

A Detailed Variational Approximation

In this section, we repeat the derivation of the variational approximation in more detail.

Since exact inference in this model is intractable, we discuss a variational approximation to the model’s true marginal likelihood and posterior in this section. Analogously to \mathbf{y} , we denote the random vectors which contain the function values of the respective functions and outputs as \mathbf{a} and \mathbf{f} . The joint probability distribution of the data can then be written as

$$\begin{aligned} p(\mathbf{y}, \mathbf{f}, \mathbf{a} | \mathbf{X}) &= p(\mathbf{f} | \mathbf{a}) \prod_{d=1}^D p(\mathbf{y}_d | \mathbf{f}_d) p(\mathbf{a}_d | \mathbf{X}), \\ \mathbf{a}_d | \mathbf{X} &\sim \mathcal{N}(\mathbf{X}, \mathbf{K}_{\mathbf{a},d} + \sigma_{\mathbf{a},d}^2 \mathbf{I}), \\ \mathbf{f} | \mathbf{a} &\sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{\mathbf{f}} + \sigma_{\mathbf{f}}^2 \mathbf{I}), \\ \mathbf{y}_d | \mathbf{f}_d &\sim \mathcal{N}(\mathbf{f}_d, \mathbf{K}_{\mathbf{g},d} + \sigma_{\mathbf{y},d}^2 \mathbf{I}). \end{aligned} \tag{9}$$

Here, we use \mathbf{K} to refer to the Gram matrix corresponding to the kernel of the respective GP. All but the CPs factorize over both the different levels of the model as well as the different outputs.

To approximate a single deep GP, Hensman and Lawrence [11] proposed nested variational compression in which every GP in the hierarchy is handled independently. While this forces a variational approximation of all intermediate outputs of the stacked processes, it has the appealing property that it allows optimization via stochastic gradient descent [10] and the variational approximation can after training be used independently of the original training data.

A.1 Augmented Model

Nested variational compression focuses on augmenting a full GP model by introducing sets of *inducing variables* \mathbf{u} with their *inducing inputs* \mathbf{Z} . Those variables are assumed to be latent observations of the same functions and are thus jointly Gaussian with the observed data.

It can be written using its marginals [22] as

$$\begin{aligned} p(\hat{\mathbf{a}}, \mathbf{u}) &= \mathcal{N}(\hat{\mathbf{a}} | \boldsymbol{\mu}_{\mathbf{a}}, \boldsymbol{\Sigma}_{\mathbf{a}}) \mathcal{N}(\mathbf{u} | \mathbf{Z}, \mathbf{K}_{\mathbf{u}\mathbf{u}}), \text{ with} \\ \boldsymbol{\mu}_{\mathbf{a}} &= \mathbf{X} + \mathbf{K}_{\mathbf{a}\mathbf{u}} \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} (\mathbf{u} - \mathbf{Z}), \\ \boldsymbol{\Sigma}_{\mathbf{a}} &= \mathbf{K}_{\mathbf{a}\mathbf{a}} - \mathbf{K}_{\mathbf{a}\mathbf{u}} \mathbf{K}_{\mathbf{u}\mathbf{u}}^{-1} \mathbf{K}_{\mathbf{u}\mathbf{a}}, \end{aligned} \tag{10}$$

where, after dropping some indices and explicit conditioning on \mathbf{X} and \mathbf{Z} for clarity, $\hat{\mathbf{a}}$ denotes the function values $a_d(\mathbf{X})$ without noise and we write the Gram matrices as $\mathbf{K}_{\mathbf{a}\mathbf{u}} = k_{\mathbf{a},d}(\mathbf{X}, \mathbf{Z})$.

While the original model in (9) can be recovered exactly by marginalizing the inducing variables, considering a specific variational approximation of the joint $p(\hat{\mathbf{a}}, \mathbf{u})$ gives rise to the desired lower bound in the next subsection. A central assumption of this approximation [22] is that given enough inducing variables at the correct location, they are a sufficient statistic for $\hat{\mathbf{a}}$, implying conditional independence of the entries of $\hat{\mathbf{a}}$ given \mathbf{X} and \mathbf{u} . We introduce such inducing variables for every GP in the model, yielding the set $\{\mathbf{u}_{\mathbf{a},d}, \mathbf{u}_{\mathbf{f},d}, \mathbf{u}_{\mathbf{g},d}\}_{d=1}^D$ of inducing variables. Note that for the CP f , we introduce one set of inducing variables $\mathbf{u}_{f,d}$ per output f_d . These inducing variables play a crucial role in sharing information between the different outputs.

