

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/325908898>

Neural-net-induced Gaussian process regression for function approximation and PDE solution

Preprint · June 2018

DOI: 10.13140/RG.2.2.35191.11681

CITATIONS

0

READS

444

3 authors, including:



Guofei Pang

Brown University

18 PUBLICATIONS **106** CITATIONS

[SEE PROFILE](#)



George Em Karniadakis

Brown University

717 PUBLICATIONS **25,550** CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Office of Naval Research and Airforce Office of Scientific Research [View project](#)



Machine Learning, AFOSR and DARPA [View project](#)

Neural-net-induced Gaussian process regression for function approximation and PDE solution

Guofei Pang, Liu Yang, and George Em Karniadakis*

Division of Applied Mathematics, Brown University, Providence, RI 02912, USA

Abstract

Neural-net-induced Gaussian process (NNGP) regression inherits both the high expressivity of deep neural networks (deep NNs) as well as the uncertainty quantification property of Gaussian processes (GPs). We generalize the current NNGP to first include a larger number of hyperparameters and subsequently train the model by maximum likelihood estimation. Unlike previous works on NNGP that targeted classification, here we apply the generalized NNGP to function approximation and to solving partial differential equations (PDEs). Specifically, we develop an analytical iteration formula to compute the covariance function of GP induced by deep NN with an error-function nonlinearity. We compare the performance of the generalized NNGP for function approximations and PDE solutions with those of GPs and fully-connected NNs. We observe that for smooth functions the generalized NNGP can yield the same order of accuracy with GP, while both NNGP and GP outperform deep NN. For non-smooth functions, the generalized NNGP is superior to GP and comparable or superior to deep NN.

Keywords: NN-induced Gaussian process, neural network, machine learning, partial differential equation, uncertainty quantification

1 Introduction

The high expressivity of deep NNs [1] can be combined with the predictive ability of GP regression by taking advantage of the apparent equivalence between GPs and fully-connected, infinitely-wide, deep NNs with random weights and biases. This equivalence was first proved by [2] for shallow NNs. In a subsequent paper, the authors of [3] obtained analytically the covariance functions of the GPs induced by the shallow NNs having two specific activation functions: error and Gaussian functions. More recently, the authors of [4] extended [3]’s work to the case of deep NNs having generic activation functions, with reference to the mean field theory of signal propagation [1, 5].

The term NNGP was coined by the authors of [4], but until now this method despite its potential merits has not been used extensively in applications. In particular, the authors of [4] validated the NNGP on two image recognition datasets: MNIST (handwritten digits data) and CIFAR-10 (color images data), and observed that NNGP outperforms finite width NNs in terms of classification accuracy. The primary objective of our work is to exploit the flexibility of NNGP in order to approximate smooth and discontinuous functions, and furthermore in assessing the effectiveness of NNGP to solve linear and nonlinear PDEs, compared to the GP and deep NN methods formulated recently in [6, 7, 8]. Specifically, so far GP regression has been successfully applied to solution of PDEs including linear and nonlinear forward problems [6] as well as inverse problems [9, 10]. Use of GP regression can bypass discretizing differential operators by properly placing GP priors, and can easily handle noisy data as well as the associated uncertainty propagation in time-dependent problems. Hence, the motivation of the current research is to evaluate if indeed NNGP possesses not only the high expressivity of deep NNs but also the uncertainty quantification property of GPs. To this end, we will perform three-way comparisons between NNGP, GP and deep NN for prototypical problems involving function approximation and PDE solution.

The present paper makes two main contributions.

First, to pursue high expressivity, we increase the number of hyperparameters of NNGP by assuming that the prior variances of weights and biases vary layer by layer and that the weight

*Corresponding author: george_karniadakis@brown.edu

variances between the input layer and the first hidden layer are different for different neurons in the input layer. We also train the resulting NNGP regression model by using maximum likelihood estimation. We note that the authors of [4] assumed only two hyperparameters: one weight variance and one bias variance, which are kept fixed over different layers; the two hyperparameters are adjusted to achieve the best performance on validation dataset. However, no training was performed in [4].

Second, we derive the analytical covariance function for the NNGP induced by the deep NN with the error-function nonlinearity. This is motivated by the observation that the analytical covariance function induced by ReLU activated NN, given in [4], cannot be used to solve the two-dimensional Poisson equation. Denoting by $k(\mathbf{x}, \mathbf{x}')$ the analytical covariance function from the ReLU case where \mathbf{x} and \mathbf{x}' are input vectors of length two, we observed that the Laplacian of the covariance function, namely, $\Delta_{\mathbf{x}}\Delta_{\mathbf{x}'}k(\mathbf{x}, \mathbf{x}')$, which has to be computed in order to solve the problem (see Eq.(18)), will tend to infinity as $\mathbf{x} \rightarrow \mathbf{x}'$.

The paper is organized as follows. In Section 2, we review the equivalence between GPs and infinitely wide NNs, which is a cornerstone of NNGP. We change the notations for weight and bias variances compared to the previous work [4] in order to show a different hyperparameter setup. Section 3 introduces the analytically tractable covariance functions induced by the deep NNs with the ReLU and error-function nonlinearities. We employed a version of Gaussian quadrature, proposed in [4], to validate numerically the two analytical covariance functions. The methodology on GP regression for function approximation and PDE solution is briefly reviewed and summarized in Section 4. Section 5 compares the accuracy of GP, NNGP and NN in function approximation and PDE solution. The uncertainty estimates of GP and NNGP are also presented in the section. Finally, we conclude in Section 5 with a brief summary.

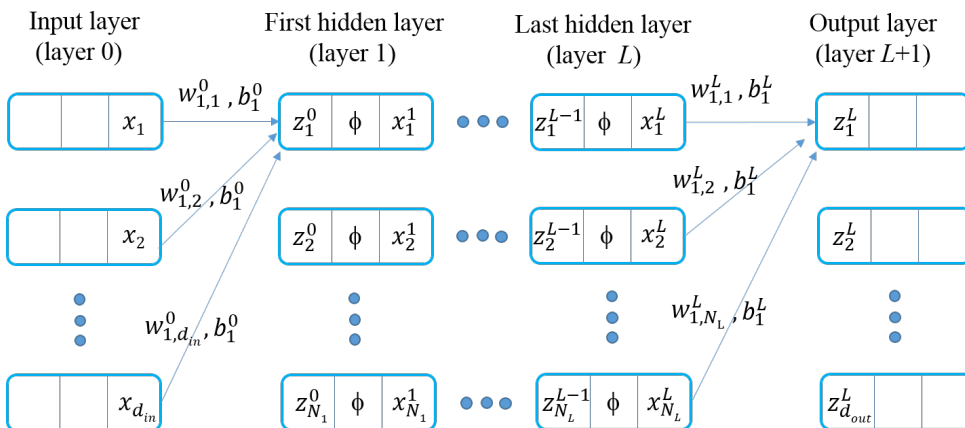


Figure 1: A fully connected neural net with L hidden layers. For each unit or neuron in hidden layers, there exist one input z_i^{l-1} and one output x_i^l with $l = 1, 1 \dots, L$. Layer 0 is the input layer. x_i is the i -th component of the input vector \mathbf{x} , and z_i^L is the i -th component of the output vector \mathbf{z} . The dimensions of input and output spaces are d_{in} and d_{out} , respectively. The width for hidden layer l is N_l . At the center of each unit is the activation function $\phi(\cdot)$ that transforms input to the corresponding output. Between two successive layers, the weight $w_{i,j}^l$ for $l = 0, 1, \dots, L$ denotes the contribution of unit j in layer l to unit i in layer $l+1$. Layer $L+1$ is the output layer. The bias b_i^l is attached to unit i in layer $l+1$ for $l = 0, 1, \dots, L$. Note that for clarity most of connecting arrows between layers are omitted.

2 Equivalence between GPs and infinitely wide NNs

In this section, we adopt nearly the same notations as [4] except those for weight and bias variances. We add the subscript l to the variances σ_w^2 and σ_b^2 in order to represent index of layer and the subscript j to σ_w^2 in order to distinguish the weights coming from different neurons in input layer.

The fully-connected neural net can be represented as in Fig.1. The equivalence relies on two assumptions: (1) Weight and bias parameters are all random variables. For $l \geq 1$ the weights $w_{i,j}^l$ are independent and identically distributed (i.i.d.) and obey the same distribution \mathcal{D}_l with zero mean and same variance, namely, $w_{i,j}^l \sim \mathcal{D}_l(0, \sigma_{w,l}^2/N_l), \forall i, j$. The biases b_i^l ($l \geq 1$) are i.i.d.

Gaussian random variables, $b_i^l \sim \mathcal{N}(0, \sigma_{b,l}^2), \forall i$. The weights and biases are independent with each other. Note that in [4] the weight and bias variances are kept fixed over different layers, i.e., $\sigma_{w,l}^2 \equiv \sigma_w^2$ and $\sigma_{b,l}^2 \equiv \sigma_b^2$, but we here assume the variances vary layer by layer and increase the number of variances to be determined. For $l = 0$, the weights need only to be independent with respect to the subscript i , i.e., $w_{ij}^0 \sim \mathcal{D}_0(0, \sigma_{w,0,j}^2), \forall i$; the biases are independent Gaussian random variables, $b_i^0 \sim \mathcal{N}(0, \sigma_{b,0}^2), \forall i$. Unlike [4], we remove the assumption of being identically distributed for the weights. These weights do not have to be identically distributed since we cannot use the Central Limit Theorem for the input of the first hidden layer, namely, z_i^0 . Actually, the width of input layer is finite and will never be infinity. Allowing weight variances to vary for different neurons in input layer, we increase again the number of variances to be determined in NNGP. (2) The widths of all hidden layers tend to infinity, i.e., $N_l \rightarrow \infty, 1 \leq l \leq L$. This allows one to use the Central Limit Theorem to deduce the equivalence relation.

It follows from [3] and [11] that for analytical derivation of the covariance function the distribution \mathcal{D}_0 must be Gaussian. The specific forms of $D_l, l \geq 1$ do not really matter, since using the Central Limit Theorem one always obtains the Gaussian distribution whatever \mathcal{D}_l is.

The i -th component of the input of the second hidden layer is computed as

$$z_i^1(\mathbf{x}) = \sum_{j=1}^{N_1} w_{ij}^1 x_j^1(\mathbf{x}) + b_i^1, \quad x_j^1(\mathbf{x}) = \phi \left(\sum_{k=1}^{d_{in}} w_{jk}^0 x_k + b_j^0 \right). \quad (1)$$

In the rest of paper, we use bold letter, say \mathbf{x} and \mathbf{K} , to represent column vector or matrix. The dependence on the input vector \mathbf{x} is emphasized in the above equations. Since weights and biases w_{jk}^0 and b_j^0 are i.i.d., the output of the first hidden layer $x_j^1(\mathbf{x})$ is also i.i.d. with respect to the subscript j . Because $z_i^1(\mathbf{x})$ is a sum of i.i.d. terms, the use of the Central Limit Theorem yields that $z_i^1(\mathbf{x})$ will be Gaussian distributed if $N_1 \rightarrow \infty$. Similarly, from the multidimensional Central Limit Theorem, any finite collection of $\{z_i^1(\mathbf{x}_1), \dots, z_i^1(\mathbf{x}_k)\}$ will form a joint multivariate Gaussian distribution, which is exactly the definition of a Gaussian process. Note that $\mathbf{x}_1, \dots, \mathbf{x}_k$ are k arbitrary input vectors.

