Readings:
K&F: 10.2, 10.3

Generalized Belief Propagation

Graphical Models – 10708
Carlos Guestrin
Carnegie Mellon University
November 12th, 2008

More details on Loopy BP

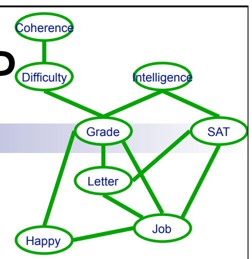

- Numerical problem:
  - messages < 1 get multiplied together as we go around the loops
  - numbers can go to zero
  - normalize messages to one:

$$\delta_{i\to j}(X_j) = \frac{1}{Z_{i\to j}} \sum_{x_i} \phi_i(x_i)\phi_{ij}(x_i, X_j) \prod_{k\in\mathcal{N}(i)-j} \delta_{k\to i}(x_i)$$

  - $Z_{i\to j}$ doesn't depend on X_j, so doesn't change the answer
- Computing node "beliefs" (estimates of probs.):

$$\hat{P}(X_i) = \frac{1}{Z_i}\phi_i(X_i) \prod_{k\in\mathcal{N}(i)} \delta_{k\to i}(X_i)$$

Loopy BP in Factor graphs

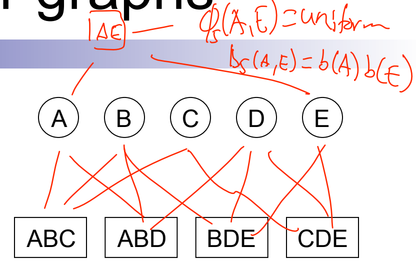

- From node i to factor j:
  - F(i) factors whose scope includes X_i

$$\delta_{i \to j}(X_i) \propto \prod_{k \in \mathcal{F}(i) - j} \delta_{k \to i}(X_i)$$

- From factor j to node i:
  - Scope[ϕ_j] = **Y** ∪ {X_i}

$$\delta_{j \to i}(X_i) \propto \sum_{\mathbf{y}} \phi_j(X_i, \mathbf{y}) \prod_{X_k \in \text{Scope}[\phi_j] - X_i} \delta_{k \to j}(x_k)$$

- Belief:
  - Node: $P(X_i) \approx b_i(X_i) \propto \prod_{\phi_j, X_i \in Scope[\phi_j]} \delta_{\phi_j \to X_i}(X_i)$
  - Factor: $\phi_j(Y)$ $P(Y) \approx b_{\phi_j}(Y) \propto \phi_j(Y) \prod_{X_i \in Y} \delta_{X_i \to \phi_j}(X_i)$

Loopy BP v. Clique trees: Two ends of a spectrum

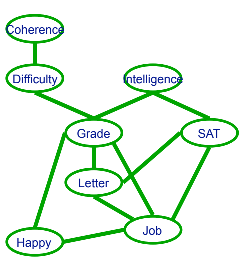

Generalize cluster graph

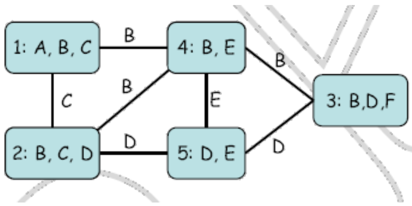

- **Generalized cluster graph**: For set of factors *F*
  - Undirected graph
  - Each node i associated with a cluster $\mathbf{C}_i$
  - *Family preserving*: for each factor $f_j \in F$, $\exists$ node i such that scope$[f_i] \subseteq \mathbf{C}_i$
  - Each edge i – j is associated with a set of variables $\mathbf{S}_{ij} \subseteq \mathbf{C}_i \cap \mathbf{C}_j$

Running intersection property

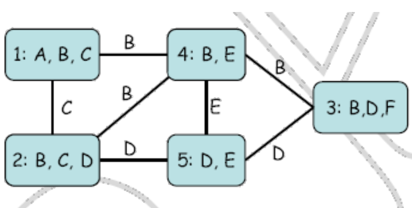

- **(Generalized) Running intersection property (RIP)**
  - Cluster graph satisfies RIP if whenever $X \in \mathbf{C}_i$ and $X \in \mathbf{C}_j$ then $\exists$ *one and only one path* from $\mathbf{C}_i$ to $\mathbf{C}_j$ where $X \in \mathbf{S}_{uv}$ for every edge (u,v) in the path