A.2 Variational Lower Bound

To derive the desired variational lower bound for the log marginal likelihood of the complete model, multiple steps are necessary. First, we will consider the innermost GPs a_d describing the alignment functions. We derive the Scalable Variational GP (SVGP), a lower bound for this model part which can be calculated efficiently and can be used for stochastic optimization, first introduced by Hensman, Fusi, and Lawrence [10]. In order to apply this bound recursively, we will both show how to propagate the uncertainty through the subsequent layers f_d and g_d and how to avoid the inter-layer cross-dependencies using another variational approximation as presented by Hensman and Lawrence [11]. While Hensman and Lawrence considered standard deep GP models, we will show how to apply their results to CPs.

The First Layer Since the inputs \mathbf{X} are fully known, we do not need to propagate uncertainty through the GPs a_d . Instead, the uncertainty about the \mathbf{a}_d comes from the uncertainty about the correct functions a_d and is introduced by the processes themselves. To derive a lower bound on the marginal log likelihood of \mathbf{a}_d , we assume a variational distribution $q(\mathbf{u}_{a,d}) \sim \mathcal{N}(\mathbf{m}_{a,d}, \mathbf{S}_{a,d})$ approximating $p(\mathbf{u}_{a,d})$ and additionally assume that $q(\hat{\mathbf{a}}_d, \mathbf{u}_{a,d}) = p(\hat{\mathbf{a}}_d | \mathbf{u}_{a,d}) q(\mathbf{u}_{a,d})$. After dropping the indices again, using Jensen’s inequality we get

$$\begin{aligned}
\log p(\mathbf{a} | \mathbf{X}) &= \log \int p(\mathbf{a} | \mathbf{u}) p(\mathbf{u}) d\mathbf{u} \\
&= \log \int q(\mathbf{u}) \frac{p(\mathbf{a} | \mathbf{u}) p(\mathbf{u})}{q(\mathbf{u})} d\mathbf{u} \\
&\geq \int q(\mathbf{u}) \log \frac{p(\mathbf{a} | \mathbf{u}) p(\mathbf{u})}{q(\mathbf{u})} d\mathbf{u} \\
&= \int \log p(\mathbf{a} | \mathbf{u}) q(\mathbf{u}) d\mathbf{u} - \int q(\mathbf{u}) \log \frac{q(\mathbf{u})}{p(\mathbf{u})} d\mathbf{u} \\
&= \mathbb{E}_{q(\mathbf{u})}[\log p(\mathbf{a} | \mathbf{u})] - \text{KL}(q(\mathbf{u}) \parallel p(\mathbf{u})),
\end{aligned} \tag{11}$$

where $\mathbb{E}_{q(\mathbf{u})}[\cdot]$ denotes the expected value with respect to the distribution $q(\mathbf{u})$ and $\text{KL}(\cdot \parallel \cdot)$ denotes the KL divergence, which can be evaluated analytically.

To bound the required expectation, we use Jensen’s inequality again together with (10) which gives

$$\begin{aligned}
\log p(\mathbf{a} | \mathbf{u}) &= \log \int p(\mathbf{a} | \hat{\mathbf{a}}) p(\hat{\mathbf{a}} | \mathbf{u}) d\hat{\mathbf{a}} \\
&= \log \int \mathcal{N}(\mathbf{a} | \hat{\mathbf{a}}, \sigma_a^2 \mathbf{I}) \mathcal{N}(\hat{\mathbf{a}} | \boldsymbol{\mu}_a, \boldsymbol{\Sigma}_a) d\hat{\mathbf{a}} \\
&\geq \int \log \mathcal{N}(\mathbf{a} | \hat{\mathbf{a}}, \sigma_a^2 \mathbf{I}) \mathcal{N}(\hat{\mathbf{a}} | \boldsymbol{\mu}_a, \boldsymbol{\Sigma}_a) d\hat{\mathbf{a}} \\
&= \log \mathcal{N}(\mathbf{a} | \boldsymbol{\mu}_a, \sigma_a^2 \mathbf{I}) - \frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a).
\end{aligned} \tag{12}$$

We apply this bound to the expectation to get

$$\mathbb{E}_{q(\mathbf{u})}[\log p(\mathbf{a} | \mathbf{u})] \geq \mathbb{E}_{q(\mathbf{u})}[\log \mathcal{N}(\mathbf{a} | \boldsymbol{\mu}_a, \sigma_a^2 \mathbf{I})] - \frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a), \text{ with} \tag{13}$$

$$\begin{aligned}
\mathbb{E}_{q(\mathbf{u})}[\log \mathcal{N}(\mathbf{a} | \boldsymbol{\mu}_a, \sigma_a^2 \mathbf{I})] &= \log \mathcal{N}(\mathbf{a} | \mathbf{K}_{au} \mathbf{K}_{uu}^{-1} \mathbf{m}, \sigma_a^2 \mathbf{I}) \\
&\quad + \frac{1}{2\sigma_a^2} \text{tr}(\mathbf{K}_{au} \mathbf{K}_{uu}^{-1} \mathbf{S} \mathbf{K}_{uu}^{-1} \mathbf{K}_{ua}).
\end{aligned} \tag{14}$$

