phi_z^v = \frac{n_{z,v} + \beta_v - 1}{\sum_{v=1}^{V}(n_{z,v} + \beta_v) - 1}, \theta_d^z = \frac{n_{d,z} + \alpha_z - 1}{\sum_{z=1}^{K}(n_{d,z} + \alpha_z) - 1}$$

Given a new document *d*, we perform inference by

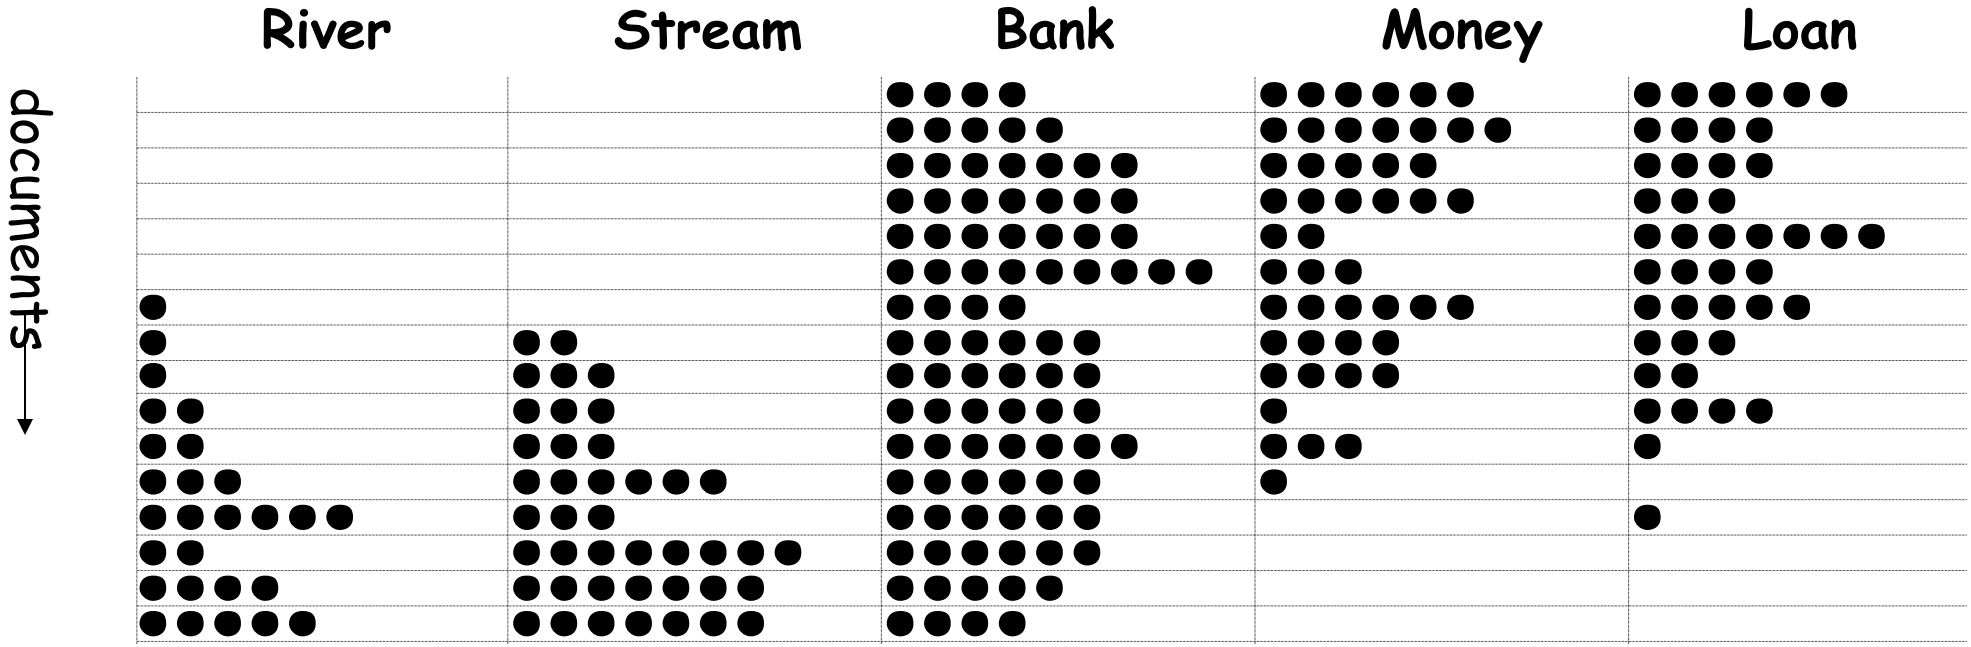
$$P(\mathbf{z}_i = j \mid \mathbf{z}_{-i}, w_i, d_i, \cdot) \propto \frac{n_k^{(t)} + n_k^{(t)'} + \beta_t}{\sum_{v=1}^{V}(n_k^{(v)} + n_k^{(t)'} + \beta_v)} \cdot \frac{n_m^{(k)} + \alpha_k}{\sum_{z=1}^{K}(n_m^z + \alpha_z)}$$

where *n'* denotes the number in the new document.

Coding

- Init()
- Estimate()
  - Foreach doc
    - Foreach word
      - Sampling()
        - » Compute probability $P(z_i|\mathbf{w}, \mathbf{z}_{-i})$;
        - » Sampling a topic *z* based on $P(z_i|\mathbf{w}, \mathbf{z}_{-i})$;
  - Compute_theta()
  - Compute_phi()

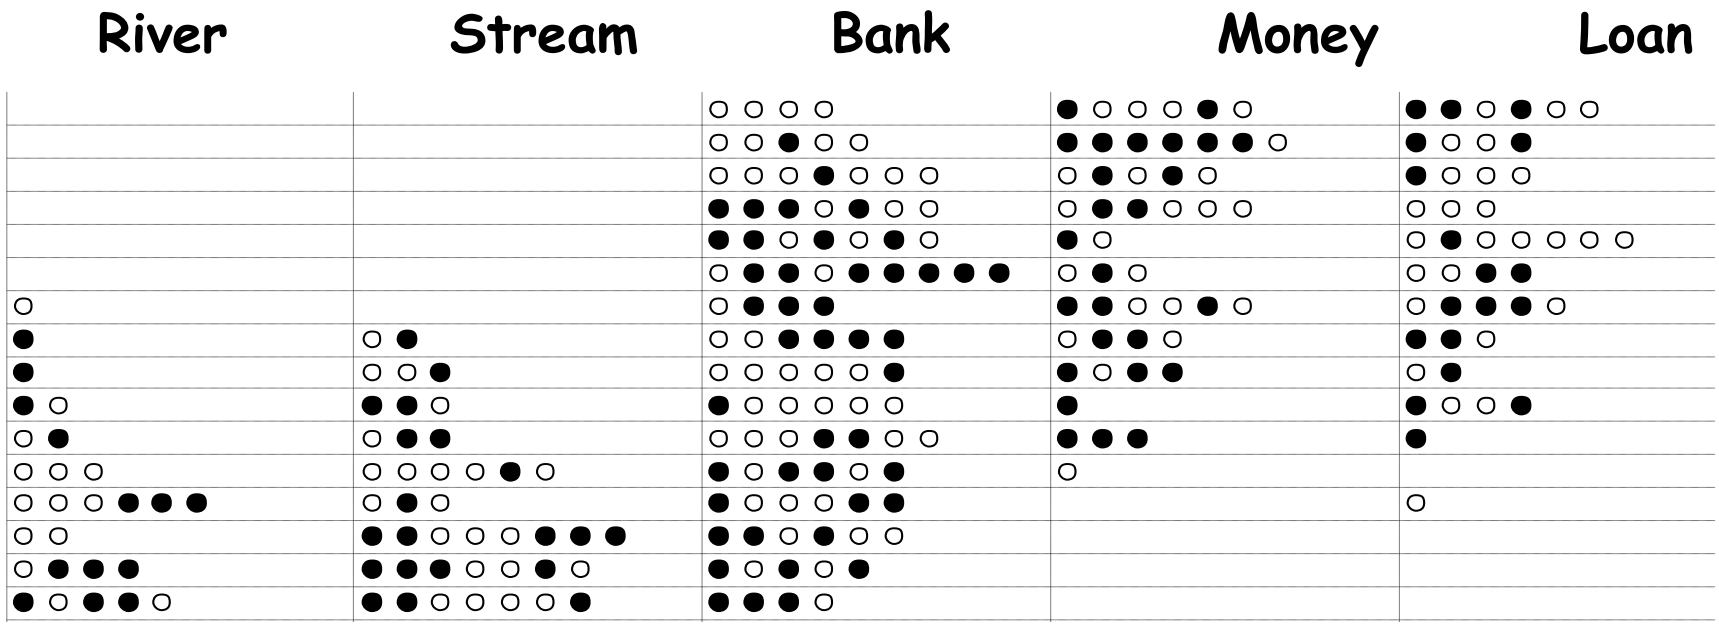
Topic Model: Inference



Can we recover the original topics and topic mixtures from this data?

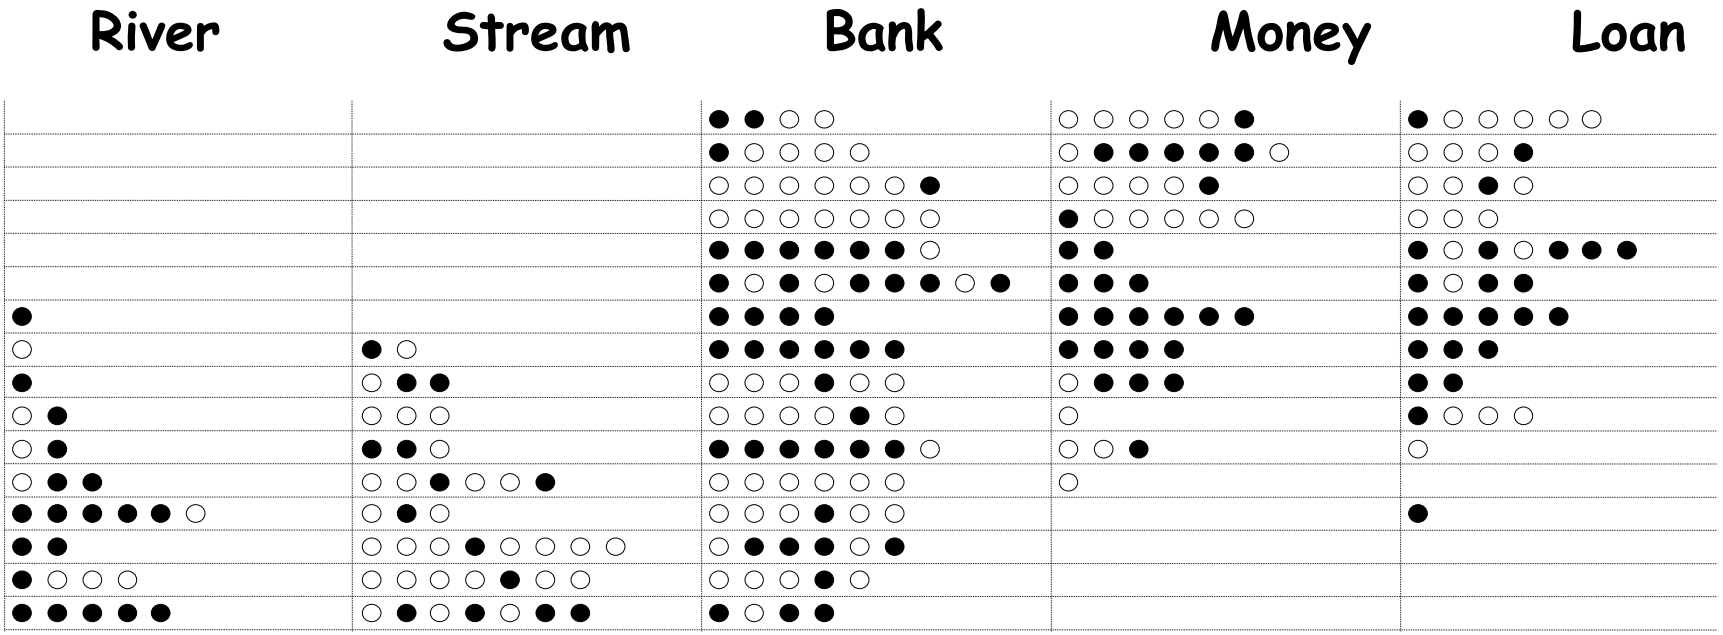
Example of Gibbs Sampling

- Assign word tokens randomly to topics (●=topic 1; ●=topic 2)

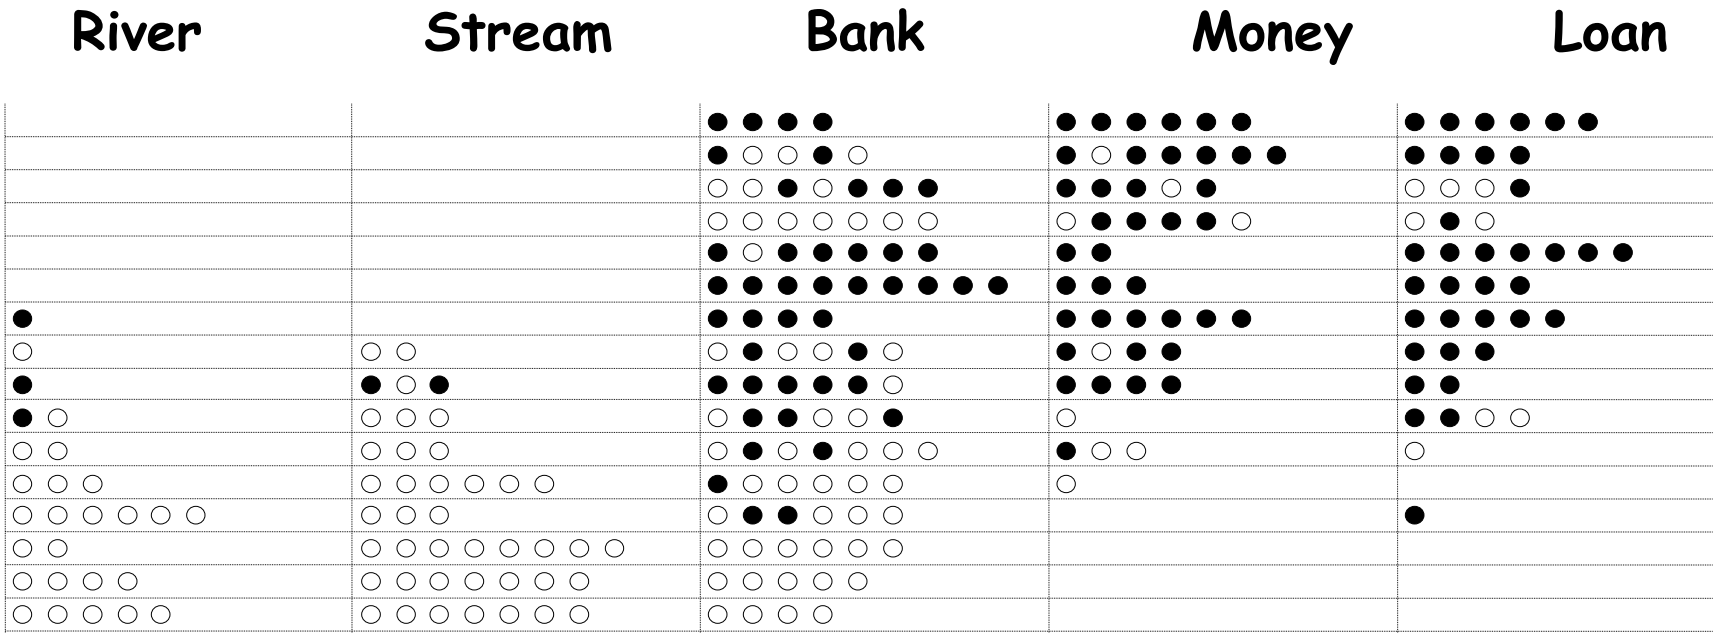


After 1 iteration

- Apply sampling equation to each word token



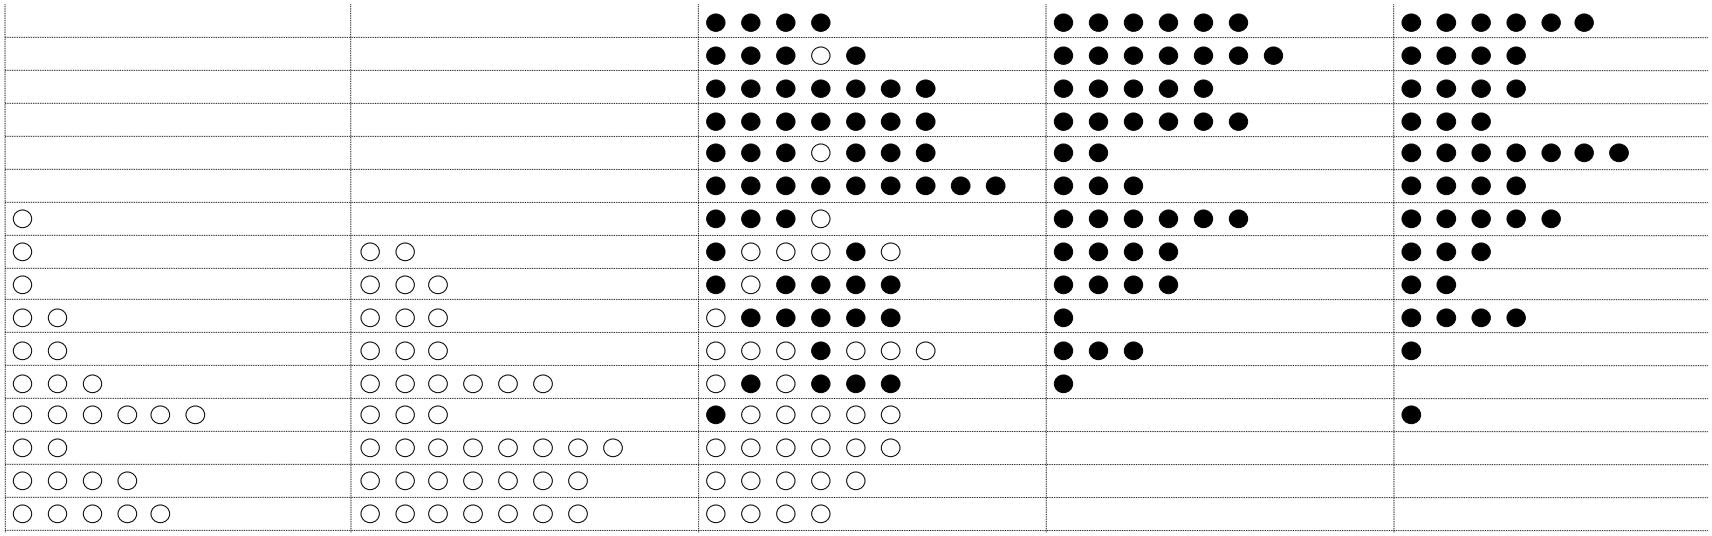
After 4 iterations



After 32 iterations

topic 1	
stream	.40
bank	.35
river	.25

topic 2	
bank	.39
money	.32
loan	.29



Parameter Estimation with Gibbs Sampling

Initialization:
  Randomly assign topic for each item in each document and initialize count matrices.
Iteration(MCMC Gibbs sampling step)
  For each Iteration
    For each document
      For each item
        assign a topic from full conditional probability with input of α and β
        update count matrices.
      End for
    End for
    Calculate document-topic and topic-item distribution with count matrices.
  End for

Hyper-parameter Estimation

- Stochastic EM (StEM) algorithm

```
Iteration(MCMC Gibbs sampling step)
  For each Iteration
    For each document
      For each item
        assign a topic from full conditional probability with
            input of α and β
        update count matrices.
      End for
    End for
    Estimate hyper-parameter α and β with input of updated
    count matrix.
End for
Calculate document-topic and topic-item distribution with count
matrices.
```

E-step

M-step

Hyper-parameter Estimation Methods

- Sampling Based Method

  Sampling new value of hyper-parameter from the posteriori distribution
  - Random Walk Metropolis-Hastings sampling
  - Adaptive Rejection Sampling

- MAP (Maximum A Posteriori) Based Method

  Obtain the new value by maximum a posterior probability using iteration methods.
  - Fixed-point iteration
  - Fixed-point by Leave-one-out Likelihood
  - Newton-Raphson iteration

The sampling based method

The full conditional for Gibbs sampling of α can be obtained by

$$p(\alpha|\vec{z}) \propto p(\vec{z}|\alpha)p(\alpha)$$
$$\propto \prod_m \frac{\Delta(\vec{n}_m + \alpha)}{\Delta_K(\alpha)} \alpha^{a-1} e^{-b\alpha}$$

Then for each parameter, we have

$$p(\alpha_i \mid \vec{\alpha}_{-i}, \vec{z}) = \frac{p(\alpha_i, \vec{\alpha}_{-i}, \vec{z})}{p(\vec{\alpha}_{-i}, \vec{z})} \propto p(\alpha_i, \vec{\alpha}_{-i}, \vec{z})$$

$$\propto p(\vec{z} \mid \vec{\alpha}) p(\vec{\alpha})$$

$$\propto \prod_m \left(\frac{\Gamma(n_{mi} + \alpha_i)}{\Gamma(\sum_k (n_{mk} + \alpha_k))} \cdot \frac{\Gamma(\sum_k \alpha_k)}{\Gamma(\alpha_i)} \right) \cdot \alpha_i^{a-1} e^{-b\alpha_i}$$

