

## A Interpretation of lengthscale parameters

First, we start with the stationary case. There exist stationary kernels  $k$  which can be represented as:

$$k(\mathbf{a}, \mathbf{b}) = \pi_k((\mathbf{a} - \mathbf{b})^\top \mathbf{\Delta}^{-1}(\mathbf{a} - \mathbf{b})), \quad (10)$$

for a given scalar function  $\pi_k$  and lengthscale parameter  $\mathbf{\Delta}$ . In this case, the lengthscale controls the spatial variation of the Gaussian process with that kernel. More concretely, in the 1D case, using scalar lengthscale  $\mathbf{\Delta} = \ell^2$  and the squared exponential kernel:

$$k_{\text{SE}}(a, b) = \exp \left[ -\frac{1}{2} \frac{(a - b)^2}{\ell^2} \right], \quad (11)$$

we have that the corresponding marginals for  $f \sim \mathcal{GP}(0(\cdot), k_{\text{SE}})$  are:

$$f(x) \sim \text{N}(0, 1), \quad (12)$$

$$\frac{d}{dx} f(x) \sim \text{N}\left(0, \frac{1}{\ell^2}\right). \quad (13)$$

So the lengthscale parameter directly controls the amplitude of the gradient's variance.

In general, non-stationary kernels do not have a corresponding concept, which attention is given to kernels that, in the neighborhood of a point  $\mathbf{x}$ , can be expressed in terms of a local lengthscale matrix  $\mathbf{\Delta}(\mathbf{x})$ . The lengthscale mixture kernels  $k_{\text{lmx}}$  described in Section 2 and our proposed kernel from Section 3,  $k_{\text{TDGP}}$  both have this local lengthscale property.

Again, assuming inputs are 1D and the base kernel is squared exponential, both kernels are:

$$k_{\text{lmx}}(a, b) = \ell(a)^{\frac{2}{4}} \ell(b)^{\frac{2}{4}} \left[ \frac{\ell(a)^2 + \ell(b)^2}{2} \right]^{-\frac{1}{2}} \exp \left[ -\frac{1}{2} \frac{(a - b)^2}{\frac{\ell(a)^2 + \ell(b)^2}{2}} \right],$$

$$k_{\text{TDGP}}(a, b) = \exp \left[ -\frac{1}{2} \left( \frac{a}{\ell(a)} - \frac{b}{\ell(b)} \right)^2 \right].$$

And in terms of derivatives:

$$\frac{d}{dx} f_{\text{lmx}}(x) \sim \text{N}\left(0, \frac{2 + \frac{d}{dx} \ell(x)^2}{2\ell(x)^2}\right), \quad (14)$$

$$\frac{d}{dx} f_{\text{TDGP}}(x) \sim \text{N}\left(0, \frac{(\ell(x) - x \frac{d}{dx} \ell(x))^2}{\ell(x)^4}\right). \quad (15)$$

Both kernels generalize Eq. (13) and, as expected, recover the stationary case when  $\frac{d}{dx} \ell(x) = 0$ .

Note that the domain of the lengthscale function  $\ell(x)$  is always the domain of the function  $f$ , meaning that as we consider deeper models, the lengthscale function is always a function of the original domain. This is unlike the general compositional case, e.g.  $f(g(h(x)))$ , where the domain of each individual function is the image of the previous function. Moreover, the relationship with the lengthscale parameter and derivative of the output function remains clear.

## B TDGP and CDGP are limits of a generalized DGP

**Theorem 3.1.** Any  $L$ -layer CDGP prior over a function  $f(\mathbf{x}) = h^L(\mathbf{h}^{L-1}(\dots \mathbf{h}^1(\mathbf{x}) \dots))$  is a special case of a TDGP prior with equal depth defined over the augmented input-space  $\tilde{\mathbf{x}} = [\mathbf{x}, 1]^\top$ . Since linear deformations  $\tilde{\mathbf{W}}\tilde{\mathbf{x}}$  in the augmented space correspond to affine transformations  $\mathbf{W}\mathbf{x} + \mathbf{d}$  in the original space, the special case of the CDGP model corresponds to a TDGP where the prior variance of the  $\mathbf{W}^\ell$  approaches zero.

