

Multi-fidelity Gaussian process surrogate modeling for regression problems in physics

Kislaya Ravi^{†*}, Vladyslav Fediukov^{‡‡*}, Felix Dietrich[†], Tobias Neckel[†], Fabian Buse[‡], Michael Bergmann[¶], Hans-Joachim Bungartz[†]

[†] School of Computation, Information and Technology, Technische Universität München, Garching b. München, Germany

[‡] Institute of System Dynamics and Control, German Aerospace Center (DLR), Weßling, Germany

[¶] Max-Planck-Institut für Plasmaphysik, Garching b. München, Germany

* These authors contributed equally to this work

March 2024

Abstract. One of the main challenges in surrogate modeling is the limited availability of data due to resource constraints associated with computationally expensive simulations. Multi-fidelity methods provide a solution by chaining models in a hierarchy with increasing fidelity, associated with lower error, but increasing cost. In this paper, we compare different multi-fidelity methods employed in constructing Gaussian process surrogates for regression. Non-linear autoregressive methods in the existing literature are primarily confined to two-fidelity models, and we extend these methods to handle more than two levels of fidelity. Additionally, we propose enhancements for an existing method incorporating delay terms by introducing a structured kernel. We demonstrate the performance of these methods across various academic and real-world scenarios. Our findings reveal that multi-fidelity methods generally have a smaller prediction error for the same computational cost as compared to the single-fidelity method, although their effectiveness varies across different scenarios.

Keywords: multi-fidelity, machine learning, Gaussian processes, physical simulations

1. Introduction

Complex simulations are often computationally expensive and time-consuming. A surrogate model is a simpler and faster model that emulates the output of a complex model as a function of input parameters. Surrogates, also known as digital twins, have various applications in design optimization [1, 2], uncertainty quantification [3, 4], real-time predictions [5, 6], etc. The amount of training data is one of the crucial factors governing the quality of the surrogate. Generating an extensive training data set is computationally infeasible for expensive simulations. In this work, we tackle the data availability limitation using Multi-fidelity [7].

The advancement of computational capabilities significantly increased the use of numerical simulation in almost every field of application. This led to the development of various simulation methods offering different levels of approximation quality and computational costs. Low-fidelity models provide estimations with reduced accuracy but demand fewer resources, whereas high-fidelity models yield precise predictions at a higher cost. These distinct fidelity levels arise from differences in physical modeling or simulation of the same phenomena, the linearization of complex physical processes, or the application of the same modeling approach with different discretization levels. The trade-off between simulation accuracy and computational cost is an ever-present challenge. Relying solely on high-fidelity models for applications that require multiple evaluations of the model is impractical due to their computational cost. Conversely, utilizing only low-fidelity models may compromise the results' accuracy. Multi-fidelity methods address this challenge by leveraging low-fidelity models to alleviate computational load and incorporating occasional high-fidelity evaluations to control result accuracy.

Multi-fidelity methods are widely used in applications like surrogate modeling, uncertainty quantification [8, 9, 10, 11], bayesian inference [12, 13, 14], optimization [15, 16, 17, 18, 19, 20], etc. They are helpful in scenarios where resource considerations and accurate estimations are important simultaneously.

This work focuses on using multi-fidelity models to build a Gaussian Process (GP) [21] surrogate. We use GP because of its probabilistic nature and additional functionalities like Bayesian optimization and uncertainty quantification. There are multiple ways to incorporate multi-fidelity in GP. The first method developed was the linear auto-regressive models [22, 23]. The linear modeling was not sufficient for more complex problems. This led to the development of non-linear auto-regression methods [24, 25]. However, the non-linear methods cannot be extended to more than two levels of fidelity and require low-fidelity evaluations during prediction. One can use Deep Gaussian Process (DGP) [26] to tackle the limitations.

This paper is organized as follows. We provide the basic theoretical background on GP, kernels, and calibration in Section 2. Then, we review different multi-fidelity GP surrogate models in Section 3. We also suggest modifications in existing non-linear autoregressive to incorporate more than two fidelities without DGP. In the same

section, we suggest using delay terms in multi-fidelity DGP. Then, we compare the performance of different methods on the academic problem and two real-world problems, namely Terramechanical problems and Plasma microturbulence simulations in Section 4. Finally, we end the paper with concluding remarks in Section 5.