Analogously,

Analogously, the full conditional for Gibbs sampling of β can be obtained by

$$p(\beta|\vec{z};\vec{w}) \propto p(\vec{z}|\beta;\vec{w})p(\beta)$$
$$\propto \prod_k \frac{\Delta(\vec{n}_k+\beta)}{\Delta_V(\beta)}\beta^{a-1}e^{-b\beta}$$

For each parameter, we have

$$p(\beta_i \mid \vec{\beta}_{-i},\vec{z};\vec{w}) = \frac{p(\beta_i,\vec{\beta}_{-i},\vec{z};\vec{w})}{p(\vec{\beta}_{-i},\vec{z};\vec{w})} \propto p(\beta_i,\vec{\beta}_{-i},\vec{z};\vec{w})$$

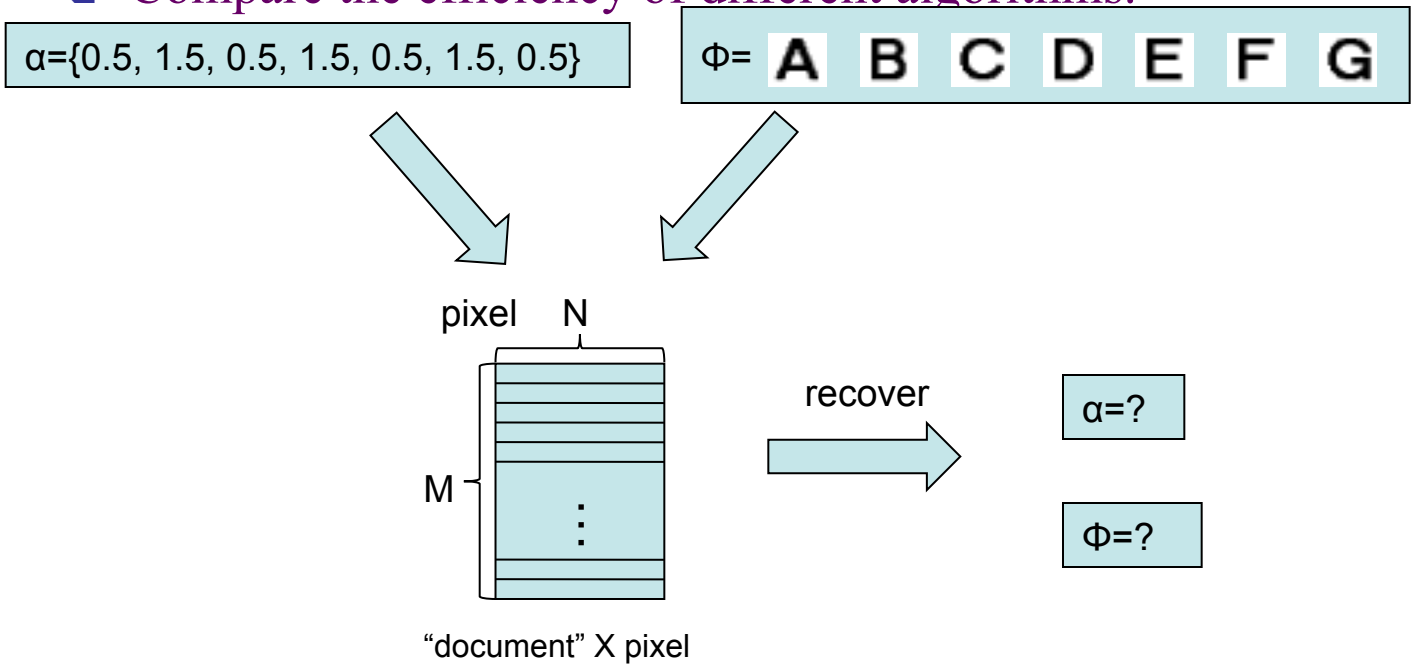
$$\propto p(\vec{z} \mid \vec{\beta};\vec{w})p(\vec{\beta};\vec{w})$$

$$\propto \prod_k \left(\frac{\Gamma(n_{ki}+\beta_i)}{\Gamma(\sum_v (n_{kv}+\alpha_v))} \cdot \frac{\Gamma(\sum_v \beta_v)}{\Gamma(\beta_i)} \right) \cdot \beta_i^{a-1}e^{-b\beta_i}$$

There are also many other methods (some non-sampling based method). For details, please refer to [Mink, 2000].

Experimental Results

- Experiment Design
  - Generate a synthetic data set using a predefined α and Φ.
  - Use different algorithms to recover both Φ and the hyperparameters α.
  - Compare the effectiveness of different learned algorithms.
  - Compare the efficiency of different algorithms.



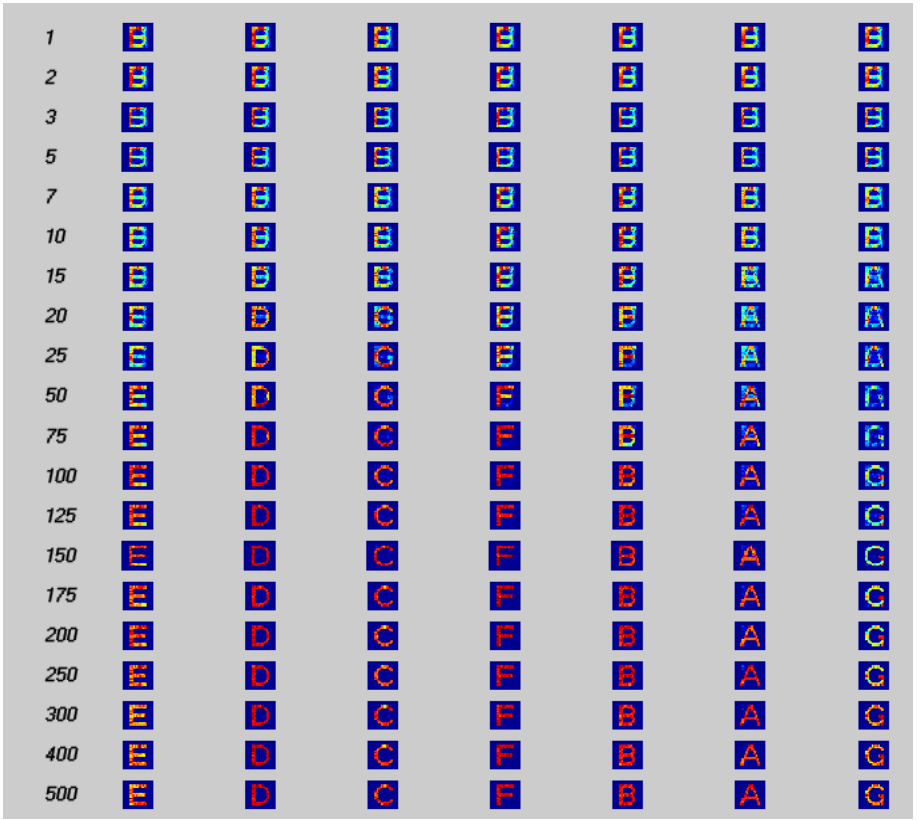
66

Performance of Recovering Φ

- Normal convergence process of topic-word (Φ) distribution



Original Φ

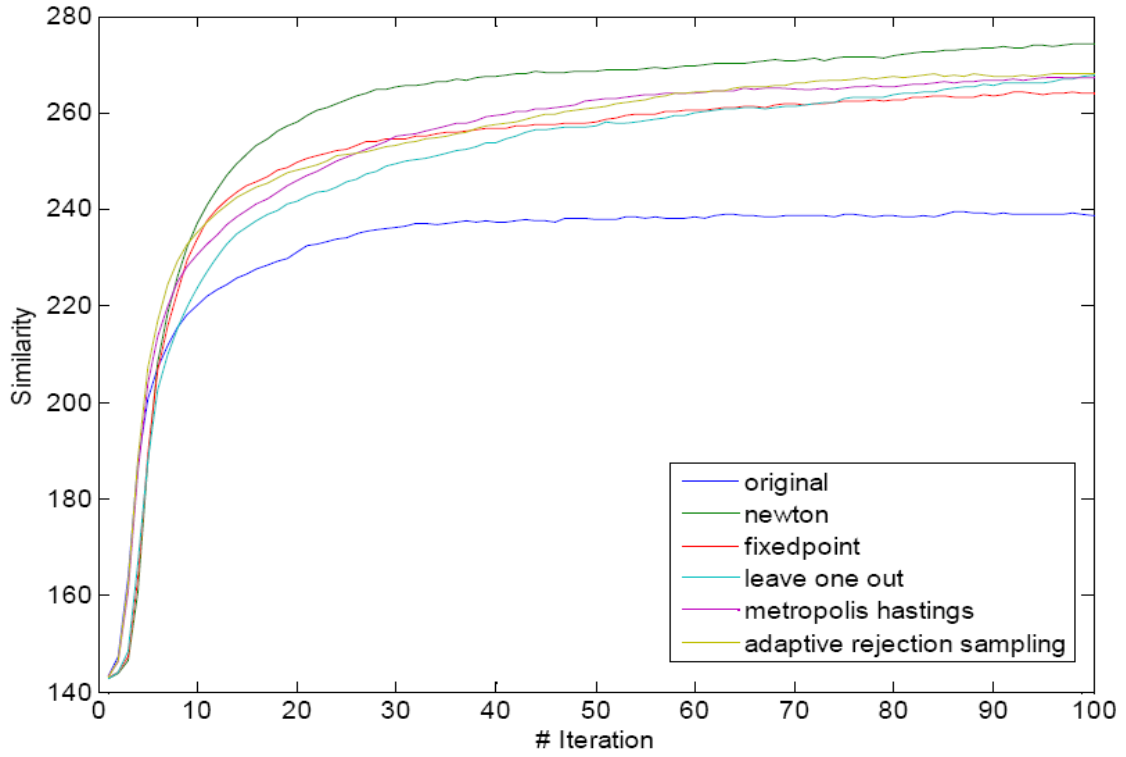


Recovered Φ

67

Performance of Recovering Φ

- Topic-item similarity traces over Gibbs sampling steps.

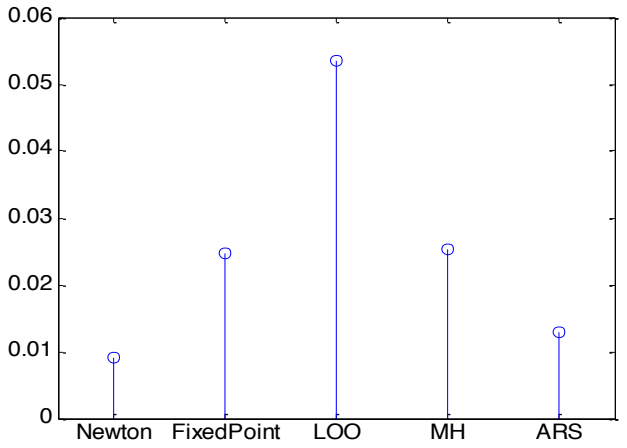


68

Performance of Recovering α

- Average estimated value with different estimation methods & the difference of between estimated value and real value.

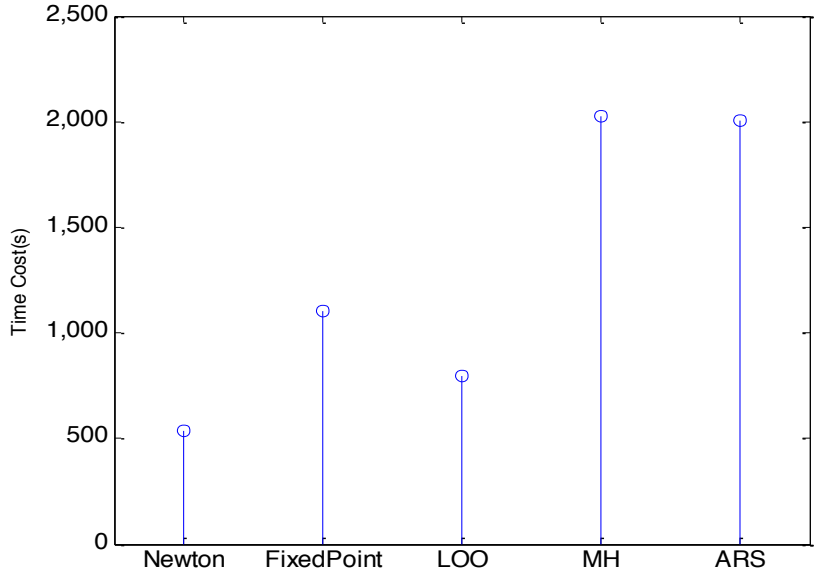
	1	2	3	4	5	6	7
Real	0.5	1.5	0.5	1.5	0.5	1.5	0.5
Newton	0.553	1.564	0.656	1.410	0.481	1.643	0.555
FixedPoint	0.488	1.476	0.533	1.374	0.462	1.478	0.892
LOO	0.495	1.195	0.444	1.258	0.576	1.076	0.683
MH	0.563	1.474	0.675	1.625	0.368	1.700	0.764
ARS	0.548	1.321	0.689	1.516	0.534	1.496	0.639



69

Efficiency Comparison

- Time Consumption comparison of each estimation methods.



70

Conclusion

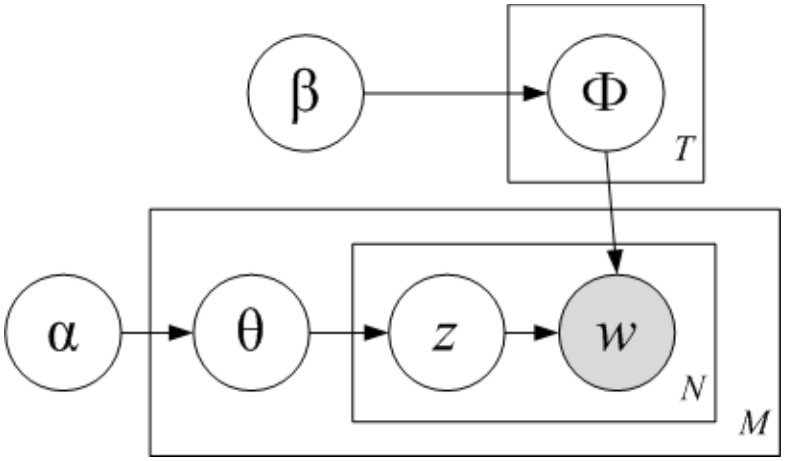
- Newton-Raphson method tends to be the best one in the experiment.
- In two sampling base methods, Metropolis-Hasting random walk is easier to implement, but the performance of adaptive rejection sampling seems a little better.
- In all the methods, the fixed point iteration using leave one out likelihood method doesn't need to calculate the special gamma function, which make it more scalable.

71

Variantal EM for Inference

- Another alternative algorithm for learning LDA

Joint Distribution



For each word, we have

$$p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \int p(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) \left(\prod_{n=1}^{N} \sum_{z_n} p(z_n \mid \boldsymbol{\theta}) p(w_n \mid z_n, \boldsymbol{\beta}) \right) d\boldsymbol{\theta}$$

$$= \frac{\Gamma(\sum_i \boldsymbol{\alpha}_i)}{\prod_i \Gamma(\boldsymbol{\alpha}_i)} \int \left(\prod_{i=1}^{k} \boldsymbol{\theta}_i^{\boldsymbol{\alpha}_i - 1} \right) \left(\prod_{n=1}^{N} \sum_{i=1}^{k} \prod_{j=1}^{V} (\boldsymbol{\theta}_i \boldsymbol{\beta}_{ij})^{w_n^j} \right) d\boldsymbol{\theta}$$

- We can write the probability of a corpus:

$$p(D \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = \prod_{d=1}^{M} \int p(\boldsymbol{\theta}_d \mid \boldsymbol{\alpha}) \left(\prod_{n=1}^{N_d} \sum_{z_{nd}} p(z_{dn} \mid \boldsymbol{\theta}_d) p(w_{dn} \mid z_{dn}, \boldsymbol{\beta}) \right) d\boldsymbol{\theta}_d$$

Variational Inference

- Here, we introduce a simple convexity-based variational algorithm for inference in LDA.
- The basic idea here is to make use of Jensen's inequality

$$\log p(w|\alpha,\beta) = \log \int \sum_z \frac{p(\theta,z,w|\alpha,\beta)q(\theta,z)}{q(\theta,z)} d\theta$$

$$\geq \log \int \sum_z \left(\frac{p(\theta,z,w|\alpha,\beta)}{q(\theta,z)} \right)^{q(\theta,z)} d\theta = \int \sum_z q(\theta,z) \log \frac{p(\theta,z,w|\alpha,\beta)}{q(\theta,z)} d\theta$$

$$= \int \sum_z q(\theta,z) \log p(\theta,z,w|\alpha,\beta) d\theta - \int \sum_z q(\theta,z) \log q(\theta,z) d\theta$$

$$= E_q\left[\log p(\theta,z,w|\alpha,\beta)\right] - E_q\left[\log q(\theta,z)\right]$$

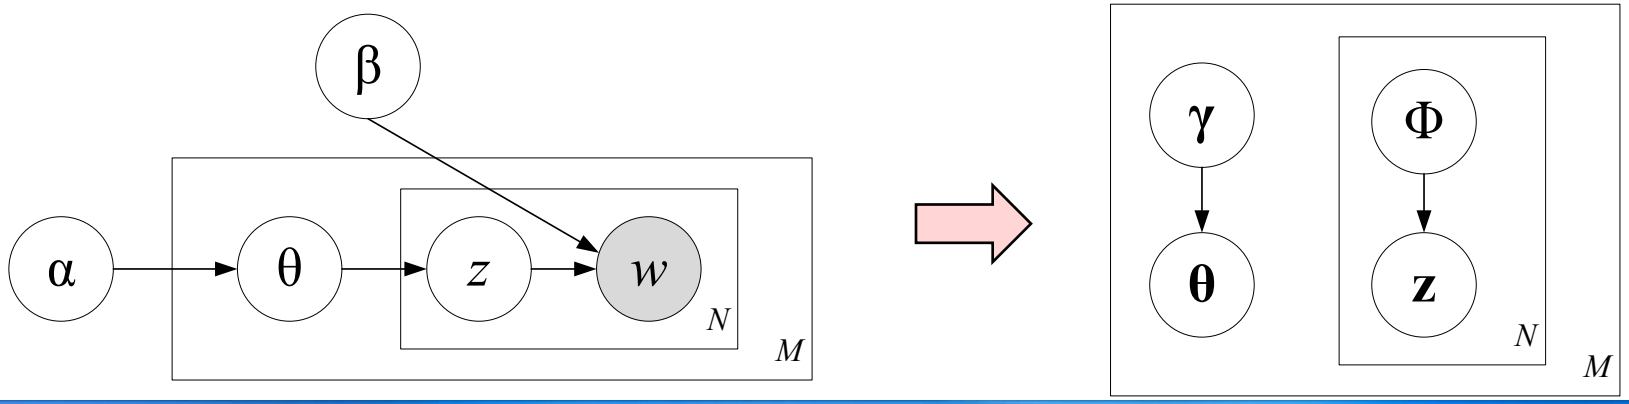
- A simple way to obtain a tractable family of lower bounds is to consider simple modifications of the original graphical model in which some of the edges and nodes are removed.

Variational Inference (cont.)

- Hence, by dropping edges between θ, z, and w, and w nodes, and also endow the resulting simplified graphical model with free variational parameters, we obtain a family of distributions on the latent variables:

$$q(\boldsymbol{\theta}, \mathbf{z} \mid \boldsymbol{\gamma}, \boldsymbol{\phi}) = q(\boldsymbol{\theta} \mid \boldsymbol{\gamma}) \prod_{n=1}^{N} q(z_n \mid \phi_n)$$

- where the Dirichlet parameter $\boldsymbol{\gamma}$ and the multinomial parameters $(\Phi_1, \ldots, \Phi_N)$ are the free variational parameters.



How to determine the vatiational parameters

- We can set up an optimization problem to determine the values of the variational parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\Phi}$.
- We can define the optimization function as minimizing the Kullback-Leibler (KL) divergence between the variational distribution and the true posterior $p(\boldsymbol{\theta},\mathbf{z}|\mathbf{w},\boldsymbol{\alpha},\boldsymbol{\beta})$:

$$(\boldsymbol{\gamma}^*,\boldsymbol{\phi}^*) = \arg\min_{(\gamma,\phi)} \mathrm{D}(q(\boldsymbol{\theta},\mathbf{z}\mid\boldsymbol{\gamma},\boldsymbol{\phi}) \,||\, p(\boldsymbol{\theta},\mathbf{z}\mid\mathbf{w},\boldsymbol{\alpha},\boldsymbol{\beta}) \qquad (5)$$

- This minimization can be achieved by an iterative fixed-point method.

