


Stat 521A
Lecture 23

Outline

- Basic issues (19.1)
- Gradient ascent for DGMs (19.2.1)
- EM for DGMs (19.2.2)
- Variational EM (19.2.4)
- MCMC for param inf in DGMs (19.3.2)
- Variational Bayes (19.3.3)

MCAR


- Let X_i be the true value of variable I , and O_i in $\{0,1\}$ be whether it is observed or not. $Y_i(O_i) = X_i$ or ?.
- Defn 19.1.6. Missing completely at random (MCAR) means $X \perp O$.
- Given MCAR, we can safely ignore the missing variables (for which $O_x=1$), since they tell us nothing about θ



$$p(\theta, \psi | Y_1, Y_2) = p(\theta | X_1) p(\psi | O_1, O_2)$$

Missing at random

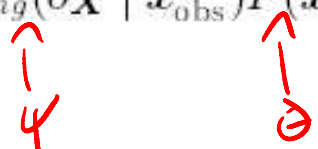
- Defn 19.1.8. Let H be hidden vars, V be visible vars, and O be observation status. Missing at random means $O \perp H \mid V$.



- Intuitively, although O may depend on some of the variables X_v , since we observe X_v , we do not learn anything new about X_h .

Benefits of MAR

- Thm 19.1.9. Given MAR, and a factored prior, $p(\theta, \psi | D) = p(\theta | X^y) p(\psi | X^y, O)$
- Pf.


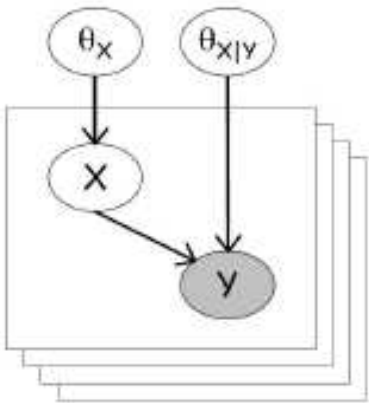
$$\begin{aligned} P_{missing}(y) &= \sum_{x_{hidden}^y} [P(x_{obs}^y, x_{hidden}^y) P_{missing}(O_X | x_{hidden}^y, x_{obs}^y)] \\ &= \sum_{x_{hidden}^y} [P(x_{obs}^y, x_{hidden}^y) P_{missing}(O_X | x_{obs}^y)] \\ &= P_{missing}(O_X | x_{obs}^y) \sum_{x_{hidden}^y} P(x_{obs}^y, x_{hidden}^y) \\ &= P_{missing}(O_X | x_{obs}^y) P(x_{obs}^y). \end{aligned}$$


Counter examples to MAR

- Collaborative filtering: people are more likely to rate movies they strongly like or dislike.
- Medicine: if a patient does not have a check mark in the “had X-ray” field, they probably don’t have any bone problems. However, if we explicitly write the “primary complaint” as the cause of which tests are performed, MAR is restored (since we observe why $O(Xray)=0$).
- Henceforth we will assume MAR.

Multimodality

- For fully observed DGMs, likelihood is convex (assuming each CPD is convex), and hence has a single global maximum.
- When we have missing data, the likelihood is a mixture of up to K^n modes, corresponding to every possible completion pattern



Proposition 19.1.10: *Assuming i.i.d. data, the likelihood can be written as*

$$L(\theta : \mathcal{D}) = \prod_m P(o[m] | \theta) = \prod_m \sum_{\mathbf{h}[m]} P(o[m], \mathbf{h}[m] | \theta).$$

Identifiability


- Sometimes we cannot uniquely identify the parameters, even given infinite data
- Eg The experimenter either tosses coin 1 or coin 2, but we don't know which. The model is

θ_H
↓
H

$$L(\theta : \mathcal{D}) = P(X = Heads)^{M[Heads]}(1 - P(X = Heads))^{M[Tails]},$$

where

$$P(X = Heads) = \theta_H \theta_{X|H=1} + (1 - \theta_H) \theta_{X|H=2}.$$



- We have eg. $p(D|\theta_H=0.5, \theta_1=0.5, \theta_2=0.5) = p(D|\theta_H=0.5, \theta_1=0.8, \theta_2=0.2)$. The problem is underconstrained.

Identifiability

- Defn 19.1.13. A parameter θ is identifiable if there is no $\theta' \neq \theta$ st $p(X|\theta)=p(X|\theta')$. A model is identifiable if all θ are identifiable.
- A mixture model cannot be identifiable since we can always arbitrarily permute the hidden labels, and the corresponding parameters.
- Hence we should not ask things like “what is the prob. X_i belongs to cluster k ” but rather “what is the prob X_i and X_i belong to the same cluster”.



Gradient descent for DGMs

- We can find a local maximum using gradient based methods.
- Consider tabular CPDs.
- Thm 19.2.1.

$$\frac{\partial}{\partial \theta_{ijk}} p(\mathbf{e}) = \frac{p(x_i = k, \mathbf{x}_{\pi_i} = j, \mathbf{e})}{\theta_{ijk}}$$

- Pf.

$$\begin{aligned} \frac{\partial}{\partial \theta_{ijk}} \prod_{i'} \theta_{i', \mathbf{x}_{i'}, x_{i'}} &= \prod_{i' \neq i} \theta_{i', \mathbf{x}_{i'}, x_{i'}} I(\mathbf{x}_i = j, x_i = k) \\ &= \frac{p(\mathbf{e})}{\theta_{ijk}} I(\mathbf{x}_i = j, x_i = k) \end{aligned}$$

Gradient descent for DGMs

- Pf contd

$$\begin{aligned}\frac{\partial}{\partial P(x | \mathbf{u})} P(e) &= \sum_{\xi: \langle \mathbf{E} \rangle = e} \frac{\partial}{\partial P(x | \mathbf{u})} P(\xi) \\ &= \sum_{\xi: \langle \mathbf{E} \rangle = e, \langle X, \text{Pa}_X \rangle = \langle x, \mathbf{u} \rangle} \frac{1}{P(x | \mathbf{u})} P(\xi) \\ &= \frac{1}{P(x | \mathbf{u})} P(x, \mathbf{u}, e).\end{aligned}$$

