## Machine Learning



Reading: Tutorial on Topic Model @ ACL12

## We are inundated with data ...



- Humans cannot afford to deal with (e.g., search, browse, or measure similarity) a huge number of text and media documents
- We need computers to help out ...


## A task:



- Say, we want to have a mapping ..., so that

- Compare similarity
- Classify contents
- Cluster/group/categorize docs
- Distill semantics and perspectives



## Representation:

- Data: Bag of Words Representation

As for the Arabian and Palestinean voices that are against the current negotiations and the so-called peace process, they are not against peace per se, but rather for their well-founded predictions that Israel would NOT give an inch of the West bank (and most probably the same for Golan Heights) back to the Arabs. An 18 months of "negotiations" in Madrid, and Washington proved these predictions. Now many will jump on me saying why are you blaming israelis for no-result negotiations. I would say why would the Arabs stall the negotiations, what do they have to loose?


- Each document is a vector in the word space
- Ignore the order of words in a document. Only count matters!
- A high-dimensional and sparse representation

Not efficient text processing tasks, e.g., search, document
classification, or similarity measure
Not effective for browsing


## Subspace analysis



- Clustering: $(0,1)$ matrix
- LSI/NMF: "arbitrary" matrices
- Topic Models: stochastic matrix
- Sparse coding: "arbitrary" sparse matrices


## An example:



## Principal Component Analysis

- The new variables/dimensions
- Are linear combinations of the original ones
- Are uncorrelated with one another - Orthogonal in original dimension space
- Capture as much of the original variance in the data as possible
- Are called Principal Components
- Orthogonal directions of greatest variance in data
- Projections along PC1 discriminate the data most along any one axis

- First principal component is the direction of greatest variability (covariance) in the data
- Second is the next orthogonal (uncorrelated) direction of greatest variability
- So first remove all the variability along the first component, and then find the next direction of greatest variability
- And so on ...


## Computing the Components

- Projection of vector $\mathbf{x}$ onto an axis (dimension) $\mathbf{u}$ is $\mathbf{u}^{\top} \mathbf{x}$
- Direction of greatest variability is that in which the average square of the projection is greatest:

| Maximize | $\mathbf{u}^{\top} \mathbf{X} \mathbf{X}^{\top} \mathbf{u}$ |
| :--- | :--- |
| s.t | $\mathbf{u}^{\top} \mathbf{u}=1$ |

Construct Langrangian $\mathbf{u}^{\top} \mathbf{X} \mathbf{X}^{\top} \mathbf{u}-\lambda \mathbf{u}^{\top} \mathbf{u}$
Vector of partial derivatives set to zero

$$
\mathbf{x} \mathbf{x}^{\top} \mathbf{u}-\lambda \mathbf{u}=\left(\mathbf{x} \mathbf{x}^{\top}-\lambda \mathbf{I}\right) \mathbf{u}=0
$$

As $\mathbf{u} \neq \mathbf{0}$ then $\mathbf{u}$ must be an eigenvector of $\mathbf{X X}^{\top}$ with eigenvalue $\lambda$

- $\lambda$ is the principal eigenvalue of the correlation matrix $\mathbf{C =} \mathbf{X X}^{\top}$
- The eigenvalue denotes the amount of variability captured along that dimension


## Computing the Components

- Similarly for the next axis, etc.
- So, the new axes are the eigenvectors of the matrix of correlations of the original variables, which captures the similarities of the original variables based on how data samples project to them

- Geometrically: centering followed by rotation

[^0]
## Eigenvalues \& Eigenvectors

- For symmetric matrices, eigenvectors for distinct eigenvalues are orthogonal

$$
S v_{\{1,2\}}=\lambda_{\{1,2\}} v_{\{1,2\}} \text {, and } \lambda_{1} \neq \lambda_{2} \Rightarrow v_{1} \bullet v_{2}=0
$$

- All eigenvalues of a real symmetric matrix are real.

$$
\text { if }|S-\lambda I|=0 \text { and } S=S^{\mathrm{T}} \Rightarrow \lambda \in \mathfrak{R}
$$

- All eigenvalues of a positive semidefinite matrix are nonnegative

$$
\forall w \in \mathfrak{R}^{n}, w^{T} S w \geq 0 \text {, then if } S v=\lambda \mathrm{v} \Rightarrow \lambda \geq 0
$$