Variational Inference

- We now discuss how to set the parameter $\boldsymbol{\gamma}$ and $\boldsymbol{\Phi}$ via an optimization procedure.
- Following Jordan et al. (1999), we have a lower bound of the log likelihood of a document using Jensen's inequality:

$$
\begin{aligned}
\boldsymbol{p}(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) &= \log \int \sum_{\mathbf{z}} p(\boldsymbol{\theta}, \mathbf{z}, \mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) d\theta \\
&= \log \int \sum_{\mathbf{z}} \frac{p(\boldsymbol{\theta}, \mathbf{z}, \mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) q(\boldsymbol{\theta}, \mathbf{z})}{q(\boldsymbol{\theta}, \mathbf{z})} d\theta \\
&> \int \sum_{\mathbf{z}} q(\boldsymbol{\theta}, \mathbf{z}) \log p(\boldsymbol{\theta}, \mathbf{z}, \mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) d\theta - \int \sum_{\mathbf{z}} q(\boldsymbol{\theta}, \mathbf{z}) \log q(\boldsymbol{\theta}, \mathbf{z}) d\theta \\
&= \mathrm{E}_q[\log p(\boldsymbol{\theta}, \mathbf{z}, \mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta})] - \mathrm{E}_q[\log q(\boldsymbol{\theta}, \mathbf{z})]
\end{aligned}
$$

Variational Inference (cont.)

- Then from the above formula, we see that the Jensen's inequality provides us with a lower bound on the log likelihood for an arbitrary variational distribution $q(\boldsymbol{\theta},\mathbf{z}|\boldsymbol{\gamma},\Phi)$.
- It can be easily verified that the difference between the left-hand side and the right-hand side of the above equation is the KL divergence between the variational posterior probability and the true posterior probability.

Variational Inference (cont.)

- That is, letting $L(\phi, \boldsymbol{\gamma}; \boldsymbol{\alpha}, \boldsymbol{\beta})$ denote the right-hand side of the above equation we have:

$$\log p(\mathbf{w} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}) = L(\boldsymbol{\gamma}, \boldsymbol{\phi}; \boldsymbol{\alpha}, \boldsymbol{\beta}) + \mathrm{D}(q(\boldsymbol{\theta}, \mathbf{z} \mid \boldsymbol{\gamma}, \boldsymbol{\phi}) \,||\, p(\boldsymbol{\theta}, \mathbf{z} \mid \mathbf{w}, \boldsymbol{\alpha}, \boldsymbol{\beta}))$$

- This means that maximizing the lower bound $L(\phi, \boldsymbol{\gamma}; \boldsymbol{\alpha}, \boldsymbol{\beta})$ w.r.t. γ and Φ is equivalent to minimizing the KL divergence between the variational posterior probability and the true posterior probability, the optimization problem in the above equation.

Variational Inference (cont.)

- We can expand the above equation

$$L(\phi,\boldsymbol{\gamma};\boldsymbol{\alpha},\boldsymbol{\beta}) = \log p(w\mid\alpha,\beta) - D\big(q(\theta,z\mid\gamma,\phi)\,\|\,p(\theta,z\mid w,\alpha,\beta)\big)$$
$$= E_q\left[\log p(\theta,z,w\mid\alpha,\beta)\right] - E_q\left[\log q(\theta,z)\right] - D\big(q(\theta,z\mid\gamma,\phi)\,\|\,p(\theta,z\mid w,\alpha,\beta)\big)$$
$$= E_q\left[\log p(\theta\mid\alpha)\right] + E_q\left[\log p(z\mid\theta)\right] + E_q\left[\log p(w\mid z,\beta)\right]$$
$$-E_q\left[\log q(\theta)\right] - E_q\left[\log q(z)\right]$$

- By extending it again, we can have

$$L(\phi,\boldsymbol{\gamma};\boldsymbol{\alpha},\boldsymbol{\beta}) = \log\Gamma\Big(\sum_{j=1}^{k}\boldsymbol{\alpha}_j\Big) - \sum_{i=1}^{k}\log\Gamma(\boldsymbol{\alpha}_i) + \sum_{i=1}^{k}(\boldsymbol{\alpha}_i - 1)\Big(\Psi(\boldsymbol{\gamma}_i) - \Psi\Big(\sum_{j=1}^{k}\boldsymbol{\gamma}_j\Big)\Big)$$
$$+ \sum_{n=1}^{N}\sum_{i=1}^{k}\boldsymbol{\phi}_{ni}\Big(\Psi(\boldsymbol{\gamma}_i) - \Psi\Big(\sum_{j=1}^{k}\boldsymbol{\gamma}_j\Big)\Big) + \sum_{n=1}^{N}\sum_{i=1}^{k}\sum_{j=1}^{V}\boldsymbol{\phi}_{ni}w_n^j\log\boldsymbol{\beta}_{ij} - \log\Gamma\Big(\sum_{j=1}^{k}\boldsymbol{\gamma}_j\Big)$$
$$+ \sum_{i=1}^{k}\log\Gamma(\boldsymbol{\gamma}_i) - \sum_{i=1}^{k}(\boldsymbol{\gamma}_i - 1)\Big(\Psi(\boldsymbol{\gamma}_i) - \Psi\Big(\sum_{j=1}^{k}\boldsymbol{\gamma}_j\Big)\Big) - \sum_{n=1}^{N}\sum_{i=1}^{k}\boldsymbol{\phi}_{ni}\log\boldsymbol{\phi}_{ni}$$

Variaitonal Multinomial

- We first maximize above Eq w.r.t. Φ_{ni}, the probability that the n-th word is generated by topic i.
- We form the Lagrangian by isolating the terms which contain Φ_{ni} and adding the appropriate Lagrange multipliers. Let β_{iv} be $p(w^v_n=1|z^i=1)$ for the appropriate v. (recall that each w_n is a vector of size V with exactly one component equal to one; we can select the unique v such that $w^v_n=1$):

$$L_{[\phi_{ni}]} = \phi_{ni}(\Psi(\gamma_i) - \Psi(\sum_{j=1}^{k}\gamma_j)) + \phi_{ni}\log\beta_{iv} - \phi_{ni}\log\phi_{ni} + \lambda_n(\sum_{j=1}^{k}\phi_{ni} - 1)$$

Variaitonal Multinomial (cont.)

- Taking derivatives w.r.t. Φ_{ni}, we obtain:

$$\frac{\partial L}{\partial \phi_{ni}} = \Psi(\boldsymbol{\gamma}_i) - \Psi(\sum_{j=1}^{k} \boldsymbol{\gamma}_j) + \log \boldsymbol{\beta}_{iv} - \log \phi_{ni} - 1 + \lambda$$

- Setting this to zero yields the maximizing value of the variational parameter Φ_{ni} :

$$\phi_{ni} \propto \boldsymbol{\beta}_{iv} \exp(\Psi(\boldsymbol{\gamma}_i) - \Psi(\sum_{j=1}^{k} \boldsymbol{\gamma}_j))$$

Variational Dirichlet

- Next we maximize equation (15) w.r.t. γ_i, the i-th component of the posterior Dirichlet parameters, the terms containing γ_i are:

$$L_{[\gamma]} = \sum_{i=1}^{k}(\alpha_i - 1)(\Psi(\gamma_i) - \Psi(\sum_{j=1}^{k}\gamma_j)) + \sum_{n=1}^{N}\phi_{ni}(\Psi(\gamma_i) - \Psi(\sum_{j=1}^{k}\gamma_j))$$

$$- \log\Gamma(\sum_{j=1}^{k}\gamma_j) + \log\Gamma(\gamma_i) - \sum_{i=1}^{k}(\gamma_i - 1)(\Psi(\gamma_i) - \Psi(\sum_{j=1}^{k}\gamma_j))$$

- By simplifying

$$L_{[\gamma]} = \sum_{i=1}^{k}(\alpha_i + \sum_{n=1}^{N}\phi_{ni} - \gamma_i)(\Psi(\gamma_i) - \Psi(\sum_{j=1}^{k}\gamma_j)) - \log\Gamma(\sum_{j=1}^{k}\gamma_j) + \log\Gamma(\gamma_i)$$

Variational Dirichlet (cont.)

- Taking the derivative w.r.t. γ_i:

$$\frac{\partial L}{\partial \boldsymbol{\gamma}_i} = \Psi'(\boldsymbol{\gamma}_i)(\boldsymbol{\alpha}_i + \sum_{n=1}^{N} \boldsymbol{\phi}_{ni} - \boldsymbol{\gamma}_i) - \Psi'(\sum_{j=1}^{k} \boldsymbol{\gamma}_j) \sum_{i=1}^{k} (\boldsymbol{\alpha}_i + \sum_{n=1}^{N} \boldsymbol{\phi}_{ni} - \boldsymbol{\gamma}_i)$$

$$= \sum_{i=1}^{k} \left(\alpha_i + \sum_{n=1}^{N} \phi_{ni} - \gamma_i \right) \left(\Psi'(\gamma_i) - \Psi'\left(\sum_{j=1}^{k} \gamma_j \right) \right)$$

- Setting this equation to zero yields a maximum at:

$$\boldsymbol{\gamma}_i = \boldsymbol{\alpha}_i + \sum_{n=1}^{N} \boldsymbol{\phi}_{ni}$$

Solve the Optimization Problem

- Derivate the KL divergence and setting them equal to zero, we obtain the following update equations:

$$\phi_{ni} \propto \beta_{iw_n} \exp\{E_q[\log(\theta_i) \mid \gamma]\}$$

$$\gamma_i = \alpha_i + \sum_{n=1}^{N} \phi_{ni}$$

where the expectation in the multinomial update can be computed as follows:

$$E_q[\log(\theta_i) \mid \gamma] = \Psi(\gamma_i) - \Psi(\sum_{j=1}^{k} \gamma_j)$$

where ψ is the first derivative of the logΓ function which is computable via Taylor approximations

Computing E[log($\theta_i|\alpha$)]

- Recall that a distribution is in the exponential family if it can be written in the form:

$$p(x \mid \eta) = h(x)\exp\{\eta^T T(x) - A(\eta)\}$$

where η is the natural parameter, T(x) is the sufficient statistic, and A(η) is the log of the normalization factor.

- As we can write the Dirichlet in this form by exponentiating the log of Eq.: $p(\boldsymbol{\theta}|\boldsymbol{\alpha})$

$$p(\boldsymbol{\theta} \mid \boldsymbol{\alpha}) = \exp\{(\sum_{i=1}^{k}(\boldsymbol{\alpha}_i - 1)\log \boldsymbol{\theta}_i) + \log\Gamma(\sum_{i=1}^{k}\boldsymbol{\alpha}_i) - \sum_{i=1}^{k}\log\Gamma(\boldsymbol{\alpha}_i)\}$$

Computing E[log($\boldsymbol{\theta}_i|\boldsymbol{\alpha}$)] (cont.)

- From this form we see that the natural parameter of the Dirichlet is $\eta_i=\alpha_i-1$ and the sufficient statistic is $T(\theta_i)=\log\theta_i$. Moreover, based on the general fact that the derivative of the log normalization factor w.r.t. the natural parameter is equal to the expectation of the sufficient statistic, we obtain:

$$\mathrm{E}[\log(\boldsymbol{\theta}_i)\,|\,\boldsymbol{\alpha}] = \Psi(\boldsymbol{\alpha}_i) - \Psi(\sum_{j=1}^{k}\boldsymbol{\alpha}_j)$$

where ψ is the digamma function, the first derivative of the log Gamma function.

Variational Inference Algorithm

(1) initialize $\phi_{ni}^0 = \frac{1}{k}$ for all i and n

(2) initialize $\gamma_i = \alpha_i + \frac{N}{k}$ for all i

(3) repeat

(4) for n=1 to N

(5) for i=1 to k

(6) $\phi_{ni}^{t+1} = \beta_{iw_n} \exp(\Psi(\gamma_i^t))$

(7) normalize ϕ_{ni}^{t+1} to sum to 1

(8) $\gamma_i^t = \alpha + \sum_{n=1}^{N} \phi_{ni}^{t+1}$

(9) until convergence

Each iteration requires O((*N*+1)*k*) operations

For a single document the iteration number is on the order of the number of words in it

Thus, the total number of operations roughly on the order

Parameter Estimation

- We can use a empirical Bayes method for parameter estimation. In particular, we wish to find parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ that maximize the marginal log likelihood:

$$\log(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \sum_{d=1}^{M} \log p(\mathbf{w}_d \mid \boldsymbol{\alpha}, \boldsymbol{\beta})$$

- The quantity $p(\mathbf{w}|\boldsymbol{\alpha}, \boldsymbol{\beta})$ can be computed by the variational inference as described above. Thus, we can find approximate empirical Bayes estimates for the LDA model via an alternating variational EM procedure that maximizes a lower bound w.r.t. the variational parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\Phi}$, and then fixed values of the variational parameters, maximizes the lower bound w.r.t. the model parameter $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$.

Variational EM

1. (E-step) For each document, find the optimizing values of the variational parameters $\{\boldsymbol{\gamma}_d^*, \boldsymbol{\phi}_d^* : d \in D\}$. This is done as described in the previous section.
2. (M-step) Maximize the resulting lower bound on the log likelihood w.r.t. the model parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$. This corresponds to finding maximum likelihood estimates with expected sufficient statistics for each document under the approximate posterior which is computed in the E-step. Actually, the update for the conditional multinomial parameter $\boldsymbol{\beta}$ can be written out analytically:

$$\boldsymbol{\beta}_{ij} \propto \sum_{d=1}^{M} \sum_{n=1}^{N_d} \boldsymbol{\phi}_{dni}^* w_{dn}^j \qquad (9)$$

The update for $\boldsymbol{\alpha}$ can be implemented using an efficient Newton -Raphson method. These two steps are repeated until converges.

Parameter Estimation

- We now consider how to obtain empirical Bayes estimates of the model parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$.
- We solve this problem by using the variational lower bound as a surrogate for the marginal log likelihood, with the variational parameters Φ and γ fixed to the values found by variational inference.
- We then obtain empirical Bayes estimates by maximizing this lower bound w.r.t. the model parameters.

Parameter Estimation (cont.)

- Recall our approach for finding the empirical Bayes estimates is based on a variational EM procedure.
- In the variational E-step, we maximize the bound $L(\phi, \gamma; \alpha, \beta)$ w.r.t. the variational parameter γ and Φ. In the M-step, which we describe in this section, we maximize the bound w.r.t. the model parameters α and β. The overall procedure can thus be viewed as coordinate ascent in L.

Conditional Multinomials

- To maximize w.r.t. β, we isolate terms and add Lagrange multipliers:

$$L_{[\boldsymbol{\beta}]} = \sum_{d=1}^{M}\sum_{n=1}^{N_d}\sum_{i=1}^{k}\sum_{j=1}^{V}\boldsymbol{\phi}_{dni} w_{dn}^{j} \log \boldsymbol{\beta}_{ij} + \sum_{i=1}^{k}\lambda_i\left(\sum_{j=1}^{V}\boldsymbol{\beta}_{ij} - 1\right)$$

- Taking the derivative w.r.t. β_{ij} and set it to zero, we have

$$\boldsymbol{\beta}_{ij} = \sum_{d=1}^{M}\sum_{n=1}^{N_d}\boldsymbol{\phi}_{dni} w_{dn}^{j}$$

Dirichlet

- First, we have

$$L_{[\boldsymbol{\alpha}]} = \sum_{d=1}^{M}\left(\log\Gamma(\sum_{j=1}^{k}\boldsymbol{\alpha}_j) - \sum_{i=1}^{k}\log\Gamma(\boldsymbol{\alpha}_i) + \sum_{i=1}^{k}(\boldsymbol{\alpha}_i - 1)(\Psi(\boldsymbol{\gamma}_{di}) - \Psi(\sum_{j=1}^{k}\boldsymbol{\gamma}_{dj}))\right)$$

- Taking derivative w.r.t. α_i, we obtain:

$$\frac{\partial L}{\partial \boldsymbol{\alpha}_i} = M(\Psi(\sum_{j=1}^{k}\boldsymbol{\alpha}_j) - \Psi(\boldsymbol{\alpha}_i)) + \sum_{d=1}^{M}\left(\Psi(\boldsymbol{\gamma}_{di}) - \Psi(\sum_{j=1}^{k}\boldsymbol{\gamma}_{dj})\right)$$

- This derivative depends on α_i, where $j<>i$, and we therefore must use an iterative method to find the maximal α. In particular, the Hessian is in the form found in equation:

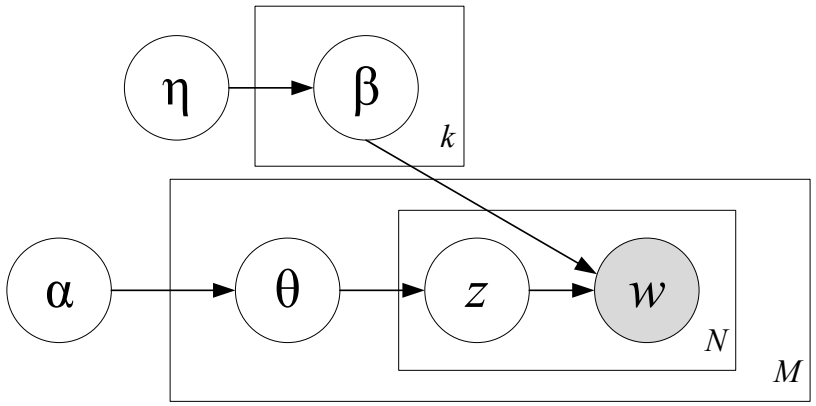
$$\frac{\partial L}{\partial \boldsymbol{\alpha}_i \boldsymbol{\alpha}_j} = \delta(i,j)M\Psi'(\boldsymbol{\alpha}_i) - \Psi'(\sum_{j=1}^{k}\boldsymbol{\alpha}_j)$$

Dirichlet (cont.)

- Finally, we can use the same algorithm to find an empirical Bayes point estimate of η, the scalar parameter for the exchangeable Dirichlet in the smoothed LDA model as mentioned above.

Smoothing

- Simple Laplace smoothing is no longer justified as a maximum a posteriori method in LDA setting.
- We can then assume that each row in $\boldsymbol{\beta}_{k \times V}$ is independently drawn from an exchangeable Dirichlet distribution. That is to treat $\boldsymbol{\beta}_i$ as random variables that are endowed with a posterior distribution, conditioned on the data.



Smoothing Model

- Thus we obtain a variational approach to Bayesian inference:

$$q(\boldsymbol{\beta}_{1:k}, \boldsymbol{\theta}_{1:M}, \mathbf{z}_{1:M} \mid \boldsymbol{\eta}, \boldsymbol{\gamma}, \boldsymbol{\phi}) = \prod_{i=1}^{k} Dir(\boldsymbol{\beta}_i \mid \boldsymbol{\eta}_i) \prod_{n=1}^{N} q_d(\boldsymbol{\theta}_d, \mathbf{z}_n \mid \phi_n, \boldsymbol{\gamma}_d)$$

where $q_d(\boldsymbol{\theta}_d, \mathbf{z}_n \mid \phi_n, \boldsymbol{\gamma}_d)$ is the variational distribution defined for LDA as above and the update for the new variational parameter $\boldsymbol{\eta}$ is as follow:

$$\boldsymbol{\eta}_{ij} \propto \boldsymbol{\eta} + \sum_{d=1}^{M} \sum_{n=1}^{N_d} \boldsymbol{\phi}^*_{dni} w^j_{dn}$$

Applications

Document Modeling

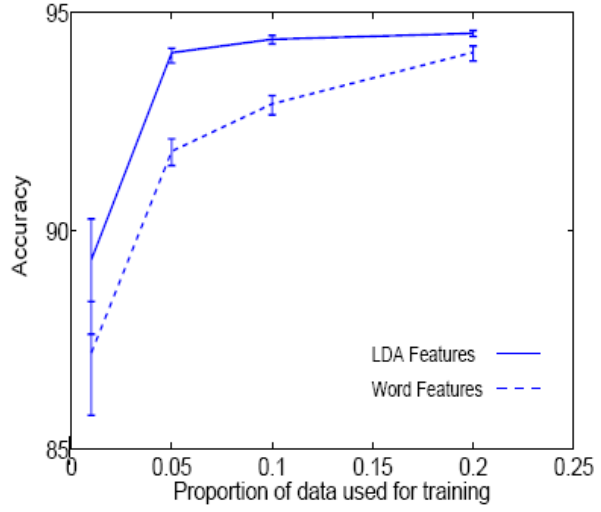
- Perplexity is used to indicate the generalization performance of a method.
- Specifically, we estimate a document modeling and use this model to describe the new data set.

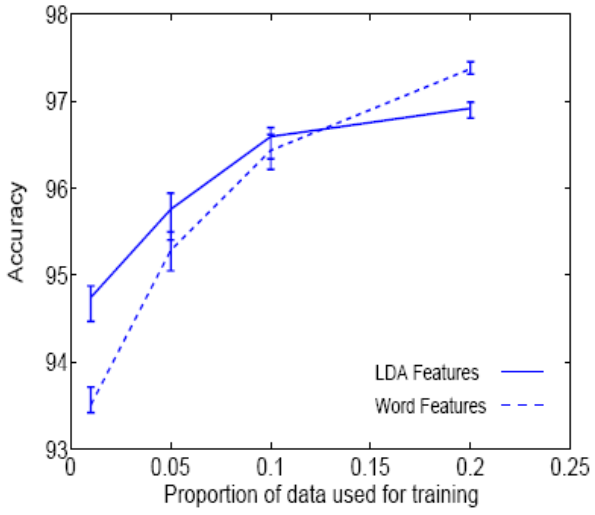
$$perplexity(D_{test}) = \exp\{-\frac{\sum_{d=1}^{M} \log p(w_d)}{\sum_{d=1}^{M} N_d}\}$$

- LDA outperforms the other models including pLSI, Smoothed Unigram, and Smoothed Mixt. Unigrams.

Document Classification

- We can use the LDA model results as the features for classifiers. In this way, say 50 topics, we can reduce the feature space by 99.6%.
- The experimental results show that such feature reduction may decrease the accuracy only a little.