- Thm 19.2.2.

$$\frac{\partial \ell(\theta : \mathcal{D})}{\partial P(x | \mathbf{u})} = \frac{1}{P(x | \mathbf{u})} \sum_{m=1}^M P(x, \mathbf{u} | o[m], \theta).$$

- Chain rule for non-tabular CPDs.

$$\frac{\partial \ell(\theta : \mathcal{D})}{\partial \theta} = \sum_{x, \mathbf{u}} \frac{\partial \ell(P_\theta : \mathcal{D})}{\partial P(x | \mathbf{u})} \frac{\partial P(x | \mathbf{u})}{\partial \theta},$$

Gradient algorithms

- Gradient requires inference to compute family marginals.
- Need to enforce positivity and sum-to-one constraints (for discrete) eg reparameterize to unconstrained form

$$P(x | u) = \frac{e^{\lambda_{x|u}}}{\sum_{x' \in \text{Val}(X)} e^{\lambda_{x'|u}}}$$

- Need to enforce positive definite – optimize wrt the cholesky factors.
- Have to specify step-size and search direction (use black-box algorithm).
- EM is much easier...

EM for DGMs

- Key intuition: if we knew the values of H , we could compute the MLEs/MAP estimates for θ easily. So we infer $H|\theta$ and then estimate $\theta|H$. For the latter, we just need the expected sufficient statistics. For tabular CPDs, this is just a table of expected counts

- E step

$$\bar{M}_{\theta^t}[x, u] = \sum_m P(x, u | o[m], \theta^t).$$

- M step

-

$$\theta_{x|u}^{t+1} = \frac{\bar{M}_{\theta^t}[x, u]}{\bar{M}_{\theta^t}[u]}$$

Pseudocode

```
1  /
2  for each  $t = 0, 1 \dots$ , until convergence
3      // E-step
4       $\{\bar{M}_t[x_i, \mathbf{u}_i]\} \leftarrow \text{Compute-Expected-Sufficient-Statistics}(\mathcal{G}, \theta^t, \mathcal{D})$ 
5      // M-step
6      for each  $i = 1, \dots, n$ 
7          for each  $x_i, \mathbf{u}_i \in \text{Val}(X_i, \text{Pa}_{X_i}^{\mathcal{G}})$ 
8               $\theta_{x_i|\mathbf{u}_i}^{t+1} \leftarrow \frac{\bar{M}_t[x_i, \mathbf{u}_i]}{\bar{M}_t[\mathbf{u}_i]}$ 
9  return  $\theta^t$ 
```

```
1  /
2      // Initialize data structures
3  for each  $i = 1, \dots, n$ 
4      for each  $x_i, \mathbf{u}_i \in \text{Val}(X_i, \text{Pa}_{X_i}^{\mathcal{G}})$ 
5           $\bar{M}[x_i, \mathbf{u}_i] \leftarrow 0$ 
6      // Collect probabilities from all instances
7  for each  $m = 1 \dots M$ 
8      Run inference on  $\langle \mathcal{G}, \theta \rangle$  using evidence  $\mathbf{o}[m]$ 
9      for each  $i = 1, \dots, n$ 
10         for each  $x_i, \mathbf{u}_i \in \text{Val}(X_i, \text{Pa}_{X_i}^{\mathcal{G}})$ 
11              $\bar{M}[x_i, \mathbf{u}_i] \leftarrow \bar{M}[x_i, \mathbf{u}_i] + P(x_i, \mathbf{u}_i | \mathbf{o}[m])$ 
12 return  $\{\bar{M}[x_i, \mathbf{u}_i] : \forall i = 1, \dots, n, \forall x_i, \mathbf{u}_i \in \text{Val}(X_i, \text{Pa}_{X_i}^{\mathcal{G}})\}$ 
```

ECDLL

- Define expected complete data log likelihood, wrt Q distribution over $\mathcal{H}|\mathcal{D}$

$$E_Q[\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)] = \sum_{\mathcal{H}} Q(\mathcal{H}) \ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)$$

- For tabular CPDs, we have

$$\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle) = \sum_{i=1}^n \sum_{(x_i, \mathbf{u}_i) \in \text{Val}(X_i, \text{Pa}_{X_i})} M_{\langle \mathcal{D}, \mathcal{H} \rangle}[x_i, \mathbf{u}_i] \log \theta_{x_i | \mathbf{u}_i}.$$

$$E_Q[\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)] = \sum_{i=1}^n \sum_{(x_i, \mathbf{u}_i) \in \text{Val}(X_i, \text{Pa}_{X_i})} E_Q[M_{\langle \mathcal{D}, \mathcal{H} \rangle}[x_i, \mathbf{u}_i]] \log \theta_{x_i | \mathbf{u}_i}.$$

$$E_Q[\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)] = \sum_{i=1}^n \sum_{(x_i, \mathbf{u}_i) \in \text{Val}(X_i, \text{Pa}_{X_i})} \bar{M}_Q[x_i, \mathbf{u}_i] \log \theta_{x_i | \mathbf{u}_i}.$$

ECDLL for exp fam

- The key to making EM simple for expfam models is that the log-likelihood is linear in the sufficient statistics

$$P(\xi | \theta) = \frac{1}{Z(\theta)} A(\xi) \exp \{ \langle \mathbf{t}(\theta), \tau(\xi) \rangle \}, \quad \tau(\langle \mathcal{D}, \mathcal{H} \rangle) = \sum_m \tau(\mathbf{o}[m], \mathbf{h}[m]).$$

$$\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle) = \langle \mathbf{t}(\theta), \tau(\langle \mathcal{D}, \mathcal{H} \rangle) \rangle + \sum_m A(\mathbf{o}[m], \mathbf{h}[m]) - \log Z(\theta).$$

$$\mathbf{E}_Q[\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)] = \langle \mathbf{t}(\theta), \mathbf{E}_Q[\tau(\langle \mathcal{D}, \mathcal{H} \rangle)] \rangle + \sum_m \mathbf{E}_Q[A(\mathbf{o}[m], \mathbf{h}[m])] - M \log Z(\theta).$$

const

Choosing Q (for E step)

- Define

$$\Phi_{\mathcal{D}}[\theta, Q] = \mathbf{E}_Q[\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)] + \mathbf{H}_Q(\mathcal{H}).$$

- Thm 19.2.5.

