CS249: ADVANCED DATA MINING

Text Data: Topic Models

Instructor: Yizhou Sun

yzsun@cs.ucla.edu

May 10, 2017

Announcements

- Course project proposal
  - Due May 8th (11:59pm)
- Homework 3 out
  - Due May 10th (11:59pm)
- Midterm Exam
  - In class May 15th

Methods to Learn

	Vector Data	Text Data	Recommender System	Graph & Network
Classification	**Decision Tree**; **Naïve Bayes**; **Logistic Regression** **SVM; NN**			Label Propagation
Clustering	**K-means; hierarchical clustering; DBSCAN; Mixture Models; kernel k-means**	**PLSA;** **LDA**	Matrix Factorization	SCAN; Spectral Clustering
Prediction	**Linear Regression** **GLM**		Collaborative Filtering	
Ranking				PageRank
Feature Representation		Word embedding		Network embedding

Text Data: Topic Models

- Text Data and Topic Models
- Revisit of Multinomial Mixture Model
- Probabilistic Latent Semantic Analysis (pLSA)
- Latent Dirichlet Allocation (LDA)
- Summary

Text Data

- Word/term
- Document
  - A sequence of words
- Corpus
  - A collection of documents

Cancer effort honored

Handsworth memorial award set up

LIFESTYLE

The geek grew up

Now Amanda Welliver helps teens bully-proof themselves

Raven in his 80s set to see again after op

Tarquin's blind, has only one eye but has mate 70 years his junior

It's man vs burger

Monster meal deals hard to digest

INDIANAPOLIS NEWS

HUNDREDS IN STATE SEE 'FLYING SAUCERS'

Franklin 'Dogfight' Alerts State Troopers

Represent a Document

- Most common way: Bag-of-Words
  - Ignore the order of words
  - keep the count

c1: *Human* machine *interface* for Lab ABC *computer* applications
c2: A *survey* of *user* opinion of *computer system response time*
c3: The *EPS user interface* management *system*
c4: *System* and *human system* engineering testing of *EPS*
c5: Relation of *user*-perceived *response time* to error measurement

m1: The generation of random, binary, unordered *trees*
m2: The intersection *graph* of paths in *trees*
m3: *Graph minors* IV: Widths of *trees* and well-quasi-ordering
m4: *Graph minors*: A *survey*

	c1	c2	c3	c4	c5	m1	m2	m3	m4
human	1	0	0	1	0	0	0	0	0
interface	1	0	1	0	0	0	0	0	0
computer	1	1	0	0	0	0	0	0	0
user	0	1	1	0	1	0	0	0	0
system	0	1	1	2	0	0	0	0	0
response	0	1	0	0	1	0	0	0	0
time	0	1	0	0	1	0	0	0	0
EPS	0	0	1	1	0	0	0	0	0
survey	0	1	0	0	0	0	0	0	1
trees	0	0	0	0	0	1	1	1	0
graph	0	0	0	0	0	0	1	1	1
minors	0	0	0	0	0	0	0	1	1

Vector space model

More Details

- Represent the doc as a vector where each entry corresponds to a different word and the number at that entry corresponds to how many times that word was present in the document (or some function of it)
  - Number of words is huge
  - Select and use a smaller set of words that are of interest
  - E.g. uninteresting words: ‘and’, ‘the’ ‘at’, ‘is’, etc. These are called stop-words
  - Stemming: remove endings. E.g. ‘learn’, ‘learning’, ‘learnable’, ‘learned’ could be substituted by the single stem ‘learn’
  - Other simplifications can also be invented and used
  - The set of different remaining words is called dictionary or vocabulary. Fix an ordering of the terms in the dictionary so that you can operate them by their index.
  - Can be extended to bi-gram, tri-gram, or so

Limitations of Vector Space Model

- Dimensionality
  - High dimensionality
- Sparseness
  - Most of the entries are zero
- Shallow representation
  - The vector representation does not capture semantic relations between words

Topics

- Topic
  - A topic is represented by a word distribution
  - Relate to an issue

universe	0.0439
galaxies	0.0375
clusters	0.0279
matter	0.0233
galaxy	0.0232
cluster	0.0214
cosmic	0.0137
dark	0.0131
light	0.0109
density	0.01

drug	0.0672
patients	0.0493
drugs	0.0444
clinical	0.0346
treatment	0.028
trials	0.0277
therapy	0.0213
trial	0.0164
disease	0.0157
medical	0.00997

cells	0.0675
stem	0.0478
human	0.0421
cell	0.0309
gene	0.025
tissue	0.0185
cloning	0.0169
transfer	0.0155
blood	0.0113
embryos	0.0111

sequence	0.0818
sequences	0.0493
genome	0.033
dna	0.0257
sequencing	0.0172
map	0.0123
genes	0.0122
chromosome	0.0119
regions	0.0119
human	0.0111

years	0.156
million	0.0556
ago	0.045
time	0.0317
age	0.0243
year	0.024
record	0.0238
early	0.0233
billion	0.0177
history	0.0148

bacteria	0.0983
bacterial	0.0561
resistance	0.0431
coli	0.0381
strains	0.025
microbiol	0.0214
microbial	0.0196
strain	0.0165
salmonella	0.0163
resistant	0.0145

male	0.0558
females	0.0541
female	0.0529
males	0.0477
sex	0.0339
reproductive	0.0172
offspring	0.0168
sexual	0.0166
reproduction	0.0143
eggs	0.0138

theory	0.0811
physics	0.0782
physicists	0.0146
einstein	0.0142
university	0.013
gravity	0.013
black	0.0127
theories	0.01
aps	0.00987
matter	0.00954

immune	0.0909
response	0.0375
system	0.0358
responses	0.0322
antigen	0.0263
antigens	0.0184
immunity	0.0176
immunology	0.0145
antibody	0.014
autoimmune	0.0128

stars	0.0524
star	0.0458
astrophys	0.0237
mass	0.021
disk	0.0173
black	0.0161
gas	0.0149
stellar	0.0127
astron	0.0125
hole	0.00824

Topic Models

- Topic modeling
  - Get topics automatically from a corpus
  - Assign documents to topics automatically
- Most frequently used topic models
  - pLSA
  - LDA

“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.