Document Clustering

- We can also use the LDA model results as the features for clustering algorithms.

Collaborative Filtering

- We can learn a model on a fully observed set of users. Then for each unobserved user, we are shown all but one of the movies preferred by that user and are asked to predict what the held-out movie is.
- Precisely, define the predictive perplexity on M test uses as:

$$predictive - perplexity(D_{test}) = \exp\{-\frac{\sum_{d=1}^{M} \log p(w_{d,N_d} \mid \mathbf{w}_{d,1:N_d-1})}{M}\}$$

Our work related to topic model

Topic Browser

Topic 13: Computing programming



Commun. ACM , IFIP Congress , IFIP Congress (2) , A
Comput. Surv. , Advances in Computers , IFIP Congre

Topic 14: Computer networks and communication



INFOCOM , IEEE/ACM Trans. Netw. , IEEE Journal on
Areas in Communications , SIGCOMM , Computer Netw
ICNP

Topic 15: Automatic reasoning



CADE , J. Autom. Reasoning , J. Log. Comput. , JELIA
Math. Artif. Intell. , KR

Topic 16: Knowledge discovery and data mining



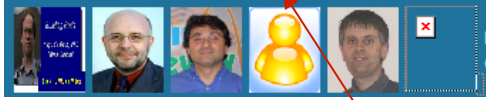
KDD , PAKDD , ICDM , PKDD , IEEE Trans. Knowl. Da
, SDM

Topic 17: Human computer interaction



CHI Extended Abstracts , CHI , Interactions , INTERA
Interacting with Computers , CHI Conference Compani

Topic 18: Evolvable Hardware



Evolvable Hardware , ICES , AHS , BioADIT , Conf.
Computing Frontiers , UC

200 topics have been discovered automatically from the academic network

Hot words in **Topic 16: Knowledge discovery and data mining**

mining:	0.041622	**data:**	0.041032
clustering:	0.014825	**discovery:**	0.010956
learning:	0.010630	**classification:**	0.010569
rules:	0.008899	**association:**	0.008655
patterns:	0.007962	**knowledge:**	0.007820
approach:	0.007636	**algorithm:**	0.007596
frequent:	0.007473	**based:**	0.007311
analysis:	0.007209	**databases:**	0.006211
efficient:	0.006129	**web:**	0.005417
text:	0.004908	**detection:**	0.004887



Hot conferences/jounals in **Topic 16: Knowledge discovery and data mining**

KDD: 0.143204
PAKDD: 0.106370
ICDM: 0.100902
PKDD: 0.082549
IEEE Trans. Knowl. Data Eng.: 0.064238
SDM: 0.047920
Intell. Data Anal.: 0.044338
SIGKDD Explorations: 0.037705
IDA: 0.033127
Knowl. Inf. Syst.: 0.030880
Data Min. Knowl. Discov.: 0.027447
DaWaK: 0.024374
CIKM: 0.022848
ICDM Workshops: 0.022806
Discovery Science: 0.020771
MLDM: 0.017677
CIDM: 0.016194
J. Intell. Inf. Syst.: 0.008055
KDID: 0.007081
DMKD: 0.007059

Active authors in **Topic 16: Knowledge discovery and data mining**


Jiawei Han


Eamonn J. Keogh


Philip S. Yu


Mohammed Javeed Zaki


Ming-Syan Chen


Wei Wang


Srinivasan Parthasarathy



Charu C. Aggarwal

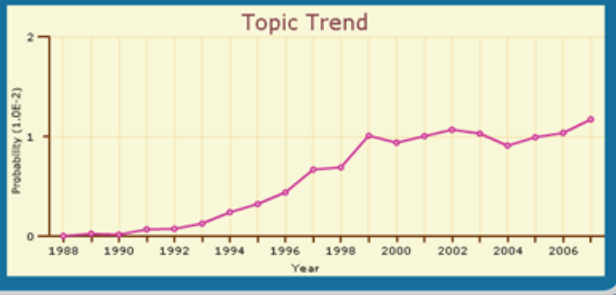

Qiang Yang


Shusaku Tsumoto

Data Mining
data stream
Data Mining Techniques
Evolving Data Streams
Mining Time Series Data

Association Rules
Association Mining
Data Streams
Mining Association Rules
data mining application

Topic Trend in **Topic 16: Knowledge discovery and data mining**



Publications in **Topic 16: Knowledge discovery and data mining**

Data Mining Techniques for Associations, Clustering and Classification
Authors: Charu C. Aggarwal, Philip S. Yu.
In Proceedings: PAKDD
Published year: 1999

Data Mining for the Web
Authors: Myra Spiliopoulou.
In Proceedings: PKDD
Published year: 1999

On string classification in data streams
Authors: Charu C. Aggarwal, Philip S. Yu.
In Proceedings: KDD
Published year: 2007

Matching in Time Series Databases
Authors: Eamonn J. Keogh, Padhraic Smyth.
In Proceedings: KDD
Published year: 1997

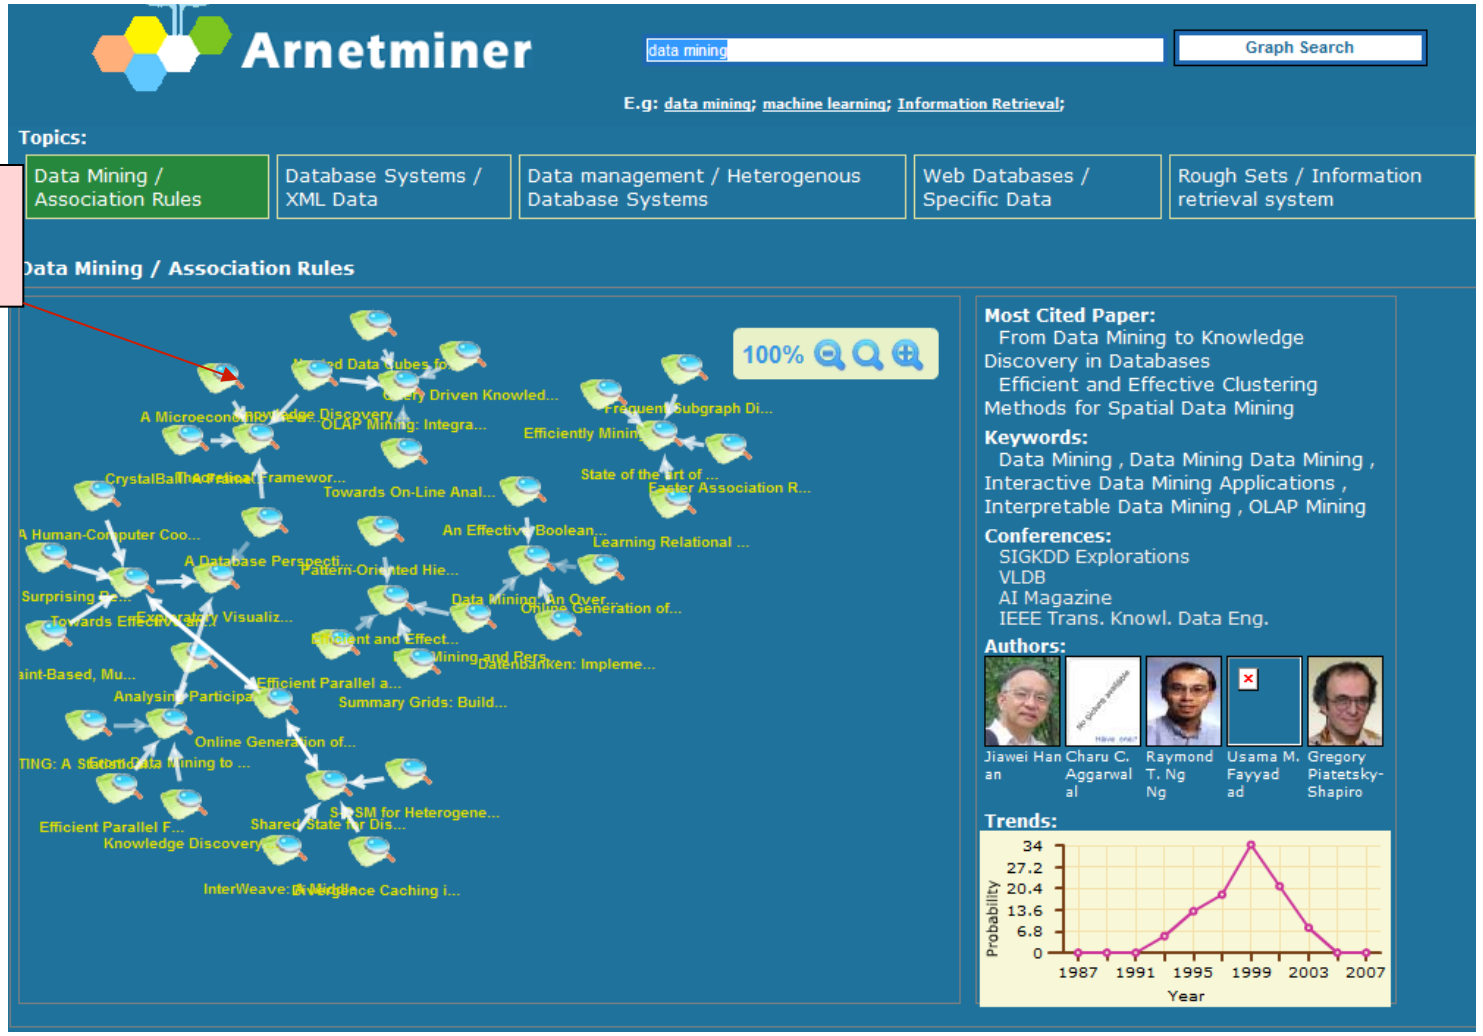
Knowledge Discovery and Data Mining: Towards a Unifying Framework
Authors: Usama M. Fayyad, Gregory Piatetsky-Shapiro, Padhraic Smyth.
In Proceedings: KDD
Published year: 1996

Efficiently mining frequent trees in a forest
Authors: Mohammed Javeed Zaki.
In Proceedings: KDD
Published year: 2002

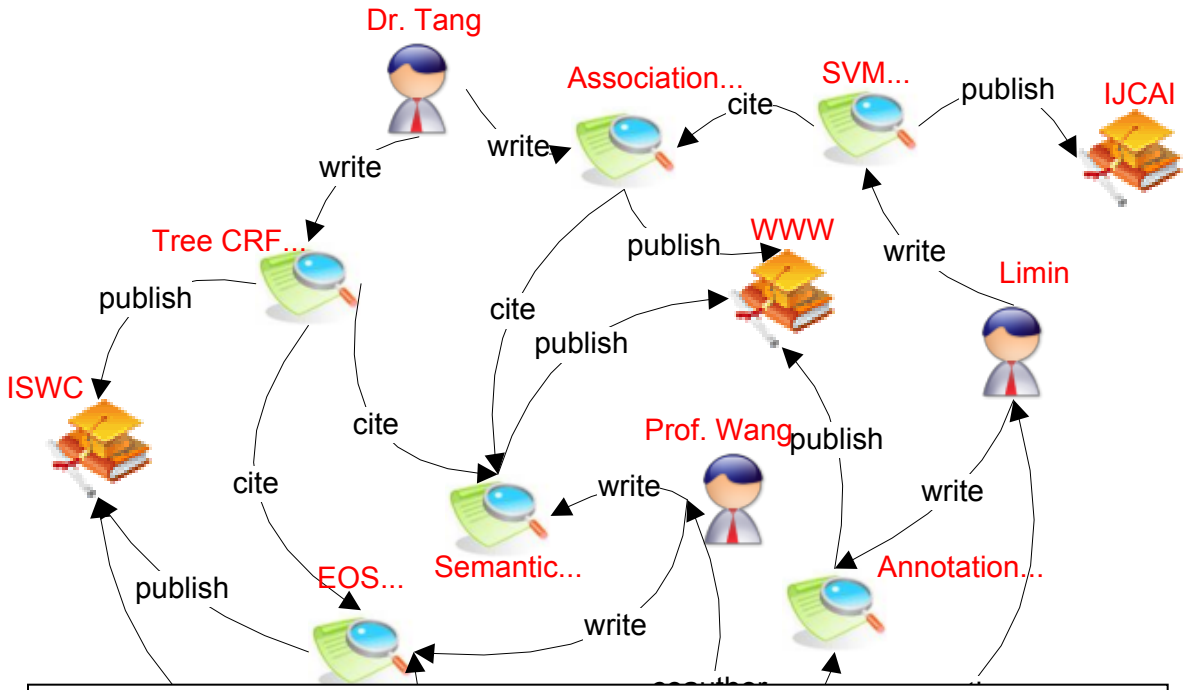
Mining of Evolving Data Streams with Privacy Preservation
Authors: Philip S. Yu.
In Proceedings: PAKDD

Graph Search

Sub graphs



Heterogeneous Academic Networks



In Academic Research Network

Heterogeneous objects:

Paper, Person, Conf./Journal

Relationships:

- Conf./Journal *publish* paper
- Paper *cite* paper
- Person *write* paper
- Person is *PC member* of Conf./Journal
- Person is *coauthor* of person

Challenges:

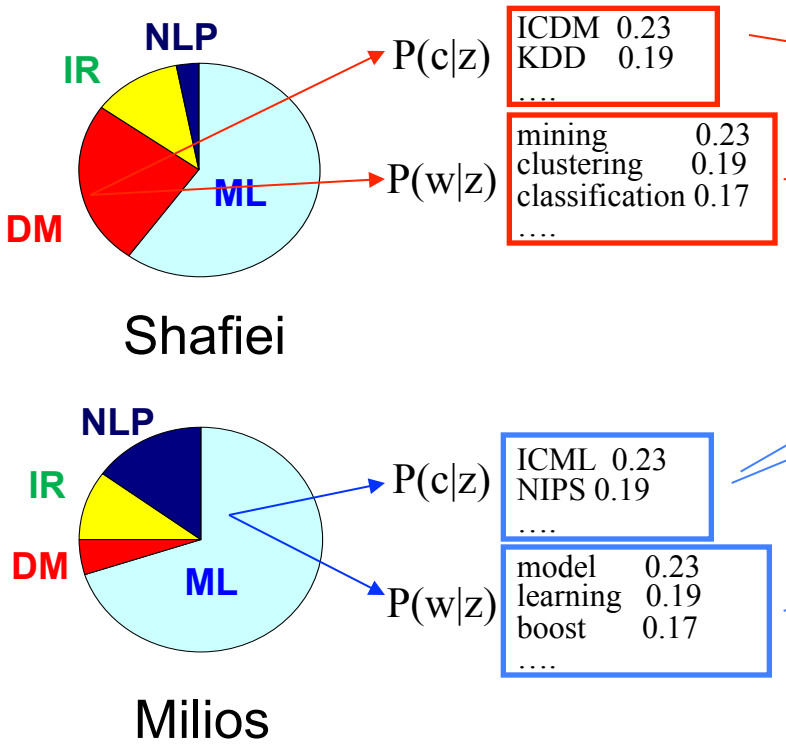
- How to model the heterogeneous objects in a unified model?

- How to model the relationships between the heterogeneous data?

- How to apply the proposed model to different applications?

Work1: Modeling social network --Generative Story

- Generative process



Paper

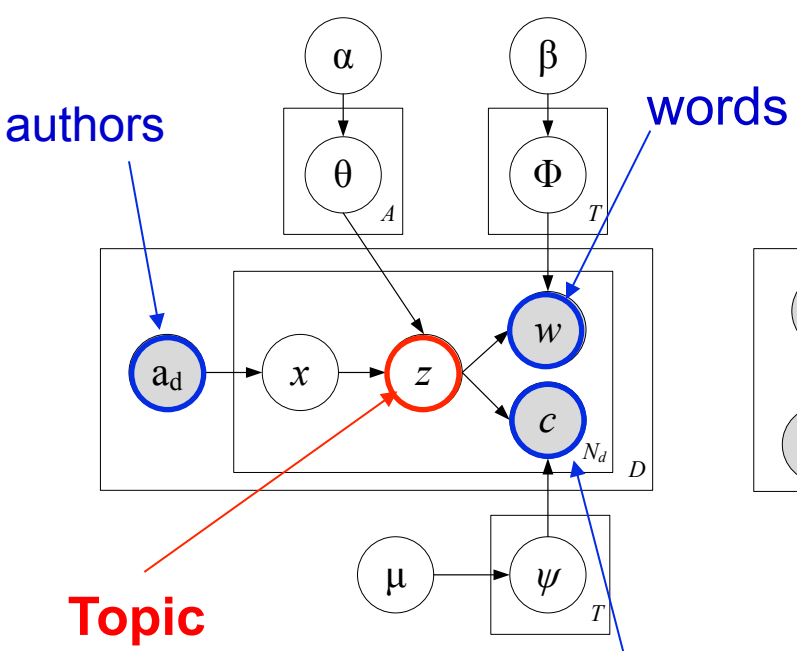
Latent Dirichlet Co-clustering

Shafiei and Milios

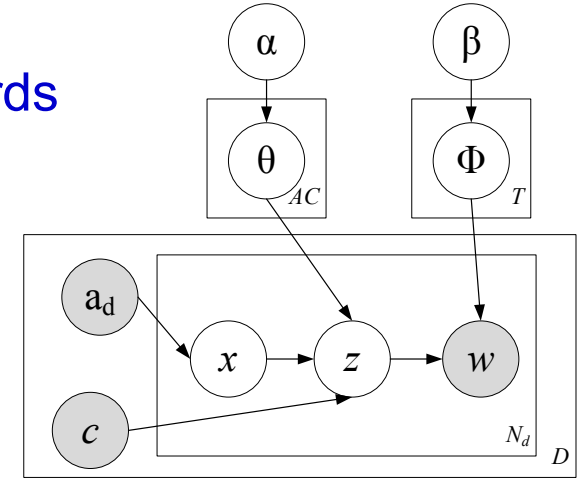
ICDM ***NIPS***

We present a generative model for clustering documents and terms. Our model is a four hierarchical bayesian model. We present efficient inference techniques based on Markow Chain Monte Carlo. We report results in document modeling, document and terms clustering …

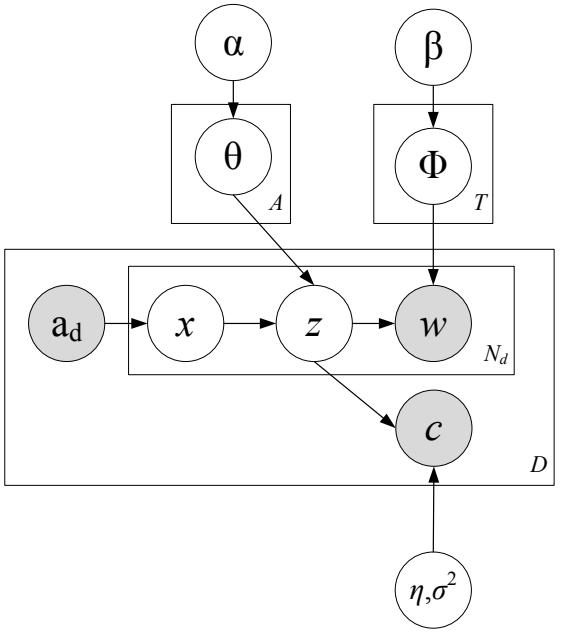
Modeling Multiple Objects in Academic network (KDD'08)



ACT1

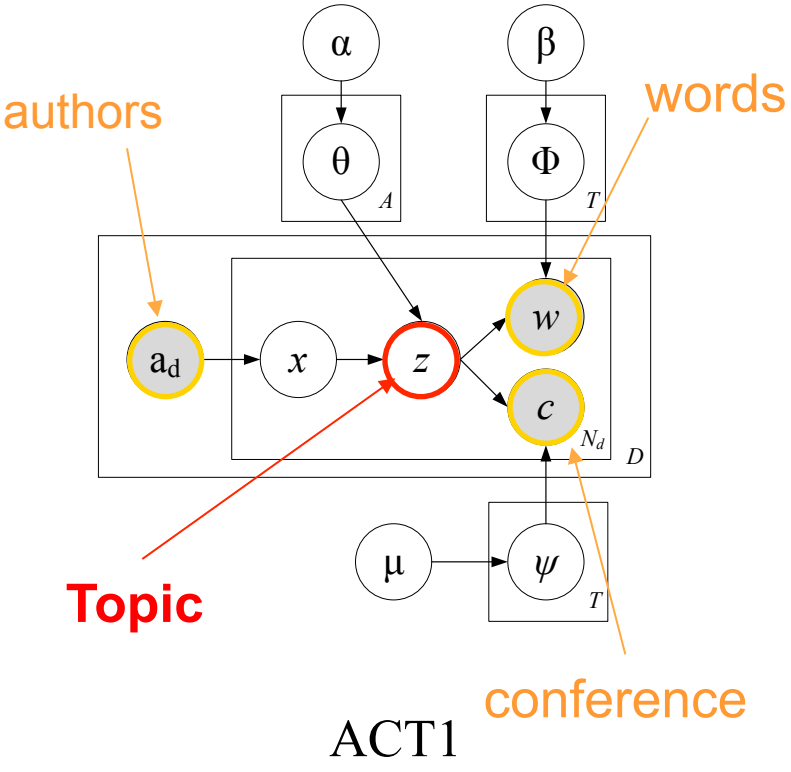


ACT2



ACT3

ACT Model 1



ACT1

Generative process:

1. For each topic z, draw ϕ_z and ψ_z respectively from Dirichlet priors β_z and μ_z;

2. For each word w_{di} in document d:

   - draw an author x_{di} from $\mathbf{a}_d$ uniformly;
   - draw a topic z_{di} from a multinomial distribution $\theta_{x_{di}}$ specific to author x_{di}, where θ is generated from a Dirichlet prior α;
   - draw a word w_{di} from multinomial $\phi_{z_{di}}$;
   - draw a conference stamp c_{di} from multinomial $\psi_{z_{di}}$.

$$P(z_{di}, x_{di} | \mathbf{z}_{-di}, \mathbf{x}_{-di}, \mathbf{w}, \mathbf{c}, \alpha, \beta, \mu) \propto \frac{m^{-di}_{x_{di} z_{di}} + \alpha_{z_{di}}}{\sum_z (m^{-di}_{x_{di} z} + \alpha_z)} \frac{n^{-di}_{z_{di} w_{di}} + \beta_{w_{di}}}{\sum_v (n^{-di}_{z_{di} v} + \beta_v)} \frac{n^{-d}_{z_{di} c_d} + \mu_{c_d}}{\sum_c (n^{-d}_{z_{di} c} + \mu_c)}$$

Expertise Search Results

Arnetminer data:
- 14,134 authors
- 10,716 papers
- 1,434 confs/journals

Evaluation measures:
- pooled relevance
- + human judgement

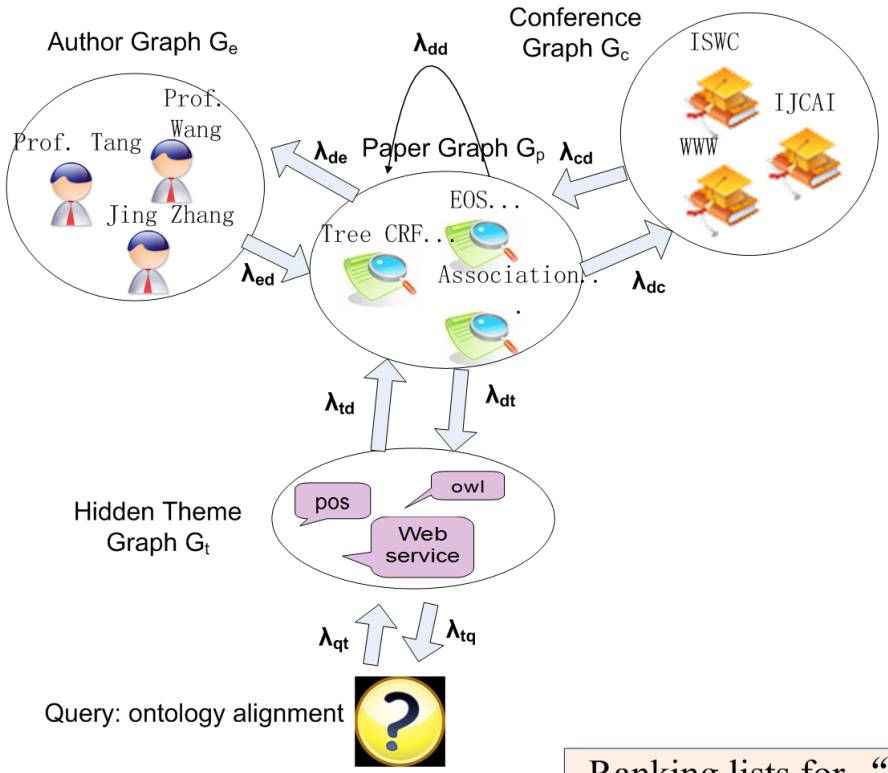
Baselines:
- Language Model (LM)
- LDA
- Author Topic (AT)

Method	Object	P@5	P@10	P@20	R-pre	MAP
LM	Paper	40.0	38.6	37.1	10.0	46.4
	Author	65.7	44.3	25.0	58.8	73.4
	Conference	51.4	32.9	21.4	47.6	63.1
	Average	52.4	38.6	27.9	38.8	61.0
LDA	Paper	31.4	48.6	42.9	13.5	45.8
AT	Paper	42.9	48.6	42.9	13.1	49.3
	Author	82.9	45.7	25.7	73.5	78.1
	Average	62.9	47.1	34.3	43.3	63.7
ACT1	Paper	42.9	45.7	43.6	16.6	51.0
	Author	91.4	50.0	26.4	80.0	89.6
	Conference	62.9	41.4	23.6	60.7	72.3
	Average	**65.7**	**45.7**	**31.2**	**52.4**	**71.0**
ACT2	Paper	42.9	47.1	39.3	15.0	47.7
	Author	74.3	50.0	25.7	69.4	80.1
	Conference	54.3	41.4	22.1	54.2	63.9
	Average	57.1	46.2	29.1	46.2	63.9
ACT3	Paper	42.9	38.6	41.4	17.1	47.0
	Author	71.4	47.1	25.7	70.0	78.7
	Conference	57.1	38.6	23.6	58.3	65.7
	Average	57.1	41.4	30.2	48.5	63.8

Work2: Integration of Links

—A Unified Probabilistic Ranking Model (ICDM'08)

Unified Probabilistic Model Combing Random Walk over Heterogeneous Network



- The unified model integrates different kinds of objects, called **heterogeneous graph**.
- The weights $p(d_i|d_j)$, $p(e_i|d_i)$, $p(c_i|d_i)$, $p(d_i|e_i)$, and $p(d_i|c_i)$ are similarly defined as uniform probability like PageRank.
- The **key difference** here is that the unified model integrates hidden theme nodes and a virtual query node into the heterogeneous graph. Hence, $p(\theta_i|d_i)$, $p(\theta_i|q)$, $p(d_i|\theta_i)$, and $p(q|\theta_i)$ are calculated based on Mixture Model.
- The rank vector is: (where M is the transition probability matrix)

$$r = Ar, A = M^T$$

- Finally, the model can rank different objects simultaneously.

$$r(d_i) = \lambda_{dd} \times \sum_{(d_j,d_i)\in E_{dd}} p(d_i \mid d_j)\cdot r(d_j) + \lambda_{ed} \times \sum_{(e_i,d_i)\in E_{de}} P(d_i \mid e_i)\cdot r(e_i)$$
$$+ \lambda_{cd} \times \sum_{(c_i,d_i)\in E_{dc}} P(d_i \mid c_i)\cdot r(c_i) + \lambda_{td} \times \sum_{(\theta_i,d_i)\in E_{dt}} P(d_i \mid \theta_i)\cdot r(\theta_i)$$

$$r(e_i) = \lambda_{de} \times \sum_{(d_j,e_i)\in E_{de}} p(e_i \mid d_j)\cdot r(d_j)$$

$$r(c_i) = \lambda_{dc} \times \sum_{(d_j,c_i)\in E_{dc}} p(c_i \mid d_j)\cdot r(d_j)$$

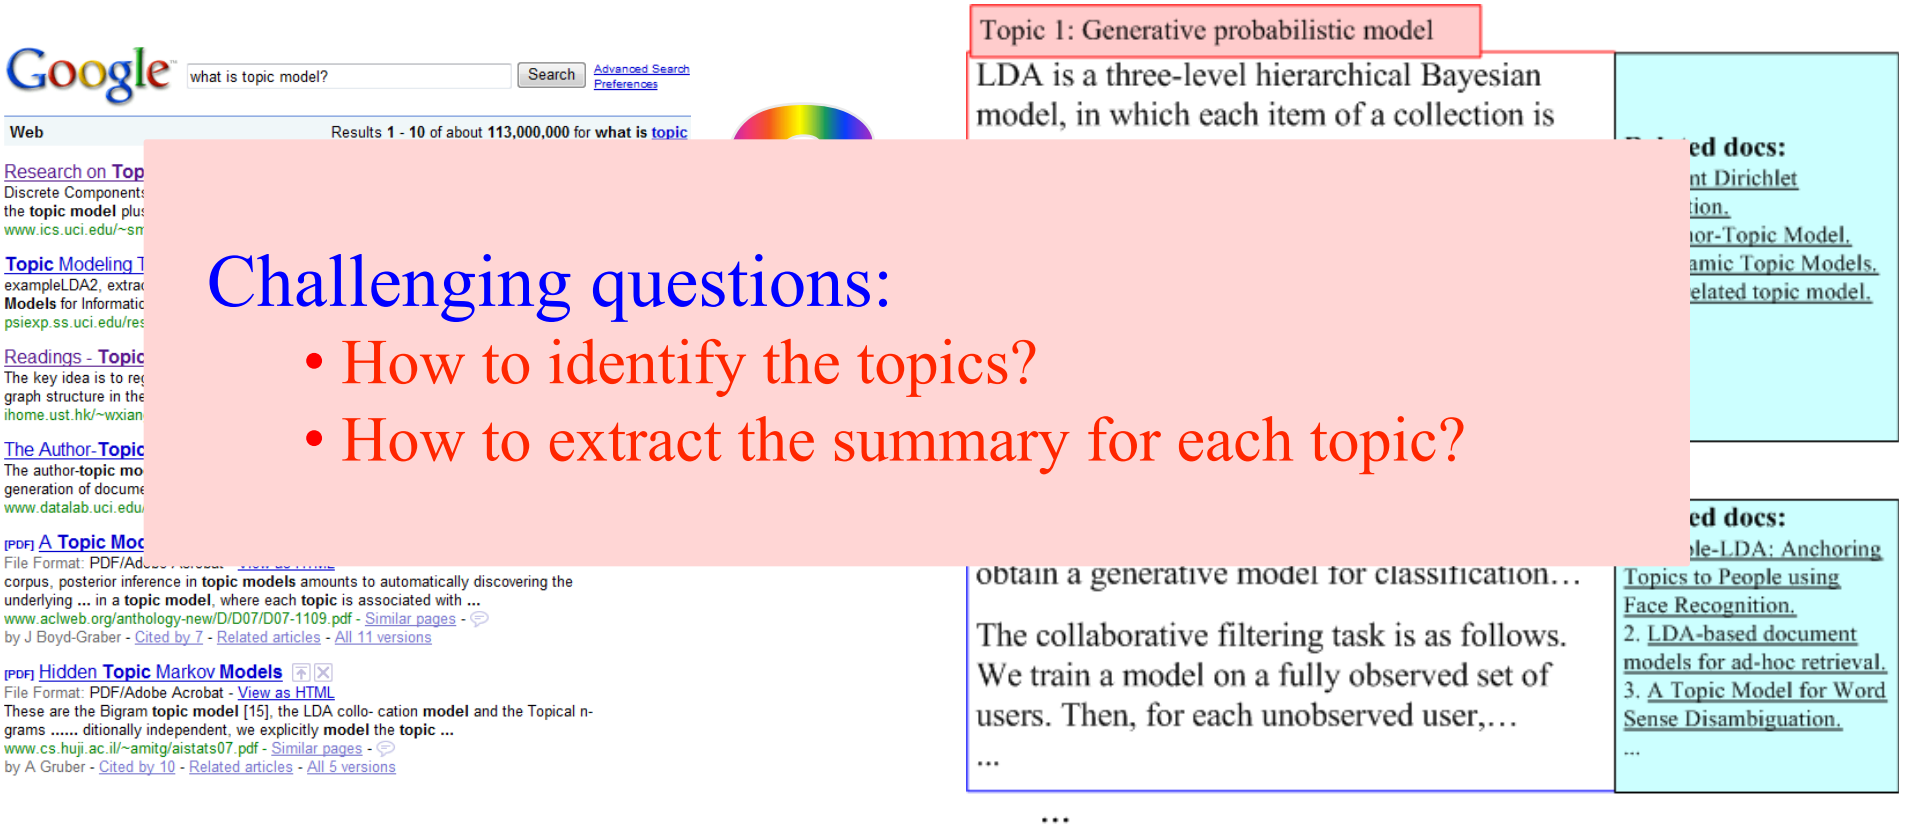
Ranking lists for "Information Extraction"

Persons	Conferences	Papers
• Raymond J. Mooney • Andrew McCallum • Fabio Ciravegna • Nicholas Kushmerick	• AAAI • ACL • IJCAI • KDD	• Conditional random fields: Probabilistic models for segmenting and labeling sequence data. • Maximum entropy markov models for information extraction and segmentation. • Wrapper induction for information extraction. • Learning information extraction rules for semi-structured and free text

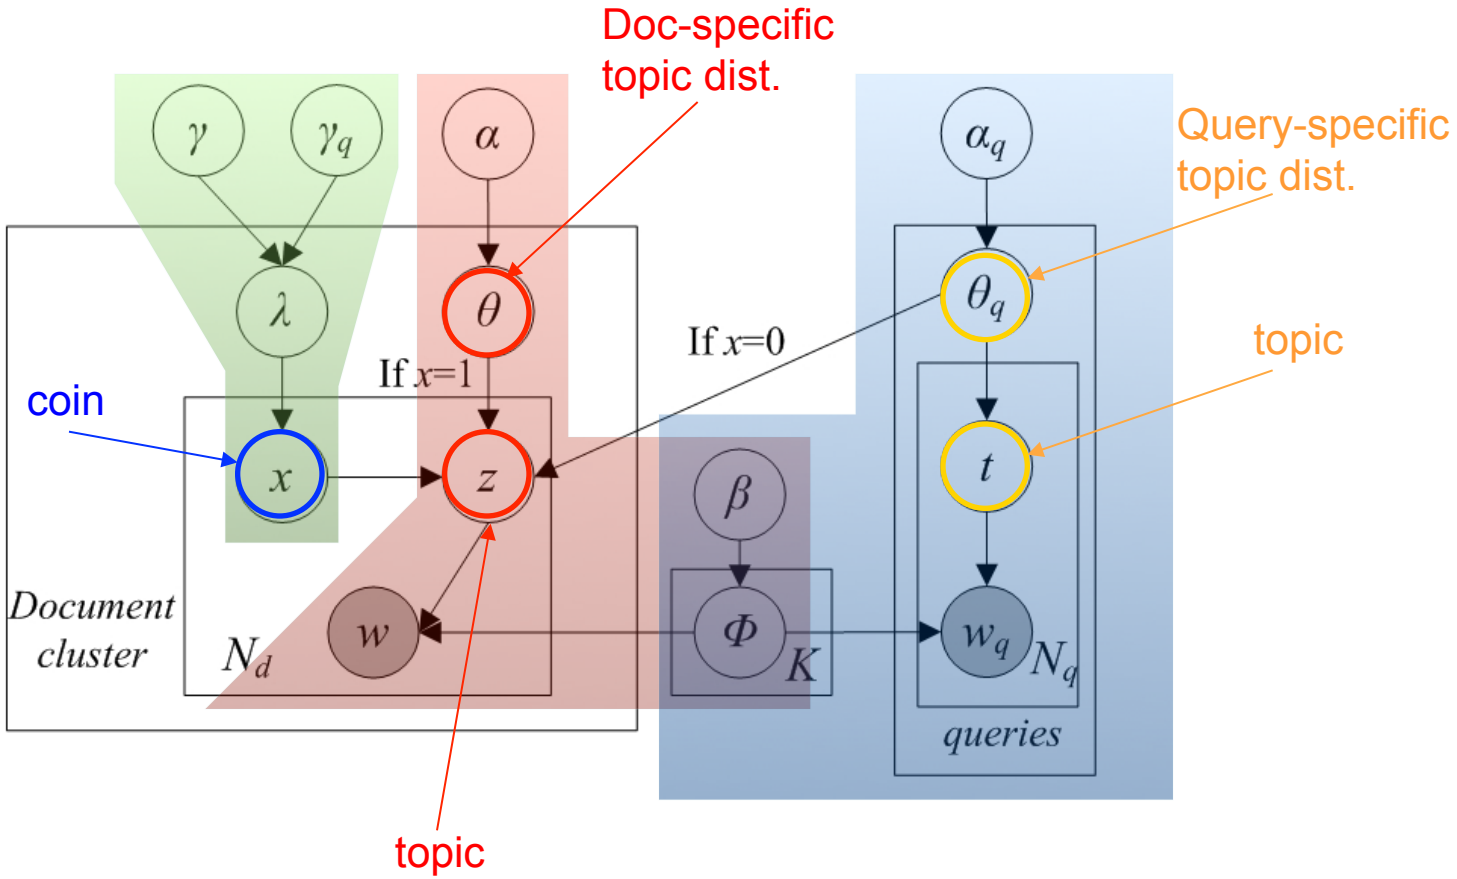
Expertise Search Results

Method	Object	P@5	P@10	P@20	R-pre	MAP
BM25	Paper	42.9	45.7	41.4	12.0	47.2
	Author	77.1	47.1	26.4	67.5	85.5
	Conference	51.4	38.6	22.9	48.8	66.0
	Average	57.1	43.8	30.2	42.8	66.2
LM	Paper	40.0	38.6	37.1	10.0	46.4
	Author	65.7	44.3	25.0	58.8	73.4
	Conference	51.4	32.9	21.4	47.6	63.1
	Average	52.4	38.6	27.9	38.8	61.0
LM+RW1	Paper	62.9	55.7	44.3	12.9	65.3
	Author	71.4	48.6	25.7	64.6	83.8
	Conference	60.0	35.7	22.1	53.6	64.6
	Average	64.8	46.7	30.7	43.7	71.2
pLSI	Paper	32.5	33.8	30	9.7	40.4
	Author	65.0	40.0	22.5	60.4	75.5
	Conference	47.5	36.3	21.3	45.1	54.1
	Average	48.3	36.7	24.6	38.4	56.7
LDA	Paper	31.4	48.6	42.9	13.5	45.8
AT	Paper	42.9	48.6	42.9	13.1	49.3
	Author	82.9	45.7	25.7	73.5	78.1
	Average	62.9	47.1	34.3	43.3	63.7
ACT1	Paper	42.9	45.7	43.6	16.6	51.0
	Author	91.4	50.0	26.4	80.0	89.6
	Conference	62.9	41.4	23.6	60.7	72.3
	Average	65.7	45.7	31.2	52.4	71.0
ACT1+RW1	Paper	68.6	61.4	50.7	17.1	66.6
	Author	80.0	51.4	27.1	77.6	87.4
	Conference	62.9	42.9	23.6	59.5	72.0
	Average	**70.5**	**51.9**	**33.8**	**51.4**	**75.4**

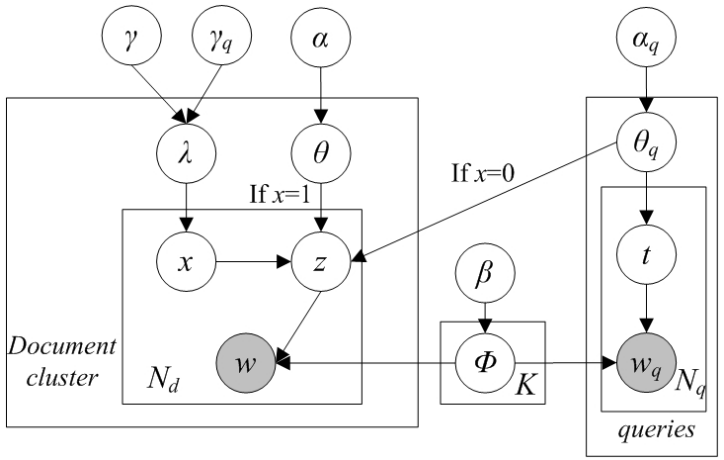
Work3: Multi-topic based Query-oriented Summarization (SDM'09)



qLDA – Query Latent Dirichlet Allocation



qLDA



1. For each topic t, draw the word distribution Φ_t from Dirichlet(β)
2. For each query q
   a) Draw a topic distribution θ_q from Dirichlet(α_q)
   b) For each word w_{qi} in query,
      i) Draw a topic t_{qi} from θ_q
      ii) Draw a word w_{qi} from Φ_{tqi}
3. For each document d
   a) Draw a topic mixture θ_d from Dirichlet(α)
   b) Draw the proportion between words associated with a specific query and those associated with its general topic $\lambda=p(x=0|d)\sim beta(\gamma_q, \gamma)$.
   c) For each word w_{di},
      i) Toss a coin x_{di} from λ with $x=bernoulli(\lambda)$
      ii) if $x=0$,
         a) Draw a topic z_{di} from θ_q
         b) Draw a word w_{di} from Φ_{zdi}
      iii) else
         a) Draw a topic z_{di} from θ_d
         b) Draw a word w_{di} from Φ_{zdi}

$$P(t_{qi}|\mathbf{t}_{-qi},\mathbf{w}_q,\alpha_q)=\frac{n_{qt_{qi}}^{-qi}+\alpha_q}{\sum_t(n_{qt}^{-qi}+\alpha_q)}\frac{n_{t_{qi}w_{qi}}^{-qi}+\beta}{\sum_v(n_{t_{qi}v}^{-qi}+\beta)}$$

$$P(z_{w_{di}},x_{w_{di}}=0|\mathbf{w}_d,\mathbf{w}_q,\mathbf{z}_{-di},\gamma,\gamma_q,\alpha_q,\beta)=\frac{n_{d0}^{-di}+\gamma_q}{n_{d0}^{-di}+n_{d1}^{-di}+\gamma_q+\gamma}\frac{n_{d0z_{w_{di}}}^{-di}+n_{qz_{w_{di}}}+\alpha_q}{\sum_z(n_{d0z}^{-di}+n_q+\alpha_q)}\frac{n_{z_{w_{di}}w_{di}}^{-di}+n_{z_{w_{di}}w_{di}}^{q}+\beta}{\sum_v(n_{z_{w_{di}}v}^{-di}+n_{z_{w_{di}}v}^{q}+\beta)}$$

$$P(z_{w_{di}},x_{w_{di}}=1|\mathbf{w}_d,\mathbf{w}_q,\mathbf{z}_{-di},\gamma,\gamma_q,\alpha,\beta)=\frac{n_{d1}^{-di}+\gamma}{n_{d0}^{-di}+n_{d1}^{-di}+\gamma_q+\gamma}\frac{n_{d1z_{w_{di}}}^{-di}+\alpha}{\sum_z(n_{d1z}^{-di}+\alpha)}\frac{n_{z_{w_{di}}w_{di}}^{-di}+n_{z_{w_{di}}w_{di}}^{q}+\beta}{\sum_v(n_{z_{w_{di}}v}^{-di}+n_{z_{w_{di}}v}^{q}+\beta)}$$

Topic Modeling with Regularization

The new objective function:

$$O_\xi(C,q) = -\xi L(C) + (1-\xi)R(C,q)$$

with

$$R(C,q) = \frac{1}{2}\sum_{d\in C} w(d,q)\sum_z (\theta_{qz} - \theta_{dz})^2$$

Algorithm 1: Parameter estimation

1. initialize the parameters (θ, θ_q, ϕ) randomly;
2. run 200 burn-in sampling iterations using LDA [5];
3. run a two-stage training process iteratively:

(a) train the model parameters (θ, θ_q, ϕ) using the objective function $O_1(C,q) = -L(C)$ with a standard Gibbs sampling procedure by setting the Dirichlet prior for each document as

$$\alpha_{dz} = \alpha + \eta\theta_{qz}$$

where θ_{qz} is the z-th dimension of the multinomial of query q with $1 \le z \le T$; α_{dz} is a parameter specific to the topic of document d; and η is a parameter.

