The Variational Coupled Gaussian Process Dynamical Model

Dmytro Velychko, Dominik Endres
Theoretical Neuroscience Group, Department of Psychology
Philipps-University, 35032 Marburg, Germany.
velychko@staff.uni-marburg.de, dominik.endres@uni-marburg.de

Introduction

Planning and execution of full-body movements is a hard control problem. Modular movement primitives (MP) have been suggested as a means to simplify it while retaining a sufficient degree of control flexibility for a wide range of task [Bizzi et al., 2008]. A particularly well-developed type of MP in robotics is the dynamical movement primitive (DMP) [Schaal, 2006]. But the form of the differential equation of a DMP remains fixed during learning, potentially reducing the representational capacity. Coupled Gaussian Process Dynamical Model (CGPDM) [Velychko et al., 2014] combines the advantages of modularity and flexibility in the dynamics, at least theoretically. Non-parametric GPs learning have high computational cost. We improve this by employing sparse variational approximations [Titsias, 2009, Titsias and Lawrence, 2010, Frigola et al., 2014] and obviating the need for sampling.

Coupled GPDMs model

Full dynamical system comprises M parts where every part is a modular MP (here 2 for simplicity). We introduce $M \times M$ dynamics in the latent space, each of which makes prediction about a part of the latent space from the previous state of some (other) part, these predictions are combined via PoE:

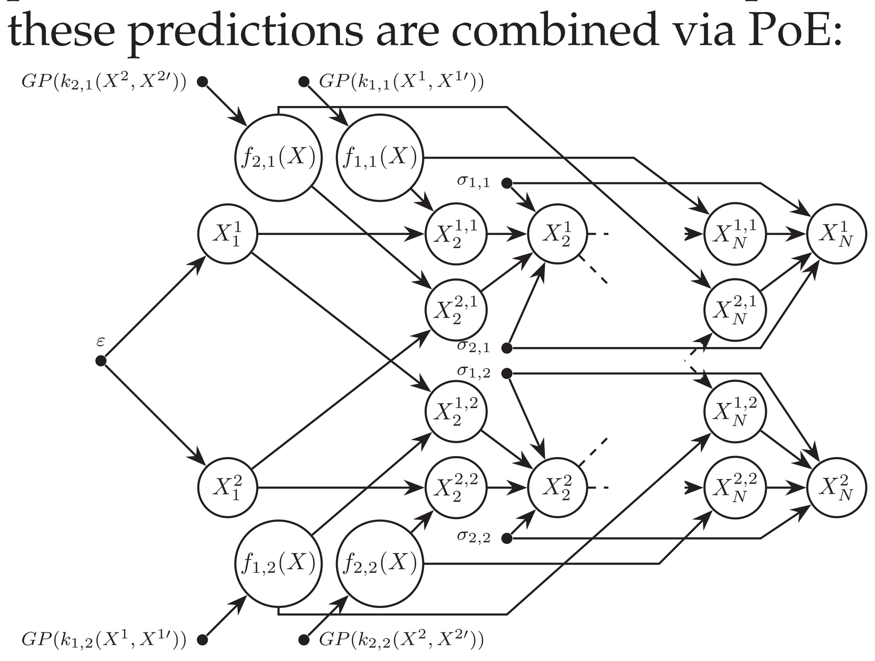

Product-of-Experts for Gaussian experts:

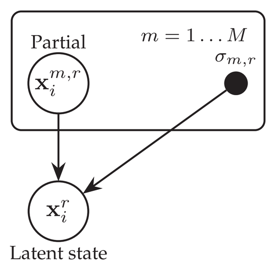

$$\sigma_r^2 = \left(\sum_m \sigma_{m,r}^{-2} \right)^{-1} \quad (1)$$

$$p(\mathbf{x}_i^r | \mathbf{x}_i^{:,r}, \sigma_{:,r}) \propto \prod_m \mathcal{N}\left(\mathbf{x}_i^r | \mathbf{x}_i^{m,r}, \sigma_{m,r}^2\right)$$

$$\Rightarrow p(\mathbf{x}_i^r | \mathbf{x}_i^{:,r}, \sigma_r) = \frac{\exp\left[-\frac{1}{2\sigma_r^2}\left(\mathbf{x}_i^r - \sigma_r^2 \sum_m \frac{\mathbf{x}_i^{m,r}}{\sigma_{m,r}^2}\right)^2\right]}{(2\pi\sigma_r^2)^{\frac{\dim(X^r)}{2}}} \quad (2)$$