Denoting by $\mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ the GP with the mean function $\mu(\cdot)$ and the covariance function $k(\cdot, \cdot)$, we learn that $z_i^1(\mathbf{x}) \sim \mathcal{GP}(\mu^1(\mathbf{x}), k^1(\mathbf{x}, \mathbf{x}'))$. Also, $\{z_1^1(\mathbf{x}), z_2^1(\mathbf{x}), \dots, z_{N_2}^1(\mathbf{x})\}$ are independent GPs with same mean and covariance functions. The mean function of GP $z_i^1(\mathbf{x})$ is zero due to the zero mean of the distribution \mathcal{D}_1 , and the covariance function is computed as

$$\begin{aligned} k^1(\mathbf{x}, \mathbf{x}') &= \mathbb{E}(z_i^1(\mathbf{x})z_i^1(\mathbf{x}')) = \mathbb{E} \left(\left(\sum_{j=1}^{N_1} w_{ij}^1 x_j^1(\mathbf{x}) + b_i^1 \right) \left(\sum_{k=1}^{N_1} w_{ik}^1 x_k^1(\mathbf{x}') + b_i^1 \right) \right) \\ &= \mathbb{E} \left(\left(\sum_{j=1}^{N_1} w_{ij}^1 \phi(z_j^0(\mathbf{x})) + b_i^1 \right) \left(\sum_{k=1}^{N_1} w_{ik}^1 \phi(z_k^0(\mathbf{x}')) + b_i^1 \right) \right) \\ &= \sum_{j=1}^{N_1} \mathbb{E}(w_{ij}^1)^2 \mathbb{E} [\phi(z_j^0(\mathbf{x}))\phi(z_j^0(\mathbf{x}'))] + \mathbb{E}(b_i^1)^2 \\ &= \frac{\sigma_{w,1}^2}{N_1} \sum_{j=1}^{N_1} \mathbb{E} [\phi(z_j^0(\mathbf{x}))\phi(z_j^0(\mathbf{x}'))] + \sigma_{b,1}^2 \\ &= \sigma_{w,1}^2 \mathbb{E} [\phi(z_j^0(\mathbf{x}))\phi(z_j^0(\mathbf{x}'))] + \sigma_{b,1}^2. \end{aligned} \quad (2)$$

In the above derivation, we use the independence relations between weights $\mathbb{E}(w_{ij}^1 w_{ik}^1) = \delta_{jk} \sigma_{w,1}^2$ (δ_{ij} is the Kronecker delta), between weight and bias $\mathbb{E}(w_{ij}^1 b_i^1) = \mathbb{E}(w^1) \mathbb{E}(b_i^1) = 0$, and between weight and activation function $\mathbb{E} [(w_{ij}^1)^2 \phi(z_j^0(\mathbf{x}))\phi(z_j^0(\mathbf{x}'))] = \mathbb{E} (w_{ij}^1)^2 \mathbb{E} [\phi(z_j^0(\mathbf{x}))\phi(z_j^0(\mathbf{x}'))]$. The last equality in Eq. (2) comes from the independence of $\phi(z_j^0(\mathbf{x}))\phi(z_j^0(\mathbf{x}'))$ with respect to j . The input of the first hidden layer is denoted by $z_j^0(\mathbf{x}) = \sum_{k=1}^{d_{in}} w_{jk}^0 x_k + b_j^0$, and the expectation $\mathbb{E}(\cdot)$ is taken over the distribution of w_{ij}^0 and b_i^0 . Actually, $z_j^0(\mathbf{x})$ is a GP only if the distribution D_0 is

Gaussian. In this case, the covariance function of the GP $z_j^0(\mathbf{x})$ is

$$\begin{aligned} k^0(\mathbf{x}, \mathbf{x}') &= \mathbb{E}(z_i^0(\mathbf{x})z_i^0(\mathbf{x}')) = \mathbb{E}\left(\left(\sum_{j=1}^{d_{in}} w_{ij}^0 x_j + b_i^0\right)\left(\sum_{k=1}^{d_{in}} w_{ik}^0 x'_k + b_i^0\right)\right) \\ &= \sum_{j=1}^{d_{in}} \mathbb{E}(w_{ij}^0)^2 x_j x'_j + \mathbb{E}(b_i^0)^2 \\ &= \sum_{j=1}^{d_{in}} \sigma_{w,0,j}^2 x_j x'_j + \sigma_{b,0}^2. \end{aligned} \quad (3)$$

The input vector components x_j and x'_j are deterministic and thus $\mathbb{E}(x_j x'_j) = x_j x'_j$.

3 Two NN-induced covariance kernels of GP

The arguments of the previous section can be extended to deeper layers by induction. For layer l (≥ 2), using the multidimensional Central Limit Theorem, we learn that $z_i^l(\mathbf{x})$ and $z_i^{l-1}(\mathbf{x})$ are both GPs. Specifically, $z_i^l(\mathbf{x})$ has zero mean function and the covariance function computed as

$$k^l(\mathbf{x}, \mathbf{x}') = \mathbb{E}(z_i^l(\mathbf{x})z_i^l(\mathbf{x}')) = \sigma_{w,l}^2 \mathbb{E}[\phi(z_i^{l-1}(\mathbf{x}))\phi(z_i^{l-1}(\mathbf{x}'))] + \sigma_{b,l}^2. \quad (4)$$

The expectation in Eq. (4) is taken over the GP governing $z_i^{l-1}(\mathbf{x})$, and this amounts to integrating over the joint distribution of two Gaussian random variables $z_i^{l-1}(\mathbf{x})$ and $z_i^{l-1}(\mathbf{x}')$ given \mathbf{x} and \mathbf{x}' . The covariance function is rewritten as

$$k^l(\mathbf{x}, \mathbf{x}') = \sigma_{w,l}^2 \int_{\mathbb{R}^2} \phi(z_i^{l-1}(\mathbf{x}))\phi(z_i^{l-1}(\mathbf{x}'))p(z_i^{l-1}(\mathbf{x}), z_i^{l-1}(\mathbf{x}'))dz_i^{l-1}(\mathbf{x})dz_i^{l-1}(\mathbf{x}') + \sigma_{b,l}^2, \quad (5)$$

where the density function of the joint distribution of $z_i^l(\mathbf{x})$ and $z_i^l(\mathbf{x}')$ is

$$p(z_i^{l-1}(\mathbf{x}), z_i^{l-1}(\mathbf{x}')) = \frac{1}{2\pi|\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}\begin{bmatrix} z_i^{l-1}(\mathbf{x}) \\ z_i^{l-1}(\mathbf{x}') \end{bmatrix}^T \Sigma^{-1} \begin{bmatrix} z_i^{l-1}(\mathbf{x}) \\ z_i^{l-1}(\mathbf{x}') \end{bmatrix}\right) \quad (6)$$

with the covariance matrix

$$\Sigma = \begin{bmatrix} \mathbb{E}(z_i^{l-1}(\mathbf{x}))^2 & \mathbb{E}(z_i^{l-1}(\mathbf{x})z_i^{l-1}(\mathbf{x}')) \\ \mathbb{E}(z_i^{l-1}(\mathbf{x}')z_i^{l-1}(\mathbf{x})) & \mathbb{E}(z_i^{l-1}(\mathbf{x}')^2) \end{bmatrix} = \begin{bmatrix} k^{l-1}(\mathbf{x}, \mathbf{x}') & k^{l-1}(\mathbf{x}, \mathbf{x}') \\ k^{l-1}(\mathbf{x}', \mathbf{x}) & k^{l-1}(\mathbf{x}', \mathbf{x}') \end{bmatrix}. \quad (7)$$

Eq. (5) is actually an iteration formula, given by

$$\begin{aligned} k^l(\mathbf{x}, \mathbf{x}') &= \sigma_{w,l}^2 F_\phi(k^{l-1}(\mathbf{x}, \mathbf{x}), k^{l-1}(\mathbf{x}', \mathbf{x}'), k^{l-1}(\mathbf{x}, \mathbf{x}')) + \sigma_{b,l}^2, l = 1, 2, \dots, L, \\ k^0(\mathbf{x}, \mathbf{x}') &= \mathbf{x}^T \mathbf{\Lambda} \mathbf{x}' + \sigma_{b,0}^2, \end{aligned} \quad (8)$$

where the diagonal matrix $\mathbf{\Lambda}$ is defined by $\mathbf{\Lambda} = \text{diag}(\sigma_{w,0,1}^2, \sigma_{w,0,2}^2, \dots, \sigma_{w,0,d_{in}}^2)$. Note that $k^{l-1}(\mathbf{x}, \mathbf{x}') = k^{l-1}(\mathbf{x}', \mathbf{x})$ due to symmetry of the kernel. The trivariate function $F_\phi(\cdot)$ represents the two-fold integral in Eq. (5), which depends on specific form of activation function $\phi(\cdot)$. For certain $\phi(\cdot)$, the function $F_\phi(\cdot)$ is analytically tractable. Two iteration formulas that are analytically derivable will be given in the next two subsections.

3.1 Kernel from the ReLU nonlinearity

Letting the activation function be a rectifier $\phi(x) = \max(x, 0)$, we can derive the analytical expression for $F_\phi(\cdot)$, and the corresponding iteration formula is [4, 11]

$$\begin{aligned} k^l(\mathbf{x}, \mathbf{x}') &= \frac{\sigma_{w,l}^2}{2\pi} \sqrt{k^{l-1}(\mathbf{x}, \mathbf{x})k^{l-1}(\mathbf{x}', \mathbf{x}')} (\sin \theta^{l-1} + (\pi - \theta^{l-1}) \cos \theta^{l-1}) + \sigma_{b,l}^2, l = 1, 2, \dots, L, \\ \theta^{l-1} &= \cos^{-1}\left(\frac{k^{l-1}(\mathbf{x}, \mathbf{x}')}{\sqrt{k^{l-1}(\mathbf{x}, \mathbf{x})k^{l-1}(\mathbf{x}', \mathbf{x}')}}\right), \\ k^0(\mathbf{x}, \mathbf{x}') &= \mathbf{x}^T \mathbf{\Lambda} \mathbf{x}' + \sigma_{b,0}^2. \end{aligned} \quad (9)$$

Here we have increased the number of undetermined variances compared to the original formula in [4] in which $\sigma_{w,l}^2 = \sigma_w^2$, $\sigma_{b,l}^2 = \sigma_b^2$, and $\mathbf{\Lambda} = \frac{\sigma_w^2}{d_{in}} \mathbf{I}$. There are totally $d_{in} + 1 + 2L$ undetermined parameters θ : $\sigma_{w,0,j}^2$ ($j = 1, 2, \dots, d_{in}$), $\sigma_{b,0}^2, \sigma_{w,l}^2$ ($l = 1, 2, \dots, L$), and $\sigma_{b,l}^2$ ($l = 1, 2, \dots, L$). These parameters (a.k.a. hyperparameters in GP regression) can be learned by using maximum likelihood estimation.

Since the trivariate function $F_\phi(\cdot)$ is not analytically tractable for a generic nonlinearity, the authors of [4] developed an efficient Gaussian quadrature scheme to approximate the two-fold integral in Eq.(5). Fig. 2 (left) compares the numerical results computed by the Gaussian quadrature scheme with the analytical results computed by Eq. (9), showing good agreement.

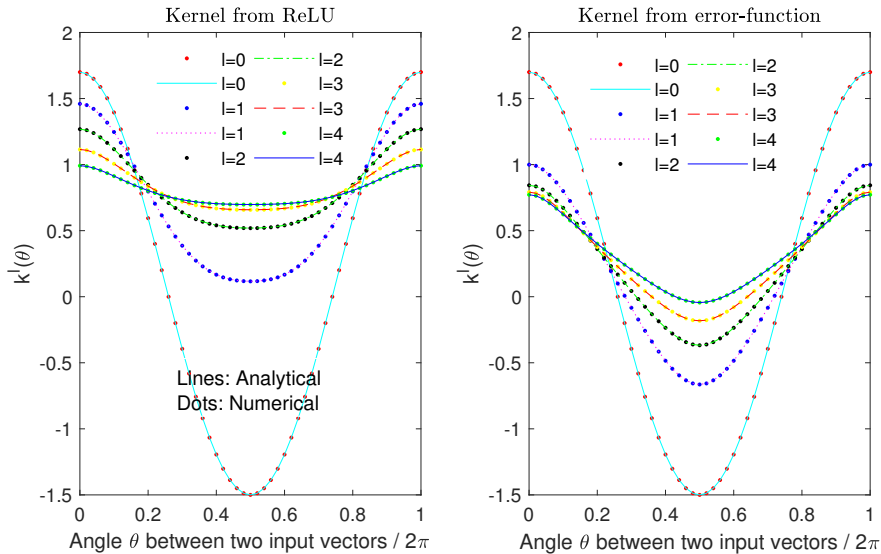


Figure 2: Comparison of the analytical and numerical iteration formulas for the covariance functions from the ReLU (left) and error-function (right) nonlinearities. The covariance function $k^l(\mathbf{x}, \mathbf{x}')$ of the GP $z_i^l(\mathbf{x})$ is computed analytically and numerically for $l = 0, 1, 2, 3, 4$ ($L = 4$). For ease of comparison, the input vectors $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^2$ ($d_{in} = 2$) are normalized to $\|\mathbf{x}\|_2 = \|\mathbf{x}'\|_2 = 1$. The angle between the two vectors is $\theta = \cos^{-1}(\mathbf{x}^T \mathbf{x}')$. All the weight variances are set to be 1.6 and the bias variances are 0.1.