Corollary 19.2.5: For any Q,

$$\begin{aligned} \ell(\theta : \mathcal{D}) &= \Phi_{\mathcal{D}}[\theta, Q] + \mathbf{D}(Q(\mathcal{H}) \| P(\mathcal{H} | \mathcal{D}, \theta)) \\ &= \mathbf{E}_Q[\ell(\theta : \langle \mathcal{D}, \mathcal{H} \rangle)] + \mathbf{H}_Q(\mathcal{H}) + \mathbf{D}(Q(\mathcal{H}) \| P(\mathcal{H} | \mathcal{D}, \theta)). \end{aligned}$$

- From (2), ECDLL is lower bound on LL.
- From (1), if $Q=p(\mathcal{H}|\mathcal{D},\theta)$, then bound is tight.
- EM alternates between optimizing Q and optimizing θ . Can do partial updates.

Convergence

- Thm 19.2.6. If we do exact EM (so $Q=p(H|D,\theta)$), then the LL never decreases

Theorem 19.2.6: *During iterations of the EM procedure of Algorithm 19.2, we have that*

$$\ell(\theta^{t+1} : \mathcal{D}) - \ell(\theta^t : \mathcal{D}) \geq \mathbf{E}_{P(\mathcal{H}|\mathcal{D},\theta^t)}[\ell(\theta^{t+1} : \mathcal{D}, \mathcal{H})] - \mathbf{E}_{P(\mathcal{H}|\mathcal{D},\theta^t)}[\ell(\theta^t : \mathcal{D}, \mathcal{H})].$$

As a consequence, we obtain that:

$$\ell(\theta^t : \mathcal{D}) \leq \ell(\theta^{t+1} : \mathcal{D}).$$

PROOF We begin with the first statement. Using Corollary 19.2.5, with the distribution $Q^t(\mathcal{H}) = P(\mathcal{H} | \mathcal{D}, \theta^t)$ we have that

$$\begin{aligned} \ell(\theta^{t+1} : \mathcal{D}) &= \mathbf{E}_{Q^t}[\ell(\theta^{t+1} : \langle \mathcal{D}, \mathcal{H} \rangle)] + H_{Q^t}(\mathcal{H}) + \mathbf{D}(Q^t(\mathcal{H}) \| P(\mathcal{H} | \mathcal{D}, \theta^{t+1})) \\ \ell(\theta^t : \mathcal{D}) &= \mathbf{E}_{Q^t}[\ell(\theta^t : \langle \mathcal{D}, \mathcal{H} \rangle)] + H_{Q^t}(\mathcal{H}) + \mathbf{D}(Q^t(\mathcal{H}) \| P(\mathcal{H} | \mathcal{D}, \theta^t)) \\ &= \mathbf{E}_{Q^t}[\ell(\theta^t : \langle \mathcal{D}, \mathcal{H} \rangle)] + H_{Q^t}(\mathcal{H}) \end{aligned}$$

The last step is justified by our choice of $Q^t(\mathcal{H}) = P(\mathcal{H} | \mathcal{D}, \theta^t)$. Subtracting these two terms, we have that

$$\ell(\theta^{t+1} : \mathcal{D}) - \ell(\theta^t : \mathcal{D}) = \mathbf{E}_{Q^t}[\ell(\theta^{t+1} : \mathcal{D}, \mathcal{H})] - \mathbf{E}_{Q^t}[\ell(\theta^t : \mathcal{D}, \mathcal{H})] + \mathbf{D}(Q^t(\mathcal{H}) \| P(\mathcal{H} | \mathcal{D}, \theta^{t+1}))$$

As the last term is non-negative, we get the desired inequality.

Convergence cont'd

Theorem 19.2.6: *During iterations of the EM procedure of Algorithm 19.2, we have that*

$$\ell(\boldsymbol{\theta}^{t+1} : \mathcal{D}) - \ell(\boldsymbol{\theta}^t : \mathcal{D}) \geq \mathbf{E}_{P(\mathcal{H}|\mathcal{D},\boldsymbol{\theta}^t)}[\ell(\boldsymbol{\theta}^{t+1} : \mathcal{D}, \mathcal{H})] - \mathbf{E}_{P(\mathcal{H}|\mathcal{D},\boldsymbol{\theta}^t)}[\ell(\boldsymbol{\theta}^t : \mathcal{D}, \mathcal{H})].$$

As a consequence, we obtain that:


$$\ell(\boldsymbol{\theta}^t : \mathcal{D}) \leq \ell(\boldsymbol{\theta}^{t+1} : \mathcal{D}).$$

To prove the second statement of the theorem, we note that $\boldsymbol{\theta}^{t+1}$ is the value of $\boldsymbol{\theta}$ that maximizes $\mathbf{E}_{P(\mathcal{H}|\mathcal{D},\boldsymbol{\theta}^t)}[\ell(\boldsymbol{\theta} : \mathcal{D}, \mathcal{H})]$. Hence the value obtained for this expression for $\boldsymbol{\theta}^{t+1}$ is at least as large as the value obtained for any other set of parameters, including $\boldsymbol{\theta}^t$. We conclude that the right-hand side of the inequality is non-negative, which implies the first statement. ■

Theorem 19.2.7: *Suppose that $\boldsymbol{\theta}^t$ is such that $\boldsymbol{\theta}^{t+1} = \boldsymbol{\theta}^t$ during EM, and $\boldsymbol{\theta}^t$ is also an interior point of the allowed parameter space. Then $\boldsymbol{\theta}^t$ is a stationary point of the log-likelihood function.*

Rate of convergence

- Initially fast, then very slow; can switch over to conjugate gradient near optimum



- EM has linear convergence rate


$$\epsilon_t = l^* - l_t$$

Although we do not go through the proof, one can show that EM has *linear convergence rate*. This means that for each domain there exists a t_0 and $\alpha < 1$ such that for all $t \geq t_0$


$$\epsilon_{t+1} \leq \alpha \epsilon_t.$$

Local maxima

- Maxima can differ a lot in quality.
- Can do multiple restart, killing off some runs early if they look bad (as in beam search).