Examples of cluster graphs

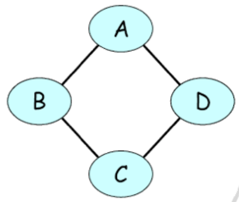

Two cluster graph satisfying RIP with different edge sets

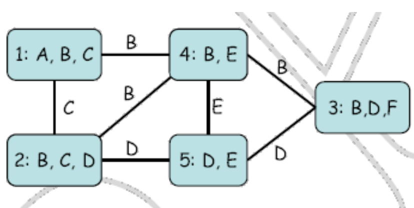

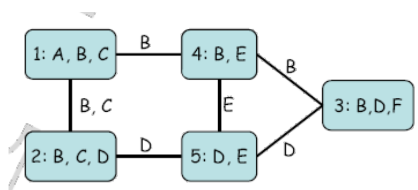

Generalized BP on cluster graphs satisfying RIP

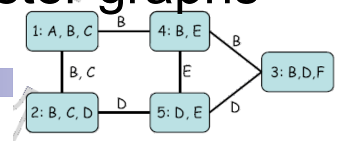

- Initialization:
  - Assign each factor ϕ to a clique $\alpha(\phi)$, Scope[ϕ]$\subseteq \mathbf{C}_{\alpha(\phi)}$
  - Initialize cliques: $\psi_i^0(\mathbf{C}_i) \propto \prod_{\phi:\alpha(\phi)=i} \phi$
  - Initialize messages: $\delta_{j\rightarrow i} = 1$
- While not converged, send messages:

$$\delta_{i\rightarrow j}(\mathbf{S}_{ij}) \propto \sum_{\mathbf{C}_i - \mathbf{S}_{ij}} \psi_i^0(\mathbf{C}_i) \prod_{k\in\mathcal{N}(i)-j} \delta_{k\rightarrow i}(\mathbf{S}_{ik})$$

- Belief:

Cluster graph for Loopy BP

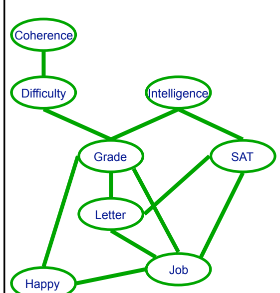

What if the cluster graph doesn't satisfy RIP

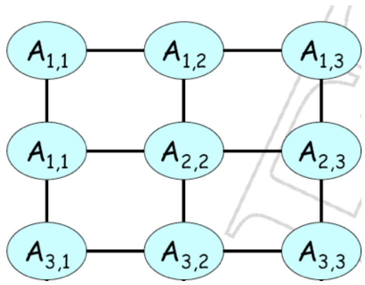

Region graphs to the rescue

- Can address generalized cluster graphs that don't satisfy RIP using *region graphs*:
  - Book: 10.3

- Example in your homework! ☺

Revisiting Mean-Fields

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

- Choice of Q:
- Optimization problem:

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

Announcements

- Recitation tomorrow
- HW5 out soon
- Will not cover relational models this semester
  - Instead, recommend Pedro Domingos' tutorial on Markov Logic
    - Markov logic is one example of a relational probabilistic model
    - November 14th from 1:00 pm to 3:30 pm in Wean 4623

Interpretation of energy functional

- Energy functional: $F[P_{\mathcal{F}}, Q] = \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi] + H_Q(\mathcal{X})$
- Exact if P=Q: $\ln Z = F[P_{\mathcal{F}}, Q] + D(Q||P_{\mathcal{F}})$
- View problem as an approximation of entropy term:

Entropy of a tree distribution

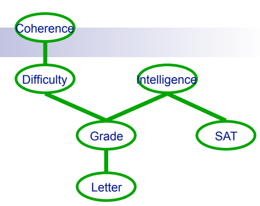

- Entropy term:
- Joint distribution:
- Decomposing entropy term:
- More generally: $H_P(\mathbf{X}) = \sum_{(i,j) \in E} H(X_i, X_j) - \sum_i (d_i - 1) H(X_i)$
  - d_i number neighbors of X_i

