PETUUM

Carnegie Mellon University School of Computer Science

Probabilistic Graphical Models

Variational Inference 1

Eric Xing

Lecture 7, February 5, 2020

Reading: see class homepage

Inference Problems in Graphical Models

- E.g.: A general undirected graphical model (MRF):

$$p(x) = \frac{1}{Z} \prod_{C \in \mathcal{C}} \psi_C(x_C)$$

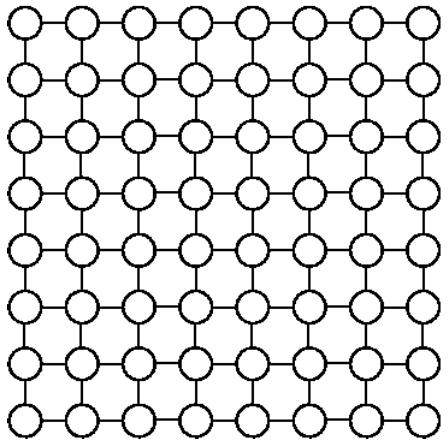

- The quantities of interest:
  - marginal distributions: $$p(x_i) = \sum_{x_j, j \neq i} p(x)$$
  - normalization constant (partition function): Z
- Exact inference: tree graph, discrete scope or known integral, …
- What if exact inference is expensive or even impossible? (when this can happen?)

Approximate Inference: The Big Picture

- Variational Inference
  - Mean-field (inner approximation)
  - Loopy Belief Propagation (outer approximation)
  - Kikuchi and variants (tighter outer approximation)
  - Expectation Propagation (reverse KL)
  - …
- Sampling
  - Monte Carlo
  - Sequential Monte Carlo (Particle Filters)
  - Markov Chain Monte Carlo
  - Hybrid Monte Carlo
  - …

Variational Methods

- "Variational": fancy name for optimization-based formulations
  - i.e., represent the quantity of interest as the solution to an optimization problem
  - *approximate* the desired solution by *relaxing/approximating* the *intractable* optimization problem

- Examples:
  - Courant-Fischer for eigenvalues: $\lambda_{\max}(A) = \max_{\|x\|_2=1} x^T A x$
  - Linear system of equations: $Ax = b, A \succ 0, x^* = A^{-1}b$
    - variational formulation:

$$x^* = \arg\min_x \left\{ \frac{1}{2} x^T A x - b^T x \right\}$$

    - for large system, apply conjugate gradient method

Variational Inference: High-level Idea

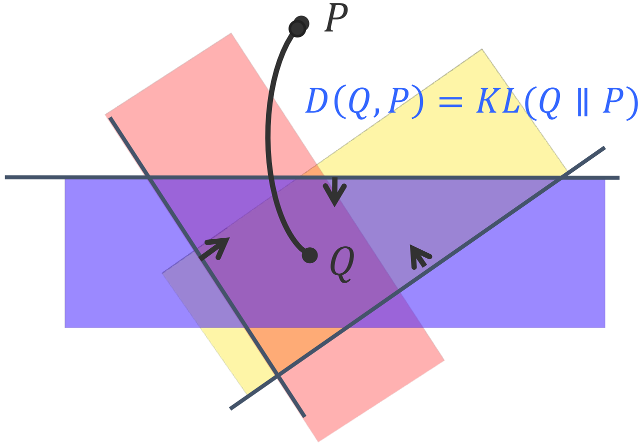

- Inference: answer queries of P
- Challenge: direct inference on P is often intractable
- Indirect approach:
  - Project P to a tractable family of distributions *Q*
  - Perform inference on the projected Q
- Projection requires a measure of distance
  - A convenient choice: KL(Q, P)
- Mean-field: Assume Q is fully factorized
  - The simplest possible family of distributions
- Example: Latent Dirichlet Allocation (LDA)

$$q^* = \arg\min_{q \in \mathcal{M}_F(G)} \langle E \rangle_q - H_q$$

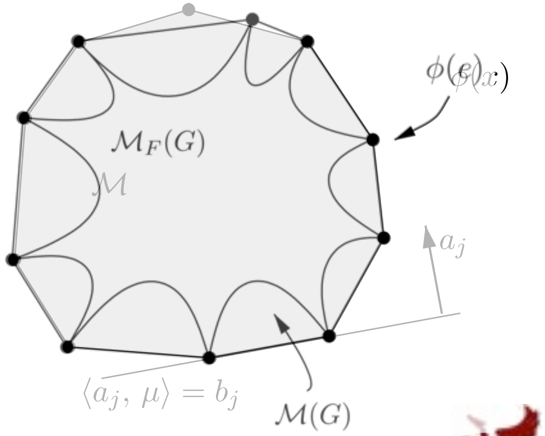

Probabilistic Topic Models

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

How to get started for a new modeling task?

Here are some important elements to consider before you start:

- Task:
  - Embedding? Classification? Clustering? Topic extraction? …
- Data representation:
  - Input and output (e.g., continuous, binary, counts, …)
- Model:
  - BN? MRF? Regression? SVM?
- Inference:
  - Exact inference? MCMC? Variational?
- Learning:
  - MLE? MCLE? Max margin?
- Evaluation:
  - Visualization? Human interpretability? Perperlexity? Predictive accuracy?

It is better to consider one element at a time!

Tasks: document embedding

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

⇒

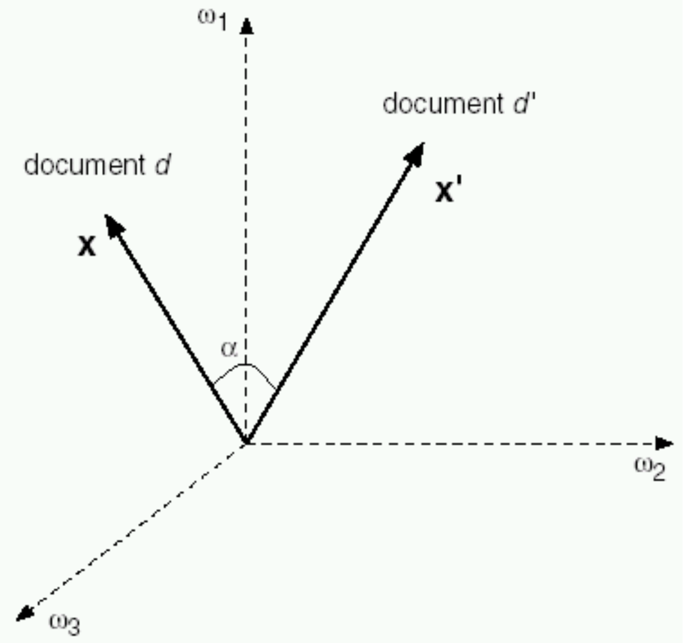

- Compare similarity
- Classify contents
- Cluster/group/categorizing
- Distill semantics and perspectives
- ..

Summarizing the data using topics

Bayesian modeling	Visual cortex	Education	Market
Bayesian	cortex	students	market
model	cortical	education	economic
inference	areas	learning	financial
models	visual	educational	economics
probability	area	teaching	markets
probabilistic	primary	school	returns
Markov	connections	student	price
prior	ventral	skills	stock
hidden	cerebral	teacher	value
approach	sensory	academic	investment

See how data changes over time

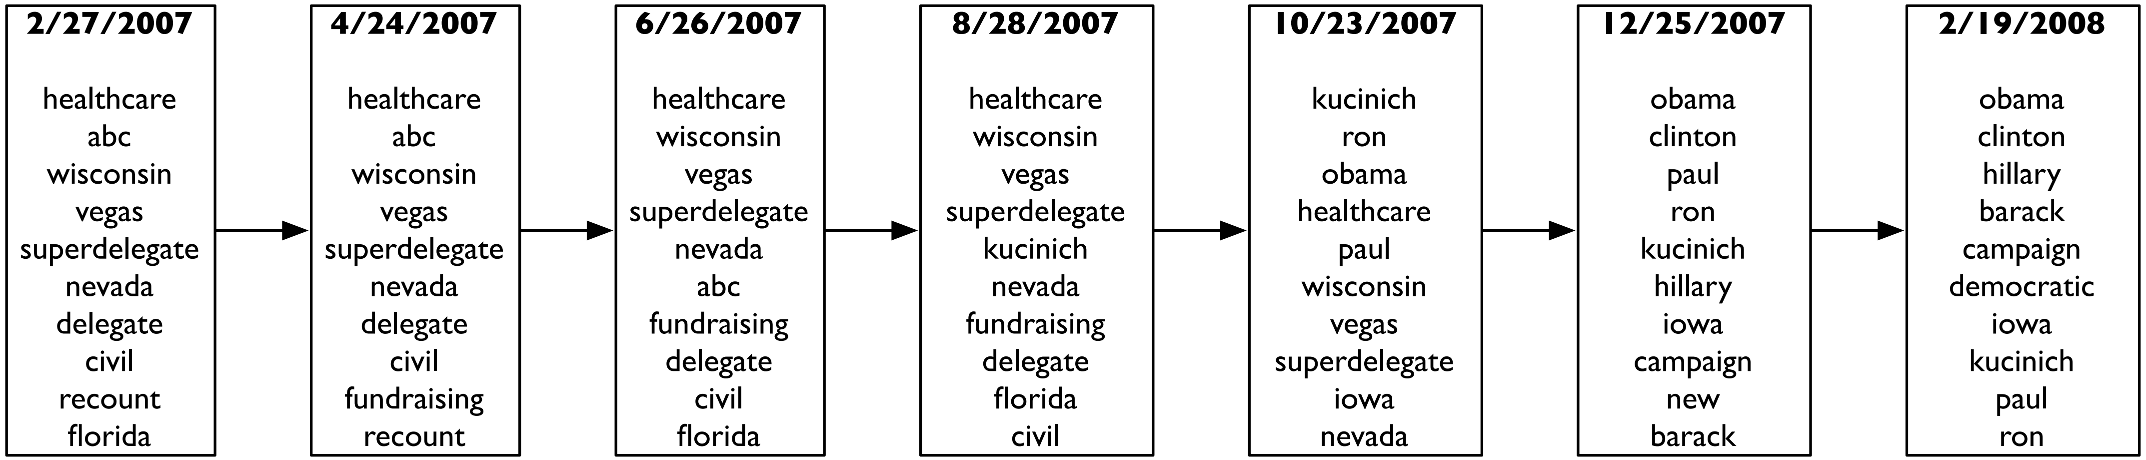

User interest modeling using topics

User interest profile (adjustable with sliders---Changing these changes recommendations.)

Weight | **User preferred topics**

1: learning machine training vector learn machines kernel learned classifiers classifier
2: online classification digital library libraries browsing classify classifying labels catalog
3: two differences active hypothesis arise difference evolved morphological modify morphology
4: experiments ability demonstrated produced contexts situations instances fail recognize string
5: features class classes subset java characteristic earlier represented defines separate
6: process making presents objective steps reports distinguish exploit maintaining select
7: algorithm signal input signals output exact performs music sound iterative
8: database databases contains version list comprehensive release stored update curated
9: applications application provide built numerous proven providing discusses tremendous presents
10: text literature discovery mining biomedical full extract discovering texts themes

http://cogito-demos.ml.cmu.edu/cgi-bin/recommendation.cgi

Representation: Bag of Words Representation

- Data:

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 $(|V| \gg D)$
  - Not efficient text processing tasks, e.g., search, document classification, or similarity measure
  - Not effective for browsing

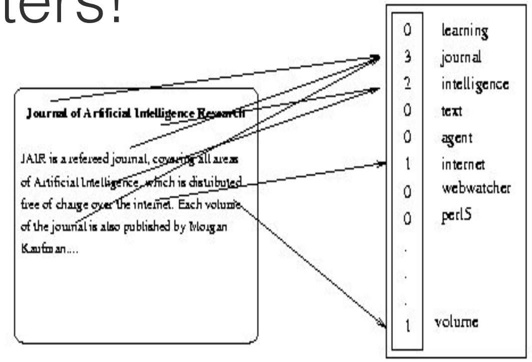

How to Model Semantic?

