

Probabilistic & Unsupervised Learning

Approximate Inference

Factored Variational Approximations and Variational Bayes

Maneesh Sahani

maneesh@gatsby.ucl.ac.uk

Gatsby Computational Neuroscience Unit, and
MSc ML/CSML, Dept Computer Science
University College London

Term 1, Autumn 2020

Intractabilities and approximations

- ▶ Inference – computational intractability
 - ▶ Factored variational approx
 - ▶ Loopy BP/EP/Power EP
 - ▶ LP relaxations/ convexified BP
 - ▶ Gibbs sampling, other MCMC
- ▶ Inference – analytic intractability
 - ▶ Laplace approximation (global)
 - ▶ Parametric variational approx
 - ▶ Message approximations (linearised, sigma-point, Laplace)
 - ▶ Assumed-density methods and Expectation-Propagation
 - ▶ (Sequential) Monte-Carlo methods
- ▶ Learning – intractable partition function
 - ▶ Sampling parameters
 - ▶ Contrastive divergence
 - ▶ Score-matching
- ▶ Model selection
 - ▶ Laplace approximation / BIC
 - ▶ Variational Bayes
 - ▶ (Annealed) importance sampling
 - ▶ Reversible jump MCMC

Not a complete list!

Intractabilities

- ▶ **Analytic intractability:** non-conjugacy prevents closed-form evaluation of likelihood or free energy

$$\ell(\theta) = \int d\mathcal{Z} P(\mathcal{Z}|\theta)P(\mathcal{X}|\mathcal{Z}, \theta) \quad \mathcal{F}(\theta, Q) = \int d\mathcal{Z} Q(\mathcal{Z}) \log P(\mathcal{X}, \mathcal{Z}|\theta) + \mathbf{H}[Q]$$

- ▶ **Computational intractability:** graphical structure prevents simplification of high-dimensional sums or integrals

$$Q(\mathbf{z}_i) \propto \int d\mathbf{z}_1 \dots d\mathbf{z}_{i-1} d\mathbf{z}_{i+1} \dots d\mathbf{z}_k \log P(\mathcal{X}, \mathcal{Z}|\theta)$$


- ▶ **Intractable normalisers:** learning requires normaliser gradient

$$\nabla_{\theta} \mathcal{F}(\theta, Q) = \nabla_{\theta} \langle E(\mathcal{X}, \mathcal{Z}|\theta) \rangle_Q - \nabla_{\theta} \log Z(\theta)$$

- ▶ **Model selection or weighting** (by marginal likelihood)

$$p(\mathcal{D}|m) = \int d\theta p(\theta|m)p(\mathcal{D}|\theta, m)$$

Computational intractability – distributed models



Consider an FHMM with M state variables taking on K values each.

- ▶ Moralisation puts simultaneous states $(s_i^{(1)}, s_i^{(2)}, \dots, s_i^{(M)})$ into a single clique
- ▶ Triangulation extends cliques to size $M + 1$
- ▶ Each state takes K values \Rightarrow sums over K^{M+1} terms.
- ▶ **Factorial prior \nrightarrow Factorial posterior (explaining away).**

Variational methods **approximate** the posterior, often in a factored form.

To see how they work, we need to review the free-energy interpretation of EM.

The Free Energy for a Latent Variable Model

Observed data $\mathcal{X} = \{\mathbf{x}_i\}$; Latent variables $\mathcal{Z} = \{\mathbf{z}_i\}$; Parameters θ .

Goal: Maximize the log likelihood wrt θ (i.e. ML learning):

$$\ell(\theta) = \log P(\mathcal{X}|\theta) = \log \int P(\mathcal{Z}, \mathcal{X}|\theta) d\mathcal{Z}$$

Any distribution, $q(\mathcal{Z})$, over the hidden variables can be used to obtain a lower bound on the log likelihood using Jensen's inequality:


$$\ell(\theta) = \log \int q(\mathcal{Z}) \frac{P(\mathcal{Z}, \mathcal{X}|\theta)}{q(\mathcal{Z})} d\mathcal{Z} \geq \int q(\mathcal{Z}) \log \frac{P(\mathcal{Z}, \mathcal{X}|\theta)}{q(\mathcal{Z})} d\mathcal{Z} \stackrel{\text{def}}{=} \mathcal{F}(q, \theta)$$

$$\begin{aligned} \int q(\mathcal{Z}) \log \frac{P(\mathcal{Z}, \mathcal{X}|\theta)}{q(\mathcal{Z})} d\mathcal{Z} &= \int q(\mathcal{Z}) \log P(\mathcal{Z}, \mathcal{X}|\theta) d\mathcal{Z} - \int q(\mathcal{Z}) \log q(\mathcal{Z}) d\mathcal{Z} \\ &= \int q(\mathcal{Z}) \log P(\mathcal{Z}, \mathcal{X}|\theta) d\mathcal{Z} + \mathbf{H}[q], \end{aligned}$$

where $\mathbf{H}[q]$ is the entropy of $q(\mathcal{Z})$.

