## Probabilistic Graphical Models

Mean Fiend Approximation

\&<br>Topic Models



Eric Xing
Lecture 15, March 5, 2014

Reading: See class website

## Variational Principle

- Exact variational formulation

$$
A(\theta)=\sup _{\mu \in \mathcal{M}}\left\{\theta^{T} \mu-A^{*}(\mu)\right\}
$$

- $\mathcal{M}$ : the marginal polytope, difficult to characterize
- $A^{*}$ : the negative entropy function, no explicit form
- Mean field method: non-convex inner bound and exact form of entropy
- Bethe approximation and loopy belief propagation: polyhedral outer bound and non-convex Bethe approximation



## Mean Field Approximation

## Mean Field Methods

- For a given tractable subgraph F, a subset of canonical parameters is

$$
\mathcal{M}(F ; \phi):=\left\{\tau \in \mathbb{R}^{d} \mid \tau=\mathbb{E}_{\theta}[\phi(X)] \text { for some } \theta \in \Omega(F)\right\}
$$

- Inner approximation

$$
\mathcal{M}(F ; \phi)^{o} \subseteq \mathcal{M}(G ; \phi)^{o}
$$

- Mean field solves the relaxed problem

$$
\max _{\tau \in \mathcal{M}_{F}(G)}\left\{\langle\tau, \theta\rangle-A_{F}^{*}(\tau)\right\}
$$

- $A_{F}^{*}=\left.A^{*}\right|_{\mathcal{M}_{F}(G)}$ is the exact dual function restricted to $\mathcal{M}_{F}(G)$


## Tractable Subgraphs

- For an exponential family with sufficient statistics $\phi$ defined
on graph G , the set of realizable mean parameter set
- For an exponential family with sufficient statistics $\phi$ de
on graph G , the set of realizable mean parameter set

$$
\mathcal{M}(G ; \phi):=\left\{\mu \in \mathbb{R}^{d} \mid \exists p \text { s.t. } \mathbb{E}_{p}[\phi(X)]=\mu\right\}
$$

- Idea: restrict $p$ to a subset of distributions associated with a



## Example: Naïve Mean Field for Ising Model

- Ising model in $\{0,1\}$ representation

$$
p(x) \propto \exp \left\{\sum_{s \in V} x_{s} \theta_{s}+\sum_{(s, t) \in E} x_{s} x_{t} \theta_{s t}\right\}
$$

- Mean parameters

$$
\begin{aligned}
\mu_{s} & =E_{p}\left[X_{s}\right]=P\left[X_{s}=1\right] \text { for all s } V \text {, and } \\
\mu_{s t} & =E_{p}\left[X_{s} X_{t}\right]=P\left[\left(X_{s}, X_{t}\right)=(1,1)\right] \text { for all }(s, t) \quad E .
\end{aligned}
$$

- For fully disconnected graph F,

$\mathcal{M}_{F}(G):=\left\{\tau \in \mathbb{R}^{|V|+|E|} \mid 0 \leq \tau_{s} \leq 1, \forall s \in V, \tau_{s t}=\tau_{s} \tau_{t}, \forall(s, t) \in E\right\}$
- The dual decomposes into sum, one for each node

$$
A_{F}^{*}(\tau)=\sum_{s \in V}\left[\tau_{s} \log \tau_{s}+\left(1-\tau_{s}\right) \log \left(1-\tau_{s}\right)\right]
$$

## Naïve Mean Field for Ising Model

- Optimization Problem

$$
\max _{\mu \in[0,1]^{m}}\left\{\sum_{s \in V} \theta_{s} \mu_{s}+\sum_{(s, t) \in E} \theta_{s} \mu_{s} \mu_{t}+\sum_{s \in V} H_{s}\left(\mu_{s}\right)\right\}
$$

- Update Rule

$$
\mu_{s} \leftarrow \sigma\left(\theta_{s}+\sum_{t \in N(s)} \theta_{s t} \mu_{t}\right)
$$