(a)




(b)

Assessing convergence


- Can check whether parameters stop changing or LL stops changing. Can be quite different.

- Recall
$$\frac{\partial \ell(\theta : \mathcal{D})}{\partial P(x | u)} = \frac{1}{P(x | u)} \sum_{m=1}^M P(x, u | o[m], \theta).$$

- If $p(x, u | o[m])$ small, gradient is small, else $O(M)$
- Hence effects of param on LL can be small or large.



(a)





(b)

Accelerating convergence


- Hard assignment EM (eg Kmeans). E step is searching over discrete assignments; this tends to converge faster (but to a worse solution).
- Hybrid EM/CG
- Over relaxation: step size > 1 .
- Stochastic EM: since $Q(H) = \prod_m Q(h_m|o_m)$, we can do inference on only a subset of the datacases (mini-batch) and then do an M step
- (Monte Carlo EM: sampling in the E step)

Example: fitting planes to 3d point clouds



$$P(X_m | C_m = k : \theta_k) \propto \mathcal{N}(d(x, p_k); 0, \sigma^2).$$

$$d(x, p_k) = |\alpha_k x - \beta_k|.$$



Variational EM

- Restrict Q distribution in E step to a tractable family, rather than $p(H|D, \theta)$

$$\max_{\theta} \max_{Q \in \mathcal{Q}} \Phi_{\mathcal{D}}[\theta, Q]$$

- Eg do mean-field in the E step, then regular M step
- Maximizes a lower bound on the LL

$$\ell(\theta : \mathcal{D}) = \max_Q \Phi_{\mathcal{D}}[\theta, Q] \geq \max_{Q \in \mathcal{Q}} \Phi_{\mathcal{D}}[\theta, Q].$$

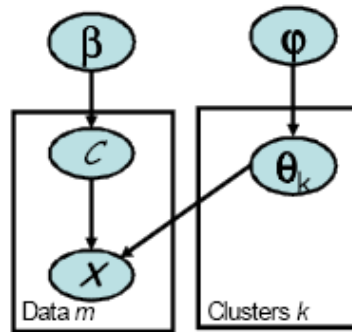


MCMC

- Can compute $p(\theta, H|D)$ using standard algorithms
- Parameter collapsed particles: sample θ , compute $p(H|D)$ analytically
- Data completion collapsed particles: sample H , compute $p(\theta|H, D)$ analytically

Marginalizing out H

- Bayesian Mixture model

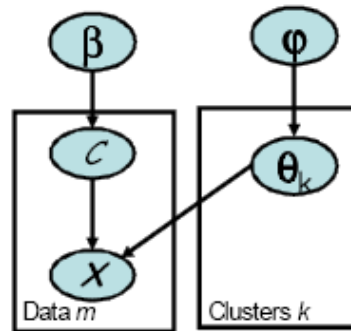


$$P(\mathcal{D} | \theta_1, \dots, \theta_K, \lambda) = \prod_{m=1}^M \left(\sum_{k=1}^K P(C[m] = c^k | \lambda) P(x[m] | C[m] = c^k, \theta_k) \right)$$

$$P(\theta_k | \lambda, \theta_{-k}, \mathcal{D}) \propto P(\mathcal{D} | \theta_1, \dots, \theta_K, \lambda).$$

Have to use MH

Marginalizing out θ



Collapsed Gibbs sampling
Cf DP mixtures

$$P(\lambda | c) = \text{Dirichlet}(\alpha_0/K + |I_1(c)|, \dots, \alpha_0/K + |I_K(c)|).$$

$$P(\theta_k | c, \mathcal{D}, \phi) = Q(\theta_k | \mathcal{D}_{I_k(c)}, \phi) \propto P(\theta_k | \phi) \prod_{m \in I_k(c)} P(x[m] | \theta_k),$$

$$P(C[m'] = k | c_{-m'}, \mathcal{D}, \phi) \propto P(C[m'] = k | \lambda, c_{-m'}) P(x[m'] | C[m'] = k, x[I_k(c_{-m'})], \phi).$$

$$P(C[m'] = k | c_{-m'}, \mathcal{D}, \phi) \propto (|I_k(c_{-m'})| + \alpha_0/K) Q(X | \mathcal{D}_{I_k(c_{-m'})}, \phi).$$

Variational Bayes

- Min $KL(Q|P)$ where we assume

$$Q(\theta, \mathcal{H}) = Q(\theta)Q(\mathcal{H}).$$

- Thm 19.3.6. If we have global param independence and $Q(\theta, \mathcal{H}) = Q(\theta) Q(\mathcal{H})$ then

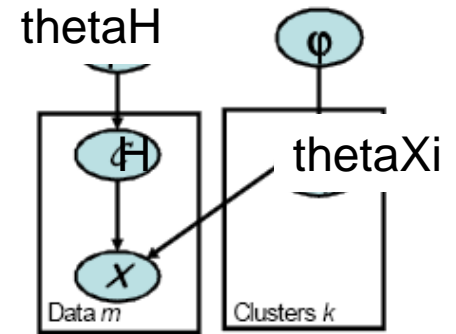
$$Q(\theta, \mathcal{H}) = \prod_i Q(\theta_{X_i|U_i}) \prod_m Q(h[m]).$$

- Hence we can optimize each $Q(h[m])$ separately – just like inference per case in the E step – and then optimize each $Q(\theta_i)$ separately – just like optimizing each family in the M step
- E step: we do inference with expected params
- M step: we fit a distribution

VB for MixBernoulli

$$Q = Q(\theta_H) \left[\prod_i Q(\theta_{X_i|H}) \right] \left[\prod_m Q(H[m]) \right].$$