So: $\mathcal{F}(q, \theta) = \langle \log P(\mathcal{Z}, \mathcal{X}|\theta) \rangle_{q(\mathcal{Z})} + \mathbf{H}[q]$

EM as Coordinate Ascent in \mathcal{F}



The E and M steps of EM

The log likelihood is bounded below by:

$$\mathcal{F}(q, \theta) = \langle \log P(\mathcal{Z}, \mathcal{X}|\theta) \rangle_{q(\mathcal{Z})} + \mathbf{H}[q] = \ell(\theta) - \mathbf{KL}[q(\mathcal{Z})||P(\mathcal{Z}|\mathcal{X}, \theta)]$$

EM alternates between:

E step: optimise $\mathcal{F}(q, \theta)$ wrt distribution over hidden variables holding parameters fixed:

$$q^{(k)}(\mathcal{Z}) := \operatorname{argmax}_{q(\mathcal{Z})} \mathcal{F}(q(\mathcal{Z}), \theta^{(k-1)}) = P(\mathcal{Z}|\mathcal{X}, \theta^{(k-1)})$$

M step: maximise $\mathcal{F}(q, \theta)$ wrt parameters holding hidden distribution fixed:

$$\theta^{(k)} := \operatorname{argmax}_{\theta} \mathcal{F}(q^{(k)}(\mathcal{Z}), \theta) = \operatorname{argmax}_{\theta} \langle \log P(\mathcal{Z}, \mathcal{X}|\theta) \rangle_{q^{(k)}(\mathcal{Z})}$$

EM Never Decreases the Likelihood

The E and M steps together never decrease the log likelihood:

$$\ell(\theta^{(k-1)}) \underset{\text{E step}}{=} \mathcal{F}(q^{(k)}, \theta^{(k-1)}) \underset{\text{M step}}{\leq} \mathcal{F}(q^{(k)}, \theta^{(k)}) \underset{\text{Jensen}}{\leq} \ell(\theta^{(k)}),$$

- ▶ The E step brings the free energy to the likelihood.
- ▶ The M-step maximises the free energy wrt θ .
- ▶ $\mathcal{F} \leq \ell$ by Jensen – or, equivalently, from the non-negativity of KL

If the M-step is executed so that $\theta^{(k)} \neq \theta^{(k-1)}$ iff \mathcal{F} increases, then the overall EM iteration will step to a new value of θ iff the likelihood increases.

Free-energy-based variational approximation

What if finding expected sufficient stats under $P(\mathcal{Z}|\mathcal{X}, \theta)$ is computationally **intractable**?

For the **generalised EM** algorithm, we argued that intractable maximisations could be replaced by gradient M-steps.

- ▶ Each step increases the likelihood.
- ▶ A fixed point of the gradient M-step must be at a mode of the expected log-joint.

For the E-step we could:

- ▶ **Parameterise** $q = q_\rho(\mathcal{Z})$ and take a gradient step in ρ .
- ▶ **Assume** some simplified form for q , usually **factored**: $q = \prod_i q_i(\mathcal{Z}_i)$ where \mathcal{Z}_i partition \mathcal{Z} , and maximise within this form.

In either case, we choose q from within a limited set \mathcal{Q} :

VE step: maximise $\mathcal{F}(q, \theta)$ wrt **constrained** latent distribution given parameters:

$$q^{(k)}(\mathcal{Z}) := \operatorname{argmax}_{q(\mathcal{Z}) \in \mathcal{Q} \leftarrow \text{Constraint}} \mathcal{F}(q(\mathcal{Z}), \theta^{(k-1)}).$$

M step: unchanged

$$\theta^{(k)} := \operatorname{argmax}_{\theta} \mathcal{F}(q^{(k)}(\mathcal{Z}), \theta) = \operatorname{argmax}_{\theta} \int q^{(k)}(\mathcal{Z}) \log p(\mathcal{Z}, \mathcal{X}|\theta) d\mathcal{Z},$$

Unlike in GEM, the fixed point may not be at an unconstrained optimum of \mathcal{F} .

KL divergence

Recall that

$$\begin{aligned} \mathcal{F}(q, \theta) &= \langle \log P(\mathcal{X}, \mathcal{Z}|\theta) \rangle_{q(\mathcal{Z})} + \mathbf{H}[q] \\ &= \langle \log P(\mathcal{X}|\theta) + \log P(\mathcal{Z}|\mathcal{X}, \theta) \rangle_{q(\mathcal{Z})} - \langle \log q(\mathcal{Z}) \rangle_{q(\mathcal{Z})} \\ &= \langle \log P(\mathcal{X}|\theta) \rangle_{q(\mathcal{Z})} - \mathbf{KL}[q||P(\mathcal{Z}|\mathcal{X}, \theta)]. \end{aligned}$$

Thus,

E step maximise $\mathcal{F}(q, \theta)$ wrt the distribution over latents, given parameters:

$$q^{(k)}(\mathcal{Z}) := \operatorname{argmax}_{q(\mathcal{Z}) \in \mathcal{Q}} \mathcal{F}(q(\mathcal{Z}), \theta^{(k-1)}).$$

is equivalent to:

E step minimise $\mathbf{KL}[q||p(\mathcal{Z}|\mathcal{X}, \theta)]$ wrt distribution over latents, given parameters:

$$q^{(k)}(\mathcal{Z}) := \operatorname{argmin}_{q(\mathcal{Z}) \in \mathcal{Q}} \int q(\mathcal{Z}) \log \frac{q(\mathcal{Z})}{p(\mathcal{Z}|\mathcal{X}, \theta^{(k-1)})} d\mathcal{Z}$$

So, in each E step, the algorithm is trying to find the best approximation to $P(\mathcal{Z}|\mathcal{X})$ in \mathcal{Q} in a KL sense. This is related to ideas in *information geometry*. It also suggests generalisations to other distance measures.

What do we lose?

What does restricting q to \mathcal{Q} cost us?

