

Probabilistic & Unsupervised Learning Approximate Inference

Expectation Propagation


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

Nonlinear state-space model (NLSSM)



$$\mathbf{z}_{t+1} = f(\mathbf{z}_t, \mathbf{u}_t) + \mathbf{w}_t$$


$$\mathbf{x}_t = g(\mathbf{z}_t, \mathbf{u}_t) + \mathbf{v}_t$$

$\mathbf{w}_t, \mathbf{v}_t$ usually still Gaussian.

Extended Kalman Filter (EKF): linearise nonlinear functions about current estimate, $\hat{\mathbf{z}}_t^f$:

$$\mathbf{z}_{t+1} \approx \underbrace{f(\hat{\mathbf{z}}_t^f, \mathbf{u}_t)}_{\tilde{\mathbf{B}}_t \mathbf{u}_t} + \underbrace{\frac{\partial f}{\partial \mathbf{z}_t} \Big|_{\hat{\mathbf{z}}_t^f}}_{\tilde{\mathbf{A}}_t} (\mathbf{z}_t - \hat{\mathbf{z}}_t^f) + \mathbf{w}_t$$

$$\mathbf{x}_t \approx \underbrace{g(\hat{\mathbf{z}}_t^{t-1}, \mathbf{u}_t)}_{\tilde{\mathbf{D}}_t \mathbf{u}_t} + \underbrace{\frac{\partial g}{\partial \mathbf{z}_t} \Big|_{\hat{\mathbf{z}}_t^{t-1}}}_{\tilde{\mathbf{C}}_t} (\mathbf{z}_t - \hat{\mathbf{z}}_t^{t-1}) + \mathbf{v}_t$$



Run the Kalman filter (smoother) on non-stationary linearised system $(\tilde{\mathbf{A}}_t, \tilde{\mathbf{B}}_t, \tilde{\mathbf{C}}_t, \tilde{\mathbf{D}}_t)$:

- ▶ Adaptively approximates non-Gaussian messages by Gaussians.
- ▶ Local linearisation depends on central point of distribution \Rightarrow approximation degrades with increased state uncertainty. May work acceptably for close-to-linear systems.

Can base EM-like algorithm on EKF/EKS (or alternatives).

Intractabilities and approximations

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

Not a complete list!

Other message approximations

Consider the forward messages on a latent chain:

$$P(\mathbf{z}_t | \mathbf{x}_{1:t}) = \frac{1}{Z} P(\mathbf{x}_t | \mathbf{z}_t) \int d\mathbf{z}_{t-1} P(\mathbf{z}_t | \mathbf{z}_{t-1}) P(\mathbf{z}_{t-1} | \mathbf{x}_{1:t-1})$$

We want to approximate the messages to retain a tractable form (i.e. Gaussian).

$$\tilde{P}(\mathbf{z}_t | \mathbf{x}_{1:t}) \approx \frac{1}{Z} P(\mathbf{x}_t | \mathbf{z}_t) \int d\mathbf{z}_{t-1} \underbrace{P(\mathbf{z}_t | \mathbf{z}_{t-1})}_{\mathcal{N}(f(\mathbf{z}_{t-1}), \mathbf{Q})} \underbrace{\tilde{P}(\mathbf{z}_{t-1} | \mathbf{x}_{1:t-1})}_{\mathcal{N}(\hat{\mathbf{z}}_{t-1}, \mathbf{V}_{t-1})}$$

- ▶ Linearisation at the peak (EKF) is only one approach.
- ▶ Laplace filter: use mode and curvature of integrand.
- ▶ Sigma-point (“unscented”) filter:
 - ▶ Evaluate $f(\hat{\mathbf{z}}_{t-1}), f(\hat{\mathbf{z}}_{t-1} \pm \sqrt{\lambda \mathbf{v}})$ for eigenvalues, eigenvectors $\hat{\mathbf{V}}_{t-1} \mathbf{v} = \lambda \mathbf{v}$.
 - ▶ “Fit” Gaussian to these $2K + 1$ points.
 - ▶ Equivalent to numerical evaluation of mean and covariance by Gaussian quadrature.
 - ▶ One form of “Assumed Density Filtering” and EP.
- ▶ Parametric variational: $\text{argmin}_{\mathcal{N}(\hat{\mathbf{z}}_t, \hat{\mathbf{V}}_t)} \mathbf{KL}[\mathcal{N}(\hat{\mathbf{z}}_t, \hat{\mathbf{V}}_t) \parallel \int d\mathbf{z}_{t-1} \dots]$. Requires Gaussian expectations of $\log f \Rightarrow$ (usually) Monte-Carlo integration (later lecture).
- ▶ The other KL: $\text{argmin}_{\mathcal{N}(\hat{\mathbf{z}}_t, \hat{\mathbf{V}}_t)} \mathbf{KL}[\int d\mathbf{z}_{t-1} \dots \parallel \mathcal{N}(\hat{\mathbf{z}}_t, \hat{\mathbf{V}}_t)]$ needs only first and second moments of nonlinear message \Rightarrow EP.

