Probabilistic & Unsupervised Learning

Expectation Propagation
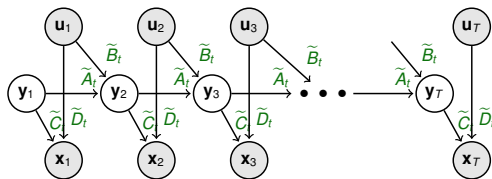
Maneesh Sahani

maneesh@gatsby.ucl.ac.uk

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

Term 1, Autumn 2015

Intractabilities and approximations

- Inference – computational intractability
    - Factored variational approx
    - Loopy BP/EP/Power EP
    - Gibbs sampling, other MCMC

- Inference – analytic intractability
    - Laplace approximation (global)
    - Parametric variational approx (for special cases).
    - Message approximations (linearised, sigma-point, Laplace)
    - Assumed-density methods and Expectation-Propagation
    - (Sequential) Monte-Carlo methods

- Learning – intractable partition function
    - Constrastive divergence
    - Sampling parameters
    - Score-matching

- Model selection
    - Laplace approximation / BIC
    - Variational Bayes
    - (Annealed) importance sampling
    - Reversible jump MCMC

Not a complete list!

Nonlinear state-space model (NLSSM)

$$\mathbf{y}_{t+1} = f(\mathbf{y}_t, \mathbf{u}_t) + \mathbf{w}_t$$
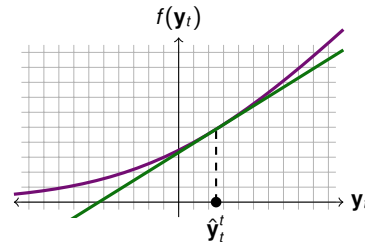$$\mathbf{x}_t = g(\mathbf{y}_t, \mathbf{u}_t) + \mathbf{v}_t$$

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

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

$$\mathbf{y}_{t+1} \approx \underbrace{f(\hat{\mathbf{y}}_t^t, \mathbf{u}_t)}_{\widetilde{B}_t \mathbf{u}_t} + \underbrace{\left.\frac{\partial f}{\partial \mathbf{y}_t}\right|_{\hat{\mathbf{y}}_t^t}}_{\widetilde{A}_t} (\mathbf{y}_t - \hat{\mathbf{y}}_t^t) + \mathbf{w}_t$$

$$\mathbf{x}_t \approx \underbrace{g(\hat{\mathbf{y}}_t^{t-1}, \mathbf{u}_t)}_{\widetilde{D}_t \mathbf{u}_t} + \underbrace{\left.\frac{\partial g}{\partial \mathbf{y}_t}\right|_{\hat{\mathbf{y}}_t^{t-1}}}_{\widetilde{C}_t} (\mathbf{y}_t - \hat{\mathbf{y}}_t^{t-1}) + \mathbf{v}_t$$

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

- 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).

Other message approximations

Consider the forward messages on a latent chain:

$$P(y_t|x_{1:t}) = \frac{1}{Z} P(x_t|y_t) \int dy_{t-1} \, P(y_t|y_{t-1}) P(y_{t-1}|x_{1:t-1})$$

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

$$\tilde{P}(y_t|x_{1:t}) \approx \frac{1}{Z} P(x_t|y_t) \int dy_{t-1} \, \underbrace{P(y_t|y_{t-1})}_{\mathcal{N}\left(f(\mathbf{y}_{t-1}), Q\right)} \underbrace{\tilde{P}(y_{t-1}|x_{1:t-1})}_{\mathcal{N}\left(\hat{\mathbf{y}}_{t-1}, V_{t-1}\right)}$$