- ▶ Recall that the free-energy is bounded above by Jensen:

$$\mathcal{F}(q, \theta) \leq \ell(\theta^{\text{ML}})$$

Thus, as long as every step increases \mathcal{F} , **convergence is still guaranteed**.

- ▶ But, since $P(\mathcal{Z}|\mathcal{X}, \theta^{(k)})$ may not lie in \mathcal{Q} , we no longer saturate the bound after the E-step. Thus, the **likelihood may not increase** on each full EM step.

$$\ell(\theta^{(k-1)}) \stackrel{\text{E step}}{\not\leq} \mathcal{F}(q^{(k)}, \theta^{(k-1)}) \stackrel{\text{M step}}{\leq} \mathcal{F}(q^{(k)}, \theta^{(k)}) \stackrel{\text{Jensen}}{\leq} \ell(\theta^{(k)}),$$

- ▶ This means we **may not (and usually won't) converge to a maximum** of ℓ .

The hope is that by *increasing a lower bound* on ℓ we will find a decent solution. [Note that if $P(\mathcal{Z}|\mathcal{X}, \theta^{\text{ML}}) \in \mathcal{Q}$, then θ^{ML} is a fixed point of the variational algorithm.]

Factored Variational E-step

The most common form of variational approximation partitions \mathcal{Z} into disjoint sets \mathcal{Z}_i with

$$\mathcal{Q} = \{q \mid q(\mathcal{Z}) = \prod_i q_i(\mathcal{Z}_i)\}.$$

In this case the E-step is itself iterative:

(Factored VE step)_i: maximise $\mathcal{F}(q, \theta)$ wrt $q_i(\mathcal{Z}_i)$ given other q_j and parameters:

$$q_i^{(k)}(\mathcal{Z}_i) := \operatorname{argmax}_{q_i(\mathcal{Z}_i)} \mathcal{F}(q_i(\mathcal{Z}_i) \prod_{j \neq i} q_j(\mathcal{Z}_j), \theta^{(k-1)}).$$

- ▶ q_i updates iterated to convergence to "complete" VE-step.
- ▶ In fact, every (VE)_i-step separately increases \mathcal{F} , so **any** schedule of (VE)_i- and M-steps will converge. Choice can be dictated by practical issues (rarely efficient to fully converge E-step before updating parameters).

Factored Variational E-step

The Factored Variational E-step has a general form.

The free energy is:


$$\begin{aligned} \mathcal{F}\left(\prod_j q_j(\mathcal{Z}_j), \theta^{(k-1)}\right) &= \left\langle \log P(\mathcal{X}, \mathcal{Z} | \theta^{(k-1)}) \right\rangle_{\prod_j q_j(\mathcal{Z}_j)} + \mathbf{H}\left[\prod_j q_j(\mathcal{Z}_j)\right] \\ &= \int d\mathcal{Z}_i q_i(\mathcal{Z}_i) \left\langle \log P(\mathcal{X}, \mathcal{Z} | \theta^{(k-1)}) \right\rangle_{\prod_{j \neq i} q_j(\mathcal{Z}_j)} + \mathbf{H}[q_i] + \sum_{j \neq i} \mathbf{H}[q_j] \end{aligned}$$

Now, taking the variational derivative of the Lagrangian (enforcing normalisation of q_i):

$$\begin{aligned} \frac{\delta}{\delta q_i} \left(\mathcal{F} + \lambda \left(\int q_i - 1 \right) \right) &= \left\langle \log P(\mathcal{X}, \mathcal{Z} | \theta^{(k-1)}) \right\rangle_{\prod_{j \neq i} q_j(\mathcal{Z}_j)} - \log q_i(\mathcal{Z}_i) - \frac{q_i(\mathcal{Z}_i)}{q_i(\mathcal{Z}_i)} + \lambda \\ (= 0) \Rightarrow q_i(\mathcal{Z}_i) &\propto \exp \left\langle \log P(\mathcal{X}, \mathcal{Z} | \theta^{(k-1)}) \right\rangle_{\prod_{j \neq i} q_j(\mathcal{Z}_j)} \end{aligned}$$

In general, this depends only on the expected sufficient statistics under q_j . Thus, again, we don't actually need the *entire* distributions, just the **relevant** expectations (now for approximate inference as well as learning).

Mean-field FHMM



$$q(\mathbf{s}_{1:T}^{1:M}) = \prod_{m,t} q_t^m(\mathbf{s}_t^m)$$

$$\begin{aligned} q_t^m(\mathbf{s}_t^m) &\propto \exp \left\langle \log P(\mathbf{s}_{1:T}^{1:M}, \mathbf{x}_{1:T}) \right\rangle_{\prod_{-(m,t)} q_{t'}^{m'}(\mathbf{s}_{t'}^{m'})} \\ &= \exp \left\langle \sum_{\mu} \sum_{\tau} \log P(\mathbf{s}_{\tau}^{\mu} | \mathbf{s}_{\tau-1}^{\mu}) + \sum_{\tau} \log P(\mathbf{x}_{\tau} | \mathbf{s}_{\tau}^{1:M}) \right\rangle_{\prod_{-(m,t)} q_{t'}^{m'}} \\ &\propto \exp \left[\underbrace{\langle \log P(\mathbf{s}_t^m | \mathbf{s}_{t-1}^m) \rangle_{q_{t-1}^m} + \langle \log P(\mathbf{x}_t | \mathbf{s}_t^{1:M}) \rangle_{\prod_{m'} q_{t'}^{m'}}}_{\alpha_t^m(i) \propto e^{\sum_j \log \Phi_{ij}^m q_{t-1}^m(j)} \cdot e^{\langle \log A_i(\mathbf{x}_t) \rangle_{q_t^{-m}}} + \underbrace{\langle \log P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m) \rangle_{q_{t+1}^m}}_{\beta_t^m(i) \propto e^{\sum_j \log \Phi_{ij}^m q_{t+1}^m(j)}} \right] \\ \text{Cf. forward-backward:} \quad &\alpha_t(i) \propto \sum_j \alpha_{t-1}(j) \Phi_{ji} \cdot A_i(\mathbf{x}_t) \quad \beta_t(i) \propto \sum_j \Phi_{ij} A_j(\mathbf{x}_{t+1}) \beta_{t+1}(j) \end{aligned}$$