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

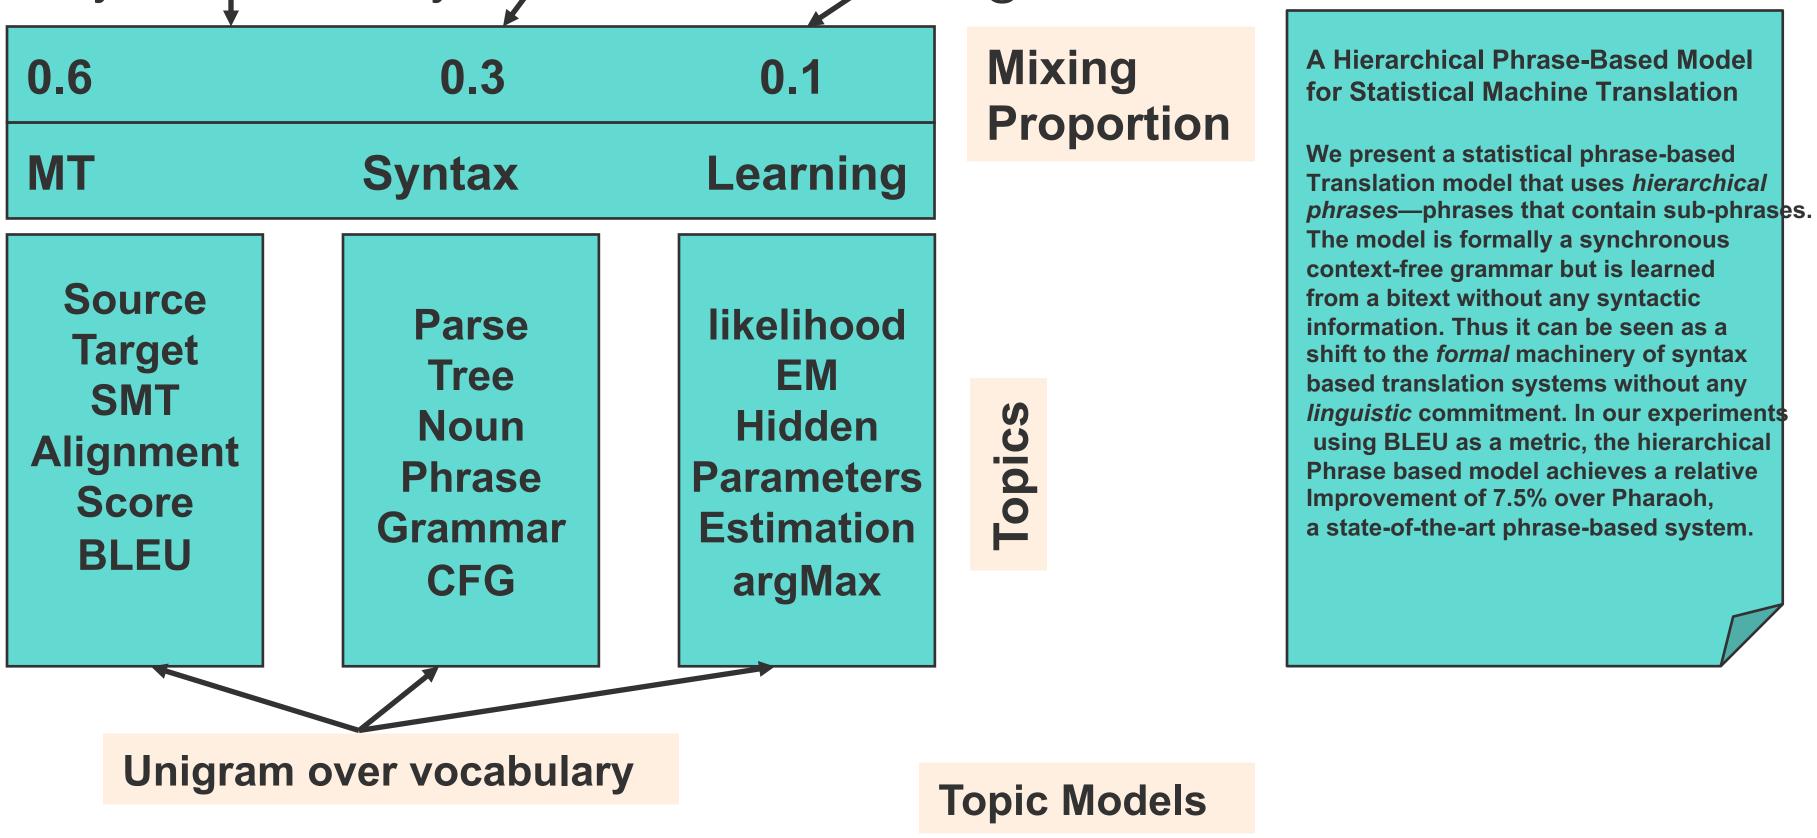

Why this is Useful?

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

0.6	0.3	0.1
MT	Syntax	Learning

Mixing 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.

- **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 …

Topic Models: The Big Picture

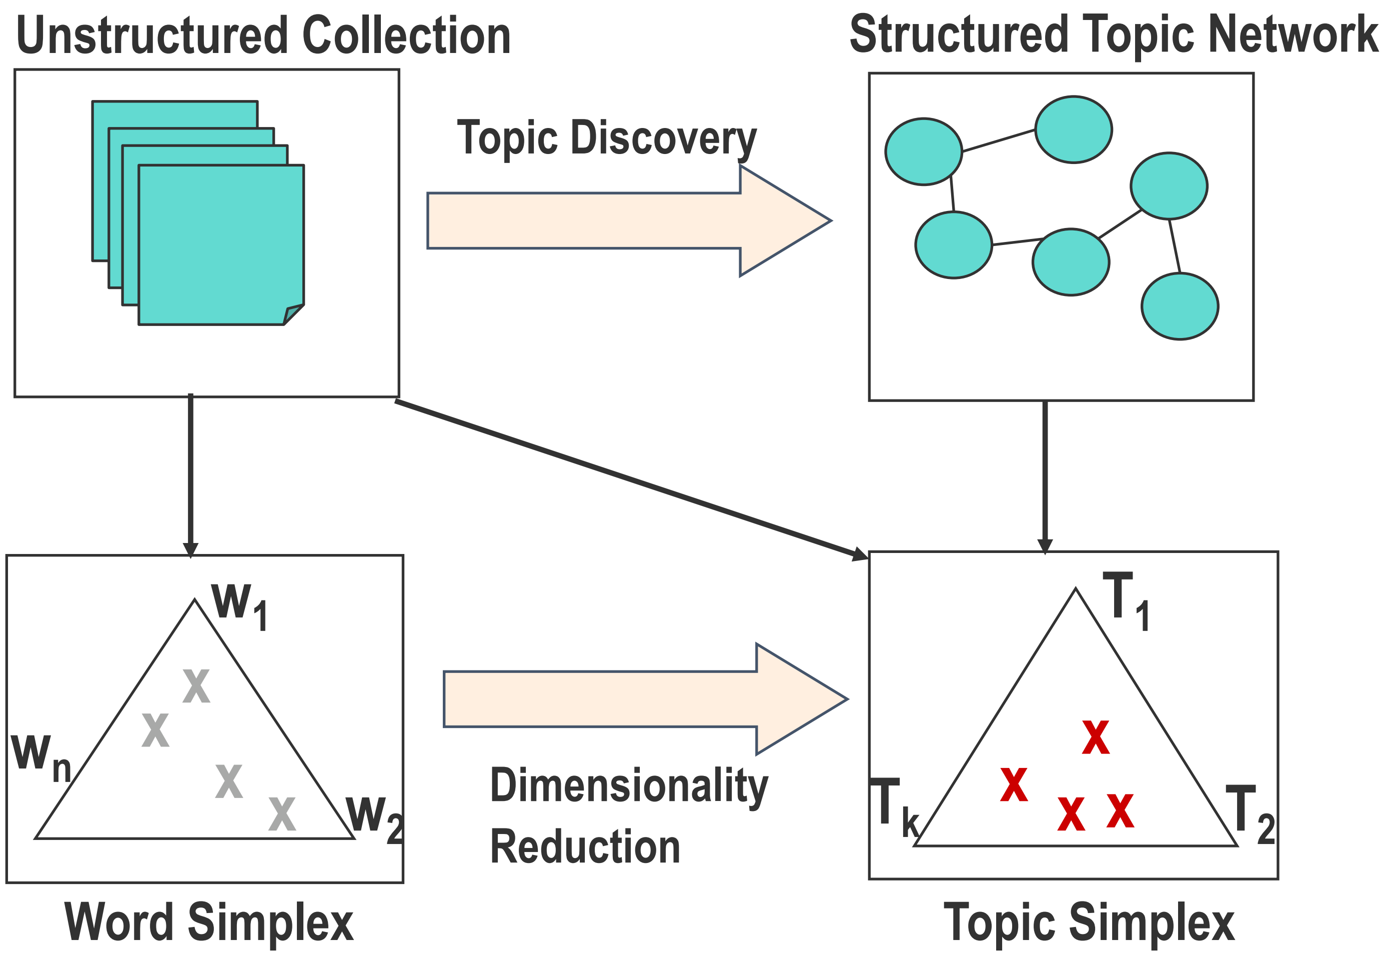

LSI versus Topic Model (probabilistic LSI)

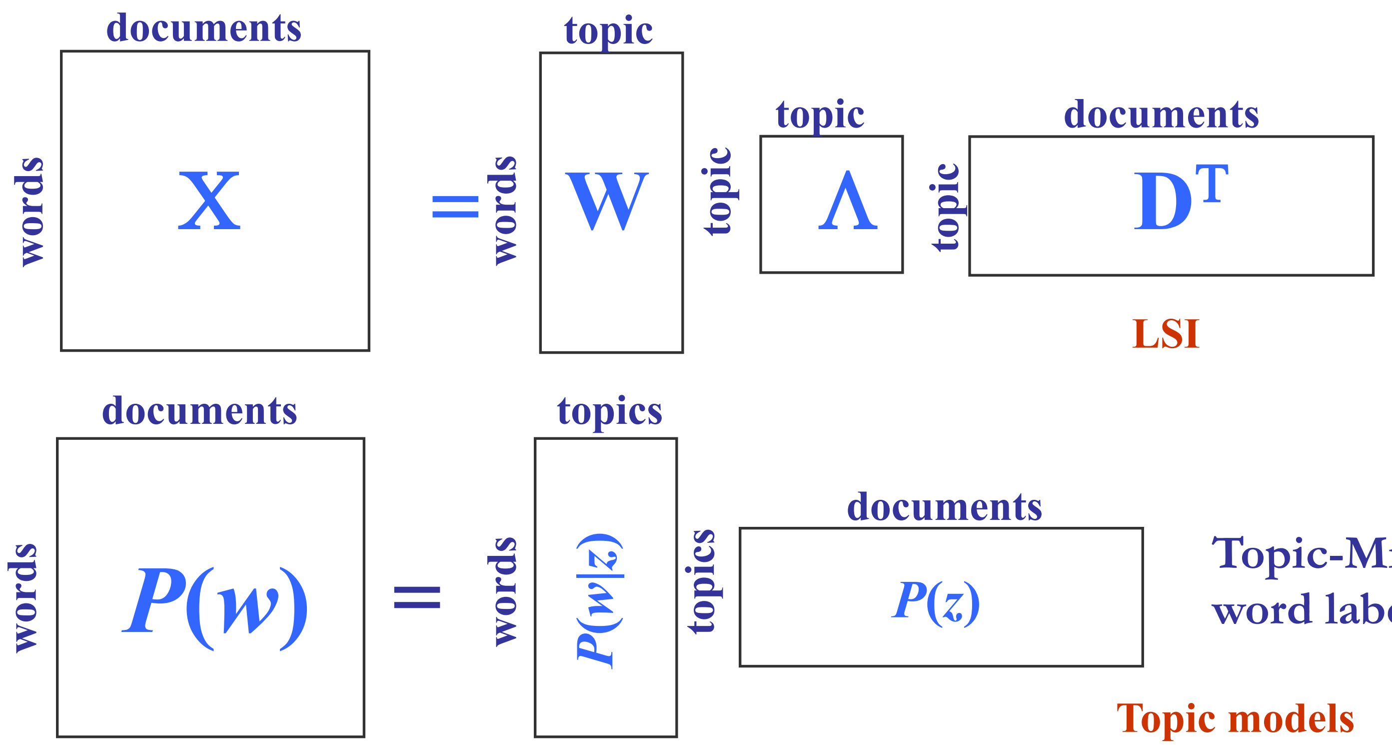

$$\vec{x} = W'\vec{d}$$

Topic-Mixing is via repeated word labeling

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**.

