Supplementary Information for Gaussian Process Prior Variational Autoencoders

Contents

1 Supplementary Methods

1.1 ELBO derivation

We assume the following posterior distribution on latent variables

$$
q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y}) = \prod_{n} \mathcal{N}\left(\boldsymbol{z}_n \mid \boldsymbol{\mu}_{\boldsymbol{\psi}}^{\boldsymbol{z}}(\boldsymbol{y}_n), \text{diag}(\boldsymbol{\sigma}^{\boldsymbol{z}}_{\boldsymbol{\psi}}^2(\boldsymbol{y}_n))\right),\tag{1}
$$

where ψ are variational parameters. The evidence lower bound (ELBO) can be derived as follows:

$$
\begin{array}{rcl}\n\log p(\boldsymbol{Y} \mid \boldsymbol{X}, \boldsymbol{W}, \boldsymbol{\phi}, \sigma_{y}^{2}, \boldsymbol{\theta}) & = & \log \int \frac{p(\boldsymbol{Y} \mid \boldsymbol{z}, \boldsymbol{\phi}, \sigma_{y}^{2}) p(\boldsymbol{Z} \mid \boldsymbol{X}, \boldsymbol{W}, \boldsymbol{\theta}, \alpha)}{q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y})} q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y}) d\boldsymbol{Z} \\
& \geq & \int \log \left(\frac{p(\boldsymbol{Y} \mid \boldsymbol{z}, \boldsymbol{\phi}, \sigma_{y}^{2}) p(\boldsymbol{Z} \mid \boldsymbol{X}, \boldsymbol{W}, \boldsymbol{\theta}, \alpha)}{q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y})} \right) q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y}) d\boldsymbol{Z} \\
& = & \mathbb{E}_{\boldsymbol{Z} \sim q_{\boldsymbol{\psi}}} \left[\sum_{n} \log \mathcal{N}(\boldsymbol{y}_{n} \mid g_{\boldsymbol{\phi}}(\boldsymbol{z}_{n}), \sigma_{y}^{2} \boldsymbol{I}_{K}) + \log p(\boldsymbol{Z} \mid \boldsymbol{X}, \boldsymbol{W}, \boldsymbol{\theta}, \alpha) \right] - \int \log q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y}) q_{\boldsymbol{\psi}}(\boldsymbol{Z} \mid \boldsymbol{Y}) d\boldsymbol{Z} \\
& = & \mathbb{E}_{\boldsymbol{Z} \sim q_{\boldsymbol{\psi}}} \left[\sum_{n} \log \mathcal{N}(\boldsymbol{y}_{n} \mid g_{\boldsymbol{\phi}}(\boldsymbol{z}_{n}), \sigma_{y}^{2} \boldsymbol{I}_{K}) + \log p(\boldsymbol{Z} \mid \boldsymbol{X}, \boldsymbol{W}, \boldsymbol{\theta}, \alpha) \right] + \frac{1}{2} \sum_{nl} \log(\boldsymbol{\sigma}^{z_{l}}_{\boldsymbol{\psi}}(\boldsymbol{y}_{n})_{l}) + \text{const}\n\end{array}
$$

1.2 Loss derivation

First, we approximate the expectation by sampling. Specifically, we sample a latent representation $\boldsymbol{Z}_{\boldsymbol{\psi}}=\left[\boldsymbol{z}_{\boldsymbol{\psi}_1}, \dots, \boldsymbol{z}_{\boldsymbol{\psi}_N}\right] \in \mathbb{R}^{N \times L}$ from the posterior as:

$$
\mathbf{z}_{\psi_n} = \mu_{\psi}^z(\mathbf{y}_n) + \boldsymbol{\epsilon}_n \odot \sigma^z_{\psi}(\mathbf{y}_n), \ \boldsymbol{\epsilon}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{L \times L}), \ n = 1, \dots, N,
$$
\n(2)

where we used the reparametrization trick to separate the noisy generation of samples from the model parameters. The resulting approximate ELBO is

ELBO
$$
\approx \sum_{n} \log \mathcal{N}(\mathbf{y}_n | g_{\phi}(\mathbf{z}_{\psi_n}), \sigma_y^2 \mathbf{I}_K) + \log p(\mathbf{Z}_{\psi} | \mathbf{X}, \mathbf{W}, \theta, \alpha) + \frac{1}{2} \sum_{nl} \log(\sigma^2 \phi(\mathbf{y}_n)_l) + \text{const}
$$