*Proof.* First, we append a bias to the input data:  $\tilde{\mathbf{x}} = [\mathbf{x} \ 1]^\top \in \mathbb{R}^{D+1}$ . Therefore, the hidden-layer matrices need to be expanded so that  $\tilde{\mathbf{W}}(\tilde{\mathbf{h}}) \in \mathbb{R}^{(Q+1) \times (D+1)}$ . Then, choose the following form for

$\tilde{\mathbf{W}}$ :

$$\tilde{\mathbf{W}}(\tilde{\mathbf{h}}) = \begin{bmatrix} \mathbf{W}(\mathbf{h}) & \mathbf{d}(\mathbf{h}) \\ \mathbf{0}_{1 \times Q} & 1 \end{bmatrix}, \quad (16)$$

where,

$$\mathbf{W}(\cdot) \in \mathbb{R}^{Q \times D} \sim \prod_{q=1}^Q \prod_{d=1}^D \mathcal{GP}(w_{qd}(\cdot) \mid 0(\cdot), k_{w_{qd}}), \quad (17)$$

$$\mathbf{d}(\mathbf{x}) \in \mathbb{R}^Q \sim \prod_{q=1}^Q \mathcal{GP}(d_q(\cdot) \mid \mu_{d_q}, k_{d_q}). \quad (18)$$

Note that  $\tilde{\mathbf{W}}(\cdot)$  still follows the TDGP prior because all of its entries are either GP distributed or limits of GP priors, like in the case of the lower row where the Dirac delta distributions can be obtained by taking the limit of the kernel variance parameter  $\sigma^2$  to zero.

Then, the  $\ell$ -th latent space of this model is:

$$\tilde{\mathbf{h}}^\ell = \tilde{\mathbf{W}}(\tilde{\mathbf{h}}^{\ell-1}) \tilde{\mathbf{x}} \quad (19)$$

$$= \begin{bmatrix} \mathbf{W}(\mathbf{h}^{\ell-1}) & \mathbf{d}(\mathbf{h}^{\ell-1}) \\ \mathbf{0}_{1 \times Q} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ 1 \end{bmatrix} \quad (20)$$

$$= \begin{bmatrix} \mathbf{W}(\mathbf{h}^{\ell-1}) \mathbf{x} + \mathbf{d}(\mathbf{h}^{\ell-1}) \cdot 1 \\ 0 \cdot \mathbf{x} + 1 \end{bmatrix} \quad (21)$$

$$= [\mathbf{W}(\mathbf{h}^{\ell-1}) \mathbf{x} + \mathbf{d}(\mathbf{h}^{\ell-1}) \quad 1]^\top \quad (22)$$

$$= [\mathbf{h}^\ell \quad 1]^\top. \quad (23)$$

By ignoring the bias dimension, we get a latent space  $\mathbf{h}$  that includes a multiplicative component  $\mathbf{W}(\cdot)$  and an additive component  $\mathbf{d}(\cdot)$ . If the prior variance of  $\mathbf{W}(\cdot)$  goes to zero, which is controlled by the kernel variance hyperparameter, then  $\mathbf{W}(\cdot) \rightarrow \mathbf{0}$ , resulting in only the additive component remaining:

$$\mathbf{h}^\ell = \mathbf{d}(\mathbf{h}^{\ell-1}), \quad (24)$$

which recovers the traditional compositional deep GP.  $\square$

Note that if  $\mu_{d_q}(\cdot)$  is also a zero function, then, if the prior variance of  $\mathbf{d}(\cdot)$  tends to zero, then  $\mathbf{d}(\cdot) \rightarrow \mathbf{0}$ , meaning that  $\mathbf{h} = \mathbf{W}(\mathbf{h}^{\ell-1}) \mathbf{x}$ , which recovers the original TDGP model.

## C Variational inference

First, the prior  $L$ -layer TDGP model as defined in Section 3 is

$$p(f^L(\cdot) \mid \mathbf{h}^{L-1}(\cdot)) = \mathcal{GP}(f \mid 0, \pi_k(\|\mathbf{h}^{L-1}(\mathbf{a}) - \mathbf{h}^{L-1}(\mathbf{b})\|)), \quad (25)$$

where

$$\mathbf{h}^\ell(\mathbf{x}) = \mathbf{W}^\ell(\mathbf{h}^{\ell-1}(\mathbf{x})) \mathbf{x}, \quad \mathbf{h}^0(\mathbf{x}) = \mathbf{x}, \quad (26)$$

$$p(\mathbf{W}^\ell(\cdot) \mid \mathbf{h}^{\ell-1}(\cdot)) = \prod_{d=1}^D \prod_{q=1}^{Q_\ell} \mathcal{GP}(w_{qd}^\ell \mid \mu_{w_{qd}^\ell}, \pi_{k_{w_{qd}^\ell}}^\ell(\|\mathbf{h}^{\ell-1}(\mathbf{a}) - \mathbf{h}^{\ell-1}(\mathbf{b})\|)). \quad (27)$$