3.2 Kernel from the error-function nonlinearity

The error-function nonlinearity can also yield an analytical iteration formula just as the ReLU nonlinearity does. However, the existing formula only holds for single-hidden layer, which is given by [3]. The present paper extends the single-hidden layer case to multi-layer case using the strategy presented in Section 2.3 of [11]. The resulting iteration formula is

$$k^l(\mathbf{x}, \mathbf{x}') = \frac{2\sigma_{w,l}^2}{\pi} \sin^{-1} \left(\frac{2k^{l-1}(\mathbf{x}, \mathbf{x}')}{\sqrt{(1 + 2k^{l-1}(\mathbf{x}, \mathbf{x})) (1 + 2k^{l-1}(\mathbf{x}', \mathbf{x}'))}} \right) + \sigma_{b,l}^2, l = 1, 2, \dots, L, \quad (10)$$

$$k^0(\mathbf{x}, \mathbf{x}') = \mathbf{x}^T \mathbf{\Lambda} \mathbf{x}' + \sigma_{b,0}^2.$$

To validate the above formula, we compare the analytical results from the above formula with the numerical results from the aforementioned Gaussian quadrature scheme, and as we see in Fig.2 (right) the new iteration formula (10) is correct.

We specifically call the GPs with the kernels that are induced iteratively by NNs “NNGP” for short. NNGP is actually a specific case of GP, but we still distinguish it from the standard GP involving an *ad hoc* kernel (e.g., square exponential and Matern kernels). An exception is the output of the first hidden layer, $x_i^1(\mathbf{x})$. The output is also a GP and its covariance function is [3, 12]

$$k(\mathbf{x}, \mathbf{x}') = \frac{2}{\pi} \sin^{-1} \left(\frac{2k^0(\mathbf{x}, \mathbf{x}')}{\sqrt{(1 + 2k^0(\mathbf{x}, \mathbf{x})) (1 + 2k^0(\mathbf{x}', \mathbf{x}'))}} \right). \quad (11)$$

We still regard the GP with this kernel as a standard GP rather than a NNGP, as the kernel $k(\mathbf{x}, \mathbf{x}')$ corresponds to an incomplete, shallow NN without output layer. In Section 5.2.2 we will compare the predictive ability of this kernel with our generic kernel $k^l(\mathbf{x}, \mathbf{x}')$ computed by formula (10).

4 GP regression

The essence of GP regression is the conditional distribution of quantities of interest, \mathbf{q} , given the known observations \mathbf{o} , where \mathbf{q} and \mathbf{o} are both Gaussian random vectors. Thanks to the analytical tractability of Gaussian distribution, we can derive analytically mean vector and covariance matrix of the conditional distribution. Supposed that \mathbf{q} and \mathbf{o} have the joint multivariate Gaussian distribution with zero mean vector and block-wise covariance matrix

$$\begin{bmatrix} \mathbf{q} \\ \mathbf{o} \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} \mathbf{K}_{qq} & \mathbf{K}_{qo} \\ \mathbf{K}_{oq} & \mathbf{K}_{oo} \end{bmatrix}\right), \quad (12)$$

the conditional distribution $p(\mathbf{q}|\mathbf{o})$ is also a multivariate Gaussian distribution with the mean vector

$$\mathbf{m}(\mathbf{q}) = \mathbf{K}_{qo}\mathbf{K}_{oo}^{-1}\mathbf{o} \quad (13)$$

and the covariance matrix

$$\mathbf{K}(\mathbf{q}) = \mathbf{K}_{qq} - \mathbf{K}_{qo}\mathbf{K}_{oo}^{-1}\mathbf{K}_{oq} \quad (14)$$

(See (A.5) and (A.6) of [12]). Zero-mean and σ_{noise}^2 -variance Gaussian white noise is assumed in observations.

Note that observations and quantities of interest are usually evaluated at specific time-space coordinates. To preserve a multivariate Gaussian distribution (12) at arbitrary time-space coordinates, the observations and the quantities of interest can be seen as samples of GPs. In other words, $\mathbf{o} = [X_o(\mathbf{x}_1), \dots, X_o(\mathbf{x}_{N_o})]^T$ is a sample of $X_o(\mathbf{x}) \sim \mathcal{GP}(0, k_o(\mathbf{x}, \mathbf{x}'))$ and $\mathbf{q} = [X_q(\mathbf{x}_1), \dots, X_q(\mathbf{x}_{N_q})]^T$ is a sample of $X_q(\mathbf{x}) \sim \mathcal{GP}(0, k_q(\mathbf{x}, \mathbf{x}'))$, where the index set \mathbf{x} represents arbitrary time and/or space coordinates. Particularly, for solving PDEs, the observations \mathbf{o} consist of samples from different GPs. Different covariance function $k_o(\cdot, \cdot)$ will capture different trends and features that the observations \mathbf{o} exhibit.

Before computing mean and covariance of conditional distribution from formulas (13) and (14), we need to first learn the parameters hidden in covariance function as well as the noise variance σ_{noise}^2 . These parameters are also known as hyperparameters and can be learned by using maximum likelihood estimation (MLE) [12, 13]. In practice, we find the optimal hyperparameters (including undetermined parameters $\boldsymbol{\theta}$ in covariance function and the noise variance σ_{noise}^2) that minimize the negative log marginal-likelihood [12]

$$\frac{1}{2}\mathbf{o}^T\mathbf{K}_{oo}^{-1}(\boldsymbol{\theta}, \sigma_{noise}^2)\mathbf{o} + \frac{1}{2}\log|\mathbf{K}_{oo}(\boldsymbol{\theta}, \sigma_{noise}^2)| + \frac{N_o}{2}\log 2\pi, \quad (15)$$

where $|\cdot|$ denotes the determinant and the covariance matrix for observations (or training data) \mathbf{K}_{oo} depends on $\boldsymbol{\theta}$ and σ_{noise}^2 .

In what follows, we will briefly review that different formulations of observations \mathbf{o} will enable the GP regression to solve different problems including function approximation and PDE solution.

4.1 Function approximation

Let $\mathbf{q} = [f(\mathbf{x})]$ and $\mathbf{o} = [f(\mathbf{x}_1) + \epsilon, \dots, f(\mathbf{x}_{N_o}) + \epsilon]^T$ where ϵ is Gaussian white noise having variance σ_{noise}^2 . The objective is to approximate the unknown scalar-valued function $f(\mathbf{x})$ given the observations \mathbf{o} . After prescribing the covariance function $k_o(\cdot, \cdot)$ for computing the covariance matrix \mathbf{K}_{oo} and then training the regression model by the MLE, we derive the function approximation by using formula (13). Moreover, the use of the variance computed from (14) yields the confidence interval. Denoting by \mathbf{X} (N_o by d_{in} matrix) the collection of the evaluation points $\{\mathbf{x}_j\}$ (usually called training inputs), we formulate the covariance matrix in (12) by:

$$\begin{aligned} \mathbf{K}_{oo} &= [k_o(\mathbf{X}, \mathbf{X}) + \sigma_{noise}^2\mathbf{I}], \\ \mathbf{K}_{qo} &= \mathbf{K}_{oq}^T = [k_o(\mathbf{x}, \mathbf{X})], \\ \mathbf{K}_{qq} &= [k_o(\mathbf{x}, \mathbf{x})]. \end{aligned} \quad (16)$$

4.2 PDE solution

First we consider the following time-independent linear PDE

$$\begin{aligned} L_{\mathbf{x}}u(\mathbf{x}) &= f(\mathbf{x}), \mathbf{x} \in \Omega \in \mathbb{R}^{d_{in}}, \\ u(\mathbf{x}) &= g(\mathbf{x}), \mathbf{x} \in \partial\Omega, \end{aligned} \quad (17)$$

where $L_{\mathbf{x}}$ is a linear differential operator. The source term $f(\cdot)$ and the boundary condition term $g(\cdot)$ are the only information we know, and thus we put them in the observations by letting $\mathbf{o} = [\mathbf{o}_u^T, \mathbf{o}_f^T]^T = [g(\mathbf{x}_{u,1}) + \epsilon_u, \dots, g(\mathbf{x}_{u,N_u}) + \epsilon_u, f(\mathbf{x}_{f,1}) + \epsilon_f, \dots, f(\mathbf{x}_{f,N_f}) + \epsilon_f]^T$, where we select N_u boundary points as the evaluation points for \mathbf{o}_u and N_f domain points as the evaluation points for \mathbf{o}_f . The variances of the Gaussian white noises ϵ_u and ϵ_f are σ_u^2 and σ_f^2 , respectively. We assume that the solution $\mathbf{q} = [u(\mathbf{x})]$ is a GP with zero mean and covariance function $k_u(\mathbf{x}, \mathbf{x}')$, and the source term is another GP that depends on u , whose covariance function is written as [6]

$$k_f(\mathbf{x}, \mathbf{x}') = L_{\mathbf{x}}L_{\mathbf{x}'}k_u(\mathbf{x}, \mathbf{x}'). \quad (18)$$

Denoting by \mathbf{X}_u (N_u by d_{in} matrix) the collection of the boundary evaluation points $\{\mathbf{x}_{u,j}\}$ and by \mathbf{X}_f (N_f by d_{in} matrix) the collection of domain evaluation points $\{\mathbf{x}_{f,j}\}$, we formulate the covariance matrix in (12) as [6]

$$\begin{aligned} \mathbf{K}_{oo} &= \begin{bmatrix} k_u(\mathbf{X}_u, \mathbf{X}_u) + \sigma_{noise,u}^2 \mathbf{I} & L_{\mathbf{x}'}k_u(\mathbf{X}_u, \mathbf{X}_f) \\ L_{\mathbf{x}}k_u(\mathbf{X}_f, \mathbf{X}_u) & L_{\mathbf{x}}L_{\mathbf{x}'}k_u(\mathbf{X}_f, \mathbf{X}_f) + \sigma_{noise,f}^2 \mathbf{I} \end{bmatrix}, \\ \mathbf{K}_{qo} &= \mathbf{K}_{oq}^T = \begin{bmatrix} k_u(\mathbf{x}, \mathbf{X}_u) & L_{\mathbf{x}'}k_u(\mathbf{x}, \mathbf{X}_f) \end{bmatrix}, \\ \mathbf{K}_{qq} &= [k_u(\mathbf{x}, \mathbf{x})]. \end{aligned} \quad (19)$$

For time-dependent problem, the above GP regression model also works as we can regard the time variable as the $(d_{in} + 1)$ -th dimension of space-time domain (See Section 4.3.1 of [6]).

