

Variational Linearized Laplace Approximation for Bayesian Deep Learning

Luis A. Ortega

Universidad Autónoma de Madrid

LUIS.ORTEGA@UAM.ES

Simón Rodríguez Santana

Universidad Pontificia Comillas

SRSANTANA@ICAI.COMILLAS.EDU

Daniel Hernández-Lobato

Universidad Autónoma de Madrid

DANIEL.HERNANDEZ@UAM.ES

Abstract

The Linearized Laplace Approximation (LLA) has been recently used to perform uncertainty estimation on the predictions of pre-trained deep neural networks (DNNs). However, its widespread application is hindered by significant computational costs, particularly in scenarios with a large number of training points or DNN parameters. Consequently, additional approximations of LLA, such as Kronecker-factored or diagonal approximate GGN matrices, are employed, potentially compromising the model’s performance. To address these challenges, we propose a new method for approximating LLA using a variational sparse Gaussian Process (GP). Our method is based on the dual RKHS formulation of GPs and retains as the predictive mean the output of the original DNN. Furthermore, it allows for efficient stochastic optimization, which results in sub-linear training time in the size of the training dataset. Specifically, its training cost is independent of the number of training points. We compare our proposed method against accelerated LLA (ELLA), which relies on the Nyström approximation, as well as other LLA variants employing the sample-then-optimize principle. Experimental results show that our method outperforms these already existing efficient variants of LLA, both in terms of the quality of the predictive distribution and in terms of total computational time.

1. Introduction

Deep neural networks (DNNs) have gained widespread popularity for addressing pattern recognition problems due to their state-of-the-art performance in predicting target values from a set of input attributes (He et al., 2016; Vaswani et al., 2017). Despite this great success, DNNs exhibit limitations when computing a predictive distribution that accounts for the confidence in the predictions. Specifically, DNNs result in weak calibration (Guo et al., 2017) and in poor reasoning regarding model uncertainty (Blundell et al., 2015). These issues become particularly critical in risk-sensitive situations like autonomous driving (Kendall and Gal, 2017) and healthcare systems (Leibig et al., 2017) among others.

The Laplace Approximation (LA) leverages the maximum a posteriori (MAP) solution, attainable via back-propagation, providing a Gaussian posterior approximation where the pre-trained optimal solution furnishes the mean. This resembles a pre-training step followed by fine-tuning, a common practice in deep learning (Daxberger et al., 2021a). LA’s weakness is the necessity of computing the Hessian at the MAP, which becomes prohibitive for large




Figure 1: Predictive distribution (mean in blue and shaded two times the standard deviation) on a toy 1D regression dataset with a 2 hidden layer MLP with 50 units trained using back-propagation. The predictive distribution of VaLLA is on par with or better than other approximations (last layer, diagonal and Kronecker factorization) and other methods (ELLA). The optimal values for the noise and the Gaussian prior are optimized by maximizing the marginal log likelihood estimate of LLA. VaLLA and ELLA use the optimal values found by LLA. VaLLA uses 20 inducing points for the predictive variances. ELLA uses 20 random locations and 20 features.

DNNs. To simplify this, the Hessian is often approximated using the generalized Gauss-Newton (GGN) matrix (Martens and Grosse, 2015). This approach aligns with the fact that the GGN Hessian estimate is exact when applying LA to a linearized DNN. This method is referred to as Linearized LA (LLA) (Immer et al., 2021). However, in the context of large DNNs, further approximations on top of GGN are required.

In this work, we broaden LLA’s usage for DNNs to cases with a substantial number of parameters and training instances. We reinterpret the LLA predictive distribution as that of a Gaussian Process (GP) (Khan et al., 2019), and use a sparse variational GP based on inducing points to approximate it (Titsias, 2009). Standard sparse GPs often alter the predictive mean and hence may deteriorate the accuracy of the pre-trained DNN, which is assumed to have minimal generalization error. To avoid this, we use a dual representation of the GP in the Reproducing Kernel Hilbert Space (RKHS). This enables the sparse approximation to be used only in the computation of the predictive variances of the GP. The result is a sparse GP approximation where the predictive mean coincides with the predictions of the pre-trained DNN, without introducing additional prediction error.

We call our method Variational LLA (VaLLA) and conduct a series of experiments to compare its performance and computational cost with that of related methods from the literature. Our comparisons include (i) a method that uses Kronecker and diagonal approximations of the GGN matrix, (ii) a method that uses a Nyström approximation of that matrix (ELLA) (Deng et al., 2022), and (iii) a method that relies on generating samples from LLA’s posterior distribution using the sample-then-optimize principle (Antorán et al.,

2023). VaLLA’s performance is comparable or better than that of these other methods and is obtained at a smaller cost. Specifically, VaLLA adopts mini-batch training, resulting in sub-linear cost w.r.t. the training set size. Furthermore, VaLLA often generates predictive distributions that closely resemble those of LLA. Figure 1 shows this in a simple 1-D toy regression problem.

2. Background

Consider the task of inferring an unknown function $f : \mathbb{R}^D \rightarrow \mathbb{R}$ based on noisy observations $\mathbf{y} = (y_1, \dots, y_N)^\top$ at corresponding locations $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_N)$. Deep learning (DL) defines a Neural Network $g : \mathbb{R}^D \rightarrow \mathbb{R}^P \rightarrow \mathbb{R}$ with P parameters so that $\theta \in \mathbb{R}^P$ s.t. $f(\cdot) \approx g(\cdot; \theta)$.

The Laplace approximation (LA) builds a Gaussian approximate posterior (Mackay, 1992). This takes the form of $q(\cdot) = N(\cdot; \hat{\mathbf{j}})$, where $\hat{\cdot}$ denotes the MAP solution, *i.e.*, $\hat{\cdot} = \arg \max \log p(\mathbf{y}; \cdot) + \log p(\cdot)$, and Σ^{-1} is the inverse of the negative Hessian of the log posterior, *i.e.*, $\Sigma^{-1} = -\nabla^2 [\log p(\mathbf{y}; \cdot) + \log p(\cdot)]|_{\cdot = \hat{\cdot}}$. Often, an isotropic Gaussian prior $p(\cdot) = N(\cdot; \mathbf{0}; \frac{1}{2} \mathbf{I}_P)$ is considered. Due to the intractability of the Hessian in large DNNs, it is common to approximate it with the generalized Gauss-Newton (GGN) matrix (Immer et al., 2021):

$$\Sigma^{-1} \approx \sum_{n=1}^N J_\wedge(\mathbf{x}_n)^\top (\mathbf{x}_n; y_n) J_\wedge(\mathbf{x}_n) + \frac{1}{2} \mathbf{I}_P; \quad (1)$$

where $J_\wedge(\mathbf{x}_n) = \nabla_{\mathbf{g}} g(\mathbf{x}_n; \cdot)|_{\mathbf{g} = \hat{\mathbf{g}}}$ and $(\mathbf{x}_n; y_n) = \nabla_{\mathbf{g}}^2 \log p(y_n; \mathbf{g})|_{\mathbf{g} = \hat{\mathbf{g}}(\mathbf{x}_n)}$. The GGN matrix is guaranteed to be positive definite, which means that $\hat{\cdot}$ need not be a maximum of the log posterior. It can be, *e.g.*, any solution found by early-stopping back-propagation.

The earlier formulation of LA suffers from underfitting (Lawrence, 2001). This is attributed to the fact that the GGN approximation is the true Hessian matrix of the linearized DNN $g_\wedge^{\text{lin}}(\mathbf{x}; \cdot) := g(\mathbf{x}; \hat{\cdot}) + J_\wedge(\mathbf{x}_n)(\cdot - \hat{\cdot})$ (Immer et al., 2021). This implies a shift between posterior inference and predictions that can be mitigated by also predicting using the linearized model:

$$p_{\text{LLA}}(y^2; \mathbf{x}^2; \mathbf{y}) = \mathbb{E}_{q(\cdot)} \left[p(y^2; g_\wedge^{\text{lin}}(\mathbf{x}^2; \cdot)) \right]; \quad (2)$$

This method is known as the linearized LA (LLA). Despite these approximations, LLA still requires the inversion of Σ^{-1} which scales cubically with the number of parameters of the DNN. A dual formulation of LLA as a Gaussian Process (GP), described in the next section, scales with cubic cost in the number of training points.

2.1. Gaussian Process (GP) Interpretation of LLA

Using the approximate posterior from LLA, *i.e.*, $q(\cdot) = N(\cdot; \hat{\mathbf{j}})$, we also obtain a GP for prediction. The mean and covariance functions, providing the predictive mean and variances of the LLA approximation, are in this case $m(\mathbf{x}) = g(\mathbf{x}; \hat{\cdot})$; and $K(\mathbf{x}; \mathbf{x}^\beta) = J_\wedge(\mathbf{x})^\top J_\wedge(\mathbf{x}^\beta)$, respectively. Using the Woodbury formula on Σ^{-1} and defining $\mathbf{K}(\mathbf{x}; \mathbf{x}^\beta) = \frac{1}{2} J_\wedge(\mathbf{x})^\top J_\wedge(\mathbf{x}^\beta)$ as the (scaled) Neural Tangent Kernel (*i.e.*, prior covariance function) of the GP (Immer et al., 2021), and $\mathbf{Q} = \frac{1}{2} \mathbf{I}_{\mathbf{x}, \mathbf{y}} + \mathbf{K}(\mathbf{X}; \mathbf{X})$, the covariance function of the GP takes the expression

$$K(\mathbf{x}; \mathbf{x}^\beta) = \mathbf{K}(\mathbf{x}; \mathbf{x}^\beta) - \mathbf{K}(\mathbf{x}; \mathbf{X}) \mathbf{Q}^{-1} \mathbf{K}(\mathbf{X}; \mathbf{x}^\beta); \quad (3)$$

This allows us to interpret the LLA approximate predictive distribution as a posterior GP in function space, with prior covariance function given by $(\cdot; \cdot)$. The bottleneck of this interpretation is the evaluation of \mathbf{Q}^{-1} , with $O(N^3 + N^2P)$ cost. In the case of classification problems with C classes, $(\mathbf{X}; \mathbf{X})$ and \mathbf{Q} have size $NC \times NC$, thus the cost is also cubic in the number of classes C or DNN outputs.

2.2. Dual formulation of Gaussian Processes in RKHS

A Reproducing Kernel Hilbert Space (RKHS) H is a Hilbert space of functions with the reproducing property: $\forall \mathbf{x} \in X, \forall f \in H$ such that $\langle f, \delta_{\mathbf{x}} \rangle = f(\mathbf{x})$. In general, H can be infinite-dimensional and approximate continuous functions on a compact set.

A zero-mean prior GP with posterior mean and covariance functions $m(\cdot)$ and $K(\cdot; \cdot)$, respectively, has a dual representation in a RKHS (Cheng and Boots, 2016). There exists $\delta_{\mathbf{x}} \in H$ and a linear semi-definite positive operator $\mathcal{K} : H \rightarrow H$ s.t. $\langle \delta_{\mathbf{x}}, \delta_{\mathbf{x}'} \rangle = K(\mathbf{x}; \mathbf{x}')$ verifying $m(\mathbf{x}) = \langle \delta_{\mathbf{x}}, m \rangle$ and $K(\mathbf{x}; \mathbf{x}') = \langle \delta_{\mathbf{x}}, \mathcal{K} \delta_{\mathbf{x}'} \rangle$.

3. Variational LLA (VaLLA)

Variational sparse GPs approximate the GP posterior using a GP parameterized by M inducing points \mathbf{Z} , each in \mathbb{R}^D , and associated process values $\mathbf{u} = f(\mathbf{Z})$ (Titsias, 2009), $p(\mathbf{f}; \mathbf{u}) = q(\mathbf{f}; \mathbf{u}) = p(\mathbf{f}|\mathbf{u})q(\mathbf{u})$, where $q(\mathbf{u}) = N(\mathbf{u}|\hat{\mathbf{m}}; \hat{\mathbf{S}})$, $\mathbf{f} = f(\mathbf{X})$ and $p(\mathbf{f}|\mathbf{u})$ is fixed. The approximate distribution $q(\mathbf{u})$ is obtained by minimizing the KL-divergence $KL(q(\mathbf{f}; \mathbf{u}) \parallel p(\mathbf{f}; \mathbf{u}))$. In practice, the minimization problem is transformed into the maximization of the lower bound of the log-marginal likelihood

$$\log p(\mathbf{y}) \approx \max_{\mathbf{Z}, \hat{\mathbf{m}}, \hat{\mathbf{S}}} \int_{\mathbf{f}, \mathbf{u}} q(\mathbf{f}; \mathbf{u}) \log \frac{p(\mathbf{y}|\mathbf{f})p(\mathbf{f}|\mathbf{u})p(\mathbf{u})}{q(\mathbf{f}; \mathbf{u})} ; \quad (4)$$

with cost $O(NM^2 + M^3)$ due to $q(\mathbf{f}; \mathbf{u}) = p(\mathbf{f}|\mathbf{u})q(\mathbf{u})$.