Variational learning

Free energy:

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

E-steps:

▶ Exact EM: $q(\mathcal{Z}) = \underset{q}{\operatorname{argmax}} \mathcal{F} = P(\mathcal{Z}|\mathcal{X}, \theta)$

▶ Saturates bound: converges to local maximum of likelihood.

▶ (Factored) variational approximation:

$$q(\mathcal{Z}) = \underset{q_1(\mathcal{Z}_1)q_2(\mathcal{Z}_2)}{\operatorname{argmax}} \mathcal{F} = \underset{q_1(\mathcal{Z}_1)q_2(\mathcal{Z}_2)}{\operatorname{argmin}} \mathbf{KL}[q_1(\mathcal{Z}_1)q_2(\mathcal{Z}_2)\|P(\mathcal{Z}|\mathcal{X}, \theta)]$$

▶ Increases bound: converges, but not necessarily to ML.

▶ Other approximations: $q(\mathcal{Z}) \approx P(\mathcal{Z}|\mathcal{X}, \theta)$

▶ Usually no guarantees, but if learning converges it may be more accurate than the factored approximation

The other KL

What about the 'other' KL ($q = \operatorname{argmin} \mathbf{KL}[P\|q]$)?

For a factored approximation the (clique) marginals obtained by minimising this KL are correct:

$$\begin{aligned} \underset{q_i}{\operatorname{argmin}} \mathbf{KL} \left[P(\mathcal{Z}|\mathcal{X}) \left\| \prod q_i(\mathcal{Z}_i|\mathcal{X}) \right. \right] &= \underset{q_i}{\operatorname{argmin}} - \int d\mathcal{Z} P(\mathcal{Z}|\mathcal{X}) \log \prod q_i(\mathcal{Z}_i|\mathcal{X}) \\ &= \underset{q_i}{\operatorname{argmin}} - \sum_j \int d\mathcal{Z} P(\mathcal{Z}|\mathcal{X}) \log q_j(\mathcal{Z}_j|\mathcal{X}) \\ &= \underset{q_i}{\operatorname{argmin}} - \int d\mathcal{Z}_i P(\mathcal{Z}_i|\mathcal{X}) \log q_i(\mathcal{Z}_i|\mathcal{X}) \\ &= P(\mathcal{Z}_i|\mathcal{X}) \end{aligned}$$

and the marginals are what we need for learning (although if factored over disjoint sets as in the variational approximation some cliques will be missing).

Perversely, this means finding the best q for this KL is intractable!

But it raises the hope that **approximate** minimisation might still yield useful results.

Approximating the posterior

Linearisation (or local Laplace, sigma-point and other such approaches) seem *ad hoc*. A more principled approach might look for an approximate q that is **closest** to P in some sense.

$$q = \underset{q \in \mathcal{Q}}{\operatorname{argmin}} D(P \leftrightarrow q)$$

Open choices:

- ▶ form of the metric D
- ▶ nature of the constraint space \mathcal{Q}

▶ Variational methods: $D = \mathbf{KL}[q\|P]$.

▶ Choosing $\mathcal{Q} = \{\text{tree-factored distributions}\}$ leads to efficient message passing.

▶ Can we use other divergences?

Approximate optimisation

The posterior distribution in a graphical model is a (normalised) product of factors:

$$P(\mathcal{Z}|\mathcal{X}) = \frac{P(\mathcal{Z}, \mathcal{X})}{P(\mathcal{X})} = \frac{1}{Z} \prod_i P(\mathcal{Z}_i | \text{pa}(\mathcal{Z}_i)) \propto \prod_{i=1}^N f_i(\mathcal{Z}_i)$$

where the \mathcal{Z}_i are not necessarily disjoint. In the language of EP the f_i are called **sites**.

Consider q with the **same** factorisation, but potentially approximated sites:

$q(\mathcal{Z}) \stackrel{\text{def}}{=} \prod_{i=1}^N \tilde{f}_i(\mathcal{Z}_i)$. We would like to minimise (at least in some sense) $\mathbf{KL}[P\|q]$.