2. Background

2.1. Gaussian process regression

Gaussian processes (GP) are an efficient and flexible tool for probabilistic predictions [21]. They provide reliable statistical inference as well as uncertainty predictions. A GP is a stochastic process in which each finite subset of variables forms a multivariate Gaussian distribution. GPs are defined by their covariance (kernel) function k and define a probability distribution over functions,

$$f(\cdot, \cdot) \sim \mathcal{GP}(0, k(\cdot, \cdot)). \quad (1)$$

The mean function is usually assumed to be zero for simplicity. In practice, the base space of the process is sampled with a finite number of N points $X = \{x_i\}_{i=1}^N$. Then, the GP is given as a finite-dimensional, multivariable Gaussian distribution of points y ,

$$y \sim \mathcal{N}(\mu, K), \quad (2)$$

where μ is a vector of means, usually zeros, and $K = (K_{ij})_{i,j=1}^N = (k(x_i, x_j))_{i,j=1}^N$ is a covariance matrix of size $N \times N$. Gaussian distributions can be conveniently employed in a Bayesian framework because they are closed under condition and conjugate distributions, which means marginal and conditional distributions of multivariate Gaussians are again Gaussian.

GP is typically used in regression, also called the “Kriging method” [27]. In this case, each training point $x_i \in X$ is assigned an additional function value y_i , and the goal is to construct function values y^* for unobserved test points $X^* = \{x^*\}$. Then, we can construct a joint distribution

$$\begin{bmatrix} f(X) \\ f(X^*) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K(X, X) & K(X, X^*) \\ K(X^*, X) & K(X^*, X^*) \end{bmatrix}\right). \quad (3)$$

After we observe the training data, the posterior consists of a narrowed distribution of functions, initially defined by the prior distribution as observed in Figure 1. The inference is done then using posterior mean and posterior covariance,

$$\mathbb{E}(f(X^*)) = K(X, X^*)^T K(X, X)^{-1} f(X), \quad (4)$$

$$\text{cov}(f(X^*)) = K(X^*, X^*) - K(X, X^*)^T K(X, X)^{-1} K(X^*, X). \quad (5)$$

Having posterior mean and covariance, we can generate predictions for unobserved points.




Figure 1: A kernel function defines the family of the functions with which we are approximating the target function. The first picture shows a prior distribution defined by a kernel. In the posterior distribution in the second plot, newly observed training data points, depicted here in black, narrow the space of potential functions from the initial family.

2.2. Kernels

K denotes a matrix built with the help of a *kernel*, which is a function used to define the covariances between the data points. This covariance is usually interpreted as a measure of similarity between data points, and thus, points x and x' that are considered similar in the input space will have target values y and y' that are also similar. The covariance also encodes certain properties of the functions we want to predict, for example, periodicity [28]. Choosing the right kernel is crucial in the modeling process with GP.

A valid covariance function must be positive semidefinite [21]. The product and sum of two valid covariance functions generate a valid covariance function. Using this property, one can generate more sophisticated covariance functions. There are multiple options for covariance functions like squared-exponential, exponential, trigonometric, etc. The task of choosing the right kernel is complicated. There is no standardized algorithm to do so except for the random search and the researcher's intuition regarding modeling. However, some authors are making promising progress by creating interesting heuristics and algorithms [29, 30, 31, 32]. In this work, we only work with squared-exponential covariance functions.

After choosing the kernel family, we evaluate the hyperparameters of the kernel by maximizing the log-likelihood. The log-likelihood L as a function of observed values of data points Y , parameterized by the hyperparameters θ as

$$\log(p(y|X, \theta)) = -\frac{1}{2}y^T K_y^{-1}y - \frac{1}{2}\log |K_y| - \frac{n}{2}\log 2\pi. \quad (6)$$

We will use a gradient-based method (LBFGS [33] with multiple starting points) to find the maximum likelihood.

2.3. Calibration

As shown in equation 6, hyperparameter optimization for Gaussian process kernels is usually performed by maximizing the marginal likelihood. This often leads to an acceptable mean squared error [21], but the posterior variance tends to be lower than the variance of the actual distribution [34, 35]. This problem requires an additional post-processing step known as calibration. There exists a variety of calibration techniques, but in all cases, they aim to adjust the predicted variance towards the actual variance in the data. Calibration is an established concept in the classification context, but in the case of regression, it is relatively new [36]. In general, calibration can be split into three categories: *quantile*, *distribution*, and *variance* calibration.