Theorem 1 (Cheng and Boots, 2016) *Using a sparse GP approximation with $q(\mathbf{f}; \mathbf{u}) = p(\mathbf{f}|\mathbf{u})q(\mathbf{u})$ is equivalent to restricting the mean and covariance functions of the dual representation in the RKHS to $\tilde{m} = \mathbf{z}(\mathbf{a})$ and $\tilde{K} = I + \mathbf{z}\mathbf{A}\mathbf{z}^T$. Where the functional $\mathbf{z} : \mathbb{R}^M \rightarrow H$ defines a linear combination of basis functions as $\mathbf{z}(\mathbf{a}) = \sum_{m=1}^M a_m \mathbf{z}_m$, with $\mathbf{a} = (a_1; \dots; a_M) \in \mathbb{R}^M$ and the functional $\mathbf{z}\mathbf{A}\mathbf{z}^T = \sum_{i=1}^M \sum_{j=1}^M z_i A_{ij} z_j^T$, defines a quadratic expression where $\mathbf{A} \in \mathbb{R}^{M \times M}$ such that $\tilde{K} \succ 0$. Proof in Appendix A.*

Theorem 1 indicates that the algorithm of Titsias (2009) optimizes a variational Gaussian measure where \tilde{m} and \tilde{K} are parameterized by a function basis $f : \mathbf{z} \in H \rightarrow \mathbf{z} \in \mathbb{R}^M$. Cheng and Boots (2017) propose to generalize this so that each of the linear operators is optimized using different bases (sets of inducing points). Let \mathbf{Z} and \mathbf{Z}' be two sets of inducing points for the mean and the variance, respectively. The parameterization of Cheng and Boots (2017) is $\tilde{m} = \mathbf{z}(\mathbf{a})$ and $\tilde{K} = (I + \mathbf{z}\mathbf{A}\mathbf{z}^T)^{-1}$. Where $\mathbf{z} : \mathbb{R}^M \rightarrow H$ and $\mathbf{z}' : \mathbb{R}^M \rightarrow H$ are defined as \mathbf{z} using \mathbf{Z} and \mathbf{Z}' , respectively. Now, there are two sets of inducing points, M for the mean and M' for the covariances, respectively. This

parameterization is a generalization and cannot be obtained using the approach of Titsias (2009). Here, q must be found by optimizing Gaussian measures (Cheng and Boots, 2016):

$$\max_{q(f)} L(q(f)) = \max_{q(f)} \mathbb{E}_q [\log p(\mathbf{y}|\mathbf{f})] - \text{KL}(q \| p); \quad (5)$$

where $\text{KL}(q \| p) = \frac{1}{2} \mathbf{a}^T \mathbf{K}^{-1} \mathbf{a} + \frac{1}{2} \log |\mathbf{I} + \mathbf{K}^{-1} \mathbf{A}| + \frac{1}{2} \text{tr}(\mathbf{K}^{-1} (\mathbf{A}^{-1} + \mathbf{K}))$; with \mathbf{K} and \mathbf{K} matrices with the prior covariances among $f(\mathbf{Z})$ and $f(\mathbf{X})$, respectively. In practice, the objective leads to null variances and ∞ -divergences with Early-Stopping are used instead of the standard ELBO (see Appendix D for further information).

3.1. Using Decoupled SGP and LLA

We use the decoupled reparameterization of sparse GPs to establish a model where the mean of the approximated posterior distribution is fixed to a pre-trained MAP solution. We denote this method as *variational LLA* (VaLLA).

Proposition 2 *If $g(\cdot; \hat{\cdot}) \in H$, then $\delta > 0$ exists a set of M inducing points \mathbf{Z} and a collection of scalar values $\mathbf{a} \in \mathbb{R}^M$ such that the dual representation of the sparse Gaussian process defined by $\tilde{\cdot} = \mathbf{Z}(\mathbf{a})$ and $\tilde{\cdot} = (\mathbf{I} + \mathbf{Z} \mathbf{A} \mathbf{Z}^T)^{-1}$ corresponds to a GP posterior approximation with mean $m^2(\mathbf{x}) = h(\mathbf{x})$ and covariance function defined as*

$$K^2(\mathbf{x}; \mathbf{x}^\theta) = K(\mathbf{x}; \mathbf{x}^\theta) - \mathbf{K}_{\mathbf{x}; \mathbf{Z}} (\mathbf{A}^{-1} + \mathbf{K}_{\mathbf{Z}})^{-1} \mathbf{K}_{\mathbf{Z}; \mathbf{x}^\theta}; \quad (6)$$

where \mathbf{Z} is a set of M inducing points, $\mathbf{A} \in \mathbb{R}^{M \times M}$, $\mathbf{K}_{\mathbf{x}; \mathbf{Z}}$ is a vector with the covariances between $f(\mathbf{x})$ and $f(\mathbf{Z})$, and h verifies $d_H(g(\cdot; \hat{\cdot}); h) < \delta$, with $d_H(\cdot; \cdot)$ the distance in the RKHS. Proof in Appendix A.

Proposition 2 implies that if $g(\cdot; \hat{\cdot}) \in H$ we can find values for \mathbf{a} and inducing points for the mean \mathbf{Z} s.t. $d_H(g(\cdot; \hat{\cdot}); h)$ can be made as small as desired. For sufficiently small δ , $h(\cdot) \approx g(\cdot; \hat{\cdot})$, and $g(\cdot; \hat{\cdot})$ can be used for prediction instead of $h(\mathbf{x})$. Thus, there is no need to optimize \mathbf{a} and \mathbf{Z} in (5), and the posterior distribution of VaLLA uses $g(\cdot; \hat{\cdot})$ as its mean function. The optimal parameters \mathbf{Z} and \mathbf{A} can be found by optimizing (5) with \mathbf{a} and \mathbf{Z} held constant. From the following proposition, computing the optimal value of \mathbf{A} has cost $O(NM^2 + M^3)$.

Proposition 3 *The value of \mathbf{A} in Proposition 2 that minimizes Equation (5) is $\mathbf{A} = \frac{1}{2} \mathbf{K}_{\mathbf{Z}; \mathbf{X}}^{-1} \mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} \mathbf{K}_{\mathbf{Z}}^{-1}$ where σ^2 is the noise variance and $\mathbf{K}_{\mathbf{X}; \mathbf{Z}}$ is a matrix with the prior covariances between $f(\mathbf{X})$ and $f(\mathbf{Z})$. If $\mathbf{Z} = \mathbf{X}$, the covariance function of the predictive distribution in (6) is equal to that of the full GP. Proof in Appendix A.*

3.2. Limitations of VaLLA

VaLLA is limited by three factors: (i) Computing the predictive distribution at each training iteration involves inverting $\mathbf{A}^{-1} + \mathbf{K}_{\mathbf{Z}}$ in (6), with cubic cost in the number of inducing points M . Therefore, VaLLA cannot accommodate a very large number of inducing points. (ii) The minimization objective requires a validation set and the use of early-stopping to

Method	ResNet-20			ResNet-32			ResNet-44			ResNet-56			Mean Rank
	ACC	NLL	ECE	ACC	NLL	ECE	ACC	NLL	ECE	ACC	NLL	ECE	
MAP	92.6	0.282	0.039	93.5	0.292	0.041	94.0	0.275	0.039	94.4	0.252	0.037	
MF-VI	92.7	0.231	0.016	93.5	0.222	0.020	93.9	0.206	0.018	94.4	0.188	0.016	
LLA Diag	92.2	0.728	0.404	92.7	0.755	0.430	92.8	0.778	0.445	92.9	0.843	0.480	
LLA KFAC	92.0	0.852	0.467	91.8	1.027	0.547	91.4	1.091	0.566	89.8	1.174	0.579	
LLA	92.6	0.269	0.034	93.5	0.259	0.033	94.0	0.237	0.028	94.4	0.213	0.022	
LLA KFAC	92.6	0.271	0.035	93.5	0.260	0.033	94.0	0.232	0.028	94.4	0.202	0.024	
ELLA	92.5	0.233	0.009	93.5	0.215	0.008	93.9	0.204	0.007	94.4	0.187	0.007	2.375
Sampled LLA	92.5	0.231	0.006	93.5	0.217	0.008	94.0	0.200	0.007	94.4	0.185	0.015	2.000
VaLLA	92.6	0.228	0.007	93.5	0.211	0.007	94.0	0.198	0.008	94.4	0.183	0.009	1.375

Table 1: Results on CIFAR10. ACC, NLL and ECE are computed using Monte-Carlo estimation. Best value highlighted in **purple** and second to best in **teal**. Sampled LLA uses 64 samples. ELLA uses $M = 2000$ points and $K = 20$. Average results over 5 different random seeds (standard deviations $< 10^{-4}$ in all cases and omitted). for Last Layer LLA. Mean rank only considers both NLL and ECE.

effectively optimize the prior variance $\frac{2}{0}$, thus further increasing training time. (iii) Mini-batch optimization involves evaluating $K_{\mathbf{x}, \mathbf{z}}$, $\delta_{\mathbf{x}} \mathcal{L} B$ and \mathbf{K} . Hence, we require efficient evaluation of the (scaled) Neural Tangent Kernel, $(;) = {}^2 J_{\wedge}(\cdot)^T J_{\wedge}(\cdot)$ and its gradients to find \mathbf{Z} . While there are libraries that use structure in the derivatives for the efficient computation of $(;)$, these are limited to a few DNN models (Novak et al., 2022). A simple but inefficient approach to evaluate $(;)$ involves computing and storing all full Jacobians in memory, for each mini-batch instance and inducing point. This is tractable in our problems, but makes VaLLA infeasible in very large problems, *e.g.*, ImageNet. Appendix H.3 shows a very efficient *layer-by-layer* method to get $K_{\mathbf{x}, \mathbf{z}}$, $\delta_{\mathbf{x}} \mathcal{L} B$ and \mathbf{K} . However, this requires computing each layer’s contribution to the Jacobian at hand, which is difficult for large and complex DNNs.

4. Experiments

We compare VaLLA with other methods using the LLA implementation by Daxberger et al. (2021a) on various ResNet models trained on Cifar10 data. VaLLA uses a batch size of 100. We extract the results for other methods from Deng et al. (2022) and use the same DNN for VaLLA. The corresponding pre-trained models are those of Deng et al. (2022) (accessible [here](#)). Table 1 shows ACC, NLL and ECE for each method, including LLA variants and a mean-field VI approach (Deng and Zhu, 2023). We compare with Antorán et al. (2023)’s approach to retrieve samples from the true GP posterior defined by LLA. Since we use the same pre-trained models, the results of all other methods are consistent with those reported by Deng et al. (2022). VaLLA with $M = 100$ outperforms other methods in most cases, always being either the best or second-best method.

Appendix G contains further experimentation of the proposed method. More precisely, regression experiments in Year, Airline and Taxi datasets are reported; along with further classification results in MNIST and FMNIST with OOD detection.

5. Conclusions

We introduced VaLLA, a method derived from the formulation of a generalized sparse GP that offers the flexibility to fix the predictive mean to any desired function in the RKHS. VaLLA excels in computing error bars for pre-trained DNNs with a vast number of parameters on extensive datasets, handling even millions of training instances. VaLLA’s applicability spans both regression and classification problems, showcasing costs independent of the number of training points N . In comparison, the Nyström approximation by [Deng et al. \(2022\)](#) incurs a linear cost in N , unless early-stopping is employed. Furthermore, VaLLA surpasses the sample-then-optimize method of [Antorán et al. \(2023\)](#) in terms of speed, while also providing predictive distributions robust to input corruptions. In essence, VaLLA stands out by delivering robust predictive distributions akin to LLA, all while maintaining noteworthy computational efficiency.