Integrate out the intermediate predictions and get a simplified model:

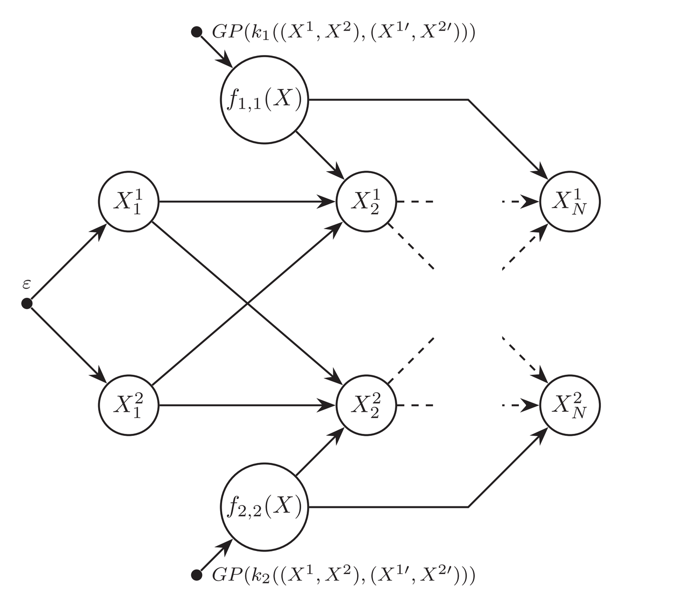

Covariance matrix of latent points $\mathbf{x}$:

$$\mathbf{\Sigma} + \mathbf{PKP}^T = \sigma_r^2 \mathbf{1}_{(N-1)\times(N-1)} + \sigma_r^4 \sum_{m=1}^{M} \frac{\mathbf{K}^m}{\sigma_{m,r}^4} \quad (3)$$

Kernel function k_r generating these covariance matrices:

$$k_r(X^{:}, X^{:\prime}) = \sigma_r^2 \delta(X^{:}, X^{:\prime}) + \sigma_r^4 \sum_{m=1}^{M} \frac{k_{m,r}(X^m, X^{m\prime})}{\sigma_{m,r}^4} \quad (4)$$

$\sigma_{m,r}$ can be increased to control the coupling. Coupling strength:

$$\sigma_{rel(i,j)}^2 = \frac{\sigma_{j,j}^4}{\sigma_{i,j}^4} \quad (5)$$

Augmented model

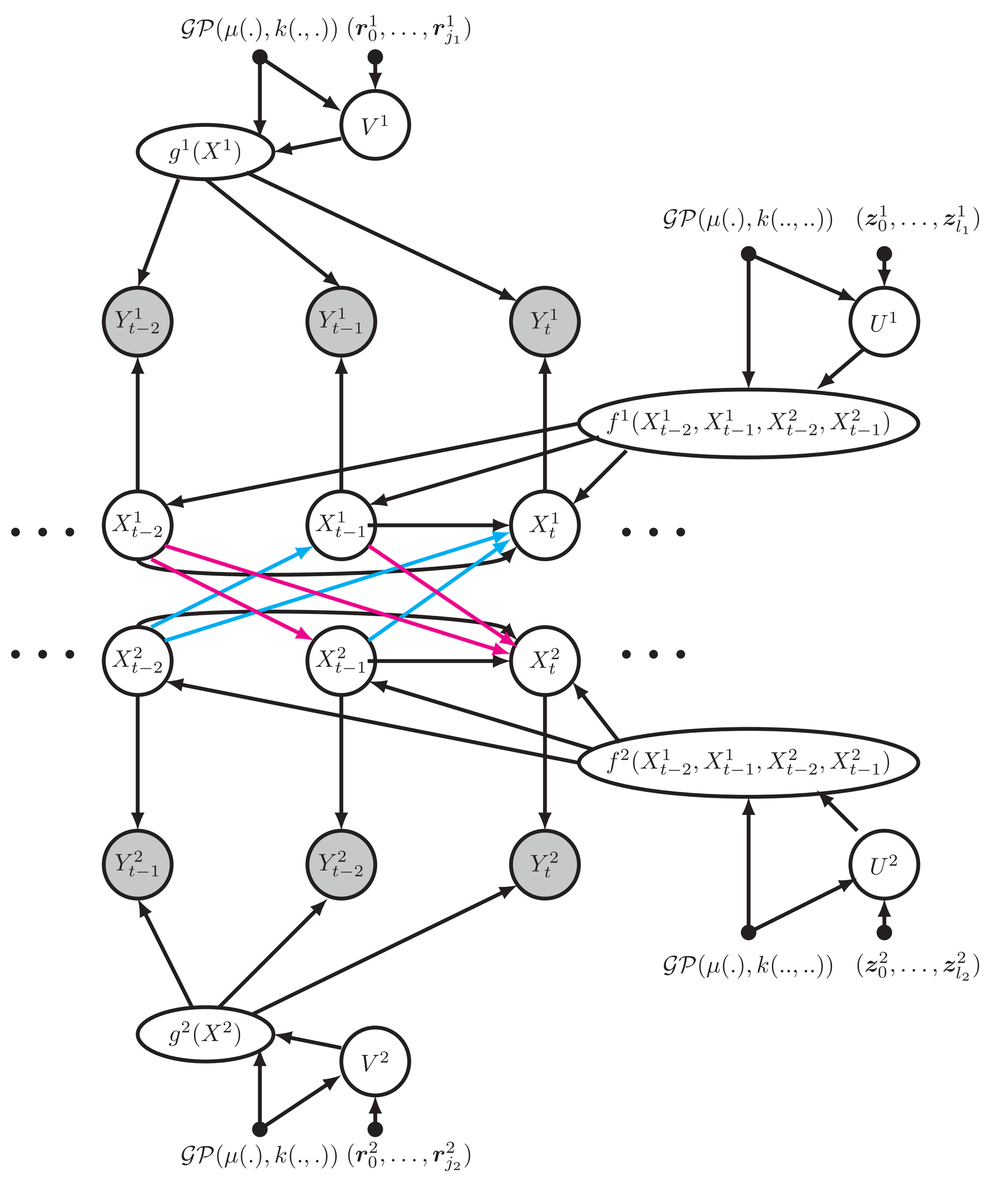

$\boldsymbol{Y}$ - observed data
$\boldsymbol{X}$ - latent points
$\boldsymbol{r}_i \to \boldsymbol{v}_i$ and $\boldsymbol{z}_j \to \boldsymbol{u}_j$ - augmenting mappings

For each part:

$$\boldsymbol{X} = \{\boldsymbol{x}_0 \dots \boldsymbol{x}_T\}; \boldsymbol{x}_t \in \mathbb{R}^Q$$
$$\boldsymbol{Y} = \{\boldsymbol{y}_0 \dots \boldsymbol{y}_T\}; \boldsymbol{y}_t \in \mathbb{R}^D$$
$$f(\boldsymbol{x}_{-t}) \sim \mathcal{GP}(0, k_f(\boldsymbol{x}_{-t}, \boldsymbol{x}'_{-t}))$$
$$\boldsymbol{f}_t = f(\boldsymbol{x}_{-t})$$
$$\boldsymbol{x}_t \sim \mathcal{N}(\boldsymbol{f}_t, \boldsymbol{I}\alpha^2)$$
$$g_d(\boldsymbol{X}) \sim \mathcal{GP}(0, k_g(\boldsymbol{x}_t, \boldsymbol{x}'_t))$$
$$\boldsymbol{g}_d = g_d(\boldsymbol{X})$$
$$\boldsymbol{Y}_{:,d} \sim \mathcal{N}(\boldsymbol{g}_d, \boldsymbol{I}\beta^2)$$
$$p(\boldsymbol{x}_0) = \mathcal{N}(\boldsymbol{x}_0 | 0, \boldsymbol{I})$$