Beta Beta Bernoulli



$$Q(\theta_H) \propto \exp \left\{ \ln P(\theta_H) + \sum_m E_{Q(H[m])} [\ln P(H[m] | \theta_H)] \right\}$$

$$Q(\theta_{X_i|H}) \propto \exp \left\{ \ln P(\theta_{X_i|H}) + \sum_m E_{Q(H[m])} [\ln P(x_i[m] | H[m], \theta_{X_i|H})] \right\}$$

$$Q(H[m]) \propto \exp \left\{ E_{Q(\theta_H)} [\ln P(H[m] | \theta_H)] + \sum_i E_{Q(\theta_{X_i|H})} [\ln P(x_i[m] | H[m], \theta_{X_i|H})] \right\}.$$

$$\ln P(\theta_H = \langle \theta_{h^0}, \theta_{h^1} \rangle) = \ln c + (\alpha_{h^0} - 1) \ln \theta_{h^0} + (\alpha_{h^1} - 1) \ln \theta_{h^1}.$$

$$E_{Q(H[m])} [\ln P(H[m] | \theta_H = \langle \theta_{h^0}, \theta_{h^1} \rangle)] = Q(H[m] = h^0) \ln \theta_{h^0} + Q(H[m] = h^1) \ln \theta_{h^1}.$$

$$Q(\theta_H = \langle \theta_{h^0}, \theta_{h^1} \rangle) \propto \exp \left\{ \left(\alpha_{h^0} + \sum_m Q(H[m] = h^0) - 1 \right) \ln \theta_{h^0} + \left(\alpha_{h^1} + \sum_m Q(H[m] = h^1) - 1 \right) \ln \theta_{h^1} \right\}$$

$$= \theta_{h^0}^{\alpha_{h^0} + \sum_m Q(H[m]=h^0) - 1} \theta_{h^1}^{\alpha_{h^1} + \sum_m Q(H[m]=h^1) - 1}.$$

$$\alpha'_{h^0} = \alpha_{h^0} + \sum_m Q(H[m] = h^0)$$

$$\alpha'_{h^1} = \alpha_{h^1} + \sum_m Q(H[m] = h^1).$$

VB update for $H[m]$

Regular E step

$$P(H[m] | x_1[m], \dots, x_n[m]) \propto P(H[m] | \theta_H) \prod_i P(x_i[m] | H[m], \theta_{X_i|H}).$$

VB version

$$\mathbf{E}_{Q(\theta_{X_i|H})} [\ln P(x_i | H[m], \theta_{X_i|H})] = \int_0^1 Q(\theta_{x_i|H[m]}) \ln \theta_{x_i|H[m]} d\theta_{x_i|H[m]}.$$

$$\mathbf{E}_{Q(\theta_{X_i|H})} [\ln P(x_i | H[m], \theta_{X_i|H})] = \varphi(\alpha'_{x_i|h}) - \varphi\left(\sum_{x'_i} \alpha'_{x'_i|h}\right)$$

where α' are the hyperparameters of the posterior approximation in $Q(\theta_{X_i|H})$ and $\varphi(z) = (\ln \Gamma(z))' = \frac{\Gamma'(z)}{\Gamma(z)}$ is the digamma function, which is equal to $\ln(z)$ plus a polynomial function of $\frac{1}{z}$. And so, for $z \gg 1$, $\varphi(z) \approx \ln(z)$. Using this approximation, we see that

$$\mathbf{E}_{Q(\theta_{X_i|H})} [\ln P(x_i | H[m], \theta_{X_i|H})] \approx \ln \frac{\alpha'_{x_i|h}}{\sum_{x'_i} \alpha'_{x'_i|h}},$$



Variational methods

- From Lecture 10:
- Minimize

$$D(Q||P) = \ln Z - F(\tilde{P}, Q)$$
$$F(\tilde{P}, Q) \stackrel{\text{def}}{=} H_Q(x) + \sum_i E_{C_i \sim Q} \ln \psi_i(C_i)$$

- This always increases the lower bound and will always converge

Mean field approximation


- Let us assume the approximate posterior is fully factorized

$$Q(x) = \prod_i Q_i(x_i)$$

- Then the objective (negative free energy) is

$$\begin{aligned} F(\tilde{P}, Q) &\stackrel{\text{def}}{=} H_Q(x) + \sum_c E_{X_c \sim Q} \ln \phi_c(X_c) \\ &= \sum_i H(Q_i) + \sum_c \sum_{x_c} \left(\prod_{i \in c} Q_i(x_{c,i}) \right) \ln \phi_c(x_c) \end{aligned}$$

- Eg 4x4 grid $O(n_e K^2)$ for energy, $O(n_e K)$ for H



$$\begin{aligned} F[\tilde{P}, Q] &= E_{\{A_{1,1}, A_{2,1}\} \sim Q} [\ln \phi(A_{1,1}, A_{2,1})] + E_{\{A_{2,1}, A_{2,2}\} \sim Q} [\ln \phi(A_{2,1}, A_{2,2})] + E_{\{A_{2,1}, A_{3,1}\} \sim Q} [\ln \phi(A_{2,1}, A_{3,1})] + \dots \\ &\quad E_Q [\ln \phi(A_{1,1}, A_{1,2})] + E_Q [\ln \phi(A_{1,2}, A_{1,3})] + E_Q [\ln \phi(A_{1,3}, A_{1,4})] + \dots \\ &\quad H_Q(A_{1,1}) + H_Q(A_{1,2}) + H_Q(A_{1,3}) + H_Q(A_{1,4}) + \dots \\ &\quad H_Q(A_{4,1}) + H_Q(A_{4,2}) + H_Q(A_{4,3}) + H_Q(A_{4,4}) \end{aligned}$$

Convexity

- Objective is concave in each arg (entropy is concave in each Q_i , expected energy is linear in Q_i)

$$F(\tilde{P}, Q) = \sum_i H(Q_i) + \sum_c \sum_{x_c} \left(\prod_{i \in c} Q_i(x_{c,i}) \right) \ln \phi_c(x_c)$$

- The set of completely factorized distributions is not convex

$$Q^3(x) = \lambda \prod_i Q^1(x_i) + (1 - \lambda) \prod_i Q_i^2(x_i) \quad \text{Not factorized}$$

- Hence we are optimizing the objective over a non-convex space, and will be subject to local maxima
- Let us derive equations that characterize the fixed points. These could correspond to saddle points or local minima, but such points are unstable and unlikely to be the result of our iterative update scheme.