Acknowledgments

Authors gratefully acknowledge the use of the facilities of Centro de Computacion Científica (CCC) at Universidad Autónoma de Madrid. The authors acknowledge financial support from project PID2022-139856NB-I00 funded by MCIN/ AEI / 10.13039/501100011033 / FEDER, UE and project PID2019-106827GB-I00 / AEI / 10.13039/501100011033 and from the Autonomous Community of Madrid (ELLIS Unit Madrid). The authors also acknowledge financial support from project TED2021-131530B-I00, funded by MCIN/AEI /10.13039/501100011033 and by the European Union NextGenerationEU PRTR.

References

- Javier Antorán, Shreyas Padhy, Riccardo Barbano, Eric T. Nalisnick, David Janz, and José Miguel Hernández-Lobato. Sampling-based inference for large linear models, with application to linearised Laplace. In *International Conference on Learning Representations*, 2023.
- Federico Bergamin, Pablo Moreno-Muñoz, Søren Hauberg, and Georgios Arvanitidis. Riemannian Laplace approximations for Bayesian neural networks. *Advances in Neural Information Processing Systems*, 2023.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622, 2015.
- Aleksandar Botev, Hippolyt Ritter, and David Barber. Practical Gauss-Newton optimisation for deep learning. In *International Conference on Machine Learning*, pages 557–565, 2017.
- Ching-An Cheng and Byron Boots. Incremental variational sparse Gaussian process regression. *Advances in Neural Information Processing Systems*, 29:4403–4411, 2016.

- Ching-An Cheng and Byron Boots. Variational inference for Gaussian process models with linear complexity. *Advances in Neural Information Processing Systems*, 30:5184–5194, 2017.
- Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace Redux - Effortless Bayesian deep learning. *Advances in Neural Information Processing Systems*, 34:20089–20103, 2021a.
- Erik Daxberger, Eric Nalisnick, James U Allingham, Javier Antoran, and Jose Miguel Hernandez-Lobato. Bayesian deep learning via subnetwork inference. In *International Conference on Machine Learning*, pages 2510–2521, 2021b.
- Zhijie Deng and Jun Zhu. Bayesadapter: Being Bayesian, inexpensively and reliably, via Bayesian fine-tuning. In *Asian Conference on Machine Learning*, pages 280–295, 2023.
- Zhijie Deng, Feng Zhou, and Jun Zhu. Accelerated linearized Laplace approximation for Bayesian deep learning. *Advances in Neural Information Processing Systems*, 35:2695–2708, 2022.
- Vincent Dutoridoir, Nicolas Durrande, and James Hensman. Sparse Gaussian processes with spherical harmonic features. In *International Conference on Machine Learning*, pages 2793–2802, 2020.
- Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American statistical Association*, 102:359–378, 2007.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330, 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- Alexander Immer, Maciej Korzepa, and Matthias Bauer. Improving predictions of Bayesian neural nets via local linearization. In *International Conference on Artificial Intelligence and Statistics*, pages 703–711, 2021.
- Pavel Izmailov, Wesley J Maddox, Polina Kirichenko, Timur Garipov, Dmitry Vetrov, and Andrew Gordon Wilson. Subspace inference for Bayesian deep learning. In *Uncertainty in Artificial Intelligence*, pages 1169–1179, 2020.
- Alex Kendall and Yarin Gal. What uncertainties do we need in Bayesian deep learning for computer vision? *Advances in Neural Information Processing Systems*, 30:5574–5584, 2017.
- Mohammad Emtiyaz E Khan, Alexander Immer, Ehsan Abedi, and Maciej Korzepa. Approximate inference turns deep networks into Gaussian processes. *Advances in Neural Information Processing Systems*, 32, 2019.

- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference for Learning Representations*, 2015.
- Neil David Lawrence. *Variational inference in probabilistic models*. PhD thesis, Citeseer, 2001.
- Jongseok Lee, Jianxiang Feng, Matthias Humt, Marcus Gerhard Müller, and Rudolph Triebel. Trust your robots! predictive uncertainty estimation of neural networks with sparse Gaussian processes. In *Conference on Robot Learning*, pages 1168–1179, 2022.
- Christian Leibig, Vaneeda Allken, Murat Seçkin Ayhan, Philipp Berens, and Siegfried Wahl. Leveraging uncertainty information from deep neural networks for disease detection. *Scientific reports*, 7:1–14, 2017.
- Jihao Andreas Lin, Javier Antorán, Shreyas Padhy, David Janz, José Miguel Hernández-Lobato, and Alexander Terenin. Sampling from Gaussian process posteriors using stochastic gradient descent. *Advances in Neural Information Processing Systems*, 36, 2024.
- David JC MacKay. The evidence framework applied to classification networks. *Neural computation*, 4:720–736, 1992.
- David John Cameron Mackay. *Bayesian methods for adaptive models*. California Institute of Technology, 1992.
- James Martens. New insights and perspectives on the natural gradient method. *The Journal of Machine Learning Research*, 21:5776–5851, 2020.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning*, pages 2408–2417, 2015.
- Roman Novak, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. Fast finite width neural tangent kernel. In *International Conference on Machine Learning*, pages 17018–17044, 2022.
- Luis A Ortega, Simón Rodríguez Santana, and Daniel Hernández-Lobato. Deep variational implicit processes. In *International Conference of Learning Representations*, 2023.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift? *Advances in Neural Information Processing Systems*, pages 13969–13980, 2019.
- Hippolyt Ritter, Aleksandar Botev, and David Barber. A scalable Laplace approximation for neural networks. In *International Conference on Learning Representations*, volume 6, 2018.
- Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep Gaussian processes. *Advances in Neural Information Processing Systems*, 30:4588–4599, 2017.

Aidan Scannell, Riccardo Mereu, Paul Chang, Ella Tamir, Joni Pajarinen, and Arno Solin. Function-space parameterization of neural networks for sequential learning. *International Conference on Learning Representations*, 2024.

Michalis Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial Intelligence and Statistics*, pages 567–574, 2009.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.

Appendix A. Proofs

Theorem 1 (*Cheng and Boots, 2016*) Using a sparse GP approximation with $q(\mathbf{f}; \mathbf{u}) = p(\mathbf{f}|\mathbf{u})q(\mathbf{u})$ is equivalent to restricting the mean and covariance functions of the dual representation in the RKHS to $\tilde{\mu} = \mathbf{z}(\mathbf{a})$ and $\tilde{\Sigma} = I + \mathbf{z}\mathbf{A}\mathbf{z}^T$. Where the functional $\mathbf{z} : \mathbb{R}^M \rightarrow H$ defines a linear combination of basis functions as $\mathbf{z}(\mathbf{a}) = \sum_{m=1}^M a_m \mathbf{z}_m$, with $\mathbf{a} = (a_1; \dots; a_M) \in \mathbb{R}^M$ and the functional $\mathbf{z}\mathbf{A}\mathbf{z}^T = \sum_{i=1}^M \sum_{j=1}^M \mathbf{z}_i \mathbf{A}_{ij} \mathbf{z}_j^T$, defines a quadratic expression where $\mathbf{A} \in \mathbb{R}^{M \times M}$ such that $\tilde{\Sigma} \succ 0$.

Proof First of all, notice that $\mathbf{z} = [\mathbf{z}_1; \dots; \mathbf{z}_M] \in H^M$, leading to

$$K_{\mathbf{z}} := \mathbf{z}\mathbf{z}^T = \begin{pmatrix} h_{\mathbf{z}_1; \mathbf{z}_1} & h_{\mathbf{z}_1; \mathbf{z}_2} & \dots & h_{\mathbf{z}_1; \mathbf{z}_M} \\ \vdots & \vdots & \ddots & \vdots \\ h_{\mathbf{z}_M; \mathbf{z}_1} & h_{\mathbf{z}_M; \mathbf{z}_2} & \dots & h_{\mathbf{z}_M; \mathbf{z}_M} \end{pmatrix} \in \mathbb{R}^{M \times M}. \quad (7)$$

Furthermore, $K_{\mathbf{x}; \mathbf{z}}$ defined as $\mathbf{v} \mapsto h_{\mathbf{x}; \mathbf{z}(\mathbf{v})}$ can be seen as a vector of \mathbb{R}^M , considering the image of any orthonormal basis of \mathbb{R}^M . In fact, let $\mathbf{e}_1; \dots; \mathbf{e}_M$ be the usual basis of \mathbb{R}^M :

$$K_{\mathbf{x}; \mathbf{z}} = \begin{pmatrix} h_{\mathbf{x}; \mathbf{z}(\mathbf{e}_1)} \\ h_{\mathbf{x}; \mathbf{z}(\mathbf{e}_2)} \\ \vdots \\ h_{\mathbf{x}; \mathbf{z}(\mathbf{e}_M)} \end{pmatrix} = \begin{pmatrix} h_{\mathbf{x}; \sum_{i=1}^M \mathbf{e}_{1;i} \mathbf{z}_i} \\ h_{\mathbf{x}; \sum_{i=1}^M \mathbf{e}_{2;i} \mathbf{z}_i} \\ \vdots \\ h_{\mathbf{x}; \sum_{i=1}^M \mathbf{e}_{M;i} \mathbf{z}_i} \end{pmatrix} = \begin{pmatrix} h_{\mathbf{x}; \mathbf{z}_1} \\ h_{\mathbf{x}; \mathbf{z}_2} \\ \vdots \\ h_{\mathbf{x}; \mathbf{z}_M} \end{pmatrix} \in \mathbb{R}^M. \quad (8)$$

Assume a variational distribution $q(\mathbf{u}) = N(\mathbf{u}|\mathbf{m}; \mathbf{S})$ with $\mathbf{m} \in \mathbb{R}^M$ and $\mathbf{S} \in \mathbb{R}^{M \times M}$. Then, defining the correspondent dual vectors as

$$\mathbf{a} = K_{\mathbf{z}}^{-1} \mathbf{m} \in \mathbb{R}^M; \quad \mathbf{A} = K_{\mathbf{z}}^{-1} \mathbf{S} K_{\mathbf{z}}^{-1} \in \mathbb{R}^{M \times M}; \quad (9)$$

the mean and covariance functions in the dual formulation of the sparse GP $q(f)$ are

$$m^2(\mathbf{x}) = h_{\mathbf{x}; \tilde{\mu}} = h_{\mathbf{x}; \mathbf{z}(K_{\mathbf{z}}^{-1} \mathbf{m})} = K_{\mathbf{x}; \mathbf{z}} K_{\mathbf{z}}^{-1} \mathbf{m}; \quad (10)$$

and

$$\begin{aligned} K^2(\mathbf{x}; \mathbf{x}^\theta) &= h_{\mathbf{x}; (\cdot | \mathbf{x}^\theta)} = h_{\mathbf{x}; \mathbf{x}^\theta} + \mathbf{z}\mathbf{A}\mathbf{z}^T \mathbf{x}^\theta = h_{\mathbf{x}; \mathbf{x}^\theta} + h_{\mathbf{x}; \mathbf{z}\mathbf{A}\mathbf{z}^T \mathbf{x}^\theta} \\ &= K(\mathbf{x}; \mathbf{x}^\theta) + K_{\mathbf{x}; \mathbf{z}} \mathbf{A} K_{\mathbf{x}^\theta; \mathbf{z}}^T = K(\mathbf{x}; \mathbf{x}^\theta) + K_{\mathbf{x}; \mathbf{z}} (K_{\mathbf{z}}^{-1} \mathbf{S} K_{\mathbf{z}}^{-1} - K_{\mathbf{z}}^{-1}) K_{\mathbf{z}; \mathbf{x}^\theta}; \end{aligned} \quad (11)$$

which is the same approximate GP posterior $p(\mathbf{f}) = \int_{\mathbb{R}} p(\mathbf{f}|\mathbf{u})q(\mathbf{u}) d\mathbf{u}$ found in Equation (6) of Titsias (2009). \blacksquare

Proposition 2 If $g(\cdot) \in H$, then $\delta > 0$ exists a set of M inducing points \mathbf{Z} and a collection of scalar values $\mathbf{a} \in \mathbb{R}^M$ such that the dual representation of the sparse Gaussian process defined by

$$\tilde{\mu} = \mathbf{z}(\mathbf{a}) \quad \text{and} \quad \tilde{\Sigma} = (I + \mathbf{z}\mathbf{A}\mathbf{z}^T)^{-1}; \quad (12)$$

corresponds to a GP posterior approximation with mean and covariance functions defined as

$$\begin{aligned} m^2(\mathbf{x}) &= h(\mathbf{x}); \\ K^2(\mathbf{x}; \mathbf{x}^\theta) &= K(\mathbf{x}; \mathbf{x}^\theta) - \mathbf{K}_{\mathbf{x}; \mathbf{Z}} (\mathbf{A}^{-1} + \mathbf{K})^{-1} \mathbf{K}_{\mathbf{Z}; \mathbf{x}^\theta}; \end{aligned} \quad (13)$$

where \mathbf{Z} is a set of M inducing points, $\mathbf{A} \in \mathbb{R}^{M \times M}$, $\mathbf{K}_{\mathbf{x}; \mathbf{Z}}$ is a vector with the covariances between $f(\mathbf{x})$ and $f(\mathbf{Z})$, and h verifies $d_H(g(\cdot; \hat{\cdot}); h) \leq \delta$, with $d_H(\cdot; \cdot)$ the distance in the RKHS.

