
Supplementary Material:

Temporal alignment and latent Gaussian process factor inference in population spike trains

Lea Duncker & Maneesh Sahani
 Gatsby Computational Neuroscience Unit
 University College London
 London, W1T 4JG
 {duncker, maneesh}@gatsby.ucl.ac.uk

1 Derivation of the svGPFA Variational Lower Bound

In order to arrive at a scalable variational inference algorithm, we make use of a sparse GP approximation. We introduce a set of inducing points $U = [\mathbf{u}_1, \dots, \mathbf{u}_K]$ for each latent process, which are each evaluated on a set of inducing point locations $Z = [z_1, \dots, z_K]$.

The joint-data likelihood of the full model, now including the inducing points, is hence

$$p(\mathcal{Y}, \mathbf{x}(\cdot), U) = p(\mathcal{Y}|\mathbf{x}(\cdot)) \prod_{k=1}^K p(x_k(\cdot)|\mathbf{u}_k)p(\mathbf{u}_k|z_k) \quad (1)$$

where we have omitted explicit conditioning on model parameters to avoid cluttered notation. A variational lower bound to the log-likelihood can be obtained by applying Jensen's inequality:

$$\log p(\mathcal{Y}) \geq \iint q(\{\mathbf{u}_k\}_{k=1}^K, \mathbf{x}(\cdot)) \log \left(\frac{p(\mathcal{Y}|\mathbf{x}(\cdot)) \prod_{k=1}^K p(x_k(\cdot)|\mathbf{u}_k)p(\mathbf{u}_k|z_k)}{q(\{\mathbf{u}_k\}_{k=1}^K, \mathbf{x}(\cdot))} \right) d\mathbf{u}_1 \dots d\mathbf{u}_K d\mathbf{x} \quad (2)$$

We choose a factorised approximating distribution of the form

$$q(\{\mathbf{u}_k\}_{k=1}^K, \mathbf{x}(\cdot)) = \prod_{k=1}^K q(\mathbf{u}_k, x_k(\cdot)) = \prod_{k=1}^K p(x_k(\cdot)|\mathbf{u}_k)q(\mathbf{u}_k) \quad (3)$$

and choose $q(\mathbf{u}_k) = \mathcal{N}(\mathbf{m}_k, S_k)$ to be multivariate Gaussian. This choice of approximating distribution allows one to write the lower bound as

$$\log p(\mathcal{Y}) \geq \iint \prod_{k=1}^K q(\mathbf{u}_k)p(x_k(\cdot)|\mathbf{u}_k) \log \left(p(\mathcal{Y}|\mathbf{x}(\cdot)) \right) d\mathbf{u}_k d\mathbf{x} - \sum_{k=1}^K \text{KL}[q(\mathbf{u}_k)||p(\mathbf{u}_k)] \quad (4)$$

The second term is the Kullback-Leibler divergence between two Gaussian distributions, which can be evaluated analytically. In order to manipulate the first term, let $h_n(\cdot) = \sum_{k=1}^K c_{n,k}x_k(\cdot) + d_n$ denote the affine transformation of latent GPs for the n -th neuron. We can obtain a marginal variational distribution over $h_n(\cdot)$ as a GP with additive structure

$$q(h_n(\cdot)) = \int \prod_{k=1}^K q(\mathbf{u}_k)p(h_n(\cdot)|\{\mathbf{u}_k\})d\mathbf{u}_k = \mathcal{GP}(\nu_n(\cdot), \sigma_n(\cdot, \cdot)) \quad (5)$$

where

$$\nu_n(t) = \sum_{k=1}^K c_{n,k} \boldsymbol{\kappa}_k(t, \mathbf{z}_k) \mathbf{K}_{zz}^{(k)-1} \mathbf{m}_k + d_n \quad (6)$$

$$\sigma_n(t, t') = \sum_{k=1}^K c_{n,k}^2 \left(\kappa_k(t, t') + \boldsymbol{\kappa}_k(t, \mathbf{z}_k) \left(\mathbf{K}_{zz}^{(k)-1} S_k \mathbf{K}_{zz}^{(k)-1} - \mathbf{K}_{zz}^{(k)-1} \right) \boldsymbol{\kappa}_k(\mathbf{z}_k, t') \right) \quad (7)$$