Mean-field approximations

If $\mathcal{Z}_i = \mathbf{z}_i$ (i.e., q is factored over all variables) then the variational technique is often called a "mean field" approximation.

- Suppose $P(\mathcal{X}, \mathcal{Z})$ has sufficient statistics that are **separable** in the latent variables: e.g. the Boltzmann machine

$$P(\mathcal{X}, \mathcal{Z}) = \frac{1}{Z} \exp \left(\sum_{ij} W_{ij} s_i s_j + \sum_i b_i s_i \right)$$

with some $s_i \in \mathcal{Z}$ and others observed.


- Expectations wrt a fully-factored q distribute over all $s_i \in \mathcal{Z}$

$$\langle \log P(\mathcal{X}, \mathcal{Z}) \rangle_{\prod q_i} = \sum_{ij} W_{ij} \langle s_i \rangle_{q_i} \langle s_j \rangle_{q_j} + \sum_i b_i \langle s_i \rangle_{q_i}$$

(where q_i for $s_i \in \mathcal{X}$ is a delta function on the observed value).

- Thus, we can update each q_i in turn given the **means** (or, in general, mean sufficient statistics) of the others.
- Each variable sees the **mean field** imposed by its neighbours, and we update these fields until they all agree.

Mean-field FHMM




$$q(\mathbf{s}_{1:T}^{1:M}) = \prod_{m,t} q_t^m(\mathbf{s}_t^m)$$

$$\begin{aligned} q_t^m(\mathbf{s}_t^m) &\propto \exp \left[\underbrace{\langle \log P(\mathbf{s}_t^m | \mathbf{s}_{t-1}^m) \rangle_{q_{t-1}^m} + \langle \log P(\mathbf{x}_t | \mathbf{s}_t^{1:M}) \rangle_{\prod_{m'} q_{t'}^{m'}}}_{\alpha_t^m(i) \propto e^{\sum_j \log \Phi_{ij}^m q_{t-1}^m(j)} \cdot e^{\langle \log A_i(\mathbf{x}_t) \rangle_{q_t^{-m}}} + \underbrace{\langle \log P(\mathbf{s}_{t+1}^m | \mathbf{s}_t^m) \rangle_{q_{t+1}^m}}_{\beta_t^m(i) \propto e^{\sum_j \log \Phi_{ij}^m q_{t+1}^m(j)}} \right] \\ \text{Cf. forward-backward:} \quad &\alpha_t(i) \propto \sum_j \alpha_{t-1}(j) \Phi_{ji} \cdot A_i(\mathbf{x}_t) \quad \beta_t(i) \propto \sum_j \Phi_{ij} A_j(\mathbf{x}_{t+1}) \beta_{t+1}(j) \end{aligned}$$

- Yields a message-passing algorithm like forward-backward
- Updates depend only on immediate neighbours in chain
- Chains couple only through joint output
- Multiple passes; messages depend on (approximate) marginals
- Evidence does not appear explicitly in backward message (cf Kalman smoothing)


Structured variational approximation

- ▶ $q(\mathcal{Z})$ need not be completely factorized.
- ▶ For example, suppose \mathcal{Z} can be partitioned into sets \mathcal{Z}_1 and \mathcal{Z}_2 such that computing the expected sufficient statistics under $P(\mathcal{Z}_1|\mathcal{Z}_2, \mathcal{X})$ and $P(\mathcal{Z}_2|\mathcal{Z}_1, \mathcal{X})$ would be tractable.
- ⇒ Then the factored approximation $q(\mathcal{Z}) = q(\mathcal{Z}_1)q(\mathcal{Z}_2)$ is tractable.
- ▶ In particular, any factorisation of $q(\mathcal{Z})$ into a product of distributions on **trees**, yields a tractable approximation.



Messages on an arbitrary graph

Consider a DAG:



$$P(\mathcal{X}, \mathcal{Z}) = \prod_k P(V_k | \text{pa}(V_k))$$

and let $q(\mathcal{Z}) = \prod_i q_i(\mathcal{Z}_i)$ for disjoint sets $\{\mathcal{Z}_i\}$.


We have that the VE update for q_i is given by $q_i^*(\mathcal{Z}_i) \propto \exp(\langle \log p(\mathcal{Z}, \mathcal{X}) \rangle_{q_{-i}(\mathcal{Z})})$ where $\langle \cdot \rangle_{q_{-i}(\mathcal{Z})}$ denotes averaging with respect to $q_j(\mathcal{Z}_j)$ for all $j \neq i$

Then:

$$\begin{aligned} \log q_i^*(\mathcal{Z}_i) &= \left\langle \sum_k \log P(V_k | \text{pa}(V_k)) \right\rangle_{q_{-i}(\mathcal{Z})} + \text{const} \\ &= \sum_{j \in \mathcal{Z}_i} \langle \log P(Z_j | \text{pa}(Z_j)) \rangle_{q_{-i}(\mathcal{Z})} + \sum_{j \in \text{ch}(\mathcal{Z}_i)} \langle \log P(V_j | \text{pa}(V_j)) \rangle_{q_{-i}(\mathcal{Z})} + \text{const} \end{aligned}$$

This defines messages that are passed between nodes in the graph. Each node receives messages from its **Markov boundary**: parents, children and parents of children (all neighbours in the corresponding factor graph).