Proof First of all, if $g(\cdot; \hat{\cdot}) \in H$, the reproducing property of the RKHS verifies that $\exists \mathbf{a} \in \mathbb{R}^M$ such that $h := \sum_{i=1}^M a_i \mathbf{z}_i = \mathbf{z}^T \mathbf{a}$ verifies

$$d_H(g(\cdot; \hat{\cdot}); h) \leq \delta. \quad (14)$$

As a result, the mean function of the approximate posterior is

$$m^2(\mathbf{x}) = h_{\mathbf{x}; \sim} = h(\mathbf{x}) - g(\mathbf{x}; \hat{\cdot}). \quad (15)$$

On the other hand, using that

$$(I + \mathbf{Z}^T \mathbf{A}^{-1} \mathbf{Z})^{-1} = I - \mathbf{Z} (\mathbf{A}^{-1} + \mathbf{Z}^T \mathbf{Z})^{-1} \mathbf{Z}^T; \quad (16)$$

the covariance function is

$$\begin{aligned} K^2(\mathbf{x}; \mathbf{x}^\theta) &= h_{\mathbf{x}; \sim}(\mathbf{x}^\theta) - h(\mathbf{x}; \mathbf{x}^\theta) + h_{\mathbf{x}; \mathbf{Z}} (\mathbf{A}^{-1} + \mathbf{Z}^T \mathbf{Z})^{-1} \mathbf{Z}^T \mathbf{x}^\theta \\ &= K(\mathbf{x}; \mathbf{x}^\theta) - \mathbf{K}_{\mathbf{x}; \mathbf{Z}} (\mathbf{A}^{-1} + \mathbf{K}_{\mathbf{Z}})^{-1} \mathbf{K}_{\mathbf{Z}; \mathbf{x}^\theta}; \end{aligned} \quad (17)$$

Where the characterization of $\mathbf{K}_{\mathbf{Z}}$ and $\mathbf{K}_{\mathbf{x}; \mathbf{Z}}$ as elements of $\mathbb{R}^{M \times M}$ and \mathbb{R}^M of Equations (7) and (8) are used. \blacksquare

Proposition 3 The value of \mathbf{A} in Proposition 2 that minimizes Equation (5) is

$$\mathbf{A} = \frac{1}{2} \mathbf{K}^{-1} \mathbf{K}_{\mathbf{Z}; \mathbf{x}} \mathbf{K}_{\mathbf{x}; \mathbf{Z}} \mathbf{K}^{-1}; \quad (18)$$

where σ^2 is the noise variance and $\mathbf{K}_{\mathbf{x}; \mathbf{Z}}$ is a matrix with the prior covariances between $f(\mathbf{X})$ and $f(\mathbf{Z})$. If $\mathbf{Z} = \mathbf{X}$, the covariance function of the predictive distribution in (6) is equal to that of the full GP.

Proof For simplicity, assume a non-inverse reparameterization where

$$q(\mathbf{f}) = N(\mathbf{f} | \mathbf{Z}^T \hat{\mathbf{A}} \mathbf{Z}; I + \mathbf{Z}^T \hat{\mathbf{A}} \mathbf{Z}) \quad (19)$$

We will find the optimal value for $\hat{\mathbf{A}}$ and compute the corresponding value for \mathbf{A} (using the inverse reparameterization). First, we will show that the true GP posterior can be written in the dual formulation as the following

$$p(\mathbf{f} | \mathbf{y}) = N(\mathbf{f} | \mathbf{X}^T (\mathbf{X}^T \mathbf{X} + \sigma^2 \mathbf{I})^{-1} \mathbf{X} \mathbf{y}; \sigma^2 (\mathbf{X}^T \mathbf{X} + \sigma^2 \mathbf{I})^{-1}); \quad (20)$$

where it verifies that the GP posterior mean function is

$$\begin{aligned} m(\mathbf{x}) &= h_{\mathbf{x}}: (\mathbf{X}^T \mathbf{X} + \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y} = (\mathbf{X}^T \mathbf{X} + \mathbf{I})^{-1} \mathbf{X}^T \mathbf{x} \mathbf{y} \\ &= (K_{\mathbf{X};\mathbf{X}} + \mathbf{I})^{-1} K_{\mathbf{X};\mathbf{x}} \mathbf{y} = K_{\mathbf{x};\mathbf{X}} (K_{\mathbf{X};\mathbf{X}} + \mathbf{I})^{-1} \mathbf{y}; \end{aligned} \quad (21)$$

Where the characterization of $K_{\mathbf{X};\mathbf{X}}$ and $K_{\mathbf{X};\mathbf{x}}$ as elements of $\mathbb{R}^{N \times N}$ and \mathbb{R}^N is used, similarly to Equations (7) and (8). Furthermore, using Woodbury matrix identity,

$$(\mathbf{X}^T \mathbf{X} + \mathbf{I})^{-1} = \mathbf{I} - \mathbf{X} (K_{\mathbf{X};\mathbf{X}} + \mathbf{I})^{-1} \mathbf{X}^T; \quad (22)$$

where again, correspondence between operators and matrices is used. This leads to the covariance function:

$$\begin{aligned} K(\mathbf{x}; \mathbf{x}^\theta) &= h_{\mathbf{x}}: (\mathbf{X}^T \mathbf{X} + \mathbf{I})^{-1} (\mathbf{x}^\theta)^T = h_{\mathbf{x}}: \mathbf{I} - \mathbf{X} (K_{\mathbf{X};\mathbf{X}} + \mathbf{I})^{-1} \mathbf{X}^T (\mathbf{x}^\theta)^T \\ &= K_{\mathbf{x};\mathbf{x}^\theta} - K_{\mathbf{x};\mathbf{X}} (K_{\mathbf{X};\mathbf{X}} + \mathbf{I})^{-1} K_{\mathbf{X};\mathbf{x}^\theta}; \end{aligned} \quad (23)$$

This mean and covariance function are exactly the ones obtained from the original GP formulation (Titsias, 2009). Using the ELBO is equivalent to the KL divergence between the true posterior and the variational approximation, we got

$$\text{KL}(q(f) \parallel p(f|\mathbf{y})) \approx \frac{1}{2} \text{tr}(\mathbf{B}^{-1} (\mathbf{I} + \mathbf{Z} \hat{\mathbf{A}} \mathbf{Z}^T)) - \frac{1}{2} \ln |\mathbf{I} + \mathbf{Z} \hat{\mathbf{A}} \mathbf{Z}^T|; \quad (24)$$

where

$$\mathbf{B} = (\mathbf{X}^T \mathbf{X} + \mathbf{I})^{-1}; \quad (25)$$

Naming $\mathbf{M} = \mathbf{I} + \mathbf{Z} \hat{\mathbf{A}} \mathbf{Z}^T$, it is important to notice that given $\mathbf{Z} \hat{\mathbf{A}} \mathbf{Z}^T = \sum_{i=1}^M \sum_{j=1}^M \mathbf{z}_i \hat{\mathbf{a}}_{ij} \mathbf{z}_j^T$, despite being an operator, \mathbf{M} can be seen as a matrix whose entries are the application of $\mathbf{I} + \mathbf{Z} (\cdot) \mathbf{Z}^T$ to the usual basis of matrices. In short

$$\mathbf{M} = \mathbf{I} + \begin{matrix} \text{O} & & & & & \\ & \mathbf{z}_1 \hat{\mathbf{a}}_{1;1} \mathbf{z}_1^T & & & & \\ & \vdots & & & & \\ & & & \mathbf{z}_M \hat{\mathbf{a}}_{M;1} \mathbf{z}_1^T & & \\ & & & \vdots & & \\ & & & & & \mathbf{z}_M \hat{\mathbf{a}}_{M;M} \mathbf{z}_M^T \end{matrix} \in \mathbb{R}^{M \times M}; \quad (26)$$

The partial derivative of \mathbf{M} w.r.t. a single position in the matrix $\hat{\mathbf{A}}$ is $\frac{\partial \mathbf{M}}{\partial \hat{\mathbf{a}}_{ij}} = (\mathbf{z}_i) (\mathbf{z}_j)^T$, where \mathbf{z}_j denotes a zero-vector with a single 1 at position j . Then, using the chain rule for matrices,

$$\frac{\partial g(\mathbf{U})}{\partial X_{ij}} = \text{tr} \left(\frac{\partial g(\mathbf{U})}{\partial \mathbf{U}} \frac{\partial \mathbf{U}}{\partial X_{ij}} \right); \quad (27)$$

we can compute the optimum in the two terms in the KL. The optimum for the logarithm term can be computed as

$$\begin{aligned} \frac{\partial \ln(j\mathbf{M}j)}{\partial \hat{\mathbf{a}}_{ij}} &= \text{tr} \left(\frac{\partial \ln(j\mathbf{M}j)}{\partial \mathbf{M}} \frac{\partial \mathbf{M}}{\partial \hat{\mathbf{a}}_{ij}} \right) \\ &= \text{tr}(\mathbf{M}^{-1} (\mathbf{z}_i) (\mathbf{z}_j)^T) \\ &= \text{tr}((\mathbf{z}_j)^T \mathbf{M}^{-1} (\mathbf{z}_i)) \\ &= (\mathbf{z}_j^T \mathbf{M}^{-1} \mathbf{z}_i); \end{aligned} \quad (28)$$

leading to

$$\frac{\partial \ln(j\mathbf{M}j)}{\partial \hat{\mathbf{A}}} = \frac{\mathbf{Z}^T \mathbf{M}^{-1} \mathbf{z}}{\mathbf{Z}^T (\mathbf{I} + \mathbf{z} \hat{\mathbf{A}} \mathbf{z}^T)^{-1} \mathbf{z}} : \quad (29)$$

On the other hand, the optimum for the trace term is

$$\begin{aligned} \frac{\partial \text{tr}(\mathbf{B}^{-1} \mathbf{M})}{\partial \hat{\mathbf{a}}_{ij}} &= \text{tr} \frac{\partial \text{tr}(\mathbf{B}^{-1} \mathbf{M})^T \partial \mathbf{M}}{\partial \mathbf{M} \partial \hat{\mathbf{a}}_{ij}} \\ &= \text{tr} \mathbf{B}^{-1} (\mathbf{z}_i) (\mathbf{z}_j)^T \\ &= \text{tr}((\mathbf{z}_j)^T \mathbf{B}^{-1} (\mathbf{z}_i)) \\ &= (\mathbf{Z}^T \mathbf{B}^{-1} \mathbf{z})_{ji}; \end{aligned} \quad (30)$$

where we used that $\mathbf{B} = \mathbf{B}^T$. As a result,

$$\begin{aligned} \frac{\partial \text{tr}(\mathbf{B}^{-1} \mathbf{M})}{\partial \hat{\mathbf{A}}} &= \frac{\mathbf{Z}^T \mathbf{B}^{-1} \mathbf{z}}{\mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z}} : \end{aligned} \quad (31)$$

Using all the derivations

$$\frac{\partial \text{KL}(q(f)j p(fjy))}{\partial \hat{\mathbf{A}}} = 0 \quad \left(\frac{\mathbf{Z}^T (\mathbf{B}^{-1} (\mathbf{I} + \mathbf{z} \hat{\mathbf{A}} \mathbf{z}^T)^{-1} \mathbf{z})}{\mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z}} = 0 \right) \quad (32)$$