$\boldsymbol{\kappa}_k(\cdot, \mathbf{z}_k)$ is a vector valued function taking a single time-point as an input argument and consisting of evaluations of the kernel $\kappa_k(\cdot, \cdot)$ at the inducing point locations. $\mathbf{K}_{zz}^{(k)}$ is the kernel Gram matrix of the k th process evaluated at the respective inducing point locations.

To obtain the final expression for the variational lower bound, we can simply rewrite the expression in equation (??):

$$\log p(\mathcal{Y}) \geq \mathbb{E}_{q(h_n)} [\log p(\mathcal{Y} | h_n(\cdot))] - \sum_{k=1}^K \text{KL}[q(\mathbf{u}_k) \| p(\mathbf{u}_k)] \quad (8)$$

2 Use of a point-process likelihood

Using a point-process likelihood in the GPFA model amounts to evaluating the expected log-likelihood in the first term in (??):

$$\mathbb{E}_{q(h_n^{(r)})} [\log p(\mathbf{t}_n^{(r)} | h_n^{(r)})] = -\mathbb{E}_{q(h_n^{(r)})} \left[\int_{\mathcal{T}} g(h_n^{(r)}(t)) dt \right] + \sum_{i_n=1}^{\Phi(n,r)} \mathbb{E}_{q(h_n^{(r)})} [\log g(h_n^{(r)}(t_i))] \quad (9)$$

We can apply Fubini's theorem to switch the order of integration in the first term:

$$\mathbb{E}_{q(h_n^{(r)})} [\log p(\mathbf{t}_n^{(r)} | h_n^{(r)})] = -\int_{\mathcal{T}} \mathbb{E}_{q(h_n^{(r)})} [g(h_n^{(r)}(t))] dt + \sum_{i_n=1}^{\Phi(n,r)} \mathbb{E}_{q(h_n^{(r)})} [\log g(h_n^{(r)}(t_i))] \quad (10)$$

Which gives the final form of the expected log-likelihood. Depending on the choice of non-linearity $g(\cdot)$, the expectation terms can either be evaluated analytically, or efficiently using Gauss-Hermite quadrature. The first term in (??) involves one-dimensional integrals, which can be computed using efficient numerical approximations such as Gauss-Legendre quadrature.

3 Condition-grouped model with time-warping

The full approximating distribution across trials is chosen to be of the form

$$\begin{aligned} & q(\{\{\{\zeta_k^{(r)}, \mathbf{u}_k^{\zeta, (r)}\}_{k=1}^K\}_{\zeta}, \tau^{(r)}, \mathbf{u}^{\tau, (r)}\}_{r=1}^R) \\ &= \prod_{r=1}^R \left(\prod_{\zeta \in \{\alpha, \beta, \gamma\}} \prod_{k=1}^K p(\zeta_k^{(r)} | \mathbf{u}_k^{\zeta, (r)}, \tau^{(r)}) q(\mathbf{u}_k^{\zeta, (r)}) \right) p(\tau^{(r)} | \mathbf{u}^{\tau, (r)}) q(\mathbf{u}^{\tau, (r)}) \end{aligned} \quad (11)$$