- Linearisation at the peak (EKF) is only one approach.
- Laplace filter: use mode and curvature of integrand.
- Sigma-point ("unscented") filter:
    - Evaluate $f(\hat{\mathbf{y}}_{t-1})$, $f(\hat{\mathbf{y}}_{t-1} \pm \sqrt{\lambda}\mathbf{v})$ for eigenvalues, eigenvectors $\hat{V}_{t-1}\mathbf{v} = \lambda \mathbf{v}$.
    - "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: argmin $\mathbf{KL}\left[\mathcal{N}\left(\hat{\mathbf{y}}_t, \hat{V}_t\right) \| \int dy_{t-1} \ldots\right]$. Requires Gaussian expectations of $\log \int \Rightarrow$ may be challenging.
- The other KL: argmin $\mathbf{KL}\left[\int dy_{t-1} \| \mathcal{N}\left(\hat{\mathbf{y}}_t, \hat{V}_t\right)\right]$ needs only first and second moments of nonlinear message $\Rightarrow$ EP.

Variational learning

Free energy:

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

E-steps:

- ▶ Exact EM: $q(\mathcal{Y}) = \underset{q}{\text{argmax}}\, \mathcal{F} = P(\mathcal{Y}|\mathcal{X}, \theta)$
  - ▶ Saturates bound: converges to local maximum of likelihood.

- ▶ (Factored) variational approximation:

  $$q(\mathcal{Y}) = \underset{q_1(\mathcal{Y}_1)q_2(\mathcal{Y}_2)}{\text{argmax}}\, \mathcal{F} = \underset{q_1(\mathcal{Y}_1)q_2(\mathcal{Y}_2)}{\text{argmin}}\, \mathbf{KL}[q_1(\mathcal{Y}_1)q_2(\mathcal{Y}_2)\|P(\mathcal{Y}|\mathcal{X},\theta)]$$
  - ▶ Increases bound: converges, but not necessarily to ML.

- ▶ Other approximations: $q(\mathcal{Y}) \approx P(\mathcal{Y}|\mathcal{X}, \theta)$
  - ▶ Usually no guarantees, but if learning converges if is frequently more accurate than the factored approximation

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.

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

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?

The other KL

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

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

$$\underset{q_i}{\text{argmin}}\, \mathbf{KL}\left[P(\mathcal{Y}|\mathcal{X})\Big\|\prod q_j(\mathcal{Y}_j|\mathcal{X})\right] = \underset{q_i}{\text{argmin}} -\int d\mathcal{Y}\, P(\mathcal{Y}|\mathcal{X}) \log \prod_j q_j(\mathcal{Y}_j|\mathcal{X})$$

$$= \underset{q_i}{\text{argmin}} -\sum_j \int d\mathcal{Y}\, P(\mathcal{Y}|\mathcal{X}) \log q_j(\mathcal{Y}_j|\mathcal{X})$$

$$= \underset{q_i}{\text{argmin}} -\int d\mathcal{Y}_i\, P(\mathcal{Y}_i|\mathcal{X}) \log q_i(\mathcal{Y}_i|\mathcal{X})$$

$$= P(\mathcal{Y}_i|\mathcal{X})$$

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.

Approximate optimisation

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

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

where the $\mathcal{Y}_i$ 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: $q(\mathcal{Y}) \overset{\text{def}}{=} \prod_{i=1}^{N} \tilde{f}_i(\mathcal{Y}_i)$.

We would like to minimise (at least in some sense) $\mathbf{KL}[P\|q]$.

Possible optimisations:

$$\underset{q(\mathcal{Y})}{\min}\, \mathbf{KL}\left[\prod_{i=1}^{N} f_i(\mathcal{Y}_i)\Big\|\prod_{i=1}^{N} \tilde{f}_i(\mathcal{Y}_i)\right] \qquad \text{(global: intractable)}$$

$$\underset{\tilde{f}_i}{\min}\, \mathbf{KL}\left[f_i(\mathcal{Y}_i)\Big\|\tilde{f}_i(\mathcal{Y}_i)\right] \qquad \text{(local, fixed: simple, inaccurate)}$$

$$\underset{\tilde{f}_i}{\min}\, \mathbf{KL}\left[f_i(\mathcal{Y}_i)\prod_{j\neq i} \tilde{f}_j(\mathcal{Y}_j)\Big\|\tilde{f}_i(\mathcal{Y}_i)\prod_{j\neq i} \tilde{f}_j(\mathcal{Y}_j)\right] \qquad \text{(local, contextual: iterative, accurate)} \leftarrow \text{EP}$$

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.

