

Projection pursuit Gaussian process regression

Gecheng Chen  and Rui Tuo

Wm Michael Barnes '64 Department of Industrial and Systems Engineering, Texas A&M University, College Station, TX, USA

ABSTRACT

A primary goal of computer experiments is to reconstruct the function given by the computer code via scattered evaluations. Traditional isotropic Gaussian process models suffer from the curse of dimensionality, when the input dimension is relatively high given limited data points. Gaussian process models with additive correlation functions are scalable to dimensionality, but they are more restrictive as they only work for additive functions. In this work, we consider a projection pursuit model, in which the nonparametric part is driven by an additive Gaussian process regression. We choose the dimension of the additive function higher than the original input dimension, and call this strategy "dimension expansion". We show that dimension expansion can help approximate more complex functions. A gradient descent algorithm is proposed for model training based on the maximum likelihood estimation. Simulation studies show that the proposed method outperforms the traditional Gaussian process models.

ARTICLE HISTORY

Received 8 April 2022
Accepted 30 August 2022

KEYWORDS

Computer experiments;
surrogate modeling;
additive Gaussian process;
neural networks;
dimension expansion

1. Introduction

Contemporary practices in engineering and physical sciences have made increasing use of (deterministic) computer simulations, in disciplines including aerospace designs, material science, and biomedical studies. One of the central research topics is to build an accurate surrogate model to emulate computer simulations. Gaussian process regression (Rasmussen and Williams, 2006; Santner *et al.*, 2003) is one of the most popular surrogate models. Various modifications and extensions of the standard Gaussian process regression models have been proposed to address the specific needs in practical situations. An incomplete list of these methods includes composite Gaussian processes (Ba and Joseph, 2012), treed Gaussian processes (Gramacy and Lee, 2008), non-stationary models (Heaton *et al.*, 2017), transformed approximately additive Gaussian processes (Lin and Roshan Joseph, 2020), etc.

Data analysis for computer simulations usually suffers from the "small data" issue, because the computer simulation runs can be highly costly. For example, each run of a typical computational fluid dynamics model for aerospace engineering takes a few days or even weeks to run (Mak *et al.*, 2018). Many computer simulations also pose the *curse of dimensionality* problem, in the sense that the input dimension is relatively high so that building an accurate surrogate model based on limited data points becomes more challenging. Classic approaches for dimension reduction in computer experiments include sensitivity analysis (Oakley and O'Hagan, 2004; Saltelli *et al.*, 2010; Durrande *et al.*, 2013), ridge approximation (Pinkus, 1997; Hokanson and Constantine, 2018; Glaws *et al.*, 2020). Variable selection for Gaussian processes models

is considered in Linkletter *et al.* (2006), Constantine *et al.* (2014) and Gu (2019). In Gaussian process regression, it is also known that some correlation structures perform better in high-dimensional scenarios (Stein, 1999). Recently, additive Gaussian process models have received considerable attention (Lebarbier, 2005; Duvenaud *et al.*, 2011; Durrande *et al.*, 2012; Tripathy *et al.*, 2016; Deng *et al.*, 2017; Delbridge *et al.*, 2020). Although these models are more scalable to the input dimension, their capability of model fitting is lower because these models can only reconstruct additive functions precisely.

In this work, we propose a novel surrogate modeling technique based on the projection pursuit methodology (Friedman and Stuetzle, 1981) and additive Gaussian process models. Gaussian Process Regression (GPR) can provide prediction variance as opposed to projection pursuit (neural networks). Additionally, unlike the conventional estimation approaches for projection pursuit (Ferraty *et al.*, 2013; Gilboa *et al.*, 2013; Li *et al.*, 2016), we suggest choosing a large number of intermediate nodes to introduce more model flexibility. Then we use the Maximum Likelihood (ML) estimation to identify the model parameters. A gradient descent algorithm is proposed to search the maximum of the likelihood function. In this work, we also find an error bound of the prediction error for Gaussian process regression with additive Matérn correlation functions. Our theoretical results show that the prediction error of additive Gaussian process models is much lower than that given by isotropic Gaussian process models for high-dimensional problems, provided that a design with nice projection properties, such as a Latin hypercube design, is adopted.

This article is organized as follows. In [Section 2](#), we review the background of GPR with isotropic and additive Matérn correlation functions. In [Section 3](#), we introduce the proposed methodology, called the Projection Pursuit Gaussian Process Regression (PPGPR). An algorithm of the proposed method is given at the end of [Section 3](#). In [Sections 4](#) and [5](#), we conduct simulation studies to demonstrate the use of the proposed method, and show that the proposed method outperforms some existing methods. In [Section 6](#), we show that the performance of the proposed method is satisfactory through a real-world application. Concluding remarks are made in [Section 7](#).

2. Review on GPR

In this section, we review a simple version of the Gaussian process emulation ([Santner et al., 2003](#)). Let Z be a stationary Gaussian process on \mathbb{R}^d with mean zero, variance σ^2 , and correlation function Φ . Given scattered evaluations $(x_1, Z(x_1)), \dots, (x_n, Z(x_n))$, one can reconstruct Z using its conditional expectation

$$\hat{Z}(x) := \mathbb{E}(Z(x)|Z(x_1), \dots, Z(x_n)) = r^T(x)K^{-1}Y, \quad (1)$$

for $x \in \mathbb{R}^d$, where $r(x) := (\Phi(x - x_1), \dots, \Phi(x - x_n))^T$, $K = (\Phi(x_j - x_k))_{jk}$ for $j = 1, \dots, n$ and $k = 1, \dots, n$, and $Y = (Z(x_1), \dots, Z(x_n))^T$.

2.1. Curse of dimensionality in GPR with isotropic Matérn correlation

The curse of dimensionality is one of the fundamental challenges in various high-dimensional statistical and machine learning problems. In this section, we review how the curse of dimensionality can affect the prediction performance of GPR.

The prediction error of the GPR is

$$Z(x) - \hat{Z}(x) = Z(x) - \mathbb{E}(Z(x)|Z(x_1), \dots, Z(x_n)),$$

which is a function of x . [Tuo and Wang \(2020\)](#) study the rate of convergence of the prediction error under different function norms, under the assumption that the Gaussian process has an isotropic Matérn correlation function ([Santner et al., 2003](#)), defined as

$$\Phi(x; \nu, \phi) = \frac{1}{\Gamma(\nu)2^{\nu-1}} (2\sqrt{\nu}\phi\|x\|)^\nu K_\nu(2\sqrt{\nu}\phi\|x\|), \quad (2)$$

where $\nu > 0$ is the *smoothness parameter*, K_ν is the modified Bessel function of the second kind, $\phi > 0$ is the scale parameter.

To explain the curse of dimensionality issue posed by the isotropic Matérn correlation functions, we refer to [Theorem 3.3 of Tuo and Wang \(2020\)](#), which states a lower bound of the maximum of the prediction error of an isotropic Gaussian process. For simplicity, we consider the expected

maximum prediction error. Suppose the input region of interest is Ω , and then the expected maximum prediction error is $\mathbb{E} \sup_{x \in \Omega} |Z(x) - \hat{Z}(x)|$. Here the expectation is taken over the randomness of the Gaussian process $Z(\cdot)$. [Theorem 3.3 of Tuo and Wang \(2020\)](#) implies

$$\mathbb{E} \sup_{x \in \Omega} |Z(x) - \hat{Z}(x)| \geq C\sigma n^{-\nu/d} \sqrt{\log n}, \quad (3)$$

for a constant C independent of n , σ and the choice of the experimental design.