- $\mu_{t}=p\left(X_{t}=1\right)=\mathbb{E}_{p}\left[X_{t}\right]$ resembles "message" sent from node $t$ to $s$
- $\left\{\mathbb{E}_{p}\left[X_{t}\right], t \in N(s)\right\}$ forms the "mean field" applied to $s$ from its neighborhood
- Also yields lower bound on log partition function

$$
K L(Q \| P)=-H_{Q}(X)-\sum_{f_{a} \in F} E_{Q} \log f_{a}\left(X_{a}\right)+\log Z
$$

## Geometry of Mean Field

- Mean field optimization is always non-convex for any exponential family in which the state space $\mathcal{X}^{m}$ is finite
- Recall the marginal polytope is a convex hull

$$
\mathcal{M}(G)=\operatorname{conv}\left\{\phi(e) ; e \in \mathcal{X}^{m}\right\}
$$

- $\mathcal{M}_{F}(G)$ contains all the extreme points
- If it is a strict subset, then it must be non-convex
- Example: two-node Ising model


$$
\mathcal{M}_{F}(G)=\left\{0 \leq \tau_{1} \leq 1,0 \leq \tau_{2} \leq 1, \tau_{12}=\tau_{1} \tau_{2}\right\}
$$

- It has a parabolic cross section along $\tau_{1}=\tau_{2}$, hence non-convex


## Cluster-based approx. to the <br> Gibbs free energy (Wiegerinck 2001, Xing et al 03,04)

Exact: $\quad G[p(X)] \quad$ (intractable)
Clusters: $G\left[\left\{q_{c}\left(X_{c}\right)\right\}\right]$


## Mean field approx. to Gibbs free energy

- Given a disjoint clustering, $\left\{\mathrm{C}_{1}, \ldots, \mathrm{C}_{\beta}\right\}$, of all variables
- Let

$$
q(\mathbf{X})=\prod_{i} q_{i}\left(\mathbf{X}_{C_{i}}\right)
$$

- Mean-field free energy

$$
\begin{aligned}
& G_{\mathrm{MF}}=\sum_{i} \sum_{\mathbf{x}_{c_{i}}} \prod_{i} q_{i}\left(\mathbf{x}_{C_{i}}\right) E\left(\mathbf{x}_{C_{i}}\right)+\sum_{i} \sum_{\mathbf{x}_{C_{i}}} q_{i}\left(\mathbf{x}_{C_{i}}\right) \ln q_{i}\left(\mathbf{x}_{C_{i}}\right) \\
& \text { e.g., } \quad G_{\mathrm{MF}}=\sum_{i j} \sum_{x_{i}, x_{j}} q\left(x_{i}\right) q\left(x_{j}\right) p\left(x_{i} x_{j}\right)+\sum_{i} \sum_{x_{i}} q\left(x_{i}\right) b\left(x_{i}\right)+\sum_{i} \sum_{x_{i}} q\left(x_{i}\right) \ln q\left(x_{i}\right) \quad \text { (naive mean field) }
\end{aligned}
$$

- 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 $q_{i}\left(x_{c}\right)$ 's.
- Variational calculus ...
- Do inference in each $q_{i}\left(x_{c}\right)$ 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:

$$
q_{i}^{*}\left(\mathbf{X}_{H, C_{i}}\right)=p\left(\mathbf{X}_{H, C_{i}} \mid \mathbf{x}_{E, C_{i}},\left\langle\mathbf{X}_{H, M B_{i}}\right\rangle_{q_{j \neq i}}\right)
$$