Resubstituting this result into (11) yields the final bound

$$\begin{aligned}
\log p(\mathbf{a} | \mathbf{X}) &\geq \log \mathcal{N}(\mathbf{a} | \mathbf{K}_{au} \mathbf{K}_{uu}^{-1} \mathbf{m}, \sigma_a^2 \mathbf{I}) - \text{KL}(q(\mathbf{u}) \parallel p(\mathbf{u})) \\
&\quad - \frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a) - \frac{1}{2\sigma_a^2} \text{tr}(\mathbf{K}_{au} \mathbf{K}_{uu}^{-1} \mathbf{S} \mathbf{K}_{uu}^{-1} \mathbf{K}_{ua}).
\end{aligned} \tag{15}$$

This bound, which depends on the hyper parameters of the kernel and likelihood $\{\boldsymbol{\theta}, \sigma_a\}$ and the variational parameters $\{\mathbf{Z}, \mathbf{m}, \mathbf{S}\}$, can be calculated in $\mathcal{O}(NM^2)$ time. It factorizes along the data points which enables stochastic optimization.

In order to obtain a bound on the full model, we apply the same techniques to the other processes. Since the alignment processes a_d are assumed to be independent, we have $\log p(\mathbf{a}_1, \dots, \mathbf{a}_D | \mathbf{X}) = \sum_{d=1}^D \log p(\mathbf{a}_d | \mathbf{X})$, where every term can be approximated using the bound in (15). However, for all subsequent layers, the bound is not directly applicable, since the inputs are no longer known but instead are given by the outputs of the previous process. It is therefore necessary to propagate their uncertainty and also handle the interdependencies between the layers introduced by the latent function values \mathbf{a} , \mathbf{f} and \mathbf{g} .

The Second and Third Layer Our next goal is to derive a bound on the outputs of the second layer

$$\log p(\mathbf{f} | \mathbf{u}_f) = \log \int p(\mathbf{f}, \mathbf{a}, \mathbf{u}_a | \mathbf{u}_f) d\mathbf{a} d\mathbf{u}_a, \quad (16)$$

that is, an expression in which the uncertainty about the different \mathbf{a}_d and the cross-layer dependencies on the $\mathbf{u}_{a,d}$ are both marginalized. While on the first layer, the different \mathbf{a}_d are conditionally independent, the second layer explicitly models the cross-covariances between the different outputs via convolutions over the shared latent processes w_r . We will therefore need to handle all of the different \mathbf{f}_d , together denoted as \mathbf{f} , at the same time.

We start by considering the relevant terms from (9) and apply (12) to marginalize \mathbf{a} in

$$\begin{aligned} \log p(\mathbf{f} | \mathbf{u}_f, \mathbf{u}_a) &= \log \int p(\mathbf{f}, \mathbf{a} | \mathbf{u}_f, \mathbf{u}_a) d\mathbf{a} \\ &\geq \log \int \tilde{p}(\mathbf{f} | \mathbf{u}_f, \mathbf{a}) \tilde{p}(\mathbf{a} | \mathbf{u}_a) \cdot \exp\left(-\frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a) - \frac{1}{2\sigma_f^2} \text{tr}(\boldsymbol{\Sigma}_f)\right) d\mathbf{a} \\ &\geq \mathbb{E}_{\tilde{p}(\mathbf{a} | \mathbf{u}_a)}[\log \tilde{p}(\mathbf{f} | \mathbf{u}_f, \mathbf{a})] - \mathbb{E}_{\tilde{p}(\mathbf{a} | \mathbf{u}_a)}\left[\frac{1}{2\sigma_f^2} \text{tr}(\boldsymbol{\Sigma}_f)\right] - \frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a), \end{aligned} \quad (17)$$

where we write $\tilde{p}(\mathbf{a} | \mathbf{u}_a) = \mathcal{N}(\mathbf{a} | \boldsymbol{\mu}_a, \sigma_a^2 \mathbf{I})$ to incorporate the Gaussian noise in the latent space. Due to our assumption that \mathbf{u}_a is a sufficient statistic for \mathbf{a} we choose

$$\begin{aligned} q(\mathbf{a} | \mathbf{u}_a) &= \tilde{p}(\mathbf{a} | \mathbf{u}_a), \text{ and} \\ q(\mathbf{a}) &= \int \tilde{p}(\mathbf{a} | \mathbf{u}_a) q(\mathbf{u}_a) d\mathbf{u}_a, \end{aligned} \quad (18)$$

and use another variational approximation to marginalize \mathbf{u}_a . This yields