Finally, as we optimize σ_y^2 on a validation set (and not on the training set in order to avoid overfitting), maximization of the approximate ELBO on the training set is equivalent to minimizing the following cost:

$$
\mathcal{L}(\boldsymbol{\phi}, \boldsymbol{\psi}, \boldsymbol{\theta}, \alpha, \sigma_y^2) = \frac{1}{K} \underbrace{\sum_i (\mathbf{y}_i - g_{\boldsymbol{\phi}}(\boldsymbol{\omega}_{\boldsymbol{\psi}_n}))^2 - \frac{\lambda}{L}}_{\text{reconstruction term}} \left[\underbrace{\log_p (\mathbf{Z}_{\boldsymbol{\psi}} \, | \, \mathbf{X}, \mathbf{W}, \boldsymbol{\theta}, \alpha)}_{\text{latent-space GP model}} + \underbrace{\frac{1}{2} \sum_{nl} \log(\boldsymbol{\sigma}^2 \hat{\boldsymbol{\psi}}(\mathbf{y}_n)_l)}_{\text{regularization}} \right], \quad (3)
$$

where K is the number of pixels in the image and λ is a trade-off parameter balancing data reconstruction and latent space prior regularization.

Selection of λ . We select the value of λ based on a standard VAE and use the same value for all compared models. Specifically, we optimize the VAE loss for different values of λ and select the value for which the VAE ELBO is maximal on a validation set. In order to compute the VAE validation ELBO we need to estimate optimal value of σ_y^2 for every value of λ in the grid. Given a certain value of $\lambda = \hat{\lambda}$, the optimal value of σ_y^2 can be estimated as

$$
\sigma_y^{2(\text{val})} = \frac{1}{N(\text{val})} \sum_{n=1}^{N(\text{val})} \left(\mathbf{y}_n^{(\text{val})} - g_{\phi_{\lambda}}(z_{\psi_{\lambda_n}}^{(\text{val})}) \right)^2, \tag{4}
$$

where $N^{(val)}$ is the number of samples in the validation set and $(\phi_{\hat{\lambda}}, \psi_{\hat{\lambda}})$ are the values of the encoder/decoder parameters after training for $\lambda = \hat{\lambda}$.

1.3 Fast low-rank computations

Parameter inference in Gaussian models scales cubically with the number of observations. To preform fast parameter inference in our Gaussian model, we use the fact that the total covariance is the sum of low-rank matrix and the identity matrix:

$$
\boldsymbol{K} = \boldsymbol{V}\boldsymbol{V}^T + \alpha \boldsymbol{I} \tag{5}
$$

where $V \in \mathbb{R}^{N \times O}$ and $O \ll N$. Using the Woodbury identity [\[1\]](#page-6-0) and the determinant lemma [\[2\]](#page-6-1), the linear system $K^{-1}M$ with $M \in \mathbb{R}^{N \times K}$ and the log determinant of K can be computed as:

$$
\mathbf{K}^{-1}\mathbf{M} = \frac{1}{\alpha}\mathbf{I} - \frac{1}{\alpha}\mathbf{V}(\alpha\mathbf{I} + \mathbf{V}^T\mathbf{V})^{-1}\mathbf{V}^T\mathbf{M}, \tag{6}
$$

$$
logdet \mathbf{K} = NLlogdet \alpha + logdet(\mathbf{I} + \frac{1}{\alpha} \mathbf{V}^T \mathbf{V}), \tag{7}
$$

which have $O(NO^2 + O^3 + ONK)$ and $O(NO^2 + O^3)$ complexities, respectively.