The alternative way to handle the time-dependent problems is numerical GP regression [7]. Consider the following time-dependent PDE

$$\begin{aligned} \frac{\partial u(\mathbf{x}, t)}{\partial t} &= \tilde{L}_{\mathbf{x}}u(\mathbf{x}, t), \mathbf{x} \in \Omega \in \mathbb{R}^{d_{in}}, \\ u(\mathbf{x}, 0) &= h(\mathbf{x}), \\ u(\mathbf{x}, t) &= g(\mathbf{x}, t), \mathbf{x} \in \partial\Omega. \end{aligned} \quad (20)$$

The differential operator $\tilde{L}_{\mathbf{x}}$ can be nonlinear (say, Burgers' equation). We denote by $L_{\mathbf{x}}$ the linearized counterpart of $\tilde{L}_{\mathbf{x}}$. For linear operator, $L_{\mathbf{x}} = \tilde{L}_{\mathbf{x}}$. The use of Euler backward scheme leads to $u^n(\mathbf{x}) - \Delta t L_{\mathbf{x}}u^n(\mathbf{x}) = u^{n-1}(\mathbf{x})$. According to the known information from initial-boundary conditions, we define the observations at n -th time step as $\mathbf{o}^n = [\mathbf{o}_n^T, \mathbf{o}_{n-1}^T]^T = [u^n(\mathbf{x}_{b,1}) + \epsilon_n, \dots, u^n(\mathbf{x}_{b,N_b}) + \epsilon_n, u^{n-1}(\mathbf{x}_1) + \epsilon_{n-1}, \dots, u^{n-1}(\mathbf{x}_{N^{n-1}}) + \epsilon_{n-1}]^T$ and $\mathbf{q}^n = [u^n(\mathbf{x})]$. Note that $u^n(\mathbf{x}_{b,j}) = g(\mathbf{x}_{b,j}, n\Delta t)$ and $u^0(\mathbf{x}_j) = h(\mathbf{x}_j)$. Thus, we actually apply GP regression at each time step. Assuming $u^n(\mathbf{x})$ is a GP with zero mean and covariance kernel $k_{u^n}(\mathbf{x}, \mathbf{x}')$ and denoting by \mathbf{X}_b^n (N_b by d_{in} matrix) the collection of the boundary evaluation points selected at n -th time step and by \mathbf{X}^{n-1} (N^{n-1} by d_{in} matrix) the collection of the domain evaluation points sampled at $(n-1)$ -th time step, we formulate the covariance matrix as [7]

$$\begin{aligned} \mathbf{K}_{o^n o^n} &= \begin{bmatrix} k_{u^n}(\mathbf{X}_b^n, \mathbf{X}_b^n) + \sigma_{noise,n}^2 \mathbf{I} & (\mathbf{I} - \Delta t L_{\mathbf{x}'})k_{u^n}(\mathbf{X}_b^n, \mathbf{X}^{n-1}) \\ (\mathbf{I} - \Delta t L_{\mathbf{x}})k_{u^n}(\mathbf{X}^{n-1}, \mathbf{X}_b^n) & (\mathbf{I} - \Delta t L_{\mathbf{x}})(\mathbf{I} - \Delta t L_{\mathbf{x}'})k_{u^n}(\mathbf{X}^{n-1}, \mathbf{X}^{n-1}) + \sigma_{noise,n-1}^2 \mathbf{I} \end{bmatrix}, \\ \mathbf{K}_{q^n o^n} &= \mathbf{K}_{o^n q^n}^T = \begin{bmatrix} k_{u^n}(\mathbf{x}, \mathbf{X}_b^n) & (\mathbf{I} - \Delta t L_{\mathbf{x}'})k_{u^n}(\mathbf{x}, \mathbf{X}^{n-1}) \end{bmatrix}, \\ \mathbf{K}_{q^n q^n} &= [k_{u^n}(\mathbf{x}, \mathbf{x})]. \end{aligned} \quad (21)$$

Generally, we let the covariance function be the same for all n , namely $k_{u^n}(\mathbf{x}, \mathbf{x}') = k_o(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta}_n)$, except with varied undetermined parameters $\boldsymbol{\theta}_n$. Considering the propagation of the uncertainty in time marching, we need to take into account the uncertainty of the predictions in previous time

step (i.e., $u^{n-1}(\mathbf{x})$ predicted at \mathbf{X}^{n-1}). We rewrite the predictive or posterior variance in (14) as [7]

$$\begin{aligned} \mathbf{K}(\mathbf{q}^n) &= \mathbf{K}_{q^n q^n} - \mathbf{K}_{q^n o^n} \mathbf{K}_{o^n o^n}^{-1} \mathbf{K}_{o^n q^n} \\ &+ \mathbf{K}_{q^n o^n} \mathbf{K}_{o^n o^n}^{-1} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{K}(u^{n-1}(\mathbf{X}^{n-1})) \end{bmatrix} \mathbf{K}_{o^n o^n}^{-1} \mathbf{K}_{o^n q^n}. \end{aligned} \quad (22)$$

5 Numerical results

We perform a three-way comparison between NNGP, NN, and GP for function approximation and PDE solution for prototypical problems. The hyperparameters of NNGP and GP are trained by minimizing the negative log marginal-likelihood function. The conjugate gradient algorithm (see the function `minimize()` in GPML Matlab code (version 3.6) [14]) is employed for the optimization. For NN, network parameters (weights and biases) are trained by the L-BFGS-B algorithm that is provided in the TensorFlow package.

Because of the non-convex nature of the objective function, to avoid trapping into the local minima, we run our optimization algorithm code ten times with different initializations. Among the ten groups of optimized hyperparameters, for GP/NNGP we choose the one yielding the smallest negative log marginal-likelihood, and for NN we select the one producing the lowest loss for NN. The initializations of hyperparameters for GP/NNGP are taken as the first ten entries of the Halton quasi-random sequence [15] and these initializations are kept deterministic. The initializations of network weights are obtained from Xavier’s initialization algorithm [16], which is provided by the TensorFlow package. Also, for each initialization at most 200 function evaluations are allowed in conjugate gradient search in GP/NNGP. For NNGP, the central finite difference scheme with the step size 10^{-4} is used in computing the gradients of covariance function with respect to hyperparameters while for GP the analytical derivation is employed. Additionally, we compute l_2 relative error $\|\mathbf{u}_{approximate} - \mathbf{u}_{exact}\|_2 / \|\mathbf{u}_{exact}\|_2$ in quantifying accuracy of function approximation and PDE solution, where \mathbf{u} is a vector formed by the function values or PDE’s solutions evaluated at test points.

5.1 Function approximation

The covariance matrices of GP/NNGP are formulated according to (16).

5.1.1 Step function

We approximate the step function $f(x) = 1, x \geq 0; f(x) = 0$, otherwise. Ten evenly distributed training inputs and 100 equispaced test points are chosen. Fig. 3 shows the approximation by the GPs with two commonly used kernels: squared exponential (SE) and Matern [12]. The first kernel can describe well the data exhibiting smooth trend while the second one is suitable for finite regularity. However for the step function, a non-smooth function, the performance of these kernels is poor. The NNGP predictions shown in Figs. 4 and 5 are much better. In the domain away from the discontinuity, the NNGP approximation is good with small uncertainty. Particularly, the ReLU-induced kernel even succeeds in capturing the discontinuity. An important observation is that the ReLU induced kernel is more suitable for non-smooth functions than the error-function induced kernel. Another observation regarding NNGP is that deepening the inducing NN does not change the approximation accuracy for the step function. Figs. 6 and 7 show the predictions of NN, which are less accurate than NNGP’s. Unlike GP/NNGP, NN does not quantify any uncertainty for approximation.

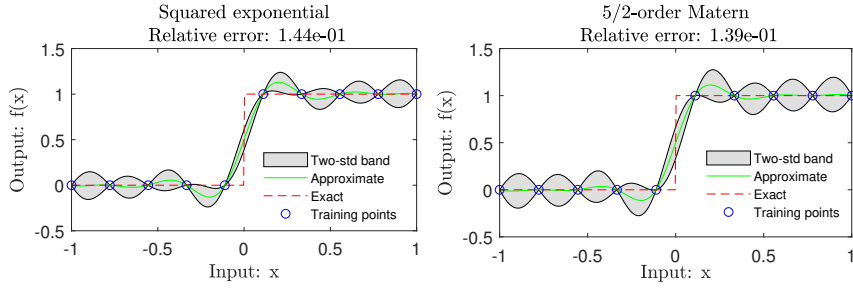


Figure 3: Approximating the step function: GPs with SE (left) and Matern (right) kernels. “Two-std band” denotes the mean vector of the conditional distribution ($m(\mathbf{q})$ computed from (13)) plus/minus twice the standard deviation (diagonals of $\sqrt{\mathbf{K}(\mathbf{q})}$ computed from (14)). The quantities of interest is $\mathbf{q} = f(\mathbf{X}_t)$ where \mathbf{X}_t (100 by 1 matrix) is the collection of test points.

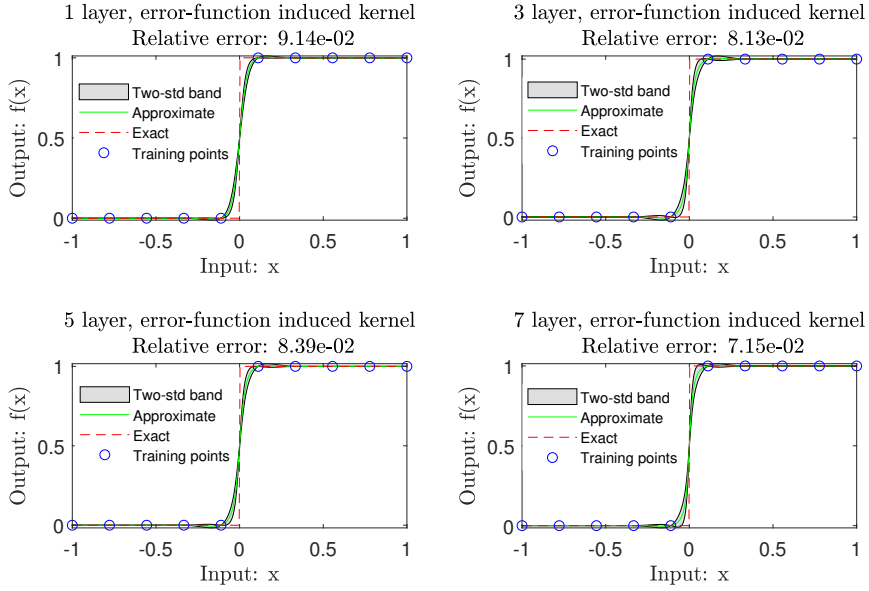


Figure 4: Approximating the step function: NNGP with error-function induced kernel. “One-layer” indicates that there is totally one hidden layer in the inducing NN ($L = 1$).

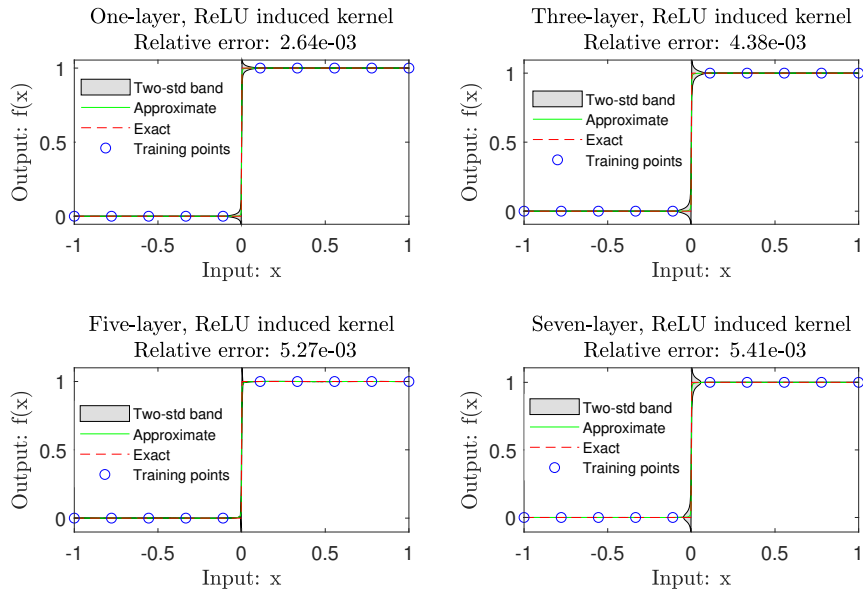


Figure 5: Approximating the step function: NNGP with ReLU induced kernel.