Variational learning

Proposal variational distributions: $q(\boldsymbol{x}_t) = \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{x}_t}, \boldsymbol{S}_{\boldsymbol{x}_t})$. $q(\boldsymbol{u})$ and $q(\boldsymbol{v})$ are unconstrained distributions. The full variational joint proposal posterior distribution is:

$$q(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{f}, \boldsymbol{v}, \boldsymbol{g}) = p(\boldsymbol{y}|\boldsymbol{g}) p(\boldsymbol{g}|\boldsymbol{x}, \boldsymbol{v}) q(\boldsymbol{v}) \left[\prod_{t=1}^{T} p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t-1}, \boldsymbol{x}_{0:t-1}, \boldsymbol{u}) \right] q(\boldsymbol{x}) q(\boldsymbol{u})$$

The ELBO:

$$\log p(\boldsymbol{y}|\boldsymbol{\theta}) \geq \mathcal{L}(\boldsymbol{\theta}) = \int_{\boldsymbol{x},\boldsymbol{u},\boldsymbol{v},\boldsymbol{g},\boldsymbol{f}} q(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{f}, \boldsymbol{v}, \boldsymbol{g}) \log \frac{p(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{f}, \boldsymbol{v}, \boldsymbol{g}, \boldsymbol{y})}{q(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{f}, \boldsymbol{v}, \boldsymbol{g})} \quad (6)$$

Separating the latent dynamics part and applying the sufficient statistics assumption $q(\boldsymbol{x}, \boldsymbol{f}, \boldsymbol{u}) = q(\boldsymbol{u}) q(\boldsymbol{x}) \prod_{t=1}^{T} p(\boldsymbol{f}_t | \boldsymbol{x}_{t-1}, \boldsymbol{u})$ [Frigola et al., 2014], we get even much lower bound:

$$\mathcal{L}(\boldsymbol{\theta}) \geq \sum_{d=1}^{D} \int_{\boldsymbol{x},\boldsymbol{v},\boldsymbol{g}_d} p(\boldsymbol{g}_d|\boldsymbol{x}, \boldsymbol{v}) q(\boldsymbol{x}) q(\boldsymbol{v}) \log \frac{p(\boldsymbol{y}_d|\boldsymbol{g}_d)}{q(\boldsymbol{v})} + \log \frac{p(\boldsymbol{u})}{q(\boldsymbol{u})} d\boldsymbol{u} + H(q(\boldsymbol{x}))$$
$$+ \int q(\boldsymbol{u}) \left[\sum_{t=1}^{T} \int q(\boldsymbol{x}_t) q(\boldsymbol{x}_{-t}) \left(\int p(\boldsymbol{f}_t|\boldsymbol{x}_{-t}, \boldsymbol{u}) \log p(\boldsymbol{x}_t|\boldsymbol{f}_t) d\boldsymbol{f}_t \right) d\boldsymbol{x}_t d\boldsymbol{x}_{-t} \right]$$

where $H(q(\boldsymbol{x}))$ is the entropy of $q(\boldsymbol{x})$. The first integral is given in [Titsias and Lawrence, 2010]. The remaining integral is derived below. The innermost integral:

$$\mathcal{A} = \int p(\boldsymbol{f}_t|\boldsymbol{x}_{-t}, \boldsymbol{u}) \log p(\boldsymbol{x}_t|\boldsymbol{f}_t) d\boldsymbol{f}_t$$
$$= \log \mathcal{N}(\boldsymbol{x}_t | \boldsymbol{K}_{\boldsymbol{x}_{-t},\boldsymbol{z}} \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \boldsymbol{u}, \boldsymbol{I}\alpha) - \frac{1}{2} \operatorname{tr}(\alpha^{-1} \boldsymbol{K}_{\boldsymbol{x}_{-t},\boldsymbol{x}_{-t}} - \alpha^{-1} \boldsymbol{K}_{\boldsymbol{x}_{-t},\boldsymbol{z}} \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \boldsymbol{K}_{\boldsymbol{x}_{-t},\boldsymbol{z}})$$