Using the Woodbury matrix identity

$$\begin{aligned} 0 &= \frac{\mathbf{Z}^T (\mathbf{B}^{-1} (\mathbf{I} + \mathbf{z} \hat{\mathbf{A}} \mathbf{z}^T)^{-1} \mathbf{z})}{\mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z}} \\ &= \frac{\mathbf{Z}^T (\mathbf{B}^{-1} (\mathbf{I} + \mathbf{z} (\hat{\mathbf{A}}^{-1} + \mathbf{K})^{-1} \mathbf{z}^T) \mathbf{z})}{\mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z}} \\ &= \frac{\mathbf{Z}^T \mathbf{B}^{-1} \mathbf{z} (\mathbf{K} + \mathbf{K} (\hat{\mathbf{A}}^{-1} + \mathbf{K})^{-1} \mathbf{K})}{\mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z}} : \end{aligned} \quad (33)$$

Thus, the value of $\hat{\mathbf{A}}$ where $\partial \text{KL} = \partial \hat{\mathbf{A}} = 0$ verifies

$$\hat{\mathbf{A}} = ((\mathbf{K}^{-1} \mathbf{K}^{-1} \mathbf{Z}^T \mathbf{B}^{-1} \mathbf{z} \mathbf{K}^{-1})^{-1} \mathbf{K})^{-1} : \quad (34)$$

Using that $\mathbf{B}^{-1} = \mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z}$, we can take further derivations on the expression of $\hat{\mathbf{A}}$ as

$$\begin{aligned} \hat{\mathbf{A}} &= ((\mathbf{K}^{-1} \mathbf{K}^{-1} \mathbf{Z}^T \mathbf{B}^{-1} \mathbf{z} \mathbf{K}^{-1})^{-1} \mathbf{K})^{-1} \\ &= ((\mathbf{K}^{-1} \mathbf{K}^{-1} \mathbf{Z}^T (\mathbf{X} \mathbf{X}^T + \mathbf{I})^{-1} \mathbf{z} \mathbf{K}^{-1})^{-1} \mathbf{K})^{-1} \\ &= ((\mathbf{K}^{-1} \mathbf{Z}^T \mathbf{K}^{-1} \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} \mathbf{K}^{-1} \mathbf{K}^{-1})^{-1} \mathbf{K})^{-1} \\ &= ((\mathbf{Z}^T \mathbf{K}^{-1} \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} \mathbf{K}^{-1})^{-1} \mathbf{K})^{-1} \\ &= ((\mathbf{Z}^T \mathbf{K}^{-1} \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} \mathbf{K}^{-1})^{-1} + \mathbf{K})^{-1} : \end{aligned} \quad (35)$$

Applying again Woodbury matrix identity:

$$\begin{aligned} \hat{\mathbf{A}} &= ((\mathbf{Z}^T \mathbf{K}^{-1} \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} \mathbf{K}^{-1})^{-1} + \mathbf{K})^{-1} \\ &= (\mathbf{K}^{-1} \mathbf{K}^{-1} (\mathbf{Z}^T \mathbf{K}^{-1} \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} \mathbf{K}^{-1} + \mathbf{K})^{-1} \mathbf{K}^{-1}) \\ &= (\mathbf{K}^{-1} (\mathbf{Z}^T \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} + \mathbf{K})^{-1}) \\ &= \mathbf{K}^{-1} + (\mathbf{Z}^T \mathbf{K}_Z ; \mathbf{X} \mathbf{K}_X ; \mathbf{Z} + \mathbf{K})^{-1} : \end{aligned} \quad (36)$$

If we substitute this value on the predictive distribution

$$\begin{aligned}
 K^?(\mathbf{x}; \mathbf{x}^\ell) &= h_{\mathbf{x}; \mathbf{M}(\mathbf{x}^\ell)} = K(\mathbf{x}; \mathbf{x}^\ell) + K_{\mathbf{x}; \mathbf{Z}} \hat{\mathbf{A}} K_{\mathbf{Z}; \mathbf{x}^\ell} \\
 &= K(\mathbf{x}; \mathbf{x}^\ell) + K_{\mathbf{x}; \mathbf{Z}} (\mathbf{K}^{-1} + (\mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} + \mathbf{K})^{-1}) K_{\mathbf{Z}; \mathbf{x}^\ell} \\
 &= K(\mathbf{x}; \mathbf{x}^\ell) - K_{\mathbf{x}; \mathbf{Z}} \mathbf{K}^{-1} K_{\mathbf{Z}; \mathbf{x}^\ell} + K_{\mathbf{x}; \mathbf{Z}} (\mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} + \mathbf{K})^{-1} K_{\mathbf{Z}; \mathbf{x}^\ell} \\
 &= K(\mathbf{x}; \mathbf{x}^\ell) - K_{\mathbf{x}; \mathbf{Z}} \mathbf{K}^{-1} K_{\mathbf{Z}; \mathbf{x}^\ell} \\
 &\quad + K_{\mathbf{x}; \mathbf{Z}} \mathbf{K}^{-1} (\mathbf{K} (\mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} + \mathbf{K})^{-1} \mathbf{K}) \mathbf{K}^{-1} K_{\mathbf{Z}; \mathbf{x}^\ell}.
 \end{aligned} \tag{37}$$

This expression coincides with the optimal sparse GP solution described by Titsias (2009), with optimal variational covariance $(\mathbf{K} (\mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} + \mathbf{K})^{-1} \mathbf{K})$. As a result, the optimal solution for VaLLA coincides with the optimal solution for standard sparse GPs. Let us now compute the optimal value of \mathbf{A} given the optimal value of $\hat{\mathbf{A}}$. First, notice that using Woodbury Matrix identity

$$(\mathbf{I} + \mathbf{Z} \mathbf{A} \mathbf{Z}^T)^{-1} = \mathbf{I} - \mathbf{Z} (\mathbf{A} + \mathbf{K})^{-1} \mathbf{Z}^T. \tag{38}$$

Therefore, the relation between \mathbf{A} and $\hat{\mathbf{A}}$ is $\hat{\mathbf{A}} = (\mathbf{A} + \mathbf{K})^{-1}$. Meaning that

$$\hat{\mathbf{A}} = \mathbf{K}^{-1} + (\mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} + \mathbf{K})^{-1} \Rightarrow \mathbf{A} = \frac{1}{2} \mathbf{K}^{-1} \mathbf{K}_{\mathbf{Z}; \mathbf{X}} \mathbf{K}_{\mathbf{X}; \mathbf{Z}} \mathbf{K}^{-1}. \tag{39}$$

Global optimum To complete the proof of the solution in (39) being not only optimal but also a maximum of the ELBO, we must test the behavior of the second derivative w.r.t. \mathbf{A} . Let us reuse the previous results, where we found that

$$\frac{\partial^2 \text{KL}(q(f) \parallel p(f|y))}{\partial \hat{a}_{ij}} = \frac{1}{2} \mathbf{Z}^T (\mathbf{B}^{-1} \mathbf{Z} - \mathbf{M}^{-1} \mathbf{Z})_i \tag{40}$$

Taking a second derivative w.r.t. another location $\hat{a}_{u,v}$ yields

$$\frac{\partial}{\partial \hat{a}_{u,v}} \frac{\partial \text{KL}(q(f) \parallel p(f|y))}{\partial \hat{a}_{ij}} = \frac{1}{2} \frac{\partial}{\partial \hat{a}_{u,v}} \mathbf{Z}^T (\mathbf{B}^{-1} \mathbf{Z} - \mathbf{M}^{-1} \mathbf{Z})_i \tag{41}$$

Considering that \mathbf{B} does not depend on $\hat{\mathbf{A}}$, the first term drops from the derivative. Thus

$$\frac{\partial}{\partial \hat{a}_{u,v}} \frac{\partial \text{KL}(q(f) \parallel p(f|y))}{\partial \hat{a}_{ij}} = \frac{1}{2} \frac{\partial}{\partial \hat{a}_{u,v}} \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z}_i \tag{42}$$

Here we aim to use the chain rule for matrices,

$$\frac{\partial \mathbf{g}(U)}{\partial \mathbf{X}_{ij}} = \text{tr} \left(\frac{\partial \mathbf{g}(U)}{\partial U} \frac{\partial U}{\partial \mathbf{X}_{ij}} \right); \tag{43}$$

Then, consider that

$$\frac{\partial \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z}_i}{\partial \mathbf{M}} = -\mathbf{M}^{-1} \mathbf{Z}_j \mathbf{Z}_i^T \mathbf{M}^{-1} \tag{44}$$

Then,

$$\frac{\partial}{\partial \hat{\mathbf{a}}_{u,v}} \text{tr} \left(\mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \right)^{-1} = \text{tr} \left(\mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \right)^{-1} \quad (45)$$

We can work on this trace to simplify the expression as

$$\begin{aligned} \frac{\partial}{\partial \hat{\mathbf{a}}_{u,v}} \text{tr} \left(\mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \right)^{-1} &= \text{tr} \left(\mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \right)^{-1} \\ &= \text{tr} \left(\mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \right)^{-1} \\ &= \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \\ &= \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z} \end{aligned} \quad (46)$$

Naming $\mathbf{Q} = \mathbf{Z}^T \mathbf{M}^{-1} \mathbf{Z}$, we got

$$\frac{\partial}{\partial \hat{\mathbf{a}}_{u,v}} \frac{\partial \text{KL}(q(f) \parallel p(f|y))}{\partial \hat{\mathbf{a}}_{i,j}} = \frac{1}{2} \mathbf{Q}_{v,j} \mathbf{Q}_{i,u} = \frac{1}{2} \mathbf{Q}_{j,v} \mathbf{Q}_{i,u} \quad (47)$$

where in the last equality we used that \mathbf{Q} is symmetric. Using the definition of \mathbf{M} we know that

$$\mathbf{M}^{-1} = \mathbf{I} - \mathbf{Z} (\mathbf{A} + \mathbf{K})^{-1} \mathbf{Z}^T \Rightarrow \mathbf{Q} = \mathbf{K} - \mathbf{K} (\mathbf{A} + \mathbf{K})^{-1} \mathbf{K} \quad (48)$$

This shows that \mathbf{Q} is the posterior covariance of a GP with prior covariances \mathbf{K} and noise covariances \mathbf{A} .

$$\frac{\partial}{\partial \hat{\mathbf{A}}} \frac{\partial \text{KL}(q(f) \parallel p(f|y))}{\partial \hat{\mathbf{A}}} = \frac{1}{2} \mathbf{Q} \otimes \mathbf{Q} \quad (49)$$

with \otimes denoting the Kronecker product and \mathbf{Q} a definite positive matrix. As the Kronecker product of two definite positive matrices is definite positive, the optimal is a minimum of the KL. Given that this second derivative is positive-definite, the value found is the global minimum of the ELBO objective. ■

Appendix B. Related Work

LA for DNNs was originally introduced by Mackay (1992), applying it to small networks using the full Hessian. MacKay (1992) also proposed an approximation similar to the generalized Gauss-Newton (GGN). The combination of scalable factorizations or diagonal Hessian approximations (Martens and Grosse, 2015; Botev et al., 2017) with the GGN approximation (Martens, 2020) played a crucial role in the resurgence of LA for modern DNNs (Ritter et al., 2018; Khan et al., 2019). Recent works aim to relax the Gaussian assumption of LLA adopting a Riemannian-Laplace approximation, where samples naturally fall into weight regions with low negative log-posterior (Bergamin et al., 2023).

To address the underfitting issue associated with LA (Lawrence, 2001), particularly when combined with the GGN approximation, Ritter et al. (2018) proposed a Kronecker factored (KFAC) LLA approximation. This approach outperforms LA with a diagonal Hessian matrix.

The GP interpretation of LLA (Khan et al., 2019) allows using GP approximate methods to speed up the computations. Immer et al. (2021) propose to use a subset of the training dataset as a scalable alternative to the true GP. Lee et al. (2022) propose a Mixture of Experts approach where each expert is trained on a different soft-margin cluster. However, the proposed clustering algorithm, although more efficient than Kernel-K-means, has linear cost w.r.t. the training set size. VaLLA, on the other hand, has sub-linear training time w.r.t. training set size due to mini-batch training. Moreover, it is not clear how to consider neighboring clusters in high dimensional input spaces. The authors only provide code for a 1-dimensional problem. Third, fitting a local GP using the data of the corresponding cluster and its neighbors is expected to overestimate the predictive variance since the model has been trained with a smaller number of training instances. This is particularly the case in datasets with millions of training instances such as Taxi. This problem is also described by Immer et al. (2021).

