Variational Inference

CS772A: Probabilistic Machine Learning

Piyush Rai

Variational Bayes (VB) or Variational Inference (VI)

- Consider a model with data $\boldsymbol{X}$ and unknowns $\boldsymbol{Z}$. Goal: Compute the posterior $p(\boldsymbol{Z}|\boldsymbol{X})$
- $\boldsymbol{Z}$ denotes all unknowns (params, latent vars, hyperparams of likelihood, prior, etc)
- Assuming $p(\boldsymbol{Z}|\boldsymbol{X})$ is intractable, VB/VI approximates it by a distr $q(\boldsymbol{Z}|\phi)$ or $q_\phi(\boldsymbol{Z})$

  Defines a class of distributions parametrized by ϕ

- We find the best approx. distr by finding ϕ s.t. its distance from $p(\boldsymbol{Z}|\boldsymbol{X})$ is minimized

  Often called variational parameters

But since we don't know $p(\boldsymbol{Z}|\boldsymbol{X})$, can we easily solve this optimization problem?

Often, we will simply write it as $\text{argmin}_q\ \text{KL}[q||p_z]$

Other measures have also been used such as reverse KL ($\text{KL}[p||q]$), and various other divergence functions defined for distributions

$$\phi^* = \text{argmin}_\phi\ \text{KL}[q_\phi(\boldsymbol{Z})||p(\boldsymbol{Z}|\boldsymbol{X})]$$

VI turns inference into optimization

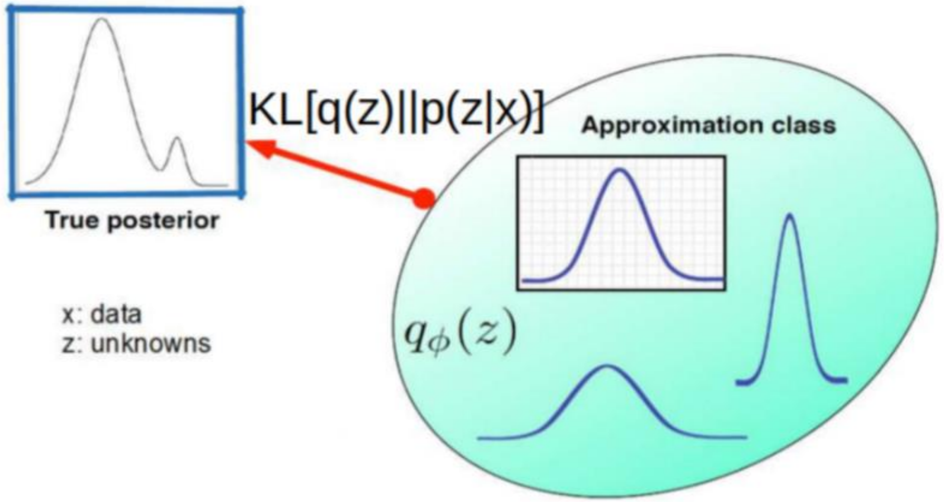

- Note: The name "variational" comes from Physics
  - Optimizing functions of distributions (KL is a func of distr)

Variational Bayes (VB) or Variational Inference (VI)

- VB/VI is based on following identity for the log marg-lik (log evidence) of a model $\boldsymbol{m}$

Similar as the identify we had in case of EM, which was defined for log of the ILL

$$\log p(\boldsymbol{X}|m) = \mathcal{L}(q) + \mathrm{KL}(q||p_z)$$

Also, unlike EM, here we don't have any distinction b/w latent variables $\boldsymbol{Z}$ and parameters $\boldsymbol{\Theta}$ (all unknowns will be denoted by $\boldsymbol{Z}$ here, and we have $\mathcal{L}(q)$, not $\mathcal{L}(q,\Theta)$)

$$\mathcal{L}(q) = \int q(\boldsymbol{Z}) \log \left\{ \frac{p(\boldsymbol{X},\boldsymbol{Z})}{q(\boldsymbol{Z})} \right\} d\boldsymbol{Z}$$

$$\mathrm{KL}(q||p_z) = -\int q(\boldsymbol{Z}) \log \left\{ \frac{p(\boldsymbol{Z}|\boldsymbol{X})}{q(\boldsymbol{Z})} \right\} d\boldsymbol{Z}$$