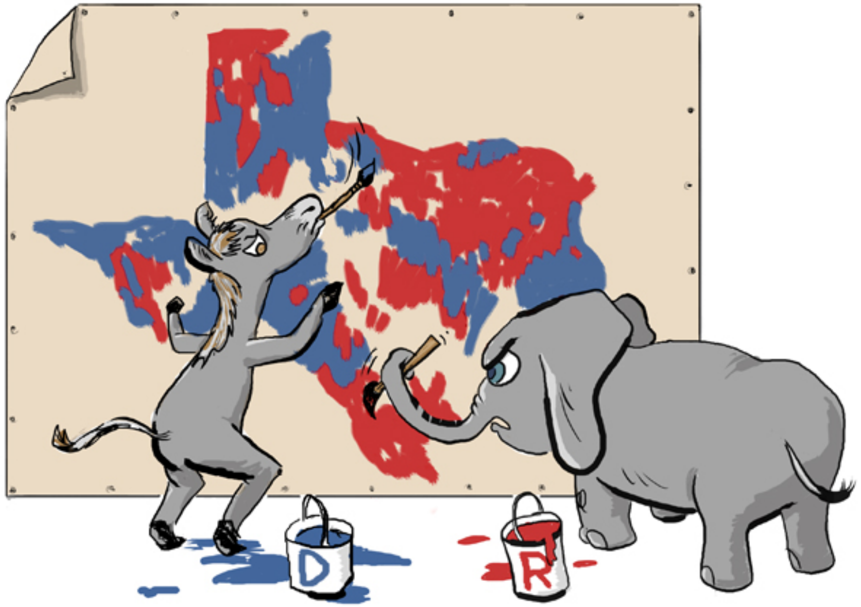

"Words" in Contexts (con'd)

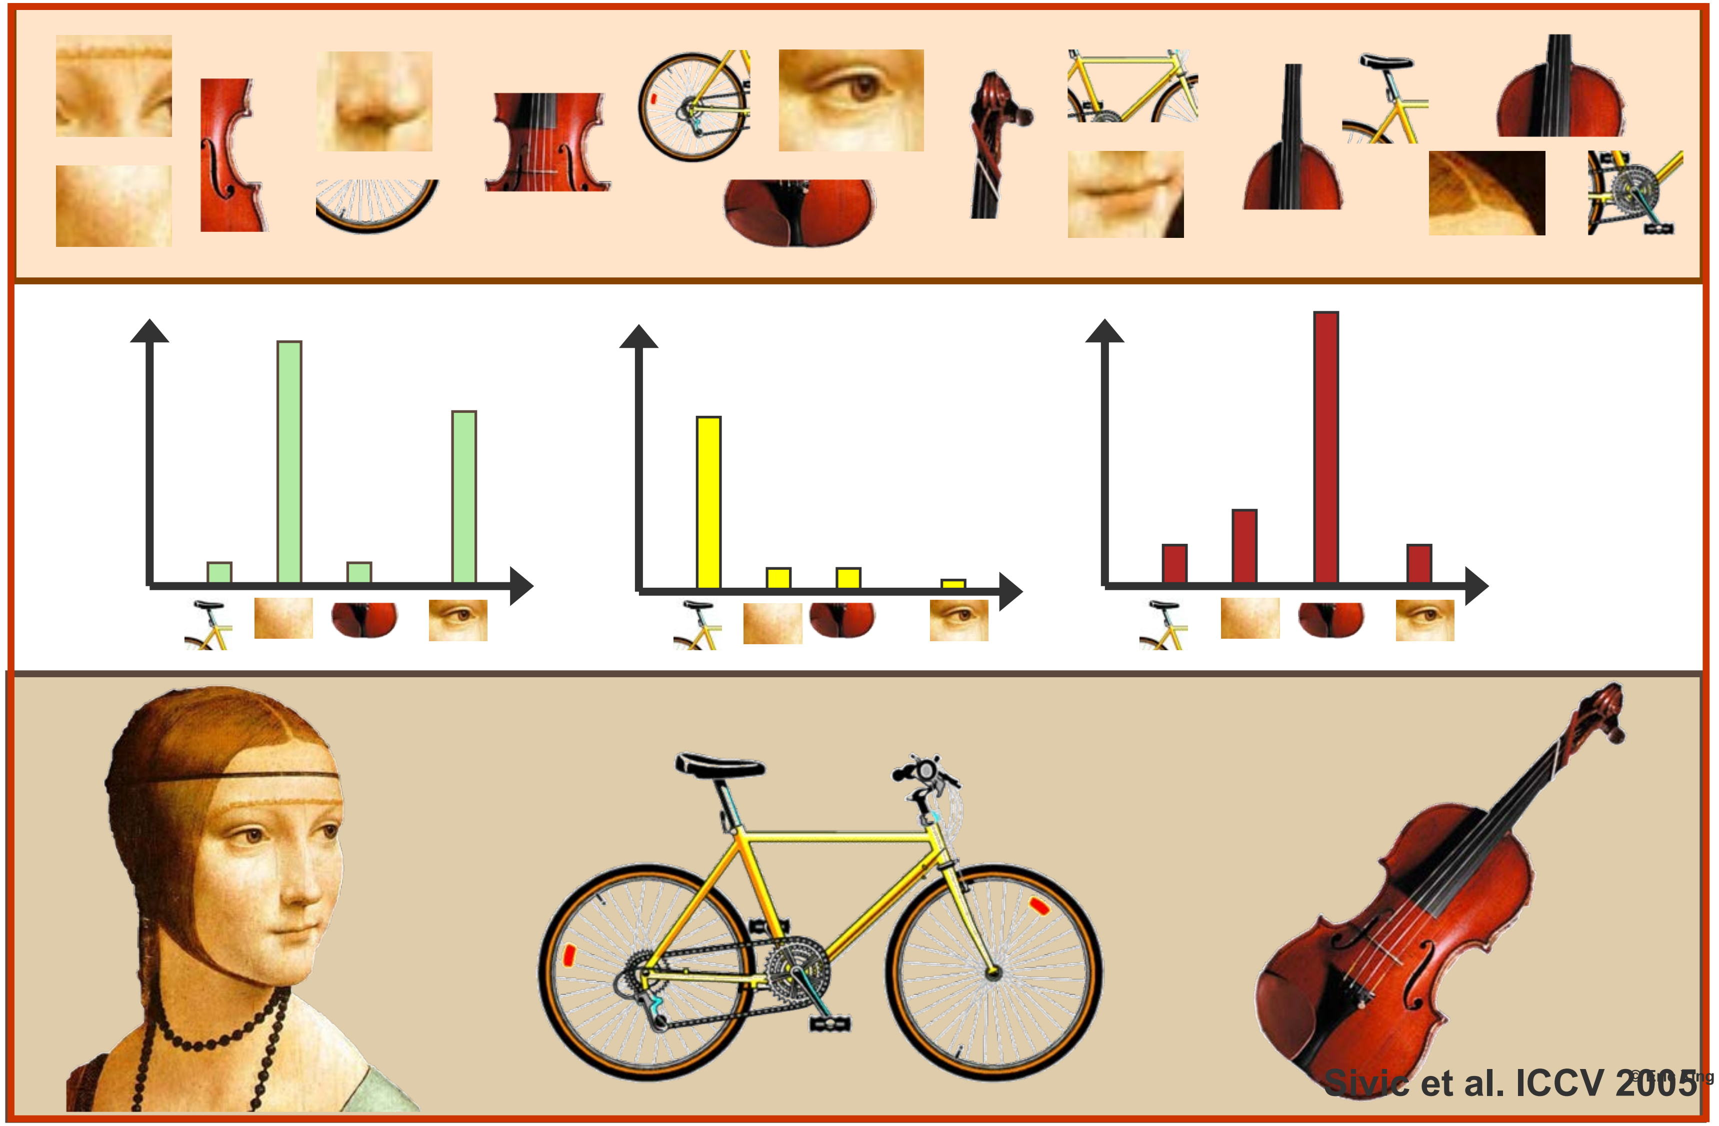

More Generally: Admixture Models

- Objects are bags of elements
- Mixtures are distributions over elements
- Objects have mixing vector θ
  - Represents each mixtures' contributions
- Object is generated as follows:
  - Pick a mixture component from θ
  - Pick an element from that component

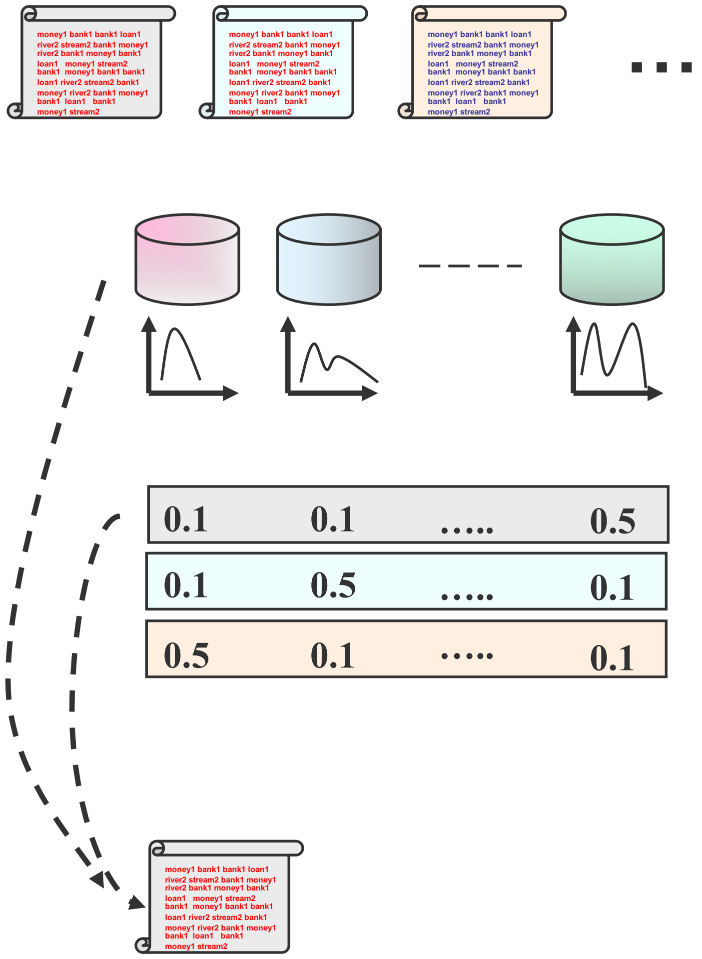

Topic Models Represented as a GM

Generating a document

– *Draw* θ from the prior

For each word n

- Draw z_n from $multinomial(\theta)$
- Draw $w_n \mid z_n, \{\beta_{1:k}\}$ from $multinomial(\beta_{z_n})$

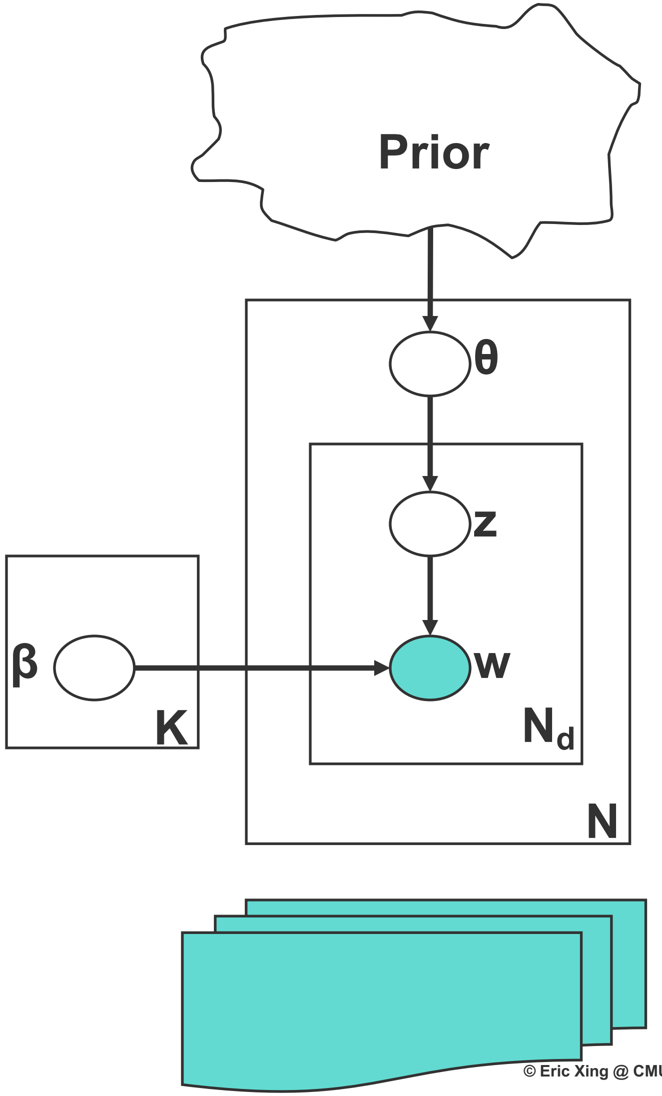

Which prior to use?