Loopy BP & Bethe approximation

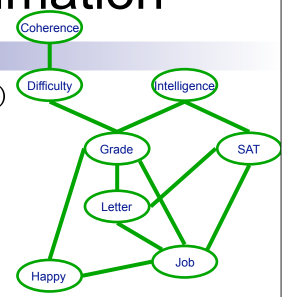

- Energy functional: $F[P_{\mathcal{F}}, Q] = \sum_{\phi \in \mathcal{F}} E_Q[\ln \phi] + H_Q(\mathcal{X})$
- **Bethe approximation of Free Energy**:
  - use entropy for trees, but loopy graphs:

$$\tilde{F}[P_{\mathcal{F}},Q] = \sum_{(i,j)\in E} E_{b_{ij}}[\ln \phi_{ij}] + \sum_{(i,j)\in E} H_{b_{ij}}(X_i,X_j) - \sum_i (d_i - 1) H_{b_i}(X_i)$$

- **Theorem**: If Loopy BP converges, resulting b_{ij} & b_i are stationary point (usually local maxima) of Bethe Free energy!

GBP & Kikuchi approximation

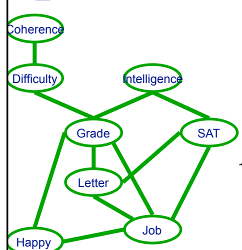

- Exact Free energy: Junction Tree

$$F[P_{\mathcal{F}}, Q] = \sum_{(i,j)\in E} E_{\pi_{ij}}[\ln \phi_{ij}] + \sum_i H_{\pi_{C_i}}(C_i) - \sum_{(i,j)\in \mathcal{T}} H_{\pi_{S_{ij}}}(S_{ij})$$

- Bethe Free energy:

$$\tilde{F}[P_{\mathcal{F}},Q] = \sum_{(i,j)\in E} E_{b_{ij}}[\ln \phi_{ij}] + \sum_{(i,j)\in E} H_{b_{ij}}(X_i,X_j) - \sum_i (d_i - 1) H_{b_i}(X_i)$$

- Kikuchi approximation: Generalized cluster graph
  - spectrum from Bethe to exact
- **Theorem**: If GBP converges, resulting b_{C_i} are stationary point (usually local maxima) of Kikuchi Free energy!

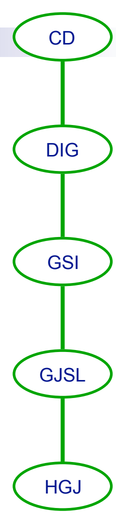

What you need to know about GBP

- Spectrum between Loopy BP & Junction Trees:
  - More computation, but typically better answers
- If satisfies RIP, equations are very simple
- General setting, slightly trickier equations, but not hard
- Relates to variational methods: Corresponds to local optima of approximate version of energy functional

Readings:
K&F: 10.2, 10.3

Parameter learning in Markov nets

Graphical Models – 10708
Carlos Guestrin
Carnegie Mellon University
November 12th, 2008

Learning Parameters of a BN

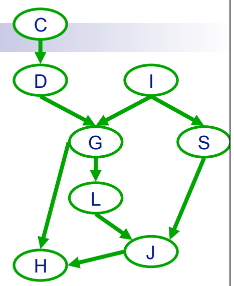

- Log likelihood decomposes:

$$\ell(\mathcal{D}:\theta) = \log P(\mathcal{D} \mid \theta) = m \sum_i \sum_{x_i, \mathbf{Pa}_{x_i}} \widehat{P}(x_i, \mathbf{Pa}_{x_i}) \log P(x_i \mid \mathbf{Pa}_{x_i})$$

- Learn each CPT independently
- Use counts

$$\widehat{P}(\mathbf{u}) = \frac{\text{Count}(\mathbf{U}=\mathbf{u})}{m}$$

Log Likelihood for MN

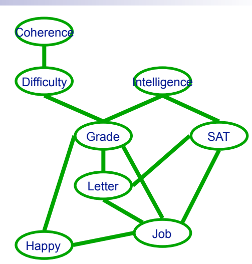

- Log likelihood of the data:

Log Likelihood doesn't decompose for MNs

$$\hat{P}(\mathbf{u}) = \frac{\text{Count}(\mathbf{U} = \mathbf{u})}{m}$$

- Log likelihood:

$$\ell(\mathcal{D} : \theta) = \log P(\mathcal{D} \mid \theta, \mathcal{G}) = m \sum_i \sum_{\mathbf{c}_i} \hat{P}(\mathbf{c}_i) \log \psi_i(\mathbf{c}_i) - m \log Z$$

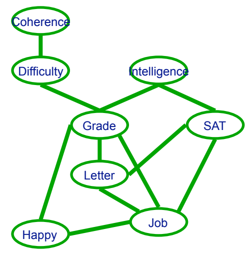

- A convex problem
  - Can find global optimum!!
- Term log Z doesn't decompose!!