- Since the log evidence $\log p(\boldsymbol{X}|m)$ is constant w.r.t $\boldsymbol{Z}$, we must have

$$\mathrm{argmin}_q \ \mathrm{KL}[q||p_z] = \mathrm{argmax}_q \ \mathcal{L}(q)$$

- Also note that since $\mathrm{KL}[q||p_z] \geq 0$, we must have $\log p(\boldsymbol{X}|m) \geq \mathcal{L}(q)$
- Therefore, $\mathcal{L}(q)$ is also known as Evidence Lower Bound (ELBO)
  - VB/VI finds the best $q(Z)$ by maximizing the ELBO w.r.t. q

VB/VI = Maximizing the ELBO

- Notation: $q(\boldsymbol{Z})$, $q(\boldsymbol{Z}|\phi)$, $q_\phi(\boldsymbol{Z})$, all refer to the same thing (the approx. distr.)
- VB/VI finds an approximating distribution $q(\boldsymbol{Z})$ that maximizes the ELBO

$$\mathcal{L}(q) = \int q(\boldsymbol{Z}) \log \left[\frac{p(\boldsymbol{X},\boldsymbol{Z})}{q(\boldsymbol{Z})}\right] d\boldsymbol{Z}$$

- Since $q(\boldsymbol{Z})$ depends on ϕ, the ELBO is essentially a function of ϕ

$$\mathcal{L}(q) = \mathcal{L}(\phi) = \mathbb{E}_q[\log p(\boldsymbol{X},\boldsymbol{Z})] - \mathbb{E}_q[\log q(\boldsymbol{Z})]$$
$$= \mathbb{E}_q[\log p(\boldsymbol{X}|\boldsymbol{Z})] - \mathrm{KL}[q(\boldsymbol{Z})||p(\boldsymbol{Z})]$$

- Thus maximizing the ELBO will give an approximating distr. $q(\boldsymbol{Z})$ which
  - Explains the data $\boldsymbol{X}$ well, i.e., gives it large probability (large $\mathbb{E}_q[\log p(\boldsymbol{X}|\boldsymbol{Z})]$)
  - Is close to the prior $p(\boldsymbol{Z})$, i.e. is simple/regularized (small $\mathrm{KL}[\log q(\boldsymbol{Z})||p(\boldsymbol{Z})]$)

Maximizing the ELBO

- The goal is to maximize the ELBO

$$\mathcal{L}(q) = \mathcal{L}(\phi) = \mathbb{E}_q[\log p(\boldsymbol{X}, \boldsymbol{Z})] - \mathbb{E}_q[\log q(\boldsymbol{Z})]$$
$$= \mathbb{E}_q[\log p(\boldsymbol{X}|\boldsymbol{Z})] - \mathrm{KL}[\log q(\boldsymbol{Z})||p(\boldsymbol{Z})]$$

- This may still be hard because
  - ELBO expression has expectations, computed which may be intractable
  - Maximizing the ELBO will require computing gradients which may not always be easy

E.g., part of the ELBO may have terms that are not differentiable

- Some of the ways to make this problem easier
  - Restricting the form of our approximation $\boldsymbol{q(Z)}$, e.g., mean-field VI
  - Using Monte-Carlo approximation of the expectation/gradient of the ELBO
- For locally conjugate models, ELBO maximization is easy
  - Closed form updates for $\boldsymbol{q(Z)}$

Some Properties of VI

- Recall that VI approximates a posterior p by finding q that minimizes $\mathrm{KL}(q||p)$

$$\mathrm{KL}(q||p) = \int q(\mathbf{Z}) \log \left[\frac{q(\mathbf{Z})}{p(\mathbf{Z}|\mathbf{X})} \right]$$