(b) fix ϕ, and re-estimate the multinomial θ to minimize O_ξ; we employ the algorithm proposed in [20] to optimize O_ξ, (i.e. run an iterative process to obtain the new θ for each document d by minimizing the objective function again.)

$$\theta_{dz}^{(n+1)} = \mu\theta_{dz}^{(n)} + (1-\mu)w(d,q)\theta_{qz}$$

where μ is a coefficient to smooth the topic distribution between documents and the query.

Results

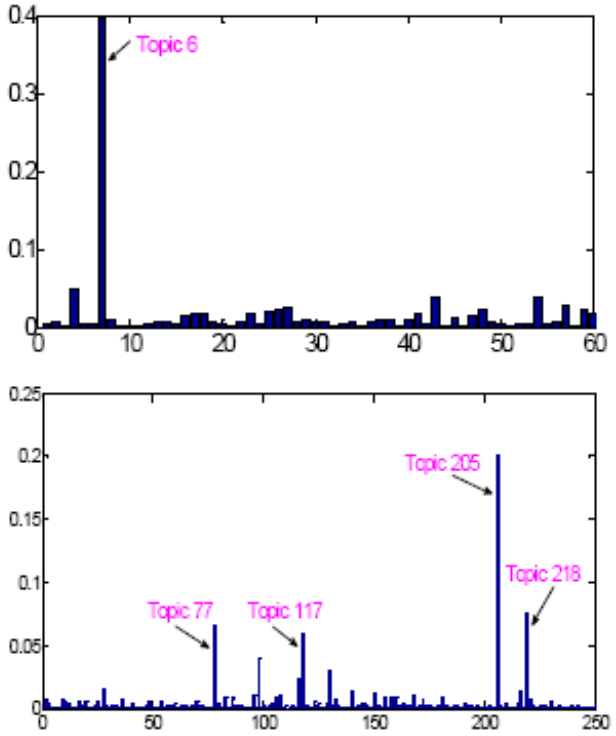


Figure 3: Topic distribution for in D357 (T=60 and T=250). The x axis denotes topics and the y axis denotes the occurrence probability of each topic in D357.

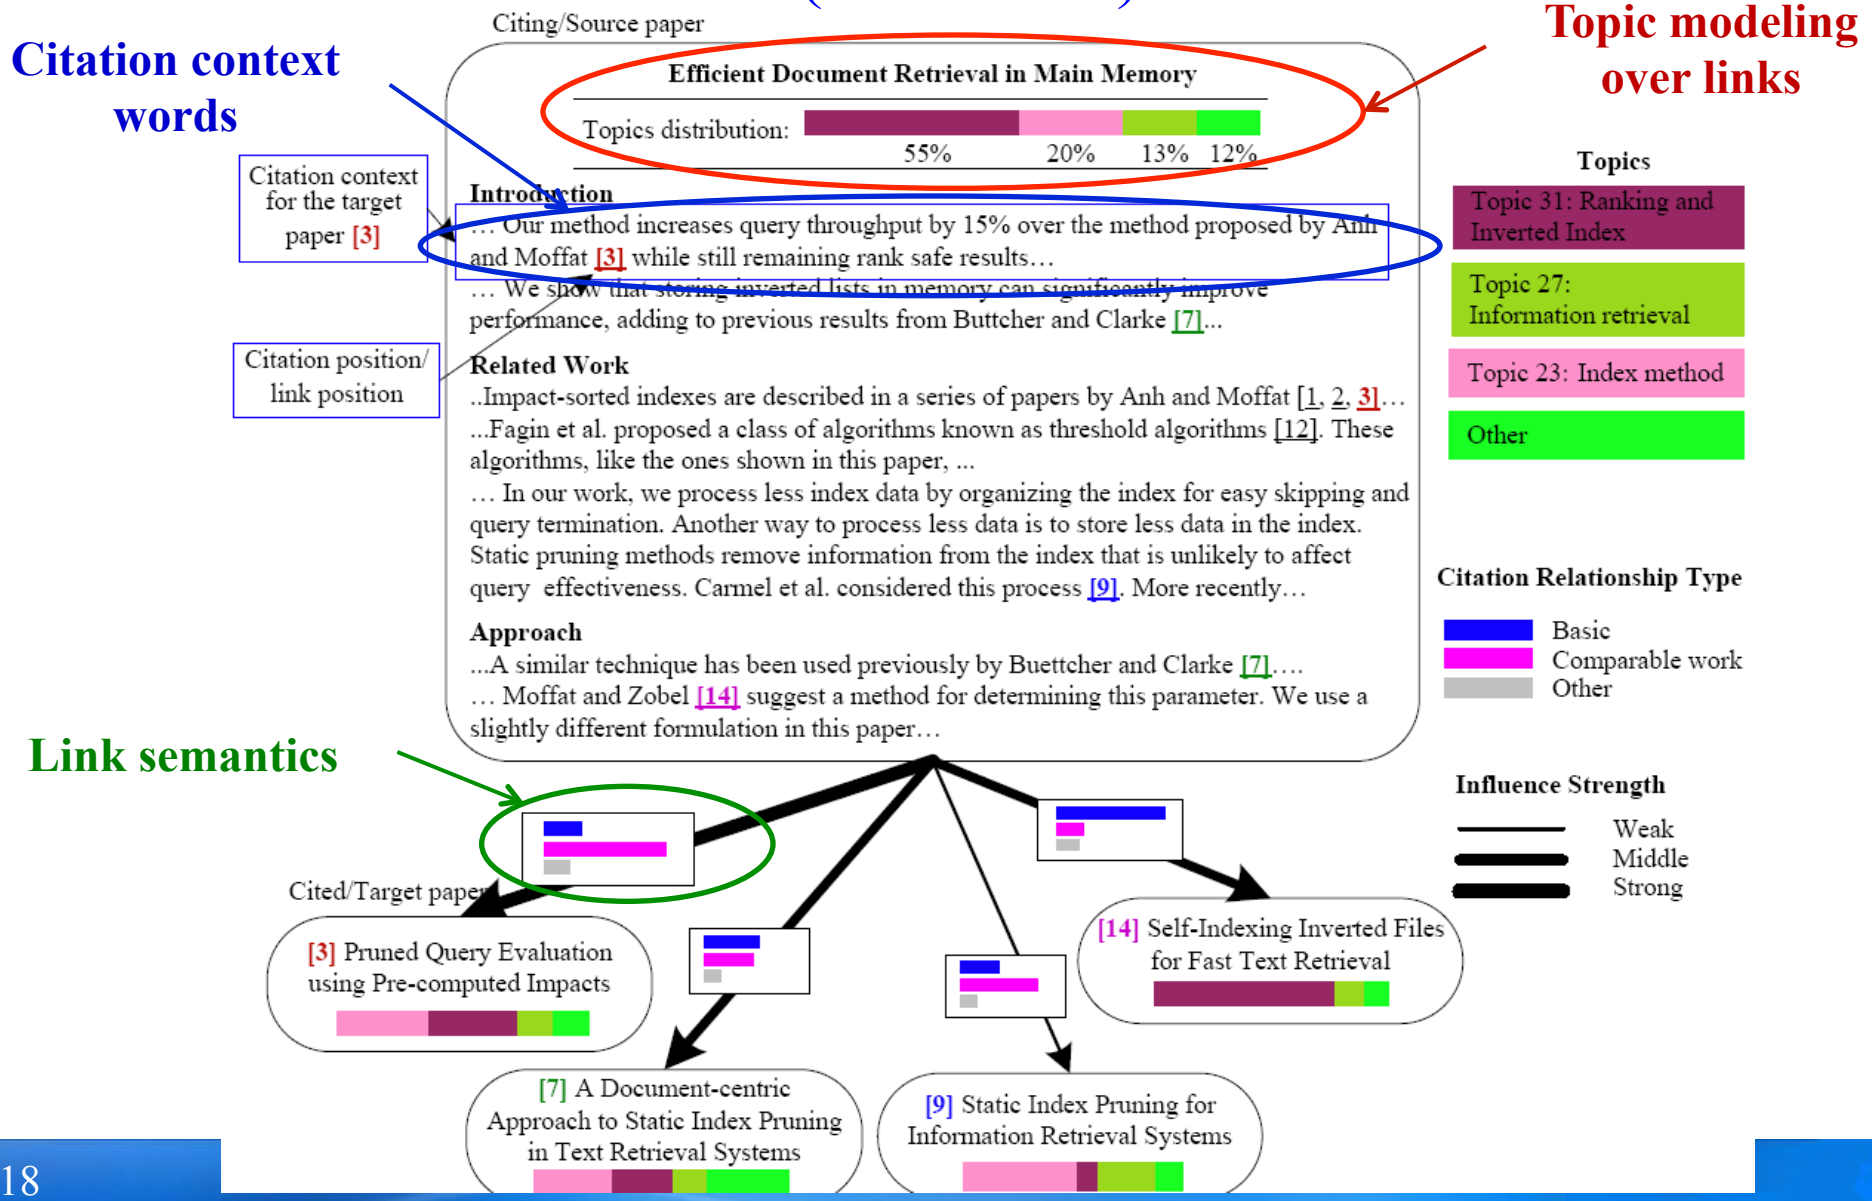
Table 6: Results of different methods on the Epinions data.

Method	Rouge1	Rouge2	RougeSU4
TF	0.28486	0.07579	0.10588
pLSI (Max Score)	0.28763	0.08031	0.10215
pLSI (Sum Score)	0.28291	0.07628	0.10072
pLSI+TF (Max Score)	0.28653	0.07956	0.10572
pLSI+TF (Sum Score)	0.29104	0.07832	0.10721
LDA (Max Score)	0.29155	0.08218	0.11245
LDA (Sum Score)	0.28761	0.07827	0.10671
LDA+TF (Max Score)	0.28483	0.07228	0.10403
LDA+TF (Sum Score)	0.29305	0.07802	0.10796
qLDA (Max Score)	0.29113	0.07765	0.10662
qLDA (Sum Score)	0.28815	0.08189	0.10803
qLDA+TF (Max Score)	0.29698	0.07430	0.10741
qLDA +TF (Sum Score)	0.30326	0.08774	0.11611
TMR (Max Score)	0.30453	0.08349	0.11590
TMR (Sum Score)	0.31271	0.08960	0.11931
TMR+TF (Max Score)	0.31647	0.08815	0.11758
TMR +TF (Sum Score)	**0.32375**	**0.09081**	**0.12894**

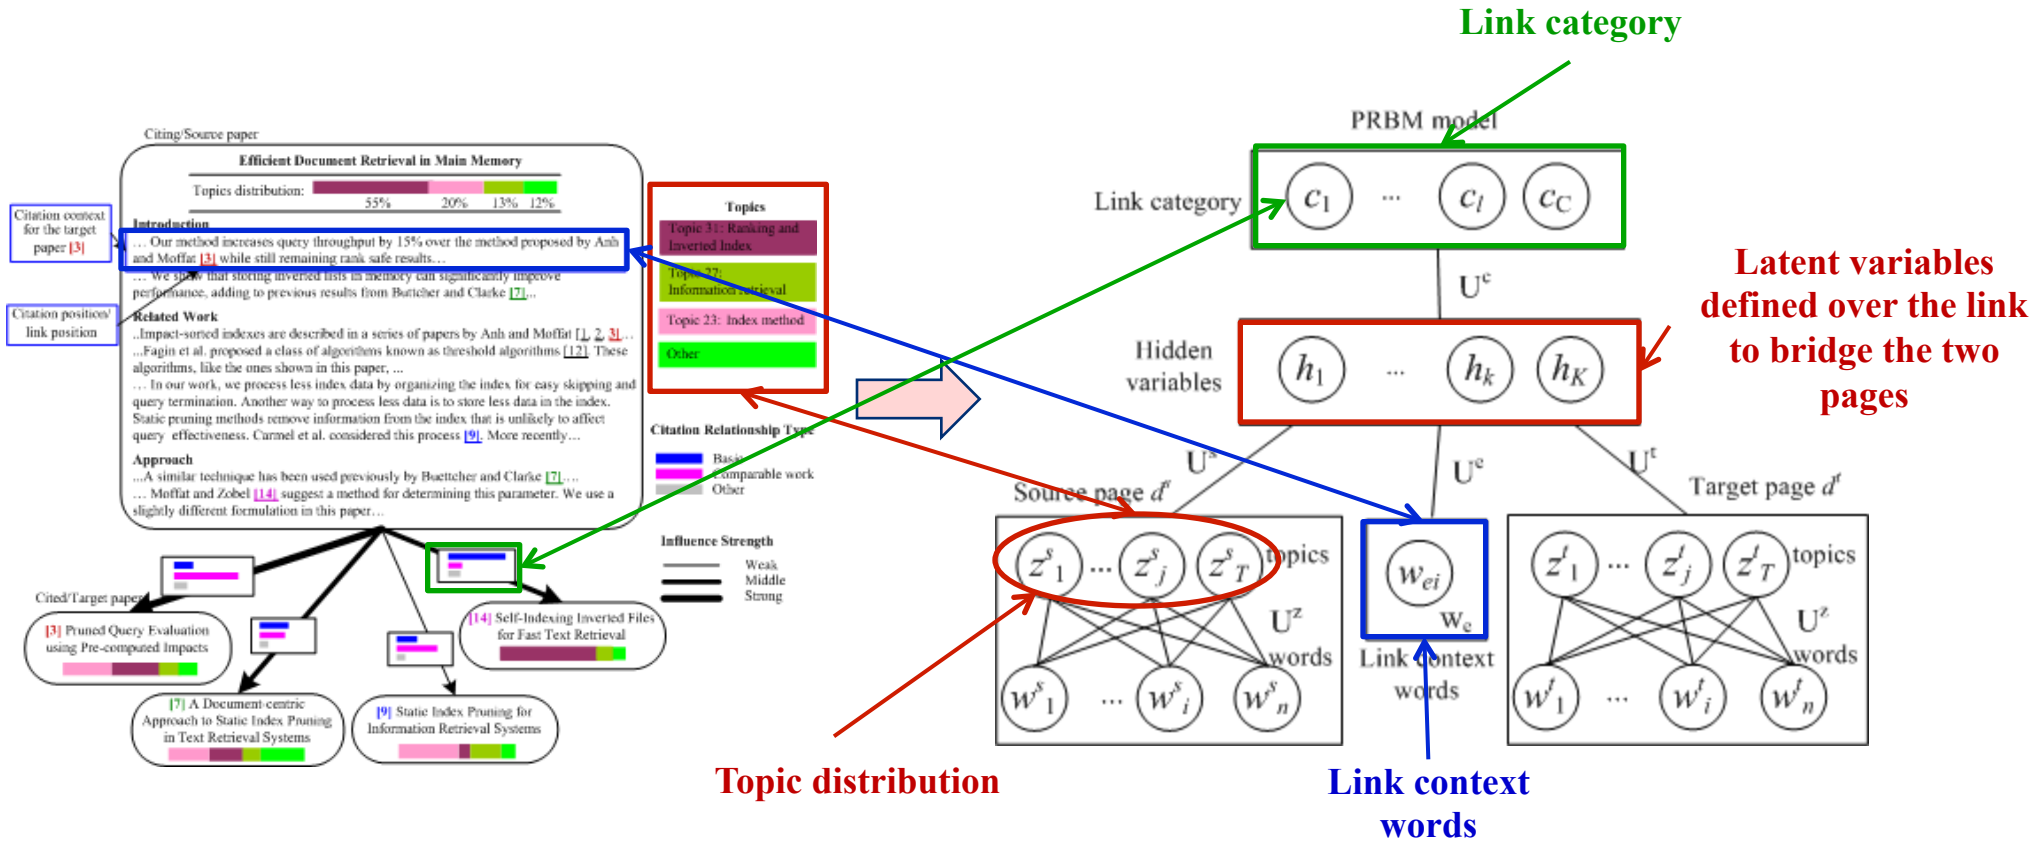
DUC 05 (document understanding conference)

Work4: Topic Distributions over Links on Web

(ICDM'09)



Pairwise Restricted Boltzmann Machines (PRBMs)



Example

Pairwise Restricted Boltzmann Machines (PRBMs)

Results: Accuracy of Link Categorization

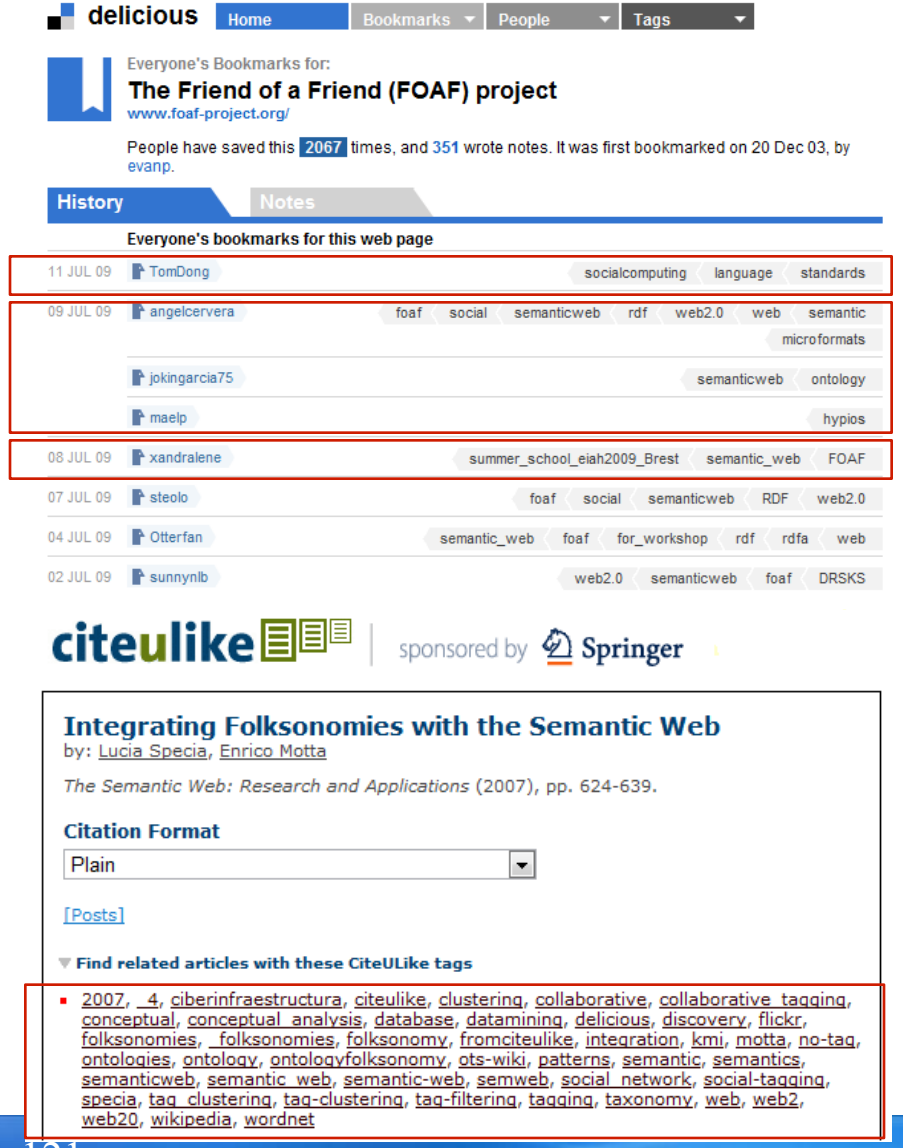
Approach	Type	Precision	Recall	F1-measure	Accuracy
SVM	drill down	64.09	53.60	58.24	88.50
	similar	77.66	83.34	80.33	76.59
	other	64.63	59.58	61.77	79.71
SVM+LDA	dı				
SVM+RBM	dı				
gPRBM	dı				
dPRBM	dı				
hPRBM	dı				