Possible optimisations:

$$\min_{\{\tilde{f}_i\}} \mathbf{KL} \left[\prod_{i=1}^N f_i(\mathcal{Z}_i) \left\| \prod_{i=1}^N \tilde{f}_i(\mathcal{Z}_i) \right. \right] \quad (\text{global: intractable})$$

$$\min_{\tilde{f}_i} \mathbf{KL} \left[f_i(\mathcal{Z}_i) \left\| \tilde{f}_i(\mathcal{Z}_i) \right. \right] \quad (\text{local, fixed: simple, inaccurate})$$

$$\min_{\tilde{f}_i} \mathbf{KL} \left[f_i(\mathcal{Z}_i) \prod_{j \neq i} \tilde{f}_j(\mathcal{Z}_j) \left\| \tilde{f}_i(\mathcal{Z}_i) \prod_{j \neq i} \tilde{f}_j(\mathcal{Z}_j) \right. \right] \quad (\text{local, contextual: iterative, accurate}) \leftarrow \text{EP}$$

Expectation? Propagation?

EP is really two ideas:

- ▶ **Approximation** of factors.
 - ▶ Usually by “projection” to exponential families.
 - ▶ This involves finding expected sufficient statistics, hence **expectation**.
- ▶ **Local** divergence minimization in the context of other factors.
 - ▶ This leads to a message passing approach, hence **propagation**.

Expectation Propagation (EP)

Input $f_1(\mathcal{Z}_1) \dots f_N(\mathcal{Z}_N)$

Initialize $\tilde{f}_1(\mathcal{Z}_1) = \operatorname{argmin}_{f \in \{\tilde{f}\}} \mathbf{KL}[f_1(\mathcal{Z}_1) \| f_1(\mathcal{Z}_1)]$, $\tilde{f}_i(\mathcal{Z}_i) = 1$ for $i > 1$, $q(\mathcal{Z}) \propto \prod_i \tilde{f}_i(\mathcal{Z}_i)$

repeat

for $i = 1 \dots N$ **do**

Delete: $q_{-i}(\mathcal{Z}) \leftarrow \frac{q(\mathcal{Z})}{\tilde{f}_i(\mathcal{Z}_i)} = \prod_{j \neq i} \tilde{f}_j(\mathcal{Z}_j)$

Project: $\tilde{f}_i^{\text{new}}(\mathcal{Z}) \leftarrow \operatorname{argmin}_{f \in \{\tilde{f}\}} \mathbf{KL}[f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i) \| f(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i)]$

Include: $q(\mathcal{Z}) \leftarrow \tilde{f}_i^{\text{new}}(\mathcal{Z}_i) q_{-i}(\mathcal{Z})$

end for

until convergence

Local updates

Each EP update involves a KL minimisation:

$$\tilde{f}_i^{\text{new}}(\mathcal{Z}) \leftarrow \operatorname{argmin}_{f \in \{\tilde{f}\}} \mathbf{KL}[f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}) \| f(\mathcal{Z}_i) q_{-i}(\mathcal{Z})] \quad \left[q_{-i}(\mathcal{Z}) \stackrel{\text{def}}{=} \prod_{j \neq i} \tilde{f}_j(\mathcal{Z}_j) \right]$$

Separate the contextual factor: $q_{-i}(\mathcal{Z}) = q_{-i}(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_{-i} | \mathcal{Z}_i)$ $[\mathcal{Z}_{-i} \stackrel{\text{def}}{=} \mathcal{Z} \setminus \mathcal{Z}_i]$

Then:

$$\begin{aligned} & \min_f \mathbf{KL}[f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}) \| f(\mathcal{Z}_i) q_{-i}(\mathcal{Z})] \\ &= \max_f \int d\mathcal{Z} f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}) \log f(\mathcal{Z}_i) q_{-i}(\mathcal{Z}) \\ &= \max_f \int d\mathcal{Z}_i d\mathcal{Z}_{-i} f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_{-i} | \mathcal{Z}_i) (\log f(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i) + \log q_{-i}(\mathcal{Z}_{-i} | \mathcal{Z}_i)) \\ &= \max_f \int d\mathcal{Z}_i f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i) (\log f(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i)) \int d\mathcal{Z}_{-i} q_{-i}(\mathcal{Z}_{-i} | \mathcal{Z}_i) \\ &= \min_f \mathbf{KL}[f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i) \| f(\mathcal{Z}_i) q_{-i}(\mathcal{Z}_i)] \end{aligned}$$

$q_{-i}(\mathcal{Z}_i)$ is sometimes called the **cavity distribution**.

Message Passing

- ▶ The cavity distribution (in a tree) can be further broken down into a product of terms from each neighbouring clique:

$$q_{-i}(\mathcal{Z}_i) = \prod_{j \in \text{ne}(i)} M_{j \rightarrow i}(\mathcal{Z}_j \cap \mathcal{Z}_i)$$

- ▶ Once the i th site has been approximated, the messages can be passed on to neighbouring cliques by marginalising to the shared variables (SSM example follows).
⇒ **belief propagation**.