- $q(\boldsymbol{Z})$ will be small where $p(\boldsymbol{Z}|\boldsymbol{X})$ is small otherwise KL will blow up
- Thus $q(\boldsymbol{Z})$ avoids low-probability regions of the true posterior

q (red) avoids regions of p (green) where the latter has low values

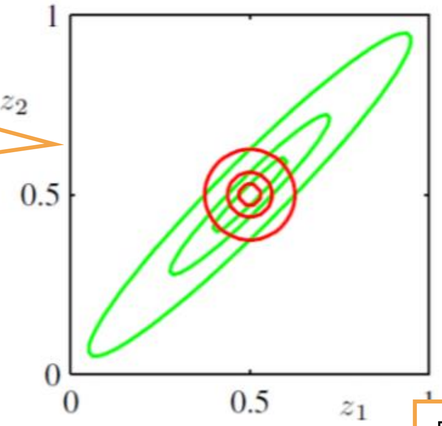

q (red) concentrated on one of the modes of p (blue)

For q to also capture the other mode, it will require crossing the low-probability region of p, thereby blowing up the KL

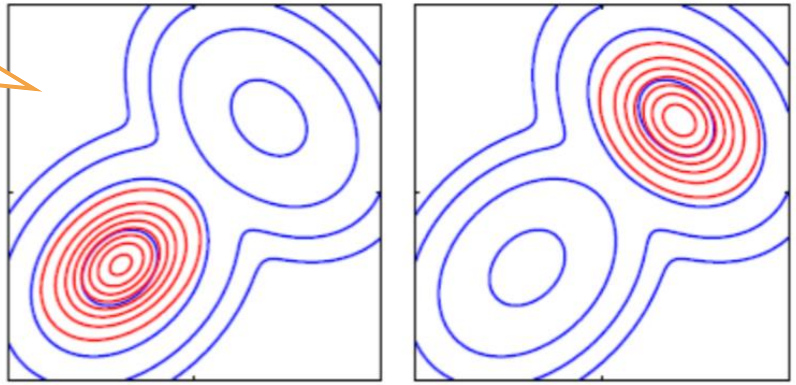

EP minimizes $\mathrm{KL}(p||q)$

- Some methods, e.g., Expectation Propagation (EP), can avoid this behavior

ELBO for Model Selection

- Recall that ELBO is a lower bound on log of model evidence $\log p(\boldsymbol{X}|m)$
- Can compute ELBO for each model m and choose the one with largest ELBO

Plot of the variational lower bound $\mathcal{L}$ versus the number K of components in the Gaussian mixture model, for the Old Faithful data, showing a distinct peak at $K = 2$ components. For each value of K, the model is trained from 100 different random starts, and the results shown as '+' symbols plotted with small random horizontal perturbations so that they can be distinguished. Note that some solutions find suboptimal local maxima, but that this happens infrequently.

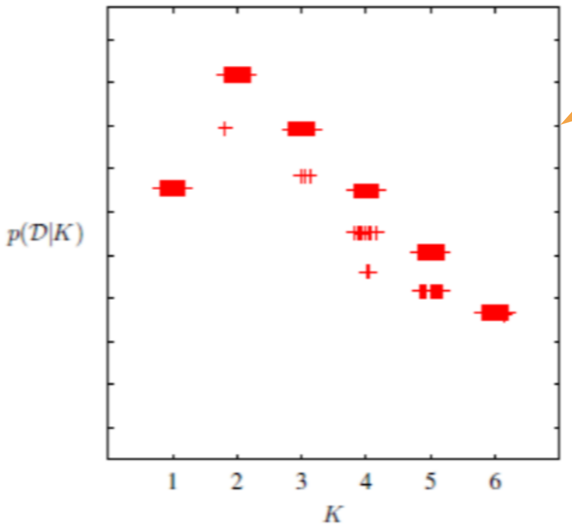

Each value of K represents a different model

- Some criticism since we are using a lower-bound but often works well in practice