Notation

- Define

$$\langle f(x_h) \rangle \stackrel{\text{def}}{=} \sum_{x_h} \left[\prod_{i \in h} Q_i(x_i) \right] f(x_h)$$

$$\langle f(x_h) \rangle_{j,k} \stackrel{\text{def}}{=} \sum_{x_h \setminus x_j} \left[\prod_{i \in h, i \neq j} Q_i(x_i) \right] f(x_h | x_j = k)$$

$$\langle f(x_h) \rangle = \sum_k Q_j(x_j = k) \langle f(x_h) \rangle_{j,k}$$

$$\ln p(x_v) \geq \sum_c \langle \ln \phi_c(x_c) \rangle + \sum_i H(Q_i)$$

$$= \sum_k Q_j(k) \sum_c \langle \ln \phi_c(x_c) \rangle_{j,k} + H(Q_j) + \sum_{i \neq j} H(Q_i)$$

We mostly follow Tommi Jaakkola's notation rather than Daphne Koller's

Mean field equations

$$\ln p(x_v) \geq \sum_k Q_j(k) \sum_c \langle \ln \phi_c(x_c) \rangle_{j,k} + H(Q_j) + \sum_{i \neq j} H(Q_i)$$

$$\stackrel{\text{def}}{=} L(Q_j)$$

$$S_{j,k} \stackrel{\text{def}}{=} \sum_{c:j \in c} \langle \ln \phi_c(x_c) \rangle_{j,k}$$

$$L(Q_j) = \sum_k Q_j(k) (S_{j,k} - \ln Q_j(k)) + C$$

$$L(Q_j, \lambda) \stackrel{\text{def}}{=} L(Q_j) + \lambda \left(\sum_{k'} Q_j(k') - 1 \right)$$


$$\frac{\partial}{\partial Q_j(k)} L(Q_j, \lambda) = S_{j,k} - \ln Q_j(k) - 1 + \lambda = 0$$

$$\begin{aligned} Q_j(k) &= \exp(S_{j,k}) \exp(\lambda - 1) \\ &= \frac{1}{Z_j} \exp\left(\sum_c \langle \ln \phi_c(x_c) \rangle_{j,k}\right) \end{aligned}$$

Example: grid

$$Q(x_i) = \frac{1}{Z_i} \exp \left\{ \sum_{\phi: X_i \in \text{Scope}[\phi]} E(\mathbf{U}_\phi - \{X_i\}) \sim Q[\ln \phi(\mathbf{U}_\phi, x_i)] \right\}$$

$$Q(a_{i,j}) = \frac{1}{Z_{i,j}} \exp \left\{ \begin{array}{l} \sum_{a_{i-1,j}} Q(a_{i-1,j}) \ln(\phi(a_{i-1,j}, a_{i,j})) + \\ \sum_{a_{i,j-1}} Q(a_{i,j-1}) \ln(\phi(a_{i,j-1}, a_{i,j})) + \\ \sum_{a_{i+1,j}} Q(a_{i+1,j}) \ln(\phi(a_{i,j}, a_{i+1,j})) + \\ \sum_{a_{i,j+1}} Q(a_{i,j+1}) \ln(\phi(a_{i,j}, a_{i,j+1})) \end{array} \right\}.$$





EM

- Suppose we want to find a MAP estimate

$$\max_{\theta} \log p(\theta) + \sum_n \log p(x_n|\theta)$$

- If we have latent variables Z we can use EM
- E step: compute expected complete data log joint

$$f(\theta, \theta_{old}) = \log p(\theta) + \sum_{n=1}^N \sum_z p(z|x_n, \theta_{old}) \log p(z, x_n|\theta)$$

- M step: set

$$\theta_{new} = \arg \max_{\theta} f(\theta, \theta_{old})$$

Variational EM

- Consider the negative free energy

$$F(x, Q, \theta) = \sum_z Q(z) \log p(x, z | \theta) + H(Q)$$

- Earlier we showed this is a lower bound on the log-likelihood

$$F(x, Q, \theta) = \ln Z(x, \theta) - D(Q || p(z|x, \theta))$$

$$\log p(x|\theta) = \ln Z = \max_Q F(x, Q, \theta) = F(x, Q^*, \theta) \geq F(x, Q, \theta)$$

- Where the bound is tight if $Q^*(z) = p(z|x, \theta)$
- E step: find $Q_n(z)$ that maximize

$$F(x_n, Q_n, \theta_{old})$$

- M step: find θ that maximize

$$\log p(\theta) + \sum_n F(x_n, Q_n, \theta)$$

Variational EM

- An exact E step is equivalent to setting

$$Q_n(z) = p(z|x_n, \theta_{old})$$

- The corresponding M step maximizes

$$\begin{aligned} \sum_n F(x_n, Q_n, \theta) &= \sum_n \left[\sum_z p(z|x_n, \theta_{old}) \log p(z, x_n|\theta) \right] + H(Q_n) \\ &= f(\theta, \theta_{old}) + \sum H(Q_n) \end{aligned}$$

- Since $H(Q_n)$ is independentⁿ of θ , this reduces to the standard EM algorithm.
- Generalized EM merely increases (not maximizes) θ in the M step.
- Similarly we can simply improve Q_n in the E step

Variational Bayes

- We can replace the point estimate of θ with a distribution and try to minimize

$$D(Q(z_{1:N}, \theta | x_{1:N}) || p(z_{1:N}, \theta | x_{1:N}))$$

- The distinction between E and M vanishes: we are just doing sequential updates of $Q(Z_n)$ and $Q(\theta)$
- This gives us the benefits of being Bayesian for the same computational speed as EM