Structured FHMM



For the FHMM we can factor the chains:

$$q(\mathbf{s}_{1:T}^{1:M}) = \prod_m q^m(\mathbf{s}_{1:T}^m)$$

$$\begin{aligned} q^m(\mathbf{s}_{1:T}^m) &\propto \exp \left\langle \log P(\mathbf{s}_{1:T}^{1:M}, \mathbf{x}_{1:T}) \right\rangle_{\prod_{-m} q^{m'}(\mathbf{s}_{1:T}^{m'})} \\ &= \exp \left\langle \sum_{\mu} \sum_t \log P(s_t^\mu | s_{t-1}^\mu) + \sum_t \log P(\mathbf{x}_t | s_t^{1:M}) \right\rangle_{\prod_{-m} q^{m'}(\mathbf{s}_{1:T}^{m'})} \\ &\propto \exp \left[\sum_t \log P(s_t^m | s_{t-1}^m) + \sum_t \left\langle \log P(\mathbf{x}_t | s_t^{1:M}) \right\rangle_{\prod_{-m} q^{m'}(\mathbf{s}_{1:T}^{m'})} \right] \\ &= \prod_t P(s_t^m | s_{t-1}^m) \prod_t e^{\langle \log P(\mathbf{x}_t | s_t^{1:M}) \rangle_{\prod_{-m} q^{m'}(\mathbf{s}_{1:T}^{m'})}} \end{aligned}$$

This looks like a standard HMM joint, with a modified likelihood term \Rightarrow cycle through multiple forward-backward passes, updating likelihood terms each time.

Non-factored variational methods

The term **variational approximation** is used whenever a bound on the likelihood (or on another estimation cost function) is optimised, but does not necessarily become tight.

Many further variational approximations have been developed, including:

- ▶ parametric forms (e.g. Gaussian) for non-linear models (later lecture)
 - ▶ closed form updates in special cases
 - ▶ numerical or sampling-based computation of expectations
 - ▶ 'recognition networks' or amortisation to estimate variational parameters
- ▶ non-free-energy-based bounds (both upper and lower) on the likelihood.

We can also see **MAP**- or **zero-temperature EM** and **recognition models** as parametric forms of variational inference.

Variational Bayes

So far, we have applied Jensen's bound and factorisations to help with integrals over latent variables.

We can do the same for integrals over parameters in order to bound the log [marginal likelihood](#) or [evidence](#).

$$\begin{aligned}\log P(\mathcal{X}|\mathcal{M}) &= \log \int \int d\mathcal{Z} d\theta P(\mathcal{X}, \mathcal{Z}|\theta, \mathcal{M})P(\theta|\mathcal{M}) \\ &= \max_Q \int \int d\mathcal{Z} d\theta Q(\mathcal{Z}, \theta) \log \frac{P(\mathcal{X}, \mathcal{Z}, \theta|\mathcal{M})}{Q(\mathcal{Z}, \theta)} \\ &\geq \max_{Q_{\mathcal{Z}}, Q_{\theta}} \int \int d\mathcal{Z} d\theta Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta) \log \frac{P(\mathcal{X}, \mathcal{Z}, \theta|\mathcal{M})}{Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta)}\end{aligned}$$

The constraint that the distribution Q must [factor](#) into the product $Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta)$ leads to the **variational Bayesian EM algorithm** or just **"Variational Bayes"**.

Some call this the "Evidence Lower Bound" (ELBO). I'm not fond of that term.

Conjugate-Exponential models

Let's focus on [conjugate-exponential](#) (**CE**) latent-variable models:

- ▶ **Condition (1)**. The [joint probability](#) over *variables* is in the [exponential family](#):

$$P(\mathcal{Z}, \mathcal{X}|\theta) = f(\mathcal{Z}, \mathcal{X}) g(\theta) \exp \left\{ \phi(\theta)^T T(\mathcal{Z}, \mathcal{X}) \right\}$$

where $\phi(\theta)$ is the vector of *natural parameters*, T are *sufficient statistics*

- ▶ **Condition (2)**. The [prior](#) over *parameters* is [conjugate](#) to this joint probability:

$$P(\theta|\nu, \tau) = h(\nu, \tau) g(\theta)^\nu \exp \left\{ \phi(\theta)^T \tau \right\}$$

where ν and τ are hyperparameters of the prior.

Conjugate priors are computationally convenient and have an intuitive interpretation:

- ▶ ν : number of pseudo-observations
- ▶ τ : values of pseudo-observations

Variational Bayesian EM ...