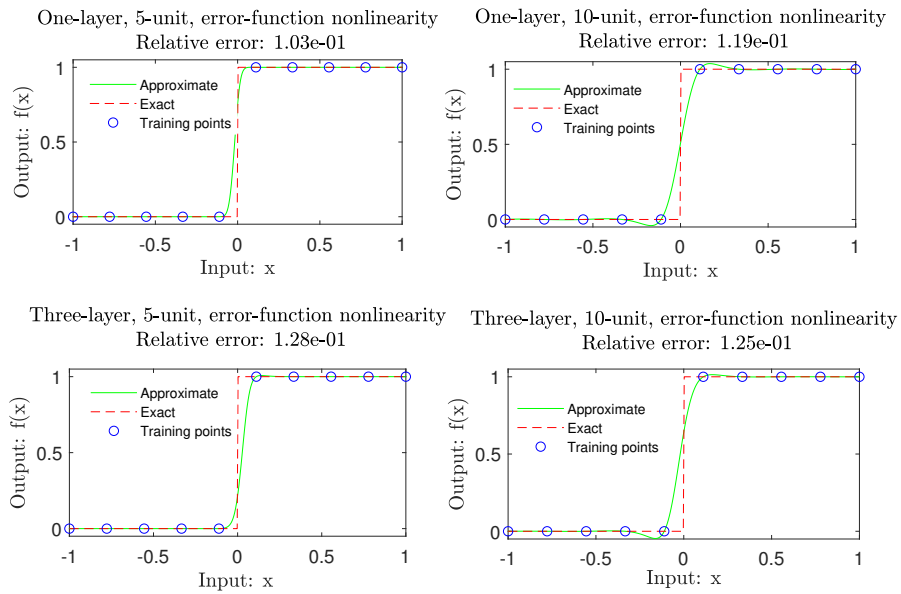


Figure 6: Approximating the step function: NN with the error-function nonlinearity (“5-unit” indicates that there exist 5 units/neurons in each hidden layer).

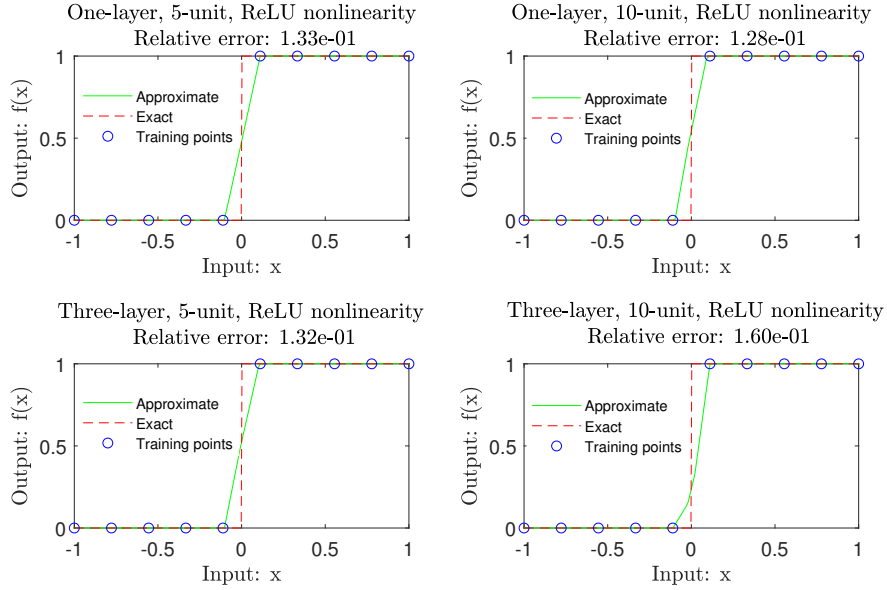


Figure 7: Approximating the step function: NN with the ReLU nonlinearity.

5.1.2 Hartmann 3D function

The Hartmann function is frequently used for testing global optimization algorithms. Here we consider the trivariate case. The function is much smoother than the step function and thus we can expect GP to perform well. To test the approximation accuracy, we first generate N ($N = 100, 200, 500$, and 1000) points by choosing the first N entries of the Halton sequence, and then randomly permute the points, followed by selecting the first 70% points as training points and the remaining 30% points as the test points.

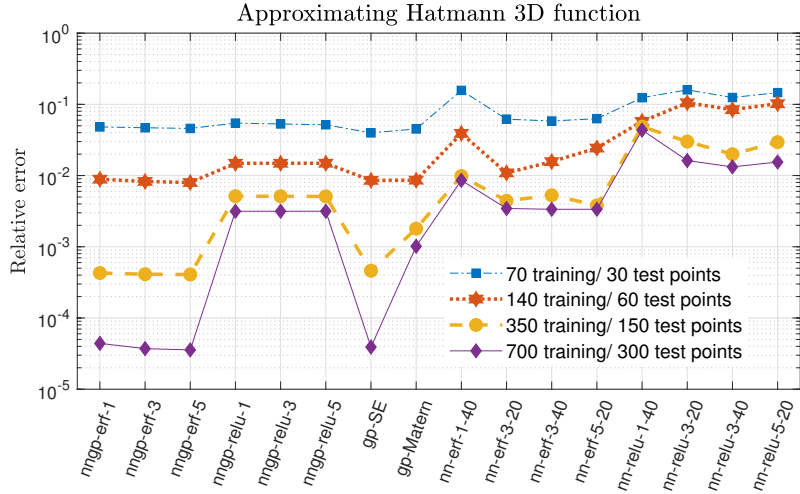


Figure 8: Approximating the Hartman 3D function: error comparison. (“nngp-erf-1” means the NNGP with the one hidden layer and error-function induced kernel; “gp-SE” denotes the GP with the squared exponential kernel; “nn-erf-3-20” indicates the three-layer, 20-unit-wide NN with the error-function nonlinearity.)

Fig. 8 compares NN, NNGP, and GP in terms of approximation accuracy. Since the Hartmann 3D function is smooth, the NNGP with error-function induced kernel and the GP with squared exponential kernel are both very accurate. ReLU is not suitable for smooth function in contrast to approximating the non-smooth step function. Additionally, the GP and NNGP both outperform NN. It is observed again that NN does not give any uncertainty estimate, whereas GP and NNGP

do. Another observation is that increasing the depth of the inducing NN in NNGP does not change the accuracy.

5.2 PDE solution

We use the proposed covariance function from error-function nonlinearity to solve the following two PDEs. We replace the covariance function $k_u(\mathbf{x}, \mathbf{x}')$ in Eq. (19) as well as $k_{u^n}(\mathbf{x}, \mathbf{x}')$ in Eq. (21) with the new kernel (10). The derivation of kernel's derivatives for the case of Poisson equation is given in Appendix.

5.2.1 2D Poisson equation

Consider the equation

$$-\Delta u(x, y) = f(x, y), (x, y) \in (0, 1)^2, \quad (23)$$

with two fabricated solutions (1) $\sin(\pi x)(y^2 + \exp(-y))$ and (2) $\sin(\pi x) \cos(2\pi(y^2 + x))$. Dirichlet boundary conditions are assumed. The second solution is more complex than the first one and thus we use more training points in the second solution case. The training inputs \mathbf{X}_f are chosen from the first N_f entries of the Halton sequence and \mathbf{X}_u is equispaced on the boundary, which are shown in Fig. 9.

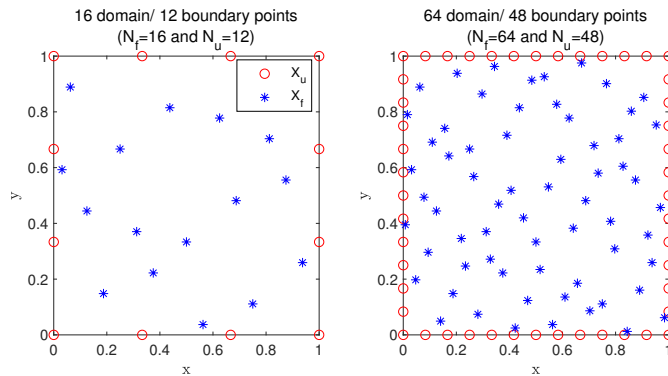


Figure 9: Solving 2D Poisson equation: Two setups of training inputs for both fabricated solutions.

The covariance matrices of GP/NNGP are obtained by (19). Figs. 10 and 11 give the error plots of NN, NNGP, and GP for fabricated solution 1 and 2, respectively. Totally 441 evenly distributed test points are taken to evaluate the relative error. The NN results are obtained according to the approach proposed in [8]. To keep a fair comparison, the same training inputs are selected in the NN case. It is observed from the figures that for different fabricated solutions, the approximation accuracy of the GP and NNGP is comparable. Also, the NN accuracy is nearly one order of magnitude lower than that of GP/NNGP.

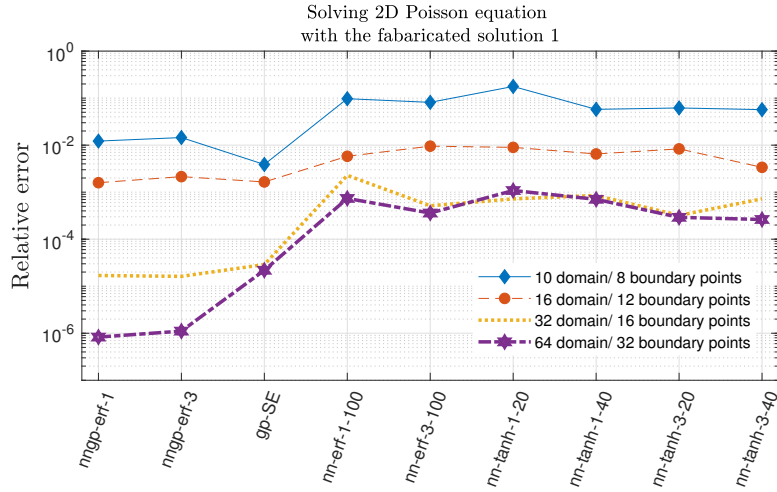


Figure 10: Solving 2D Poisson equation with the exact solution $u(x, y) = \sin(\pi x)(y^2 + \exp(-y))$. (“nn-tanh-3-20” represents the three-layer, 20-unit-wide NN with the hyperbolic tangent nonlinearity).

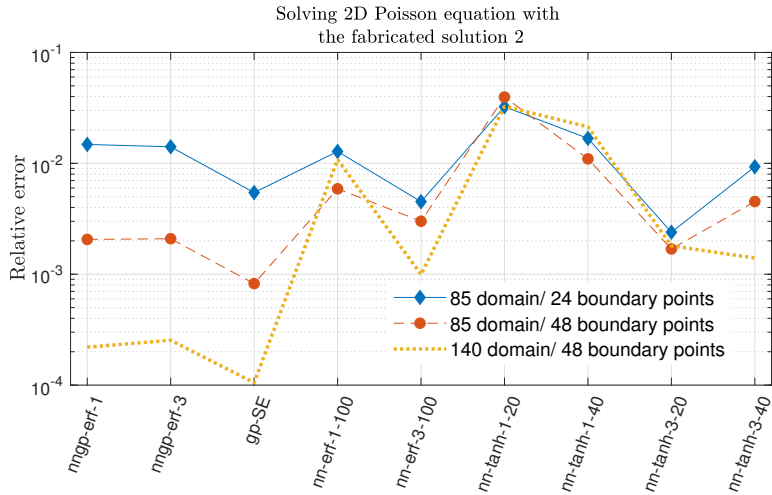


Figure 11: Solving 2D Poisson equation with the exact solution $u(x, y) = \sin(\pi x) \cos(2\pi(y^2 + x))$.

Fig. 12 shows the uncertainty estimates of GP and one-layer NNGP evaluated on the cut line $y = x$ of the square domain, for the case of fabricated solution 2. We see that for both methods uncertainty is reduced with increasing number of training points. Moreover, the uncertainty is strongly correlated with the prediction error as smaller uncertainty corresponds to lower error. Note that the uncertainty at the endpoints of the cut line, namely, $(x, y) = (0, 0)$ and $(x, y) = (1, 1)$, is exactly zero, due to the boundary condition.