Deng et al. (2022) proposed a Nyström approximation of the true GP covariance matrix by using $M \ll N$ points chosen at random from the training set. The method, called ELLA, has cost $O(NM^3)$. ELLA also requires computing the costly Jacobian vectors required in VaLLA, but does not need their gradients. Unlike VaLLA, the Nyström approximation needs to visit each instance in the training set. However, as stated by Deng et al. (2022), ELLA suffers from over-fitting. An early-stopping strategy, using a validation set, is proposed to alleviate it. In this case, ELLA only considers a subset of the training data. ELLA does not allow for hyper-parameter optimization, unlike VaLLA. The prior variance $\frac{2}{0}$ must be tuned using grid search and a validation set, which increases training time significantly.

The recent work of Scannell et al. (2024) proposes a similar approach to VaLLA, where an inducing point sparse approach is used to construct a GP from a pre-trained DNN. However, two main points differentiate this work from our approach: (i) the pre-trained DNN is not kept as the posterior mean of the model, potentially losing prediction performance and also departing from LLA’s post-hoc nature and goal; (ii) instead of using mini-batches to optimize variational parameters, they perform a full iteration over the training data to find optimal variational parameters. Thus, this results in a potentially slower method than VaLLA, which due to early-stopping and stochastic optimization, can avoid iterating over the full dataset.

Samples from a GP posterior can be efficiently computed using stochastic optimization, eluding the explicit inversion of the kernel matrix (Lin et al., 2024). This approach can be extended to LLA to generate samples from the GP posterior, avoiding the $O(N^3)$ cost (Antorán et al., 2023). However, this method cannot provide an estimate of the log-marginal likelihood for hyper-parameter optimization. To address this limitation, Antorán et al. (2023) propose using the *EM-algorithm*, where samples are generated (E-step) and hyper-parameters are optimized afterwards (M-step) iteratively. The EM algorithm significantly increases computational cost, as generating a single sample is as expensive as training the original DNN on the full data. Finally, the method of Antorán et al. (2023) only considers classification problems.

Appendix C. MAP solution in Hilbert Space

Proposition 2 assumes that $g(\cdot; \hat{\cdot}) \in H$. In practice, this need not be the case. Covariance functions such as squared exponential or Matérn are recognized for encompassing the entire space of continuous functions. However, whether $g(\cdot; \hat{\cdot}) \in H$ holds in general remains unknown. We assume that if $g(\cdot; \hat{\cdot}) \notin H$, then H is sufficiently expressive to include a close approximation to $g(\cdot; \hat{\cdot})$. Consequently, $g(\cdot; \hat{\cdot})$ can be used as the sparse GP posterior mean.

Whether the map solution is in the Hilbert space might be difficult (if not impossible) to know. However, there are cases where it can be theoretically shown; for example in a linear model. Let the map solution be a linear model as

$$g(\mathbf{x}; (\mathbf{w}; b)) = \mathbf{w}^T \mathbf{x} + b;$$

Then, the features (Jacobians) are

$$(\mathbf{x}) = (\mathbf{x}^T; 1)^T;$$

With a single inducing point \mathbf{z} and scalar value $a \in \mathbb{R}$, the mean function would be

$$m(\mathbf{x}) = a (\mathbf{x})^T (\mathbf{z}) = a(\mathbf{x}^T \mathbf{z} + 1);$$

This recovers the MAP solution if $a = b$ and $\mathbf{z} = \mathbf{w} = a$. However, if the model is not linear but has a linear last layer as

$$g(\mathbf{x}; (\mathbf{w}; b; \cdot)) = \mathbf{w}^T h(\mathbf{x}) + b;$$

where h is a non-linear function that depends on parameters \cdot . Then, the features (Jacobians) are

$$(\mathbf{x}) = (h(\mathbf{x})^T; 1; (r h(\mathbf{x}))^T \mathbf{w}^T);$$

Here it might be difficult to check if there exists a combination that yields the map solution as the mean function.




Figure 2: Predictive distribution (mean and two times the standard deviation) on a toy 1D regression dataset with a 2 hidden layer MLP with 50 units. Training points are shown in black and the validation set is shown in orange. The first column shows the obtained predictive distribution using early-stopping and a validation set, with 5 and 10 inducing points. The second column shows the results obtained without early-stopping. LLA and ELLA are shown in the last column.

Appendix D. Over-fitting and Early Stopping

In Section 3 we mentioned the fact that the standard maximization of the ELBO leads to null predictive variances. The optimal value for the prior variance is infinite as a result of the mean being fixed to the optimal MAP solution. In a regression scenario with Gaussian noise with variance σ^2 the first term in the r.h.s. of (5) becomes:

$$\prod_{i=1}^N \frac{\log(2\sigma^2)}{2} - \frac{(y_i - g(\mathbf{x}_i; \hat{\theta}))^2}{2\sigma^2} - \frac{K^2(\mathbf{x}_i; \mathbf{x}_i)}{2\sigma^2} \quad (50)$$

where $(y_i - g(\mathbf{x}_i; \hat{\theta}))^2$ is constant. Maximizing (50) w.r.t. σ^2 results in the prior covariances, $\frac{1}{2} J_{\theta}^{\top}(\mathbf{x}) J_{\theta}(\mathbf{x}^{\theta})$, tending to 0. This makes posterior covariances $K^2(\mathbf{x}_i; \mathbf{x}_i)$ also tend to 0, effectively cancelling the last term in (50). The KL term in (5) is also optimal and 0 if $\sigma^2 \rightarrow 0$.

We circumvent this by applying β -divergences, which are not ill-defined in this learning setup; allowing the optimization of the prior. However, the use of this optimization objective is not perfect and we faced the fact that it tends to over-fit the prior variance to the training data. The middle column of Figure 2 (middle) shows the obtained predictive distribution (two times standard deviation) learned from VaLLA using the black points as training data. The MAP solution is obtained using a 2 hidden layer MLP with 50 hidden units and *tanh* activation, optimized to minimize the RMSE of the training data for 10000 iterations with Adam and learning rate 10^{-3} . VaLLA on the other hand is trained for 20000 iterations. As one may see in the image, the prior variance is fitted to the data to the point where the uncertainty does not increase in the middle gap of the data.

In this experiment, VaLLA optimizes hyper-parameters along with the variational objective. LLA optimizes the prior variance and likelihood variance by maximizing the marginal




Figure 3: Predictive distribution (two times the standard deviation) on a toy 1D regression dataset with a 2 hidden layer MLP with 50 units. The obtained results for 5, 10 and 20 inducing points for VaLLA are shown in the first row. ELLA’s predictive distribution with the same amount of samples from the training data is shown in the second row. LLA’s predictive distribution is shown in dotted orange.

log likelihood and ELLA uses LLA’s optimal hyper-parameters. In this situation there are two simple courses of action: use the ELBO and choose the prior variance by cross-validation; or perform early-stopping with the -divergences objective, using a validation set to stop training before over-fitting. The latter may not work as it assumes that there is a point during training where the prior variance truly explains the underlying data without over-fitting. However, as the prior variance is set to a relatively large value compared to the optimal one (which is small and leads to over-fitting), this method resulted in great performance for VaLLA. The left column of Figure 2 shows the obtained predictive distribution (two times standard deviation) learned from VaLLA, in this case, using the black points as training data and the orange points as the validation set. In the experiments, we computed the NLL of the validation set every 100 training iterations and stopped training when it worsens. This also allowed us to save computational time.

Appendix E. Increasing Inducing Points

Using the optimal covariance in Proposition 3, if the set of inducing points equals the training points $\mathbf{Z} = \mathbf{X}$, the posterior distribution of VaLLA equals that of the exact LLA Gaussian Process. This suggest that increasing the number of inducing points would lead to better uncertainty estimations. In this section, we aim to show how close is the predictive distribution of VaLLA to that of LLA when we increase the number of inducing points.

Figure 3 shows the obtained predictive distribution of VaLLA (first row) and ELLA (second row) for $M = 5$, $M = 10$ and $M = 20$ inducing points/samples. The initial and final locations of the inducing points are also shown for VaLLA. The posterior distribution obtained by LLA is shown in dotted orange. The MAP solution is obtained using a 2 hidden layer MLP with 50 hidden units and \tanh activation, optimized to minimize the RMSE of the training data for 12000 iterations with Adam and learning rate 10^{-3} . VaLLA on the

Algorithm 1: VaLLA’s training loop with $\epsilon = 1$

Require: Pre-trained MAP solution f , input batch $\mathbf{X}; \mathbf{y}$, m number of inducing points and T iterations.

```

Z  $kmeans(\mathbf{X}; m)$  {Initialize inducing points}
L  $\mathbf{I}$  {Initialize Cholesky decomposition of  $\mathbf{A}$ }
for  $i \in \{0; \dots; T - 1\}$  do
     $(\mathbf{X}_b; \mathbf{y}_b) \leftarrow get\_batch()$  {Get mini-batch of data}

     $\mathbf{J}_x \leftarrow compute\_jacobian(\mathbf{X}_b)$  {Compute Jacobians}
     $\mathbf{J}_z \leftarrow compute\_jacobian(\mathbf{Z})$ 

     $\mathbf{K}_x \leftarrow \int_0^1 \mathbf{J}_x \mathbf{J}_x^T$  {Compute Kernels}
     $\mathbf{K}_{xz} \leftarrow \int_0^1 \mathbf{J}_x \mathbf{J}_z^T$ 
     $\mathbf{K}_z \leftarrow \int_0^1 \mathbf{J}_z \mathbf{J}_z^T$ 
     $\mathbf{A} \leftarrow \mathbf{L}\mathbf{L}^T$  {Compute Variational Matrix}

     $Q\_mean \leftarrow f(\mathbf{X}_b)$  {Compute posterior mean}
     $Q\_var \leftarrow \mathbf{K}_x - \mathbf{K}_{xz}(\mathbf{A}^{-1} + \mathbf{K}_z)^{-1} \mathbf{K}_{xz}^T$  {Compute posterior covariance matrix}

     $KL \leftarrow compute\_KL(\mathbf{A}; \mathbf{K}_z)$  {Compute Kullback-Leibler divergence}
     $NLL \leftarrow compute\_NLL(\mathbf{y}_n; Q\_mean; Q\_var)$  {Compute Negative Log-likelihood}
     $loss \leftarrow \frac{len(\mathbf{X})}{len(\mathbf{X}_b)} NLL + KL$ 
    Optimize parameters by minimizing  $loss$ .
end for
    
```

other hand is trained for 30000 iterations. For this experiment, VaLLA and ELLA use the optimal prior variance and likelihood variance obtained by optimizing LLA’s marginal log likelihood. As one may see in the image, it is clear that one of the main differences between the two methods is that VaLLA tends to over-estimate the variance whereas ELLA tends to infra-estimate it, compared to LLA. Furthermore, the value of M for which the model is closer to the LLA posterior is lower for VaLLA than for ELLA. As we increase M , VaLLA’s predictive distribution becomes closer and closer to that of LLA.

The initial and final position of the inducing locations is also shown in the figure. For this experiments, the initial values are computed using K-Means. It can be seen how VaLLA is capable of tuning the inducing locations and move them from one cluster of points to another as needed. This is one of the main advantages of this method compared to ELLA.

Appendix F. Pseudocode

Algorithm 1 shows the structure of VaLLA’s training loop, where no Early-Stopping is considered. Using the kernels and \mathbf{A} it is easy to compute the KL in Eq. 5. As a result, the training loop is easy to implement, since $q(f)$, given by $k \in \mathcal{C}^{\wedge}$ and $k \in \mathcal{f}^{-q}$ are easily computable as detailed in the algorithm.

Appendix G. Further Experiments

We compare VaLLA with other methods using the LLA implementation by [Daxberger et al. \(2021a\)](#). VaLLA utilizes a batch size of 100. For regression, MNIST and FMNIST problems, we train our own DNN (standard multi-layer perceptron), which is stored for reproducibility. In the CIFAR10 experiments with ResNet, we extract the results for other methods from [Deng et al. \(2022\)](#) and use the same DNN for VaLLA. Hyper-parameters in all LLA variants (diagonal, KFAC, last-layer LLA) are optimized based on the marginal log-likelihood estimate. Additional experimental details are given in Appendix [H](#).