VI and Convergence

- VI is guaranteed to converge to a local optima (just like EM)
- Therefore proper initialization is important (just like EM)
  - Can sometimes run multiple times with different initializations and choose the best run

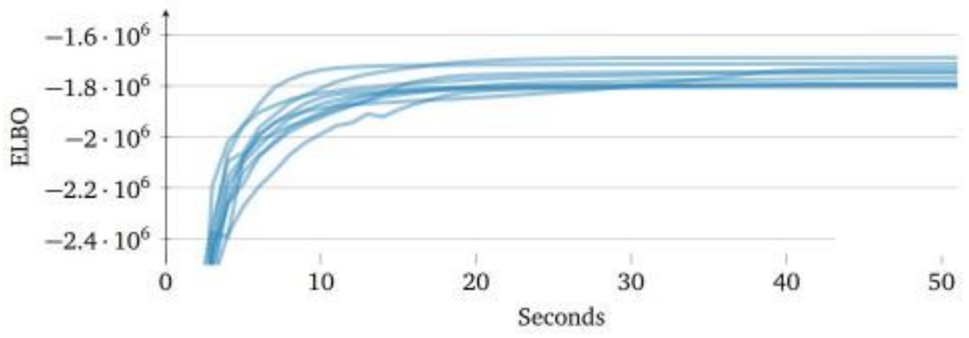

Different initializations may lead to different optima

- ELBO increases monotonically with iterations
  - Can thus monitor the ELBO to assess convergence

Variational Inference and Expectation Maximization

- VI can be seen as a generalization of the EM algorithm
- In VI, there is no distinction between parameters Θ and latent variables $\boldsymbol{Z}$
  - Also recall that EM finds CP of $\boldsymbol{Z}$ and point estimate for Θ
  - VI treats all unknowns identically and infers posterior for all
- VI can be used within an EM algorithm if the E step is intractable
- E step is intractable if the CP of latent variables given params is intractable
- This version of EM is known as Variational EM (VEM)
- If we only care about point estimates of the parameters, VEM is widely used if the CP of latent variables is intractable

Mean-Field VI

The name "mean-field" comes from statistical physics literature

- One of the simplest ways for doing VB/VI
- Assumes unknowns $\boldsymbol{Z}$ can be partitioned into $\boldsymbol{M}$ groups $\boldsymbol{Z}_1, \boldsymbol{Z}_2, \ldots, \boldsymbol{Z}_M$, s.t.,

$$q(\boldsymbol{Z}|\phi) = \prod_{i=1}^{M} q(\boldsymbol{Z}_i|\phi_i)$$

As a shorthand, often written as $q = \prod_{i=1}^{M} q_i$ where $q_i = q(Z_i|\phi_i)$

- Learning the optimal $\boldsymbol{q}$ reduces to learning the optimal $\boldsymbol{q_1, q_2, \ldots, q_M}$
- Groups usually chosen based on model's structure, e.g., in Bayesian linear regression

$$p(\boldsymbol{w}, \beta, \lambda|X, y) \approx q(Z|\phi) = q(\boldsymbol{w}, \beta, \lambda|\phi) = q(\boldsymbol{w}|\phi_w)p(\beta|\phi_\beta)p(\lambda|\phi_\lambda)$$

- Mean-field is a very restrictive assumption. Ignores the correlations among unknowns
  - Less restrictive versions also exist, such as structured mean-field (factorization is still there but only among groups of unknowns)

Deriving Mean-Field VI Updates

- With $q = \prod_{i=1}^{M} q_i$, what's the optimal q_i when we do $\text{argmax}_q \mathcal{L}(q)$?
- Note that under this mean-field assumption, the ELBO simplifies to

$$\mathcal{L}(q) = \int q(\mathbf{Z}) \log \left[\frac{p(\mathbf{X},\mathbf{Z})}{q(\mathbf{Z})}\right] d\mathbf{Z} = \int \prod_i q_i \left[\log p(X,Z) - \sum_i \log q_i\right] d\mathbf{Z}$$