- ▶ In loopy graphs, we can use **loopy belief propagation**. In that case

$$q_{-i}(\mathcal{Z}_i) = \prod_{j \in \text{ne}(i)} M_{j \rightarrow i}(\mathcal{Z}_j \cap \mathcal{Z}_i)$$

becomes an approximation to the **true** cavity distribution (or we can recast the approximation directly in terms of messages ⇒ later lecture).

- ▶ For some approximations (e.g. Gaussian) may be able to compute true loopy cavity using approximate sites, even if computing exact message would have been intractable.
- ▶ In either case, message updates can be scheduled in any order.
- ▶ No guarantee of convergence (but see “power-EP” methods).

EP for a NLSSM



$$P(\mathbf{z}_i | \mathbf{z}_{i-1}) = \phi_i(\mathbf{z}_i, \mathbf{z}_{i-1}) \quad \text{e.g. } \exp(-\|\mathbf{z}_i - h_s(\mathbf{z}_{i-1})\|^2 / 2\sigma^2)$$

$$P(\mathbf{x}_i | \mathbf{z}_i) = \psi_i(\mathbf{z}_i) \quad \text{e.g. } \exp(-\|\mathbf{x}_i - h_o(\mathbf{z}_i)\|^2 / 2\sigma^2)$$

Then $f_i(\mathbf{z}_i, \mathbf{z}_{i-1}) = \phi_i(\mathbf{z}_i, \mathbf{z}_{i-1})\psi_i(\mathbf{z}_i)$. As ϕ_i and ψ_i are non-linear, inference is not generally tractable.

Assume $\tilde{f}_i(\mathbf{z}_i, \mathbf{z}_{i-1})$ is Gaussian. Then,

$$q_{-i}(\mathbf{z}_i, \mathbf{z}_{i-1}) = \int \prod_{\substack{\mathbf{z}_1 \dots \mathbf{z}_{i-2} \\ \mathbf{z}_{i+1} \dots \mathbf{z}_n}} \prod_{i' \neq i} \tilde{f}_{i'}(\mathbf{z}_{i'}, \mathbf{z}_{i'-1}) = \underbrace{\int \prod_{i' < i} \tilde{f}_{i'}(\mathbf{z}_{i'}, \mathbf{z}_{i'-1})}_{\alpha_{i-1}(\mathbf{z}_{i-1})} \underbrace{\int \prod_{i' > i} \tilde{f}_{i'}(\mathbf{z}_{i'}, \mathbf{z}_{i'-1})}_{\beta_i(\mathbf{z}_i)}$$

with both α and β Gaussian.

$$\tilde{f}_i(\mathbf{z}_i, \mathbf{z}_{i-1}) = \operatorname{argmin}_{f \in \mathcal{N}} \mathbf{KL}[f(\mathbf{z}_i, \mathbf{z}_{i-1})\psi_i(\mathbf{z}_i)\alpha_{i-1}(\mathbf{z}_{i-1})\beta_i(\mathbf{z}_i) \| f(\mathbf{z}_i, \mathbf{z}_{i-1})\alpha_{i-1}(\mathbf{z}_{i-1})\beta_i(\mathbf{z}_i)]$$

Moment Matching

Each EP update involves a KL minimisation:

$$\tilde{f}_i^{\text{new}}(\mathcal{Z}) \leftarrow \operatorname{argmin}_{f \in \{\}} \mathbf{KL}[f_i(\mathcal{Z}_i)q_{-i}(\mathcal{Z}) \| f(\mathcal{Z}_i)q_{-i}(\mathcal{Z})]$$

Usually, both $q_{-i}(\mathcal{Z}_i)$ and \tilde{f} are in the same exponential family. Let $q(x) = \frac{1}{Z(\theta)} e^{\mathbf{T}(x) \cdot \theta}$. Then

$$\begin{aligned} \operatorname{argmin}_q \mathbf{KL}[p(x) \| q(x)] &= \operatorname{argmin}_\theta \mathbf{KL}\left[p(x) \left\| \frac{1}{Z(\theta)} e^{\mathbf{T}(x) \cdot \theta} \right.\right] \\ &= \operatorname{argmin}_\theta - \int dx p(x) \log \frac{1}{Z(\theta)} e^{\mathbf{T}(x) \cdot \theta} \\ &= \operatorname{argmin}_\theta - \int dx p(x) \mathbf{T}(x) \cdot \theta + \log Z(\theta) \\ \frac{\partial}{\partial \theta} &= - \int dx p(x) \mathbf{T}(x) + \frac{1}{Z(\theta)} \frac{\partial}{\partial \theta} \int dx e^{\mathbf{T}(x) \cdot \theta} \\ &= -\langle \mathbf{T}(x) \rangle_p + \frac{1}{Z(\theta)} \int dx e^{\mathbf{T}(x) \cdot \theta} \mathbf{T}(x) \\ &= -\langle \mathbf{T}(x) \rangle_p + \langle \mathbf{T}(x) \rangle_q \end{aligned}$$

So minimum is found by **matching sufficient stats** or **moment matching**.