The lower bound in (3) shows that the uniform error of a GPR predictor with an isotropic Matérn correlation is no less than a multiple of $n^{-\nu/d} \sqrt{\log n}$. This rate grows dramatically as d increases with a fixed ν . Therefore, when a Gaussian process model with an isotropic Matérn correlation is considered, its prediction suffers from the curse of dimensionality, in the sense that, for a high-dimensional problem, acquiring extra data points cannot improve the prediction accuracy as effectively as in lower-dimensional problems.

In GPR, the curse of dimensionality is inevitable if the underlying function is indeed a realization of a Gaussian process with isotropic Matérn correlation. The reason behind this is that the reproducing kernel Hilbert spaces generated by these correlation functions are too large in high-dimensional circumstances. Fortunately, in most real applications, we confront much “simpler” high-dimensional functions. These functions admit a certain “sparse representation”, and therefore, at least theoretically, can be recovered at a much higher rate of convergence. In [Section 2.2](#), we examine a special and simple structure of this kind.

2.2. Additive models: Accuracy and limitations

A scalable GPR approach proceeds by equipping an additive correlation function. Denote $x = (x_{(1)}, \dots, x_{(d)})$. We consider the following function:

$$\Phi(x) = \frac{1}{d} \sum_{j=1}^d \Phi_1(x_{(j)}), \quad (4)$$

where Φ_1 denotes a one-dimensional correlation function. It is easily seen that Φ is positive definite if Φ_1 is positive definite. Thus, one can consider Gaussian process models with correlation (4). This approach is called the additive Gaussian process regression ([Leadbetter, 2005; Duvenaud et al., 2011; Deng et al., 2017](#)).

Compared with isotropic models, additive models are much more scalable to the dimensionality. It can be shown that the rate of convergence of the uniform error is independent of d . Specifically, if Φ_1 is a Matérn correlation function with smoothness ν , the uniform prediction error in (3) can have a rate of convergence $O(n^{-\nu} \sqrt{\log n})$; see our theoretical results in the [Supplementary Materials](#).

Despite the above advantages, the limitations of additive models are also evident. Only additive functions, i.e., the functions that can be decomposed as the sum of functions such that each of them relies on only one entry of x , can be accurately

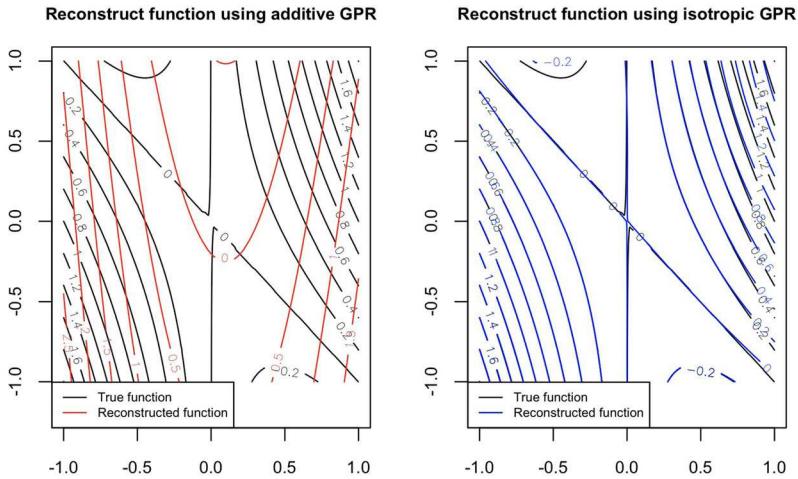


Figure 1. Contour plots of $f(x, y) = xy + x^2$ and the reconstructed functions by additive and isotropic GPR using a same 25-point random design between -1 and 1 . It can be seen that the isotropic model has a much better prediction performance.

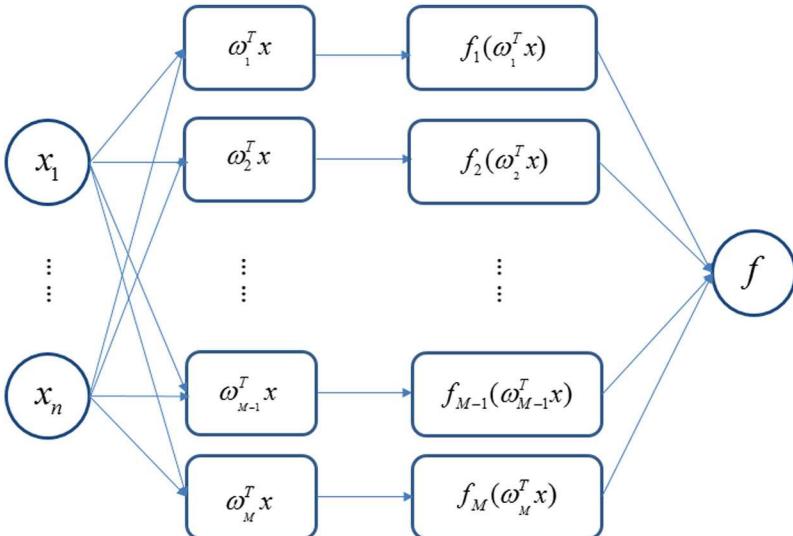


Figure 2. Network structure of PPGPR.

reconstructed. This assumption is *not* true for most of the practical problems. Consider a two-dimensional input (x, y) . A simple non-additive function is $f(x, y) = xy + x^2$. Figure 1 shows that the additive model cannot fit this function well, while the isotropic model works in this case.

3. PPGPR

In this section, we propose a general approach to reconstruct multi-dimensional functions that admits more complicated sparse representations. To this end, we consider a model which is more flexible than additive Gaussian process models. Specifically, we employ the projection pursuit regression method (Friedman and Stuetzle, 1981) to model the underlying function as

$$y(x) = f(w_1^T x, w_2^T x, \dots, w_M^T x), \quad (5)$$

where w_1, \dots, w_M are unknown vectors, M is a positive integer, and f is an additive function in the sense that f can be written as

$$f(w_1^T x, w_2^T x, \dots, w_M^T x) = f_1(w_1^T x) + f_2(w_2^T x) + \dots + f_M(w_M^T x), \quad (6)$$

with unknown univariate functions f_1, \dots, f_M . In other words, this model first applies a linear transformation on the input space, and then use an additive function to fit the responses.

A projection pursuit model can be represented by a four-layer network shown in Figure 2, which is similar to a neural network model. Neural networks have been widely used to enhance the precision of nonparametric regression (Psichogios and Ungar, 1992; Hinton and Salakhutdinov, 2006; LeCun *et al.*, 2015; Goodfellow *et al.*, 2016); Khoo *et al.* (2017) and Tripathy and Bilionis (2018) employ deep neural networks to reduce the dimension of data; Wilson *et al.* (2011) combine neural networks with GPR method to tackle multi-task problems. The main difference between the projection pursuit method and neural networks lies in the activation functions. In neural networks, the activation functions are chosen as a fixed function, such as rectified linear unit (ReLU) functions. In contrast, the projection pursuit

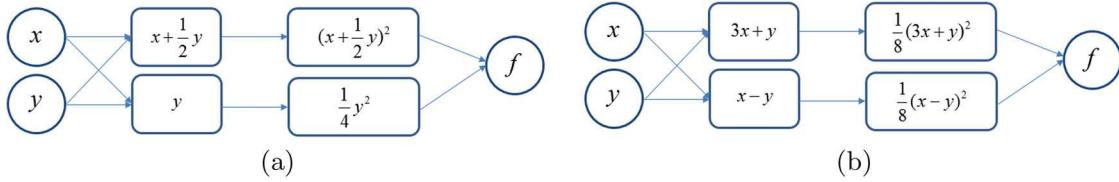


Figure 3. Two different representations of $f(x,y) = xy + x^2$ via projection pursuit.

method uses estimated activation functions. In this work, we call the two hidden layers the *transformation layers*.