## Eigen/diagonal Decomposition

- Let $\mathbf{S} \in \mathbb{R}^{m \times m}$ be a square matrix with $\boldsymbol{m}$ linearly independent eigenvectors (a "non-defective" matrix)
- Theorem: Exists an eigen decomposition

$$
\mathbf{S}=\mathbf{U} \sqrt{\mathbf{U}^{-1}} \text { diagonal }
$$

| Unique |
| :---: |
| for |
| distinc |
| t eigen- |
| values |

(cf. matrix diagonalization theorem)

- Columns of $\boldsymbol{U}$ are eigenvectors of $\boldsymbol{S}$
- Diagonal elements of $\boldsymbol{\Lambda}$ are eigenvalues of $\mathbf{S}$

$$
\boldsymbol{\Lambda}=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{m}\right), \quad \lambda_{i} \geq \lambda_{i+1}
$$

## PCs, Variance and Least-Squares

- The first PC retains the greatest amount of variation in the sample
- The $k^{\text {th }}$ PC retains the $k$ th greatest fraction of the variation in the sample
- The $\mathrm{k}^{\text {th }}$ largest eigenvalue of the correlation matrix C is the variance in the sample along the $k^{\text {th }}$ PC
- The least-squares view: PCs are a series of linear least squares fits to a sample, each orthogonal to all previous ones


## The Corpora Matrix

$\mathbf{X}=$|  | Doc 1 | Doc 2 | Doc 3 | $\cdots$ |
| :---: | :---: | :---: | :---: | :---: |
| Word 1 | 3 | 0 | 0 | $\ldots$ |
| Word 2 | 0 | 8 | 1 | $\ldots$ |
| Word 3 | 0 | 1 | 3 | $\ldots$ |
| Word 4 | 2 | 0 | 0 | $\ldots$ |
| Word 5 | 12 | 0 | 0 | $\cdots$ |
|  | 0 | 0 | 0 | $\cdots$ |

## Singular Value Decomposition

For an $m \times n$ matrix A of rank $r$ there exists a factorization (Singular Value Decomposition = SVD) as follows:


The columns of $U$ are orthogonal eigenvectors of $A A^{T}$.
The columns of $V$ are orthogonal eigenvectors of $A^{\top} A$.
Eigenvalues $\lambda_{1} \ldots \lambda_{r}$ of $A A^{T}$ are the eigenvalues of $A^{\top} A$.

$$
\begin{aligned}
& \sigma_{i}=\sqrt{\lambda_{i}} \\
& \Sigma=\operatorname{diag}\left(\sigma_{1} \ldots \sigma_{r}\right)<\text { singular values. }
\end{aligned}
$$

## SVD and PCA

- The first root is called the prinicipal eigenvalue which has an associated orthonormal ( $\mathbf{u}^{\top} \mathbf{u}=1$ ) eigenvector $\mathbf{u}$
- Subsequent roots are ordered such that $\lambda_{1}>\lambda_{2}>\ldots>\lambda_{M}$ with rank(D) non-zero values.
- Eigenvectors form an orthonormal basis i.e. $\mathbf{u}_{\mathrm{i}}^{\top} \mathbf{u}_{\mathrm{j}}=\delta_{\mathrm{ij}}$
- The eigenvalue decomposition of $\mathbf{X} \mathbf{X}^{\top}=\mathbf{U} \mathbf{\Sigma} \mathbf{U}^{\top}$
- where $\mathbf{U}=\left[\mathbf{u}_{1}, \mathbf{u}_{2}, \ldots, \mathbf{u}_{\mathrm{M}}\right]$ and $\boldsymbol{\Sigma}=\operatorname{diag}\left[\lambda_{1}, \lambda_{2}, \ldots, \lambda_{\mathrm{M}}\right]$
- Similarly the eigenvalue decomposition of $\mathbf{X}^{\top} \mathbf{X}=\mathbf{V} \boldsymbol{\Sigma} \mathbf{V}^{\top}$
- The SVD is closely related to the above $\mathbf{X}=\mathbf{U} \Sigma^{1 / 2} \mathbf{V}^{\top}$
- The left eigenvectors U, right eigenvectors V,
- singular values $=$ square root of eigenvalues.