NLSSM EP message updates

$$\tilde{f}_i(\mathbf{z}_i, \mathbf{z}_{i-1}) = \operatorname{argmin}_{f \in \mathcal{N}} \mathbf{KL}[f(\mathbf{z}_i, \mathbf{z}_{i-1})q_{-i}(\mathbf{z}_i, \mathbf{z}_{i-1}) \| f(\mathbf{z}_i, \mathbf{z}_{i-1})q_{-i}(\mathbf{z}_i, \mathbf{z}_{i-1})] = \operatorname{argmin}_{f \in \mathcal{N}} \mathbf{KL}[\underbrace{\phi_i(\mathbf{z}_i, \mathbf{z}_{i-1})\psi_i(\mathbf{z}_i)}_{\tilde{P}} \| \underbrace{q_{-i}(\mathbf{z}_i, \mathbf{z}_{i-1})}_{\tilde{P}}]$$

$$\tilde{P}(\mathbf{z}_{i-1}, \mathbf{z}_i) = \operatorname{argmin}_{P \in \mathcal{N}} \mathbf{KL}[\tilde{P}(\mathbf{z}_{i-1}, \mathbf{z}_i) \| P(\mathbf{z}_{i-1}, \mathbf{z}_i)] \quad \tilde{f}_i(\mathbf{z}_i, \mathbf{z}_{i-1}) = \frac{\tilde{P}(\mathbf{z}_{i-1}, \mathbf{z}_i)}{\alpha_{i-1}(\mathbf{z}_{i-1})\beta_i(\mathbf{z}_i)}$$

$$\alpha_i(\mathbf{z}_i) = \int \prod_{\mathbf{z}_1 \dots \mathbf{z}_{i-1}, i' < i+1} \tilde{f}_{i'}(\mathbf{z}_{i'}, \mathbf{z}_{i'-1}) = \int_{\mathbf{z}_{i-1}} \alpha_{i-1}(\mathbf{z}_{i-1}) \tilde{f}_i(\mathbf{z}_i, \mathbf{z}_{i-1}) = \frac{1}{\beta_i(\mathbf{z}_i)} \int_{\mathbf{z}_{i-1}} \tilde{P}(\mathbf{z}_{i-1}, \mathbf{z}_i)$$

$$\beta_{i-1}(\mathbf{z}_{i-1}) = \int \prod_{\mathbf{z}_{i+1} \dots \mathbf{z}_n, i' > i} \tilde{f}_{i'}(\mathbf{z}_{i'}, \mathbf{z}_{i'-1}) = \int_{\mathbf{z}_i} \beta_i(\mathbf{z}_i) \tilde{f}_i(\mathbf{z}_i, \mathbf{z}_{i-1}) = \frac{1}{\alpha_{i-1}(\mathbf{z}_{i-1})} \int_{\mathbf{z}_i} \tilde{P}(\mathbf{z}_{i-1}, \mathbf{z}_i)$$



Numerical issues

How do we calculate $\langle \mathbf{T}(x) \rangle_p$?

Often analytically tractable, but even if not requires a (relatively) low-dimensional integral:

► Quadrature methods.


- Classical Gaussian quadrature (same Gauss, but nothing to do with the distribution) gives an iterative version of Sigma-point methods.
- Positive definite joints, but not guaranteed to give positive definite messages.
- Heuristics include skipping non-positive-definite steps, or damping messages by interpolation or exponentiating to power < 1 .
- Other quadrature approaches (e.g. GP quadrature) may be more accurate, and may allow formal constraint to pos-def cone.

► Laplace approximation.

- Equivalent to Laplace propagation.
- As long as messages remain positive definite will converge to global Laplace approximation.

EP for Gaussian process classification

EP provides a successful framework for Gaussian-process modelling of non-Gaussian observations (e.g. for classification).