VB for univariate Gaussian

$$q_j^*(\mathbf{Z}_j) = \frac{\exp(\mathbb{E}_{i \neq j}[\ln p(\mathbf{X}, \mathbf{Z})])}{\int \exp(\mathbb{E}_{i \neq j}[\ln p(\mathbf{X}, \mathbf{Z})]) d\mathbf{Z}_j}. \quad \ln q_j^*(\mathbf{Z}_j) = \mathbb{E}_{i \neq j}[\ln p(\mathbf{X}, \mathbf{Z})] + \text{const.}$$

$$p(\mathcal{D}|\mu, \tau) = \left(\frac{\tau}{2\pi}\right)^{N/2} \exp\left\{-\frac{\tau}{2} \sum_{n=1}^N (x_n - \mu)^2\right\}. \quad \begin{aligned} p(\mu|\tau) &= \mathcal{N}(\mu|\mu_0, (\lambda_0\tau)^{-1}) \\ p(\tau) &= \text{Gam}(\tau|a_0, b_0) \end{aligned}$$

$$q(\mu, \tau) = q_\mu(\mu)q_\tau(\tau).$$

Gaussian

$$\begin{aligned} \ln q_\mu^*(\mu) &= \mathbb{E}_\tau [\ln p(\mathcal{D}|\mu, \tau) + \ln p(\mu|\tau)] + \text{const.} \\ &= -\frac{\mathbb{E}[\tau]}{2} \left\{ \lambda_0(\mu - \mu_0)^2 + \sum_{n=1}^N (x_n - \mu)^2 \right\} + \text{const.} \end{aligned}$$

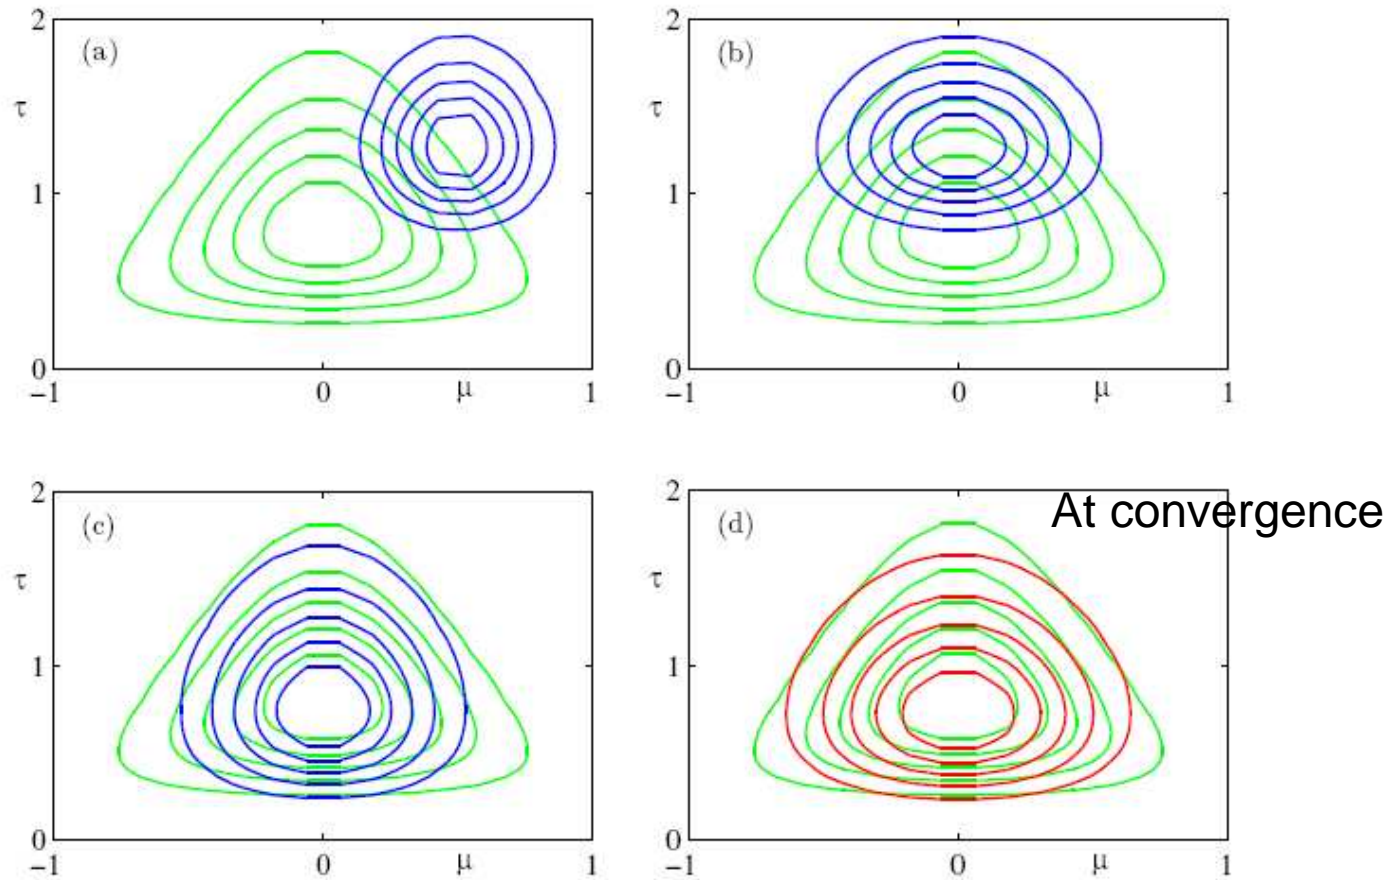
$$\begin{aligned} \mu_N &= \frac{\lambda_0\mu_0 + N\bar{x}}{\lambda_0 + N} \\ \lambda_N &= (\lambda_0 + N)\mathbb{E}[\tau]. \end{aligned}$$

$$\begin{aligned} \ln q_\tau^*(\tau) &= \mathbb{E}_\mu [\ln p(\mathcal{D}|\mu, \tau) + \ln p(\mu|\tau)] + \ln p(\tau) + \text{const.} \\ &= (a_0 - 1) \ln \tau - b_0\tau + \frac{N+1}{2} \ln \tau \\ &\quad - \frac{\tau}{2} \mathbb{E}_\mu \left[\sum_{n=1}^N (x_n - \mu)^2 + \lambda_0(\mu - \mu_0)^2 \right] + \text{const.} \end{aligned}$$

Gamma

$$\begin{aligned} a_N &= a_0 + \frac{N+1}{2} \\ b_N &= b_0 + \frac{1}{2} \mathbb{E}_\mu \left[\sum_{n=1}^N (x_n - \mu)^2 + \lambda_0(\mu - \mu_0)^2 \right]. \end{aligned}$$