The integrals over $d\boldsymbol{x}_t d\boldsymbol{x}_{-t}$:

$$\mathcal{B} = \int q(\boldsymbol{x}_t) q(\boldsymbol{x}_{-t}) \mathcal{A} d\boldsymbol{x}_t d\boldsymbol{x}_{-t} = -\log(\mathcal{Z}_{\boldsymbol{I}\alpha}) - \frac{1}{2} \alpha^{-1} \boldsymbol{u}^T \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \Psi_2(\boldsymbol{x}_{-t}) \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \boldsymbol{u}$$
$$+ \alpha^{-1} \boldsymbol{\mu}_{\boldsymbol{x}_t}^T \Psi_1(\boldsymbol{x}_{-t}) \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \boldsymbol{u} - \frac{1}{2} \alpha^{-1} \boldsymbol{\mu}_{\boldsymbol{x}_t}^T \boldsymbol{\mu}_{\boldsymbol{x}_t} - \frac{1}{2} \operatorname{tr}(\alpha^{-1} \Psi_0(\boldsymbol{x}_{-t}) - \alpha^{-1} \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \Psi_2(\boldsymbol{x}_{-t}))$$

where $\mathcal{Z}_\Sigma = \sqrt{|2\pi\Sigma|}$ is a normalization coefficient for multivariate Gaussian with Σ covariance. Then the sum over T can be written as a quadratic form:

$$\mathcal{C} = \sum_{t=1}^{T} \mathcal{B} = -\frac{1}{2} \boldsymbol{u}^T \mathcal{F} \boldsymbol{u} + \boldsymbol{u}^T \mathcal{G} + \mathcal{H} = -\frac{1}{2} (\boldsymbol{u} - \mathcal{F}^{-1}\mathcal{G})^T \mathcal{F} (\boldsymbol{u} - \mathcal{F}^{-1}\mathcal{G}) + \frac{1}{2} \mathcal{G}^T \mathcal{F}^{-1} \mathcal{G} + \mathcal{H}$$

$$\mathcal{F} = \alpha^{-1} \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \left[\sum_{t=1}^{T} \Psi_2(\boldsymbol{x}_{-t}) \right] \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \qquad \mathcal{G} = \left(\alpha^{-1} \left[\sum_{t=1}^{T} \boldsymbol{\mu}_{\boldsymbol{x}_t}^T \Psi_1(\boldsymbol{x}_{-t}) \right] \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \right)^T$$

$$\mathcal{H} = \sum_{t=1}^{T} \left[-\log(\mathcal{Z}_{\boldsymbol{I}\alpha}) - \frac{1}{2} \alpha^{-1} (\boldsymbol{\mu}_{\boldsymbol{x}_t}^T \boldsymbol{\mu}_{\boldsymbol{x}_t} + \operatorname{tr}(\boldsymbol{S}_{\boldsymbol{x}_t})) - \frac{1}{2} \operatorname{tr}(\alpha^{-1}(\Psi_0(\boldsymbol{x}_{-t}) - \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}}^{-1} \Psi_2(\boldsymbol{x}_{-t}))) \right]$$

The dynamics ELBO, also accounting for the optimal variational $q(\boldsymbol{u})$:

$$\mathcal{L}_{dyn}(\boldsymbol{\theta}) \geq \log \int p(\boldsymbol{u}) exp(\mathcal{C}) d\boldsymbol{u} + H(q(\boldsymbol{x})) = -\log \mathcal{Z}(\mathcal{F}^{-1} + \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}})$$
$$-\frac{1}{2} \mathcal{G}^T \mathcal{F}^{-1} (\mathcal{F}^{-1} + \boldsymbol{K}_{\boldsymbol{z},\boldsymbol{z}})^{-1} \mathcal{F}^{-1} \mathcal{G} + \log \mathcal{Z}(\mathcal{F}^{-1}) + \frac{1}{2} \mathcal{G}^T \mathcal{F}^{-1} \mathcal{G} + \mathcal{H} + H(q(\boldsymbol{x}))$$

Due to the factorising property of the PoE kernel, for multiple interacting parts total ELBO is the sum of ELBOs for each part.