- Suppose we wish to find the optimal q_j given all other q_i's $(i \neq j)$ as fixed, then

$$\mathcal{L}(q) = \int q_j \left[\int \log p(X,Z) \prod_{i\neq j} q_i \, dZ_i\right] dZ_j - \int q_j \log q_j dZ_j + \text{const w.r.t. } q_j$$

$$= \int q_j \log \hat{p}(X,Z_j) \, dZ_j - \int q_j \log q_j Z_j$$

$$= -\text{KL}(q_j || \hat{p})$$

$\log \hat{p}(X,Z_j) = \mathbb{E}_{i\neq j}[\log p(\mathbf{X},\mathbf{Z})]$ + const

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

- Thus $q_j^* = \text{argmax}_{q_j} \mathcal{L}(q) = \text{argmin}_{q_j} \text{KL}(q_j || \hat{p}) = \hat{p}(X, Z_j)$

Deriving Mean-Field VI Updates

- So we saw that the optimal $\boldsymbol{q_j}$ when doing mean-field VI is

$$q_j^*(\boldsymbol{Z}_j) = \frac{\exp(\mathbb{E}_{i\neq j}[\log p(\boldsymbol{X},\boldsymbol{Z})])}{\int \exp(\mathbb{E}_{i\neq j}[\log p(\boldsymbol{X},\boldsymbol{Z})]\, d\boldsymbol{Z}_j}$$

- Note: Can often just compute the numerator and recognize denominator by inspection
- Important: For locally conjugate models, $q_j^*(\boldsymbol{Z}_j)$ will have the same form as prior $p(Z_j)$
  - Only the distribution parameters will be different
- Important: For estimating $\boldsymbol{q_j}$ the required expectation depends on other $\{q_i\}_{i\neq j}$
  - Thus we use an alternating update scheme for these (akin to ALT-OPT, Gibbs sampling, etc)
- Guaranteed to converge (to a local optima)
  - We are basically solving a sequence of concave maximization problems
  - Reason: $\mathcal{L}(q) = \int q_j \log \hat{p}(X, Z_j)\, Z_j \; - \; \int q_j \log q_j Z_j$ is concave in $\boldsymbol{q_j}$

The Mean-Field VI Algorithm

- Also known as Co-ordinate Ascent Variational Inference (CAVI) Algorithm
- Input: Model in form of priors and likelihood, or joint $p(\boldsymbol{X},\boldsymbol{Z})$, Data $\boldsymbol{X}$
- Output: A variational distribution $q(Z) = \prod_{j=1}^{M} q_j(\boldsymbol{Z}_j)$
- Initialize: Variational distributions $q_j(\boldsymbol{Z}_j)$, $j = 1,2, \dots M$
- While the ELBO has not converged
  - For each $j = 1,2, \dots M$, set

$$q_j(\boldsymbol{Z}_j) \propto \exp(\mathbb{E}_{i \neq j}[\log p(\boldsymbol{X},\boldsymbol{Z})])$$

  - Compute ELBO $\mathcal{L}(q) = \mathbb{E}_q[\log p(\boldsymbol{X},\boldsymbol{Z})] - \mathbb{E}_q[\log q(\boldsymbol{Z})]$

Mean-Field VI: A Simple Example

- Consider data $\mathbf{X} = \{x_1, x_2, \ldots, x_N\}$ from a one-dim Gaussian $\mathcal{N}(\mu, \tau^{-1})$
- Assume the following normal-gamma prior on μ and τ

$$p(\mu|\tau) = \mathcal{N}(\mu|\mu_0, (\lambda_0\tau)^{-1}) \quad p(\tau) = \text{Gamma}(\tau|a_0, b_0)$$