Recall:

- ▶ A GP defines a **multivariate Gaussian** distribution on any finite subset of random vars $\{g_1 \dots g_n\}$ drawn from a (usually uncountable) potential set indexed by "inputs" \mathbf{x}_i .
- ▶ The Gaussian parameters depend on the inputs: $(\boldsymbol{\mu} = [\mu(\mathbf{x}_i)], \Sigma = [K(\mathbf{x}_i, \mathbf{x}_j)])$.
- ▶ If we think of the g s as function values, a GP provides a prior over functions.
- ▶ In a GP regression model, noisy observations y_i are conditionally independent given g_i .
- ▶ No parameters to learn (though often hyperparameters); instead, we make predictions on test data directly: [assuming $\boldsymbol{\mu} = 0$, and matrix Σ incorporates diagonal noise]

$$P(y' | \mathbf{x}', D) = \mathcal{N}(\Sigma_{\mathbf{x}', X} \Sigma_{X, X}^{-1} \mathbf{z}, \Sigma_{\mathbf{x}', \mathbf{x}'} - \Sigma_{\mathbf{x}', X} \Sigma_{X, X}^{-1} \Sigma_{X, \mathbf{x}'})$$


EP GP prediction



- ▶ Once approximate site potentials have stabilised, they can be used to make predictions.
- ▶ Introducing a test point changes K , but does not affect the *marginal* $P(g_1 \dots g_n)$ (by consistency of the GP).
- ▶ The unobserved output factor provides no information about g' (\Rightarrow constant factor on g')
- ▶ Thus no change is needed to the approximating potentials \tilde{f}_i .
- ▶ Predictions are obtained by marginalising the approximation: [let $\tilde{\Psi} = \text{diag}[\tilde{\psi}_1^2 \dots \tilde{\psi}_n^2]$]

$$P(y' | \mathbf{x}', D) = \int dg' P(y' | g') \mathcal{N}(g' | K_{\mathbf{x}', X} (K_{X, X} + \tilde{\Psi})^{-1} \tilde{\boldsymbol{\mu}}, K_{\mathbf{x}', \mathbf{x}'} - K_{\mathbf{x}', X} (K_{X, X} + \tilde{\Psi})^{-1} K_{X, \mathbf{x}'})$$

GP EP updates



- ▶ We can write the GP joint on g_i and y_i as a factor graph:

$$P(g_1 \dots g_n, y_1, \dots, y_n) = \underbrace{\mathcal{N}(g_1 \dots g_n | \mathbf{0}, K)}_{f_0(G)} \prod_i \underbrace{\mathcal{N}(y_i | g_i, \sigma_i^2)}_{f_i(g_i)}$$

- ▶ The same factorisation applies to non-Gaussian $P(y_i | g_i)$ (e.g. $P(y_i=1) = 1/(1 + e^{-g_i})$).
- ▶ EP: approximate **non-Gaussian** $f_i(g_i)$ by **Gaussian** $\tilde{f}_i(g_i) = \mathcal{N}(\tilde{\mu}_i, \tilde{\psi}_i^2)$.
- ▶ $q_{-i}(g_i)$ can be constructed by the usual GP marginalisation. If $\Sigma = K + \text{diag}[\tilde{\psi}_1^2 \dots \tilde{\psi}_n^2]$

$$q_{-i}(g_i) = \mathcal{N}(\Sigma_{i, -i} \Sigma_{-i, -i}^{-1} \tilde{\boldsymbol{\mu}}_{-i}, K_{i, i} - \Sigma_{i, -i} \Sigma_{-i, -i}^{-1} \Sigma_{-i, i})$$

- ▶ The EP updates thus require calculating Gaussian expectations of $f_i(g)g^{\{1,2\}}$:

$$\tilde{f}_i^{\text{new}}(g_i) = \mathcal{N}\left(\int dg q_{-i}(g) f_i(g) g, \int dg q_{-i}(g) f_i(g) g^2 - (\tilde{\mu}_i^{\text{new}})^2\right) / q_{-i}(g_i)$$

Normalisers

- ▶ As long as our approximating class is a tractable exponential family, normalisers can be computed as needed.
- ▶ Consider an approximating class written

$$\tilde{f}_i(\mathcal{Z}_i) \propto e^{\mathbf{T}(\mathcal{Z}_i) \cdot \boldsymbol{\theta}_i - \Phi(\boldsymbol{\theta}_i)}$$

i.e., define a single sufficient statistic vector on all latents, setting entries in $\boldsymbol{\theta}_i$ to 0 for suff stat functions that take cliques other than \mathcal{Z}_i .

- ▶ Then

$$q(\mathcal{Z}) \propto \prod_i \tilde{f}_i \propto e^{\mathbf{T}(\mathcal{Z}) \cdot \sum \boldsymbol{\theta}_i - \sum \Phi(\boldsymbol{\theta}_i)}$$

and so we can simply renormalise at the end as usual:

$$q(\mathcal{Z}) = e^{\mathbf{T}(\mathcal{Z}) \cdot \sum \boldsymbol{\theta}_i - \Phi(\sum \boldsymbol{\theta}_i)}.$$

- ▶ However, to compute an approximation to the likelihood $\prod_i f_i(\mathcal{Z}_i)$ we need to keep track of the site integrals.