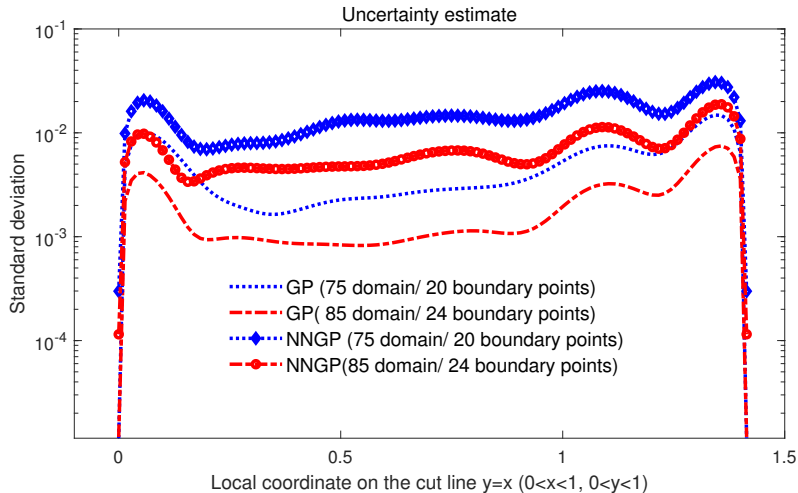


Figure 12: Solving 2D Poisson equation with the fabricated solution 2: uncertainty estimates for GP and one-layer NNGP evaluated on the cut line $y = x$. The y -axis represents the standard deviation of the conditional or posterior distribution computed by (14). The relative errors evaluated on the cut line are 0.049, 0.029, 0.022, and 0.0081 from the top curve to the bottom one. It is seen that the uncertainty is strongly correlated with the prediction error.

5.2.2 1D Burgers' equation

Consider the equation [7]

$$\frac{\partial u(x, t)}{\partial t} + u(x, t) \frac{\partial u(x, t)}{\partial x} = \frac{0.01}{\pi} \frac{\partial^2 u(x, t)}{\partial x^2}, (x, t) \in [0, 1]^2, \quad (24)$$

with $u(-1, t) = u(1, t) = 0$. The initial condition is $u(x, 0) = -\sin(\pi x)$. After the linearization through replacing the nonlinear term $u^n \frac{\partial u^n}{\partial x}$ by $\mu^{n-1} \frac{\partial u^n}{\partial x}$, where μ^n is the mean vector computed for the previous time step by (13), we derive the differential operator $L_x = \frac{0.01}{\pi} \frac{\partial^2}{\partial x^2} - \mu^{n-1} \frac{\partial}{\partial x}$. The covariance matrices of GP/NNGP are formulated by (21).

The time step size is fixed to be $\Delta t = 0.01$. To evaluate the relative error at each time step, we place 400 equispaced test points in the space domain $[-1, 1]$. We also solve the same equation using the NN approach proposed in [8]. One main difference between the NN approach and the numerical GP/NNGP regression is the sampling strategy for training inputs. For NN, we sample the training inputs in the time-space domain $(x, t) \in [0, 1]^2$, whereas for GP/NNGP, we sample the training inputs merely in the space domain and then perform time-marching. To make the comparison as fair as possible, we sample 10000 training points for NN, as in GP/NNGP we have at most 100 (number of time step) \times 100 (number of training inputs in space) = 10000 (number of time-space sampling points). The Latin hypercube sampling strategy is adopted in sampling the training inputs in NN, GP, and NNGP. To accelerate training of numerical GP/NNGP, the initial guess of hyperparameters for current time step is taken as the optimized hyperparameters attained at previous time step. This strategy is reasonable since for a small time step two successive GPs are highly correlated and therefore the respective hyperparameters should be close.

The kernel (11), rather than the SE considered in solving Poisson equation, is taken in GP. We choose a non-stationary kernel instead of a stationary one, because the solution exhibits discontinuity that cannot be well captured by stationary kernel. Additionally, the NN with four hidden layers of 40-unit width and the hyperbolic tangent nonlinearity is employed in the NN approach.

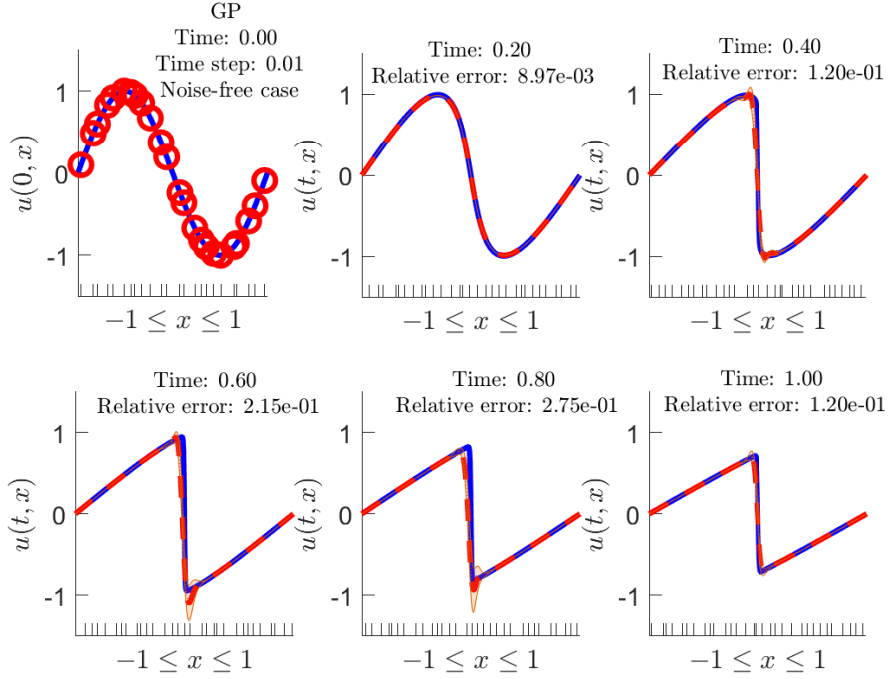


Figure 13: Solving 1D Burgers' equation with noise-free initial condition: GP with kernel (11). The blue solid curve is the exact solution computed according to [17], the red circle is the training points at $n = 0$ (namely $N^0 = 24$), and the red dashed curve is the numerical solution. The x-axis ticks denote the locations of 31 randomly sampled training points at $n \geq 1$ ($N^n = 31$). The orange shadowed region is the two-standard-deviation band.

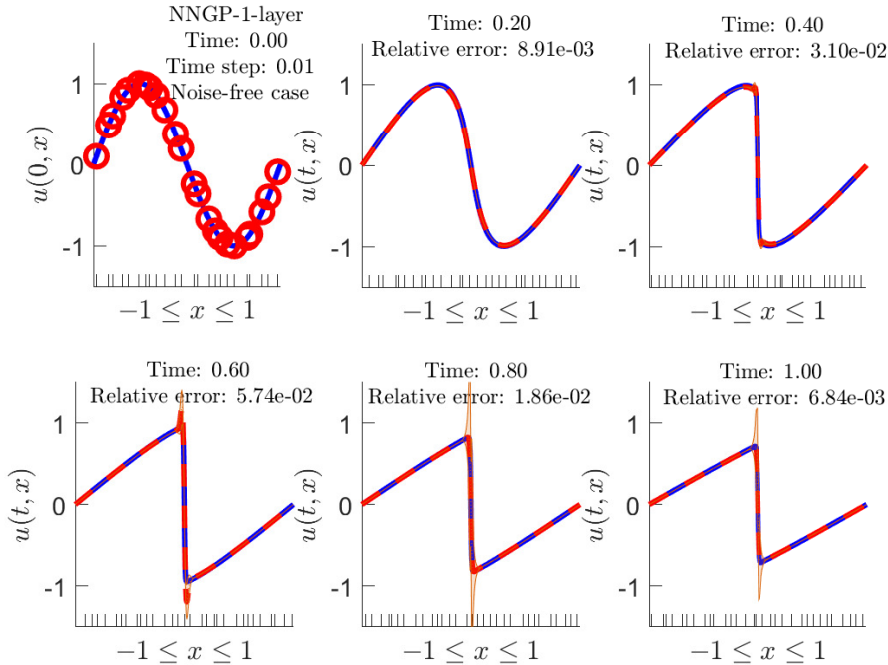


Figure 14: Solving 1D Burgers' equation with noise-free initial condition: NNGP with one-layer, error-function induced kernel. Number of training points at $n = 0$ is $N^0 = 24$ and for $n \geq 1$ is $N^n = 31$.

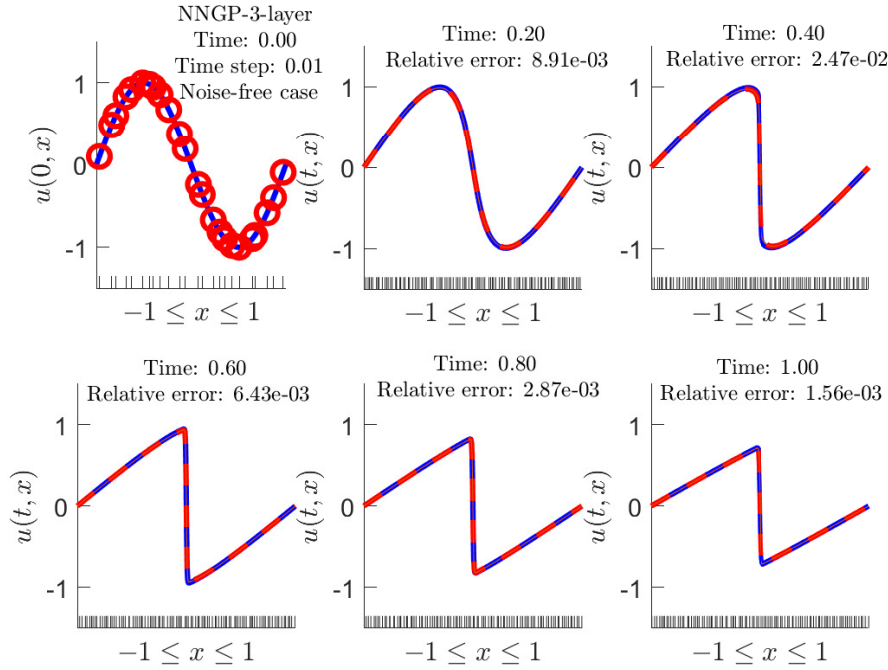


Figure 15: Solving 1D Burgers' equation with noise-free initial condition: NNGP with three-layer, error-function induced kernel. Number of training points at $n = 0$ is $N^0 = 24$ and at $n \geq 1$ is $N^n = 101$.

We first consider the case where the initial condition is noise-free. Figs. 13, 14, and 15 show the numerical solutions computed by the GP, the NNGP with one-hidden-layer ($L = 1$), and the NNGP with three-hidden-layer ($L = 3$), respectively. We can see that the NNGP has higher solution accuracy than the GP. Importantly, different from previous examples, increasing the depth of inducing NN in NNGP improves the results in the current example. The high accuracy of NNGP is attributed to a larger number of hyperparameters compared to the GP case. More hyperparameters presumably implies higher expressivity for function approximation, which means that we can use the kernel to approximate a wider spectrum of functions. For simple function or solution (as in the Hartmann 3D function and Poisson equation examples), the advantage of higher expressivity is not fully realized, but for complex ones, such as the step function and the Burgers' equation here, the advantage can be clearly seen.