## How Many PCs?

- For n original dimensions, sample covariance matrix is nxn, and has up to $n$ eigenvectors. So $n$ PCs.
- Where does dimensionality reduction come from?

Can ignore the components of lesser significance.


You do lose some information, but if the eigenvalues are small, you don't lose much

- $\quad \mathrm{n}$ dimensions in original data
- calculate n eigenvectors and eigenvalues
- choose only the first $p$ eigenvectors, based on their eigenvalues
- final data set has only p dimensions © Eric Xing @ CMU, 2006-2015

$K$ is the number of singular values used
© Eric Xing @ CMU, 2006-2015


## Summary: <br> Latent Semantic Indexing ${ }_{\text {(Deervester e tal, 1990) }}$




T
( $\mathrm{m} \times \mathrm{k}$ )

$\Lambda$
( $\mathrm{k} \times \mathrm{k}$ )

$\mathrm{D}^{\top}$
( $\mathrm{k} \times \mathrm{n}$ )

$$
\vec{W}=\sum_{k=1}^{K} d_{k} \lambda_{k} \vec{T}_{k}
$$

- LSI does not define a properly normalized probability distribution of observed and latent entities
- Does not support probabilistic reasoning under uncertainty and data fusion


## Connecting Probability Models to Data

(Generative Model)
P(Data | Parameters)


Probabilistic
Real World
Model
Data


P(Parameters | Data)
(Inference)

## Latent Semantic Structure in GM

Distribution over words

$$
P(\mathbf{w})=\sum_{\ell} P(\mathbf{w}, \ell)
$$

Inferring latent structure

$$
P(\ell \mid \mathrm{w})=\frac{P(\mathrm{w} \mid \ell) P(\ell)}{P(\mathrm{w})}
$$

## How to Model Semantics?

- Q: What is it about?
- A: Mainly MT, with syntax, some learning



## Topic Models

A Hierarchical Phrase-Based Model for Statistical Machine Translation

We present a statistical phrase-based Translation model that uses hierarchical phrases-phrases that contain sub-phrases. The model is formally a synchronous context-free grammar but is learned from a bitext without any syntactic information. Thus it can be seen as a shift to the formal machinery of syntax based translation systems without any linguistic commitment. In our experiments using BLEU as a metric, the hierarchical Phrase based model achieves a relative Improvement of $7.5 \%$ over Pharaoh, a state-of-the-art phrase-based system.

## Why this is Useful?

- Q: What is it about?
- A: Mainly MT, with syntax, some learning

- Q: give me similar document?
- Structured way of browsing the collection
- Other tasks
- Dimensionality reduction
- TF-IDF vs. topic mixing proportion
- Classification, clustering, and more ...

AdMixing Proportion

## A Hierarchical Phrase-Based Model for Statistical Machine Translation

We present a statistical phrase-based Translation model that uses hierarchical phrases-phrases that contain sub-phrases. The model is formally a synchronous context-free grammar but is learned from a bitext without any syntactic information. Thus it can be seen as a shift to the formal machinery of syntax based translation systems without any linguistic commitment. In our experiments using BLEU as a metric, the hierarchical Phrase based model achieves a relative Improvement of $7.5 \%$ over Pharaoh, a state-of-the-art phrase-based system.

## Words in Contexts

## . "It was a nice shot."



## Words in Contexts (con'd)

- the opposition Labor Party fared even worse, with a predicted 35 Seats, seven less than last election.