The regression is *quantile calibrated*, if

$$\frac{1}{N} \sum_{n=1}^N \mathbf{I}\{y_n < F_n^{-1}(p)\} \rightarrow p \text{ for all } p \in [0, 1] \text{ as } N \rightarrow \infty \quad (7)$$

for the set of data $(x_n, y_n)_{n=1}^N$, where F_n stands for predicted cumulative distribution at x_n and F_n^{-1} the corresponding quantile function. Intuitively, $X\%$ of the confidence interval should capture the ground truth values in $X\%$ predictions [37, 38]. Quantile calibration is considered a global calibration, as it concentrates only on the marginal distribution without considering a calibration of distribution around each exact prediction.

Distribution calibration resembles the calibration approach in classification tasks because it is also focused on local calibration. In the classification case, instances are grouped by their predicted probability and they are considered calibrated only if each of that group matches the indicated probability. Assume X as input variables, Y as a target variable, and the regression model $f : \mathbb{X} \rightarrow \mathbb{S}_Y$ or intuitively, the regression model is mapping from input features into a space of distributions and for each test instance predicted a particular distribution. Denoting s as an arbitrary distribution predicted by the regression model, the regression model f is distribution calibrated if

$$p(Y = y | f_{Y|X} = s) = s(y) \text{ for all } y \in Y, s \in S, \quad (8)$$

or intuitively, for all predicted instances with a particular distribution s , all instances on average with such prediction should follow this distribution s [37].

Variance calibration aims to make the model predict its own error, i.e. for all instances where variance $\sigma^2(x)$ takes a certain value ω , the squared predicted error $\mathbb{E}_{X,Y}[\mu(x) - y]$ will match this predicted variance [39]. Given $\mu(x)$ is a predicted mean and $\sigma^2(x)$ is a predicted variance, then the regression model is variance calibrated if

$$\mathbb{E}_{X,Y}[(\mu(x) - y)^2 | \sigma^2(x) = \omega] = \omega. \quad (9)$$

Different methods approach uncertainty calibration tasks in different ways, therefore it is beneficial to test several of them on given data.

3. Methodology

3.1. Multi-fidelity

In this paper, our focus is on exploring various multi-fidelity Gaussian process surrogate modeling methods. In the upcoming sub-sections, we delve into the motivation behind each method, elucidate their formulation, and critically evaluate their limitations. Additionally, we propose enhancements for some methods within the corresponding sub-sections to address identified shortcomings.

3.1.1. Linear Auto-regressive model We start with the Auto-Regressive Model (AR1), a straightforward linear auto-regressive model introduced by Kennedy and O'Hagan [22]. In this model, we formulate the joint distribution of all fidelities, building on the fundamental concept of expressing the model for a particular fidelity as the sum of a linear scaling of the previous fidelity and an additive correction term.

We explain AR1 using a two-fidelity case. This process can be easily extended to cases with more than two fidelities. Let us consider the following random variables:

$$\begin{aligned} u_l &\sim \mathcal{N}(0, K_l) \Rightarrow \mathbb{E}[u_l] = 0 \quad \text{and} \quad \mathbb{E}[u_l u_l'] = K_l(X, X'), \\ u_\delta &\sim \mathcal{N}(0, K_\delta) \Rightarrow \mathbb{E}[u_\delta] = 0 \quad \text{and} \quad \mathbb{E}[u_\delta u_\delta'] = K_\delta(X, X'). \end{aligned}$$

We are modeling all the surrogates and the additive correction term using GP. u_l and u_δ represent the samples drawn from the low-fidelity and the additive correction GP. Let u_h represent a realization of the GP that approximates a high-fidelity function f_h . $\rho \in \mathbb{R}$ is a linear-scaling parameter. We assume an additive relationship between the consecutive fidelities so that

$$u_h = \rho u_l + u_\delta. \tag{10}$$

The expected value of the surrogate of the high-fidelity model is assumed to be zero,

$$\mathbb{E}[u_h] = \mathbb{E}[\rho u_l + u_\delta] = \rho \mathbb{E}[u_l] + \mathbb{E}[u_\delta] = 0. \tag{11}$$

The covariance of the surrogate of the high-fidelity model is then given by

$$\begin{aligned} \text{Cov}(u_h, u_h') &= \mathbb{E}[u_h u_h'] - \underbrace{\mathbb{E}[u_h] \mathbb{E}[u_h']}_{=0} \\ &= \mathbb{E}[(\rho u_l + u_\delta)(\rho u_l' + u_\delta')] = \rho^2 \mathbb{E}[u_l u_l'] + \mathbb{E}[u_\delta u_\delta'] = \rho^2 K_l + K_\delta. \end{aligned}$$