- Posterior is also normal-gamma due to the jointly conjugate prior
- Let's anyway verify this by trying mean-field VI for this model
- With mean-field assumption on the variational posterior $q(\mu, \tau) = q_\mu(\mu)q_\tau(\tau)$

$$\begin{aligned}\log q_\mu^*(\mu) &= \mathbb{E}_{q_\tau}[\log p(\mathbf{X}, \mu, \tau)] + \text{const} \\ \log q_\tau^*(\tau) &= \mathbb{E}_{q_\mu}[\log p(\mathbf{X}, \mu, \tau)] + \text{const}\end{aligned}$$

- In this example, the log-joint $\log p(\mathbf{X}, \mu, \tau) = \log p(\mathbf{X}|\mu, \tau) + \log p(\mu|\tau) + \log p(\tau)$. Thus

$$\begin{aligned}\log q_\mu^*(\mu) &= \mathbb{E}_{q_\tau}[\log p(\mathbf{X}|\mu, \tau) + \log p(\mu|\tau)] + \text{const} \qquad \text{(only keeping terms that involve } \mu) \\ \log q_\tau^*(\tau) &= \mathbb{E}_{q_\mu}[\log p(\mathbf{X}|\mu, \tau) + \log p(\mu|\tau) + \log p(\tau)] + \text{const}\end{aligned}$$

Mean-Field VI: A Simple Example

- Substituting $p(\mathbf{X}|\mu,\tau) = \prod_{n=1}^{N} p(x_n|\mu,\tau)$ and $p(\mu|\tau)$, we get

$$\log q_\mu^*(\mu) = \mathbb{E}_{q_\tau}[\log p(\mathbf{X}|\mu,\tau) + \log p(\mu|\tau)] + \text{const}$$
$$= -\frac{\mathbb{E}_{q_\tau}[\tau]}{2}\left\{\sum_{n=1}^{N}(x_n-\mu)^2 + \lambda_0(\mu-\mu_0)^2\right\} + \text{const}$$

- (Verify) The above is log of a Gaussian. This $q_\mu^* = \mathcal{N}(\mu|\mu_N,\lambda_N)$ with

$$\mu_N = \frac{\lambda_0\mu_0 + N\bar{x}}{\lambda_0 + N} \quad \text{and} \quad \lambda_N = (\lambda_0 + N)\mathbb{E}_{q_\tau}[\tau]$$

This update depends on q_τ

- Proceeding in a similar way (verify), we can show that $q_\tau^* = \text{Gamma}(\tau|a_N, b_N)$

$$a_N = a_0 + \frac{N+1}{2} \quad \text{and} \quad b_N = b_0 + \frac{1}{2}\mathbb{E}_{q_\mu}\left[\sum_{n=1}^{N}(x_n-\mu)^2 + \lambda_0(\mu-\mu_0)^2\right]$$

This update depends on q_μ

- Note: Updates of q_μ^* and q_τ^* depend on each other (hence alternating updates needed)

Mean-Field VI: A Closer Look

- Since $\log q_j^*(\boldsymbol{Z}_j) = \mathbb{E}_{i\neq j}[\log p(\boldsymbol{X},\boldsymbol{Z})]$ + const = $\mathbb{E}_{i\neq j}[\log p(\boldsymbol{X},\boldsymbol{Z}_j,\boldsymbol{Z}_{-j})]$ + const

$$\log q_j^*(\mathbf{Z}_j) = \mathbb{E}_{i\neq j}[\log p(\mathbf{Z}_j|\mathbf{X},\mathbf{Z}_{-j})] + \text{const}$$

For any model

- Thus opt variational distr $q_j^*(\boldsymbol{Z}_j)$ basically requires expectations of CP $p(\boldsymbol{Z}_j|\boldsymbol{X},\boldsymbol{Z}_{-j})$
- For locally conjugate models, CP can be easily found and is usually an exp-fam distr

$$p(\mathbf{Z}_j|\mathbf{X},\mathbf{Z}_{-j}) = h(\mathbf{Z}_j)\exp\left[\eta(\mathbf{X},\mathbf{Z}_{-j})^\top\mathbf{Z}_j - A(\eta(\mathbf{X},\mathbf{Z}_{-j}))\right]$$

