Readings:
K&F: 11.1, 11.5

Mean Field and Variational Methods

First approximate inference

Graphical Models – 10708
Carlos Guestrin
Carnegie Mellon University

November 1st, 2006

Approximate inference overview

- So far: VE & junction trees
  - exact inference
  - exponential in tree-width
- There are many many many many approximate inference algorithms for PGMs
- We will focus on three representative ones:
  - sampling
  - variational inference
  - loopy belief propagation and generalized belief propagation
- There will be a special recitation by Pradeep Ravikumar on more advanced methods

Approximating the posterior v. approximating the prior

- Prior model represents entire world
  - world is complicated
  - thus prior model can be very complicated
- Posterior: after making observations
  - sometimes can become much more sure about the way things are
  - sometimes can be approximated by a simple model
- First approach to approximate inference: **find *simple* model that is "close" to posterior**
- Fundamental problems:
  - **what is close?**
  - **posterior is intractable result of inference, how can we approximate what we don't have?**

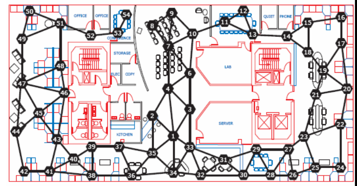

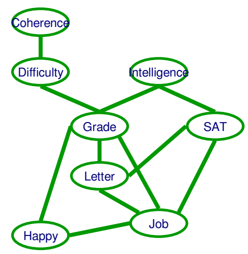

KL divergence: Distance between distributions

- Given two distributions *p* and *q* KL divergence:
- D(p||q) = 0 iff p=q
- Not symmetric – p determines where difference is important
  - p(x)=0 and q(x)≠0
  - p(x)≠0 and q(x)=0

Find simple approximate distribution

- Suppose *p* is intractable posterior
- Want to find simple *q* that approximates *p*
- KL divergence not symmetric
- D(p||q)
  - true distribution p defines support of diff.
  - the "correct" direction
  - will be intractable to compute
- D(q||p)
  - approximate distribution defines support
  - tends to give overconfident results
  - will be tractable

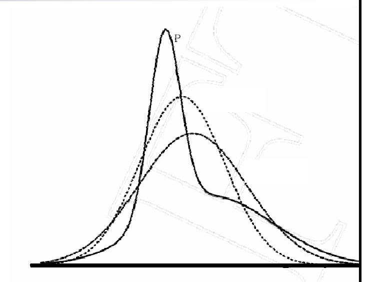

Back to graphical models

- Inference in a graphical model:
  - P(**x**) =
  - want to compute $P(X_i|\mathbf{e})$
  - our *p*:
- What is the simplest *q*?
  - every variable is independent:
  - mean field approximation
  - can compute any prob. very efficiently

D(p||q) for mean field – KL the right way

- p:
- q:
- D(p||q)=

D(q||p) for mean field – KL the reverse direction

- p:
- q:
- D(p||q)=

What you need to know so far

- Goal:
  - Find an efficient distribution that is close to posterior
- Distance:
  - measure distance in terms of KL divergence
- Asymmetry of KL:
  - D(p||q) ≠ D(q||p)
- Computing right KL is intractable, so we use the reverse KL

Reverse KL & The Partition Function

Back to the general case

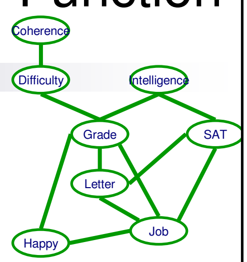

- Consider again the defn. of D(q||p):
  - p is Markov net P_F
- **Theorem**: $\ln Z = F[P_{\mathcal{F}}, Q] + D(Q||P_{\mathcal{F}})$
- where energy functional:

$$F[P_{\mathcal{F}}, Q] = \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi] + H_Q(\mathcal{X})$$

Understanding Reverse KL, Energy Function & The Partition Function

$\ln Z = F[P_{\mathcal{F}}, Q] + D(Q||P_{\mathcal{F}})$ $\quad F[P_{\mathcal{F}}, Q] = \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi] + H_Q(\mathcal{X})$

- Maximizing Energy Functional ⇔ Minimizing Reverse KL
- **Theorem**: Energy Function is lower bound on partition function
  - Maximizing energy functional corresponds to search for tight lower bound on partition function

Structured Variational Approximate Inference

$\ln Z = F[P_{\mathcal{F}}, Q] + D(Q||P_{\mathcal{F}})$

$F[P_{\mathcal{F}}, Q] = \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi] + H_Q(\mathcal{X})$

- Pick a family of distributions *Q* that allow for exact inference
  - e.g., fully factorized (mean field)
- Find Q∈*Q* that maximizes $F[P_{\mathcal{F}}, Q]$
- For mean field

Optimization for mean field

$$\max_Q F[P_{\mathcal{F}}, Q] = \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi] + \sum_j H_{Q_j}(X_j), \quad \forall i, \ \sum_{x_i} Q_j(x_i) = 1$$

- Constrained optimization, solved via Lagrangian multiplier
  - ∃ λ, such that optimization equivalent to:
  - Take derivative, set to zero
- **Theorem**: Q is a stationary point of mean field approximation iff for each *i*:

$$Q_i(x_i) = \frac{1}{Z_i} \exp\left\{ \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi \mid x_i] \right\}$$