## A possible generative process of a document



```
DOCUMENT 1: money \({ }^{1}\) bank \(^{1}\) bank \({ }^{1}\) Ioan \({ }^{1}\) river \(^{2}\) stream \(^{2}\) bank \({ }^{1}\) money \({ }^{1}\) river \(^{2}\) bank \(^{1}\) money \({ }^{1}\) bank \({ }^{1}\) loan \({ }^{1}\) money \({ }^{1}\) stream \({ }^{2}\) bank \(^{1}\) money \({ }^{1}\) bank \(^{1}\) bank \(^{1}\) loan \(^{1}\) river \(^{2}\) stream \(^{2}\) bank \({ }^{1}\) money \({ }^{1}\) river \(^{2}\) bank \(^{1}\) money \({ }^{1}\) bank \({ }^{1}\) loan \({ }^{1}\) bank \({ }^{1}\) money \({ }^{1}\) stream \({ }^{2}\)
DOCUMENT 2: river \(^{2}\) stream \({ }^{2}\) bank \({ }^{2}\) stream \({ }^{2}\) bank \(^{2}\) money \({ }^{1}\) loan \({ }^{1}\) river \(^{2}\) stream \({ }^{2}\) loan \(^{1}\) bank \(^{2}\) river \(^{2}\) bank \(^{2}\) bank \({ }^{1}\) stream \({ }^{2}\) river \(^{2}\) loan \(^{1}\) bank \(^{2}\) stream \(^{2}\) bank \(^{2}\) money \({ }^{1}\) loan \({ }^{1}\) river \(^{2}\) stream \({ }^{2}\) bank \(^{2}\) stream \({ }^{2}\) bank \(^{2}\) money \({ }^{1}\) river \(^{2}\) stream \(^{2}\) loan \({ }^{1}\) bank \(^{2}\) river \(^{2}\) bank \(^{2}\) money \(^{1}\) bank \(^{1}\) stream \(^{2}\) river \(^{2}\) bank \(^{2}\) stream \({ }^{\mathbf{2}}\) bank \(^{\mathbf{2}}\) money \(^{1}\)
```

admixing weight vector $\theta$ (represents all (distributions over components' elements) contributions)

Bayesian approach: use priors Admixture weights $\sim \operatorname{Dirichlet}(\alpha)$ Mixture components $\sim \operatorname{Dirichlet}(\Gamma)$


## Probabilistic LSI

- A "generative" model
- Models each word in a document as a sample from a mixture model.
- Each word is generated from a single topic, different words in the document may be generated from different topics.
- A topic is characterized by a distribution over words.
- Each document is represented as a list of admixing proportions for the components (i.e. topic vector $\theta$ ).



## Latent Dirichlet Allocation

Blei, Ng and Jordan (2003)

Essentially a Bayesian pLSI:

$$
\begin{aligned}
\boldsymbol{\theta} & \sim \operatorname{Dir}(\boldsymbol{\alpha}) \\
\mathbf{z}_{n} & \sim \operatorname{Mult}(\boldsymbol{\theta}) \\
\mathbf{w}_{n} & \sim p\left(\mathbf{w}_{n} \mid \mathbf{z}_{n}, \boldsymbol{\beta}\right)
\end{aligned}
$$



## LDA

- Generative model
- Models each word in a document as a sample from a mixture model.
- Each word is generated from a single topic, different words in the document may be generated from different topics.
- A topic is characterized by a distribution over words.
- Each document is represented as a list of admixing proportions for the components (i.e. topic vector).
- The topic vectors and the word rates each follows a Dirichlet prior --- essentially a Bayesian pLSI



## Topic Models = Mixed Membership Models = Admixture

Generating a document

- Draw $\theta$ from the prior

For each word $n$

- Draw $z_{n}$ from multinomial $(\theta)$
- Draw $w_{n} \mid z_{n},\left\{\beta_{1: k}\right\}$ from multinomial $\left(\beta_{z_{n}}\right)$




## Choices of Priors

- Dirichlet (LDA) (Blei et al. 2003)
- Conjugate prior means efficient inference
- Can only capture variations in each topic's intensity independently
- Logistic Normal (CTM=LoNTAM)
 (Blei \& Lafferty 2005, Ahmed \& Xing 2006)
- Capture the intuition that some topics are highly correlated and can rise up in intensity together
- Not a conjugate prior implies hard inference

- Nested CRP (Blei et al 2005)
- Defines hierarchy on topics



## Generative Semantic of LoNTAM

Generating a document

## Draw $\theta$ from the prior

For each word $n$

- Draw $Z_{n}$ from multinomia $l(\theta)$
- Draw $w_{n} \mid z_{n},\left\{\beta_{1: k}\right\}$ from multinomia $l\left(\beta_{z_{n}}\right)$