1.4 Implementation of low-memory stochastic backprobagation

For a single latent dimension the GP prior introduces the following term in the ELBO:

$$
\log p(\mathbf{z}_{\psi} | \mathbf{X}, \mathbf{W}, \boldsymbol{\theta}, \alpha) = -\frac{1}{2} \mathbf{z}_{\psi}^T \mathbf{K}_{\theta}(\mathbf{X}, \mathbf{V})^{-1} \mathbf{z}_{\psi} - \frac{1}{2} \log \det \mathbf{K}_{\theta}(\mathbf{X}, \mathbf{V})
$$
(8)

In the equation above, z_{ψ} functionally depends on the image space representations and thus a naive computation of full gradient descent would require loading the entire high-dimensional dataset in memory, which would be unfeasible form many high-dimensional image datasets. To overcome this limitation, we recast computations in a form where full-matrix operations only take place in the low dimensional space, while the dependency on the nested derivatives, which involve high memory loads, is linearized and mini-batch operations are therefore allowed. This can be achieved by considering a first-order Taylor expansion of the GP prior term. Specfically, using that $K_{\theta} = V_{\theta} V_{\theta}^T + \alpha I$, and collecting all parameters in $\xi = {\psi, \theta, \alpha}$, we can rewrite the term in Eq. [\(8\)](#page-1-2) in this functional form $f(z(\xi), V(\xi), \alpha(\xi))$. The first-order Taylor expansion of $f(z(\xi), V(\xi), \alpha(\xi))$ around $(z_{\xi_0}, V_{\xi_0}, \alpha_{\xi_0})$ is:

$$
f(\mathbf{z}(\xi), \mathbf{V}(\xi), \alpha(\xi)) \approx \mathbf{a}^T \mathbf{z}(\xi) + \text{tr}\left(\mathbf{B}^T \mathbf{V}(\xi)\right) + c\alpha(\xi) + \text{const.}
$$
\n(9)

where

$$
a = \left(\frac{\partial f}{\partial z}\right)_{\xi_0} = \left(K^{-1}z\right)_{\xi_0} \tag{10}
$$

$$
\boldsymbol{B} = \left(\frac{\partial f}{\partial \boldsymbol{V}}\right)_0 = \left(-\boldsymbol{K}^{-1} \boldsymbol{z} \boldsymbol{z}^T \boldsymbol{K}^{-1} \boldsymbol{V} + \boldsymbol{K}^{-1} \boldsymbol{V}\right)_{\boldsymbol{\xi}_0}
$$
(11)

$$
c = \left(\frac{\partial f}{\partial \alpha}\right)_{\xi_0} = \frac{1}{2} \left(-z^T \mathbf{K}^{-1} \mathbf{K}^{-1} z + \text{tr}(\mathbf{K}^{-1})\right)_{\xi_0}
$$
(12)

Note that this approximation, applied at every step of gradient descent, locally preserves the gradients. Thus we can engineer the following low-memory four-step procedure for full gradient descent:

- produce and store latent noise realizations (ϵ) used for the reparametrization trick;
- obtain latent variable representations z , combining the outputs of the encoder and the noise ϵ . We employ mini-batch forward propagation for this step;
- evaluate a, B and c across all samples. Note that this step has low-memory requirements as it only involes low-dimensional representations;
- exploit the local Taylor expansion as a proxy for our optimization. As this function is linear in the data, we can accumulate its gradient in a mini-batch fashion, by passing mini-batches of ϵ , Y , a and B :
- update the parameters ξ using the full gradients as in standard gradient descent.

Finally, we note that in our specific setting a, B and c can be computed linearly on the number of sample because of the low rank structure of K (see previous section).

1.5 Out of sample predictions.

We here derive an approximate predictive posterior for GPPVAE. Specifically, given training images Y , object representations X and view representations V the predictive posterior for image y_* representing object x_* in view v_* is

$$
p(\mathbf{y}_{\star} | \mathbf{x}_{\star}, \mathbf{v}_{\star}, \mathbf{Y}, \mathbf{X}, \mathbf{V}) = \frac{p(\mathbf{y}_{\star}, \mathbf{Y} | \mathbf{x}_{\star}, \mathbf{v}_{\star}, \mathbf{X}, \mathbf{V})}{p(\mathbf{Y} | \mathbf{X}, \mathbf{V})}
$$
(13)

$$
= \frac{1}{p(Y|X,V)} \int p(y_*|z_*)p(Y|Z) \underbrace{p(z_*,Z|x_*,v_*,X,V)}_{\text{joint distribution}} dz_* dZ \tag{14}
$$

$$
= \frac{1}{p(Y|X,V)} \int p(y_{\star}|z_{\star}) p(Y|Z) \underbrace{p(z_{\star}|x_{\star},v_{\star},Z,X,V)}_{\text{predictive posterior}} \underbrace{p(Z|X,V)}_{\text{GP prior on training}} dz_{\star} dZ \quad (15)
$$