Under this approximation, the variational lower bound to the log-likelihood becomes

$$\begin{aligned} \log p(\mathcal{Y}) &\geq \sum_{r=1}^R \sum_{n=1}^N \mathbb{E}_{q(h_n^{(r)})} [\log p(\mathbf{y}_n^{(r)} | h_n^{(r)})] - \sum_{r=1}^R \text{KL}[q(\mathbf{u}^{\tau, (r)}) \| p(\mathbf{u}^{\tau, (r)})] \\ &\quad - \sum_{k=1}^K \text{KL}[q(\mathbf{u}_k^{\alpha}) \| p(\mathbf{u}_k^{\alpha})] - \sum_{\ell=1}^L \sum_{k=1}^K \text{KL}[q(\mathbf{u}_k^{\beta, (\ell)}) \| p(\mathbf{u}_k^{\beta, (\ell)})] \\ &\quad - \sum_{r=1}^R \sum_{k=1}^K \text{KL}[q(\mathbf{u}_k^{\gamma, (r)}) \| p(\mathbf{u}_k^{\gamma, (r)})] \end{aligned} \quad (12)$$

Where

$$\begin{aligned}
q(h_n^{(r)}(t)) &= \int d\mathbf{u}_k^\alpha d\mathbf{u}_k^{\beta, \ell(r)} d\mathbf{u}_k^{\gamma, (r)} d\mathbf{u}^{\tau, (r)} d\tau^{(r)} \\
&\prod_{k=1}^K q(\mathbf{u}_k^\alpha) q(\mathbf{u}_k^{\beta, \ell(r)}) q(\mathbf{u}_k^{\gamma, (r)}) p(h_n^{(r)}(t) | \{\mathbf{u}_k^{\zeta, (r)}\}_{k, \zeta}, \tau^{(r)}) p(\tau^{(r)} | \mathbf{u}^{\tau, (r)}) q(\mathbf{u}^{\tau, (r)})
\end{aligned} \tag{13}$$

Letting

$$q(\tau^{(r)}) = \int d\mathbf{u}^{\tau, (r)} p(\tau^{(r)} | \mathbf{u}^{\tau, (r)}) q(\mathbf{u}^{\tau, (r)}) \tag{14}$$

We can marginalise out the inducing points and evaluate $q(h_n^{(r)}(t))$ as an additive Gaussian Process with mean and covariance function:

$$\begin{aligned}
\nu_n^{(r)}(t) &= \sum_{\zeta, k} c_{n, k}^\zeta \Psi_{k, 1}^{\zeta, (r)}(t, \mathbf{z}_k^\zeta) \mathbf{K}_{zz}^{\zeta, (k)}{}^{-1} \mathbf{m}_k^{\zeta, (r)} + d_n \\
\sigma_n^{(r)}(t, t) &= \sum_{\zeta, k} c_{n, k}^\zeta{}^2 \left(\Psi_{k, 0}^{\zeta, (r)}(t) + \text{Tr} \left[\left(\mathbf{K}_{zz}^{\zeta, (k)}{}^{-1} S_k^{\zeta, (r)} \mathbf{K}_{zz}^{\zeta, (k)}{}^{-1} - \mathbf{K}_{zz}^{\zeta, (k)}{}^{-1} \right) \Psi_{k, 2}^{\zeta, (r)}(t, \mathbf{z}_k^\zeta) \right] \right)
\end{aligned} \tag{15}$$

where

$$\begin{aligned}
\Psi_{k, 0}^{\zeta, (r)}(t) &= \mathbb{E}_{q(\tau^{(r)})} \left[\boldsymbol{\kappa}_k^\zeta(\tau^{(r)}(t), \tau^{(r)}(t)) \right] \\
\Psi_{k, 1}^{\zeta, (r)}(t, \mathbf{z}_k^\zeta) &= \mathbb{E}_{q(\tau^{(r)})} \left[\boldsymbol{\kappa}_k^\zeta(\tau^{(r)}(t), \mathbf{z}_k^\zeta) \right] \\
\Psi_{k, 2}^{\zeta, (r)}(t, \mathbf{z}_k^\zeta) &= \mathbb{E}_{q(\tau^{(r)})} \left[\boldsymbol{\kappa}_k^\zeta(\mathbf{z}_k^\zeta, \tau^{(r)}(t)) \boldsymbol{\kappa}_k^\zeta(\tau^{(r)}(t), \mathbf{z}_k^\zeta) \right]
\end{aligned} \tag{16}$$