GMF algorithm: Iterate over each $q_{i}$

## A generalized mean field algorithm [xing etal. UA1 2003]



## A generalized mean field algorithm [xing etal. UA1 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.

## The naive mean field approximation

- Approximate $p(\mathbf{X})$ by fully factorized $q(\mathbf{X})=\mathrm{P}_{i} q_{i}\left(X_{i}\right)$
- For Boltzmann distribution $p(X)=\exp \left\{\sum_{i<j} q_{i j} X_{i} X_{j}+q_{i o} X_{i}\right\} / Z$ :
mean field equation:

$$
\begin{aligned}
q_{i}\left(X_{i}\right) & =\exp \left\{\theta_{i 0} X_{i}+\sum_{j \in \mathscr{N}_{i}} \theta_{i j} X_{i}\left\langle X_{j}\right\rangle_{q_{j}}+A_{i}\right\} \\
& =p\left(X_{i} \mid\left\{\left\langle X_{j}\right\rangle_{q_{j}}: j \in \mathcal{N}_{i}\right\}\right)
\end{aligned}
$$



- $\left\langle X_{j}\right\rangle_{q_{j}}$ resembles a "message" sent from node $j$ to $i$
- $\left\{\left\langle X_{j}\right\rangle_{q_{j}}: j \in \mathcal{N}_{i}\right\}$ forms the "mean field" applied to $X_{i}$ from its neighborhood


## Example 1: Generalized MF approximations to Ising models




Cluster marginal of a square block $C_{k}$ :

$$
q\left(X_{C_{k}}\right) \propto \exp \left\{\sum_{i, j \in C_{k}} \theta_{i j} X_{i} X_{j}+\sum_{i \in C_{k}} \theta_{i 0} X_{i}+\sum_{\substack{i \in C_{k}, \in \in \in B_{k}, k \in \in \in B_{k}}} \theta_{i j} X_{i}\left\langle X_{j}\right\rangle_{q\left(X_{C_{k}}\right)}\right\}
$$

Virtually a reparameterized Ising model of small size.

## GMF approximation to Ising models



Singleton marginal error


© Eric Xing @ CMU, 2005-2014
Attractive coupling: positively weighted
Repulsive coupling: negatively weighted

## Example 2: Sigmoid belief network




## Example 3: Factorial HMM




## Automatic Variational Inference


fHMM
(a)


Mean field approx.
(b)


Structured variational approx.

- Currently for each new model we have to
- derive the variational update equations
- write application-specific code to find the solution
- Each can be time consuming and error prone
- Can we build a general-purpose inference engine which automates these procedures?


## 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?

- 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.)


## 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:

- 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 $(|V| \gg D)$

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


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.

## Topic Models

## Why this is Useful?

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

| 0.6 | $\downarrow$ |  |
| :---: | :---: | :---: |
|  | 0.3 | 0.1 |
| MT | Syntax | Learning | | Mixing |
| :--- |
| Proportion |

- 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

Unstructured Collection


Word Simplex

Structured Topic Network


Dimensionality Reduction

## LSI versus Topic Model (probabilistic LSI)



## 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)



## Admixture Models

- Objects are bags of elements

- Mixtures are distributions over element

$\xrightarrow{n}$

- Objects have mixing vector $\theta$
- Represents each mixtures' contributions
- Object is generated as follows:
- Pick a mixture component from $\theta$
- Pick an element from that component

| 0.1 | 0.1 | $\ldots .$. | 0.5 |
| :--- | :--- | :--- | :--- |
| 0.1 | 0.5 | $\ldots .$. | 0.1 |
| 0.5 | 0.1 | $\ldots .$. | 0.1 |

## Topic Models

## Generating a document

> - Draw $\theta$ from the prior
> For each word $n$
> $\quad$ - Draw $z_{n}$ from multinomial $(\theta)$
> - Draw $w_{n} \mid z_{n},\left\{\beta_{1: k}\right\}$ from multinomial $\left(\beta_{z_{n}}\right)$