Choices of Priors

- Dirichlet (LDA) (Blei et al. 2003)
  - Conjugate prior means efficient inference
  - Can only capture variations in each topic's intensity independently

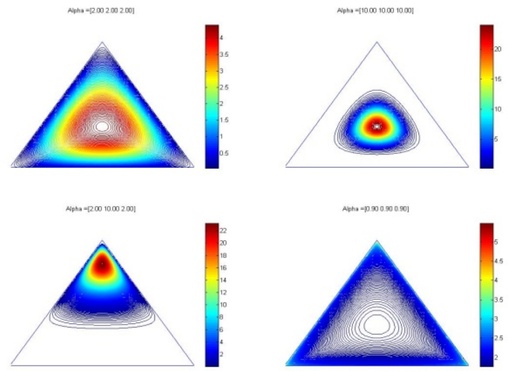

- Logistic Normal (CTM=LoNTAM) (Blei & Lafferty 2005, Ahmed & Xing 2006)
  - Capture the intuition that some topics are highly correlated and can rise up in intensity together
  - Not a conjugate prior implies hard inference

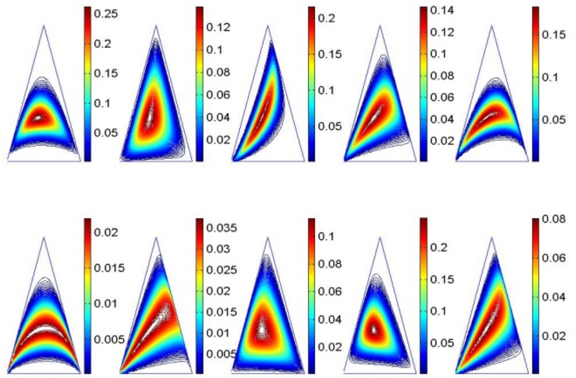

Generative Semantic of LoNTAM

Generating a document

– *Draw* θ from the prior

For each word n

- Draw z_n from $multinomial(\theta)$
- Draw $w_n \mid z_n, \{\beta_{1:k}\}$ from $multinomial(\beta_{z_n})$

$$\theta \sim LN_K(\mu, \Sigma)$$

$$\gamma \sim N_{K-1}(\mu, \Sigma) \qquad \gamma_K = 0$$

$$\theta_i = \exp\left\{\gamma_i - \log\left(1 + \sum_{i=1}^{K-1} e^{\gamma_i}\right)\right\}$$

$$C(\gamma) = \log\left(1 + \sum_{i=1}^{K-1} e^{\gamma_i}\right)$$

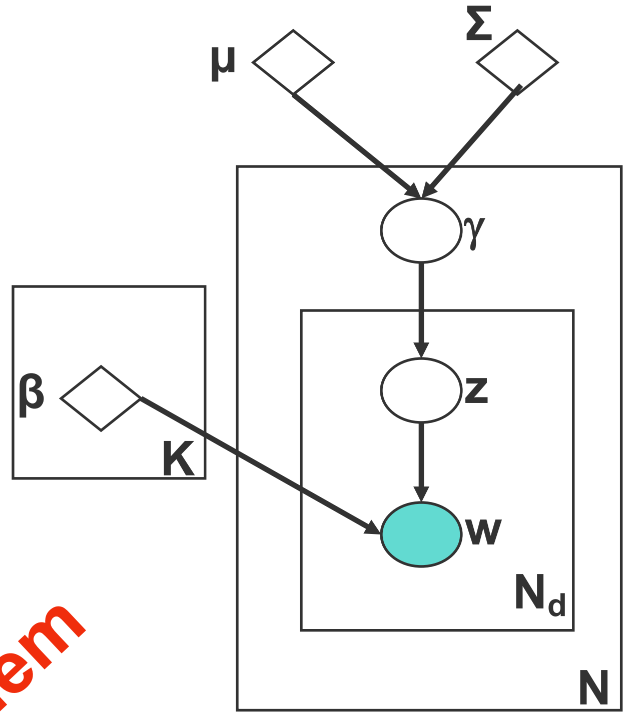

Problem

- Log Partition Function
- Normalization Constant

Posterior inference

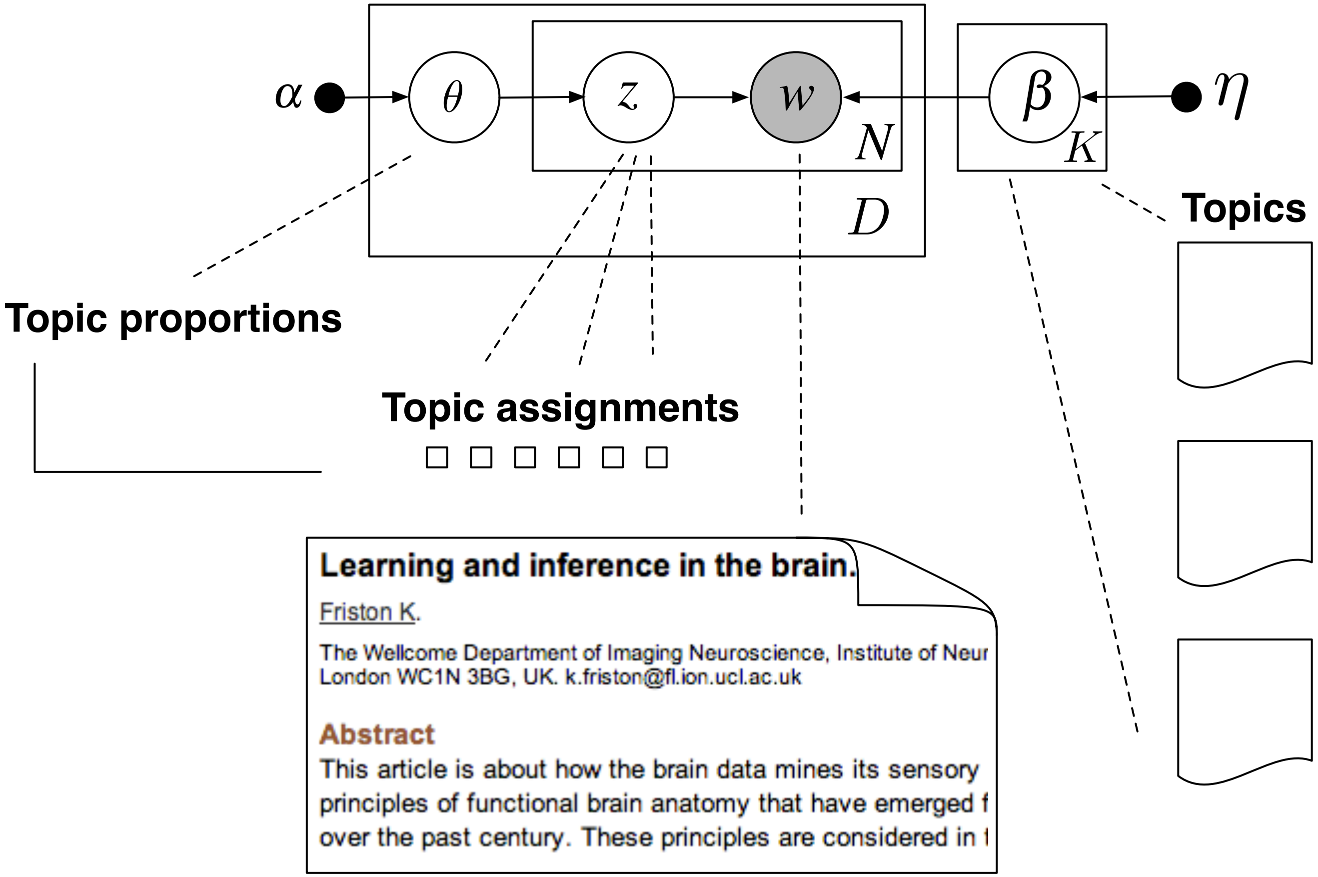

Posterior inference results

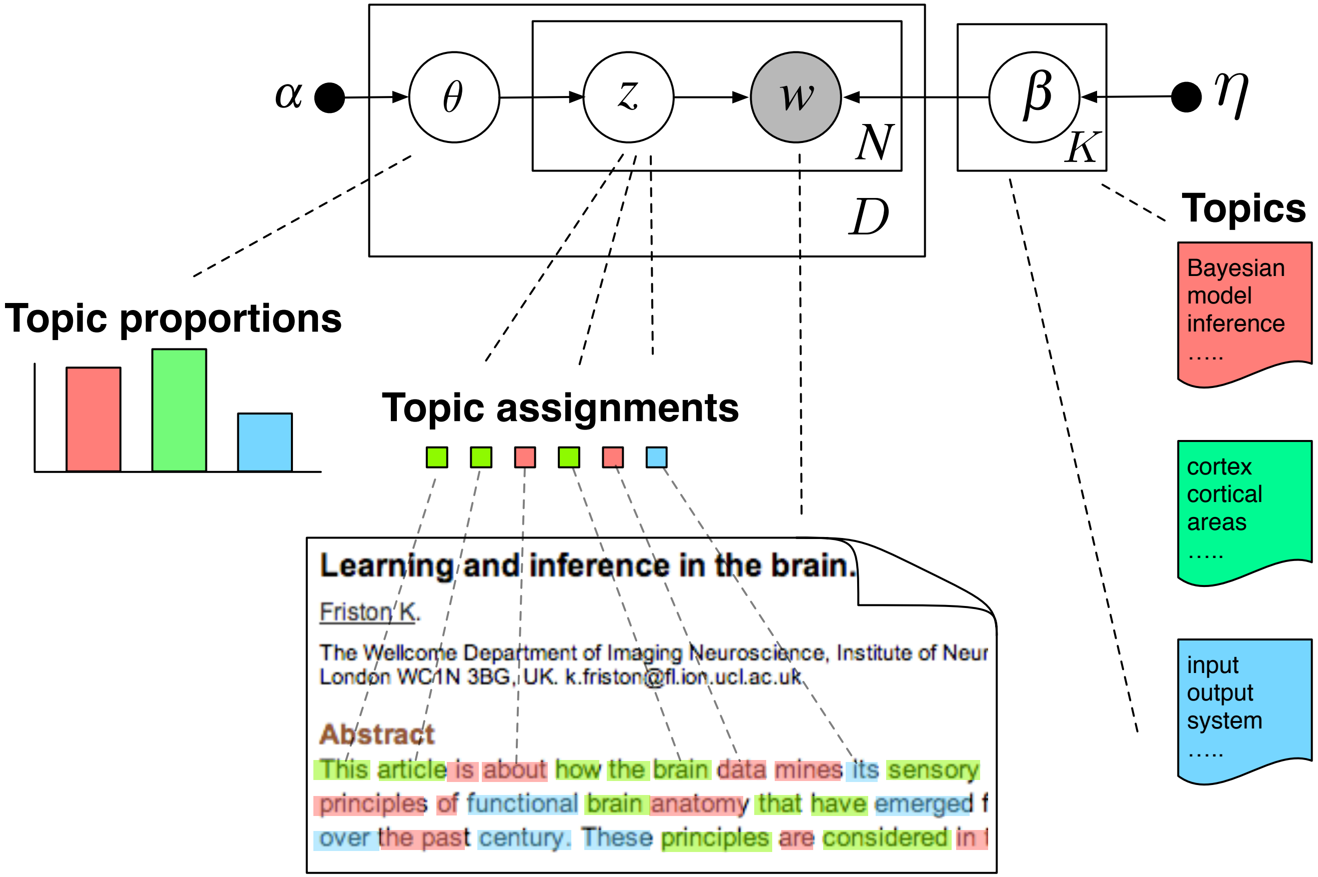

Joint likelihood of all variables

$$p(\beta, \theta, \boldsymbol{z}, \boldsymbol{w}) = \prod_{k=1}^{K} p(\beta_k|\eta) \prod_{d=1}^{D} p(\theta_d|\alpha) \prod_{n=1}^{N} p(z_{dn}|\theta_d) p(w_{dn}|z_{dn}, \beta)$$

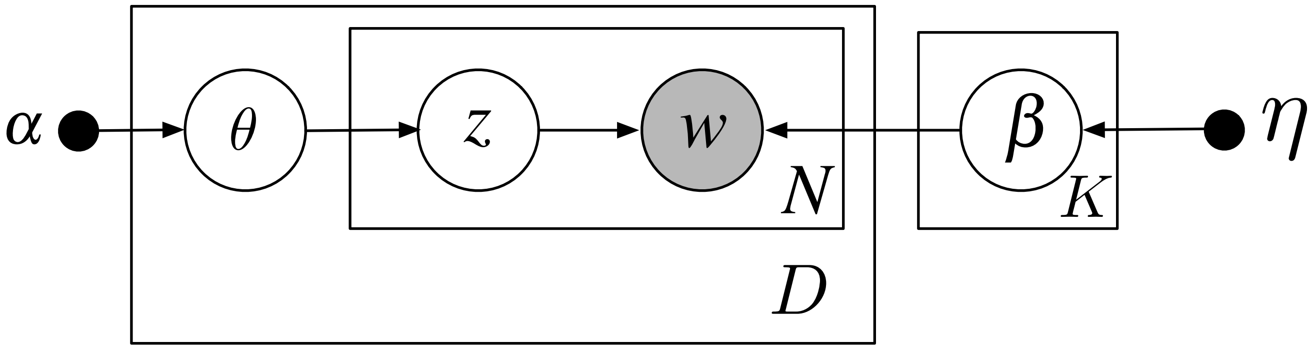

We are interested in computing the posterior, and the data likelihood!

Inference and Learning are both intractable

- A possible query:

$$p(\theta_n \mid D) = ?$$

$$p(z_{n,m} \mid D) = ?$$

  - Close form solution?

$$p(\theta_n \mid D) = \frac{p(\theta_n, D)}{p(D)}$$

$$= \frac{\sum_{\{z_{n,m}\}} \int \left(\prod_n \left(\prod_m p(w_{n,m} \mid \beta_{z_n}) p(z_{n,m} \mid \theta_n) \right) p(\theta_n \mid \alpha) \right) p(\beta \mid \eta) d\theta_{-i} d\beta}{p(D)}$$

$$p(D) = \sum_{\{z_{n,m}\}} \int \cdots \int \left(\prod_n \left(\prod_m p(x_{n,m} \mid \beta_{z_n}) p(z_{n,m} \mid \theta_n) \right) p(\theta_n \mid \alpha) \right) p(\beta \mid \eta) d\theta_1 \cdots d\theta_N d\beta$$

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

- Learning: What to learn? What is the objective function?