We now introduce inducing points for each GP layer in this process. The last-layer process has inducing points  $\mathbf{u}^L$  and each hidden layer has inducing points  $\mathbf{V}^\ell$  defined as follows:

$$\mathbf{u}^L = f^L(\mathbf{Z}^L), \quad \mathbf{V}^\ell = \mathbf{W}^\ell(\mathbf{Z}^\ell), \quad (28)$$

where we also introduce  $L$  sets of pseudo-inputs  $\{\mathbf{Z}^\ell \in \mathbb{R}^{m_\ell \times Q_\ell} \mid \ell \in [1, L]\}$ . Finally, we define the variational distribution as follows:

$$q(f^L, \mathbf{u}, \{\mathbf{W}^\ell, \mathbf{V}^\ell\}) = p(f(\cdot) \mid \mathbf{u}^L) q(\mathbf{u}^L) \prod_{\ell=1}^{L-1} \prod_{d=1}^D \prod_{q=1}^{Q_\ell} p(\mathbf{w}_{qd}^\ell(\cdot) \mid \mathbf{v}_{qd}^\ell) q(\mathbf{v}_{qd}^\ell), \quad (29)$$

where

$$q(\mathbf{u}^L) = \mathcal{N}(\mathbf{u}^L \mid \check{\boldsymbol{\mu}}_u^L, \check{\boldsymbol{\Sigma}}_u^L), \quad q(\mathbf{v}_{qd}^\ell) = \mathcal{N}(\mathbf{v}_{qd}^\ell \mid \check{\boldsymbol{\mu}}_{v_{qd}}^\ell, \check{\boldsymbol{\Sigma}}_{v_{qd}}^\ell). \quad (30)$$

The main simplification of this ELBO is to make each layer conditionally independent of each other when conditioned on the set of inducing variables.

### C.1 Simplification for efficiency

In order to simplify this model, first, we will make each row of  $\mathbf{W}^\ell(\cdot)$  share the same kernel and kernel hyperparameters, this means that the variational posterior covariance of  $\mathbf{V}^\ell$  are also shared, i.e.,  $\check{\boldsymbol{\Sigma}}_{v_{qd}}^\ell = \check{\boldsymbol{\Sigma}}_{v_{qd'}}^\ell$  for every  $d$  and  $d'$ , so we will represent this as  $\check{\boldsymbol{\Sigma}}_{v_q}^\ell$ . Secondly, to compute some expectations in closed form, following Titsias and Lázaro-Gredilla [6], we will consider all kernels to be squared exponential kernels. So that  $\pi_k(r) = \sigma_f^2 \exp[-0.5r^2]$  and  $\pi_{k_{w_{qd}}}^\ell(r) = \sigma_q^2 \exp[-0.5r^2]$ .

And then, the prior model and variational distributions become:

$$\begin{aligned} p(f^L(\cdot) \mid \mathbf{h}^{L-1}(\cdot)) &= \mathcal{GP}(f \mid 0, \sigma_f^2 \exp[-0.5\|\mathbf{h}^{L-1}(\mathbf{a}) - \mathbf{h}^{L-1}(\mathbf{b})\|^2]), \\ p(w_{qd}^\ell(\cdot) \mid \mathbf{h}^{\ell-1}(\cdot)) &= \mathcal{GP}(w_{qd}^\ell \mid \mu_{w_{qd}}^\ell, \sigma_{w_q}^2 \exp[-0.5\|\mathbf{h}^{\ell-1}(\mathbf{a}) - \mathbf{h}^{\ell-1}(\mathbf{b})\|^2]), \\ q(\mathbf{u}^L) &= \mathcal{N}(\mathbf{u}^L \mid \check{\boldsymbol{\mu}}_u^L, \check{\boldsymbol{\Sigma}}_u^L), \\ q(\mathbf{v}_{qd}^\ell) &= \mathcal{N}(\mathbf{v}_{qd}^\ell \mid \check{\boldsymbol{\mu}}_{v_{qd}}^\ell, \check{\boldsymbol{\Sigma}}_{v_q}^\ell). \end{aligned}$$