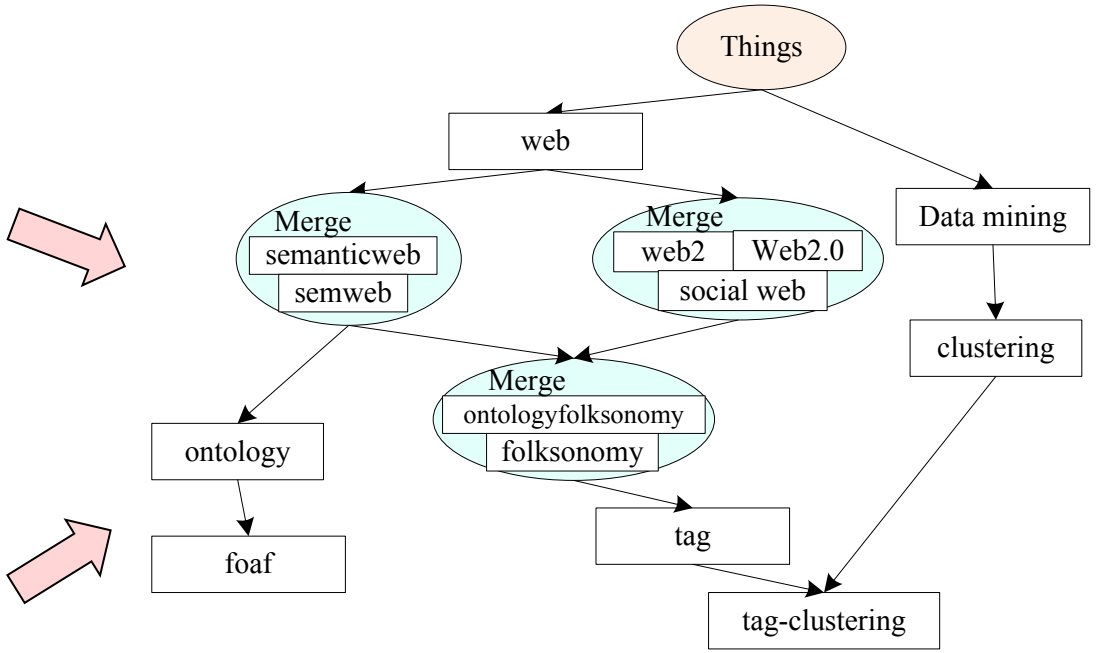
Topics (z^s) of Citing Papers	Topics (z^t) of Cited Papers	Influential Strength
Drill down (Basic theory)		
Topic 37 (Probabilistic model)	Topic 51 (Maximum entropy)	5.05
Topic 37 (Probabilistic model)		
Topic 13 (Semantic Web)		
Sin		
Topic 7 (Recommendation)		
Topic 21 (Frequent pattern learning)		
Topic 26 (Ranking & Inverted Index)		
Topic 26 (Ranking & Inverted Index)		
Topic 26 (Ranking & Inverted Index)		
Topic 9 (Clustering)		

proach

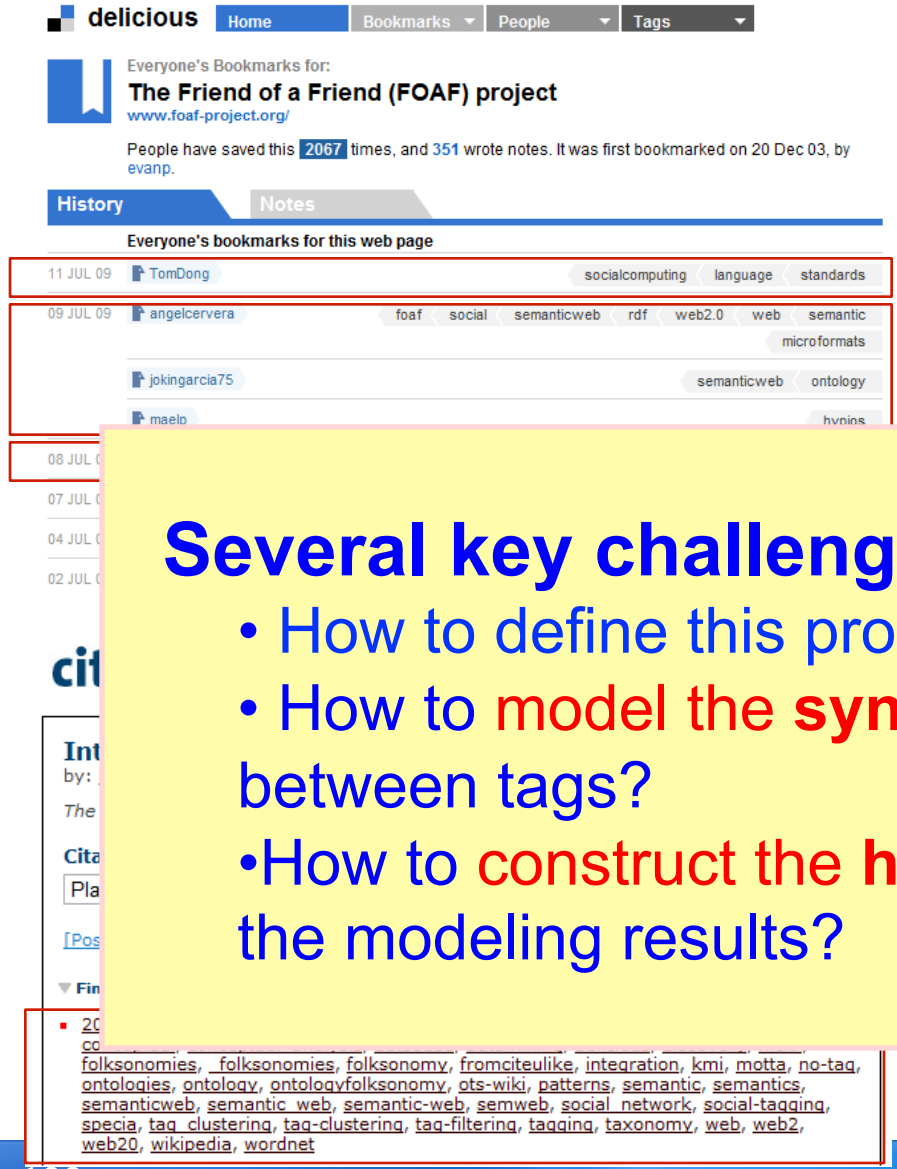
Source Paper	Target Paper	Influence
Drill Down (Basic theory)		
Latent dirichlet allocation	A variational bayesian framework for graphical models	2.36
A sentiment-aware model for predicting sales	Probabilistic latent semantic analysis	1.93
Mining and summarizing customer reviews	Foundations of statistical natural language processing	3.25
Similar (Comparable work)		
Constraint-driven clustering	Max-min d-cluster formation in wireless ad hoc networks	2.00
Hmms and coupled hmms for multi-channel eeg classification	Gaussian observation hidden markov models for eeg analysis	0.51
Scalable collaborative filtering using cluster-based smoothing	Evaluating collaborative filtering recommender systems	1.51
Other		
Latent dirichlet allocation	Overview of the first text retrieval conference	3.93
Utility scoring of product reviews	Introduction to modern information retrieval	3.16
Sentiment analyzer: extracting sentiments about a given topic	Finding parts in very large corpora	2.71

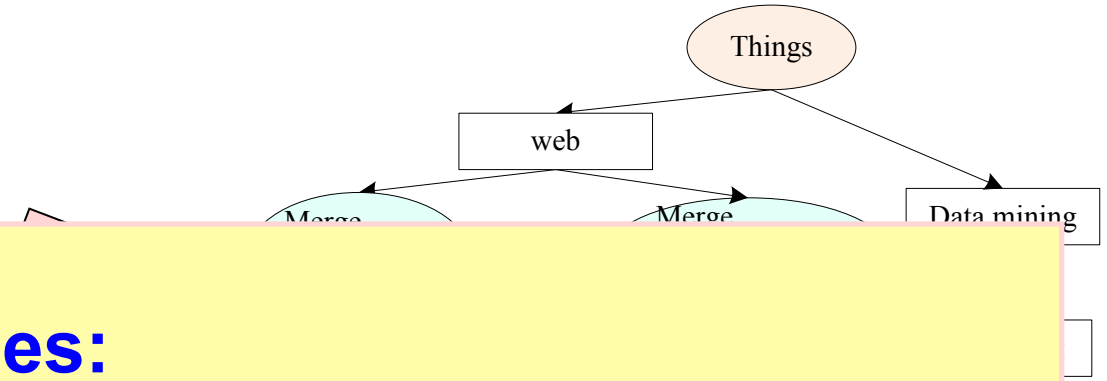
Work5: Towards Ontology Learning from Folksonomies (IJCAI'09)





Towards Ontology Learning from Folksonomies



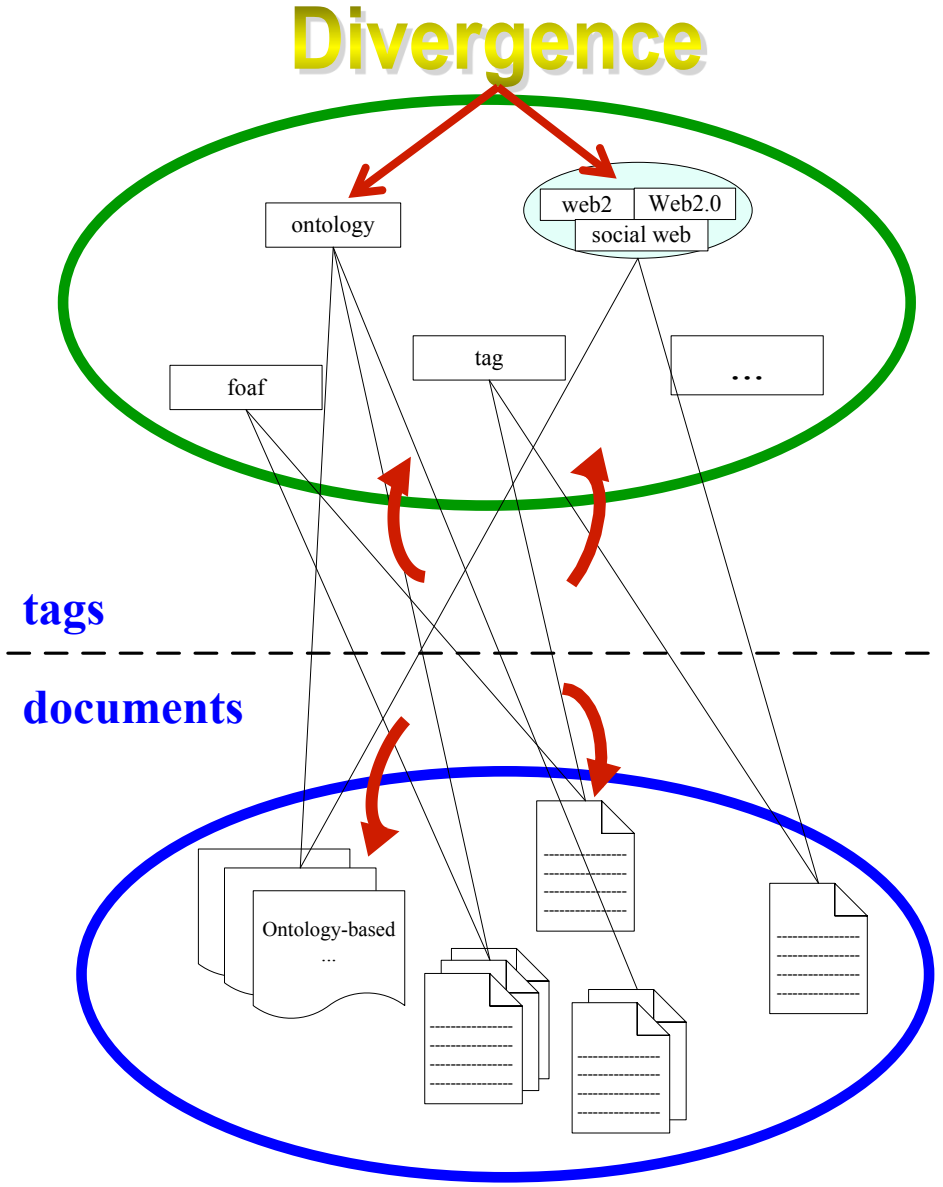


Several key challenges:

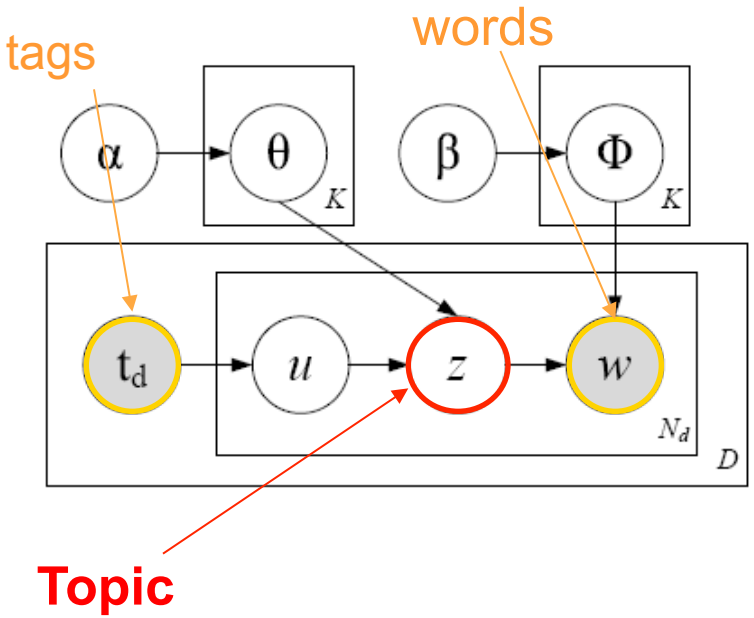
- How to define this problem in a **principled** way?
- How to model the **synonym/hypernym/homonym** between tags?
- How to construct the **hierarchical ontology** according to the modeling results?

Our Solution

1. Use topic to model tags and documents.
2. Define four divergence measures to estimate the difference between tags.
3. Present an algorithm to construct the hierarchical structure from the tags.



Tag-Topic (TT) Models



Tag-Topic (TT) Models

Generative process:

1. For each topic z, draw a multinomial distribution ϕ_z from a Dirichlet prior β_z;
2. For each word w_{di} in the document d:
   (a) draw a tag u_{di} from t_d uniformly;
   (b) draw a topic from a multinomial distribution $\theta_{u_{di}}$ specific to the tag u_{di}, where θ is generated from a Dirichlet prior α;
   (c) draw a word w_{di} from the multinomial distribution $\phi_{z_{di}}$.

$$P(z_{di}, u_{di}|\mathbf{z}_{-di}, \mathbf{u}_{-di}, \mathbf{w}, \mathbf{t}, \alpha, \beta) = \frac{P(\mathbf{z}, \mathbf{u}, \mathbf{w}, \mathbf{t}|\alpha, \beta)}{P(\mathbf{z}_{-di}, \mathbf{u}_{-di}, \mathbf{w}, \mathbf{t}|\alpha, \beta)}$$
$$\propto \frac{m_{uz}^{-di} + \alpha_z}{\sum_z (m_{uz}^{-di} + \alpha_z)} \frac{n_{zv}^{-di} + \beta_v}{\sum_v (n_{zv}^{-di} + \beta_v)}$$

Topic Smoothing

The new objective function:

$$O_\epsilon(D,G) = -\epsilon L(D) + (1-\epsilon) R(D,G)$$

with

$$R(D,G) = \frac{1}{2} \sum_{(t_i,t_j)\in G} w(t_i,t_j) \sum_{k=1}^{K} (\theta_{t_i Z_k} - \theta_{t_j Z_k})^2$$

Smoothing term

Log-likelihood of the tag-topic (TT) model.

Algorithm: Parameter estimation

1. initialize the parameters (θ, ϕ) randomly;
2. run 200 burn-in iterations using Equation (4);
3. run a two-stage training iteratively:

(a) train the model parameters (θ, ϕ) using the objective function $O_1(D,G) = -L(D)$ with a standard Gibbs sampling procedure (Equation (4)) by setting the Dirichlet prior for each tag t_i as

$$\alpha_{t_i k} = \epsilon\alpha + (1-\epsilon)\frac{K}{|G_{t_i}|} \sum_{(t_i,t_j)\in G} \theta_{t_j k}$$

where θ_{tk} is the k-th dimension of the multinomial of tag t with $1 \le k \le K$, $|G_t|$ is the number of neighbors of tag t in the tag network G;

(b) fix ϕ, and re-estimate the multinomial θ to minimize O_ϵ; we employ the algorithm proposed in [Mei *et al.*, 2008] to optimize O_ϵ, (i.e. run an iterative process to obtain the new θ for each tag t by minimizing the objective function again.)

$$\theta_{t_i k}^{(n+1)} = \gamma\theta_{t_i k}^{(n)} + (1-\gamma)\frac{\sum_{(t_i,t_j)\in G} w(t_i,t_j)\theta_{t_j k}^{(n)}}{\sum_{(t_i,t_j)\in G} w(t_i,t_j)}$$

where γ is a coefficient to smooth the topic distribution among the neighbors.

Divergence Estimation

- Tag divergence

$$diverg(t_i, t_j) = D_{KL}(\theta_{t_i} \| \theta_{t_j}) = \sum_{z=1}^{K} \theta_{t_i z} log \frac{\theta_{t_i z}}{\theta_{t_j z}}$$

Estimated topic distribution

- Hypernym-divergence

$$hyper - diverg(t_i, t_j) = \sum_{z=1}^{K} \frac{P(t_i|z_k) - P(t_j|z_k)}{P(t_i|z_k)}$$

Posterior probability derived from the topic modeling results

- Merging-divergence

$$merg - diverg(t_i, t_j) = \frac{1}{2}(diverg(t_i, t_j) + diverg(t_j, t_i))$$

- Keep-divergence

$$keep - diverg(t_i, t_j) = \frac{1}{2}(merg - diverg(t_i, t_j) + (max_k(diverg(t_i, t_k) - diverg(t_i, t_k))))$$

Hierarchical Structure Construction

1. initialize the tree **O** by taking all tags as the leaf nodes;
2. initialize a concept set **A** with all tags;
3. do{

Step 1.

(a) find a pair of tags (or virtual tag/concepts) (t_i, t_j) with an operation from A that minimizes the objective function (10);

Step 2.

(b) execute the operation:

i. for Subordinate, assign t_i as the hypernym of t_j, remove t_j from **A**, and move ti as the parent node of t_j in **O**;

ii. for Merge, create a virtual concept c_{ij} by combining t_i and t_j, add it to **A** as well to **O** as a leaf node, remove t_i and t_j from both **A** and **O**, and calculate the divergences of the new concept c_{ij} with all others tags in **A**.

iii. for Keep, create a virtual concept c_{ij} as the hypernym of both t_i and t_j, add c_{ij} as the parent node of t_i and t_j in **O**, add c_{ij} to **A**, remove t_i and t_j from **A**, and calculate the divergences of the new concept c_{ij} with all others tags in **A**.

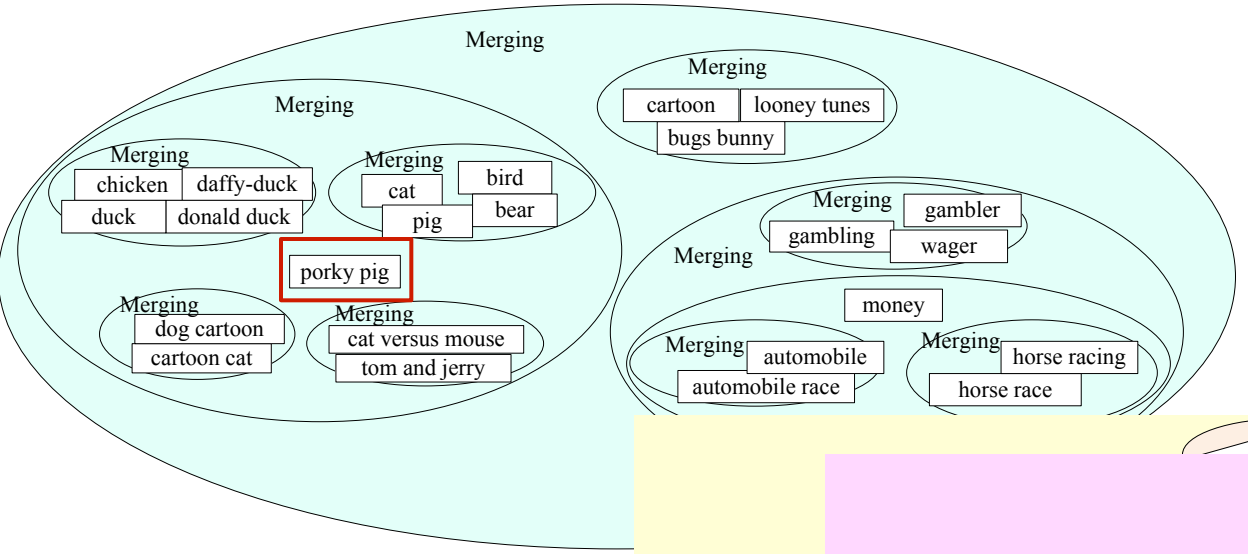
4. }until $|\mathbf{A}| \leq 1$; //$|\mathbf{A}|$ is the number of elements in **A**.
5. return the hierarchical tree **O**.