When $M=1$, the projection pursuit model reduces to a single index model, which provides a parsimonious way to implement multivariate non-parametric regression. By imposing suitable priors on the parameters, Choi *et al.* (2011), Gramacy and Lian (2012) and Hu *et al.* (2013) use the Bayesian approach to estimate the parameters of the single index model. In Wang *et al.* (2010), a dimension reduction method is applied to choose the number of nodes and then the link function is estimated using GPR. In this work, we consider projection pursuit models with $M \gg 1$, which are much more flexible than single index models.

Given a sufficiently large M , it is known that the projection pursuit model can approximate any continuous function arbitrarily well (Hastie *et al.*, 2009). For example, the non-additive function $f(x,y) = xy + x^2$ can be represented by projection pursuit as shown in Figure 3. Figure 3 also shows that the representation is not unique.

The non-uniqueness of the projected pursuit representation suggests that each of the “directions” w_i may not be essential. In contrast, these vectors exhibit a “synergistic effect”, so that they need to be estimated jointly. Consider the example shown in Figure 3(a). Taking the direction $x + y/2$ along is not helpful in obtaining the underlying function $xy + x^2$; this direction makes sense only when it is paired by the direction y . This phenomenon differs from the classic results in linear models, in which the significant directions (usually defined by the principal components) are fixed, and their importance is ordered by the corresponding eigenvalues.

Understanding this difference between the linear and nonlinear models helps build a better projection pursuit regression model. Traditionally, the projection pursuit method is usually regarded as a dimension reduction approach (Ferraty *et al.*, 2013; Gilboa *et al.*, 2013), and greedy algorithms are usually applied to identify w_i ’s (James and Silverman, 2005; Muller and Yao, 2008; Gilboa *et al.*, 2013). These strategies have the following deficiencies: (i) it is often hard to accurately approximate the underlying functions through dimension reduction ($M \ll d$). For example, the function $f(x,y) = xy + x^2$ cannot be recovered through a one-dimensional factor. (ii) Greedy algorithms, which proceed by picking the current “most significant” direction in each step, cannot perform well when there is no order of importance in the directions, as in the example shown in Figure 3. In this work, we propose a method, which conducts a dimension expansion ($M \geq d$) to improve the approximation power substantially.

When $M \geq d$, the projection pursuit model is in general non-identifiable; see Figure 3 for an example. The learning outcome on w_i ’s are meaningless, and we only focus on the prediction of the underlying response at untried input points. Our numerical experience shows that as long as M is large enough, the prediction performance of the proposed method is not heavily dependent on the specific value of M . We recommend choosing M close to, but slightly less than, the sample size n .

In this work, we propose a novel approach, called the PPGPR. To reconstruct the underlying function, we need to: (i) estimate the weight parameters $w = (w_1, w_2, \dots, w_M)$; (ii) reconstruct the combination function f given w using Gaussian process regression (Rasmussen and Williams, 2006; Santner *et al.*, 2003). Recall that the design matrix is denoted as $X = (x_1, x_2, \dots, x_n)^T$, $x_i \in \mathbb{R}^d$ for $i = 1, 2, \dots, n$, and the response as $Y = (f(x_1), f(x_2), \dots, f(x_n))^T$. Now we employ the idea of GPR to assume that f is a realization of a Gaussian process. Specifically, we assume that the Gaussian process has mean zero and an additive correlation function (4). We believe that the mean zero assumption is not too restrictive because the model is already non-identifiable.

The training of the proposed method proceeds by an iterative approach. First, we choose an initial weight parameter w . Then we compute the initial correlation matrix

$$K_w = \left(\frac{1}{M} \sum_{k=1}^M \Phi(w_k^T (x_i - x_j)) \right)_{ij}$$

based on the initial w . Next, we invoke (1) to reconstruct the underlying function f as

$$\hat{f}(x) = r^T(w^T x)(K_w + \delta I)^{-1} Y, \quad (7)$$

where δ is a nugget term to enhance the numerical stability.

Our goal is to seek for w^* which maximizes the log-likelihood function of GPR (Santner *et al.*, 2003), that is,

$$\min_w l(w) = \min_w (Y^T(K_w + \delta I)^{-1} Y + \log \det(K_w + \delta I)). \quad (8)$$

We refer $l(w)$ to the *model loss*. The gradient of $l(w)$ with respect to w_k is

$$\begin{aligned} \frac{\partial l(w)}{\partial w_k} = & -\frac{1}{M} \sum_{i=1}^n \sum_{j=1}^n (Y^T K_w^{-1} \frac{\partial K_w}{\partial w_k} K_w^{-1} Y \\ & + \text{Tr}(K_w^{-1}))(x_i - x_j)^T, \end{aligned} \quad (9)$$

for $k = 1, 2, \dots, M$. The derivative of the matrix K_w can be computed using the following facts. The derivative of the Matérn correlation function is (Wendland, 2004)

$$\frac{\partial}{\partial x} \Phi(x; \nu, \phi) = -\frac{2\nu\phi^2 x}{\nu-1} \Phi\left(\sqrt{\frac{\nu}{\nu-1}}x; \nu-1, \phi\right).$$

Then the gradient descent method can be applied here to find the minimizer via iteratively updating

$$w_k \leftarrow w_k - \eta \frac{\partial l(w)}{\partial w_k},$$

where η is the step length for the gradient descent algorithm, and is referred to as the *learning rate* in the rest of this article.

When the algorithm converges or a stopping criterion is met, one can again reconstruct the underlying function using (7). **Algorithm 1** lists the detailed steps of the proposed training method, each iteration (epoch) includes calculating the gradient for all weights and renewing the weights. To avoid overfitting, an early-stopping criterion (Prechelt, 1998) should be implemented when choosing P (the number of epochs).

In addition to P , there are other hyper-parameters in the proposed methodology, including M, η and the hyper-parameters of the covariance function. We refer the activity of adjusting these parameters to the *tuning process*. Below is a list of our general recommendations for tuning:

- The proposed method does not use the ML estimators (Santner *et al.*, 2003) to estimate the hyper-parameters of the GP covariance because the ML estimators are likely to overfit with relatively small sample sizes (Santner *et al.*, 2003).
- Determining a proper learning rate η through cross-validation such that it maintains a stable training process (i.e., the model loss decreases neither too sharply nor too slowly).
- Increasing the size of representation nodes M until the performance on the testing points starts to deteriorate. In practice, we recommend considering M in the range $[4d, 8d]$ in a d -dimensional problem.
- Adopting early stopping policies (Prechelt, 1998) in the training process when choosing P to avoid overfitting.
- Using cross-validation to choose the hyper-parameters of the covariance function.

More discussion regarding the tuning process is provided through a numerical study in Section 4.1.