The covariance between u_l and u_h is

$$\begin{aligned} \text{Cov}(u_h, u_l') &= \mathbb{E}[u_h u_l'] - \underbrace{\mathbb{E}[u_h] \mathbb{E}[u_l']}_{=0} = \mathbb{E}[(\rho u_l + u_\delta) u_l'] \\ &= \rho \mathbb{E}[u_l u_l'] + \mathbb{E}[u_\delta u_l'] = \rho K_l + \underbrace{\mathbb{E}[u_\delta] \mathbb{E}[u_l']}_{=0} = \rho K_l. \end{aligned}$$

The joint distribution of samples drawn from the low-fidelity surrogate and the high-fidelity surrogate is written

$$\begin{bmatrix} u_h \\ u_l \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} K_\delta + \rho^2 K_l & \rho K_l \\ \rho K_l & K_l \end{bmatrix} \right). \tag{12}$$

Using the maximum likelihood method, one can calculate the value of ρ and the kernel hyperparameters. We can also extend the joint distribution in equation 12 to incorporate more than two fidelities that will lead to covariance matrices with sparse blocks. Le Gratiet [23] took advantage of this block structure and decreased the complexity of the method from $\mathcal{O}\left(\left(\sum_{l=1}^L N_l\right)^3\right)$ to $\sum_{l=1}^L \mathcal{O}(N_l^3)$. He also improved the performance using polynomial terms for scaling (ρ) instead of a constant term.

Nevertheless, the efficacy of the Auto-Regressive Model (AR1) is limited by the assumption of a linear dependency between the low-fidelity and high-fidelity functions. This assumption might prove inadequate in scenarios characterized by non-linear transformations.

3.1.2. Non-linear auto-regressive models Let us consider a two-fidelity case. A non-linear transformation on the low-fidelity function to obtain the high-fidelity function is written as

$$f_h(x) = g(x, f_l(x)), \quad (13)$$

where $g : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ is a function representing the non-linear transformation applied on the low-fidelity function to obtain the high-fidelity model. When attempting to model both the low-fidelity and high-fidelity functions using GPs, writing the joint distribution of the surrogates, as presented in equation 12, becomes analytically challenging for a general transformation g . To overcome this difficulty, a common assumption in studies such as [25, 24] is that the low-fidelity function can be evaluated at zero cost. This eliminates the need for a surrogate for the low-fidelity function. Consequently, the non-linear transformation depicted in equation 13 transforms into creating a surrogate in a $d + 1$ dimensional space. In our work, we delve into the modeling of g using a Gaussian Process (GP). This approach enables us to effectively navigate the challenges associated with the joint distribution of surrogates, providing a viable solution for incorporating non-linear transformations within the multi-fidelity surrogate modeling framework.

Let us consider the following pedagogical examples of low-fidelity and high-fidelity functions from [25]. The high-fidelity model f_h and the lower-fidelity model f_l are defined through

$$\begin{aligned} f_h(x) &= (x - \sqrt{2}) \sin^2(8\pi x), \\ f_l(x) &= \sin(8\pi x). \end{aligned}$$

Figure 2 illustrates the relationship between f_l and f_h . Modeling the transformation g as described in equation 13 is synonymous with learning the manifold depicted in Figure 2. Crucially, this manifold lacks oscillations, signifying that learning it necessitates relatively fewer high-fidelity function evaluations. This observation underscores the advantage of multi-fidelity modeling in this example, as it allows for a learning process by leveraging the complicated feature to be captured in the low-fidelity function, thereby reducing the requirement of high-fidelity function evaluations.




Figure 2: It is difficult to approximate the function $f_h(x)$ (orange) using only x as an input, because it is highly oscillatory. Including the function $f_l(x)$ (blue) as an additional input means only the green, smooth manifold must be approximated, which can be done accurately with only a few data points. The example is adapted from [25].