Coordinate maximization of the VB free-energy [lower bound](#)

$$\mathcal{F}(Q_{\mathcal{Z}}, Q_{\theta}) = \int \int d\mathcal{Z} d\theta Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta) \log \frac{P(\mathcal{X}, \mathcal{Z}, \theta|\mathcal{M})}{Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta)}$$

leads to **EM-like** updates:

$$\begin{aligned}Q_{\mathcal{Z}}^*(\mathcal{Z}) &\propto \exp \langle \log P(\mathcal{Z}, \mathcal{X}|\theta) \rangle_{Q_{\theta}(\theta)} && E\text{-like step} \\ Q_{\theta}^*(\theta) &\propto P(\theta) \exp \langle \log P(\mathcal{Z}, \mathcal{X}|\theta) \rangle_{Q_{\mathcal{Z}}(\mathcal{Z})} && M\text{-like step}\end{aligned}$$

Maximizing \mathcal{F} is equivalent to minimizing KL-divergence between the *approximate posterior*, $Q(\theta)Q(\mathcal{Z})$ and the *true posterior*, $P(\theta, \mathcal{Z}|\mathcal{X})$.

$$\begin{aligned}\log P(\mathcal{X}) - \mathcal{F}(Q_{\mathcal{Z}}, Q_{\theta}) &= \log P(\mathcal{X}) - \int \int d\mathcal{Z} d\theta Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta) \log \frac{P(\mathcal{X}, \mathcal{Z}, \theta)}{Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta)} \\ &= \int \int d\mathcal{Z} d\theta Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta) \log \frac{Q_{\mathcal{Z}}(\mathcal{Z})Q_{\theta}(\theta)}{P(\mathcal{Z}, \theta|\mathcal{X})} = KL(Q||P)\end{aligned}$$

Conjugate-Exponential examples

In the **CE** family:

- ▶ Gaussian mixtures
- ▶ factor analysis, probabilistic PCA
- ▶ hidden Markov models and factorial HMMs
- ▶ linear dynamical systems and switching models
- ▶ discrete-variable belief networks

Other as yet undreamt-of models combinations of Gaussian, Gamma, Poisson, Dirichlet, Wishart, Multinomial and others.

Not in the **CE** family:

- ▶ Boltzmann machines, MRFs (no simple conjugacy)
- ▶ logistic regression (no simple conjugacy)
- ▶ sigmoid belief networks (not exponential)
- ▶ independent components analysis (not exponential)

Note: one can often approximate such models with a suitable choice from the **CE** family.

Conjugate-exponential VB

Given an iid data set $\mathcal{D} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, if the model is **CE** then:

- ▶ $Q_\theta(\theta)$ is also **conjugate**, i.e.

$$\begin{aligned} Q_\theta(\theta) &\propto P(\theta) \exp \left\langle \sum_i \log P(\mathbf{z}_i, \mathbf{x}_i | \theta) \right\rangle_{Q_Z} \\ &= h(\nu, \tau) g(\theta)^\nu e^{\phi(\theta)^\top \tau} g(\theta)^n e^{\langle \log r(\mathcal{Z}, \mathcal{X}) \rangle_{Q_Z}} e^{\phi(\theta)^\top \langle \sum_i \mathbb{T}(\mathbf{z}_i, \mathbf{x}_i) \rangle_{Q_Z}} \\ &\propto h(\tilde{\nu}, \tilde{\tau}) g(\theta)^{\tilde{\nu}} e^{\phi(\theta)^\top \tilde{\tau}} \end{aligned}$$

with $\tilde{\nu} = \nu + n$ and $\tilde{\tau} = \tau + \sum_i \langle \mathbb{T}(\mathbf{z}_i, \mathbf{x}_i) \rangle_{Q_Z} \Rightarrow$ **only need to track $\tilde{\nu}, \tilde{\tau}$.**

- ▶ $Q_Z(\mathcal{Z}) = \prod_{i=1}^n Q_{\mathbf{z}_i}(\mathbf{z}_i)$ takes the **same form** as in the E-step of regular EM

$$\begin{aligned} Q_{\mathbf{z}_i}(\mathbf{z}_i) &\propto \exp \langle \log P(\mathbf{z}_i, \mathbf{x}_i | \theta) \rangle_{Q_\theta} \\ &\propto f(\mathbf{z}_i, \mathbf{x}_i) e^{\langle \phi(\theta) \rangle_{Q_\theta}^\top \mathbb{T}(\mathbf{z}_i, \mathbf{x}_i)} = P(\mathbf{z}_i | \mathbf{x}_i, \bar{\phi}(\theta)) \end{aligned}$$

with **natural parameters** $\bar{\phi}(\theta) = \langle \phi(\theta) \rangle_{Q_\theta} \Rightarrow$ **inference unchanged from regular EM.**

VB and model selection

- ▶ Variational Bayesian EM yields an **approximate posterior** Q_θ over model parameters.
- ▶ It also yields an **optimised lower bound** on the model evidence

$$\max_{\mathcal{F}_M} \mathcal{F}_M(Q_Z, Q_\theta) \leq P(\mathcal{D} | \mathcal{M})$$

- ▶ These lower bounds can be **compared** amongst models to learn the right (structure, connectivity ... of the) model
- ▶ If a continuous domain of models is specified by a hyperparameter η , then the VB free energy depends on that parameter:

$$\mathcal{F}(Q_Z, Q_\theta, \eta) = \iint d\mathcal{Z} d\theta Q_Z(\mathcal{Z}) Q_\theta(\theta) \log \frac{P(\mathcal{X}, \mathcal{Z}, \theta | \eta)}{Q_Z(\mathcal{Z}) Q_\theta(\theta)} \leq P(\mathcal{X} | \eta)$$

A **hyper-M** step maximises the current bound wrt η :

$$\eta \leftarrow \operatorname{argmax}_\eta \iint d\mathcal{Z} d\theta Q_Z(\mathcal{Z}) Q_\theta(\theta) \log P(\mathcal{X}, \mathcal{Z}, \theta | \eta)$$