Algorithm 1 Training steps for transformation weight w

Input: design matrix $X = (x_1, x_2, \dots, x_n)$, response $Y = (y_1, y_2, \dots, y_n)$, initialized weight $w = (w_1, w_2, \dots, w_M)$, correlation function Φ , learning rate η , number of iterations P

Output: transformation weight w

for p in $1 : P$ do

2: $X' \leftarrow w^T X$

$K_w \leftarrow \Phi(X', X')$

4: for k in $1 : M$ do

$$grad_k \leftarrow -\frac{1}{M} \sum_{i=1}^n \sum_{j=1}^n \left(Y^T K_w^{-1} \frac{\partial K_w}{\partial w_k} K_w^{-1} Y + Tr(K_w^{-1}) \right) (x_i - x_j)^T$$

6: $w_k \leftarrow w_k - \eta \cdot grad_k$

end for

8: end for

4. Simulation studies

In this section, we examine the performance of the proposed method via simulation studies. Based on four numerical experiments, we will provide some guidelines for parameter tuning for PPGPR in Section 4.1. In Section 4.2, we compare the proposed method with some other prevailing algorithms and show the advantages of the proposed method.

4.1. Choice of tuning parameters

In this section, we study how the choice of the hyper-parameters of PPGPR can affect its prediction performance. Recall that the hyper-parameters include the learning rate η , the size of nodes M in the transformation layers, the number of epochs (iterations) P , the choice of the correlation function (Matérn or Gaussian) and smoothness parameter ν if a Matérn correlation is used.

In the rest of this subsection, we will use the Borehole function (Harper and Gupta, 1983) as the test function to study the performance of the proposed PPGPR under different choices of hyper-parameters. The Borehole function is defined as

$$y = \frac{2\pi T_u (H_u - H_l)}{\log\left(\frac{r}{r_w}\right) \left[1 + \frac{T_u}{T_l} + \frac{2LT_u}{\log\left(\frac{r}{r_w}\right) r_w^2 K_w} \right]},$$

with the ranges for the eight variables given by $r_w \in (0.05, 0.15)$, $r \in (100, 50000)$, $T_u \in (63070, 115600)$, $H_u \in (900, 1110)$, $T_l \in (63.1, 116)$, $H_l \in (700, 820)$, $L \in (1120, 1680)$ and $K_w \in (9855, 12045)$. Halton sequences¹ (Halton, 1964) with 40 samples are used as the training set inputs and 500 random samples are used as the testing set inputs. We consider different choices of the tuning parameters and compare the corresponding prediction performance in terms of the Mean Absolute Percentage Error (MAPE) (Makridakis, 1993):

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{\hat{y}_i - y_i}{y_i} \right|, \quad (10)$$

where $n=500$ is the size of testing samples; \hat{y} and y denote the predictive value and true value of a testing sample, respectively.

The details of the numerical experiments are described in Sections 4.1.1–4.1.3. We choose 10^{-6} as the nugget term of (8) in this section to avoid some numerical instability, see Peng and Wu (2014) for more guidance on choosing nugget terms.

4.1.1. Learning rate η and number of representation nodes M

In this experiment, a Matérn correlation function with $\nu = 2.5$ is used and training epochs $P = 150$. We examine the performance of PPGPR under different learning rates and different node sizes in the transformation layers.

¹Halton sequences are deterministic low discrepancy sequences used to generate points in space for numerical experiments. The Halton sequences can be generated efficiently by the R package SDdraw.

Figure 4 shows the MAPE of PPGPR under different learning rates with respect to the size of representation nodes. It can be seen that when $\eta = 10^{-10}$, the MAPE is much higher than those in the other three situations. For $\eta = 10^{-8}$, the model reaches its best performance when $M = 28$. The models with $M = 35$ have lower MAPE when $\eta = 10^{-7}$ and $\eta = 10^{-9}$. In general, the models with $\eta = 10^{-9}$ perform slightly better and more stably.

According to Hastie *et al.* (2009), the PPGPR model can approximate any continuous functions as $M \rightarrow \infty$ for an appropriate choice of kernel function. The Matérn kernels are within this class because the reproducing kernel Hilbert space generated by any Matérn kernel contains all polynomials. This explains why the performance of PPGPR grows as M increases when M is small. However, when M is above 35, the MAPE becomes worse for most of the curves in Figure 4, which may be due to overfitting because there are too many hidden nodes. In practice, we suggest employing cross-validation to select the optimal M .

Figure 5 shows four curves generated with a common initial w and different learning rates when $M = 35$. Each of them

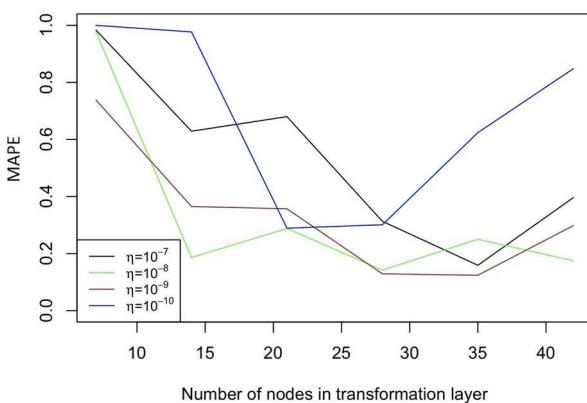


Figure 4. MAPE under different learning rates and size of nodes in transformation layers.

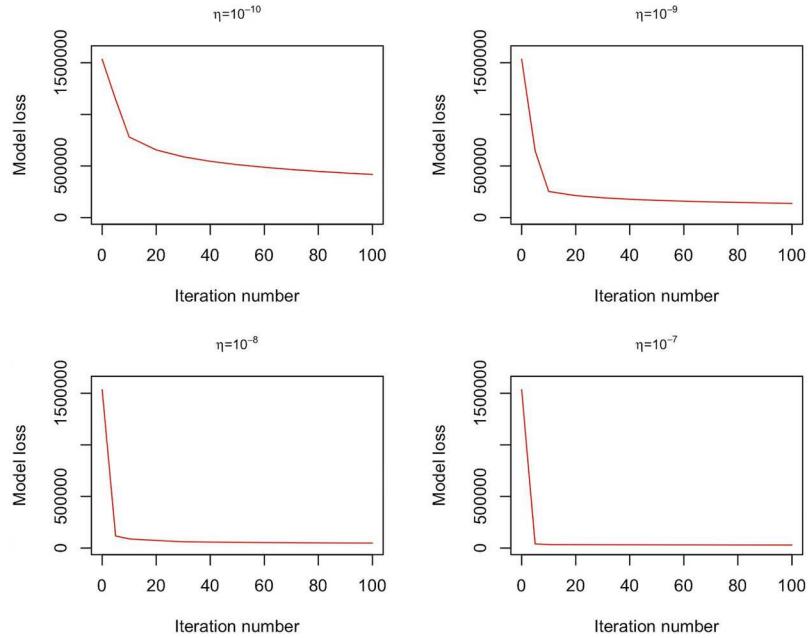


Figure 5. Model loss with different learning rate.

shows the relationship between the model loss defined in (8) and the number of iteration. From Figure 5, we find that, 10^{-10} is too low as a learning rate, because the model loss is still high (about 5×10^5) even after 100 iterations. This observation is also confirmed by the MAPE results in Figure 4, in which the MAPE for $M = 35$ corresponding to $\eta = 10^{-10}$ is much higher than those in the other ones. The model loss curves for the other three learning rates are similar. We believe that the choice of $\eta = 10^{-9}$ gives a slightly better result than those given by $\eta = 10^{-8}$ or $\eta = 10^{-7}$, because the model loss curve under $\eta = 10^{-9}$ decreases more smoothly than the other two, which implies a more stable learning process (Lawrence and Giles, 2000). According to Keskar *et al.* (2016), flat minima might have higher generalization than sharp minima. In addition, a too small model loss after training might result in overfitting which will be shown in Section 4.1.3. Figure 4 also implies that $\eta = 10^{-9}$ gives the best MAPE when $M = 35$. In practice, the optimal learning rate relies on the underlying function. Therefore, we recommend tuning η via cross-validation.