Approximate Inference

- Variational Inference
  - Mean field approximation (Blei et al.)
  - Expectation propagation (Minka et al.)
  - Variational 2nd-order Taylor approximation (Xing)
- Markov Chain Monte Carlo
  - Gibbs sampling (Griffiths et al)

Variational Inference

- Consider a generative model $p_\theta(\boldsymbol{x}|\boldsymbol{z})$, and prior $p(\boldsymbol{z})$
  - Joint distribution: $p_\theta(\boldsymbol{x},\boldsymbol{z}) = p_\theta(\boldsymbol{x}|\boldsymbol{z})p(\boldsymbol{z})$
- Assume variational distribution $q_\phi(\boldsymbol{z}|\boldsymbol{x})$
- Objective: Maximize lower bound for log likelihood

$$
\begin{aligned}
&\log p(\boldsymbol{x}) \\
&= KL\left(q_\phi(\boldsymbol{z}|\boldsymbol{x}) \,||\, p_{\boldsymbol{\theta}}(\boldsymbol{z}|\boldsymbol{x})\right) + \int_{\boldsymbol{z}} q_\phi(\boldsymbol{z}|\boldsymbol{x}) \log \frac{p_\theta(\boldsymbol{x},\boldsymbol{z})}{q_\phi(\boldsymbol{z}|\boldsymbol{x})} \\
&\geq \int_{\boldsymbol{z}} q_\phi(\boldsymbol{z}|\boldsymbol{x}) \log \frac{p_\theta(\boldsymbol{x},\boldsymbol{z})}{q_\phi(\boldsymbol{z}|\boldsymbol{x})} \\
&:= \mathcal{L}(\boldsymbol{\theta},\boldsymbol{\phi};\boldsymbol{x})
\end{aligned}
$$

  - Equivalently, minimize free energy

$$F(\boldsymbol{\theta},\phi;\boldsymbol{x}) = -\log p(\boldsymbol{x}) + KL(q_\phi(\boldsymbol{z}|\boldsymbol{x}) \,||\, p_{\boldsymbol{\theta}}(\boldsymbol{z}|\boldsymbol{x}))$$

Variational Inference

Maximize the variational lower bound:

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \boldsymbol{x}) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)] + KL\left(q_\phi\,(z|x)||p(z)\right)$$
$$= \log p(\boldsymbol{x}) - KL(q_\phi(\boldsymbol{z}|\boldsymbol{x})\,||\,p_{\boldsymbol{\theta}}(\boldsymbol{z}|\boldsymbol{x}))$$

- E-step: maximize $\mathcal{L}$ w.r.t. $\boldsymbol{\phi}$, with $\boldsymbol{\theta}$ fixed

$$\max_\phi \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \boldsymbol{x})$$

  - If closed form solutions exist:

$$q_\phi^*(z|x) \propto \exp[\log p_\theta(x, z)]$$

- M-step: maximize $\mathcal{L}$ w.r.t. $\boldsymbol{\theta}$, with $\boldsymbol{\phi}$ fixed

$$\max_\theta \mathcal{L}(\boldsymbol{\theta}, \boldsymbol{\phi}; \boldsymbol{x})$$

Mean-field assumption (in topic models)

- True posterior

$$p(\beta, \theta, \boldsymbol{z} | \boldsymbol{w}) = \frac{p(\beta, \theta, \boldsymbol{z}, \boldsymbol{w})}{p(\boldsymbol{w})}$$

- Break the dependency using the fully factorized distribution

$$q(\beta, \theta, \boldsymbol{z}) = \prod_k q(\beta_k) \prod_d q(\theta_d) \prod_n q(z_{dn})$$

- Mean-field family usually does NOT include the true posterior.

Mean Field Approximation

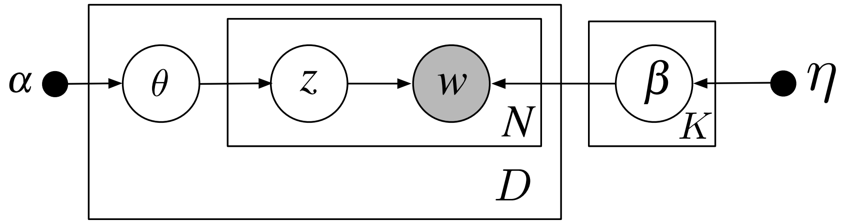

- Parametric form for each marginal factor in $q(\beta, z, \theta \mid \lambda, \phi, \gamma)$:

$$q(\beta_k \mid \lambda_k) = \text{Dirichlet}(\beta_k \mid \lambda_k)$$
$$q(\theta_d \mid \gamma_d) = \text{Dirichlet}(\theta_d \mid \gamma_d)$$
$$q(z_{dn} \mid \phi_{dn}) = \text{Multinomial}(z_{dn} \mid \phi_{dn})$$

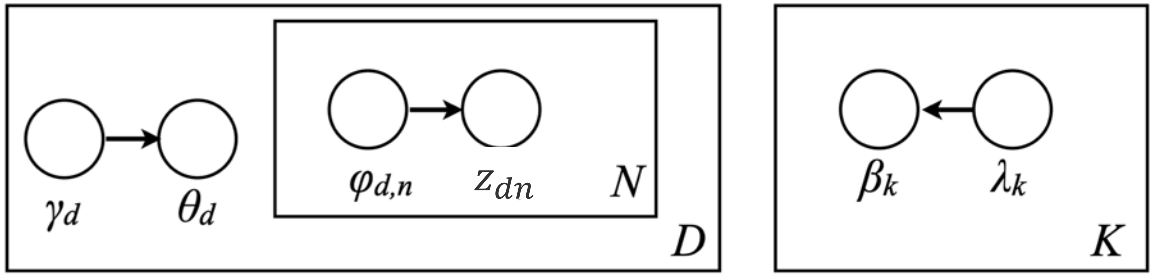

- Learning parameters of the variational distribution (E-step):

$$\gamma^\star, \lambda^\star, \phi^\star = \arg\min_{\gamma,\lambda,\phi} \text{KL}(q(\beta, \theta, \mathbf{z} \mid \gamma, \phi) \parallel p(\beta, \theta, \mathbf{z} \mid \mathbf{w}, \alpha, \eta))$$

- For LDA, we can compute the optimal MF approximation in closed form.

Update each marginal

- Update: $$q(\theta_d) \propto \exp\left\{\mathbb{E}_{\prod_n q(z_{dn})}\left[\log p(\theta_d|\alpha) + \sum_n \log p(z_{dn}|\theta_d)\right]\right\}$$

- Where in LDA: $$p(\theta_d|\alpha) \propto \exp\left\{\sum_{k=1}^{K}(\alpha_k - 1)\log\theta_{dk}\right\} \text{ --Dirichlet}$$
$$p(z_{dn}|\theta_d) = \exp\left\{\sum_{k=1}^{K} 1[z_{dn} = k]\log\theta_{dk}\right\} \text{ --Multinomial}$$

- And we obtain: $$q(\theta_d) \propto \exp\left\{\sum_{k=1}^{K}\left(\sum_{n=1}^{N} q(z_{dn} = k) + \alpha_k - 1\right)\log\theta_{dk}\right\}$$

This is also a Dirichlet — the same as its prior!

Update each marginal

- Similarly to $q(\theta_d \mid \gamma_d)$, we obtain optimal parameters $\phi^{\star}_{dn}$ for $q(z_{dn} \mid \phi_{dn})$:

$$q(z_d n = k \mid \phi_{dn}) = \phi_{dn}(k) = \beta_k(w_{dn}) \exp\left\{ \Psi(\gamma_d(k)) - \Psi(\sum_{j=1}^{K} \gamma_d(j)) \right\}$$

- And optimal parameters $\lambda^{\star}_k$ for $q(\beta_k \mid \lambda_k)$:

$$\lambda_k(j) = \eta(j) + \sum_{d=1}^{D} \sum_{n=1}^{N_d} \phi^{\star}_{dn}(k) 1[w_{dn} = j]$$

- Iterating these equations to convergence yields the MF approximation to the posterior distribution.

Coordinate ascent algorithm for LDA

Initialize variational topics $q(\beta_k)$, $k = 1, \ldots, K$.

repeat

  for each document $d \in \{1, 2, \ldots, D\}$ **do**

    Initialize variational topic assigments $q(z_{dn})$, $n = 1, \ldots, N$

    repeat

      Update variational topic proportions $q(\theta_d)$

      Update variational topic assigments $q(z_{dn})$, $n = 1, \ldots, N$

    until Change of $q(\theta_d)$ is small enough

  end for

  Update variational topics $q(\beta_k)$, $k = 1, \ldots, K$.

until Lower bound $L(q)$ converges

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

Supplementary

Supplementary: More on strategies in VI

- Alternative approximation scheme
- How to evaluate: empirical (ground truth unknown) vs. simulation (ground truth known)
- Comparison (of what)
- Building blocks

Recall Choices of Priors

- Dirichlet (LDA) (Blei et al. 2003)
  - Conjugate prior means efficient inference
  - Can only capture variations in each topic's intensity independently

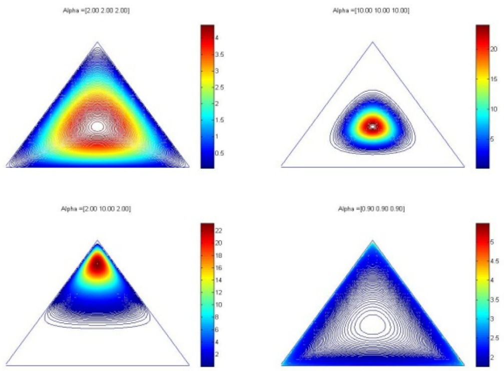

- Logistic Normal (CTM=LoNTAM) (Blei & Lafferty 2005, Ahmed & Xing 2006)
  - Capture the intuition that some topics are highly correlated and can rise up in intensity together
  - Not a conjugate prior implies hard inference

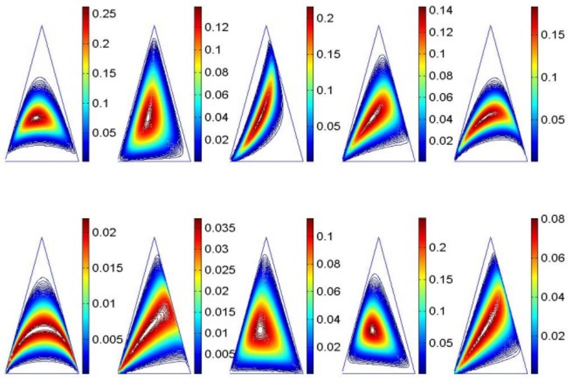

Choice of q() does matter

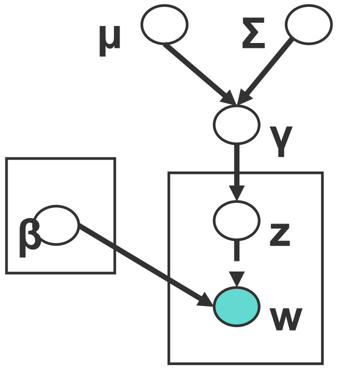

$$P(\gamma,\{z\}|D)$$

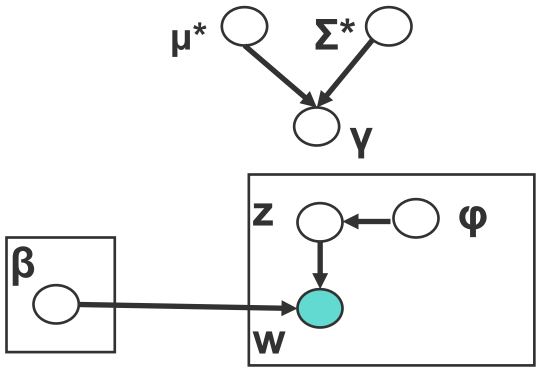

$$q(\gamma, z_{1:n}) = q(\gamma|\mu^*, \Sigma^*)\prod q(z_n|\phi_n)$$

Log Partition Function

$$\log\left(1+\sum_{i=1}^{K-1} e^{\gamma_i}\right)$$

Σ* is full matrix

Multivariate Quadratic Approx.

Closed Form Solution for μ*, Σ*

Ahmed&Xing

Σ* is assumed to be diagonal

Tangent Approx.