- Using the above, we can rewrite the optimal variational distribution as follows

$$\log q_j^*(\mathbf{Z}_j) = \mathbb{E}_{i\neq j}\left[\log\left(h(\mathbf{Z}_j)\exp\left[\eta(\mathbf{X},\mathbf{Z}_{-j})^\top\mathbf{Z}_j - A(\eta(\mathbf{X},\mathbf{Z}_{-j}))\right]\right)\right] + \text{const}$$

$$\Longrightarrow q_j^*(\mathbf{Z}_j) \propto h(\mathbf{Z}_j)\exp\left[\mathbb{E}_{i\neq j}[\eta(\mathbf{X},\mathbf{Z}_{-j})]^\top\mathbf{Z}_j\right] \quad \text{(verify)}$$

For locally conjugate model

- Thus, with local conj, we just require expectation of nat. params. of CP of $\boldsymbol{Z}_j$

VI by Computing ELBO Gradients

Modern VI methods (e.g., those used in Bayesian deep learning) use this idea (more later)

- Can also do VI by computing ELBO's gradient and doing gradient ascent/descent
- Gradient based approach is broadly applicable, not just for mean-field VI
  1. Assume $q(\boldsymbol{Z})$ to be from some family of distributions with variational parameters $\boldsymbol{\phi}$
  2. Write down the full ELBO expression (will give us a function of var parameters $\boldsymbol{\phi}$)

$$\mathcal{L}(q) = \mathcal{L}(\phi) = \mathbb{E}_q[\log p(\mathbf{X}, \mathbf{Z})] - \mathbb{E}_q[\log q(\mathbf{Z})]$$
$$= \int q(\mathbf{Z}) \log p(\mathbf{X}|\mathbf{Z}) d\mathbf{Z} + \int q(\mathbf{Z}) \log p(\mathbf{Z}) d\mathbf{Z} - \int q(\mathbf{Z}) \log q(\mathbf{Z}) d\mathbf{Z}$$

  3. Compute ELBO gradients, i.e., $\nabla_{\boldsymbol{\phi}} \mathcal{L}(\boldsymbol{\phi})$ and use gradient methods to find optimal $\boldsymbol{\phi}$
- Step 2 may be simplified due to the problem structure or the form of $q(\boldsymbol{Z})$
  - i.i.d. observations simplify $\log p(\boldsymbol{X}|\boldsymbol{Z})$; conditionally independent priors simplify $\log p(\boldsymbol{Z})$
  - Locally-conjugate models
  - The mean-field assumption simplifies $q(\boldsymbol{Z})$ as $q = \prod_{i=1}^{M} q_i$
    - Moreover, the last term reduces to sum of entropies of q_i's (which usually has known forms)

Mean-Field VI by Taking ELBO's Gradients

- Mean-field assumption $q(\boldsymbol{Z}|\phi) = \prod_{i=1}^{M} q(\boldsymbol{Z}_i|\phi_i)$ results in following optimal distribution

$$q_j^*(\boldsymbol{Z}_j) = \frac{\exp(\mathbb{E}_{i\neq j}[\log p(\boldsymbol{X},\boldsymbol{Z})])}{\int \exp(\mathbb{E}_{i\neq j}[\log p(\boldsymbol{X},\boldsymbol{Z})]\, d\boldsymbol{Z}_j}$$

This approach is applicable even if we don't have mean-field assumption

Note that here we do not have to assume the form of this variational distribution. We simply compute the RHS and find what it is (in the locally-conjugate case, it will be the same distribution as the prior)

- Alternatively, we can take ELBO's partial deriv w.r.t. $\phi_1, \phi_2, \dots, \phi_M$ to find their optimal values
- Consider a Bayesian linear regression model