### C.2 Two-layer model

Finally, we will work on the model with a single hidden layer  $\mathbf{W}(\cdot)$  and one output layer  $f(\cdot)$ . Again, the prior model is simplified to:

$$p(f(\cdot) \mid \mathbf{h}(\cdot)) = \mathcal{GP}(f \mid 0, \sigma_f^2 \exp[-0.5\|\mathbf{h}(\mathbf{a}) - \mathbf{h}(\mathbf{b})\|^2]), \quad (31)$$

where

$$p(\mathbf{W}(\cdot) \mid \mathbf{h}(\cdot)) = \prod_{d=1}^D \prod_{q=1}^Q \mathcal{GP}\left(w_{qd} \mid \mu_{w_{qd}}^\ell, \sigma_{w_q}^2 \exp\left[-0.5\left\|\frac{\mathbf{a}}{l} - \frac{\mathbf{b}}{l}\right\|^2\right]\right), \quad (32)$$

$$\mathbf{h}(\mathbf{x}) = \mathbf{W}(\mathbf{x}) \mathbf{x}. \quad (33)$$

Accordingly, the variational distribution becomes:

$$q(\mathbf{u}) = \mathcal{N}(\mathbf{u} \mid \check{\boldsymbol{\mu}}_u, \check{\boldsymbol{\Sigma}}_u), \quad q(\mathbf{v}_{qd}) = \mathcal{N}(\mathbf{v}_{qd} \mid \check{\boldsymbol{\mu}}_{v_{qd}}, \check{\boldsymbol{\Sigma}}_{v_q}). \quad (34)$$

### C.3 Evidence lower bound (ELBO)

As in Titsias and Lázaro-Gredilla [6] and Titsias [4], we will consider the marginals of  $f(\cdot)$  and  $\mathbf{W}(\cdot)$  evaluated at the training data  $(\mathbf{X}, \mathbf{y})$ ,  $\mathbf{f} \in \mathbb{R}^n = \mathbf{f}(\mathbf{X})$  and  $\mathbf{W} \in \mathbb{R}^{n \times Q \times D} = \mathbf{W}(\mathbf{X})$ . Following the definition of the ELBO, we have the following lower bound on the evidence:

$$p(\mathbf{y}) = \langle p(\mathbf{y} \mid \mathbf{f}) \rangle_{p(\mathbf{f}, \mathbf{W})} \quad (35)$$

$$\geq \boxed{\langle \log p(\mathbf{y} \mid \mathbf{f}) \rangle_{q(\mathbf{f}, \mathbf{u}, \mathbf{W}, \mathbf{v})} - \text{KL}(q(\mathbf{u}) \parallel p(\mathbf{u}))} - \sum_{d=1}^D \sum_{q=1}^Q \text{KL}(q(\mathbf{v}_{dq}) \parallel p(\mathbf{v}_{dq})). \quad (36)$$

Given our simplifying assumptions from before and choice of variational distribution, the terms inside the **blue box** have the same form as the ELBO of Titsias and Lázaro-Gredilla [6], therefore the value of the **blue box** with optimal  $q(\mathbf{u})$  is:

$$\begin{aligned} \blacksquare &= -\frac{1}{2\sigma^2} (\mathbf{y}^\top \mathbf{y} + \psi_0 - \text{Tr}[\mathbf{K}_u^{-1} \boldsymbol{\Psi}_2]) - \frac{n}{2} \ln(2\pi\sigma^2) \\ &\quad + \frac{1}{2\sigma^2} \mathbf{y}^\top \boldsymbol{\Psi}_1 \boldsymbol{\Sigma}^{-1} \boldsymbol{\Psi}_1^\top \mathbf{y} + \frac{m_u}{2} \ln(\sigma^2) - \frac{1}{2} \ln|\boldsymbol{\Sigma}| + \frac{1}{2} \ln|\mathbf{K}_u|, \end{aligned} \quad (37)$$

where the  $\Psi$ -statistics [5, 6] are defined as:

$$\psi_0 = \langle \text{Tr } \mathbf{K}_f \rangle_{q(\mathbf{W}, \mathbf{V})} = n\sigma_f^2, \quad \Psi_1 = \langle \mathbf{K}_{fu} \rangle_{q(\mathbf{W}, \mathbf{V})}, \quad \Psi_2 = \langle \mathbf{K}_{fu}^\top \mathbf{K}_{fu} \rangle_{q(\mathbf{W}, \mathbf{V})}, \quad (38)$$

and the optimal  $q(\mathbf{u})$  is:

$$q(\mathbf{u}) = N\left(\mathbf{u} \mid \mathbf{K}_u [\sigma^2 \mathbf{K}_u + \Psi_2]^{-1} \Psi_1^\top \mathbf{y}, \sigma^2 \mathbf{K}_u (\sigma^2 \mathbf{K}_u + \Psi_2)^{-1} \mathbf{K}_u\right) \quad (39)$$

#### C.4 Computing the $\Psi$ -statistics

The trick for computing the  $\Psi$  statistics is to show that each entry of the matrices only depends on a specific  $\mathbf{W}(\mathbf{x}_i)$  and can be expressed as a product in the rows  $q$ , therefore allowing us to marginalize  $q(\mathbf{W}, \mathbf{V})$  to  $\prod_{d=1}^D q(\mathbf{w}_{iqd}, \mathbf{v}_{iqd})$ . So, starting with  $\Psi_1$ :

$$[\Psi_1]_{ij} = \langle [\mathbf{K}_{fu}]_{ij} \rangle_{q(\mathbf{W}, \mathbf{V})} \quad (40)$$

$$= \left\langle \sigma_f \exp \left[ -\frac{1}{2} (\mathbf{W}_i \mathbf{x}_i - \mathbf{z}_j^{(1)}) (\mathbf{W}_i \mathbf{x}_i - \mathbf{z}_j^{(1)})^\top \right] \right\rangle_{q(\mathbf{W})}, \quad (41)$$

so we can marginalize  $q(\mathbf{W})$ ,

$$= \left\langle \sigma_f \exp \left[ -\frac{1}{2} (\mathbf{W}_i \mathbf{x}_i - \mathbf{z}_j^{(1)}) (\mathbf{W}_i \mathbf{x}_i - \mathbf{z}_j^{(1)})^\top \right] \right\rangle_{q(\mathbf{W}_i)} \quad (42)$$

$$= \left\langle \sigma_f \exp \left[ -\frac{1}{2} \sum_{q=1}^Q (\mathbf{w}_{iq}^\top \mathbf{x}_i - z_{jq}^{(1)})^2 \right] \right\rangle_{q(\mathbf{W}_i)} \quad (43)$$

$$= \sigma_f \prod_{q=1}^Q \left\langle \exp \left[ -\frac{1}{2} (\mathbf{w}_{iq}^\top \mathbf{x}_i - z_{jq}^{(1)})^2 \right] \right\rangle_{q(\mathbf{w}_{iq})}. \quad (44)$$

This is the same situation as in Appendix B.1 of [6], so:

$$[\Psi_1]_{ij} = \sigma_f \prod_{q=1}^Q \left( \mathbf{x}_i^\top \tilde{\Sigma}_q \mathbf{x}_i + 1 \right)^{-\frac{1}{2}} \exp \left[ -\frac{1}{2} \frac{(\tilde{\boldsymbol{\mu}}_{iq}^\top \mathbf{x}_i - z_{jq}^{(1)})^2}{\left( \mathbf{x}_i^\top \tilde{\Sigma}_q \mathbf{x}_i + 1 \right)} \right], \quad (45)$$

where  $\tilde{\boldsymbol{\mu}}_{iq}$  and  $\tilde{\Sigma}_q$  are the mean and covariance of  $\prod_{d=1}^D q(\mathbf{w}_{iqd})$ . Now, for  $\Psi_2$ :

$$[\Psi_2]_{jk} = \left\langle [\mathbf{K}_{fu}^\top \mathbf{K}_{fu}]_{jk} \right\rangle_{q(\mathbf{W}, \mathbf{V})} \quad (46)$$

$$= \left\langle \sum_{i=1}^n [\mathbf{K}_{fu}]_{ij} [\mathbf{K}_{fu}]_{ik} \right\rangle_{q(\mathbf{W}, \mathbf{V})} \quad (47)$$

$$= \sum_{i=1}^n \left\langle [\mathbf{K}_{fu}]_{ij} [\mathbf{K}_{fu}]_{ik} \right\rangle_{q(\mathbf{W}_i)}. \quad (48)$$