$$
= \int p(\mathbf{y}_{\star} | \mathbf{z}_{\star}) \underbrace{p(\mathbf{z}_{\star} | \mathbf{x}_{\star}, \mathbf{v}_{\star}, \mathbf{Z}, \mathbf{X}, \mathbf{V})}_{\text{predictive posterior}} \underbrace{\frac{p(\mathbf{Y} | \mathbf{Z})p(\mathbf{Z} | \mathbf{X}, \mathbf{V})}{p(\mathbf{Y} | \mathbf{X}, \mathbf{V})}}_{\text{posterior on } \mathbf{Z}} dz_{\star} d\mathbf{Z}
$$
(16)

$$
\approx \int p(\mathbf{y}_\star \,|\, \mathbf{z}_\star) \underbrace{p(\mathbf{z}_\star \,|\, \mathbf{x}_\star, \mathbf{v}_\star, \mathbf{Z}, \mathbf{X}, \mathbf{V})}_{\text{predictive posterior}} \underbrace{q(\mathbf{Z} \,|\, \mathbf{Y})}_{\text{approx. post. on } \mathbf{Z}} dz_\star d\mathbf{Z}
$$
(17)

where we dropped dependencies from model parameters to simplify the notation.

2 Supplementary Figures

Figure S1 Neural network architecture used in the MNIST experiments. (a) Encoder architecture; (b) Decoder architecture. The same encoder/decoder architectures are used for both VAE and GPPVAE. For CVAE, we still use the same architecture but we provide view representations (rotation angles) as inputs to both the encoder and the decoder. For the encoder, we provide rotation angles to the first layer (by adding additional channels to the input image) and again before the dense layer (by stacking them to the vector produced by the last convolution). Similarly, for the decoder, we provide rotation angles as inputs to the first layer (by stacking them with the 16-dimensional latent variable) and again before the first convolution layer (by adding them as additional channels to the 8x8x4 tensor fed to the first convolution layer). For a fair comparison with GPPVAE, we account for periodicity in the view representation by providing both the angle sine and cosine, i.e. $\mathbf{w} = [\sin \phi, \cos \phi]$, where ϕ is the rotation angle.

Figure S2 Neural network architecture used in the face dataset experiments. (a) Encoder architecture; (b) Decoder architecture. The same encoder/decoder architectures are used for both VAE and GPPVAE. For CVAE, we still use the same architecture but we provide view information as one hot encoding of the 9 different poses to both the encoder and the decoder networks. For the encoder, we provide view representations to the first layer (by adding additional channels to the input image) and again before the dense layer (by stacking them to the vector produced by the last convolution). Similarly, for the decoder, we provide view representations as inputs to the first layer (by stacking them with the 16-dimensional latent variable) and again before the first convolution layer (by adding them as additional channels to the 8x8x4 tensor fed to the first convolution layer).

The trade-on parameter was selected as to maximize vAE EEBO. We here show the vAE van-
lower bound as a function of the trade-off parameter for the MNIST (a) and the face dataset (b). Figure S3 Selection of the trade-off parameter between fit-to-data and latent-space model. The trade-off parameter was selected as to maximize VAE ELBO. We here show the VAE validation

Figure S4 For CVAE and LIVAE, we also considered the alternative strategy of selecting the value of the trade-off parameter λ that maximizes prediction performance on the validation set. We refer to these additional methods as CVAE-opt and LIVAE-opt. We here show the selection of λ for CVAE-opt (a) and LIVAE-opt (b) and their performance (c-d) for the experiments in MNIST.

Figure S6 Out-of-sample predictions obtained from GPPVAE-joint when considering the perceptual loss introduced in [\[3\]](#page-6-2).

References

- [1] Henderson, H. V. & Searle, S. R. On deriving the inverse of a sum of matrices. Siam Review 23, 53–60 (1981).
- [2] Harville, D. A. Matrix algebra from a statistician's perspective, vol. 1 (Springer, 1997).
- [3] Hou, X., Shen, L., Sun, K. & Qiu, G. Deep feature consistent variational autoencoder. In Applications of Computer Vision (WACV), 2017 IEEE Winter Conference on, 1133–1141 (IEEE, 2017).