Text Data: Topic Models

- Text Data and Topic Models
- Revisit of Multinomial Mixture Model
- Probabilistic Latent Semantic Analysis (pLSA)
- Latent Dirichlet Allocation (LDA)
- Summary

Review of Multinomial Distribution

- Select n data points from K categories, each with probability p_k
  - n trials of independent categorical distribution
    - E.g., get 1-6 from a dice with 1/6
  - When n=1, it is also called discrete distribution
- When K=2, binomial distribution
  - n trials of independent Bernoulli distributio
  - E.g., flip a coin to get heads or tails

Multinomial Mixture Model

- For documents with bag-of-words representation
  - $\boldsymbol{x}_d = (x_{d1}, x_{d2}, \dots, x_{dN})$, x_{dn} is the number of words for nth word in the vocabulary
- Generative model
  - For each document
    - Sample its cluster label $z \sim discrete(\boldsymbol{\pi})$
      - $\boldsymbol{\pi} = (\pi_1, \pi_1, \dots, \pi_K)$, π_k is the proportion of jth cluster
      - $p(z = k) = \pi_k$
    - Sample its word vector $\boldsymbol{x}_d \sim multi(\boldsymbol{\beta}_z)$
      - $\boldsymbol{\beta}_z = (\beta_{z1}, \beta_{z2}, \dots, \beta_{zN})$, β_{zn} is the parameter associate with nth word in the vocabulary
      - $p(\boldsymbol{x}_d | z = k) = \frac{(\sum_n x_{dn})!}{\prod_n x_{dn}!} \prod_n \beta_{kn}^{x_{dn}} \propto \prod_n \beta_{kn}^{x_{dn}}$

Likelihood Function

- For a set of M documents

$$L = \prod_d p(\boldsymbol{x}_d) = \prod_d \sum_k p(\boldsymbol{x}_d, z = k)$$

$$= \prod_d \sum_k p(\boldsymbol{x}_d | z = k) p(z = k)$$

$$\propto \prod_d \sum_k p(z = k) \prod_n \beta_{kn}^{x_{dn}}$$

Mixture of Unigrams

- For documents represented by a sequence of words
  - $\boldsymbol{w}_d = (w_{d1}, w_{d2}, \dots, w_{dN_d})$, N_d is the length of document d, w_{dn} is the word at the nth position of the document
- Generative model
  - For each document
    - Sample its cluster label $z \sim discrete(\boldsymbol{\pi})$
      - $\boldsymbol{\pi} = (\pi_1, \pi_1, \dots, \pi_K)$, π_k is the proportion of jth cluster
      - $p(z = k) = \pi_k$
    - For each word in the sequence
      - Sample the word $w_{dn} \sim discrete(\boldsymbol{\beta}_z)$
      - $p(w_{dn} | z = k) = \beta_{k w_{dn}}$

Likelihood Function

- For a set of M documents

$$L = \prod_d p(\boldsymbol{w}_d) = \prod_d \sum_k p(\boldsymbol{w}_d, z = k)$$

$$= \prod_d \sum_k p(\boldsymbol{w}_d | z = k) p(z = k)$$

$$= \prod_d \sum_k p(z = k) \prod_n \beta_{k w_{dn}}$$

Text Data: Topic Models

- Text Data and Topic Models
- Revisit of Multinomial Mixture Model
- Probabilistic Latent Semantic Analysis (pLSA)
- Latent Dirichlet Allocation (LDA)
- Summary

Notations

- Word, document, topic
  - w, d, z
- Word count in document
  - $c(w, d)$
- Word distribution for each topic (β_z)
  - $\beta_{zw}: p(w|z)$
- Topic distribution for each document (θ_d)
  - $\theta_{dz}: p(z|d)$ (Yes, soft clustering)

Issues of Mixture of Unigrams

- All the words in the same documents are sampled from the same topic

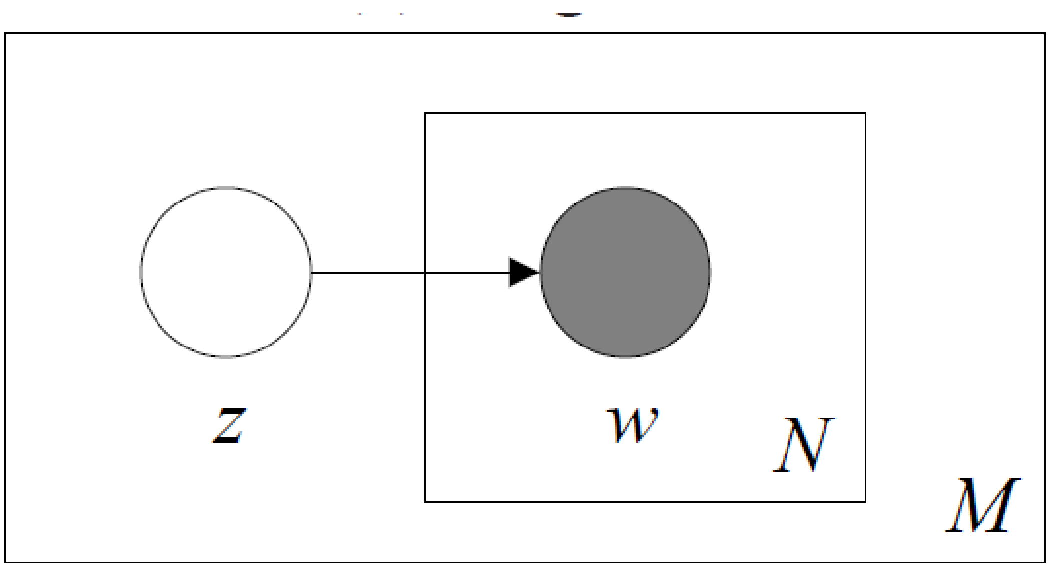

- In practice, people switch topics during their writing

Illustration of pLSA

"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.