VB for univariate Gaussian



Green = exact posterior (NormalGamma), blue = factorized approximation

VB for mixtures of Gaussians

Inference

$$q(\mathbf{Z}, \pi, \mu, \Lambda) = q(\mathbf{Z})q(\pi, \mu, \Lambda).$$

$$\ln q^*(\mathbf{Z}) = \mathbb{E}_{\pi, \mu, \Lambda}[\ln p(\mathbf{X}, \mathbf{Z}, \pi, \mu, \Lambda)] + \text{const.}$$

$$\ln q^*(\mathbf{Z}) = \sum_{n=1}^N \sum_{k=1}^K z_{nk} \ln \rho_{nk} + \text{const.}$$


$$\ln \rho_{nk} = \mathbb{E}[\ln \pi_k] + \frac{1}{2} \mathbb{E}[\ln |\Lambda_k|] - \frac{D}{2} \ln(2\pi) - \frac{1}{2} \mathbb{E}_{\mu_k, \Lambda_k}[(\mathbf{x}_n - \mu_k)^\top \Lambda_k (\mathbf{x}_n - \mu_k)]$$

$$q^*(\mathbf{Z}) \propto \prod_{n=1}^N \prod_{k=1}^K \rho_{nk}^{z_{nk}}.$$



Multinomial (soft responsibilities), as in EM, except we used expected parameters rather than plug-in

Model



$$p(\mathbf{X}, \mathbf{Z}, \pi, \mu, \Lambda) = p(\mathbf{X}|\mathbf{Z}, \mu, \Lambda)p(\mathbf{Z}|\pi)p(\pi)p(\mu, \Lambda)$$

$$p(\mathbf{X}|\mathbf{Z}, \mu, \Lambda) = \prod_{n=1}^N \prod_{k=1}^K \mathcal{N}(\mathbf{x}_n | \mu_k, \Lambda_k^{-1})^{z_{nk}}$$


$$p(\mathbf{Z}|\pi) = \prod_{n=1}^N \prod_{k=1}^K \pi_k^{z_{nk}}.$$

$$p(\pi) = \text{Dir}(\pi | \alpha_0) = C(\alpha_0) \prod_{k=1}^K \pi_k^{\alpha_0 - 1}$$

$$p(\mu, \Lambda) = \prod_{k=1}^K \mathcal{N}(\mu_k | \mathbf{m}_0, (\beta_0 \Lambda_k)^{-1}) \mathcal{W}(\Lambda_k | \mathbf{W}_0, \nu_0)$$

Automatic model selection

- Recall $\pi \sim \text{Dir}(\alpha)$. If $\alpha \ll 1$, we prefer skewed π and hence sparse z .



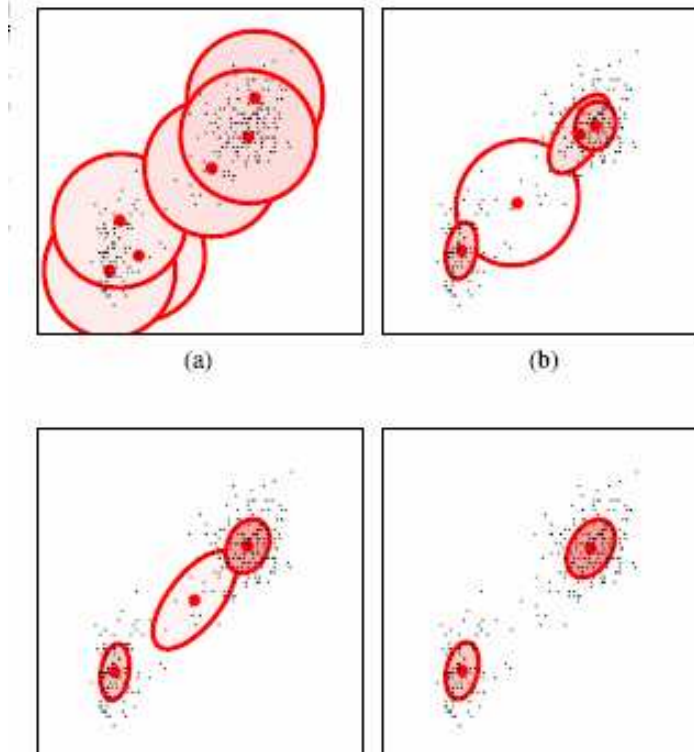
- MAP estimate from regular EM is

$$\hat{\pi}_k = \frac{\sum_n r_{nk} + \alpha_k - 1}{\sum_k (r_{nk} + \alpha_k - 1)} = \frac{N_k + \alpha - 1}{N + K\alpha - K}$$

- Posterior mean estimate from VB is

$$\hat{\pi}_k = \frac{\sum_n r_{nk} + \alpha_k}{\sum_k (r_{nk} + \alpha_k)} = \frac{N_k + \alpha}{N + K\alpha} \rightarrow \frac{\alpha}{N + K\alpha} \rightarrow 0$$

Selecting K with one run of VB



Variational message passing

- Consider a DAG model

$$p(\mathbf{x}) = \prod_i p(\mathbf{x}_i | \text{pa}_i)$$

- The mean field equations are



$$\ln q_j^*(\mathbf{x}_j) = \mathbb{E}_{i \neq j} \left[\sum_i \ln p(\mathbf{x}_i | \text{pa}_i) \right] + \text{const.}$$

- The only terms that depend on \mathbf{x}_j are in \mathbf{x}_j 's Markov blanket
- If all CPDs have conjugate-exponential form, the VB updates can be converted into a msg passing algorithm
- VIBES software (John Winn)



Structured variational approx

- Rather than assuming Q is fully factorized, we can use any structure for which computing the expectations of $\ln \phi_c$ and the entropy is tractable



$$Q(\mathcal{X}) = \frac{1}{Z_Q} \prod_{j=1}^J \psi_j$$

$\phi = \text{model}, \psi = \text{approx}$

Corollary 11.5.13: *If $Q(\mathcal{X}) = \frac{1}{Z_Q} \prod_j \psi_j$, then the potential ψ_j is a stationary point of the energy functional if and only if:*

$$\psi_j(c_j) \propto \exp \left\{ \mathbb{E}_Q [\ln \tilde{P}_\Phi | c_j] - \sum_{k \neq j} \mathbb{E}_Q [\ln \psi_k | c_j] \right\}. \quad (11.59)$$