A refinement of the method involves incorporating a well-structured kernel into the GP model, as suggested by Perdikaris et al. [25]. The proposed kernel structure is given by:

$$k(x, x') = k_\rho(x, x'; \theta_\rho) k_f(f_l(x), f_l(x'); \theta_f) + k_\delta(x, x'; \theta_\delta). \quad (14)$$

Here, θ_ρ , θ_f , and θ_δ denote the kernel hyperparameters, which are determined by maximizing the likelihood, as discussed in the previous section. This kernel structure closely mirrors the Auto-Regressive Model (AR1) formulation described in equation 10, assigning each kernel to an individual part of the transformation. Specifically, k_ρ models the scaling term, k_f is responsible for transforming the low-fidelity model, and k_δ handles the additive correction term. The formulation above gives the method its name: Non-linear Auto-Regressive Gaussian Process (NARGP).

Introducing this well-defined prior aids the model in accurately fitting the underlying relationships. A drawback of this formulation is the increase in the number of hyperparameters, which necessitates careful consideration during the optimization

process. Despite this challenge, the structured nature of the kernel contributes to a more informed and effective surrogate modeling approach.

We follow Lee et al. [24] and assume sufficient smoothness of the high-fidelity function f_h . We can then use Taylor expansion to approximate the value of f_h close to x by

$$f_h(x + \Delta x) = f_h(x) + \sum_{i=1}^d \frac{\partial}{\partial x_i} f_h(x) \Delta x_i + \mathcal{O}(\Delta x^2). \quad (15)$$

Substituting equation 13 into equation 15, we obtain

$$\begin{aligned} f_h(x + \Delta x) &= g(x, f_l(x)) + \sum_{i=1}^d \left(\frac{\partial g}{\partial x_i} g(x, f_l(x)) \right) \Delta x_i + \mathcal{O}(\Delta x^2) \\ &= g(x, f_l(x)) + \sum_{i=1}^d \left(\frac{\partial g}{\partial x_i} + \frac{\partial g}{\partial f_l} \frac{\partial f_l}{\partial x_i} \right) \Delta x_i + \mathcal{O}(\Delta x^2). \end{aligned}$$

The Taylor expansion converges if

- (i) $\left\| \frac{\partial g(\dots)}{\partial x} \right\| \leq c_g$ for $c_g \in \mathbb{R}^+$
- (ii) $\left\| \frac{\partial f_l(\dots)}{\partial x} \right\| \leq c_l$ for $c_l \in \mathbb{R}^+$

The Taylor expansion also motivates us to incorporate the derivative of low-fidelity in equation 13. Lee et al. [24] suggest writing the non-linear transformation as

$$f_h(x) = g_d \left(x, f_l(x), \frac{\partial f_l}{\partial x} \right). \quad (16)$$

In many legacy simulators, the gradients are unavailable [40]. In cases where gradients are unavailable, one can resort to finite difference techniques to approximate the derivative. Lee et al. [24] suggest replacing the derivative in equation 16 with the delay term to accommodate such scenarios. This delay term evaluates the low-fidelity function after a small time step $\tau \in \mathbb{R}$. For a higher-dimensional input variable, we add a delay along each dimension, defining τ_i as a vector of zeros with τ at the i^{th} position. The transformed formulation is now expressed as

$$f_h(x) = g_d(x, f_l(x), f_l(x + \tau_1), f_l(x + \tau_2), \dots, f_l(x + \tau_d)). \quad (17)$$

This method is called the Gaussian Process with Delay Fusion (GPDF) proposed in [24]. Additionally, we propose an enhancement to this method by modifying the kernel to mimic the Taylor expansion using a structured composite kernel similar to equation 14

$$k(x, x') = k_\rho(x, x'; \theta_\rho) k_f((f_l(x), f_l(x + \tau)), (f_l(x'), f_l(x' + \tau)); \theta_f) + k_\delta(x, x'; \theta_\delta). \quad (18)$$

This improved method is called the Gaussian Process with Delay Fusion and Composite kernel (GPDFC).

The formulations of NARGP, GPDF, and GPDFC come with certain limitations that need to be addressed for broader applicability:

- (i) **Nested Evaluation Requirements:**

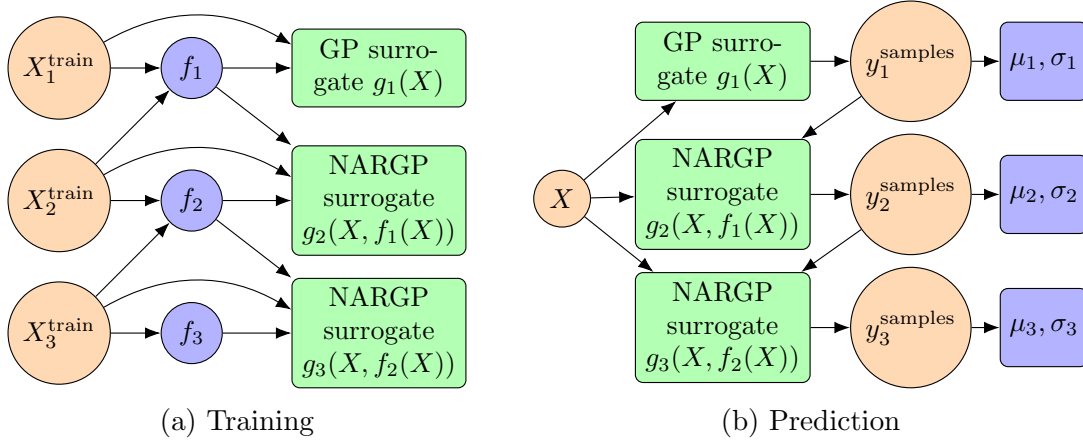


Figure 3: Flowchart showing the training and prediction steps of NARGP surrogate for a three-fidelity scenario. Training the NARGP model for a particular level involves evaluating the model of that fidelity level and the previous fidelity level. We draw samples to make predictions for a particular level, which involves drawing samples from the surrogate of the previous level and feeding them as input to the surrogate of that level. We can then use the samples to obtain the required statistical moments of the predictions by marginalizing the samples.

- NARGP: Requires low-fidelity model evaluations during training and prediction at precisely the same parameters as the high-fidelity model.
- GPDPF and GPDPFC: Share similar constraints and involve additional low-fidelity function evaluations at delay points.

These limitations were initially disregarded under the assumption of infinitely cheap low-fidelity function evaluations. However, in practice, even low-fidelity function evaluations have nonzero cost, and so these constraints can also become significant. In some cases, we are simply given data collected from the low and high-fidelity models. Gathering any further data might not be possible, and for such cases, ensuring nested training data points may also not be possible. Moreover, the possible absence of low-fidelity data at the prediction parameter point also serves as a hurdle in implementing the methods mentioned in data-driven multi-fidelity problems [26].

- (ii) **Limited to Two Fidelity Cases:** Modern NARGP, GPDPF, and GPDPFC formulations are tailored for scenarios involving only two fidelity models, restricting their application in situations with more than two fidelities.
- (iii) **Overfitting:** As shown in [26], NARGP is prone to overfitting, which leads to poor generalization.

To overcome the limitation of application to a two-fidelity case and extend the framework to accommodate multiple fidelity levels, we can explore two scenarios:

- **Nested training data:** In the case of nested training data, where $(X_{\ell+1}^{\text{train}} \subset$

X_ℓ^{train}), we propose training a Gaussian Process surrogate for the lowest fidelity and multi-fidelity Gaussian Process surrogates for higher levels. We train each multi-fidelity surrogate using the previous level, as shown in Figure 3a for three fidelity NARGP surrogates. These layers are then cascaded, forming a stacked surrogate for multiple fidelities. During prediction, we draw samples from the surrogate of the previous level to obtain prediction samples for the next level and marginalize it to derive the posterior prediction for the given level. Let us represent y_ℓ as posterior prediction at parameter X . We can evaluate the posterior distribution of $y_{\ell+1}$ as

$$p(y_{\ell+1}|X) = \int p(y_{\ell+1}|y_\ell, X)p(y_\ell|X)dy_\ell. \quad (19)$$

One can write the marginalization step after cascading as

$$p(y_{\ell+1}|X) = \int_{y_\ell} \int_{y_{\ell-1}} \cdots \int_{y_1} p(y_{\ell+1}|y_\ell, X)p(y_\ell|y_{\ell-1}, X) \cdots p(y_2|y_1, X)p(y_1|X)dy_\ell dy_{\ell-1} \cdots dy_1.$$

We can perform this integration using the Monte Carlo method. The prediction steps for three fidelity stacked NARGP surrogates are shown in Figure 3b.

Every layer in the stacked architecture is a GP. So, the samples drawn from GP will stay close to the posterior mean when the posterior predicted variance is small. If we ensure that there are enough data points in each layer such that the maximum of the posterior variance across the parameter space is smaller than a cut-off limit then we can use the posterior predicted mean of the previous layer as input for the next layer. In this way, we can bypass the Monte Carlo step. An efficient way to ensure small posterior variance is by using the adaptivity algorithm, which we will discuss later.