Again, following [6]:

$$\begin{aligned} [\Psi_2]_{jk} &= \sigma_f^2 \exp \left[ -\frac{1}{4} \sum_{q=1}^Q (z_{jq}^{(1)} - z_{kq}^{(1)})^2 \right] \\ &\quad \times \sum_{i=1}^n \prod_{q=1}^Q \left( 2\mathbf{x}_i^\top \tilde{\Sigma}_q \mathbf{x}_i + 1 \right)^{-\frac{1}{2}} \exp \left[ -\frac{(\tilde{\boldsymbol{\mu}}_{iq}^\top \mathbf{x}_i - \tilde{z}_q)^2}{2\mathbf{x}_i^\top \tilde{\Sigma}_q \mathbf{x}_i + 1} \right], \end{aligned} \quad (49)$$

where  $\tilde{z}_q = \frac{z_{jq}^{(1)} + z_{kq}^{(1)}}{2}$ .

## D Details on experiments

**Setting.** For each experiment, we performed cross-validation (five folds for bathymetry and 10 folds for UCI). We compare the generalization error using MRAE and compare the uncertainty quantification using NLPD. We used the Adam optimizer, and the following schedule (we varied Adam step size and likelihood variance):

- For 500 epochs: step size 0.1, and likelihood variance fixed to 0.01,
- For 1500 epochs: step size 0.01, and likelihood variance fixed to 0.01,
- For 5000 epochs: step size 0.001, and likelihood variance is trainable.

**Architecture.** The architectural details of each model are:

**SGPR** 50 inducing points for the output process and an ARD-squared exponential kernel. Inference is done by using the optimal  $q(\mathbf{u})$  as described by Titsias [4].

**DKL** 50 inducing points for the output process and an ARD-squared exponential kernel. For the deep kernel, we use an MLP with architecture  $[D, 500, 50, D]$ , where  $D$  is the dimension of the inputs and a final BatchNorm layer. All hidden-layer activations are ReLU. Inference is done by using the optimal  $q(\mathbf{u})$  as described by Titsias [4].

**CDGP** 50 inducing points for the output process and 25 for the latent space process. All layers use an ARD-squared exponential kernel. The dimension of the hidden layer is set to  $D$ . We use doubly stochastic inference [3] with whitened variables, i.e. we reparametrize  $\mathbf{u}$  as  $\mathbf{K}_u^{-\frac{1}{2}}\mathbf{u}$ .

**DNSGP** 50 inducing points for the output process and 25 for the lengthscale matrix space process. All layers use an ARD-squared exponential kernel and the lengthscale matrix process is set to be diagonal with warping function  $\exp(\mathbf{h} + s)$ , where  $s$  is a learnable scalar. We use doubly stochastic inference [3] with whitened variables, i.e. we reparametrize  $\mathbf{u}$  as  $\mathbf{K}_u^{-\frac{1}{2}}\mathbf{u}$ .

**TDGP** 50 inducing points for the output process and 25 for the inverse lengthscale matrix space process. The size of the inverse lengthscale matrix  $\mathbf{W}$  is set to  $Q \times D$ , where  $Q = D$  and each row  $q$  of shares the same kernel. Our variational posterior distribution in  $q(\mathbf{V})$  is set to mean-field where  $q(\mathbf{V}) = \prod_{i=1}^n \prod_{q=1}^Q \prod_{d=1}^D q(v_{iqd})$ .

### D.1 Synthetic experiment

**Data.** We generated a synthetic dataset by definition a composite function  $f = g \circ h$ , with  $g : \mathbb{R} \rightarrow \mathbb{R}$  and  $h : \mathbb{R}^2 \rightarrow \mathbb{R}$  are non-linear functions. In this context,  $h$  acts as a “funnel” inducing a 1D manifold. These functions are defined as:

$$h(\mathbf{x}) = 2x_0 \sin(x_0\pi) + 2 \cos(x_0\pi) \tag{50}$$

$$g(z) = \frac{\sin(z)}{z} - z^2. \tag{51}$$

Then, we uniformly sample  $\mathbf{x}$  in the interval  $[-1, 1] \times [-1, 1]$  and split 50/50 for train and validation.