## Outcomes from a topic model

- The "topics" $\beta$ in a corpus:

| comp.graphics | T 59 | T 104 | T 31 |
| :---: | :---: | :---: | :---: |
|  | image <br> jpeg <br> color <br> file <br> gif <br> images <br> format <br> bit <br> files <br> display | ftp pub graphics mail version tar file information send server | card monitor dos video apple windows drivers vga cards graphics |
| sci.electronics | T 30 | T 84 | T 44 |
|  | $\begin{gathered} \text { power } \\ \text { ground } \\ \text { wire } \\ \text { circuit } \\ \text { supply } \\ \text { voltage } \\ \text { current } \\ \text { wiring } \\ \text { signal } \\ \text { cable } \end{gathered}$ | water energy air nuclear loop hot cold cooling heat temperature | sale price offer shipping sell interested mail condition email cd |


|  |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | T 42 | T 78 | T 47

- There is no name for each "topic", you need to name it!
- There is no objective measure of good/bad
- The shown topics are the "good" ones, there are many many trivial ones, meaningless ones, redundant ones, ... you need to manually prune the results
- How many topics? ...


## Outcomes from a topic model

- The "topic vector" $\theta$ of each doc

- Create an embedding of docs in a "topic space"
- Their no ground truth of $\theta$ to measure quality of inference
- But on $\theta$ it is possible to define an "objective" measure of goodness, such as classification error, retrieval of similar docs, clustering, etc., of documents
- But there is no consensus on whether these tasks bear the true value of topic models ...


## Outcomes from a topic model

- The per-word topic indicator $z$ :


#### Abstract

The William Randolph Hearst Fowdation will give $\$ 1.25$ milliom to Lincoln Center, Metropolitan Opera Co., New York Philharmonic and Juilliard School. "Our board felt that we had a real opportmity 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 Fomdation President Randolph A. Hearst said Monday in amomeing 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 tanght, will get $\$ 250,000$. The Hearst Foundation, a leading supporter of the Lincoln Center Consolidated Corporate Fimd, will make its usual ammal $\$ 100,000$ domation, too.


- Not very useful under the bag of word representation, because of loss of ordering
- But it is possible to define simple probabilistic linguistic constraints (e.g, bi-grams) over z and get potentially interesting results [Griffiths, Steyvers, Blei, \& Tenenbaum, 2004]



## Outcomes from a topic model

- The topic graph $S$ (when using CTM):

- Kind of interesting for understanding/visualizing large corpora


## Outcomes from a topic model

- Topic change trends
"Theoretical Physics"

"Neuroscience"



## The Big Picture

Unstructured Collection


Word Simplex

Structured Topic Network


Topic Space
(e.g, a Simplex)

## Computation on LDA

- Inference
- Given a Document D
- Posterior: $P(\Theta \mid \mu, \Sigma, \beta, D)$
- Evaluation: $P(D \mid \mu, \Sigma, \beta)$
- Learning

- Given a collection of documents $\left\{D_{i}\right\}$
- Parameter estimation

$$
\underset{(\mu, \Sigma, \beta)}{\arg \max } \sum \log \left(P\left(D_{i} \mid \mu, \Sigma, \beta\right)\right)
$$

## Exact Bayesian inference on LDA is intractable

- A possible query:

$$
\begin{aligned}
& p\left(\theta_{n} \mid D\right)=? \\
& p\left(z_{n, m} \mid D\right)=?
\end{aligned}
$$

- Close form solution? $\quad p\left(\boldsymbol{\theta}_{\boldsymbol{n}} \mid D\right)=\frac{p\left(\theta_{n}, D\right)}{p(D)}$

$$
=\frac{\sum_{\left\{z_{n, m}\right\}} \int\left(\prod_{n}\left(\prod_{m} p\left(x_{n, m} \mid \boldsymbol{\beta}_{z_{n}}\right) p\left(z_{n, m} \mid \boldsymbol{\theta}_{n}\right)\right) p\left(\boldsymbol{\theta}_{n} \mid \alpha\right)\right) p(\phi \mid G) d \boldsymbol{\theta}_{-n} d \boldsymbol{\beta}}{p(D)}
$$

$$
p(D)=\sum_{\left\{z_{n, m}\right\}} \int \cdots \int\left(\prod_{n}\left(\prod_{m} p\left(x_{n, m} \mid \beta_{z_{n}}\right) p\left(z_{n, m} \mid \theta_{n}\right)\right) p\left(\theta_{n} \mid \alpha\right)\right) p(\beta \mid G) d_{\theta_{1}} \cdots d_{\theta_{N}} d_{\beta}
$$