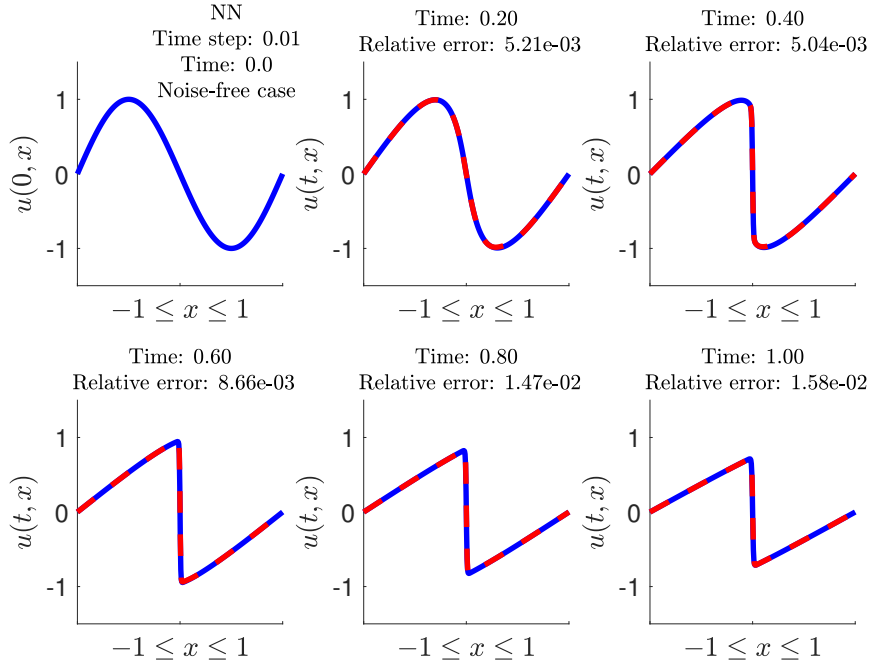


Figure 16: Solving 1D Burgers' equation with noise-free initial condition: NN with four hidden layers of 40-unit width as well as hyperbolic tangent nonlinearity. The maximum error appears near the discontinuity point.

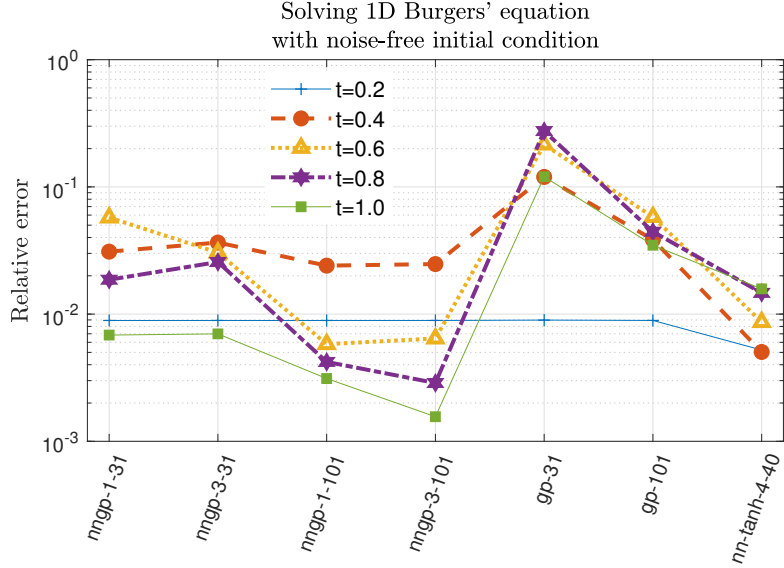


Figure 17: Solving 1D Burgers' equation with noise-free initial condition: Accuracy comparison of NN, NNGP, and GP. “nngp-1” means the NNGP with one hidden layer and error-function induced kernel. “31” and “101” denote the numbers of training points $N^n = 31$ and 101 for $n \geq 1$, respectively. It is fixed that $N^0 = 24$ for all the NNGP examples. Kernel (11) is adopted in the GP. “nn-tanh-4-40” represents the NN with four hidden layers of 40-unit width as well as hyperbolic tangent nonlinearity.

The NN results are shown in Fig.16. Although NN can derive sufficiently accurate solutions, it does not give any uncertainty estimate. Fig. 17 collects the errors from GP, NNGP, and NN for noise-free case. The NNGP and NN achieve errors of the same order of magnitude, while the GP performs worst. This is because NNGP and NN both have a large set of parameters to be tuned and thus possess higher expressivity.

Next we consider the case where the initial condition is contaminated by Gaussian white noise of zero mean and variance 0.15^2 . The NN approach for the noise case is beyond the scope of the present paper, since without proper regularization methods (such as dropout [18], early stopping, and weighted L1/L2 penalty), NN will easily encounter overfitting. Unlike NN, GP and NNGP inherently include the model complexity penalty in the negative log-marginal likelihood and have less risk in overfitting (see the discussion in Section 5.4.1 of the book [12]).

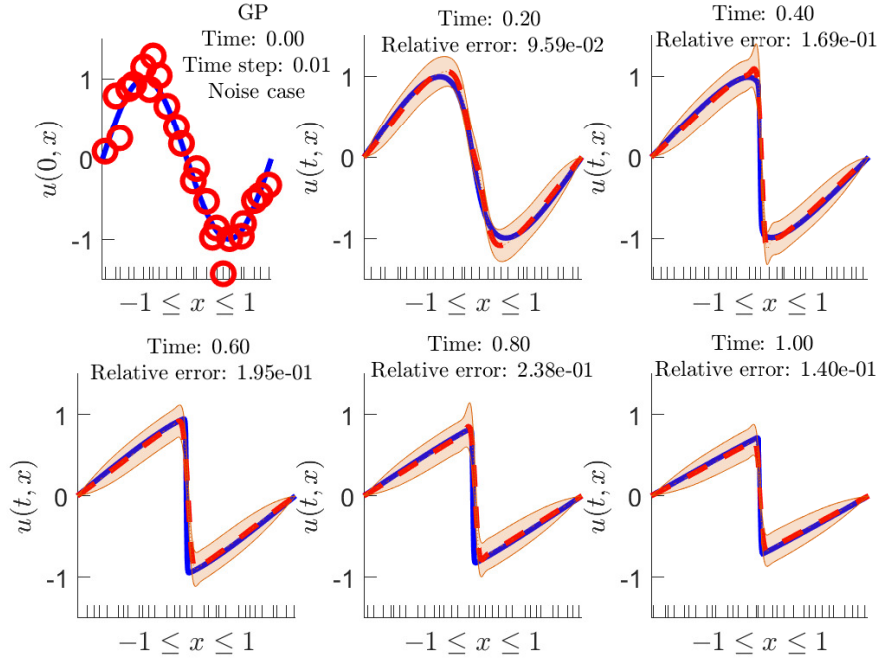


Figure 18: Solving 1D Burgers' equation with noisy initial condition (noise variance $\sigma_{noise}^2 = 0.15^2$): GP with kernel (11). Number of training points at $n = 0$ is $N^0 = 24$ and at $n \geq 1$ is $N^n = 31$.

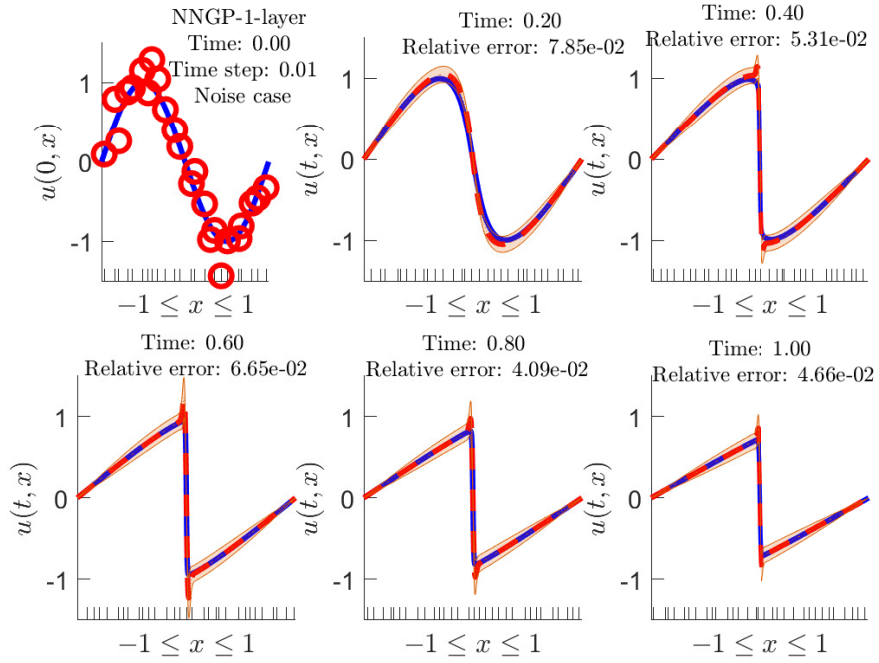


Figure 19: Solving 1D Burgers' equation with noisy initial condition (noise variance $\sigma_{noise}^2 = 0.15^2$): NNGP with the one hidden layer and error-function induced kernel. Number of training points at $n = 0$ is $N^0 = 24$ and at $n \geq 1$ is $N^n = 31$.

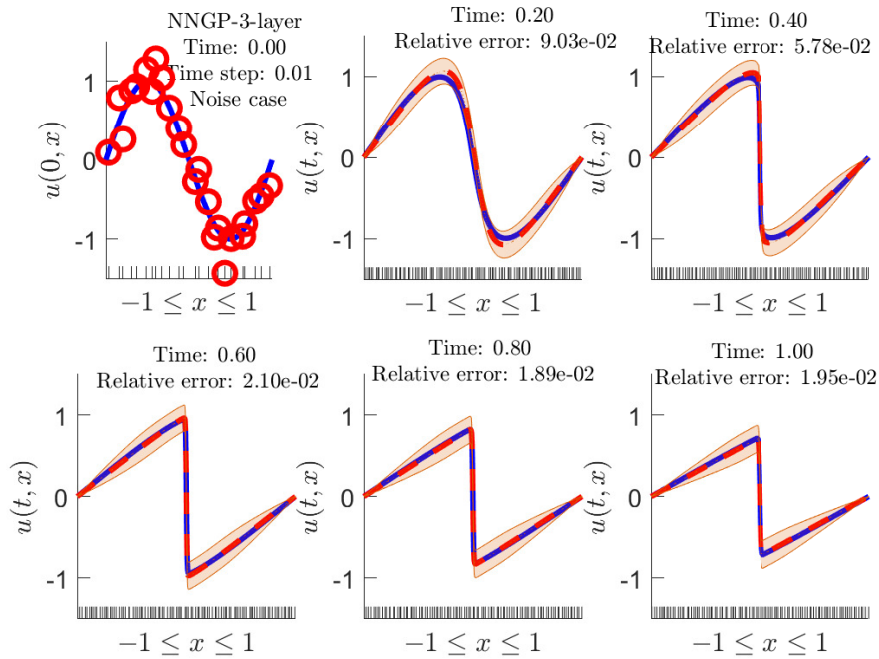


Figure 20: Solving 1D Burgers' equation with noisy initial condition (noise variance $\sigma_{noise}^2 = 0.15^2$): NNGP with the three hidden layers and error-function induced kernel. Number of training points at $n = 0$ is $N^0 = 24$ and at $n \geq 1$ is $N^n = 101$.

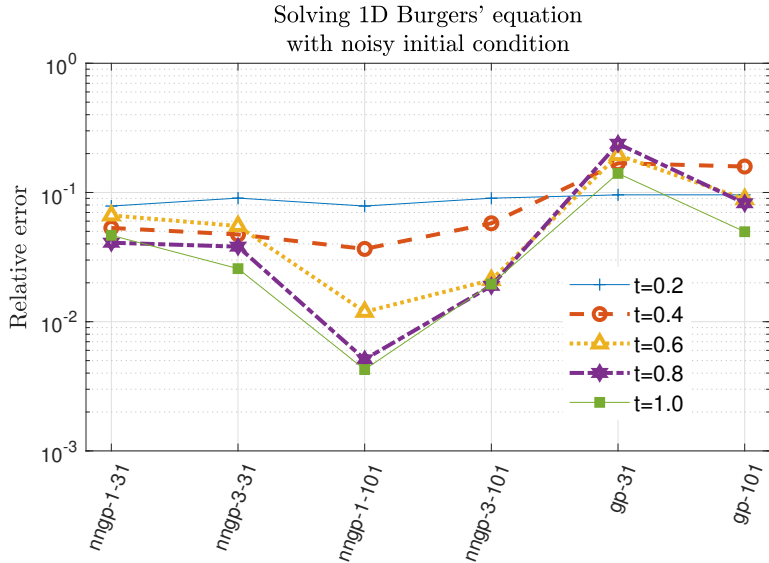


Figure 21: Solving 1D Burgers' equation with noisy initial condition (noise variance $\sigma_n^2 = 0.15^2$): Accuracy comparison of GP and NNGP.