Local updates

Each EP update involves a KL minimisation:

$$\tilde{f}_i^{\mathrm{new}}(\mathcal{Y}) \leftarrow \operatorname*{argmin}_{f \in \{\tilde{f}\}} \mathbf{KL}[f_i(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y}) \| f(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y})] \qquad \left[q_{\neg i}(\mathcal{Y}) \stackrel{\mathrm{def}}{=} \prod_{j \neq i} \tilde{f}_j(\mathcal{Y}_j) \right]$$

Write $q_{\neg i}(\mathcal{Y}) = q_{\neg i}(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_{\neg i} | \mathcal{Y}_i)$. Then: $\qquad [\mathcal{Y}_{\neg i} \stackrel{\mathrm{def}}{=} \mathcal{Y} \backslash \mathcal{Y}_i]$

$$\min_f \mathbf{KL}[f_i(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y}) \| f(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y})]$$

$$= \max_f \int d\mathcal{Y}_i d\mathcal{Y}_{\neg i} \, f_i(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}) \log f(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y})$$

$$= \max_f \int d\mathcal{Y}_i d\mathcal{Y}_{\neg i} \, f_i(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_{\neg i} | \mathcal{Y}_i) \big(\log f(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_i) + \log q_{\neg i}(\mathcal{Y}_{\neg i} | \mathcal{Y}_i) \big)$$

$$= \max_f \int d\mathcal{Y}_i \, f_i(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_i) \big(\log f(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_i) \big) \int d\mathcal{Y}_{\neg i} \, q_{\neg i}(\mathcal{Y}_{\neg i} | \mathcal{Y}_i)$$

$$= \min_f \mathbf{KL}[f_i(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_i) \| f(\mathcal{Y}_i) q_{\neg i}(\mathcal{Y}_i)]$$

$q_{\neg i}(\mathcal{Y}_i)$ is sometimes called the cavity distribution.

Expectation Propagation (EP)

Input $f_1(\mathcal{Y}_1) \ldots f_N(\mathcal{Y}_N)$

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

repeat

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

        Delete: $q_{\neg i}(\mathcal{Y}) \leftarrow \dfrac{q(\mathcal{Y})}{\tilde{f}_i(\mathcal{Y}_i)} = \prod_{j \neq i} \tilde{f}_j(\mathcal{Y}_i)$

        Project: $\tilde{f}_i^{\mathrm{new}}(\mathcal{Y}) \leftarrow \operatorname*{argmin}_{f \in \{\tilde{f}\}} \mathbf{KL}[f_i(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y}_i) \| f(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y}_i)]$

        Include: $q(\mathcal{Y}) \leftarrow \tilde{f}_i^{\mathrm{new}}(\mathcal{Y}_i) \, q_{\neg i}(\mathcal{Y})$

    end for

until convergence

Message Passing

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

$$q_{\neg i}(\mathcal{Y}_i) = \prod_{j \in \mathrm{ne}(i)} M_{j \to i}(\mathcal{Y}_j \cap \mathcal{Y}_i)$$

- ► Once the ith site has been approximated, the messages can be passed on to neighbouring cliques by marginalising to the shared variables (SSM example follows). $\Rightarrow$ belief propagation.

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

$$q_{\neg i}(\mathcal{Y}_i) = \prod_{j \in \mathrm{ne}(i)} M_{j \to i}(\mathcal{Y}_j \cap \mathcal{Y}_i)$$

becomes an approximation to the **true** cavity distribution (or we can recast the approximation directly in terms of messages $\Rightarrow$ 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
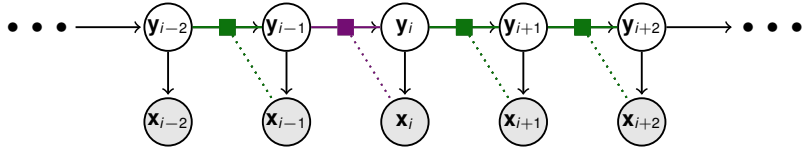
$P(\mathbf{y}_i|\mathbf{y}_{i-1}) = \phi_i(\mathbf{y}_i, \mathbf{y}_{i-1})$ *e.g.* $\exp(-\|\mathbf{y}_i - h_s(\mathbf{y}_{i-1})\|^2/2\sigma^2)$