Understanding fixed point equation

$$Q_i(x_i) = \frac{1}{Z_i} \exp\left\{ \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi \mid x_i] \right\}$$

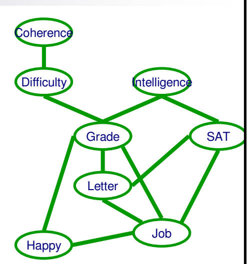

Simplifying fixed point equation

$$Q_i(x_i) = \frac{1}{Z_i} \exp\left\{ \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi \mid x_i] \right\}$$

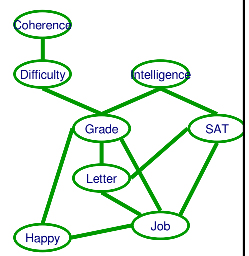

Q_i only needs to consider factors that intersect X_i

- **Theorem**: The fixed point:

$$Q_i(x_i) = \frac{1}{Z_i} \exp\left\{ \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi \mid x_i] \right\}$$

is equivalent to:

$$Q_i(x_i) = \frac{1}{Z_i} \exp\left\{ \sum_{\phi_j : X_i \in \text{Scope}[\phi_j]} E_Q[\ln \phi_j(\mathbf{U}_j, x_i)] \right\}$$

  - where the Scope[ϕ_j] = $\mathbf{U}_j \cup \{X_i\}$

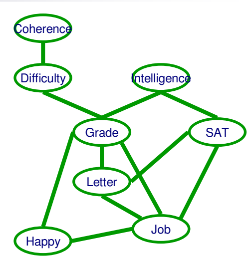

There are many stationary points!

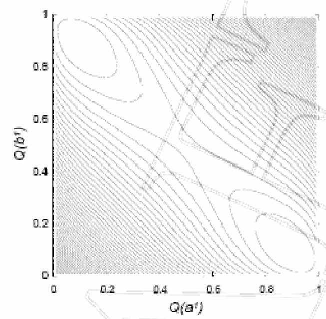

Figure 11.18 An example of a multi-modal mean field energy functional landscape. In this network, $P(a, b) = 0.25 - \epsilon$ if $a \neq b$ and ϵ if $a = b$. The axes correspond to the mean field marginal for A and B and the contours show equi-values of the energy functional.

Very simple approach for finding one stationary point

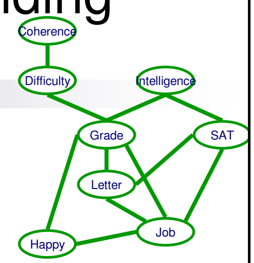

- Initialize Q (e.g., randomly or smartly)
- Set all vars to unprocessed
- Pick unprocessed var X_i
  - update Q_i:

$$Q_i(x_i) = \frac{1}{Z_i} \exp \left\{ \sum_{\phi_j : X_i \in \text{Scope}[\phi_j]} E_Q[\ln \phi_j(\mathbf{U}_j, x_i)] \right\}$$

  - set var i as processed
  - if Q_i changed
    - set neighbors of X_i to unprocessed
- Guaranteed to converge

More general structured approximations

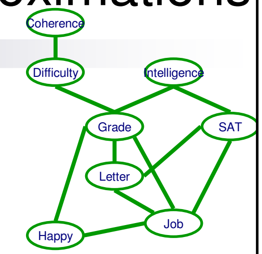

- Mean field very naïve approximation
- Consider more general form for Q
  - assumption: exact inference doable over Q
- **Theorem**: stationary point of energy functional:

$$\psi_j(\mathbf{c}_j) \propto \exp\left\{\sum_{\phi\in\mathcal{F}} E_Q[\ln\phi \mid \mathbf{c}_j] - \sum_{\psi\in\mathcal{Q}\backslash\{\psi_j\}} E_Q[\ln\psi \mid \mathbf{c}_j]\right\}$$

Computing update rule for general case

$$\psi_j(\mathbf{c}_j) \propto \exp\left\{\sum_{\phi\in\mathcal{F}} E_Q[\ln\phi \mid \mathbf{c}_j] - \sum_{\psi\in\mathcal{Q}\backslash\{\psi_j\}} E_Q[\ln\psi \mid \mathbf{c}_j]\right\}$$

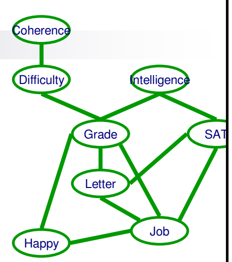

- Consider one ϕ:

Structured Variational update requires inferece

$$\psi_j(\mathbf{c}_j) \propto \exp\left\{ \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi \mid \mathbf{c}_j] - \sum_{\psi \in \mathcal{Q} \setminus \{\psi_j\}} E_Q[\ln \psi \mid \mathbf{c}_j] \right\}$$

- Compute marginals wrt Q of cliques in original graph and cliques in new graph, for all cliques
- What is a good way of computing all these marginals?
- Potential updates:
  - sequential: compute marginals, update ψ_j, recompute marginals
  - parallel: compute marginals, update all ψ's, recompute marginals

What you need to know about variational methods

- Structured Variational method:
  - select a form for approximate distribution
  - minimize reverse KL
- Equivalent to maximizing energy functional
  - searching for a tight lower bound on the partition function
- Many possible models for *Q*:
  - independent (mean field)
  - structured as a Markov net
  - cluster variational
- Several subtleties outlined in the book