Numerical Optimization to fit μ*, Diag(Σ*)

Blei&Lafferty

Tangent Approximation

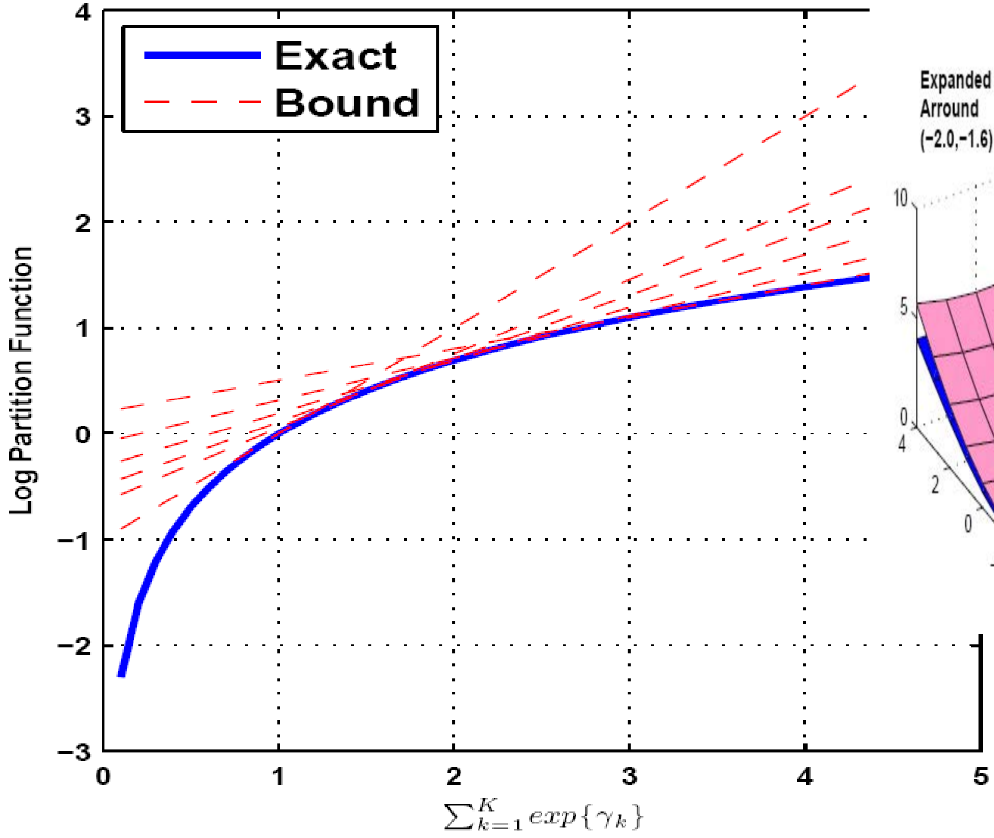

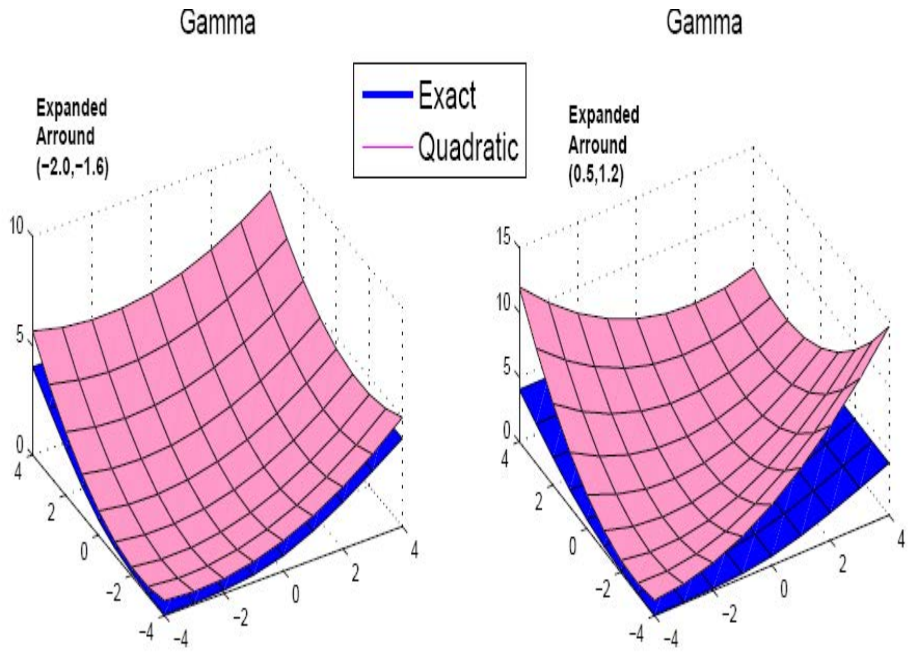

How to evaluate?

- Empirical Visualization: e.g., topic discovery on New York Times

The 5 most frequent topics from the HDP on the *New York Times*.

game	life	film	book	wine
season	know	movie	life	street
team	school	show	books	hotel
coach	street	life	novel	house
play	man	television	story	room
points	family	films	man	night
games	says	director	author	place
giants	house	man	house	restaurant
second	children	story	war	park
players	night	says	children	garden

How to evaluate?

- Test on Synthetic Text where ground truth is known:

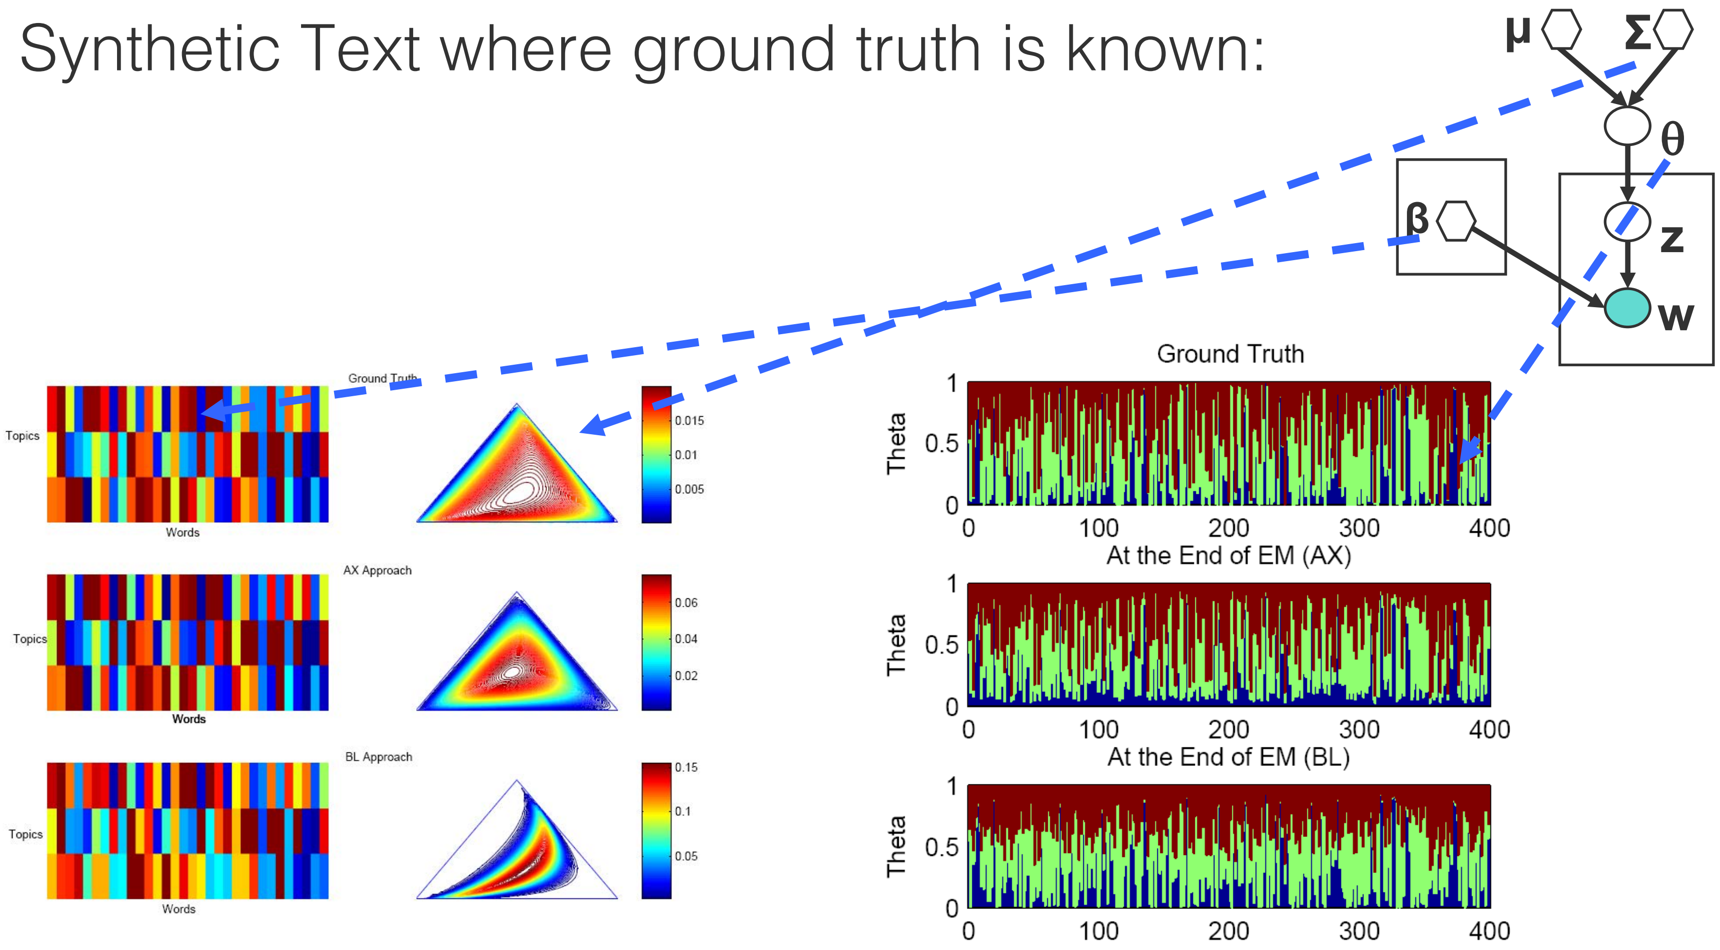

Comparison: accuracy and speed

- L2 error in topic vector est. and # of iterations
  - Varying Num. of Topics
  - Varying Voc. Size
  - Varying Num. Words Per Document