4.1.2. Effects of correlation function type and parameters

In this experiment we examine the performance of PPGPR under different correlation functions and smoothness parameters with $\eta = 10^{-9}$ and $P = 150$.

Figure 6 shows the MAPE for PPGPR with the Matérn correlation functions under different M and ν with $\phi = 1$. It can be seen that when $\nu = 2.5$ (green line), the model performs better than other choices. Under $\nu = 2.5$, the best prediction performance is achieved when $M = 35$. Generally, with a larger ν , the reconstructed function would be smoother, which may lead to overfitting; with a smaller ν , the reconstructed function would be less smooth, which may result in instability or underfitting. Figure 7 shows the MAPE for PPGPR with Gaussian correlation functions under different M and ϕ . We can see that, when $M = 35$,

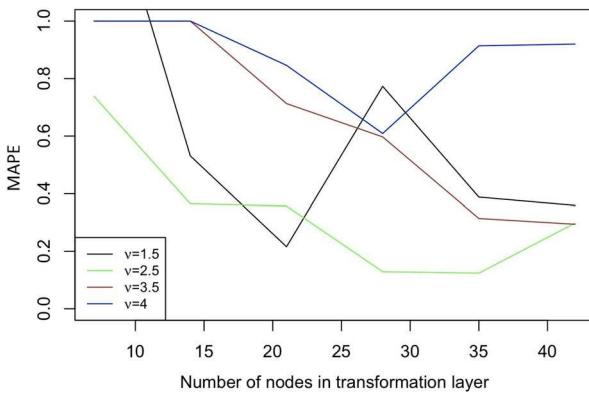


Figure 6. MAPE under different ν and M for Matérn correlation functions.

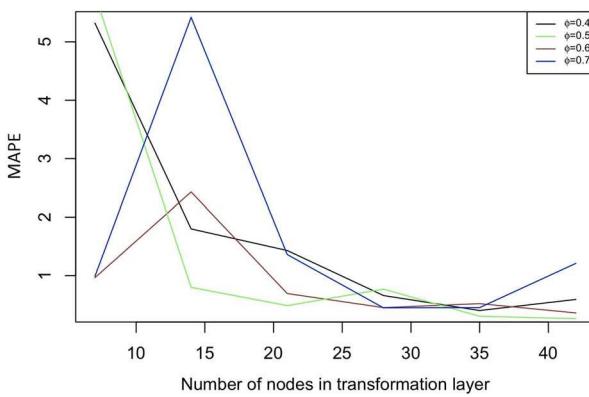


Figure 7. MAPE under different ϕ and M for Gaussian correlation functions.

the green line ($\phi = 0.5$) reaches its lowest MAPE, which is slightly better than the MAPE under other M and ϕ in this experiment. This experiment shows that Matérn correlation functions with $\nu = 2.5$ seem to be an appropriate choice of the correlation functions. We also recommend using cross-validation to determine the optimal correlation function if computational resource permits. Table 1 shows the numerical values of the lowest MAPE of PPGPR under the above Matérn and Gaussian correlation functions.

4.1.3. Training epochs P

In this experiment the model loss and the prediction error of PPGPR during the training process are monitored. Here we use a Matérn correlation function with $\nu = 2.5$ and $\eta = 10^{-8}$, $M = 21$.

Figures 8(a) and 8(b) plot the model loss and prediction error against the training epochs, respectively. We can see from Figure 8(a) that the model loss is monotonically decreasing as M increases. This implies that the proposed gradient descent algorithm works in a desired way. However, Figure 8(b) shows that the prediction error is *not* a monotonic function in the model loss. The model achieves its best performance when $P = 220$, and as P further increases, the prediction error increases. This phenomenon has been observed in other network structures such as neural networks. In a typical neural network training process, a slower early-stopping criterion with 4% (i.e., stopping the training process when the relative generalization improvement is less than 4%) could be used to avoid overfitting caused by an overshot training

Table 1. Best MAPE for PPGPR with Gaussian and Matérn correlation functions.

	M	ϕ	ν	MAPE
Matérn	35	1	2.5	0.124
Gaussian	42	0.5	—	0.263

process (Prechelt, 1998). We suggest adopting a similar approach in training the proposed PPGPR model.

4.2. Numerical comparisons

In this section we compare PPGPR with GPR, Neural Network (NN), SVR (Supporting Vector Regression) and GBDT (Gradient Boosting Decision Trees) using four test functions: OTL circuit function (Ben-Ari and Steinberg, 2007), Borehole function (Harper and Gupta, 1983), Wingweight function (Forrester *et al.*, 2008) and Welch function (Welch *et al.*, 1992). The training set is chosen as Halton series (Halton, 1964) with length $p = 5 \times d$, where d is the dimension of the input space, and the size of testing set is 500. The implementation details of five methods for these experiments are shown below:

- SVR: Matérn correlation with $\nu = 2.5$.
- GBDT: Gaussian distribution and 100 trees.
- NN (deep learning): For the OTL circuit function, it has structure (6, 12, 24, 12, 1) (meaning the node size of input layer is 6, the second layer has 12 nodes and so on) with learning rate 0.01 and 150 epochs. For the Borehole function, it has structure (8, 16, 32, 1) with learning rate 0.01 and 150 epochs. For the Wingweight function, it has structure (10, 20, 30, 20, 1) with learning rate 0.1 and 200 epochs. For the Welch function, it has structure (10, 20, 30, 20, 1) with learning rate 0.1 and 200 epochs.
- GPR (with isotropic and product correlation functions): We use the Dicekriging package (Roustant *et al.*, 2012) with isotropic and product Matérn correlation and smoothness $\nu = 2.5$ to compute the predictive results. The product correlation is defined as $K(x) = \prod_{i=1}^n \Phi_1(x_{(j)})$, where $\Phi_1(x_{(j)})$ is the same as in (4).
- PPGPR: For OTL circuit function, Matérn correlation with $\nu = 2.5$, $M = 42$, $\eta = 10^{-9}$, $P = 150$, for Borehole function, Matérn correlation with $\nu = 2.5$, $M = 35$, $\eta = 10^{-9}$, $P = 150$, for Wingweight function, Matérn correlation with $\nu = 2.5$, $M = 35$, $\eta = 10^{-10}$, $P = 150$, for Welch function, Matérn correlation with $\nu = 2.5$, $M = 7$, $\eta = 10^{-6}$, $P = 200$

The MAPE of each method above is given in Table 2. It can be seen that the performances of SVR and GBDT are inferior in most cases, which can be explained because these approaches may require more training data (Smola and Schölkopf, 2004; Ke *et al.*, 2017). The only exception is the case of the Welch function, where the SVR and the PPGPR result in comparable results. We have tried our best to tune the parameters of the NN, in order to obtain the best achievable results. It is worth noting that the parameter tuning for NN is time-consuming. In contrast, the tuning process of PPGPR is much easier, because it has only one hidden layer. Also, PPGPR outperforms NN in