The Variational Bayesian EM algorithm

EM for MAP estimation

Goal: maximize $P(\theta | \mathcal{X}, m)$ wrt θ

E Step: compute

$$Q_Z(\mathcal{Z}) \leftarrow p(\mathcal{Z} | \mathcal{X}, \theta)$$

M Step:

$$\theta \leftarrow \operatorname{argmax}_\theta \int d\mathcal{Z} Q_Z(\mathcal{Z}) \log P(\mathcal{Z}, \mathcal{X}, \theta)$$

Variational Bayesian EM

Goal: maximise bound on $P(\mathcal{X} | m)$ wrt Q_θ

VB-E Step: compute

$$Q_Z(\mathcal{Z}) \leftarrow p(\mathcal{Z} | \mathcal{X}, \bar{\phi})$$

VB-M Step:

$$Q_\theta(\theta) \leftarrow \exp \int d\mathcal{Z} Q_Z(\mathcal{Z}) \log P(\mathcal{Z}, \mathcal{X}, \theta)$$

Properties:

- ▶ Reduces to the EM algorithm if $Q_\theta(\theta) = \delta(\theta - \theta^*)$.
- ▶ \mathcal{F}_m increases monotonically, and incorporates the model complexity penalty.
- ▶ Analytical parameter distributions (but not constrained to be Gaussian).
- ▶ VB-E step has same complexity as corresponding E step.
- ▶ We can use the junction tree, belief propagation, Kalman filter, etc, algorithms in the VB-E step of VB-EM, but **using expected natural parameters**, $\bar{\phi}$.

ARD for unsupervised learning

Recall that ARD (automatic relevance determination) was a hyperparameter method to select relevant or useful inputs in regression.

- ▶ A similar idea used with variational Bayesian methods can learn a **latent dimensionality**.
- ▶ Consider factor analysis:

$$\mathbf{x} \sim \mathcal{N}(\Lambda \mathbf{z}, \Psi) \quad \mathbf{z} \sim \mathcal{N}(0, I) \quad \text{with a column-wise prior } \Lambda_{:i} \sim \mathcal{N}(0, \alpha_i^{-1} I)$$

- ▶ The VB free energy is

$$\mathcal{F}(Q_Z(\mathcal{Z}), Q_\Lambda(\Lambda), \Psi, \alpha) = \langle \log P(\mathcal{X}, \mathcal{Z} | \Lambda, \Psi) + \log P(\Lambda | \alpha) + \log P(\Psi) \rangle_{Q_Z Q_\Lambda} + \dots$$

and so hyperparameter optimisation requires

$$\alpha \leftarrow \operatorname{argmax}_\alpha \langle \log P(\Lambda | \alpha) \rangle_{Q_\Lambda}$$

- ▶ Now Q_Λ is Gaussian, with the same form as in linear regression, but with **expected moments of \mathbf{z}** appearing in place of the inputs.
- ▶ Optimisation wrt the distributions, Ψ and α in turn causes some α_i to diverge as in regression ARD.
- ▶ In this case, these parameters select "relevant" **latent dimensions**, effectively learning the dimensionality of \mathbf{z} .

Augmented Variational Methods

In our examples so far, the approximate variational distribution has been over the “natural” latent variables (and parameters) of the generative model.

Sometimes it may be useful to introduce additional latent variables, solely to achieve computational tractability.

Two examples are GP regression and the GPLVM.

Variational Sparse GP approximations

Write F for the (smooth) GP function values that underlie Y (so $Y \sim \mathcal{N}(F, \sigma^2 I)$).

Introduce **latent** measurements U at inputs Z (and integrate over U).

The likelihood can be written

$$P(Y|X) = \iint dF dU P(Y, F, U|X, Z) = \iint dF dU P(Y|F)P(F|U, X, Z)P(U|Z)$$

Now, both U and F are latent, so we introduce a variational distribution $q(F, U)$ to form a free-energy.

$$\mathcal{F}(q(F, U), \theta) = \left\langle \log \frac{P(Y|F)P(F|U, X, Z)P(U|Z)}{q(F, U)} \right\rangle_{q(F, U)}$$

Now, choose the variational form $q(F, U) = P(F|U, X, Z)q(U)$. That is, fix $F|U$ without reference to Y – so information about Y will need to be “compressed” into $q(U)$.

Then

$$\begin{aligned} \mathcal{F}(q(F, U), \theta, Z) &= \left\langle \log \frac{P(Y|F)P(F|U, X, Z)P(U|Z)}{P(F|U, X, Z)q(U)} \right\rangle_{P(F|U)q(U)} \\ &= \left\langle \log P(Y|F) \right\rangle_{P(F|U)} + \log P(U|Z) - \log q(U) \end{aligned}$$

Sparse GP approximations

GP predictions:

$$y'|X, Y, \mathbf{x}' \sim \mathcal{N} \left(K_{\mathbf{x}'X} (K_{XX} + \sigma^2 I)^{-1} Y, K_{\mathbf{x}'\mathbf{x}'} - K_{\mathbf{x}'X} (K_{XX} + \sigma^2 I)^{-1} K_{X\mathbf{x}'} + \sigma^2 \right)$$

Evidence (for learning kernel hyperparameters):

$$\log P(Y|X) = -\frac{1}{2} \log |2\pi(K_{XX} + \sigma^2 I)| - \frac{1}{2} Y (K_{XX} + \sigma^2 I)^{-1} Y^\top$$

Computing either form requires inverting the $N \times N$ matrix K_{XX} , in $\mathcal{O}(N^3)$ time.