- Sum in the denominator over $T^{n}$ terms, and integrate over $n k$-dimensional topic vectors


## Approximate Inference

- Variational Inference
- Mean field approximation (Blei et al)
- Expectation propagation (Minka et al)
- Variational $2^{\text {nd }}$-order Taylor approximation (Ahmed and Xing)
- Markov Chain Monte Carlo
- Gibbs sampling (Griffiths et al)


## Collapsed Gibbs sampling

(Tom Griffiths \& Mark Steyvers)

- Collapsed Gibbs sampling
- Integrate out $\theta$

For variables $\mathbf{z}=z_{1}, z_{2}, \ldots, z_{n}$
Draw $z_{i}^{(t+1)}$ from $P\left(z_{i} \mid \mathbf{z}_{-i}, \mathbf{w}\right)$

$$
\begin{gathered}
z_{-i}=z_{1}^{(t+1)}, z_{2}^{(t+1)}, \ldots, z_{i-1}(t+1), z_{i+1}^{(t)}, \ldots, z_{n}^{(t)} \\
\left\{z^{(1)}, z^{(2)}, \ldots, z^{(T)}\right\} \\
\theta=\frac{1}{T} \sum_{t} z^{(t)}
\end{gathered}
$$



## Gibbs sampling

- Need full conditional distributions for variable
- Since we only sample $z$ we need

$$
P\left(z_{i}=j \mid \mathbf{z}_{-i}, \mathbf{w}\right) \propto P\left(w_{i} \mid z_{i}=j, \mathbf{z}_{-i}, \mathbf{w}_{-i}\right) P\left(z_{i}=j \mid \mathbf{z}_{-i}\right)
$$

 number of times word $\boldsymbol{w}$ assigned to topic $\boldsymbol{j}$
$n_{j}^{(d)} \quad$ number of times topic $\boldsymbol{j}$ used in document $\boldsymbol{d}$

## Gibbs sampling

|  |  |  | ite |
| :---: | :---: | :---: | :---: |
|  |  |  | $\mathbf{1}$ |
| $i$ | $W_{i}$ | $d_{i}$ | $Z_{i}$ |
| 1 | MATHEMATICS | 1 | 2 |
| 2 | KNOWLEDGE | 1 | 2 |
| 3 | RESEARCH | 1 | 1 |
| 4 | WORK | 1 | 2 |
| 5 | MATHEMATICS | 1 | 1 |
| 6 | RESEARCH | 1 | 2 |
| 7 | WORK | 1 | 2 |
| 8 | SCIENTIFIC | 1 | 1 |
| 9 | MATHEMATICS | 1 | 2 |
| 10 | WORK | 1 | 1 |
| 11 | SCIENTIFIC | 2 | 1 |
| 12 | KNOWLEDGE | 2 | 1 |
| $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ |
| 50 | JOY | 5 | 2 |

## Gibbs sampling

|  |  | iteration |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  |  | $\mathbf{1}$ | $\mathbf{2}$ |
| $i$ | $W_{i}$ | $d_{i}$ | $Z_{i}$ | $Z_{i}$ |
| 1 | MATHEMATICS | 1 | 2 | $?$ |
| 2 | KNOWLEDGE | 1 | 2 |  |
| 3 | RESEARCH | 1 | 1 |  |
| 4 | WORK | 1 | 2 |  |
| 5 | MATHEMATICS | 1 | 1 |  |
| 6 | RESEARCH | 1 | 2 |  |
| 7 | WORK | 1 | 2 |  |
| 8 | SCIENTIFIC | 1 | 1 |  |
| 9 | MATHEMATICS | 1 | 2 |  |
| 10 | WORK | 1 | 1 |  |
| 11 | SCIENTIFIC | 2 | 1 |  |
| 12 | KNOWLEDGE | 2 | 1 |  |
| $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ |  |
| $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ |  |
| $\cdot$ | $\cdot$ | $\cdot$ | - |  |
| 50 | JOY | 5 | 2 |  |