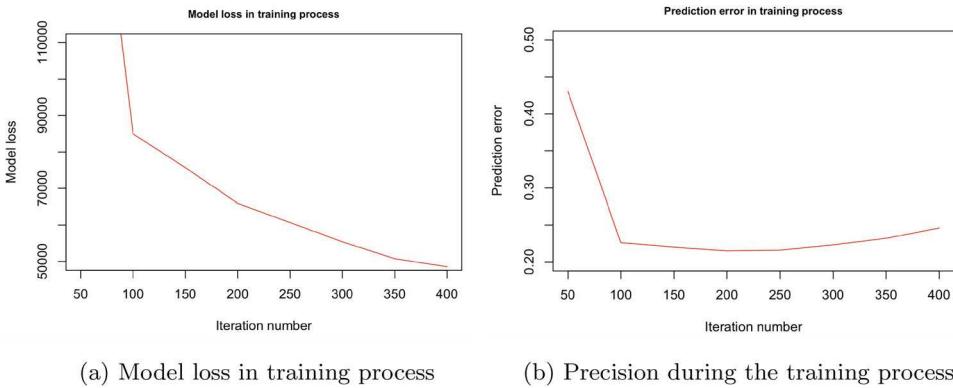


Figure 8. Model loss and precision during the training process.

Table 2. MAPEs of SVR, GBDT, NN, GPR with isotropic and product correlations and PPGPR for three functions. PPGPR outperforms all other methods.

	OTL circuit ($d = 6$)	Borehole ($d = 8$)	Wingweight ($d = 10$)	Welch ($d = 20$)
SVR	0.121	0.792	0.127	0.989
GBDT	0.130	0.407	0.142	1.778
NN	0.0334	0.222	0.240	1.113
GPR(iso)	0.0182	0.204	0.0224	1.334
GPR(pro)	0.0162	0.134	0.0199	1.058
PPGPR	0.0139	0.124	0.0184	0.994

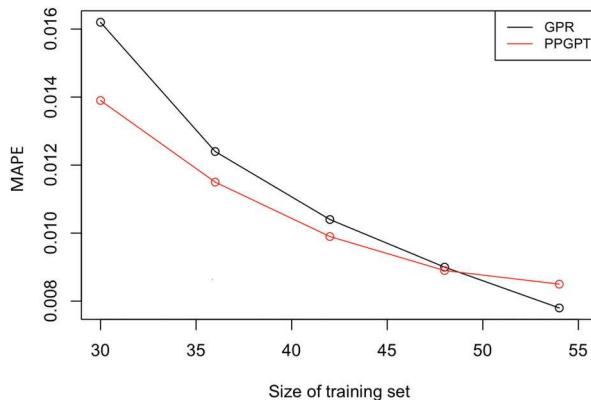


Figure 9. MAPEs of GPR and PPGPR with different size of training set for OTL function.

all three experiments. Moreover, PPGPR can beat GPR with isotropic and product correlation functions because the curse of dimensionality has less impact on PPGPR. Note that GPR with isotropic kernels performs worse than GPR with product kernels. This is not surprising in view of the slow rate of convergence for isotropic kernels shown in Section 2.1. The rate of convergence for product kernels under a general condition is not well-established, but they are known to outperform the isotropic kernels in high-dimensional circumstances (Sacks *et al.*, 1989).

Additionally, we compare the performance of PPGPR and GPR with product kernel when the size of training set changes. Figure 9 shows the MAPEs of the proposed PPGPR and GPR with product kernels for OTL function, when the number of training set varies. It can be seen that when the size of training samples is less than 48 ($8d$) the PPGPR works much better than GPR. When the size of training set increases

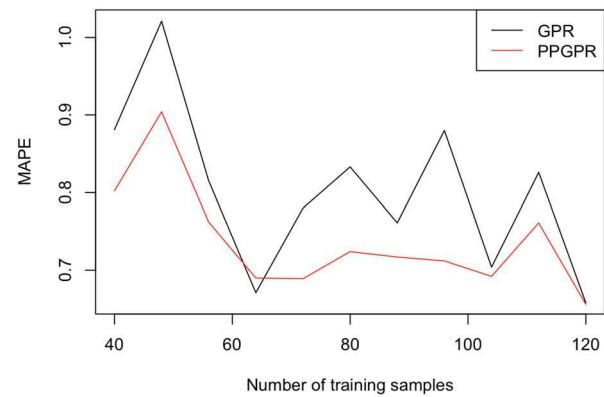


Figure 10. MAPEs of PPGPR and GPR for Dette Pepelyshev (2010) curved function under Latin hypercube designs with different sample sizes.

the MAPEs of both methods decrease and the MAPE of GPR decreases faster than PPGPR. The results in Figure 9 can prove that the proposed PPGPR is highly suitable for a sparse learning environment but when enough training samples are available the PPGPR is not recommended.

4.3. Performance of GPR and PPGPR under Latin hypercube designs with different sizes

We compare the performance of GPR and PPGPR under Latin hypercube designs (Helton and Davis, 2003) with different sample sizes. We choose the Dette Pepelyshev (2010) curved function (Dette and Pepelyshev, 2010) as the underlying function. The R package `lhs` is used to generate the Latin hypercube designs using the maximin criterion (Joseph and Hung, 2008). The size of the testing set is 500. Figure 10 shows the MAPEs of GPR and PPGPR under the sample sizes from 40 to 120. It can be easily seen that the PPGPR has lower MAPEs than GPR most of the time. GPR has a lower MAPE only when the sample size is 63. Figure 10 proves the superiority of the proposed PPGPR over GPR under the Latin hypercube design with different sample sizes.

5. More numerical studies

We conduct more numerical studies to examine the computational cost of the proposed method, and the effect of initial values of the weight w . We also compare the PPGPR with a

Table 3. Assumed design range for HE case.

	$m(\text{kg/s})$	$T_{in}(\text{K})$	$k(\text{W/mK})$	$T_{wall}(\text{K})$
Lower Bound	0.00055	270.00	202.4	330
Upper Bound	0.001	303.15	360.0	400

new additive Gaussian model proposed in Delbridge *et al.* (2020). We defer these results to the [Supplementary Materials](#).

6. Approximated heat exchanger case study

In this section, we apply the proposed method PPGPR on a Heat Exchanger (HE) application introduced by Qian *et al.* (2006). The HE data in Qian *et al.* (2006) have two fidelities, known as detailed data (high fidelity) and approximated data (low fidelity). Because this work considers only the surrogate modeling for single-fidelity datasets, we use only the approximated data to implement the proposed method. The main objective of this application is to explore the impact of four factors, including the mass flow rate of entry air m , the temperature of entry air T_{in} , the temperature of the heat source T_{wall} and the solid material thermal conductivity M , on the total rate of steady state heat transfer y_a achieved by a heat exchanger. All design points live in a hypercube whose upper and lower bounds are shown in [Table 3](#). We follow the treatment in Qian *et al.* (2006) to partition the dataset into a training set of 64 samples and a testing set of 14 samples.