- **Non-nested training data:** In situations with non-nested training data, an alternative is to employ the Deep Gaussian Process (DGP).

3.1.3. Deep Gaussian Processes Drawing inspiration from the stacked architecture of multilayer neural networks, Deep Gaussian Processes (DGP) extend a single GP by composing multiple GP with each other. Each node within DGP represents a Gaussian process. The fundamental concept of deep architectures lies in their ability to model complex functions through combinations of simpler ones, with the output of one layer serving as the input for the next [41]. The primary advantage of Deep Gaussian Processes (DGP) over a single Gaussian Process (GP) lies in the distribution of feature learning across different layers. In a DGP, each layer is dedicated to learning distinct features of the model. This is in contrast to a single GP, where all features are consolidated into one, potentially leading to a complex and intricate model, especially when using traditional kernels.

DGP can effectively capture various aspects or representations of the underlying system by distributing the learning process across multiple layers. DGP emerges as a




Figure 4: Flowchart showing the training and prediction steps of Deep Gaussian process surrogate for three fidelity scenarios. Training the DGP model for a particular level involves evaluating the model of that fidelity level and the surrogate of the previous fidelity level, which further involves evaluating the one-level lower surrogate. This creates a cascading effect until we reach the lowest fidelity surrogate. The prediction process is similar to the full GP, which involves drawing and marginalizing samples.

suitable candidate for modeling multi-fidelity scenarios. With each fidelity level, DGP gradually incorporates new features, enhancing the model's capabilities as it approaches high-fidelity representations. This progressive refinement allows DGP to effectively capture the intricacies of complex systems.

We model the transformation similar to equation 13 with one difference. We replace the actual function evaluation of the previous layer with the previous GP layer

$$f_\ell(X) = g_\ell(X, g_{\ell-1}(\cdots)). \quad (20)$$

The analytic form of the posterior of the GP surrogate g_ℓ is intractable because one of its parameters ($g_{\ell-1}$) is sampled from a posterior distribution. To sample from the posterior of g_ℓ , we use variational inference for each layer of DGP [42]. We use M_ℓ inducing points to approximate each layer representing a fidelity level. Typically, the number of inducing points is significantly smaller than the number of training data points ($M_\ell \ll N_\ell$). This decreases the computational requirements of GP training and prediction steps [43, 44, 45]. This approach is also known as the Sparse Variational Gaussian Process (SVGP) [46]. We discuss the modeling approach using DGP for multi-fidelity surrogate modeling in [26, 47].

Let us assume that we have a set of locations of the training parameters as

$$\mathbf{X} = \{X_1^{\text{train}}, X_2^{\text{train}}, \dots, X_L^{\text{train}}\}$$

and the corresponding function evaluations

$$\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_L\} = \{f_1(X_1^{\text{train}}), f_2(X_2^{\text{train}}), \dots, f_L(X_L^{\text{train}})\}.$$

Note that we do not impose the limitation that the location of the training parameters should be nested. Let \mathbf{g}_ℓ^t represent the samples drawn from the ℓ^{th} layer at the training parameter X_t^{train} . We assume that M_ℓ inducing points are sufficient to represent the distribution at layer ℓ . We represent the location of the inducing point by \mathbf{Z}_ℓ . Let \mathbf{u}_ℓ be the sample drawn from the ℓ^{th} layer at the inducing points \mathbf{Z}_ℓ . We represent the kernel function and the mean function of the prior at layer ℓ using k_ℓ and m_ℓ respectively. The unnormalized posterior distribution at ℓ^{th} layer is

$$p(\mathbf{y}_\ell, \mathbf{g}_\ell^\ell, \mathbf{u}_\ell | \mathbf{X}_\ell, \mathbf{g}_{\ell-1}^\ell, \mathbf{Z}_\ell) = \underbrace{p(\mathbf{y}_\ell | \mathbf{g}_\ell^\ell, \mathbf{X}_\ell)}_{\text{likelihood}} \underbrace{p(\mathbf{g}_\ell^\ell | \mathbf{u}_\ell, \mathbf{g}_{\ell-1}^\ell, \mathbf{X}_\ell, \mathbf{Z}_\ell) p(\mathbf{u}_\ell | \mathbf{Z}_\ell)}_{\text{DGP prior}}. \quad (21)$$

Figure 4a shows the training process for a three-fidelity scenario. The calculation of posterior at any layer needs samples from the previous layer at the training points, which in turn requires further evaluation from previous layers. For example, the calculation of the posterior of the third layer needs to draw samples from the second layer at X_3^{train} , which also requires samples from the first layer at X_3^{train} . We represent this hidden calculation with a dashed arrow in Figure 4a.