## Gibbs sampling

|  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | 1 | 2 |  |
| $i$ | $W_{i}$ | $d_{i}$ | $Z_{i}$ | $\mathrm{Z}_{i}$ |  |
| 1 | MATHEMATICS | 1 | 2 | . |  |
| 2 | KNOWLEDGE | 1 | 2 |  |  |
| 3 | RESEARCH | 1 | 1 |  |  |
| 4 | WORK | 1 | 2 |  |  |
| 5 | MATHEMATICS | 1 | 1 |  |  |
| 6 | RESEARCH | 1 | 2 |  |  |
|  | WORK | 1 | 2 |  |  |
| 8 | SCIENTIFIC | 1 | 1 |  |  |
| 9 | MATHEMATICS | 1 | 2 |  |  |
| 10 | WORK | 1 | 1 |  |  |
| 11 | SCIENTIFIC | 2 | 1 |  |  |
| 12 | KNOWLEDGE | 2 | 1 |  |  |
| . | . | . | . |  |  |
| . | . | . |  |  |  |
| 50 | JOY | 5 | 2 |  |  |
|  |  |  |  |  | $\frac{n_{-i, j}^{\left(w_{i}\right)}+\boldsymbol{G}}{n_{-i, j}^{(\cdot)}+W \boldsymbol{G}} \frac{n_{-i, j}^{\left(d_{i}\right)}+\alpha}{n_{-i, .}^{\left(d_{i}\right)}+T \alpha}$ |

## Gibbs sampling

|  |  |
| :---: | :---: |
| $i$ | $w_{i}$ |
| 1 | MATHEMATICS |
| 2 | KNOWLEDGE |
| 3 | RESEARCH |
| 4 | WORK |
| 5 | MATHEMATICS |
| 6 | RESEARCH |
| 7 | WORK |
| 8 | SCIENTIFIC |
| 9 | MATHEMATICS |
| 10 | WORK |
| 11 | SCIENTIFIC |
| 12 | KNOWLEDGE |
| $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ |
| . | JOY |


| iteration |  |
| :--- | ---: |
| $\mathbf{1}$ | $\mathbf{2}$ |
| $Z_{i}$ | $Z_{i}$ |
| 2 | $?$ |
| 2 |  |
| 1 |  |
| 2 |  |
| 1 |  |
| 2 |  |
| 2 |  |
| 1 |  |
| 2 |  |
| 1 |  |
| 1 |  |
| 1 |  |
| $\cdot$ |  |
| $\cdot$ |  |
|  |  |
| 2 |  |
|  |  |

$$
\begin{aligned}
& \qquad P\left(z_{i}=j \mid \mathbf{z}_{-i}, \mathbf{w}\right) \propto \frac{n_{-i, j}^{\left(w_{i}\right)}+\boldsymbol{G}}{n_{-i, j}^{(\cdot)}+W \boldsymbol{G}} \frac{n_{-i, j}^{\left(d_{i}\right)}+\alpha}{n_{-i, .}^{\left(d_{i}\right)}+T \alpha} \\
& \text { © Eric Xing @ Сми, 2000-2015 }
\end{aligned}
$$

## Gibbs sampling

|  |  |
| :---: | :---: |
| $i$ |  |
| 1 | $w_{i}$ |
| 2 | MATHEMATICS |
| 3 | KNOWLEDGE |
| 4 | RESEARCH |
| 5 | WORK |
| 6 | MATHEMATICS |
| 7 | WORK |
| 8 | SCIENTIFIC |
| 9 | MATHEMATICS |
| 10 | WORK |
| 11 | SCIENTIFIC |
| 12 | KNOWLEDGE |
| $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ |
| . | JOY |

## Gibbs sampling

|  |  |
| :---: | :---: |
| $i$ |  |
| $i$ | $W_{i}$ |
| 2 | MATHEMATICS |
| 3 | KNOWLEDGE |
| 4 | RESEARCH |
| 5 | WORK |
| 6 | MATHEMATICS |
| 7 | RESEARCH |
| 8 | WORK |
| 9 | MATENTIFIC |
| 10 | WORATICS |
| 11 | SCIENTIFIC |
| 12 | KNOWLEDGE |
| $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ |
| $\dot{5}$ | . |
| 50 | JOY |