### D.2 Bathymetry case study

**Data.** The selected subset of GEBCO data covers the Andes mountain range, ocean, and land as an example of a non-stationary task (longitude in the range from  $-80$  to  $-60$ , and latitude in the range from  $-20$  to  $-10$ ). We randomly subsampled 1000 data points.

**Computational resources.** Running all the models took 2.5 hours using an NVIDIA A100 GPU. More fine-grained time measurements are presented in Table 3.

Table 3: Training and evaluation time (seconds) in GEBCO dataset (avg±std). Lower is better.

	Train	Evaluation
SGP	23.4 ± 1.5	0.03 ± 0.01
DKL	384.6 ± 8.5	0.76 ± 1.28
CDGP	570.8 ± 0.6	0.14 ± 0.17
DNSGP	596.4 ± 29.8	0.06 ± 0.01
TDGP	126.7 ± 0.8	0.12

Table 4: Training time (seconds) for the benchmark datasets (avg±std). Lower is better.

	housing	concrete	energy	wine_red
SGP	41.4 ± 0.2	41.0 ± 0.7	41.5 ± 1.2	42.6 ± 2.0
DKL	379.1 ± 3.7	379.0 ± 3.5	376.7 ± 2.3	386.1 ± 13.4
CDGP	601.5 ± 9.4	592.1 ± 4.1	592.9 ± 3.5	615.4 ± 23.5
DNSGP	668.0 ± 75.9	621.8 ± 2.2	620.8 ± 1.9	650.8 ± 24.0
TDGP	572.4 ± 12.4	563.7 ± 1.8	484.3 ± 1.4	928.6 ± 42.1

### D.3 Benchmark datasets

**Data.** We used a subset of datasets from the UCI repository: housing<sup>3</sup>, concrete<sup>4</sup>, wine-red<sup>5</sup>, and energy datasets.<sup>6</sup>

Housing dataset has 506 samples with 13 features; concrete dataset has 1030 samples with 8 features; wine-red has 1599 samples with 11 features; energy dataset has 768 samples with 8 features.

**Computational resources.** Training all models in all datasets for ten folds took 25.55 hours using an NVIDIA TITAN RTX GPU. Per dataset time measurements are presented in Table 4

## E Computational and test performance as a function of width

In the experiments of Section 4, the width of the hidden layer  $Q$  for TDGP was always set to match the dimension of the input  $D$ . As seen in Fig. 10, after optimization of the hyperparameters, the effective width of the layer for all datasets was always much smaller than  $D$ . Therefore, it is reasonable to expect that a wider model wouldn't increase the model's performance.

Nevertheless, we conduct an additional experiment to explore the performance penalty of increasing  $Q$  up to  $D$  in terms of computational resources and test accuracy. In the chosen housing dataset, Fig. 10 shows that  $Q \approx 2 < D$  is the effective width of an optimized network. Therefore, we re-run this experiment with values of  $Q$  ranging from 1 to  $D$  as shown in Fig. 11.

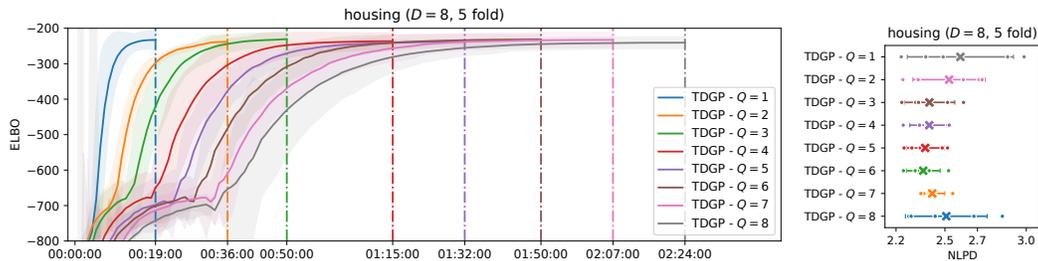


Figure 11: Training curves (left) and test metrics (right) for the housing dataset with a variable width  $Q$

<sup>3</sup><https://archive.ics.uci.edu/ml/machine-learning-databases/housing/>

<sup>4</sup><https://archive.ics.uci.edu/ml/datasets/Concrete+Compressive+Strength>

<sup>5</sup><https://archive.ics.uci.edu/ml/datasets/Wine+Quality, red variant>

<sup>6</sup><https://archive.ics.uci.edu/ml/datasets/energy+efficiency>

As discussed, we observe a linear increase in training time as the model’s width increases. In terms of predictive performance, the best widths are 3 to 6; in theory, we wouldn’t expect a performance drop above a certain minimum width, as the effective width is a trained variable, however, as stated in our limitation, we expected the increased number of variables to optimize to add more complexity to the optimization landscape and, therefore, increase the difficulty in finding the best set of hyperparameters.