$$\begin{aligned} \log p(\mathbf{f} | \mathbf{u}_f) &= \log \int p(\mathbf{f}, \mathbf{u}_a | \mathbf{u}_f) d\mathbf{u}_a \\ &= \log \int p(\mathbf{f} | \mathbf{u}_f, \mathbf{u}_a) p(\mathbf{u}_a) d\mathbf{u}_a \\ &\geq \int q(\mathbf{u}_a) \log \frac{p(\mathbf{f} | \mathbf{u}_f, \mathbf{u}_a) p(\mathbf{u}_a)}{q(\mathbf{u}_a)} d\mathbf{u}_a \\ &= \mathbb{E}_{q(\mathbf{u}_a)}[\log p(\mathbf{f} | \mathbf{u}_a, \mathbf{u}_f)] - \text{KL}(q(\mathbf{u}_a) \| p(\mathbf{u}_a)) \\ &\geq \mathbb{E}_{q(\mathbf{u}_a)}[\mathbb{E}_{\tilde{p}(\mathbf{a} | \mathbf{u}_a)}[\log \tilde{p}(\mathbf{f} | \mathbf{u}_f, \mathbf{a})]] - \text{KL}(q(\mathbf{u}_a) \| p(\mathbf{u}_a)) \\ &\quad - \frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a) - \mathbb{E}_{q(\mathbf{u}_a)}\left[\mathbb{E}_{\tilde{p}(\mathbf{a} | \mathbf{u}_a)}\left[\frac{1}{2\sigma_f^2} \text{tr}(\boldsymbol{\Sigma}_f)\right]\right] \\ &\geq \mathbb{E}_{q(\mathbf{a})}[\log \tilde{p}(\mathbf{f} | \mathbf{u}_f, \mathbf{a})] - \text{KL}(q(\mathbf{u}_a) \| p(\mathbf{u}_a)) \\ &\quad - \frac{1}{2\sigma_a^2} \text{tr}(\boldsymbol{\Sigma}_a) - \frac{1}{2\sigma_f^2} \mathbb{E}_{q(\mathbf{a})}[\text{tr}(\boldsymbol{\Sigma}_f)], \end{aligned} \quad (19)$$

where we apply Fubini's theorem to exchange the order of integration in the expected values. The expectations with respect to $q(\mathbf{a})$ involve expectations of kernel matrices, also called Ψ -statistics, in the same way as in [8] and are given by

$$\begin{aligned} \psi_f &= \mathbb{E}_{q(\mathbf{a})}[\text{tr}(\mathbf{K}_{ff})], \\ \boldsymbol{\Psi}_f &= \mathbb{E}_{q(\mathbf{a})}[\mathbf{K}_{fu}], \\ \boldsymbol{\Phi}_f &= \mathbb{E}_{q(\mathbf{a})}[\mathbf{K}_{uf} \mathbf{K}_{fu}]. \end{aligned} \quad (20)$$

These Ψ -statistics can be computed analytically for multiple kernels, including the squared exponential kernel. In Appendix A.3 we show closed-form solutions for these Ψ -statistics for the implicit kernel defined in the CP layer. To obtain the final formulation of the desired bound for $\log p(\mathbf{f} | \mathbf{u}_f)$

we substitute (20) into (19) and get the analytically tractable bound

$$\begin{aligned}
\log p(\mathbf{f} | \mathbf{u}_f) &\geq \log \mathcal{N}\left(\mathbf{f} \mid \Psi_f \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1} \mathbf{m}_f, \sigma_f^2 \mathbf{I}\right) - \text{KL}(q(\mathbf{u}_a) \parallel p(\mathbf{u}_a)) - \frac{1}{2\sigma_a^2} \text{tr}(\Sigma_a) \\
&\quad - \frac{1}{2\sigma_f^2} \left(\psi_f - \text{tr}\left(\Psi_f \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1}\right) \right) \\
&\quad - \frac{1}{2\sigma_f^2} \text{tr}\left(\left(\Phi_f - \Psi_f^\top \Psi_f\right) \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1} \left(\mathbf{m}_f \mathbf{m}_f^\top + \mathbf{S}_f\right) \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1}\right)
\end{aligned} \tag{21}$$

The uncertainties in the first layer have been propagated variationally to the second layer. Besides the regularization terms, $\mathbf{f} | \mathbf{u}_f$ is a Gaussian distribution. Because of their cross dependencies, the different outputs \mathbf{f}_d are considered in a common bound and do not factorize along dimensions. The third layer warpings \mathbf{g}_d however are conditionally independent given \mathbf{f} and can therefore be considered separately. In order to derive a bound for $\log p(\mathbf{y} | \mathbf{u}_g)$ we apply the same steps as described above, resulting in the final bound, which factorizes along the data, allowing for stochastic optimization methods:

$$\begin{aligned}
\log p(\mathbf{y} | \mathbf{X}) &\geq \sum_{d=1}^D \log \mathcal{N}\left(\mathbf{y}_d \mid \Psi_{g,d} \mathbf{K}_{\mathbf{u}_{g,d} \mathbf{u}_{g,d}}^{-1} \mathbf{m}_{g,d}, \sigma_{y,d}^2 \mathbf{I}\right) - \sum_{d=1}^D \frac{1}{2\sigma_{a,d}^2} \text{tr}(\Sigma_{a,d}) \\
&\quad - \frac{1}{2\sigma_f^2} \left(\psi_f - \text{tr}\left(\Phi_f \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1}\right) \right) - \sum_{d=1}^D \frac{1}{2\sigma_{y,d}^2} \left(\psi_{g,d} - \text{tr}\left(\Phi_{g,d} \mathbf{K}_{\mathbf{u}_{g,d} \mathbf{u}_{g,d}}^{-1}\right) \right) \\
&\quad - \sum_{d=1}^D \text{KL}(q(\mathbf{u}_{a,d}) \parallel p(\mathbf{u}_{a,d})) - \text{KL}(q(\mathbf{u}_f) \parallel p(\mathbf{u}_f)) - \sum_{d=1}^D \text{KL}(q(\mathbf{u}_{y,d}) \parallel p(\mathbf{u}_{y,d})) \\
&\quad - \frac{1}{2\sigma_f^2} \text{tr}\left(\left(\Phi_f - \Psi_f^\top \Psi_f\right) \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1} \left(\mathbf{m}_f \mathbf{m}_f^\top + \mathbf{S}_f\right) \mathbf{K}_{\mathbf{u}_f \mathbf{u}_f}^{-1}\right) \\
&\quad - \sum_{d=1}^D \frac{1}{2\sigma_{y,d}^2} \text{tr}\left(\left(\Phi_{g,d} - \Psi_{g,d}^\top \Psi_{g,d}\right) \mathbf{K}_{\mathbf{u}_{g,d} \mathbf{u}_{g,d}}^{-1} \left(\mathbf{m}_{g,d} \mathbf{m}_{g,d}^\top + \mathbf{S}_{g,d}\right) \mathbf{K}_{\mathbf{u}_{g,d} \mathbf{u}_{g,d}}^{-1}\right)
\end{aligned} \tag{22}$$

A.3 Convolution Kernel Expectations

In Section 2 we assumed the latent processes w_r to be white noise processes and the smoothing kernel functions $T_{d,r}$ to be squared exponential kernels, leading to an explicit closed form formulation for the covariance between outputs shown in (3). In this section, we derive the Ψ -statistics for this generalized squared exponential kernel needed to evaluate (22).

The uncertainty about the first layer is captured by the variational distribution of the latent alignments \mathbf{a} given by $q(\mathbf{a}) \sim \mathcal{N}(\boldsymbol{\mu}_a, \Sigma_a)$, with $\mathbf{a} = (\mathbf{a}_1, \dots, \mathbf{a}_d)$. Every aligned point in \mathbf{a} corresponds to one output of \mathbf{f} and ultimately to one of the \mathbf{y}_d . Since the closed form of the multi output kernel depends on the choice of outputs, we will use the notation $\hat{f}(\mathbf{a}_n)$ to denote $f_d(\mathbf{a}_n)$ such that \mathbf{a}_n is associated with output d .

For notational simplicity, we only consider the case of one single latent process w_r . Since the latent processes are independent, the results can easily be generalized to multiple processes. Then, ψ_f is given by

$$\begin{aligned}
\psi_f &= \mathbb{E}_{q(\mathbf{a})}[\text{tr}(\mathbf{K}_f \mathbf{f} \mathbf{f})] \\
&= \sum_{n=1}^N \mathbb{E}_{q(\mathbf{a}_n)} \left[\text{cov} \left[\hat{f}(\mathbf{a}_n), \hat{f}(\mathbf{a}_n) \right] \right] \\
&= \sum_{n=1}^N \int \text{cov} \left[\hat{f}(\mathbf{a}_n), \hat{f}(\mathbf{a}_n) \right] q(\mathbf{a}_n) d\mathbf{a}_n \\
&= \sum_{n=1}^N \hat{\sigma}_{nm}^2.
\end{aligned} \tag{23}$$

Similar to the notation $\hat{f}(\cdot)$, we use the notation $\hat{\sigma}_{nn'}$ to mean the variance term associated with the covariance function $\text{cov}[\hat{f}(\mathbf{a}_n), \hat{f}(\mathbf{a}_{n'})]$. The expectation $\Psi_f = \mathbb{E}_{q(\mathbf{a})}[\mathbf{K}_f \mathbf{u}]$ connecting the alignments and the pseudo inputs is given by