Computing likelihoods – keeping track of normalisers

- ▶ Define unnormalised ExpFam approximating sites $\tilde{f}_i = \tilde{C}_i e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_i}$.

Write $\boldsymbol{\theta} = \sum \boldsymbol{\theta}_j$ for the natural parameters of $q(\mathcal{Z})$ and $\boldsymbol{\theta}_{-i} = \sum_{j \neq i} \boldsymbol{\theta}_j$ for the natural parameters of $q_{-i}(\mathcal{Z})$.

Let $\Phi(\boldsymbol{\theta}) = \log \int e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}}$ be the (tractable) ExpFam log normaliser.

- ▶ Now, at each EP step minimise the “unnormalised KL”:

$$\text{KL}[p||q] = \int dx p(x) \log \frac{p(x)}{q(x)} + \int dx (q(x) - p(x))$$

This matches the zeroth moment of $f_i(\mathcal{Z}_i) q_{-i}(\mathcal{Z})$ as well as the expected sufficient statistics as before. That is:

$$\int \tilde{C}_i e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_i} \prod_{-i} \tilde{C}_j e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_j} = \int f_i(\mathcal{Z}_i) \prod_{-i} \tilde{C}_j e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_j} \Rightarrow \tilde{C}_i = e^{\Phi_i(\boldsymbol{\theta}_{-i}) - \Phi(\boldsymbol{\theta})}$$

where Φ_i is the log-normaliser of the “tilted” ExpFam $\hat{P}_i(\mathcal{Z}) \propto f_i(\mathcal{Z}_i) e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}}$.

- ▶ The likelihood approximation is then:

$$\log \int \prod_{i=1}^N f_i(\mathcal{Z}_i) \approx \log \int \prod_{i=1}^N \tilde{f}_i(\mathcal{Z}_i) = \Phi(\boldsymbol{\theta}) + \sum \log \tilde{C}_i \stackrel{\text{def}}{=} \tilde{\ell}$$

EP log-likelihood optimisation for learning

Let true potentials f_i depend on model (hyper)parameters η .

We have

$$\nabla_{\eta} \tilde{\ell} = \nabla_{\eta} \Phi(\boldsymbol{\theta}) + \sum_{i=1}^N \nabla_{\eta} \log \tilde{C}_i = \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} + \sum_{i=1}^N \nabla_{\eta} \log \tilde{C}_i \quad (*)$$

using the standard ExpFam moment-generating result with mean parameters

$$\boldsymbol{\mu} = \langle \mathbf{T}(\mathcal{Z}) \rangle_{q(\mathcal{Z})}$$

Now, zeroth-moment matching implies that at EP convergence:

$$\log \tilde{C}_i = \Phi_i(\boldsymbol{\theta}_{-i}) - \Phi(\boldsymbol{\theta}) \Rightarrow \nabla_{\eta} \log \tilde{C}_i = \nabla_{\eta} \Phi_i(\boldsymbol{\theta}_{-i}) - \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} \quad (**)$$

but $\Phi_i(\boldsymbol{\theta}_{-i}) = \log \int f_i(\mathcal{Z}_i) e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_{-i}}$ depends on η in two ways: *directly* through f_i and *indirectly* through the converged $\boldsymbol{\theta}_{-i}$.

$$\begin{aligned} \nabla_{\eta} \Phi_i(\boldsymbol{\theta}_{-i}) &= \partial_{\boldsymbol{\theta}_{-i}} \Phi_i(\boldsymbol{\theta}_{-i}) \cdot \nabla_{\eta} \boldsymbol{\theta}_{-i} + e^{-\Phi_i(\boldsymbol{\theta}_{-i})} \int \nabla_{\eta} f_i(\mathcal{Z}_i) e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_{-i}} \\ &= \langle \mathbf{T}(\mathcal{Z}) \rangle_{\hat{P}_i} \cdot \nabla_{\eta} \boldsymbol{\theta}_{-i} + \int \nabla_{\eta} \log f_i(\mathcal{Z}_i) f_i(\mathcal{Z}_i) e^{\mathbf{T}(\mathcal{Z}) \cdot \boldsymbol{\theta}_{-i} - \Phi_i(\boldsymbol{\theta}_{-i})} \\ &= \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta}_{-i} + \langle \nabla_{\eta} \log f_i(\mathcal{Z}_i) \rangle_{\hat{P}_i} \end{aligned} \quad (***)$$

by EP moment matching at convergence!