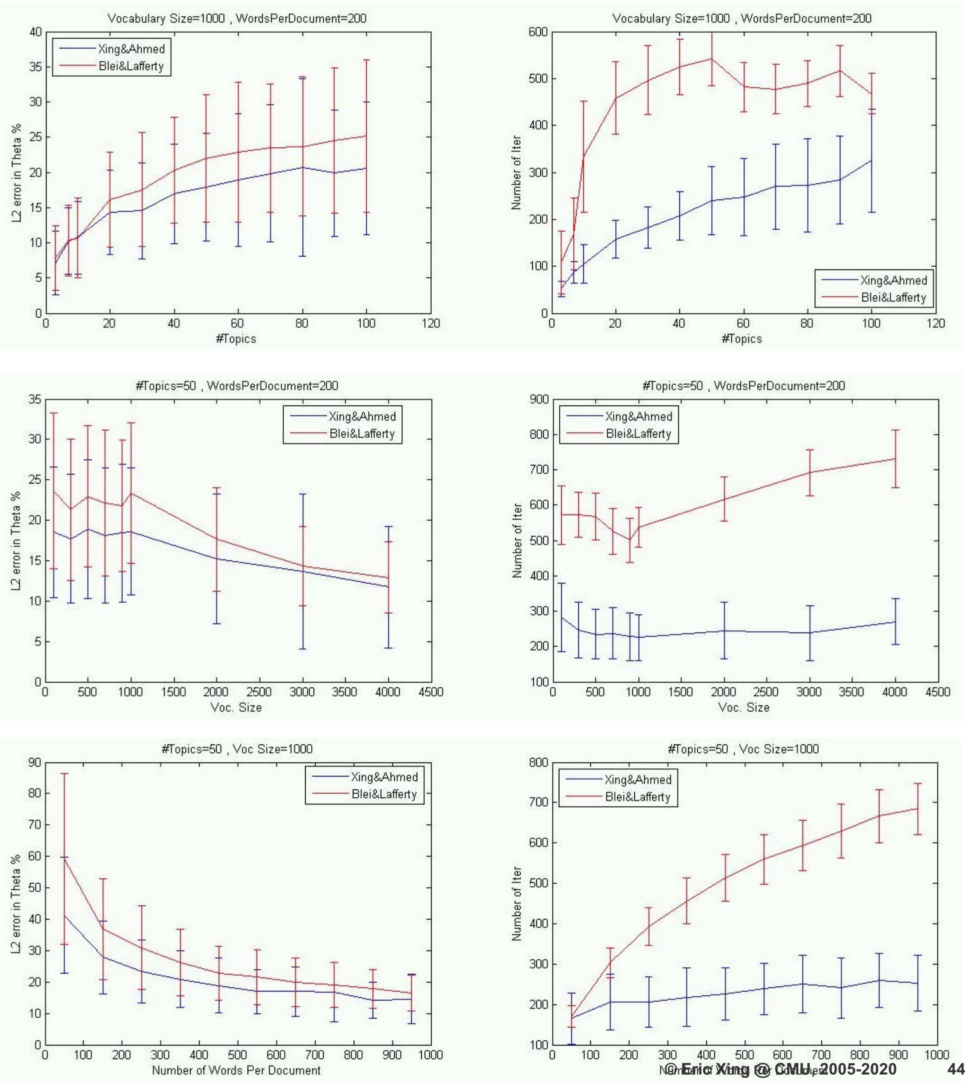

Comparison: perplexity

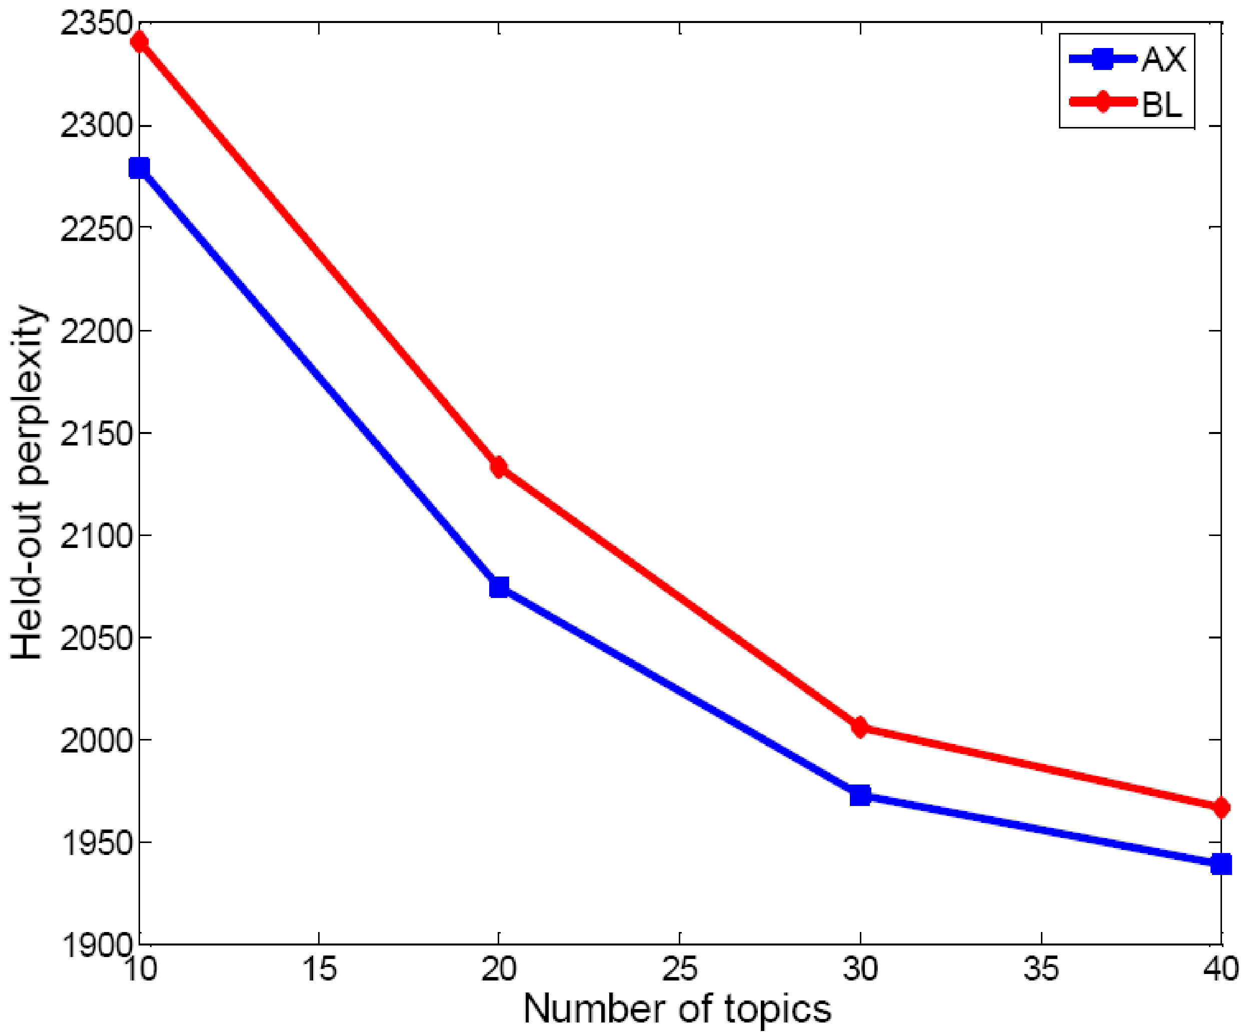

Classification Result on PNAS collection

- PNAS abstracts from 1997-2002
  - 2500 documents
  - Average of 170 words per document
- Fitted 40-topics model using both approaches
- Use low dimensional representation to predict the abstract category
  - Use SVM classifier
  - 85% for training and 15% for testing

Classification Accuracy

Category	Doc	BL	AX
Genetics	21	61.9	61.9
Biochemistry	86	65.1	77.9
Immunology	24	70.8	66.6
Biophysics	15	53.3	66.6
Total	146	64.3	72.6

-Notable Difference
-Examine the low dimensional representations below

What makes topic models useful --- The Zoo of Topic Models!

- It is a building block of many models.

Williamson et al. 2010

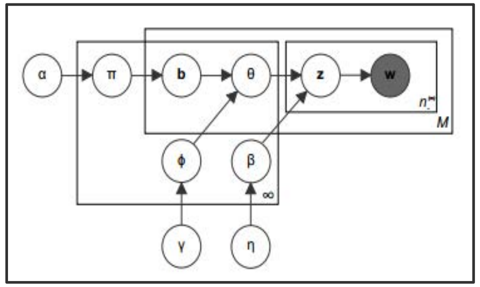

Chang & Blei, 2009

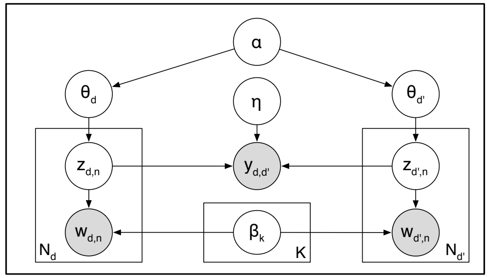

Titov & McDonald, 2008

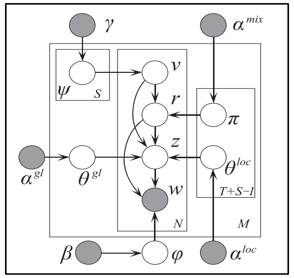

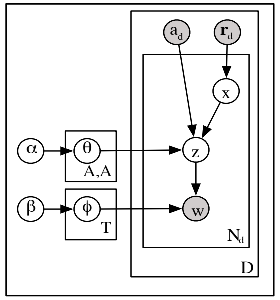

McCallum et al. 2007

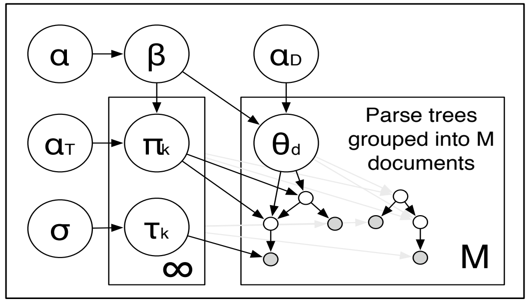

Boyd-Graber & Blei, 2008

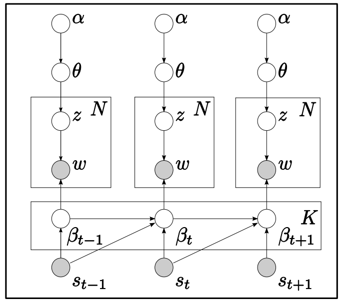

Wang & Blei, 2008

More on Mean Field Approximation

The naive mean field approximation

- Approximate $p(\mathbf{X})$ by fully factorized $q(\mathbf{X})=\prod_i q_i(X_i)$
- For Boltzmann distribution $p(X)=\exp\{\sum_{i<j} q_{ij}X_iX_j+q_{io}X_i\}/Z$:

mean field equation:

$$q_i(X_i) = \exp\left\{\theta_{i0}X_i + \sum_{j\in\mathcal{N}_i}\theta_{ij}X_i\langle X_j\rangle_{q_j} + A_i\right\}$$

$$= p(X_i \mid \{\langle X_j\rangle_{q_j} : j \in \mathcal{N}_i\})$$

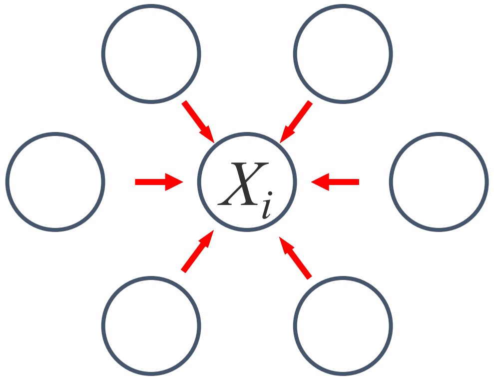

- $\langle X_j\rangle_{q_j}$ resembles a "message" sent from node j to i
- $\{\langle X_j\rangle_{q_j} : j \in \mathcal{N}_i\}$ forms the "mean field" applied to X_i from its neighborhood

Cluster-based approx. to the Gibbs free energy

(Wiegerinck 2001, Xing *et al* 03,04)

Exact: $G[p(X)]$ *(intractable)*

Clusters: $G[\{q_c(X_c)\}]$

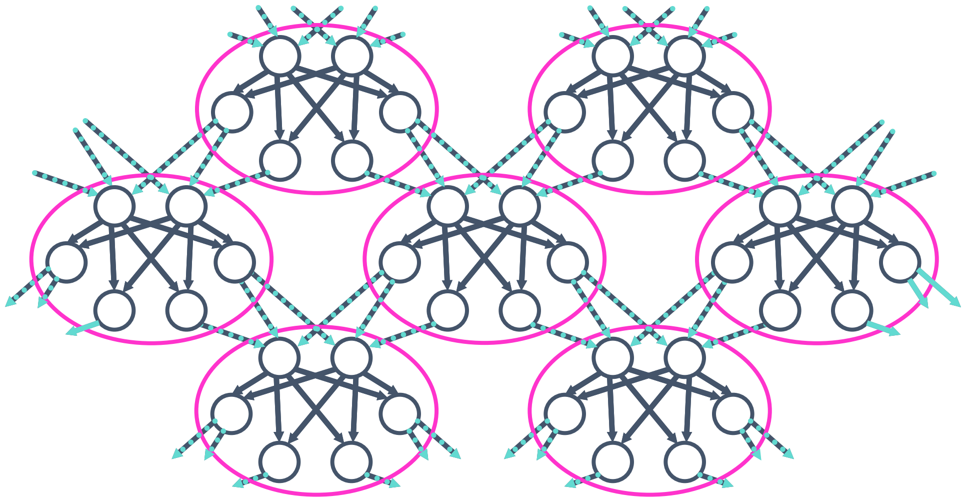

Mean field approx. to Gibbs free energy

- Given a disjoint clustering, $\{C_1, \dots, C_I\}$, of all variables
- Let

$$q(\mathbf{X}) = \prod_i q_i(\mathbf{X}_{C_i}),$$

- Mean-field free energy

$$G_{\mathrm{MF}} = \sum_i \sum_{\mathbf{x}_{C_i}} \prod_i q_i(\mathbf{x}_{C_i}) E(\mathbf{x}_{C_i}) + \sum_i \sum_{\mathbf{x}_{C_i}} q_i(\mathbf{x}_{C_i}) \ln q_i(\mathbf{x}_{C_i})$$

e.g., $G_{\mathrm{MF}} = \sum_{i<j} \sum_{x_i x_j} q(x_i) q(x_j) \phi(x_i x_j) + \sum_i \sum_{x_i} q(x_i) \phi(x_i) + \sum_i \sum_{x_i} q(x_i) \ln q(x_i)$ (naïve mean field)

  - Will **never** equal to the exact Gibbs free energy no matter what clustering is used, but it does **always** define a lower bound of the likelihood
- Optimize each $q_i(x_c)$'s.
  - Variational calculus …
  - Do inference in each $q_i(x_c)$ using any tractable algorithm

The Generalized Mean Field theorem

Theorem: The optimum GMF approximation to the cluster marginal is isomorphic to the cluster posterior of the original distribution given internal evidence and its generalized mean fields:

$$q_i^*(\mathbf{X}_{H,C_i}) = p(\mathbf{X}_{H,C_i} \mid \mathbf{x}_{E,C_i}, \langle \mathbf{X}_{H,MB_i} \rangle_{q_{j \neq i}})$$

GMF algorithm: Iterate over each q_i

A generalized mean field algorithm

[xing *et al*. UAI 2003]

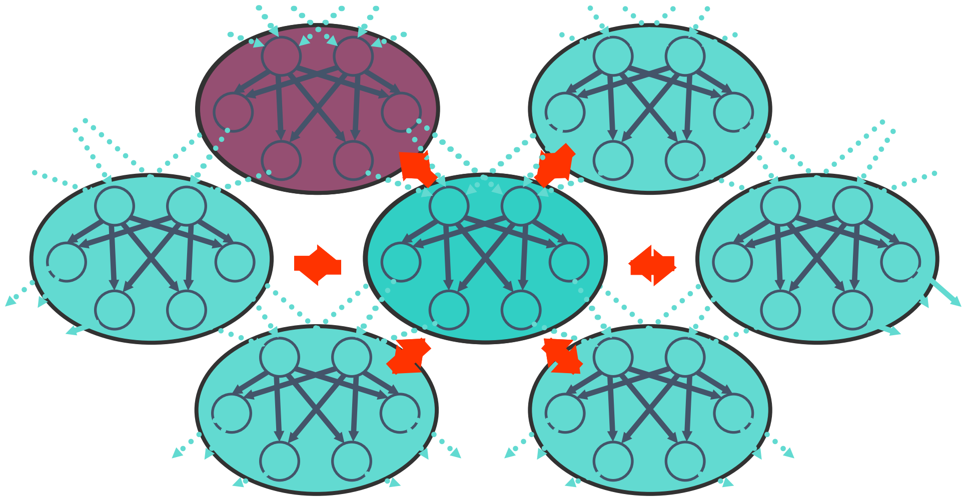

A generalized mean field algorithm

[xing *et al.* UAI 2003]