Numerical solutions computed by the GP and NNGP are plotted in Figs. 18, 19, and 20. NNGP still has higher solution accuracy than the GP due to the higher expressivity. Data noise amplifies uncertainty represented by the orange shadowed region. GP and NNGP can handle noise well, because noise variance can be directly learned from the training data and the corresponding uncertainty is quantified by the conditional (or posterior) distribution. Fig. 21 compares the solution accuracy of GP and NNGP. For long-term simulations, the accuracy of the NNGP is roughly one order higher than that of GP.

For NNGP, increasing the depth, L , of inducing NN does not guarantee the increase of accuracy. For example, one-layer NNGP with 101 training points outperforms three-layer NNGP with 101 training points, as is shown in Fig. 21. A larger L amounts to a larger number of hyperparameters, which merely increases the possibility in fitting complex functions better. Actually, L can also be seen as another hyperparameter of NNGP.

6 Summary

A larger number of hyperparameters enables NNGP to achieve higher or comparable accuracy for both smooth and non-smooth functions in comparison with GP, which can be seen from Table 1. The deep NN that induces NNGP contributes its prior variances of network parameters (weights and biases) as well as its depth to the hyperparameter list of NNGP, which endows NNGP with high expressivity. On the other hand, NNGP is able to estimate uncertainty of predictions, which is crucial to noisy-data handling and active learning [19]. Unlike NN, Bayesian NN can provide uncertainty estimate. However, the conventional Bayesian NN [20] could be time-consuming to train due to approximation of a high-dimensional integral over weight space.

Table 1: Comparison of GP, NNGP, and NN (Accuracy) for function approximation and PDE solution.

Function/Solution	Feature	Accuracy comparison	Kernel/nonlinearity		
			GP	NNGP	NN
Step function	Non-smooth	NNGP > {GP, NN}	SE/Matern	erf/ReLU	erf/ReLU
Hartmann 3D	Smooth	{NNGP, GP} > NN	SE/Matern	erf/ReLU	erf/ReLU
Poisson eq.	Smooth	{NNGP, GP} > NN	SE	erf	erf/tanh
Burgers' eq.	Non-smooth	{NNGP, NN} > GP	kernel (11)	erf	tanh

Due to the need for inverting the covariance matrix, NNGP has cubic time complexity for training (see Table 2). In Table 2, for GP and NNGP $N_{training}$ is the size of training set, and m_{GP} and m_{NNGP} are numbers of evaluations of negative log marginal-likelihood in conjugate gradient algorithm for GP and NNGP, respectively. For NN, the accurate estimate of time complexity for training is still an open question [21]. Generally, the training of a fully-connected NN is faster than that of GP, because one does not need to invert a matrix. For each training point, the forward and backward propagation only need linear cost, namely $O(N_{weight})$, where N_{weight} is the total number of weights in the network. $N_{training}$ for NN means batch size; in this paper we fed the whole training set to optimization algorithm and thus the batch size is exactly the size of training set. m_{NN} is number of iterations in optimization algorithm. In numerical examples of this paper, training set size does not exceed 1000. However, for very large datasets, GP and NNGP will be less attractive compared to NN. In future work, we intend to leverage scalable GPs recently developed in [22, 23, 24] to tackle large dataset problems.

Table 2: Comparison of GP, NNGP, and NN (whether to estimate uncertainty and the computational cost for training) for function approximation and PDE solution.

	GP	NNGP	NN
Uncertainty	✓	✓	×
Cost	$O(m_{GP}N_{training}^3)$	$O(m_{NNGP}N_{training}^3)$	$O(m_{NN}N_{weight}N_{training})$

Acknowledgements

We thank Mr. Dongkun Zhang and Dr. Maziar Raissi for useful discussions. This work was supported by AFOSR (FA9550-17-1-0013) and NSF (DMS-1736088). The first author was also supported by the National Natural Science Foundation of China (11701025). The second author was also supported by a gift from Swiss company BH-Robotics.

References

- [1] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *Advances in neural information processing systems*, pages 3360–3368, 2016.
- [2] Radford M Neal. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995.
- [3] Christopher KI Williams. Computing with infinite networks. In *Advances in neural information processing systems*, pages 295–301, 1997.

- [4] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- [5] Haim Sompolinsky, Andrea Crisanti, and Hans-Jurgen Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259, 1988.
- [6] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Inferring solutions of differential equations using noisy multi-fidelity data. *Journal of Computational Physics*, 335:736–746, 2017.
- [7] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Numerical gaussian processes for time-dependent and non-linear partial differential equations. *arXiv preprint arXiv:1703.10230*, 2017.
- [8] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017.
- [9] Guofei Pang, Paris Perdikaris, Wei Cai, and George Em Karniadakis. Discovering variable fractional orders of advection–dispersion equations from field data using multi-fidelity bayesian optimization. *Journal of Computational Physics*, 348:694–714, 2017.
- [10] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Machine learning of linear differential equations using gaussian processes. *Journal of Computational Physics*, 348:683–693, 2017.
- [11] Youngmin Cho and Lawrence K Saul. Kernel methods for deep learning. In *Advances in neural information processing systems*, pages 342–350, 2009.
- [12] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian processes for machine learning*. the MIT press, 2006.
- [13] Loic Le Gratiet. *Multi-fidelity Gaussian process regression for computer experiments*. PhD thesis, Université Paris-Diderot-Paris VII, 2013.
- [14] Carl Edward Rasmussen and Hannes Nickisch. Gaussian processes for machine learning (gpml) toolbox. *Journal of Machine Learning Research*, 11(Nov):3011–3015, 2010.
- [15] Ladislav Kocis and William J Whiten. Computational investigations of low-discrepancy sequences. *ACM Transactions on Mathematical Software (TOMS)*, 23(2):266–294, 1997.
- [16] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010.
- [17] Cea Basdevant, M Deville, P Haldenwang, JM Lacroix, J Ouazzani, R Peyret, P Orlandi, and AT Patera. Spectral and finite difference solutions of the burgers equation. *Computers & fluids*, 14(1):23–41, 1986.
- [18] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [19] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- [20] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [21] Le Song, Santosh Vempala, John Wilmes, and Bo Xie. On the complexity of learning neural networks. *arXiv preprint arXiv: 1707.04615*, 2017.

- [22] James Hensman, Nicolo Fusi, and Neil D Lawrence. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- [23] Sivaram Ambikasaran, Daniel Foreman-Mackey, Leslie Greengard, David W Hogg, and Michael O’Neil. Fast direct methods for gaussian processes. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):252–265, 2016.
- [24] Alexander Litvinenko, Ying Sun, Marc G Genton, and David Keyes. Likelihood approximation with hierarchical matrices for large spatial datasets. *arXiv preprint arXiv:1709.04419*, 2017.

Appendix: Derivatives of NNGP kernel from the error-function nonlinearity

Absorbing the constant coefficients $\frac{2}{\pi}$ and 2 into the variances, the kernel derived from the error-function nonlinearity, namely (10), can be further simplified to

$$k^l(\mathbf{x}, \mathbf{x}') = \sigma_{w,l}^2 \sin^{-1} \left(\frac{k^l(\mathbf{x}, \mathbf{x}')}{\sqrt{(1 + k^{l-1}(\mathbf{x}, \mathbf{x})) (1 + k^{l-1}(\mathbf{x}', \mathbf{x}'))}} \right) + \sigma_{b,l}^2, l = 1, 2, \dots, L, \quad (\text{A1})$$

$$k^0(\mathbf{x}, \mathbf{x}') = \mathbf{x}^T \Lambda \mathbf{x}' + \sigma_{b,0}^2.$$

To solve PDEs, we need to compute the derivatives of the kernel $k^l(\mathbf{x}, \mathbf{x}')$. We take the 2D Poisson equation as an example. We need to know the explicit forms of $-\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)k^l(\mathbf{x}, \mathbf{x}')$ and $\left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}\right)\left(\frac{\partial^2}{\partial x'^2} + \frac{\partial^2}{\partial y'^2}\right)k^l(\mathbf{x}, \mathbf{x}')$ where $\mathbf{x} = [x, y]^T$ and $\mathbf{x}' = [x', y']^T$. The partial derivatives up to fourth order needs to be derived. Denoting by the trivariate function $F_\phi(\cdot)$ the arcsin() term in the above iteration formula, the use of chain rule yields

$$\begin{aligned} \frac{\partial k^l(\mathbf{x}, \mathbf{x}')}{\partial x} &= \sigma_{w,l}^2 \left[\frac{\partial F_\phi}{\partial k^{l-1}(\mathbf{x}, \mathbf{x})} \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x})}{\partial x} + \frac{\partial F_\phi}{\partial k^{l-1}(\mathbf{x}, \mathbf{x}')} \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x}')}{\partial x} + \frac{\partial F_\phi}{\partial k^{l-1}(\mathbf{x}', \mathbf{x}')} \times 0 \right], \\ \frac{\partial^2 k^l(\mathbf{x}, \mathbf{x}')}{\partial x^2} &= \sigma_{w,l}^2 \left[\left(\frac{\partial^2 F_\phi}{\partial^2 k^{l-1}(\mathbf{x}, \mathbf{x})} \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x})}{\partial x} + \frac{\partial^2 F_\phi}{\partial k^{l-1}(\mathbf{x}, \mathbf{x}') \partial k^{l-1}(\mathbf{x}, \mathbf{x})} \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x}')}{\partial x} \right) \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x})}{\partial x} \right. \\ &\quad + \frac{\partial F_\phi}{\partial k^{l-1}(\mathbf{x}, \mathbf{x})} \frac{\partial^2 k^{l-1}(\mathbf{x}, \mathbf{x})}{\partial x^2} \\ &\quad + \left(\frac{\partial^2 F_\phi}{\partial k^{l-1}(\mathbf{x}, \mathbf{x}) \partial k^{l-1}(\mathbf{x}, \mathbf{x}')} \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x})}{\partial x} + \frac{\partial^2 F_\phi}{\partial^2 k^{l-1}(\mathbf{x}, \mathbf{x}')} \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x}')}{\partial x} \right) \frac{\partial k^{l-1}(\mathbf{x}, \mathbf{x}')}{\partial x} \\ &\quad \left. + \frac{\partial F_\phi}{\partial k^{l-1}(\mathbf{x}, \mathbf{x}')} \frac{\partial^2 k^{l-1}(\mathbf{x}, \mathbf{x}')}{\partial x^2} \right] \\ &\dots\dots\dots \end{aligned} \quad (\text{A2})$$

Generally, the iteration formulas for kernel and its derivatives can be written as

$$\frac{\partial^{i+j} k^l(\mathbf{x}, \mathbf{x}')}{\partial x^i \partial x'^j} = F_{i,j} \left(\frac{\partial^{i'+j'} k^{l-1}(\mathbf{x}, \mathbf{x}')}{\partial x^{i'} \partial x'^{j'}}, \frac{\partial^{i'+j'} k^{l-1}(\mathbf{x}, \mathbf{x})}{\partial x^{i'} \partial x'^{j'}}, \frac{\partial^{i'+j'} k^{l-1}(\mathbf{x}', \mathbf{x}')}{\partial x^{i'} \partial x'^{j'}}, 0 \leq i' \leq i, 0 \leq j' \leq j \right), \quad (\text{A3})$$

where $i, j \in 0, 1, 2$ and $F_{0,0} = F_\phi$. The initial values for the iteration are the known kernel k^0 and its derivatives. It should be noted that the explicit form of the function $F_{i,j}$ will become rather lengthy and ugly for higher order derivatives, but fortunately, we have derived for the readers these lengthy formulas using Maple. The Matlab code computing the covariance functions and its partial derivatives up to fourth order can be downloaded at https://github.com/Pang1987/nngp_kernel_derivative.