Generative Model for pLSA

- Describe how a document is generated probabilistically
  - For each position in d, $n = 1, \ldots, N_d$
    - Generate the topic for the position as
      $z_n \sim discrete(\boldsymbol{\theta}_d), i.e., p(z_n = k) = \theta_{dk}$
      (Note, 1 trial multinomial, i.e., discrete distribution)
    - Generate the word for the position as
      $w_n \sim discrete(\boldsymbol{\beta}_{z_n}), i.e., p(w_n = w) = \beta_{z_n w}$

Graphical Model

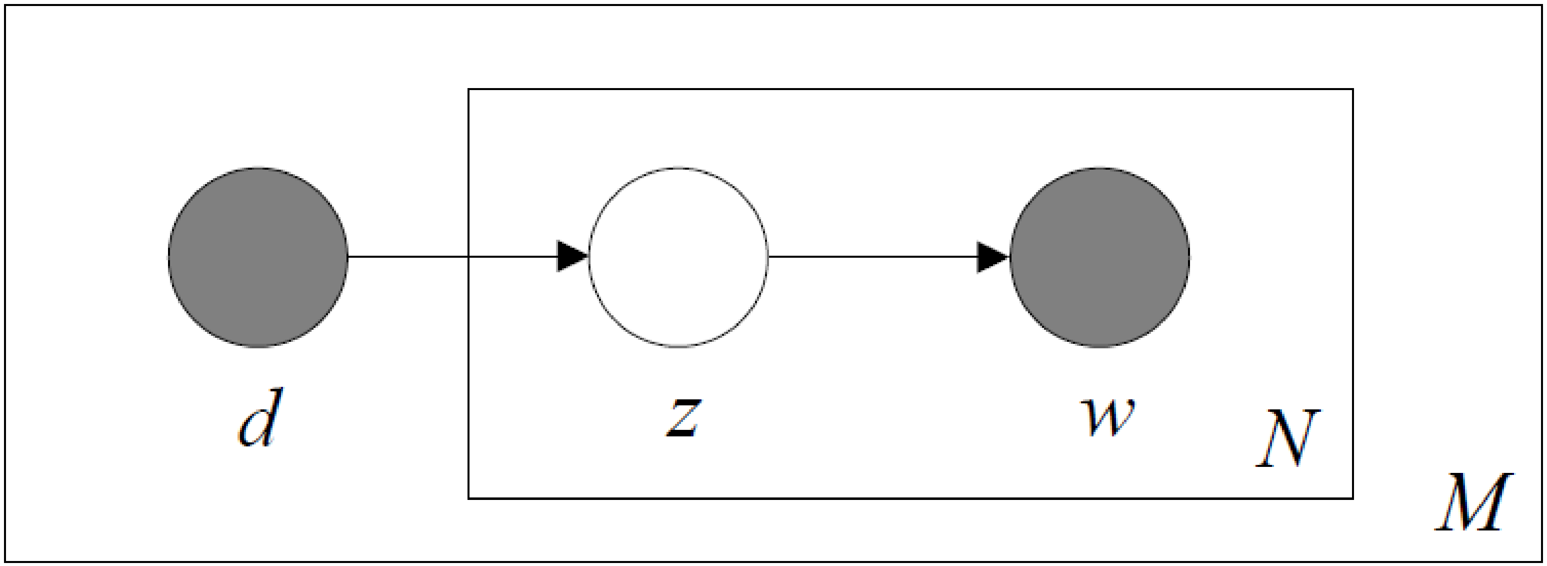

Note: Sometimes, people add parameters such as θ *and* β into the graphical model

The Likelihood Function for a Corpus

- Probability of a word

$$p(w|d) = \sum_k p(w, z = k|d) = \sum_k p(w|z = k)p(z = k|d) = \sum_k \beta_{kw}\theta_{dk}$$

- Likelihood of a corpus

$$\begin{aligned} & \prod_{d=1} P(w_1, \cdots, w_{N_d}, d|\theta, \beta, \pi) \\ = & \prod_{d=1} P(d) \left\{ \prod_{n=1}^{N_d} \left(\sum_k P(z_n = k|d, \theta_d) P(w_n|\beta_k) \right) \right\} \\ = & \prod_{d=1} \pi_d \left\{ \prod_{n=1}^{N_d} \left(\sum_k \theta_{dk}\beta_{kw_n} \right) \right\} \end{aligned}$$

π_d *is usually considered as uniform*, i.e., 1/M

Re-arrange the Likelihood Function

- Group the same word from different positions together

$$\max logL = \sum_{dw} c(w,d) log \sum_{z} \theta_{dz} \beta_{zw}$$

$$s.t. \sum_{z} \theta_{dz} = 1 \; and \; \sum_{w} \beta_{zw} = 1$$

Optimization: EM Algorithm

- Repeat until converge
  - E-step: for each word in each document, calculate is conditional probability belonging to each topic

$$p(z|w,d) \propto p(w|z,d)p(z|d) = \beta_{zw}\theta_{dz}\ (i.e., p(z|w,d) = \frac{\beta_{zw}\theta_{dz}}{\sum_{z'}\beta_{z'w}\theta_{dz'}})$$