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

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

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

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

with both α and β Gaussian.

$$\tilde{f}_i(\mathbf{y}_i, \mathbf{y}_{i-1}) = \underset{f\in\mathcal{N}}{\mathrm{argmin}}\,\mathbf{KL}\big[\phi_i(\mathbf{y}_i, \mathbf{y}_{i-1})\psi_i(\mathbf{y}_i)\alpha_{i-1}(\mathbf{y}_{i-1})\beta_i(\mathbf{y}_i)\big\|f(\mathbf{y}_i, \mathbf{y}_{i-1})\alpha_{i-1}(\mathbf{y}_{i-1})\beta_i(\mathbf{y}_i)\big]$$

Moment Matching

Each EP update involves a KL minimisation:

$$\tilde{f}_i^{\mathrm{new}}(\mathcal{Y}) \leftarrow \underset{f\in\{\tilde{f}\}}{\mathrm{argmin}}\,\mathbf{KL}[f_i(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y})\|f(\mathcal{Y}_i)q_{\neg i}(\mathcal{Y})]$$

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

$$\begin{aligned}
\underset{q}{\mathrm{argmin}}\,\mathbf{KL}\big[p(x)\big\|q(x)\big] &= \underset{\theta}{\mathrm{argmin}}\,\mathbf{KL}\Big[p(x)\Big\|\frac{1}{Z(\theta)}e^{\mathsf{T}(x)\cdot\theta}\Big]\\
&= \underset{\theta}{\mathrm{argmin}}-\int dx\, p(x)\log\frac{1}{Z(\theta)}e^{\mathsf{T}(x)\cdot\theta}\\
&= \underset{\theta}{\mathrm{argmin}}-\int dx\, p(x)\mathsf{T}(x)\cdot\theta + \log Z(\theta)\\
\frac{\partial}{\partial\theta} &= -\int dx\, p(x)\mathsf{T}(x) + \frac{1}{Z(\theta)}\frac{\partial}{\partial\theta}\int dx\, e^{\mathsf{T}(x)\cdot\theta}\\
&= -\langle\mathsf{T}(x)\rangle_p + \frac{1}{Z(\theta)}\int dx\, e^{\mathsf{T}(x)\cdot\theta}\mathsf{T}(x)\\
&= -\langle\mathsf{T}(x)\rangle_p + \langle\mathsf{T}(x)\rangle_q
\end{aligned}$$

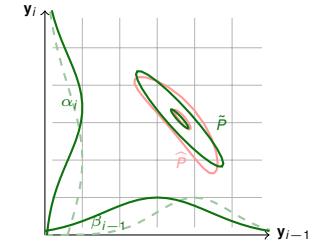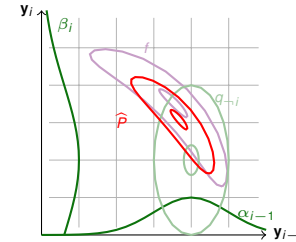So minimum is found by matching sufficient stats. This is usually moment matching.

NLSSM EP message updates

$$\tilde{f}_i(\mathbf{y}_i, \mathbf{y}_{i-1}) = \underset{f\in\mathcal{N}}{\mathrm{argmin}}\,\mathbf{KL}\big[f(\mathbf{y}_i, \mathbf{y}_{i-1})q_{\neg i}(\mathbf{y}_i, \mathbf{y}_{i-1})\big\|f(\mathbf{y}_i, \mathbf{y}_{i-1})q_{\neg i}(\mathbf{y}_i, \mathbf{y}_{i-1})\big] = \underset{f\in\mathcal{N}}{\mathrm{argmin}}\,\mathbf{KL}\big[\phi_i(\mathbf{y}_i, \mathbf{y}_{i-1})\psi$$