Learning

EP yields approximate *inferential* posteriors. To learn (hyper)parameters we can use:

- ▶ Approximate Bayesian inference (analogous to VB)
 - ▶ may be difficult to construct a coherent normalisable exponential family approximation on both latents and parameters.
- ▶ Approximate EM – maximize $\langle \log P(\mathcal{X}, \mathcal{Z}) \rangle_{q_{EP}(\mathcal{Z})}$.
 - ▶ Practical, but no coherent cost function (unlike variational inference), so no guarantee of convergence even if EP itself converges.
- ▶ Direct maximisation of EP log-likelihood estimate.
 - ▶ Consistent, although convergence guarantees still difficult.
 - ▶ Seems challenging as we need to differentiate through (iteration-based) dependence of approximate $q(\mathcal{Z})$ and \tilde{C}_i s.
 - ▶ However, proves to be simpler than it sounds.

EP log-likelihood optimisation for learning

So putting it all together:

$$\begin{aligned} \nabla_{\eta} \tilde{\ell} &= \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} + \sum_{i=1}^N \nabla_{\eta} \log \tilde{C}_i \quad (*) \\ &= \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} + \sum_{i=1}^N \left(\nabla_{\eta} \Phi_i(\boldsymbol{\theta}_{-i}) - \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} \right) \quad (**) \\ &= \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} + \sum_{i=1}^N \left(\boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta}_{-i} - \boldsymbol{\mu} \cdot \nabla_{\eta} \boldsymbol{\theta} + \langle \nabla_{\eta} \log f_i(\mathcal{Z}_i) \rangle_{\hat{P}_i} \right) \quad (***) \\ &= \boldsymbol{\mu} \cdot \nabla_{\eta} \left(\boldsymbol{\theta} + \sum_{i=1}^N (\boldsymbol{\theta}_{-i} - \boldsymbol{\theta}) \right) + \sum_{i=1}^N \langle \nabla_{\eta} \log f_i(\mathcal{Z}_i) \rangle_{\hat{P}_i} \\ &= \boldsymbol{\mu} \cdot \nabla_{\eta} \left(\sum_{i=1}^N \boldsymbol{\theta}_i + \sum_{i=1}^N (\boldsymbol{\theta}_{-i} - \boldsymbol{\theta}) \right) + \sum_{i=1}^N \langle \nabla_{\eta} \log f_i(\mathcal{Z}_i) \rangle_{\hat{P}_i} \\ &= \boldsymbol{\mu} \cdot \nabla_{\eta} \sum_{i=1}^N (\boldsymbol{\theta} - \boldsymbol{\theta}) + \sum_{i=1}^N \langle \nabla_{\eta} \log f_i(\mathcal{Z}_i) \rangle_{\hat{P}_i} \\ &= \sum_{i=1}^N \langle \nabla_{\eta} \log f_i(\mathcal{Z}_i) \rangle_{\hat{P}_i} \end{aligned}$$

and the gradient can be computed provided EP converges.

A final generalisation: alpha divergences and Power EP

- ▶ Alpha divergences $D_\alpha[p||q] = \frac{1}{\alpha(1-\alpha)} \int dx \alpha p(x) + (1-\alpha)q(x) - p(x)^\alpha q(x)^{1-\alpha}$

$$D_{-1}[p||q] = \frac{1}{2} \int dx \frac{(p(x) - q(x))^2}{p(x)}$$

$$\lim_{\alpha \rightarrow 0} D_\alpha[p||q] = \mathbf{KL}[q||p] \quad \text{Note: } \lim_{\alpha \rightarrow 0} \frac{(p(x)/q(x))^\alpha}{\alpha} = \log \frac{p(x)}{q(x)}$$

$$D_{\frac{1}{2}}[p||q] = 2 \int dx (p(x)^{\frac{1}{2}} - q(x)^{\frac{1}{2}})^2$$

$$\lim_{\alpha \rightarrow 1} D_\alpha[p||q] = \mathbf{KL}[p||q]$$

$$D_2[p||q] = \frac{1}{2} \int dx \frac{(p(x) - q(x))^2}{q(x)}$$

- ▶ Local (EP) minimisation gives fixed-point updates that blend messages (to power α) with previous site approximations.

$$\tilde{f}_i^{\text{new}} = \operatorname{argmin}_{f \in \{f\}} \mathbf{KL}[f_i(\mathcal{Z}_i)^\alpha \tilde{f}_i(\mathcal{Z}_i)^{1-\alpha} q_{-i}(\mathcal{Z}) || f(\mathcal{Z}_i) q_{-i}(\mathcal{Z})]$$

- ▶ Small changes (for $\alpha < 1$) lead to more stable updates, and more reliable convergence.