  - M-step: given the conditional distribution, find the parameters that can maximize the expected likelihood

$$\beta_{zw} \propto \sum_d p(z|w,d)c(w,d)\ (i.e., \beta_{zw} = \frac{\sum_d p(z|w,d)c(w,d)}{\sum_{w',d} p\left(z|w',d\right)c(w',d)})$$

$$\theta_{dz} \propto \sum_w p(z|w,d)c(w,d)\ \ (i.e., \theta_{dz} = \frac{\sum_w p(z|w,d)c(w,d)}{N_d})$$

Text Data: Topic Models

- Text Data and Topic Models
- Revisit of Multinomial Mixture Model
- Probabilistic Latent Semantic Analysis (pLSA)
- Latent Dirichlet Allocation (LDA)
- Summary

Limitations of pLSA

- Not a proper generative model
  - $\boldsymbol{\theta}_d$ is treated as a parameter
  - Cannot model new documents

- Solution:
  - Make it a proper generative model by adding priors to θ and β

Review of Conjugate Prior

- Model:
  - $p(x|\theta)$
- Prior:
  - $p(\theta|\alpha)$
- Posterior:
  - $p(\theta|x,\alpha) \propto p(\theta,x|\alpha) = p(x|\theta)p(\theta|\alpha)$

- Conjugate prior:
  - If $p(\theta|\alpha)$ and $p(\theta|x,\alpha)$ belong to the same distribution family (with different parameters)

Dirichlet-Multinomial Conjugate

- Dirichlet distribution: $\boldsymbol{\theta} \sim Dirichlet(\boldsymbol{\alpha})$
  - $i.e., p(\boldsymbol{\theta}|\boldsymbol{\alpha}) = \frac{\Gamma(\sum_k \alpha_k)}{\prod_k \Gamma(\alpha_k)} \prod_k \theta_k^{\alpha_k - 1}$, where $\alpha_k > 0$
    - $E(\theta_k) = \frac{\alpha_k}{\sum_k' \alpha_{k'}}$
    - $Var(\theta_k) = \frac{\alpha_k(\alpha_0 - \alpha_k)}{\alpha_0^2(\alpha_0 + 1)}$, where $\alpha_0 = \sum_k \alpha_k$

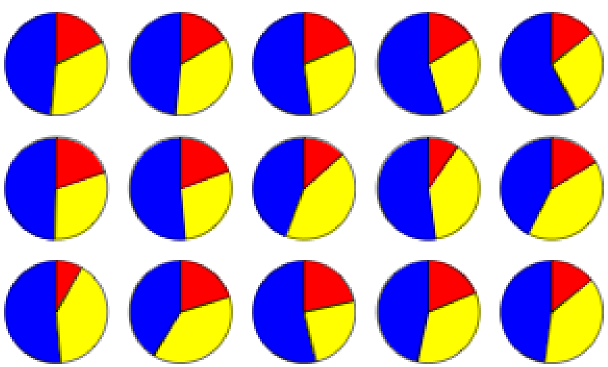

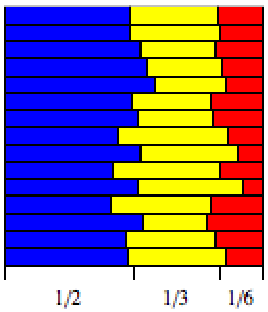

$Example: \boldsymbol{\theta} \sim Dirichlet(\boldsymbol{\alpha}), where\ \boldsymbol{\alpha}/\alpha_0 = (\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$