## F Expressivity of the prior with increasing depth

The TDGP model as defined in Section 3 places a zero-mean prior on all the layers. This is in contrast with the standard CDGP model, which as shown in Duvenaud et al. [1], suffers prior collapse under this assumption. Figure 2 shows this effect by plotting different samples from CDGP and TDGP priors with zero mean as the model depth increases. Nonetheless, as shown in Salimbeni and Deisenroth [2], another way to visualize this pathology is to plot samples of the covariance matrix as the number of layers increase.

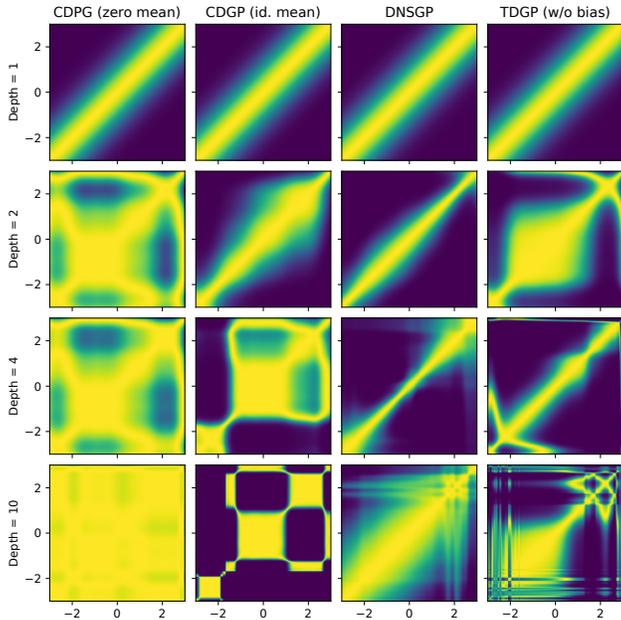


Figure 12: Samples from the prior covariance matrix for different layers and models

As shown in Fig. 12, as the model depth increases, the covariance matrix of CDGP with zero mean eventually saturates, i.e. all points are high correlated, which leads to the flat priors shown in Fig. 2. We can also see that, as discussed in Salimbeni and Deisenroth [3], changing the zero-mean prior to one with an linear mean function fixes this pathology, as well as using a zero-mean DNSGP or a zero-mean TDGP model. This is further evidence that our model

## G Societal and broader impact

Gaussian processes are popular methods for spatiotemporal modeling in, e.g., climatology, geoscience, public health, and ecology. This work proposes TDGP, a novel formulation for deep GPs that preserves the performance of compositional deep GPs while significantly improving their interpretability. We believe the inherent interpretability of TDGP priors will make it easier for applied researchers to encode their subjective knowledge, consequently improving the data efficiency of their models and reducing predictive uncertainties. Additionally, we do not foresee any negative societal impact stemming directly from this work.

## References

- [1] David Duvenaud, Oren Rippel, Ryan Adams, and Zoubin Ghahramani. “Avoiding pathologies in very deep networks”. In: *Artificial Intelligence and Statistics (AISTATS)*. 2014.
- [2] Hugh Salimbeni and Marc Peter Deisenroth. “Deeply non-stationary Gaussian processes”. In: *2nd Workshop on Bayesian Deep Learning (NeurIPS)*. 2017.
- [3] Hugh Salimbeni and Marc Peter Deisenroth. “Doubly Stochastic Variational Inference for Deep Gaussian Processes”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2017.
- [4] Michalis K. Titsias. “Variational Learning of Inducing Variables in Sparse Gaussian Processes”. In: *Artificial Intelligence and Statistics (AISTATS)*. 2009.
- [5] Michalis K. Titsias and Neil D. Lawrence. “Bayesian Gaussian Process Latent Variable Model”. In: *Artificial Intelligence and Statistics (AISTATS)*. 2010.
- [6] Michalis K. Titsias and Miguel Lázaro-Gredilla. “Variational Inference for Mahalanobis Distance Metrics in Gaussian Process Regression”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2013.