In this section, we compare the performance of GPR with isotropic and product correlations, Transformed Approximately Additive Gaussian Process Regression (TAAG) proposed in Lin and Roshan Joseph (2020), and the proposed PPGPR. In Lin and Roshan Joseph (2020) the performance was assessed in terms of the Root Mean Square Error (RMSE), defined as

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2}, \quad (11)$$

where \hat{y}_i is the predicted value and y_i means the true value for every sample, n stands for the size of testing set. Therefore, we consider RMSE of all the candidate methods. The implementation details of the these methods are as follows:

- GPR (with isotropic and product correlation functions): We use the Dicekriging package (Roustant *et al.*, 2012) with isotropic and product Matérn correlation and smoothness $\nu = 2.5$ to compute the predictive results.
- TAAG: The result in Lin and Roshan Joseph (2020) is refered here.
- PPGPR: We use the Matérn correlation with smoothness $\nu = 2.5$ and $M = 28$, $\eta = 10^{-9}$.

The results of these three methods are shown in [Table 4](#). It can be seen that the proposed method has a lower RMSE than other methods.

Table 4. RMSEs of GPR with isotropic and product correlations, TAAG and the PPGPR.

	$GPR(iso)$	$GPR(pro)$	TAAG	PPGPR
RMSE	4.20	4.26	2.08	1.82

7. Discussion

In this article, we propose a projection pursuit approach based on GPR to fit deterministic computer outputs. The proposed method has a better model prediction and generalization power when the input dimension is high, and the sample size is small.

Despite its advantages, the proposed method has a few issues to be addressed in future investigations. First, PPGPR involves quite a few hyper-parameters. Although we have provided a few guidelines regarding the choice of these hyper-parameters, how to better choose or tune these parameters requires further investigation. Second the current algorithm can only handle moderate data sets, due to its high computational cost. We believe that this issue can be mitigated by implementing the following techniques: (i) parallel or GPU computation, (ii) the recent advances in scalable GP inference and prediction (Liu *et al.*, 2020; Katzfuss and Guinness, 2021; Chen *et al.*, 2022).

In practice, uncertainty quantification is often of importance in addition to a point estimation. Note that (7) can be regarded as an original GPR with an additive kernel function. In view of this, the corresponding confidence intervals can be obtained following a standard GPR technique. However, our numerical experience implies that, the confident bands provided by the above approach are much wider than those generated by the usual GPR methods. This deficiency may be due to the lack of identifiability of the proposed models, as discussed in [Section 3](#). Uncertainty quantification for the proposed model using alternative approaches should be considered in a future work.

Supplementary materials

In the [Supplementary Materials](#), we present an upper bound of uniform prediction error of GPR with an additive correlation function, which implies a promising rate of convergence of additive Gaussian process models. Also, more numerical studies are included in the [Supplementary Materials](#).

Acknowledgments

The authors are grateful to the AE and two reviewers for very helpful comments and suggestions.

Funding

This work is supported by NSF grants DSM-1914636 and CCF-1934904, and 2022 Texas A&M Institute of Data Science Career Initiation Fellow Program.

Notes on contributors

Gecheng Chen is currently a PhD candidate in the Industrial and Systems Engineering Department of Texas A&M University. His research interests include machine learning and computer experiments.

Rui Tuo received his PhD in Statistics at University of Chinese Academy of Sciences in 2013. Currently, he is an assistant professor in the Department of Industrial and Systems Engineering and an adjunct assistant professor in the Department of Statistics at Texas A&M University. Tuo's research interest mainly lies in data science for computer simulations, including design and analysis of computer experiments and uncertainty quantification. Recently, he has also been working on theory and computational techniques for nonparametric models and Gaussian process regression.

ORCID

Gecheng Chen  <http://orcid.org/0000-0002-4466-6416>

References

Ba, S. and Joseph, V.R. (2012) Composite Gaussian process models for emulating expensive functions. *The Annals of Applied Statistics*, **6**(4), 1838–1860.

Ben-Ari, E.N. and Steinberg, D.M. (2007) Modeling data from computer experiments: An empirical comparison of kriging with mars and projection pursuit regression. *Quality Engineering*, **19**(4), 327–338.

Chen, H., Ding, L. and Tuo, R. (2022) Kernel packet: An exact and scalable algorithm for Gaussian process regression with Matérn correlations. *Journal of Machine Learning Research*, **23**(127), 1–32.

Choi, T., Shi, J.Q. and Wang, B. (2011) A Gaussian process regression approach to a single-index model. *Journal of Nonparametric Statistics*, **23**(1), 21–36.

Constantine, P.G., Dow, E. and Wang, Q. (2014) Active subspace methods in theory and practice: applications to kriging surfaces. *SIAM Journal on Scientific Computing*, **36**(4), A1500–A1524.

Delbridge, I., Bindel, D. and Wilson, A.G. (2020) Randomly projected additive Gaussian processes for regression, in *International Conference on Machine Learning*, Vol. 119, pp. 2453–2463, PMLR.

Deng, X., Lin, C.D., Liu, K.-W. and Rowe, R. (2017) Additive Gaussian process for computer models with qualitative and quantitative factors. *Technometrics*, **59**(3), 283–292.

Dette, H. and Pepelyshev, A. (2010) Generalized Latin hypercube design for computer experiments. *Technometrics*, **52**(4), 421–429.

Durrande, N., Ginsbourger, D. and Roustant, O. (2012) Additive covariance kernels for high-dimensional Gaussian process modeling. *Annales de la Faculté des sciences de Toulouse: Mathématiques*, **21**, 481–499.

Durrande, N., Ginsbourger, D., Roustant, O. and Carraro, L. (2013) Anova kernels and rkhs of zero mean functions for model-based sensitivity analysis. *Journal of Multivariate Analysis*, **115**, 57–67.

Duvenaud, D.K., Nickisch, H. and Rasmussen, C. (2011) Additive Gaussian processes, in *Advances in Neural Information Processing Systems*, Vol. 24, pp. 226–234, MIT Press.

Ferraty, F., Goia, A., Salinelli, E. and Vieu, P. (2013) Functional projection pursuit regression. *Test*, **22**(2), 293–320.

Forrester, A., Sobester, A. and Keane, A. (2008) *Engineering Design via Surrogate Modelling: A Practical Guide*. John Wiley & Sons.

Friedman, J.H. and Stuetzle, W. (1981) Projection pursuit regression. *Journal of the American Statistical Association*, **76**(376), 817–823.

Gilboa, E., Saatci, Y. and Cunningham, J. (2013) Scaling multidimensional Gaussian processes using projected additive approximations, in *International Conference on Machine Learning*, Vol. 28, pp. 454–461, PMLR.

Glaws, A., Constantine, P.G. and Cook, R.D. (2020) Inverse regression for ridge recovery: A data-driven approach for parameter reduction in computer experiments. *Statistics and Computing*, **30**(2), 237–253.

Goodfellow, I., Bengio, Y. and Courville, A. (2016) *Deep Learning*, MIT Press, Cambridge, MA.

Gramacy, R.B. and Lee, H.K.H. (2008) Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, **103**(483), 1119–1130.

Gramacy, R.B. and Lian, H. (2012) Gaussian process single-index models as emulators for computer experiments. *Technometrics*, **54**(1), 30–41.

Gu, M. (2019) Jointly robust prior for Gaussian stochastic process in emulation, calibration and variable selection. *Bayesian Analysis*, **14**(3), 857–885.

Halton, J.H. (1964) Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the ACM*, **7**(12), 701–702.

Harper, W.V. and Gupta, S.K. (1983) Sensitivity/uncertainty analysis of a borehole scenario comparing Latin hypercube sampling and deterministic sensitivity approaches. Office of Nuclear Waste Isolation, Battelle Memorial Institute.