One proposal to make this more efficient is to find (or select) a smaller set of possibly fictitious measurements U at inputs Z such that

$$P(y'|Z, U, \mathbf{x}') \approx P(y'|X, Y, \mathbf{x}').$$

What values should U and Z take?

Variational Sparse GP approximations

$$\mathcal{F}(q(U), \theta, Z) = \left\langle \log P(Y|F) \right\rangle_{P(F|U)} + \log P(U|Z) - \log q(U) \Big|_{q(U)}$$

Now $P(F|U)$ is fixed by the generative model (rather than being subject to free optimisation). So we can evaluate that expectation:

$$\begin{aligned} \left\langle \log P(Y|F) \right\rangle_{P(F|U)} &= \left\langle -\frac{1}{2} \log |2\pi\sigma^2 I| - \frac{1}{2\sigma^2} \text{Tr} \left[(Y - F)(Y - F)^\top \right] \right\rangle_{P(F|U)} \\ &= -\frac{1}{2} \log |2\pi\sigma^2 I| - \frac{1}{2\sigma^2} \text{Tr} \left[(Y - \langle F \rangle_{P(F|U)})(Y - \langle F \rangle_{P(F|U)})^\top \right] - \frac{1}{2\sigma^2} \text{Tr} [\Sigma_{F|U}] \\ &= \log \mathcal{N}(Y | K_{XZ} K_{ZZ}^{-1} U, \sigma^2 I) - \frac{1}{2\sigma^2} \text{Tr} [K_{XX} - K_{XZ} K_{ZZ}^{-1} K_{ZX}] \end{aligned}$$

So,

$$\begin{aligned} \mathcal{F}(q(U), \theta, Z) &= \left\langle \log \mathcal{N}(Y | K_{XZ} K_{ZZ}^{-1} U, \sigma^2 I) + \log P(U|Z) - \log q(U) \right\rangle_{q(U)} \\ &\quad - \frac{1}{2\sigma^2} \text{Tr} [K_{XX} - K_{XZ} K_{ZZ}^{-1} K_{ZX}]. \end{aligned}$$

Variational Sparse GP approximations

$$\mathcal{F}(q(U), \theta, Z) = \left\langle \log \frac{\mathcal{N}(Y|K_{XZ}K_{ZZ}^{-1}U, \sigma^2 I) P(U|Z)}{q(U)} \right\rangle_{q(U)} - \frac{1}{2\sigma^2} \text{Tr} [K_{XX} - K_{XZ}K_{ZZ}^{-1}K_{ZX}].$$

The expectation is the free energy of a PPCA-like model with normal prior $U \sim \mathcal{N}(0, K_{ZZ})$ and loading matrix $K_{XZ}K_{ZZ}^{-1}$. The maximum of this free energy is the log-likelihood (achieved with q equal to the posterior under the PPCA-like model).

This gives

$$\mathcal{F}(q^*(U), \theta, Z) = \log \mathcal{N}(Y|0, K_{XZ}K_{ZZ}^{-1}K_{ZZ}K_{ZZ}^{-1}K_{ZX} + \sigma^2 I) - \frac{1}{2\sigma^2} \text{Tr} [K_{XX} - K_{XZ}K_{ZZ}^{-1}K_{ZX}].$$

Note that we have eliminated all terms in K_{XX}^{-1} .

We can optimise the free energy numerically with respect to Z and θ to adjust the GP prior and quality of variational approximation.

A similar approach can be used to learn X if they are unobserved (*i.e.* in the GPLVM).

Assume $q(X, F, U) = q(X)P(F|X, U)q(U)$. Then $\mathcal{F} = \langle \log P(Y, F, U|X) \log P(X) \rangle_{q(U)q(X)}$ which simplifies into tractable components in much the same way as above.

A few references

- ▶ Jordan, Ghahramani, Jaakkola, Saul, 1999. [An introduction to variational methods for graphical models](#). *Machine Learning* **37**:183–233.
- ▶ Attias, 2000. [A variational Bayesian framework for graphical models](#). *NIPS 12*. <http://www.gatsby.ucl.ac.uk/publications/papers/03-2000.ps>
- ▶ Beal, 2003. [Variational algorithms for approximate Bayesian inference](#). *PhD thesis*, Gatsby Unit, UCL. <http://www.cse.buffalo.edu/faculty/mbeal/thesis/>
- ▶ Winn, 2003. [Variational message passing and its applications](#). *PhD thesis*, Cambridge. <http://johnwinn.org/Publications/Thesis.html>; also **VIBES** software for conjugate-exponential graphs.

Some complexities:

- ▶ MacKay, 2001. [A problem with variational free energy minimization](#). <http://www.inference.phy.cam.ac.uk/mackay/minima.pdf>
- ▶ Turner, MS, 2011. [Two problems with variational expectation maximisation for time-series models](#). In Barber, Cemgil, Chiappa, eds., *Bayesian Time Series Models*. <http://www.gatsby.ucl.ac.uk/~maneesh/papers/turner-sahani-2010-ildn.pdf>
- ▶ Berkes, Turner, MS, 2008. [On sparsity and overcompleteness in image models](#). *NIPS 20*. <http://www.gatsby.ucl.ac.uk/~maneesh/papers/berkes-etal-2008-nips.pdf>
- ▶ Giordano, Broderick, Jordan, 2015. [Linear response methods for accurate covariance estimates from mean field variational Bayes](#). *NIPS*