## Gibbs sampling

|  |  |
| :---: | :---: |
| $i$ | $w_{i}$ |
| 1 | MATHEMATICS |
| 2 | KNOWLEDGE |
| 3 | RESEARCH |
| 4 | WORK |
| 5 | MATHEMATICS |
| 6 | RESEARCH |
| 7 | WORK |
| 8 | SCIENTIFIC |
| 9 | MATHEMATICS |
| 10 | WORK |
| 11 | SCIENTIFIC |
| 12 | KNOWLEDGE |
| $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ |
| . | $\cdot$ |
| 50 | JOY |


| iteration |  |
| :---: | ---: |
| $\mathbf{1}$ | $\mathbf{2}$ |
| $Z_{i}$ | $Z_{i}$ |
| 2 | 2 |
| 2 | 1 |
| 1 | 1 |
| 2 | $?$ |
| 1 |  |
| 2 |  |
| 2 |  |
| 1 |  |
| 2 |  |
| 1 |  |
| 1 |  |
| 1 |  |
| $\cdot$ |  |
| . |  |
| 2 |  |

$$
\begin{aligned}
& \qquad P\left(z_{i}=j \mid \mathbf{z}_{-i}, \mathbf{w}\right) \propto \frac{n_{-i, j}^{\left(w_{i}\right)}+\boldsymbol{G}}{n_{-i, j}^{(\cdot)}+W \boldsymbol{G}} \frac{n_{-i, j}^{\left(d_{i}\right)}+\alpha}{n_{-i, \cdot}^{\left(d_{i}\right)}+T \alpha} \\
& \qquad \text { Eric Xing @ CMU, 2006-2015 }
\end{aligned}
$$

## Gibbs sampling

| $i$ | $W_{i}$ | $d_{i}$ |
| :---: | :---: | :---: |
| 1 | MATHEMATICS | 1 |
| 2 | KNOWLEDGE | 1 |
| 3 | RESEARCH | 1 |
| 4 | WORK | 1 |
| 5 | MATHEMATICS | 1 |
| 6 | RESEARCH | 1 |
| 7 | WORK | 1 |
| 8 | SCIENTIFIC | 1 |
| 9 | MATHEMATICS | 1 |
| 10 | WORK | 1 |
| 11 | SCIENTIFIC | 2 |
| 12 | KNOWLEDGE | 2 |
| . | . | . |
| . |  | . |
| 50 | JOY | 5 |

$$
\begin{aligned}
& \text { iteration } \\
& 1 \\
& 2 \\
& P\left(z_{i}=j \mid \mathbf{z}_{-i}, \mathbf{w}\right) \propto \frac{n_{-i, j}^{\left(w_{i}\right)}+\boldsymbol{G}}{n_{-i, j}^{(\cdot)}+W \boldsymbol{G}} \frac{n_{-i, j}^{\left(d_{i}\right)}+\alpha}{n_{-i, .}^{\left(d_{i}\right)}+T \alpha}
\end{aligned}
$$

## Gibbs sampling



## Learning a TM

- Maximum likelihood estimation:

$$
\left\{\beta_{1}, \beta_{2}, \ldots, \beta_{K}\right\}, \alpha=\underset{(\alpha, \beta)}{\arg \max } \sum \log \left(P\left(D_{i} \mid \alpha, \beta\right)\right)
$$

- Need statistics on topic-specific word assignment (due to $z$ ), topic vector distribution (due to $\theta$ ), etc.
- E.g,, this is the formula for topic $k$ :

$$
\beta_{k}=\frac{1}{\sum_{d} N_{d}} \sum_{d=1}^{D} \sum_{d_{n}=1}^{N_{d}} \delta\left(z_{d, d_{n}}, k\right) w_{d, d_{n}}
$$

- These are hidden variables, therefore need an EM algorithm (also known as data augmentation, or DA, in Monte Carlo paradigm)
- This is a "reduce" step in parallel implementation


## Conclusion

- GM-based topic models are cool
- Flexible
- Modular
- Interactive
- There are many ways of implementing topic models
- unsupervised
- supervised
- Efficient Inference/learning algorithms
- GMF, with Laplace approx. for non-conjugate dist.
- MCMC
- Many applications
- Word-sense disambiguation
- Image understanding
- Network inference


[^0]:    - Linear transformation