Correspond to a divergence

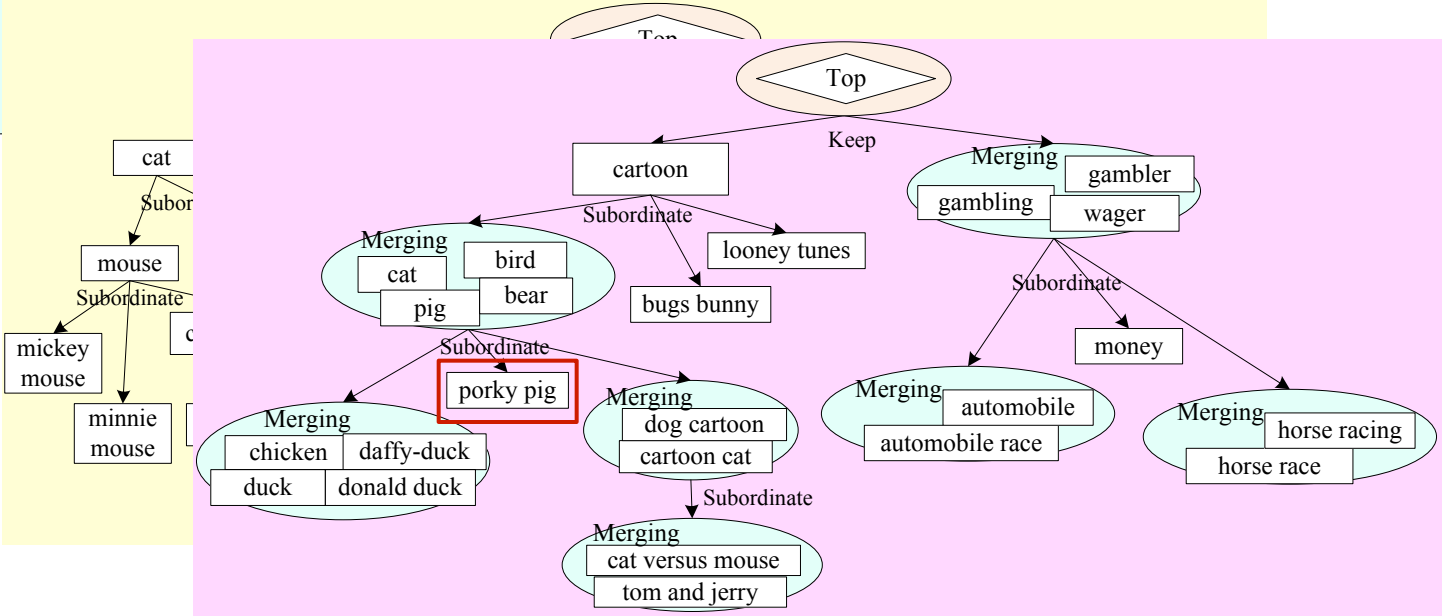
$$argmin_{i \neq j} Operation(t_i, t_j) + \lambda log(N)$$

Penalty to the complex of the generated hierarchy

Case Study—Movie



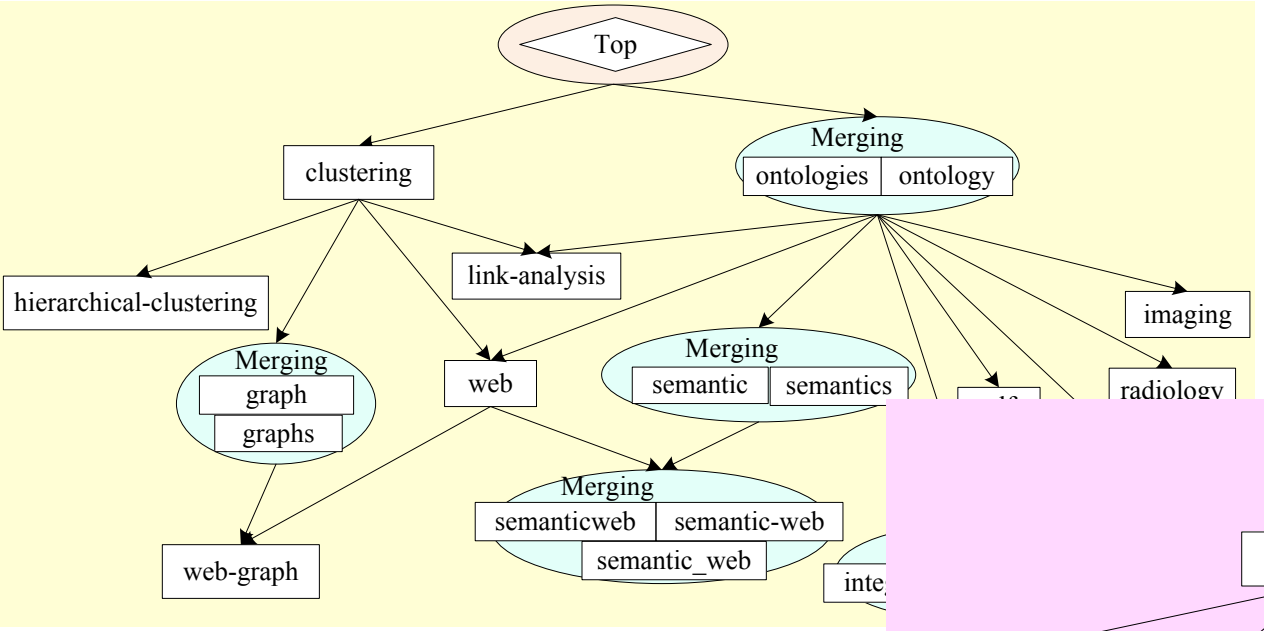
By clustering



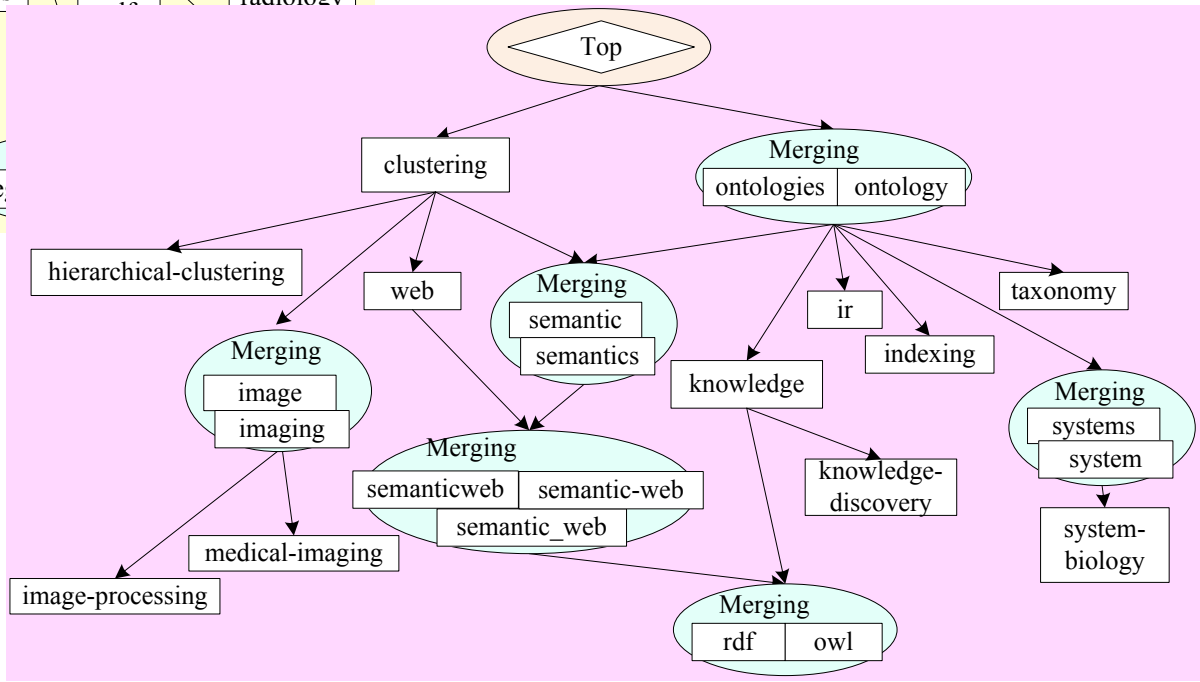
By TT

By TT with smoothing

Case Study—Paper



By TT



By TT with smoothing

Suggested Reading List

Reading List

Red color-good theoretical analysis

Blue color-good applications.

@ARTICLE{Blei:03,
AUTHOR = "David~M. Blei and Andrew~Y. Ng and Michael~I. Jordan",
TITLE = "Latent Dirichlet Allocation",
JOURNAL = "Journal of Machine Learning Research",
VOLUME = {3},
PAGES = {993-1022},
YEAR = {2003} }

@INPROCEEDINGS{ l. griffiths:integrating,
AUTHOR = "Thomas L. Griffiths and Mark Steyvers and David M. Blei and Joshua B. Tenenbaum",
TITLE = "Integrating Topics and Syntax",
booktitle = "NIPS'04“
YEAR = {2004} }

@INPROCEEDINGS{Blei2:06,
AUTHOR = "David~M. Blei and John~D. Lafferty",
TITLE = "Dynamic Topic Models",
BOOKTITLE = "Proceedings of the 23th International Conference on Machine Learning (ICML'06)",
PAGES = "113-120",
YEAR = {2006} }

@INPROCEEDINGS{Blei:07,
AUTHOR = "David~M. Blei and Jon~D. McAuliffe",
TITLE = "Supervised Topic Models",
BOOKTITLE = "Proceedings of Advances in Neural Information Processing Systems",
YEAR = {2007} }

Reading List (cont.)

```
@INPROCEEDINGS{Griffiths:04PNAS,
   AUTHOR = "Thomas~L. Griffiths and Mark Steyvers",
   TITLE = "Finding Scientific Topics",
   BOOKTITLE = "Proceedings of the National Academy of Sciences",
   PAGES = "5228-5235",
   YEAR = {2004}   }

@INPROCEEDINGS{Griffiths:04NIPS,
   AUTHOR = "Thomas~L. Griffiths and Mark Steyvers and David~M. Blei and Joshua~B",
   TITLE = "Integrating Topics and Syntax",
   BOOKTITLE = "Proceedings of Advances in Neural Information Processing Systems",
   YEAR = {2004}   }

@INPROCEEDINGS{Gruber:07,
   AUTHOR = "Amit Gruber and Michal Rosen-Zvi and Yair Weiss",
   TITLE = "Hidden Topic Markov Models",
   BOOKTITLE = "Proceedings of Artificial Intelligence and Statistics (AISTATS'07)",
   YEAR = {2007}   }
@ARTICLE{McCallum:07,
   AUTHOR = "Andrew McCallum and Xuerui Wang and Andres Corrada-Emmanuel",
   TITLE = "Topic and Role Discovery in Social Networks with Experiments on Enron and Academic Email",
   JOURNAL = {Journal of Artificial Intelligence Research (JAIR)},
   VOLUME = {30},
   PAGES = {249-272},
   YEAR = {2007} }
```

Reading List (cont.)

```
@inproceedings{Mei:07,
    author = {Qiaozhu Mei and Xu Ling and Matthew Wondra and Hang Su and ChengXiang Zhai},
    title = {Topic sentiment mixture: modeling facets and opinions in weblogs},
    year = {2007},
    pages = {171--180},
    booktitle = {Proceedings of WWW'07},
%     booktitle = {Proceedings of the 16th international conference on World Wide Web (WWW'07)},
 }

@INPROCEEDINGS{Mei:08WWW,
   AUTHOR = "Qiaozhu Mei and Deng Cai and Duo Zhang and ChengXiang Zhai",
   TITLE = "Topic Modeling with Network Regularization",
   PAGES = "101-110",
   YEAR = {2008},
   BOOKTITLE = "Proceedings of the 17th International World Wide Web Conference (WWW'08)",
}

@INPROCEEDINGS{Mei:08ACL,
   AUTHOR = "Qiaozhu Mei and ChengXiang Zhai",
   TITLE = "Generating impact-cased summaries for scientific literature",
   BOOKTITLE ={Proceedings of the 46nd Annual Meeting of the Association for Computational Linguistics (ACL'08)},
   PAGES = "816-824",
   YEAR = {2008}}
```

Reading List (cont.)

```
@INPROCEEDINGS{Mika:05,
        author = {Peter Mika},
        title = {Ontologies Are Us: A Unified Model of Social Networks and Semantics},
        BOOKTITLE = {Proceedings of the 4th International Semantic Web Conference (ISWC'05)},
        pages = {522--536},
        year = {2005}
}

@INPROCEEDINGS{Mimno:07,
   AUTHOR = "David Mimno and Andrew McCallum",
   TITLE = "Expertise Modeling for Matching Papers with Reviewers",
   PAGES = "500-509",
   YEAR = {2007},
   BOOKTITLE = "KDD'07",
}

@INPROCEEDINGS{Minka:02,
   AUTHOR = "Thomas Minka and John Lafferty",
   TITLE = "Expectation-Propagation for the Generative Aspect Model",
   BOOKTITLE = "Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI'2002)",
   PAGES = "352-359",
   YEAR = {2002}}
```

Reading List (cont.)

@INPROCEEDINGS{Minka:03,
AUTHOR = "Thomas Minka",
TITLE = "Estimating a Dirichlet Distribution",
BOOKTITLE = "Technique Report, http://research.microsoft.com/~minka/papers/dirichlet/",
YEAR = {2003} }

@INPROCEEDINGS{Nallapati:07,
AUTHOR = "Ramesh~M. Nallapati and Susan Ditmore and John~D. Lafferty and Kin Ung",
TITLE = "Multiscale Topic Tomography",
PAGES = "520-529",
YEAR = {2007},
BOOKTITLE = "KDD'2007",
}
@INPROCEEDINGS{Nallapati:08,
AUTHOR = "Ramesh~M. Nallapati and Amr Ahmed and Eric Xing and William~W. Cohen",
TITLE = "Joint Latent Topic Models for Text and Citations",
BOOKTITLE = "Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'2008)",
Year = “2008”}

@inproceedings{Newman:07,
Author = {David Newman and Arthur Asuncion and Padhraic Smyth and Max Welling},
Title = {Distributed inference for Latent Dirichlet Allocation},
Year = {2007},
Booktitle = {Proceedings of NIPS'07},
}

Reading List (cont.)

```
@inproceedings{Porteous:08,
    author = {Ian Porteous and David Newman and Alexander Ihler and Arthur Asuncion and Padhraic Smyth and Max Welling},
    title = {Fast collapsed gibbs sampling for latent dirichlet allocation},
    booktitle = {Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD'08)},
    year = {2008},
    pages = {569--577},
 }
@INPROCEEDINGS{Nie:06,
   AUTHOR = "Lan Nie and Brian~D. Davison and Xiaoguang Qi",
   TITLE = "Topical Link Analysis for Web Search",
   PAGES = "91-98",
   YEAR = {2006},
   BOOKTITLE = "SIGIR'2006",
}
@INPROCEEDINGS{Nodelman:02,
   AUTHOR = "Uri Nodelman and Christian~R. Shelton and Daphne Koller",
   TITLE = "Continuous Time Bayesian Networks",
   BOOKTITLE = "Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI'2002)",
   PAGES = "378-387",
   YEAR = {2002}}
@INPROCEEDINGS{Rosen-Zvi:04,
   AUTHOR = "Michal Rosen-Zvi and Thomas Griffiths and Mark Steyvers and Padhraic Smyth",
   TITLE = "The Author-Topic Model for Authors and Documents",
   YEAR = {2004},
   BOOKTITLE = "UAI'04“,}
```

Reading List (cont.)

```
@INPROCEEDINGS{Steyvers:04,
   AUTHOR = "Mark Steyvers and Padhraic Smyth and Thomas Griffiths",
   TITLE = "Probabilistic Author-Topic Models for Information Discovery",
   YEAR = {2004},
   PAGES = {306-315},
   BOOKTITLE = "KDD'04“,}
@INPROCEEDINGS{Teh:04,
   AUTHOR = "Yee~Whye Teh and Michael~I. Jordan and Matthew~J. Beal and David~M. Blei",
   TITLE = "Hierarhical Dirichlet Processes",
   BOOKTITLE = "Technical Report 653, Department of Statistics, UC Berkeley",
   YEAR = {2004}}
@INPROCEEDINGS{Ahmed:09KDD,
   AUTHOR = "Amr Ahmed and Eric P. Xing and William W. Cohen and Robert F. Murphy",
   TITLE = "Structured correspondence topic models for mining captioned figures in biological literature",
   booktitle = "Proceeding of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD'09)"
   PAGES = {39-48},
   YEAR = {2009}
}
@inproceedings{Leskovec:09KDD,
 author = {Leskovec, Jure and Backstrom, Lars and Kleinberg, Jon},
 title = {Meme-tracking and the dynamics of the news cycle},
 booktitle = {Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD'09)},
 year = {2009},
 pages = {497--506},
 }
```

Reading List (cont.)

```
@INPROCEEDINGS{Wallach:06,
   AUTHOR = "Hanna~M. Wallach",
   TITLE = "Topic Modeling: Beyond Bag-of-Words",
   BOOKTITLE = "Proceedings of the 23rd international conference on Machine learning (ICML'06)",
   PAGES = "977-984",
   YEAR = {2006}}

@INPROCEEDINGS{Wang:06,
   AUTHOR = "Xuerui Wang and Andrew McCallum",
   TITLE = "Topics over Time: A Non-Markov Continuous-Time Model of Topical Trends",
   BOOKTITLE = "Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD'06)",
   PAGES = "424-433",
   YEAR = {2006}}
@inproceedings{Wang:07,
    author = {Xuerui Wang and Chris Pal and Andrew McCallum},
    title = {Generalized component analysis for text with heterogeneous attributes},
    booktitle = {Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD'07)},
    year = {2007},
    pages = {794--803},}
@INPROCEEDINGS{Wei:06,
   AUTHOR = "Xing Wei and W. Bruce Croft",
   TITLE = "LDA-Based Document Models for Ad-hoc Retrieval",
   BOOKTITLE = "Proceedings of the 29th ACM SIGIR International Conference on Information Retrieval (SIGIR'06)",
   PAGES = "178-185",
   YEAR = {2006}}
```

Reading List (cont.)

```
@inproceedings{Zhai:04,
    author = {ChengXiang Zhai and Atulya Velivelli and Bei Yu},
    title = {A cross-collection mixture model for comparative text mining},
    year = {2004},
    pages = {743--748},
    booktitle = {KDD'04},
}
@INPROCEEDINGS{Cohn:01,
   AUTHOR = {David Cohn and Thomas Hofmann},
   TITLE = {The Missing Link - A Probabilistic Model of Document Content and Hypertext Connectivity},
   BOOKTITLE = {Proceedings of the 13th Neural Information Processing Systems (NIPS'01)},
   YEAR = {2001}}
@INPROCEEDINGS{Hofmann:99,
   AUTHOR = "Thomas Hofmann",
   TITLE = "Probabilistic Latent Semantic Indexing",
   PAGES = "50-57",
   YEAR = {1999},
   BOOKTITLE = "SIGIR'99“,}
@INPROCEEDINGS{Hofmann:03,
   AUTHOR = "Thomas Hofmann",
   TITLE = "Collaborative Filerting via Gaussian Probabilistic Latent Semantic Analysis",
   BOOKTITLE = "Proceedings of the 26nd International Conference on Research and Development in Information Retrieval (SIGIR'03)",
   PAGES = "259-266",
YEAR = {2003}   }
```

Reading List (cont.)

@INPROCEEDINGS{Tang:08ICDM,
AUTHOR = "Jie Tang and Ruoming Jin and Jing Zhang",
TITLE = "A Topic Modeling Approach and its Integration into the Random Walk Framework for Academic Search",
YEAR = {2008},
BOOKTITLE = "ICDM'08",}
@INPROCEEDINGS{Tang:09SDM,
AUTHOR = "Jie Tang and Limin Yao and Dewei Chen",
TITLE = "Multi-topic based Query-oriented Summarization",
YEAR = {2009},
BOOKTITLE = "Proceedings of 2009 SIAM International Conference on Data Mining (SDM'2009)“,}
@INPROCEEDINGS{Tang:09KDD,
AUTHOR = "Jie Tang and Jimeng Sun and Chi Wang and Zi Yang",
TITLE = "Social Influence Analysis in Large-scale Networks",
YEAR = {2009},
PAGES = "807--816",
BOOKTITLE = "SIGKDD'09“,}
@INPROCEEDINGS{Tang:09ICDM,
AUTHOR = "Jie Tang and Jing Zhang and Jeffrey Xu Yu and Zi Yang and Keke Cai and Rui Ma and Li Zhang and Zhong Su",
TITLE = "Topic Distributions over Links on Web",
YEAR = {2009},
BOOKTITLE = "Proceedings of 2009 IEEE International Conference on Data Mining (ICDM'09)“, }
@inproceedings{Tang:09IJCAI,
author = {Tang, Jie and Leung, Ho-fung and Luo, Qiong and Chen, Dewei and Gong, Jibin},
title = {Towards ontology learning from folksonomies},
booktitle = {Proceedings of the 21st international jont conference on Artifical intelligence (IJCAI'09)},
year = {2009},
pages = {2089--2094}, }

Thanks!

Q&A

HP: http://keg.cs.tsinghua.edu.cn/persons/tj/

http://arnetminer.org