Toy dataset

ARD RBF kernel, 2-nd order dynamics, 4 inducing inputs for dynamics mapping (circles), 8 inducing inputs for $\boldsymbol{X} \to \boldsymbol{Y}$ mapping (pluses)

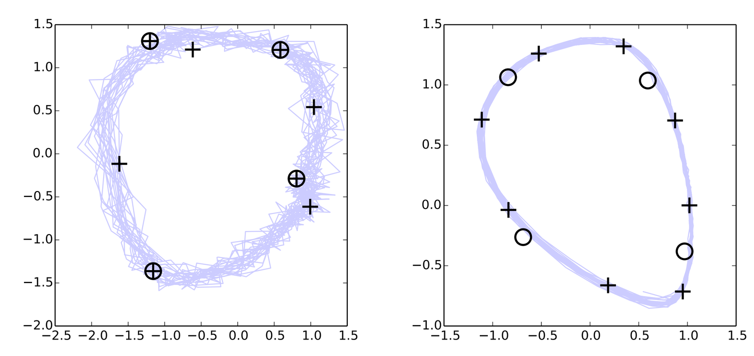

(a) (b)

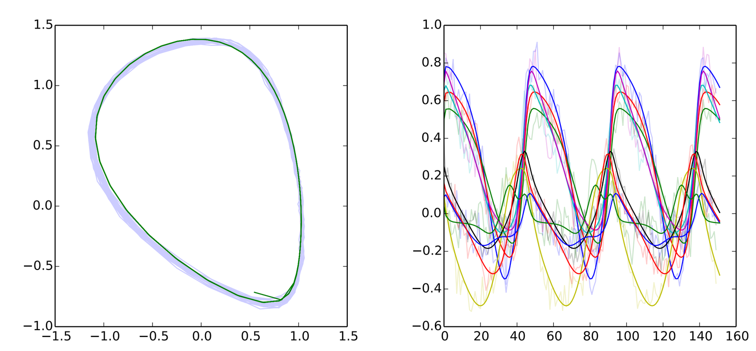

(c) (d)

PCA-initialised latents (a), optimized latent and inducing points (b). Mean-prediction was used to generate trajectories in latent (c) and observed (d) spaces overlaid with the noisy training data. Explained variance on the training data is ≈ 0.93, and ≈ 0.91 on testing data not used for training.

Summary

- Presented the variational learning for coupled dynamics in latent space, which are a type of modular MPs
- Examples with synthetic data

Future work

- Check the advantage of the variational learning vs. MAP
- Train on real MoCap datasets
- Sensory-motor integration as interaction in latent space dynamics
- Deep models

Acknowledgements

We acknowledge funding from DFG under IRTG 1901 'The Brain in Action' and the European Union Seventh Framework Program (FP7/2007 - 2013) under grant agreement no 611909 (KoroiBot)

References

[Bizzi et al., 2008] Bizzi, E., Cheung, V., d'Avella, A., Saltiel, P., and Tresch, M. (2008). Combining modules for movement. *Brain Research Reviews*, 57(1):125 – 133. Networks in Motion.

[Frigola et al., 2014] Frigola, R., Chen, Y., and Rasmussen, C. (2014). Variational gaussian process state-space models. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., and Weinberger, K., editors, *Advances in Neural Information Processing Systems 27*, pages 3680–3688. Curran Associates, Inc.

[Schaal, 2006] Schaal, S. (2006). Dynamic movement primitives -a framework for motor control in humans and humanoid robotics. In Kimura, H., Tsuchiya, K., Ishiguro, A., and Witte, H., editors, *Adaptive Motion of Animals and Machines*, pages 261–280. Springer Tokyo.

[Titsias, 2009] Titsias, M. K. (2009). Variational learning of inducing variables in sparse gaussian processes. In Dyk, D. A. V. and Welling, M., editors, *AISTATS*, volume 5 of *JMLR Proceedings*, pages 567–574. JMLR.org.

[Titsias and Lawrence, 2010] Titsias, M. K. and Lawrence, N. D. (2010). Bayesian gaussian process latent variable model. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, pages 844–851.

[Velychko et al., 2014] Velychko, D., Endres, D., Taubert, N., and Giese, M. A. (2014). Coupling Gaussian process dynamical models with product-of-experts kernels. In *Proceedings of the 24th International Conference on Artificial Neural Networks, LNCS 8681*, pages 603–610. Springer.