$$\tilde{P}(\mathbf{y}_{i-1}, \mathbf{y}_i) = \underset{P\in\mathcal{N}}{\mathrm{argmin}}\,\mathbf{KL}\big[\widehat{P}(\mathbf{y}_{i-1}, \mathbf{y}_i)\big\|P(\mathbf{y}_{i-1}, \mathbf{y}_i)\big] \qquad \tilde{f}_i(\mathbf{y}_i, \mathbf{y}_{i-1}) = \frac{\tilde{P}(\mathbf{y}_{i-1}, \mathbf{y}_i)}{\alpha_{i-1}(\mathbf{y}_{i-1})\beta_i(\mathbf{y}_i)}$$

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

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

Numerical issues

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

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 the most succesful framework for Gaussian-process modelling of non-Gaussian observations (*e.g.* for classification).
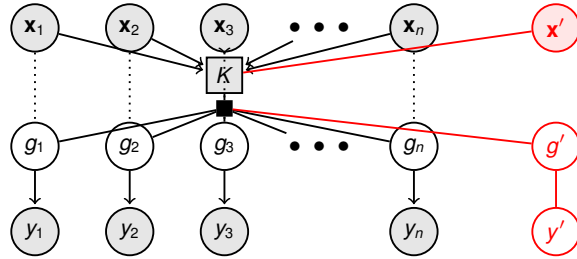

Recall:

- ► A GP defines a multivariate Gaussian distribution on any finite subset of random vars $\{g_1 \ldots g_n\}$ 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 gs 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 $\mu = 0$, and matrix Σ incorporates diagonal noise]

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

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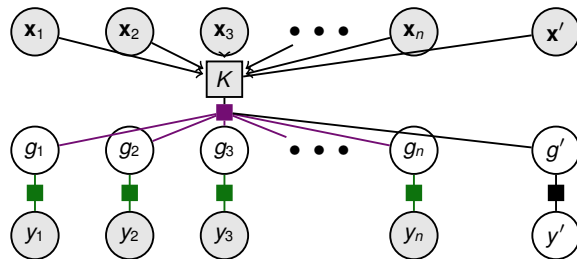
- ► Once appoximate site potentials have stabilised, they can be used to make predictions.
- ► Introducing a test point changes K, but does not affect the *marginal* $P(g_1 \ldots g_n)$ (by consistency of the GP).
- ► The unobserved output factor provides no information about g' ($\Rightarrow$ constant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

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

$$\left. K_{x',x'} - K_{x',X}(K_{X,X} + \tilde{\Psi})^{-1}K_{X,x'}\right)$$

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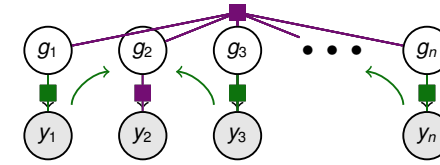
- ► We can write the GP joint on g_i and y_i as a factor graph:

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

- ► 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 $\tilde{f}_i(g_i) = \mathcal{N}\left(\tilde{\mu}_i, \tilde{\psi}_i^2\right)$.
- ► $q_{\neg i}(g_i)$ can be constructed by the usual GP marginalisation. If $\Sigma = K + \text{diag}\left[\tilde{\psi}_1^2 \ldots \tilde{\psi}_n^2\right]$

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

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

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

Normalisers

- ► Approximate sites determined by moment matching are naturally normalised.

- ► For posteriors, sufficient to normalise product after convergence.
  - ► Often straightforward for exponential family approximations.

- ► To compute likelihood need to keep track of site integrals:
  - ► minimising "unnormalised KL":

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

  incorporates normaliser into each $\tilde{f}$ (match zeroth moment, along with suff stats).

as well as the overall normaliser of $\prod_i \tilde{f}_i(\mathcal{Y}_i)$.

Alpha divergences and Power EP

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

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

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

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

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

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

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

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

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