Derivative of Log Likelihood for MNs

$$\hat{P}(\mathbf{u}) = \frac{\text{Count}(\mathbf{U} = \mathbf{u})}{m}$$

$$\ell(\mathcal{D} : \theta) = \log P(\mathcal{D} \mid \theta, \mathcal{G}) = m \sum_i \sum_{\mathbf{c}_i} \hat{P}(\mathbf{c}_i) \log \psi_i(\mathbf{c}_i) - m \log Z$$

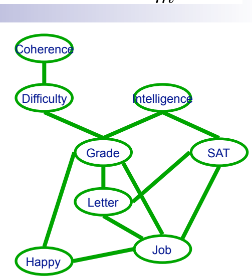

Derivative of Log Likelihood for MNs 2

$$\hat{P}(\mathbf{u}) = \frac{\text{Count}(\mathbf{U} = \mathbf{u})}{m}$$

$$\ell(\mathcal{D} : \theta) = \log P(\mathcal{D} \mid \theta, \mathcal{G}) = m \sum_i \sum_{\mathbf{c}_i} \hat{P}(\mathbf{c}_i) \log \psi_i(\mathbf{c}_i) - m \log Z$$

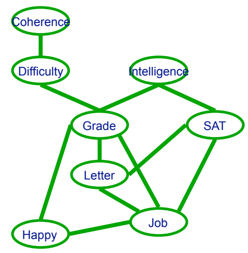

Derivative of Log Likelihood for MNs

$$\hat{P}(\mathbf{u}) = \frac{\text{Count}(\mathbf{U} = \mathbf{u})}{m}$$

$$\ell(\mathcal{D} : \theta) = \log P(\mathcal{D} \mid \theta, \mathcal{G}) = m \sum_i \sum_{\mathbf{c}_i} \hat{P}(\mathbf{c}_i) \log \psi_i(\mathbf{c}_i) - m \log Z$$

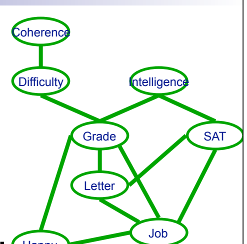

- Derivative:

$$\frac{\partial \ell}{\partial \psi_i(\mathbf{c}_i)} = \frac{m\hat{P}(\mathbf{c}_i)}{\psi_i(\mathbf{c}_i)} - \frac{mP^{\psi}_{\mathcal{F}}(\mathbf{c}_i)}{\psi_i(\mathbf{c}_i)}$$

- Computing derivative requires inference:
- Can optimize using gradient ascent
  - Common approach
  - Conjugate gradient, Newton's method,…
- Let's also look at a simpler solution

Iterative Proportional Fitting (IPF)

$$\hat{P}(\mathbf{u}) = \frac{\text{Count}(\mathbf{U} = \mathbf{u})}{m}$$

$$\frac{\partial \ell}{\partial \psi_i(\mathbf{c}_i)} = \frac{m\hat{P}(\mathbf{c}_i)}{\psi_i(\mathbf{c}_i)} - \frac{mP^{\psi}_{\mathcal{F}}(\mathbf{c}_i)}{\psi_i(\mathbf{c}_i)}$$

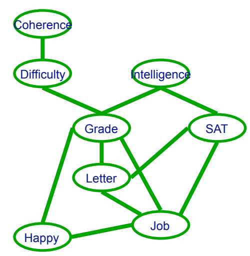

- Setting derivative to zero:
- Fixed point equation:
- Iterate and converge to optimal parameters
  - Each iteration, must compute:

What you need to know about learning MN parameters?

- BN parameter learning easy
- MN parameter learning doesn't decompose!
- Learning requires inference!
- Apply gradient ascent or IPF iterations to obtain optimal parameters