## 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 0 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 \(\left(\beta_{z_{n}}\right)\)
```

$$
\begin{aligned}
& \theta \sim L N_{K}(\mu, \Sigma) \\
& \gamma \sim N_{K-1}(\mu, \Sigma) \quad \gamma_{K}=0 \\
& \theta_{i}=\exp \left\{\gamma_{i}-\log \left(1+\sum_{i=1}^{K-1} e^{\gamma_{i}}\right)\right\} \\
& C(\gamma)=\log \left(1+\sum_{i=1}^{K-1} e^{\gamma_{i}}\right)
\end{aligned}
$$

## Posterior inference



## Posterior inference results



## Joint likelihood of all variables

$$
p(\beta, \theta, \boldsymbol{z}, \boldsymbol{w})=\prod_{k=1}^{K} p\left(\beta_{k} \mid \eta\right) \prod_{d=1}^{D} p\left(\theta_{d} \mid \alpha\right) \prod_{n=1}^{N} p\left(z_{d n} \mid \theta_{d}\right) p\left(w_{d n} \mid z_{d n}, \beta\right)
$$



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

## Inference and Learning are both 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? $\quad p\left(\theta_{n} \mid D\right)=\frac{p\left(\theta_{n}, D\right)}{p(D)}$

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

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

- 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 $2^{\text {nd }}$-order Taylor approximation (Xing)
- Markov Chain Monte Carlo
- Gibbs sampling (Griffiths et al)


## Mean-field assumption

- True posterior

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

- Break the dependency using the fully factorized distribution

$$
q(\beta, \theta, \boldsymbol{z})=\prod_{k} q\left(\beta_{k}\right) \prod_{d} q\left(\theta_{d}\right) \prod_{n} q\left(z_{d n}\right)
$$

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


## Update each marginals

- Update

$$
q\left(\theta_{d}\right) \propto \exp \left\{\mathbb{E}_{\Pi_{n} q\left(z_{d n}\right)}\left[\log p\left(\theta_{d} \mid \alpha\right)+\sum_{n} \log p\left(z_{d n} \mid \theta_{d}\right)\right]\right\}
$$

- In LDA,

$$
\begin{aligned}
& \text { A, } \quad p\left(\theta_{d} \mid \alpha\right) \propto \exp \left\{\sum_{k=1}^{K}\left(\alpha_{k}-1\right) \log \theta_{d k}\right\}-- \text { Dirichlet } \\
& p\left(z_{d n} \mid \theta_{d}\right)=\exp \left\{\sum_{k=1}^{K} 1\left[z_{d n}=k\right] \log \theta_{d k}\right\}-- \text { Multinomial }
\end{aligned}
$$

- We obtain

$$
q\left(\theta_{d}\right) \propto \exp \left\{\sum_{k=1}^{K}\left(\sum_{n=1}^{N} q\left(z_{d n}=k\right)+\alpha_{k}-1\right) \log \theta_{d k}\right\}
$$

## This is also a Dirichlet---the same as its prior!

## Coordinate ascent algorithm for LDA

1: Initialize variational topics $q\left(\beta_{k}\right), k=1, \ldots, K$.
2: repeat
3: for each document $d \in\{1,2, \ldots, D\}$ do
4: $\quad$ Initialize variational topic assigments $q\left(z_{d n}\right), n=1, \ldots, N$
5: repeat
6:
7:
8: until Change of $q\left(\theta_{d}\right)$ is small enough
end for
Update variational topics $q\left(\beta_{k}\right), k=1, \ldots, K$.
1: until Lower bound $L(q)$ converges

## Choice of $q()$ does matter



## Tangent Approximation

Gamma


## 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?



## 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

$\left.$| 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 |$\quad$|  |
| :--- |$\quad$| -Notable Difference |
| :--- |
| -Examine the Iow dimensional |
| representations below | \right\rvert\,

## 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

## 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


## Summary on VI

- Variational methods in general turn inference into an optimization problem via exponential families and convex duality
- The exact variational principle is intractable to solve; there are two distinct components for approximations:
- Either inner or outer bound to the marginal polytope
- Various approximation to the entropy function
- Mean field: non-convex inner bound and exact form of entropy
- BP: polyhedral outer bound and non-convex Bethe approximation
- Kikuchi and variants: tighter polyhedral outer bounds and better entropy approximations (Yedidia et. al. 2002)