$$\begin{aligned} \Psi_f &= \mathbb{E}_{q(\mathbf{a})}[\mathbf{K}_f \mathbf{u}], \text{ with} \\ (\Psi_f)_{ni} &= \int \text{cov}[\hat{f}(\mathbf{a}_n), \hat{f}(\mathbf{Z}_i)] q(\mathbf{a}_n) d\mathbf{a}_n \\ &= \hat{\sigma}_{ni}^2 \sqrt{\frac{(\Sigma_{\mathbf{a}})_{nn}^{-1}}{\hat{\ell}_{ni} + (\Sigma_{\mathbf{a}})_{nn}^{-1}}} \cdot \exp\left(-\frac{1}{2} \frac{(\Sigma_{\mathbf{a}})_{nn}^{-1} \hat{\ell}_{ni}}{(\Sigma_{\mathbf{a}})_{nn}^{-1} + \hat{\ell}_{ni}} ((\boldsymbol{\mu}_{\mathbf{a}})_n - \mathbf{Z}_i)^2\right) \end{aligned} \quad (24)$$

where $\hat{\ell}_{ni}$ is the combined length scale corresponding to the same kernel as $\hat{\sigma}_{ni}$. Lastly, $\Phi_f = \mathbb{E}_{q(\mathbf{a})}[\mathbf{K}_{\mathbf{u}f} \mathbf{K}_f \mathbf{u}]$ connects alignments and pairs of pseudo inputs with the closed form

$$\begin{aligned} \Phi_f &= \mathbb{E}_{q(\mathbf{a})}[\mathbf{K}_{\mathbf{u}f} \mathbf{K}_f \mathbf{u}], \text{ with} \\ (\Phi_f)_{ij} &= \sum_{n=1}^N \int \text{cov}[\hat{f}(\mathbf{a}_n), \hat{f}(\mathbf{Z}_i)] \cdot \text{cov}[\hat{f}(\mathbf{a}_n), \hat{f}(\mathbf{Z}_j)] q(\mathbf{a}_n) d\mathbf{a}_n \\ &= \sum_{n=1}^N \hat{\sigma}_{ni}^2 \hat{\sigma}_{nj}^2 \sqrt{\frac{(\Sigma_{\mathbf{a}})_{nn}^{-1}}{\hat{\ell}_{ni} + \hat{\ell}_{nj} + (\Sigma_{\mathbf{a}})_{nn}^{-1}}} \cdot \exp\left(-\frac{1}{2} \frac{\hat{\ell}_{ni} \hat{\ell}_{nj}}{\hat{\ell}_{ni} + \hat{\ell}_{nj}} (\mathbf{Z}_i - \mathbf{Z}_j)^2\right) \\ &\quad - \frac{1}{2} \frac{(\Sigma_{\mathbf{a}})_{nn}^{-1} (\hat{\ell}_{ni} + \hat{\ell}_{nj})}{(\Sigma_{\mathbf{a}})_{nn}^{-1} + \hat{\ell}_{ni} + \hat{\ell}_{nj}} \cdot \left((\boldsymbol{\mu}_{\mathbf{a}})_n - \frac{\hat{\ell}_{ni} \mathbf{Z}_i + \hat{\ell}_{nj} \mathbf{Z}_j}{\hat{\ell}_{ni} + \hat{\ell}_{nj}} \right)^2 \end{aligned} \quad (25)$$

Note that the Ψ -statistics factorize along the data and we only need to consider the diagonal entries of $\Sigma_{\mathbf{a}}$. If all the data belong to the same output, the Ψ -statistics of the standard squared exponential kernel can be recovered as a special case. It is used to propagate the uncertainties through the output-specific warpings g .

A.4 Approximative Predictions

Using the variational lower bound in (5), our model can be fitted to data, resulting in appropriate choices of the kernel hyper parameters and variational parameters. Now assume we want to predict approximate function values $g_{d,\star}$ for previously unseen points $\mathbf{X}_{d,\star}$ associated with output d , which are given by $g_{d,\star} = g_d(f_d(a_d(\mathbf{X}_{d,\star})))$.

Because of the conditional independence assumptions in the model, other outputs $d' \neq d$ only have to be considered in the shared layer f . In this shared layer, the belief about the different outputs and the shared information and is captured by the variational distribution $q(\mathbf{u}_f)$. Given $q(\mathbf{u}_f)$, the different outputs are conditionally independent of one another and thus, predictions for a single dimension in our model are equivalent to predictions in a single deep GP with nested variational compression as presented by Hensman and Lawrence [11].

B Joint models for wind experiment

In the following, we show plots with joint predictions for the models discussed in Section 4.2. Similar to Section 4.1, we trained a standard GP in Figure 6, a multi-output GP in Figure 7, a deep GP in Figure 8 and our model in Figure 9. All models were trained until convergence and multiple runs result in very similar models. For all models we used RBF kernels or dependent RBF kernels where applicable.

Each plot shows the data in gray and two mean predictions and uncertainty bands. The first violet uncertainty band is the result of the variational approximation of the respective model. The second green or blue posterior is obtained via sampling. For both the GP and MO-GP, we used the SVGP approximation [12] and since the models are shallow, the approximation is almost exact.