Likelihood; Prior on $\boldsymbol{w}$; λ assumed fixed; Prior on variance of Gaussian likelihood

$$y_i \sim \text{Normal}(x_i^T w, \alpha^{-1}), \quad w \sim \text{Normal}(0, \lambda^{-1} I), \quad \alpha \sim \text{Gamma}(a, b)$$

Needed in ELBO; Joint distribution on data and unknowns

$$p(y, w, \alpha|x) = p(\alpha)p(w)\prod_{i=1}^{N} p(y_i|x_i, w, \alpha)$$

Assumed variational posterior with mean-field assumption

$$q(w, \alpha) = q(\alpha)q(w) = \text{Gamma}(\alpha|a', b')\text{Normal}(w|\mu', \Sigma')$$

Note that in this approach, we have to assume a form for each variational distribution. It is common to assume them to have the same form as the respective priors

- Now doing VI amounts to maximizing ELBO to find the optimal variational params a', b', μ', Σ'

Mean-Field VI by Taking ELBO's Gradients

- The ELBO is

For the Bayesian linear regression model, instead of $p(\boldsymbol{X}, \boldsymbol{Z})$, it will be of the form $p(\boldsymbol{y}, \boldsymbol{Z}|\boldsymbol{X})$

$$\mathcal{L}(q) = \mathcal{L}(\phi) = \mathbb{E}_q[\log p(\boldsymbol{X}, \boldsymbol{Z})] - \mathbb{E}_q[\log q(\boldsymbol{Z})] = \mathbb{E}_q[\log p(\boldsymbol{Z})] + \mathbb{E}_q[\log p(\boldsymbol{X}|\boldsymbol{Z})] - \mathbb{E}_q[\log q(\boldsymbol{Z})]$$
$$= \int q(\boldsymbol{Z}) \log p(\boldsymbol{Z}) d\boldsymbol{Z} + \int q(\boldsymbol{Z}) \log p(\boldsymbol{X}|\boldsymbol{Z}) d\boldsymbol{Z} + \int q(\boldsymbol{Z}) \log q(\boldsymbol{Z}) d\boldsymbol{Z}$$

- Thus the ELBO in the Bayesian linear regression model will be (assuming i.i.d. obs)

$$\mathcal{L}(a', b', \mu', \Sigma') = \int q(\alpha) \ln p(\alpha) d\alpha + \int q(w) \ln p(w) dw + \sum_{i=1}^{N} \int\int q(\alpha) q(w) \ln p(y_i|x_i, w, \alpha) dw d\alpha - \int q(\alpha) \ln q(\alpha) d\alpha - \int q(w) \ln q(w) dw$$

Expectations of the log of the prior

Expectations of the log of the likelihood

Expectations of the log of the var. distributions (= their entropies)

- Substituting the priors, likelihoods, and variational distributions

$$\mathcal{L}(a', b', \mu', \Sigma') = (a-1)(\psi(a') - \ln b') - b\frac{a'}{b'} + \text{constant} - \frac{\lambda}{2}(\mu'^T \mu' + \text{tr}(\Sigma')) + \text{constant} + \frac{N}{2}(\psi(a') - \ln b') - \sum_{i=1}^{N} \frac{1}{2}\frac{a'}{b'}\left((y_i - x_i^T \mu')^2 + x_i^T \Sigma' x_i\right) + \text{constant}$$
$$+ a' - \ln b' + \ln \Gamma(a') + (1 - a')\psi(a') + \frac{1}{2} \ln |\Sigma'| + \text{constant}$$

Digamma function (log of gamma function)

- Can now maximize the above ELBO w.r.t. a', b', μ', Σ' in an alternating fashion
- For most models, ELBO or its gradients won't have a simple form (methods like BBVI, reparam trick etc will be needed in those cases)

Coming Up Next

- VI for latent variable models with local and global unknowns
- VI for non-conjugate models
  - Mostly such methods rely on computing approximations of ELBO and/or its gradients