More Examples

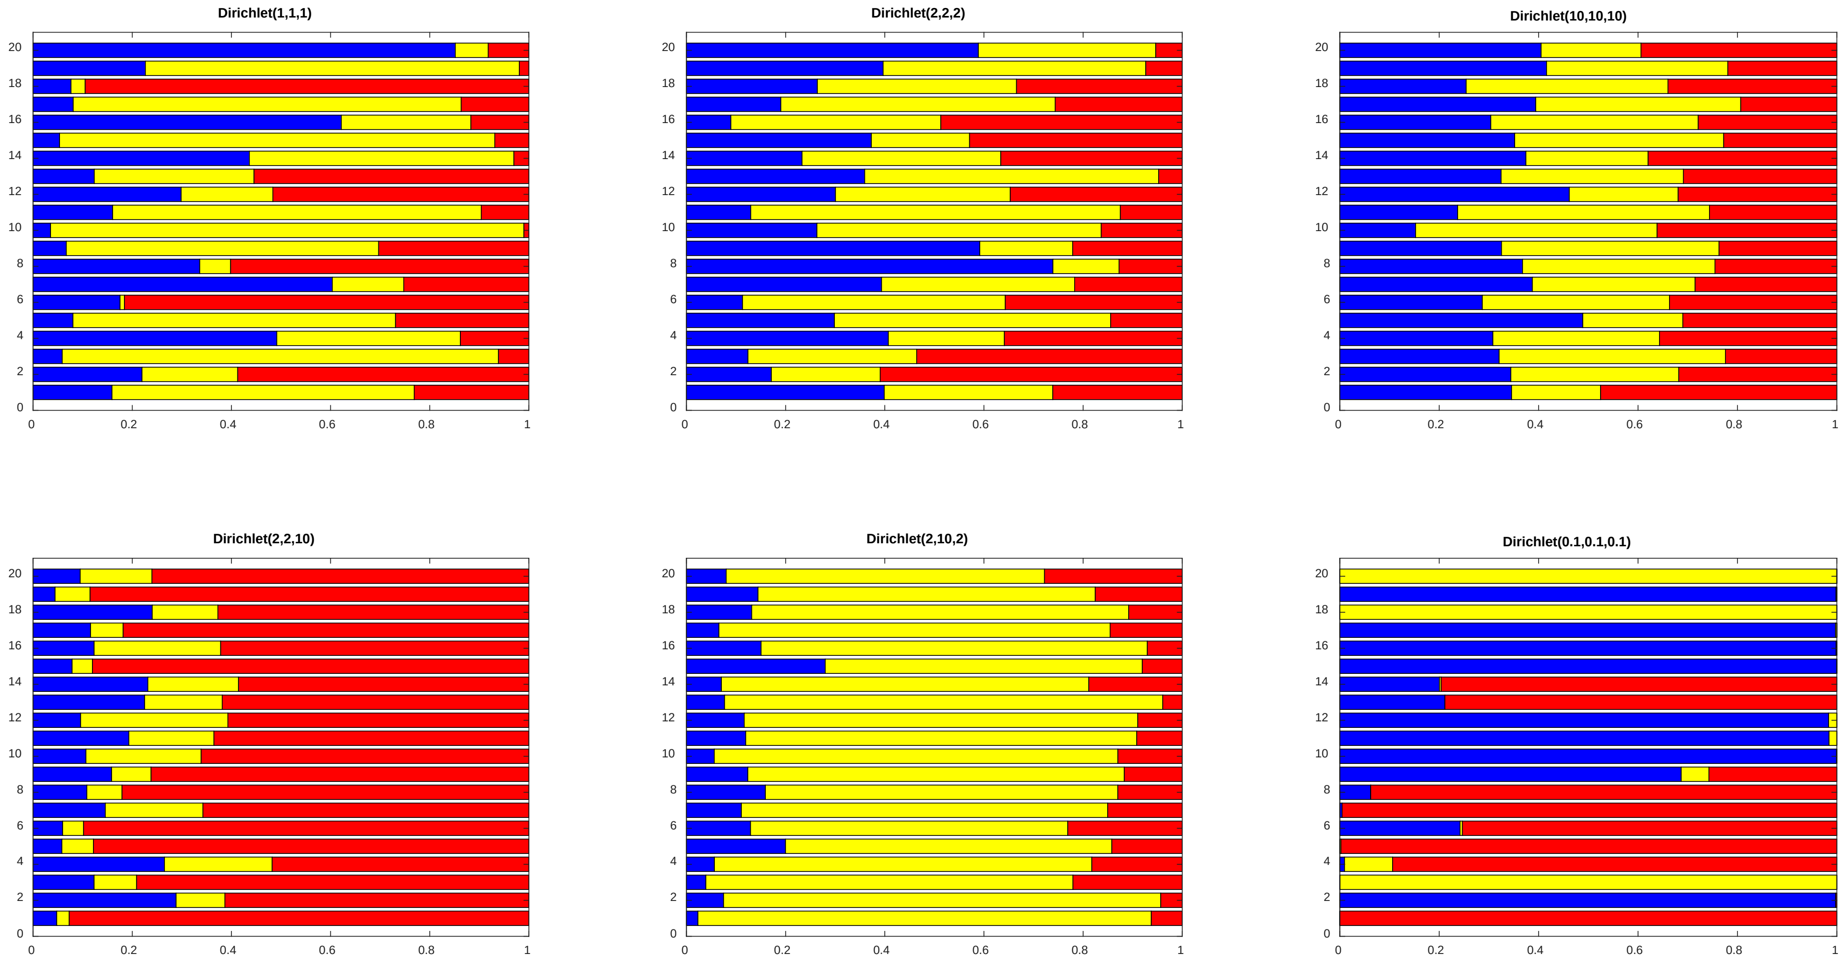

Simplex View

- $\alpha = (2,3,4)$

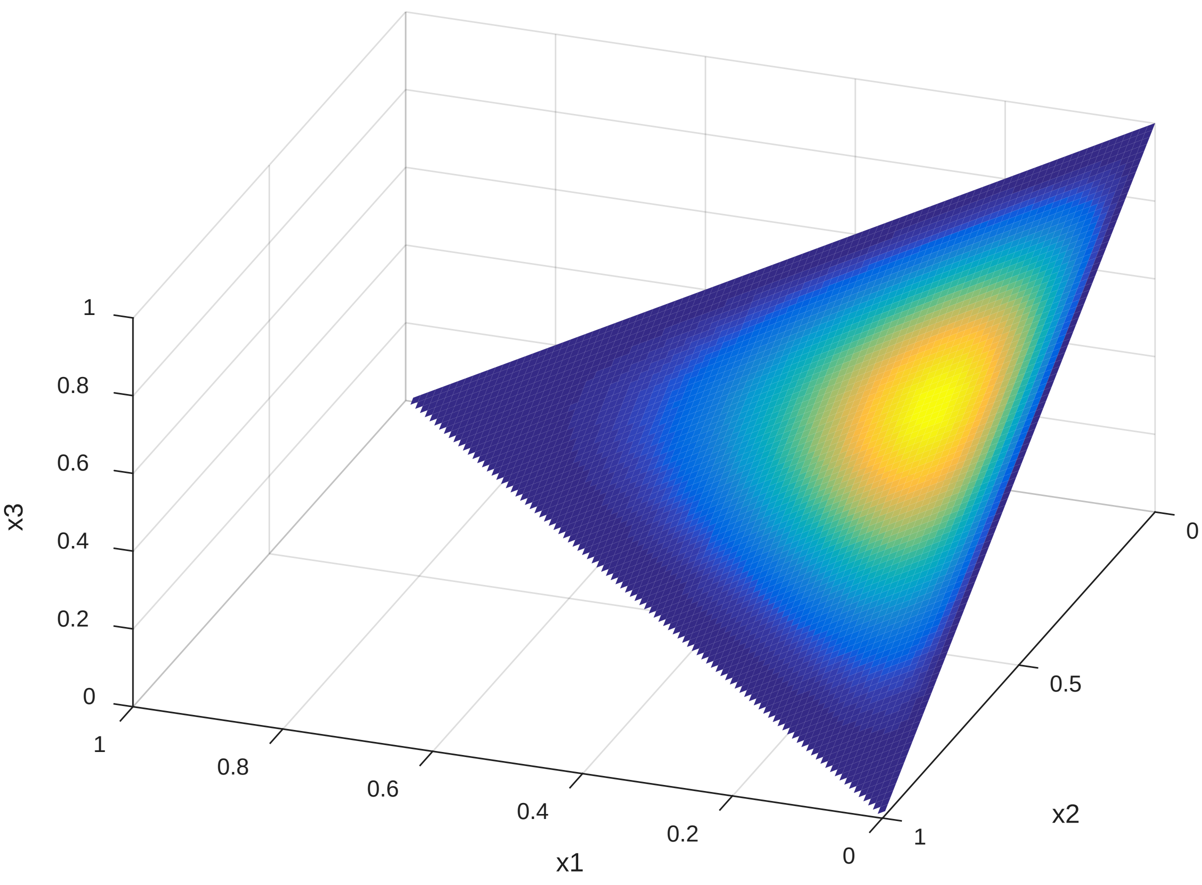