Figure 8 showcases the difficulty of training a deep GP model and the shortcomings of the nested variational compression. The violet variational approximation is used for training and approximates

the data comparatively well. As discussed above, the deep GP cannot share information, so the test sets cannot be predicted. However, as discussed in more detail in [12], the approximation tends to underestimate uncertainties when propagating them through the different layers and because of this, uncertainties obtained via sampling tend to vary considerably more. Because during model selection sample performance does not matter, the true posterior can be (and in this case is) considerably different.

Our approach in principle has the same problem as the deep GP. However, because of the strong interpretability of the different parts of the hierarchy, uncertainties within the model are never placed arbitrarily and because of this, the variational posteriors and true posteriors look much more similar. They tend to disagree in places when there is high uncertainty about the alignment.

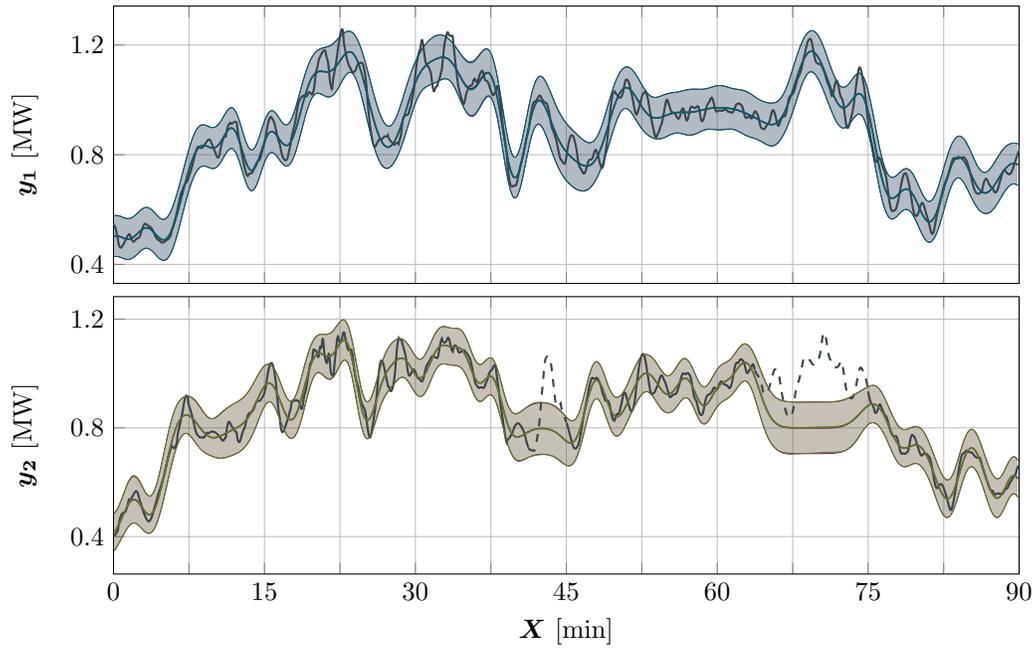


Figure 6: GP

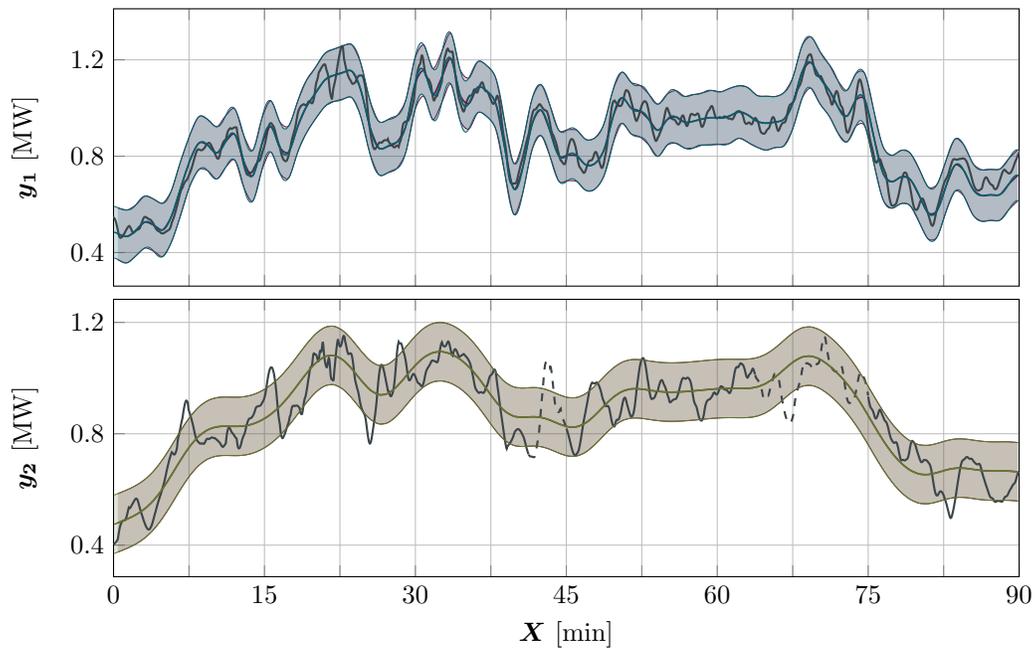


Figure 7: MO-GP

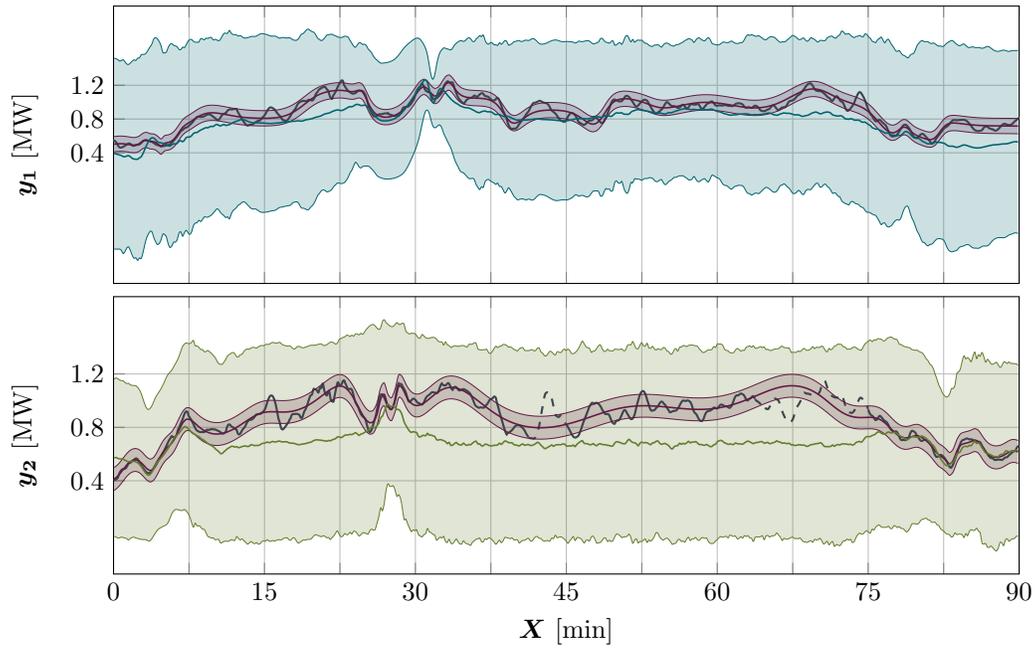


Figure 8: DGP

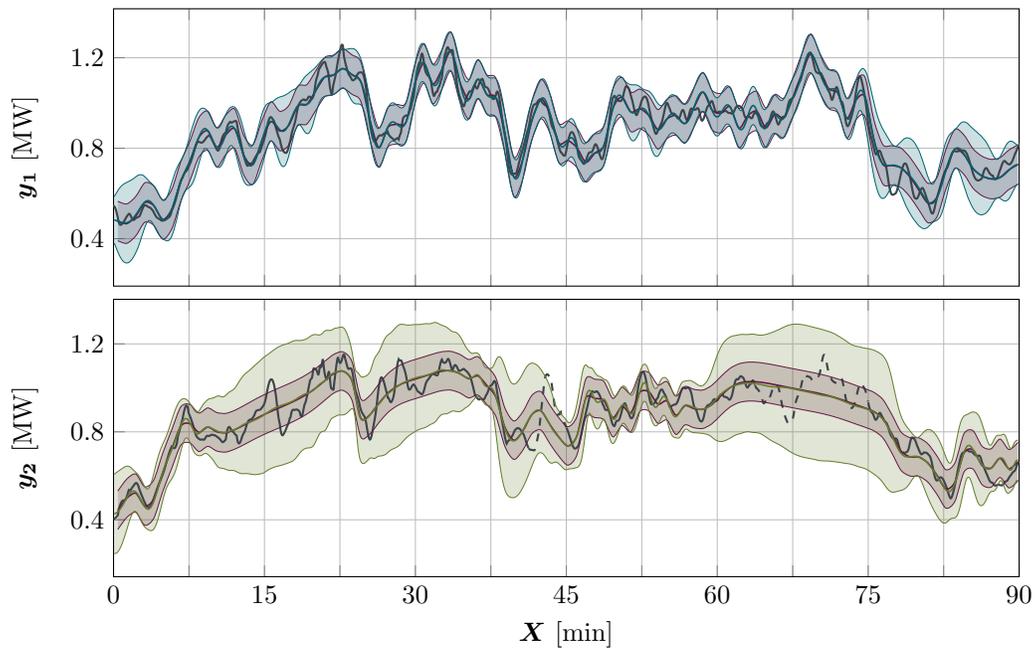


Figure 9: AMO-GP (Ours)