G.1. Synthetic Regression

We compare the predictive distribution of VaLLA with that of LLA (which is considered the optimal method), other LLA variants and ELLA, on the 1-D regression problem of [Izmailov et al. \(2020\)](#). In ELLA and VaLLA, we use the optimal hyper-parameters from LLA. The results in Figure [1](#) illustrate that VaLLA’s predictive distribution closely aligns with that of LLA. Figure [3](#) (see Appendix [E](#)) depicts the predictive distributions of VaLLA and ELLA for varying numbers of inducing points and points in the Nyström approximation, respectively. It shows that VaLLA converges to the true posterior faster than ELLA, with VaLLA tending to over-estimate the predictive variance while ELLA under-estimates it. In Figure [2](#) (see Appendix [E](#)) we observe the effect of tuning the prior variance in VaLLA in another toy 1-D problem, with and without early-stopping. Notably, early stopping, using a validation set, prevents overly small predictive variances in VaLLA. Finally, we observe that when VaLLA estimates the prior variance by maximizing the χ^2 -divergence, it tends to underestimate LLA’s predictive variance.

G.2. Airline, Year and Taxi Regression Problems

We carry out experiments on large regression datasets. (i) The *Year* dataset (UCI) with 515;345 instances and 90 features. We use the original train/test splits. (ii) The *US flight delay (Airline)* dataset ([Dutordoir et al., 2020](#)). Following [Ortega et al. \(2023\)](#) we use the first 700;000 instances for training and the next 100;000 for testing. 8 features are considered: month, day of month, day of week, plane age, air time, distance, arrival time and departure time. (iii) The *Taxi dataset*, with data recorded on January, 2023 ([Salimbeni and Deisenroth, 2017](#)). 9 attributes are considered: time of day, day of week, day of month, month, PULocationID, DOLocationID, distance and duration. We filter trips shorter than 10 seconds and larger than 5 hours, resulting in 3 million instances. The first 80% is used as train data, the next 10% as validation data, and the last (10%) as test data. In all experiments, a 3-layer DNN with 200 units, *tanh* activations and L2 regularization is considered. VaLLA and ELLA use 100 inducing points and 100 random points, respectively.

The table displayed on the l.h.s. of Figure [4](#) presents the averaged results over 5 random seeds. LLA is not considered here due to intractability. Negative log likelihood (NLL), continuous ranked probability score (CRPS) ([Gneiting and Raftery, 2007](#)) and a centered quantile metric (CQM), described below, are reported. We observe that VaLLA performs best according to NLL and CQM, while it gives worse results in terms of CRPS compared to the other methods.

Model	Airline			Year			Taxi		
	NLL	CRPS	CQM	NLL	CRPS	CQM	NLL	CRPS	CQM
MAP	5.087	18.436	0.158	3.674	5.056	0.164	3.763	3.753	0.227
LLA Diag	5.096	18.317	0.144	3.650	4.957	0.122	3.714	3.979	0.270
LLA KFAC	5.097	18.317	0.144	3.650	4.955	0.121	3.705	3.977	0.270
LLA	5.097	18.319	0.144	3.650	4.954	0.120	3.718	3.975	0.270
LLA KFAC	5.097	18.317	0.144	3.650	4.954	0.120	3.718	3.976	0.270
ELLA	5.086	18.437	0.158	3.674	5.056	0.164	3.753	3.754	0.227
VaLLA 100	4.923	18.610	0.109	3.527	5.071	0.084	3.287	3.968	0.188
VaLLA 200	4.918	18.615	0.107	3.493	5.026	0.076	3.280	3.993	0.188

Figure 4: (left) Results on regression datasets. (right) Illustration of CQM on Taxi. Average results across 5 different random seeds (standard deviations always $< 10^{-4}$ and omitted). Best value highlighted in **purple** and second to best in **teal**. for Last Layer LLA.

Centered Quantile Metric (CQM). In regression, CQM assesses the calibration of the predictions, extending the Expected Calibration Error (ECE) to regression problems with Gaussian predictions *with the same mean* but different variance. CQM calculates the *centered interval around the mean* with probability mass $\mathcal{Z}(0; 1)$. For Gaussian predictions $N(\mu(\mathbf{x}); \sigma^2(\mathbf{x}))$, the open interval is defined as $I(\mathbf{x}; \mu) = (\mu - \sigma \Phi^{-1}(\frac{1-\alpha}{2}); \mu + \sigma \Phi^{-1}(\frac{1-\alpha}{2}))$, where $\Phi^{-1}(\cdot) = \Phi^{-1}(\cdot; \mu; \sigma^2)$ with $\Phi(\cdot; \mu; \sigma^2)$ the CDF of a Gaussian with mean $\mu(\mathbf{x})$ and variance $\sigma^2(\mathbf{x})$. The fraction of test points falling inside the interval is then computed. If the predictive distribution is well calibrated, $\int_0^1 \mathbb{1}_{y \in I(\mathbf{x}; \mu)} dP(\mathbf{x}, y) = \alpha$. Formally,

$$\text{CQM} = \int_0^1 \mathbb{P}_{(\mathbf{x}, y)} [y \notin I(\mathbf{x}; \mu)] d\alpha \quad (51)$$

All methods utilize the pre-trained DNN solution as predictive mean. Thus, the differences in $I(\mathbf{x}; \mu)$ stem from the predictive variance. Evaluating the integrand in (51) on a grid of values allows us to visually interpret the uncertainty estimation of each method. The r.h.s. of Figure 4 shows $\mathbb{P}_{(\mathbf{x}, y)} [y \notin I(\mathbf{x}; \mu)]$ for several models on the Taxi dataset. In general, all methods tend to over-estimate the actual predictive variance, as evidenced by the values above the diagonal. The l.h.s. of Figure 4 shows CQM estimated using trapezoid integration with 11 points. We refer to Appendix I for more details on the CQM metric.

G.3. Image Classification Problems

MNIST and FMNIST. We employ a 2-layer fully connected DNN with 200 units in each layer and *tanh* activations. In VaLLA we considered 100 and 200 inducing points, while in ELLA, 2000 random points are used. The Out-of-distribution (OOD) detection ability of each method is evaluated using the entropy of the predictive distribution as a score. We compute the area under the ROC curve (AUC) of the binary problem that distinguishes between instances from pairs of datasets MNIST/FMNIST and FMNIST/MNIST (Immer et al., 2021). Moreover, in FMNIST we also assess the robustness of the predictive distribution by rotating the test images up to 180 degrees and computing the ECE and NLL on rotated images (Ovadia et al., 2019).

Model	ACC	NLL	ECE	BRIER	OOD-AUC
MAP	97.6	0.076	0.008	0.036	0.905
LLA Diag	97.4	0.143	0.072	0.053	0.922
LLA KFAC	97.5	0.094	0.029	0.041	0.949
LLA	97.6	0.081	0.015	0.037	0.909
LLA KFAC	97.6	0.081	0.015	0.037	0.909
ELLA	97.6	0.076	0.008	0.036	0.905
Sampled LLA	97.6	0.087	0.026	0.040	0.954
VaLLA 100	97.7	0.076	0.010	0.036	0.916
VaLLA 200	97.7	0.075	0.010	0.035	0.921

Figure 5: (left) MNIST experiments. Results averaged over 5 different random seeds (standard deviations $< 10^{-4}$ in all cases and omitted). (right) Box-plots of training times in seconds. ELLA considers 10 prior values chosen using a validation set. Sampled-LLA uses 8 EM steps and 32 samples. Best value is highlighted in **purple** and second to best in **teal**. for Last Layer LLA.

Model	ACC	NLL	ECE	BRIER	OOD-AUC
MAP	86.6	0.373	0.009	0.193	0.874
LLA Diag	86.2	0.397	0.043	0.201	0.914
LLA KFAC	86.5	0.377	0.014	0.194	0.932
LLA	86.6	0.373	0.008	0.193	0.882
LLA KFAC	86.6	0.373	0.008	0.193	0.880
ELLA	86.6	0.373	0.008	0.193	0.874
VaLLA 100	87.4	0.335	0.011	0.182	0.923
VaLLA 200	87.6	0.332	0.013	0.181	0.933

Figure 6: (left) Results on FMNIST. Results are averaged over 5 different random seeds (standard deviations are lower than 10^{-4} and omitted). Best value is highlighted in **purple** and second to best in **teal**. for Last Layer LLA. (right) ECE and NLL for rotated FMNIST.

The left table in Figure 5 shows the results on MNIST. VaLLA gives better uncertainty estimates in terms of NLL and the Brier score, but performs less effectively in terms of ECE. Remarkably, VaLLA improves prediction accuracy (ACC) due to the approximation of Daxberger et al. (2021b) to compute class probabilities in multi-class problems. In terms of OOD-AUC VaLLA outperforms the MAP solution but lags behind other methods *s.a.* Sampled-LLA or LLA with Kronecker approximations. Figure 5 (right) illustrates the training times for each method, with VaLLA being faster than ELLA, Sampled-LLA or Last-Layer LLA.

Finally, the left table in Figure 6 displays the results on FMNIST. Here, VaLLA excels in prediction accuracy and provides the best uncertainty estimates in terms of NLL and the Brier score. Although it does not perform as well in terms of ECE, the differences are small. VaLLA also achieves the best results in OOD-AUC. Figure 6 (right) shows VaLLA holds better performance in terms of ECE and NLL as the test images’ corruption increases (rotation level), indicating the greater robustness of VaLLA’s predictive distribution.

CIFAR10 and ResNet. Among with Table 1 that shows ACC, NLL and ECE in the ResNet-Cifar10 experiment, Figure 7 shows the NLL on the perturbed test set with 5

o `B iBQM H GBM2 `Bx2/ G TH +2 TT`QtBK iBQM

6B;m`2_0p mHib QM +Q``mTi2/ *A6 _Ry rBi? _2bL2i64x KTK2H2/ GG m
M/ 1GG nM=2000 MK =20 X

BM+`2 bBM; 192p2H2 QQ``mT2BQM b- UyxkX 1 +? #Qt@THQi bmKK `Bx2
i?2 i2bi LGG 7Q` 2 +? BMI2Mb Bv+ G2p2FiB-Q bX H+2 `2bmHib ; BM ?B
o GG ob `Q#mbi T`2/B+iBp2 /Bbi`B#miBQM- +?B2pBM; HbQ HQR2` L
K2i?Q/bX

TT2M/Bt >X 1tT2`BK2Mi H .2i BHb

aQm`+2 +Q/2 7Q` i?2 +QM/m+i2/ 2tT2`BK2Mib + M #2 ++2bb2/ BM i?
`2TQbB?QiTv,ff;Bi?m#X+QKfGm/pBMbfo `B iBQX H@GG fi`22fK BM

>XRX J S bQHmiBQMb

6Q``2;`2bbBQM T`Q#H2Kb Uu2 `- B`HBM2 M/ h tB / i b2ibV- j@H
r b mb2/ r20?mMBib BM 2 +? H v2`X h?2 QTIBK H r2B;?ib `2 Q#i BM2
i?2 _Ja1 m2000;Bi2` iBQMb Q710i+M/bBx2QTIBKBB2f;K M/ "
kyR8rBi? H2 `MBV; `Mi/2r2B;?i /20+2X

6Q` JLAah M/ 6JLAah 2tT2`BK2Mib- k@H v2` 7mHHv200QMM2+i2/ LI
mMBib BM 2 +? H v2`X h?2 QTIBK H r2B;?ib `2 Q#i BM2000#v KBMBKBx
Bi2` iBQMb Q710i+M/ bBx2QTIBKBB2f;K M/'yR8rBi? H2 `MBM;` i2
10³ M/ r2B;?i /20+3X

>XkX G TH +2 GB#` `v

h?2 G TH +2 Hbt#`2`y2U 2+kHkX r b mb2/ iQ T2`7Q`K H bi@H v2`- E6 *
/B ;QM H TT`QtBK iBQMb Q7 i?2 GG K2i?Q/ M/ QTIBKBx2 i?2 T`BQ`
h?2 H ii2` Bb /QM2 #v QTIBKBxBM; i?2 HQ; K `;BM H HBF2HB?QQ/ Q7
HQ;nK `;BM HnHBK2H?Q/Q7000 Bi2` iBQMb rBi? i?2 / K QTIBKBx2` I
H2 `MBM0³X2

>XjX 1{+B2Mi E2`M2H *QKTmi iBQM 7Q` JGS

AM i?Bb b2+iBQM r2 /Bb+mbb M 2{+B2Mi BKTH2K2Mi iBQM 7Q` +QKT
E2`M2Hx0X 6B`bi Q7 HH- i F2 BMiQ ++QmMi i? ii?2 +QKTmi iBQM Q