More Examples in the Simplex View

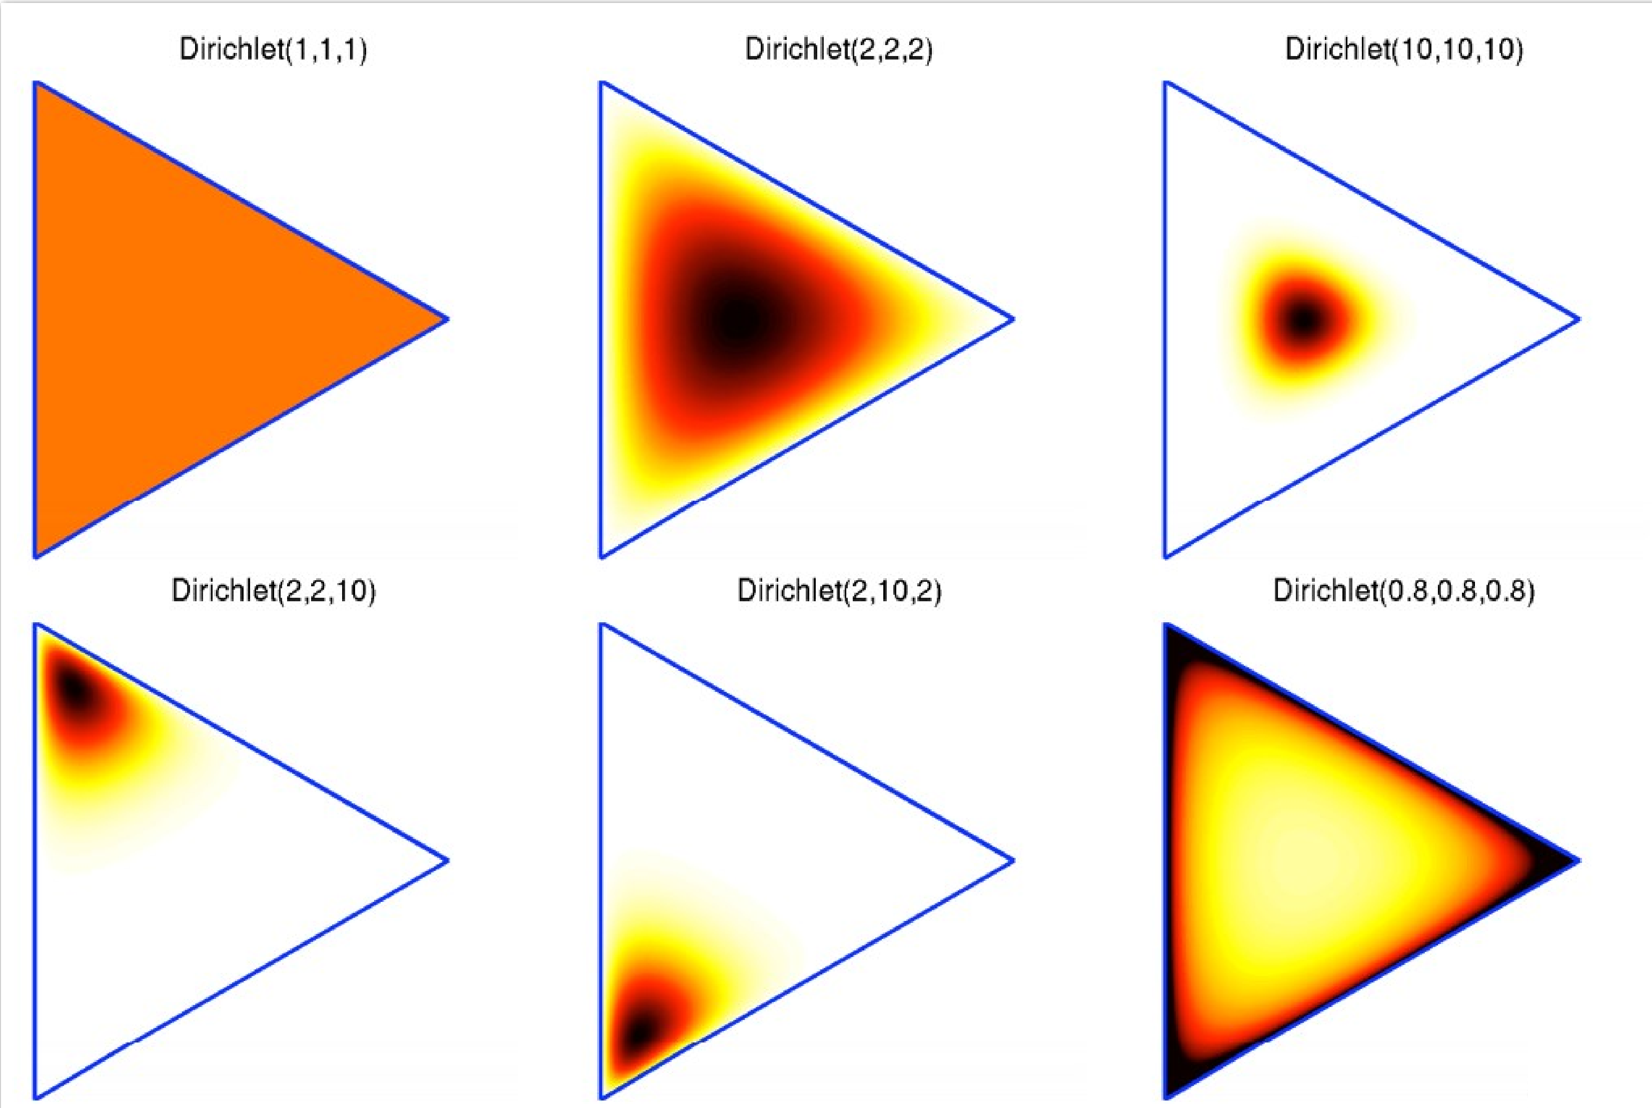

Dirichlet-Multinomial Conjugate

- Posterior
  - $p(\boldsymbol{\theta}|\boldsymbol{\alpha}, z_1, z_2, \dots z_n) \propto p(z_1, \dots, z_n|\boldsymbol{\theta})p(\boldsymbol{\theta}|\boldsymbol{\alpha})$