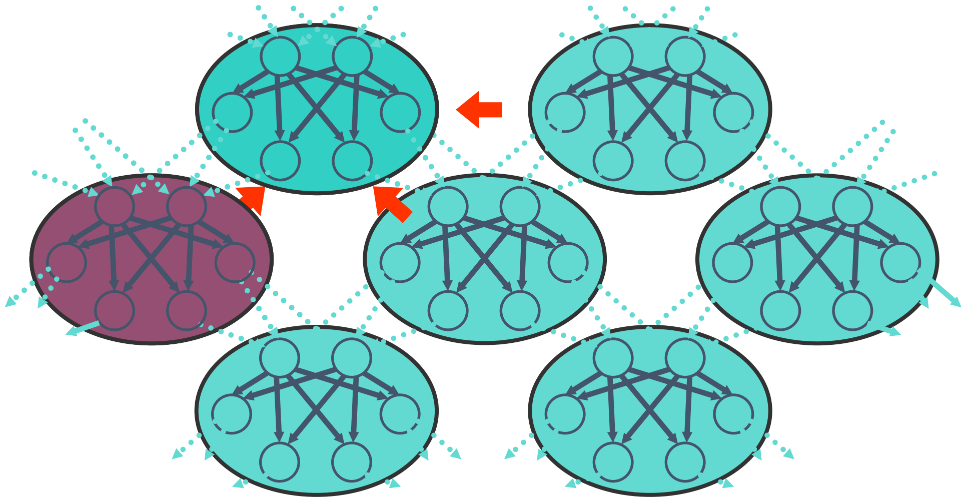

Convergence theorem

Theorem: The GMF algorithm is guaranteed to converge to a local optimum, and provides a lower bound for the likelihood of evidence (or partition function) the model.

Example 1: Generalized MF approximations to Ising models

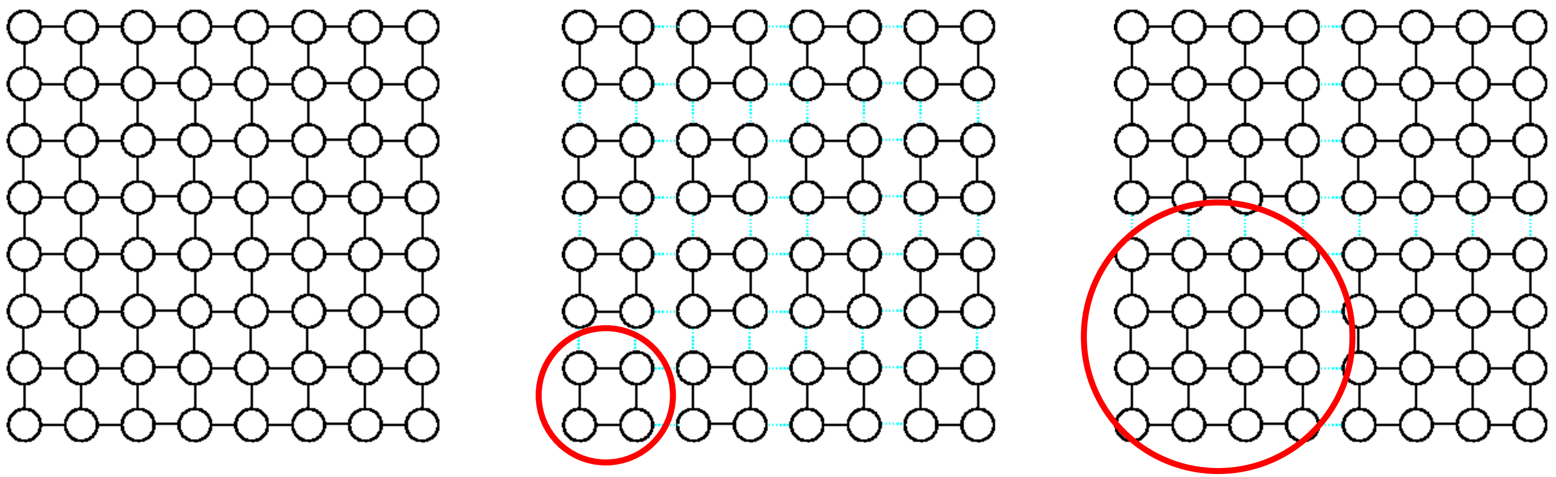

Cluster marginal of a square block C_k:

$$q(X_{C_k}) \propto \exp\left\{ \sum_{i,j \in C_k} \theta_{ij} X_i X_j + \sum_{i \in C_k} \theta_{i0} X_i + \sum_{\substack{i \in C_k, j \in MB_k, \\ k' \in MBC_k}} \theta_{ij} X_i \langle X_j \rangle_{q(X_{C_{k'}})} \right\}$$

Virtually a reparameterized Ising model of small size.

GMF approximation to Ising models

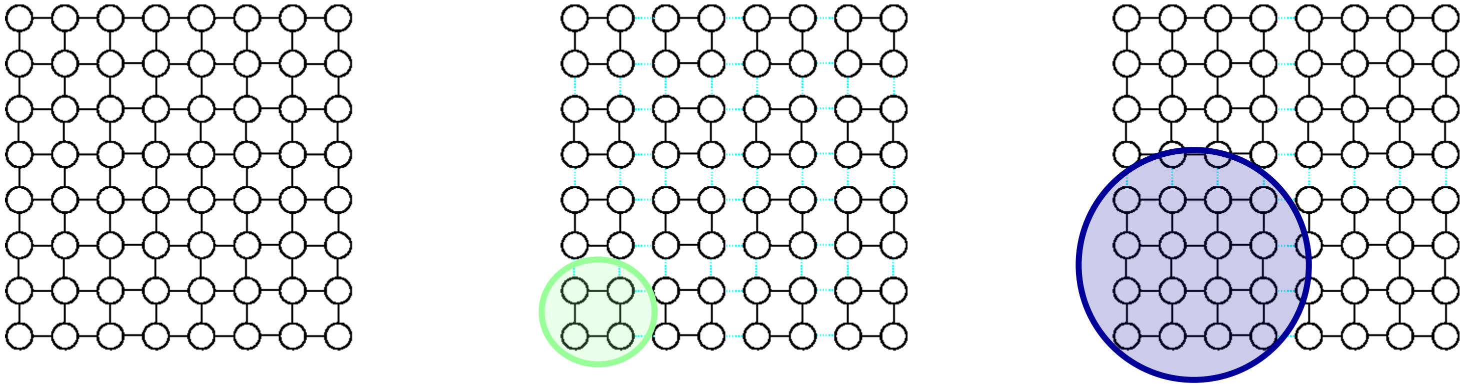

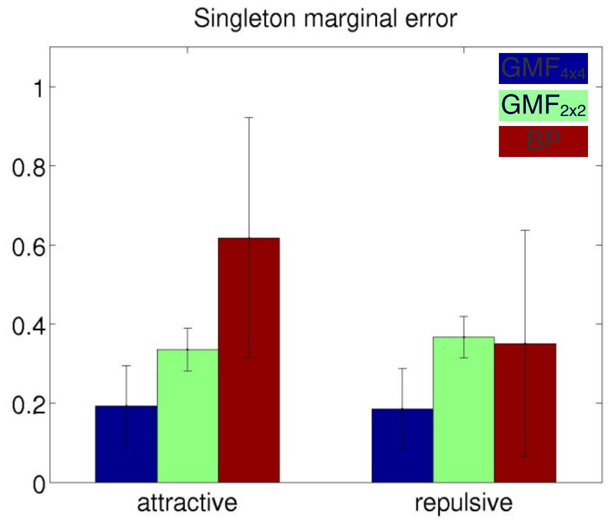

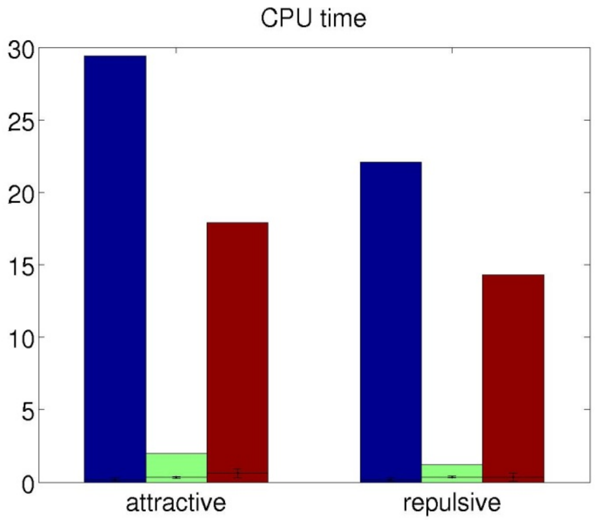

Attractive coupling: positively weighted
Repulsive coupling: negatively weighted

Example 2: Sigmoid belief network

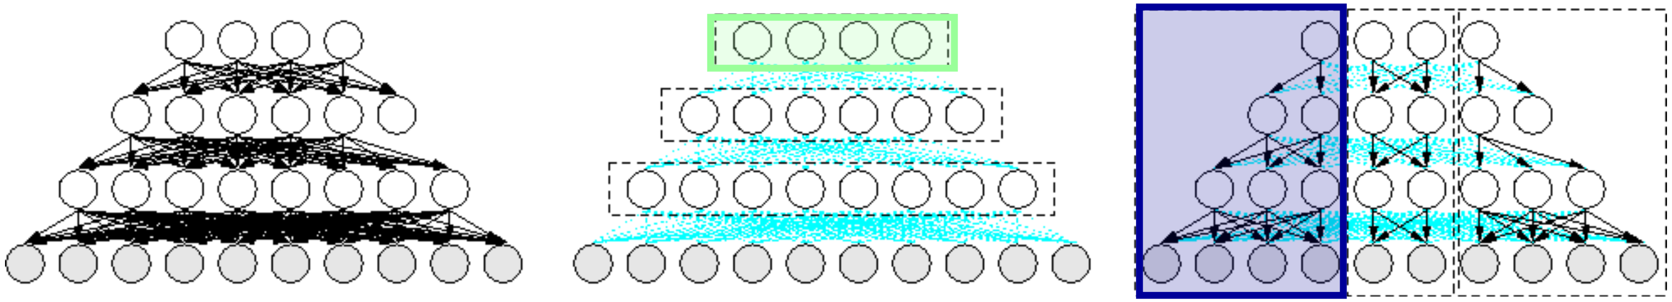

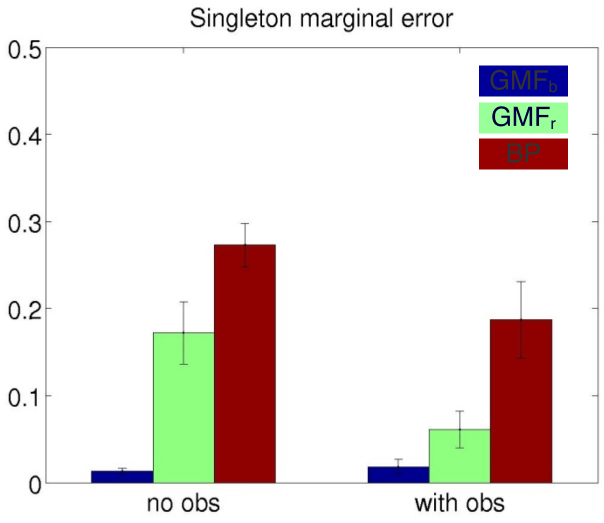

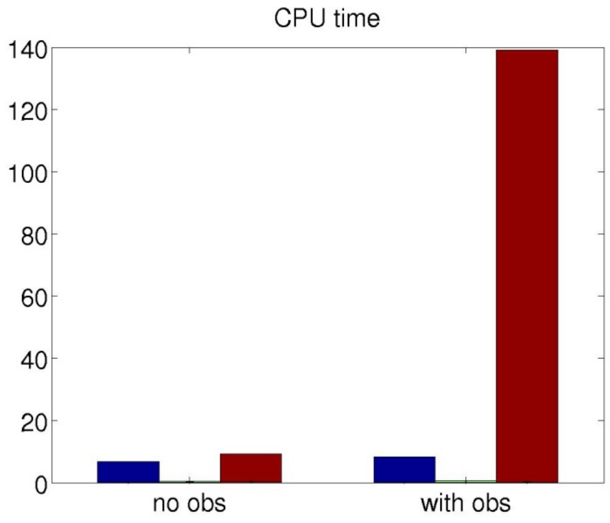

Example 3: Factorial HMM