P`i2; a Mi M >2`M{M/2x@GQ# iQ

`2/m+2/ iQ bmKK iBQM QM i?2 MmK#2` Q7 T ` K2i2`b Q7 i?2 KQ/2H,

$$(x; x^0) = \int_0^2 J_{\wedge}(x)^T J_{\wedge}(x^0) = \int_0^2 \frac{\partial}{\partial s} g(x; \wedge) \frac{\partial}{\partial s} g(x^0; \wedge); \quad U8kV$$

PM2 Q7 i?2 HBKBi iBQM b Q7 +QKTmi BM; j /BK2MbBQM H i2MbQ` Q7 U# i+? bBx2- MmK#2` Q7 +H bb2b- MmK#2` i?2 F2`M2H b bmK HHQrb iQ bBKTHB7v i?2`2[mB`2/+QKTmi iBQM b ? p2 iQ biQ`2 BM K2KQ`v i?2 C +Q#B MbVX *QM bB/2` MQR JGS b

$$g(x; \wedge) = h_L a_{H_L} \quad a_{h_1}(x); \quad U8jV$$

r?2`2 2 +? 7maMB+iBQM M@HBM2` +iBp iBQM 7mM +iBQM BM2` +? 7m 7mM+iBQM Q7 i?2 7Q`K

$$h_1(x) = W_1^T x + b_1; \quad U89V$$

qBi? i?2 b b bmTTQb2/ iQ #2 7mHHv @+QM M2H i22`M X m1` H?M27riQ 2F Q7 T`iB H/2`Bp iBp2b Q7 i?2 M2m` H M2irQ`F`2

$$\frac{\partial}{\partial W_{j,i}} g(x; \wedge) \quad M / \frac{\partial}{\partial b_j} g(x; \wedge) \quad 8l = 1; \dots; L; \quad U88V$$

M/ i?2 F2`M2H Bb +QKTmi2/ bBKTHv #v //BM; i?2 T`QVr Bbi Q7 i?2 b2 / bm# @BM/2t /2MQ@BM; bNKv2B KBjHB`h v b m# @BM/2t /2MQiBM; 2 +? +QKT Q7 i?2 #B b p2+iQ`T`K2Q`bBK Bv2`Hv- 2 +? QmiTmi Q7 i? iH v2`X

AM 7 +i- mbBM; i?2 bi`m+im`2 Q7 i?2 KQ/2H M/ i?2 +o`h BM`mH2- i? QmiTmi Q7 i?2 M2irQj`F; i`hXr`XBx?ii?2`K2i2`i`Q`H i?2`Bb,

$$\frac{\partial}{\partial W_{j,i}} g_0(x; \wedge) = \frac{\partial}{\partial h} g_0(x; \wedge)^T \frac{\partial}{\partial W_{j,i}} h_1; \quad U8eV$$

r?2`2 2 +? Q7 i?2 irQ p2+iQ`b BM i?2`X?XbX ? b H2M; i? 2[m H iQ i?2 H v2X AM 7 +i

$$\frac{\partial}{\partial W_{j,i}} h_1 = 1_i \quad a(h_{1-1})_i; \quad U8dV$$

r?2`2 a(h_{1-1})_i +Q``2bTQM/b iQ i?2 BMH m2b Q7Q?`Q`p2`^)+ M HbQ #2 +QKTmi2/ mbBM; i?2 +? BM`mH2,

$$\frac{\partial}{\partial h} g_0(x; \wedge) = \frac{\partial}{\partial h_{k+1}} g_0(x; \wedge) \frac{\partial}{\partial h} h_{k+1} = \frac{\partial}{\partial h_{k+1}} g_0(x; \wedge) W_1^T / B(a^0(h_1)); \quad U83V$$

r?B+? + M #2 2 bBHv +QKTmi2/ #v # +F@T`QT ; iBM; i?2 /2`Bp iBp2b + M #2 2 bBHv /QM2 7Q` i?2 #B X2b Q722 m`H i?2`/2`Bp iBp2b QMHv /2 QM # +F@T`QT ; Bm(x; i)27Q`2 +? H v2`- i?2 p Hm2 QW; i; 2 T`K2i2`b M/ i?2 T`QT ; i2/ QmiTmi h_{i-2}; h_{i-1} H2p2Hm i2/ ii?2 MQM@HBM2` +iBp a() M/ Bib /2`Baf(i)X ph2`Bb K2 Mb i? i- B7 r2 biQ`2 i?2 BMi2`K2/B i2 Qmi H v2; U; ; h_L 1 V QM i?2 7Q`r`/T bb Q7 i?2 KQ/2H- #v mbBM; bBM; H2 # + M +QKT`h_{i-2}; ^) 7Q`2 +? H v2`X

u2 `

B`HBM2

h tB

6B;m`2A3H,Hmbi` iBQM Q7 *ZJ- 7Q` 2 +? K2i?Q/- QM i?2 `2;`2bbBQM /
 B`HBM2 UKB//H2V M/ h tB U`B;?iVX h?2 J S bQHmiBQM Bb ;Bp
 M2irQ`F rBi? j ?B//2 200hm2 Bb bQ 7M/?+iBp iBQM bX

*`BiB+ HHv- ;B`g_o(x; ^) --+?2 + M // i?2 +QM i`B#miBQM Q7 2 +? H v2`
 F2`M2H- 6v`KMA;Mji?Bb T`Q+2bb- r2 + M bT2/@mT i?2 +QKTmi iBQM b #
 i?2 /2`Bp iBp2bX 6Q` 2v`K2Q#b`BM 2ji? ii?2 /2`Bp iBp2? b bBKTH2 7Q
 Bb p2+iQ` Q7 QM2b iBK2b b+ H `X 6m`i?2`KQj2-i?22Q2mBTmMQ /2T2
 mMBi +Q` `2bTQM/BM W_{i,j} X? 2? 2`E;?Q` 2- 7Q` i xQ M M⁰ b ii M2+2 2`M2H
 +QM i`B#miBQM UB;MQ`BM; i?2 T`BQ` p `B M+2 T `o KM b2 Bb ,+Q` `2bTQM

$$\begin{aligned} \frac{\partial}{\partial W_{j,i}} g_o(x; \wedge) \frac{\partial}{\partial W_{j,i}} g_o^0(x^0; \wedge) &= \frac{\partial}{\partial h} g_o(x; \wedge)^T \frac{\partial}{\partial W_{j,i}} h_l \frac{\partial}{\partial h} g_o(x^0; \wedge)^T \frac{\partial}{\partial W_{j,i}} h_l \\ &= \frac{\partial}{\partial h} g_o(x; \wedge)^T \frac{\partial}{\partial W_{j,i}} h_l \frac{\partial}{\partial W_{j,i}} h_l^T \frac{\partial}{\partial h} g_o^0(x^0; \wedge) \\ &= \frac{\partial}{\partial h} g_o(x; \wedge)^T \mathbf{1}_l \mathbf{a}(h_l \mathbf{1}(x))_i \mathbf{a}(h_l \mathbf{1}(x^0))_i \mathbf{1}_l^T \frac{\partial}{\partial h} g_o^0(x^0; \wedge) \\ &= s_{o,x}^l \mathbf{a}(h_l \mathbf{1}(x))_i \mathbf{a}(h_l \mathbf{1}(x^0))_i s_{o^0,x^0}^l; \\ &= s_{o,x}^l s_{o^0,x^0}^l \mathbf{a}(h_l \mathbf{1}(x))_i \mathbf{a}(h_l \mathbf{1}(x^0))_i; \end{aligned} \quad \text{U8NV}$$

rBi s_{o,x}^l = $\frac{\partial}{\partial h} g_o(x; \wedge)^T \mathbf{1}_l$ b+ H `X aBKBH ` bBKTHB}+ iBQM2K-Q++m` BM i?
 +QM pQHmiBQM H H v2`X
 amKKBM; mT- #v mbBM; i?Bb K2i?Q/- HH i?2 `2[mB`2/ F2`M2H K i`
 2{+B2MiHv +QKTmi2/- 7Q` KBMB@# i+? Q7 / i TQBMib M/ b2i Q7
 bBKBH ` +Qbi b i? i Q7 H2iiBM; i?2 KBMB@# i+? Q` i?2 BM/m+BM; TQ
 /Bb /p Mi ; 2 Bb- ?Qr2p2`- i? ii?2 /2b+`B#2/ +QKTmi iBQM b rBHH ? p2
 7Q` 2 +? /Bz2`2Mi .LL `+?Bi2+im`2X h?Bb #2+QK2b i2/BQmb BM i?2
 rBi? +QKTHB+ i2/ H v2`b- b /2 pxx B#2/ BM a2+iBQM

TT2M/Bt AX 6m`i?2` M HvbBb Q7 i?2 Zm MiBH2 J2i`B+
 b bi i2/ BM a2+iBQM TQb2/ M2r K2i`B+ 7Q` `2;`2bbBQM T`Q#H2Kb
 2ti2M/b 1*1 iQ` 2;`2bbBQM T`Q#H2Kb rBi? : mbb BBM? Ti?2/ B+K2p2 /Bbi`

P`i2; a Mi M >2`M{M/2x@GQ# iQ

K2 M h?Bb FBM/ Q7 K2i`B+ Bb /2bB` #H2 7Q` GG K2i?Q/b b HH Q7
i?2 QTIBK H J S bQHmiBQM b i?2 T`2/B+iBp2 K2 M Q7 i?2 KQ/2HX h
T`2/B+iBp2 p `B M+2X 6Q`K HHv- *ZJ(0;0)K7 2nT2 Q7 Q# BHB?vi? iTQB Mib
7 HH B MiQ i?2 T`2/B+iBp2 +2Mi2`2/ BM?2` p M/Q7HTV`QM; # B H B i M B M; Bb
B7 i?2 KQ/2H 2tTH BMb i?2 / 100% B H 2?2 T Q; B Mib r B H H 7 100% B M b B / 2 i?2
+2Mi2`2/ [m MiBH2 B Mi2`p HX h?mb- i?2 K2i`B+ /2}M2/ b

$$*ZJ = \int_0^{Z_1} P(x^?, y^?) [y^? 2 I(x^?;)] d; \quad UeyV$$

b?QmH/ #2 `Qm?2`Hv?2 KQ/2H T`2/B+iBp2 /Bbi`B#miBQM Bb bBKBH `
;Bp2M #v i?2 Q#b2`p2// i X
6B;m3?Qrb i?2 2pQ H(m;B)Q`2 Q7;)) rX`XiXQ` i?2 #2bi T2`7Q`KBM;
KQ/2Hb BM i?2 `2; `2bbBQM T`Q#H2KbX *ZJ+Q` `2bTQM/b iQ i?2 `2 #
My = x U#H +F HBM2VX h?Bb };m`2 HHQrb iQ `;m2 i? i UBM ;2M2` H
2biBK iBM; i?2 T`2/B+iBp2 p `B M+2 b i?2v `2; BpBM; p Hm2b #Qp2 i
bT2+B}+ p 2(0;2)Q7?2 `2TQ`i2/ T`Q# #BHBiB2bK`22 MB;M?i?i? -QM
p2` ;2- i?2`2 `2 KQ`2(x; Q B M i M B 2 v b?QmH/X h? i Bb- i?2 T`2/B+i2/ B
Bb H `;2` i? M Bi b?QmH/- r?B+? + M QMHv K2 M i? i i?2 p `B M+2 Bb
;2QK2i`B+ H T2`bT2+iBp2- Bi Bb +H2 `i? i *OZ JMB H Q r 2 v 0;3;M i2` i? M
BM/2T2M/2MiHv Q7 i?2 KQ/2H M// i b2i mb2/X
AM 7 +i- i?Bb };m`2 HHQriQ pBbm HHv bim/v i?2 H2p2H Q7 Qp2`fB
/B+iBQM mM+2`i BMiv- 7Q` 2 +? /X; `2Q` Q7 K Q M 2 / 2 B M 2 ? 2 u2 ` / i b2i
U6B; 6V 2 2 b22 i? i o GG bHB; ?iHv Qp2` @2biBK 22(0;0:7)2r m B H +2`i BMiv 7
Bi BM7` @2biBK i2b Bi 7QXH `;2` p Hm2b Q7