$$\propto \prod_k \theta_k^{\alpha_k + c_k - 1}$$

where c_k is the number of z that takes value of k

The Graphical Model of LDA

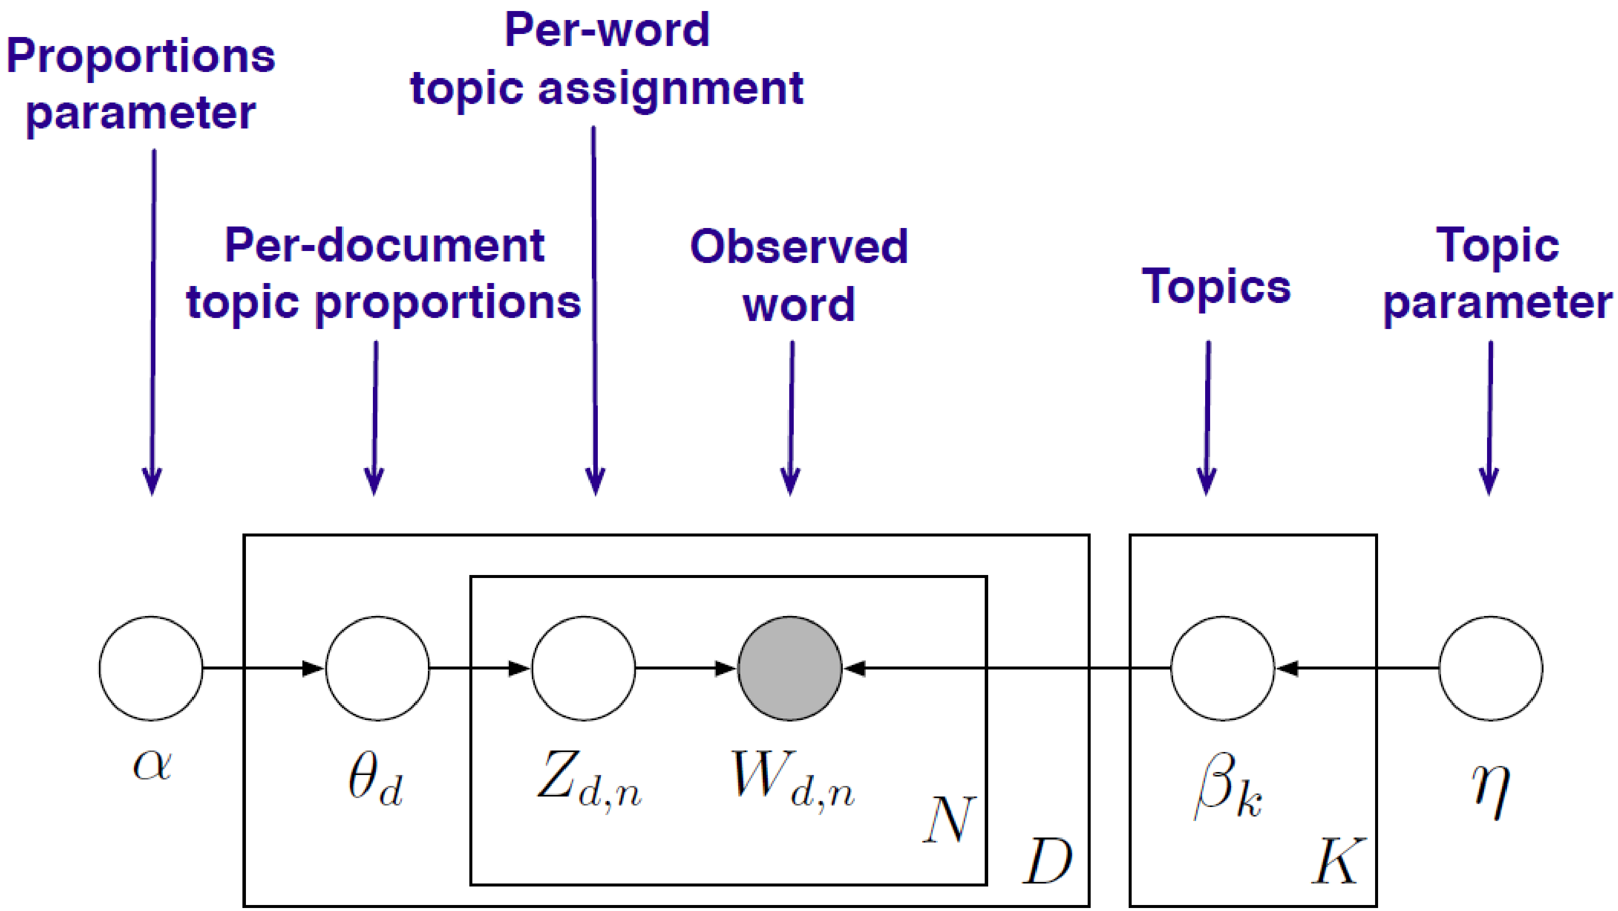

$$\boldsymbol{\theta_d \sim Dirichlet(\alpha)}$$
$$\boldsymbol{\beta_k \sim Dirichlet(\eta)}$$

Posterior Inference for LDA

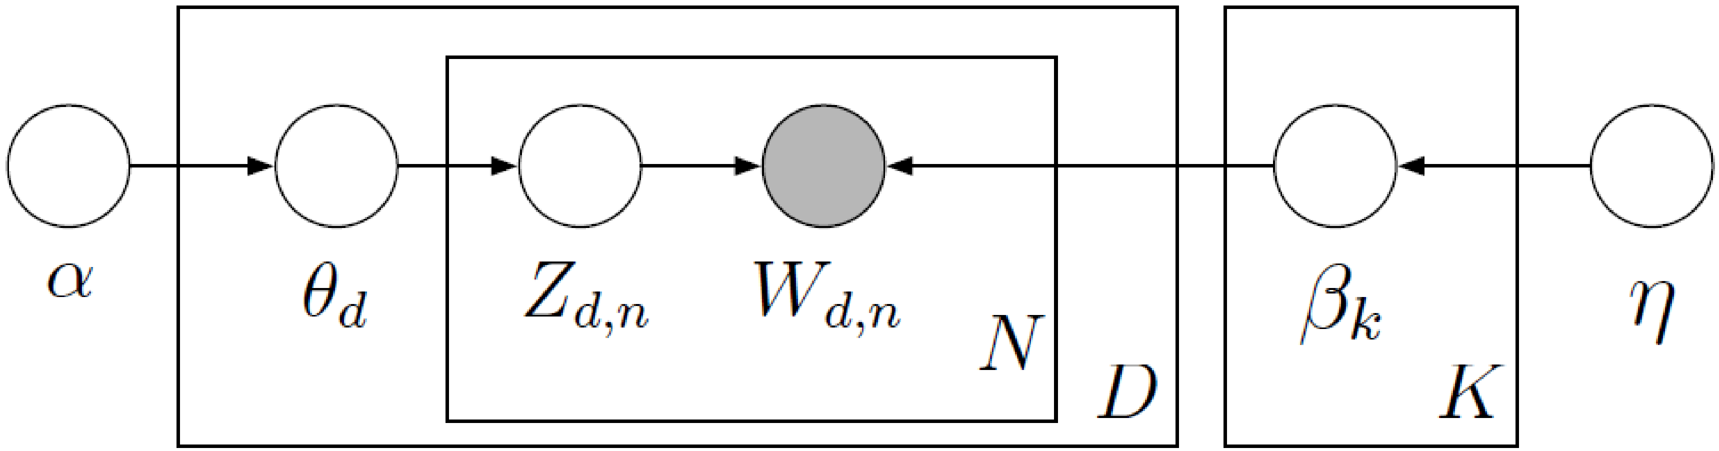

- Joint distribution of latent variables and documents is

$$\prod_{i=1}^{K} p(\beta_i \mid \eta) \prod_{d=1}^{D} p(\theta_d \mid \alpha) \left(\prod_{n=1}^{N} p(z_{d,n} \mid \theta_d) p(w_{d,n} \mid \beta_{1:K}, z_{d,n}) \right)$$

Posterior Inference

- Posterior of the latent variables

$$p(\beta, \theta, \mathbf{z} | \mathbf{w})$$

$$= \frac{p(\beta, \theta, \mathbf{z}, \mathbf{w})}{\boxed{\int_\beta \int_\theta \sum_{\mathbf{z}} p(\beta, \theta, \mathbf{z}, \mathbf{w})}}$$

$marginal\ p(\boldsymbol{w})\ is\ intractable!$

- Solutions
  - Gibbs sampling
  - Variational inference

References for LDA

- Blei et al., Latent Dirichlet Allocation, JLMR 3 (2003)
- Griffiths et al., Finding Scientific Topics, PNAS (2004)

Text Data: Topic Models

- Text Data and Topic Models
- Revisit of Multinomial Mixture Model
- Probabilistic Latent Semantic Analysis (pLSA)
- Latent Dirichlet Allocation (LDA)
- Summary

Summary

- Basic Concepts
  - Word/term, document, corpus, topic
  - How to represent a document
- Mixture of unigrams
- pLSA
  - Generative model
  - Likelihood function
  - EM algorithm
- LDA
  - Dirichlet-multinomial conjugate
  - Posterior inference