Hastie, T., Tibshirani, R., Friedman, J.H. and Friedman, J.H. (2009) The elements of statistical learning: data mining, inference, and prediction, Vol. 2, pp. 1–758. Springer, New York.

Heaton, M.J., Christensen, W.F. and Terres, M.A. (2017) Nonstationary Gaussian process models using spatial hierarchical clustering from finite differences. *Technometrics*, **59**(1), 93–101.

Helton, J.C. and Davis, F.J. (2003) Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems. *Reliability Engineering & System Safety*, **81**(1), 23–69.

Hinton, G.E. and Salakhutdinov, R.R. (2006) Reducing the dimensionality of data with neural networks. *Science*, **313**(5786), 504–507.

Hokanson, J.M. and Constantine, P.G. (2018) Data-driven polynomial ridge approximation using variable projection. *SIAM Journal on Scientific Computing*, **40**(3), A1566–A1589.

Hu, Y., Gramacy, R.B. and Lian, H. (2013) Bayesian quantile regression for single-index models. *Statistics and Computing*, **23**(4), 437–454.

James, G.M. and Silverman, B.W. (2005) Functional adaptive model estimation. *Journal of the American Statistical Association*, **100**(470), 565–576.

Joseph, V.R. and Hung, Y. (2008) Orthogonal-maximin Latin hypercube designs. *Statistica Sinica*, **18**, 171–186.

Katzfuss, M. and Guinness, J. (2021) A general framework for vecchia approximations of Gaussian processes. *Statistical Science*, **36**(1), 124–141.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q. and Liu, T.-Y. (2017) Lightgbm: A highly efficient gradient boosting decision tree, in *Advances in Neural Information Processing Systems*, Vol. 30, pp. 3146–3154, MIT Press.

Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M. and Tang, P.T.P. (2016) On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*.

Khoo, Y., Lu, J. and Ying, L. (2017) Solving parametric PDE problems with artificial neural networks. *arXiv preprint arXiv:1707.03351*.

Lawrence, S. and Giles, C.L. (2000) Overfitting and neural networks: Conjugate gradient and backpropagation, in *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, Volume 1, IEEE Press, Piscataway, NJ, pp. 114–119.

Lebarbier, É. (2005) Detecting multiple change-points in the mean of Gaussian process by model selection. *Signal Processing*, **85**(4), 717–736.

LeCun, Y., Bengio, Y. and Hinton, G. (2015) Deep learning. *Nature*, **521**(7553), 436–444.

Li, C.-L., Kandasamy, K., Póczos, B. and Schneider, J. (2016) High dimensional Bayesian optimization via restricted projection pursuit models, in *Artificial Intelligence and Statistics*, Vol. 51, pp. 884–892, PMLR.

Lin, L.-H. and Roshan Joseph, V. (2020) Transformation and additivity in Gaussian processes. *Technometrics*, **62**(4), 525–535.

Linkletter, C., Bingham, D., Hengartner, N., Higdon, D. and Ye, K.Q. (2006) Variable selection for Gaussian process models in computer experiments. *Technometrics*, **48**(4), 478–490.

Liu, H., Ong, Y.-S., Shen, X. and Cai, J. (2020) When Gaussian process meets big data: A review of scalable gps. *IEEE Transactions on Neural Networks and Learning Systems*, **31**(11), 4405–4423.

Mak, S., Sung, C.-L., Wang, X., Yeh, S.-T., Chang, Y.-H., Joseph, V.R., Yang, V. and Wu, C.F.J. (2018) An efficient surrogate model for emulation and physics extraction of large eddy simulations. *Journal of the American Statistical Association*, **113**(524), 1443–1456.

Makridakis, S. (1993) Accuracy measures: Theoretical and practical concerns. *International Journal of Forecasting*, **9**(4), 527–529.

Müller, H.-G. and Yao, F. (2008) Functional additive models. *Journal of the American Statistical Association*, **103**(484), 1534–1544.

Oakley, J.E. and O'Hagan, A. (2004) Probabilistic sensitivity analysis of complex models: A Bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **66**(3), 751–769.

Peng, C.-Y. and Wu, C.F.J. (2014) On the choice of nugget in kriging modeling for deterministic computer experiments. *Journal of Computational and Graphical Statistics*, **23**(1), 151–168.

Pinkus, A. (1997) Approximating by ridge functions. *Surface Fitting and Multiresolution Methods*, 279–292.

Prechelt, L. (1998) Early stopping—but when? in G.B. Orr, K.R. Müller (eds.), *Neural Networks: Tricks of the Trade*, Lecture Notes in Computer Science, Vol. 1524, Springer, Berlin, Heidelberg.

Psichogios, D.C. and Ungar, L.H. (1992) A hybrid neural network-first principles approach to process modeling. *AICHE Journal*, **38**(10), 1499–1511.

Qian, Z., Seepersad, C.C., Joseph, V.R., Allen, J.K. and Jeff Wu, C.F. (2005) Building surrogate models based on detailed and approximate simulations. *ASME Journal of Mechanical Design*, **128**(4), 668–677.

Rasmussen, C. and Williams, C. (2006) *Gaussian Processes for Machine Learning*, MIT Press, Cambridge, MA.

Roustant, O., Ginsbourger, D. and Deville, Y. (2012) Dicekriging, diceoptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization. *Journal of Statistical Software*, **51**(1), 1–55.

Sacks, J., Welch, W., Mitchell, T. and Wynn, H. (1989) Design and analysis of computer experiments. *Statistical Science*, **4**(4), 409–423.

Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M. and Tarantola, S. (2010) Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. *Computer Physics Communications*, **181**(2), 259–270.

Santner, T.J., Williams, B.J. and Notz, W.I. and Williams, B.J. (2013) *The Design and Analysis of Computer Experiments*, Vol. 1, Springer, New York.

Smola, A.J. and Schölkopf, B. (2004) A tutorial on support vector regression. *Statistics and Computing*, **14**(3), 199–222.

Stein, M.L. (1999) *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Science & Business Media, Berlin/Heidelberg, Germany.

Tripathy, R., Bilionis, I. and Gonzalez, M. (2016) Gaussian processes with built-in dimensionality reduction: Applications to high-dimensional uncertainty propagation. *Journal of Computational Physics*, **321**, 191–223.

Tripathy, R.K. and Bilionis, I. (2018) Deep UQ: Learning deep neural network surrogate models for high dimensional uncertainty quantification. *Journal of Computational Physics*, **375**, 565–588.

Tuo, R. and Wang, W. (2020) Kriging prediction with isotropic Matérn correlations: Robustness and experimental designs. *Journal of Machine Learning Research*, **21**, 1–38.

Wang, J.-L., Xue, L., Zhu, L. and Chong, Y.S. (2010) Estimation for a partial-linear single-index model. *The Annals of Statistics*, **38**(1), 246–274.

Welch, W.J., Buck, R.J., Sacks, J., Wynn, H.P., Mitchell, T.J. and Morris, M.D. (1992) Screening, predicting, and computer experiments. *Technometrics*, **34**(1), 15–25.

Wendland, H. (2004) *Scattered Data Approximation*. Vol. 17. Cambridge University Press, Cambridge, England.

Wilson, A.G., Knowles, D.A. and Ghahramani, Z. (2011) Gaussian process regression networks. *arXiv preprint arXiv:1110.4411* .