One can combine all the layers by multiplying the posterior from individual layers from equation 21 to obtain the joint posterior distribution for all layer $p(\mathbf{y}, \{\mathbf{g}_\ell^t, \mathbf{u}_\ell\})$. Analytical inference from the posterior is intractable. To approximate the distribution, we assume that the prior $q(\mathbf{u}_\ell)$ of \mathbf{u}_ℓ is a Gaussian distribution with mean $\tilde{\mathbf{m}}_\ell$ and variance $\tilde{\mathbf{S}}_\ell$. Now, the approximated DGP prior is

$$q(\mathbf{g}_\ell^\ell, \mathbf{u}_\ell | \mathbf{X}_\ell, \mathbf{g}_{\ell-1}^\ell, \mathbf{Z}_\ell) = p(\mathbf{g}_\ell^\ell | \mathbf{u}_\ell, \mathbf{g}_{\ell-1}^\ell, \mathbf{X}_\ell, \mathbf{Z}_\ell) q(\mathbf{u}_\ell). \quad (22)$$

If \mathbf{u}_ℓ is marginalized out of the equation 22, then the distribution becomes a Gaussian distribution with the following mean and variance:

$$\mu_\ell(x) = m_\ell(x) + \alpha(x)^T (\tilde{\mathbf{m}}_\ell - m_\ell(x)), \quad (23)$$

$$\Sigma_\ell(x, x') = k_\ell(x, x') - \alpha(x)^T (k_\ell(\mathbf{Z}_\ell, \mathbf{Z}_\ell) - \tilde{\mathbf{S}}_\ell) \alpha(x'). \quad (24)$$

We can use the before approximating the posterior at each layer and then extend this to all the layers of DGP to obtain the approximate posterior distribution $q(\mathbf{y}, \{\mathbf{g}_\ell^\ell, \mathbf{u}_\ell\}_{\ell=1}^L)$. We want to approximate the exact posterior distribution in equation 21 by an approximate distribution $q(\mathbf{y}, \{\mathbf{g}_\ell^\ell, \mathbf{u}_\ell\}_{\ell=1}^L)$. We train the DGP by minimizing the Kullback-Leibler (KL) divergence $\text{KL}[q||p]$ between the variational posterior distribution q and the true posterior distribution p . An exact evaluation of the KL divergence is not feasible. So, we convert the minimization of the KL divergence to the maximization of Evidence Lower Bound (ELBO) (\mathcal{L}) whose evaluation is feasible. ELBO is evaluated as

$$\mathcal{L} = \sum_{\ell=1}^L \mathbb{E}_{q(\mathbf{g}_\ell^\ell, \mathbf{u}_\ell)} [\log \underbrace{p(\mathbf{y}_\ell | \mathbf{g}_\ell^\ell)}_{\text{likelihood}}] + KL[q(\mathbf{u}_\ell) || p(\mathbf{u}_\ell | \mathbf{Z}_\ell)]. \quad (25)$$

One can refer to [47] for detailed proof. We optimize for the variational parameters $(\mathbf{m}_\ell, \mathbf{S}_\ell, \mathbf{Z}_\ell)$ and the hyperparameters of the kernels (k_ℓ) using the Adam optimization method.

The prediction process in Deep Gaussian Processes (DGP) follows a recursive transfer of results from the previous layer into the next layer. We draw samples at each layer from the approximate distribution defined by equation 22. We fuse the prediction parameter X^* in intermediate layers with the sample drawn from the previous layer. We then use the fused parameter to draw a sample from that layer. We repeat this process until we have S samples for each layer. We use the collected samples to compute the posterior mean and variance through a Monte Carlo approximation. Figure 4b visually illustrates these prediction steps for three fidelity models, depicting the recursive transfer and sampling process at each layer.

One can improve the surrogate by choosing a more structured kernel. Cutajar et al. [26] suggest using the kernel defined in equation 14. We call this method the Non-linear Auto-Regressive Deep Gaussian Process (NARDGP). We suggest using the delay term to further add useful information. We call the method Deep Gaussian Process with Delay Fusion (DGPDPF). We further improve upon the suggested method using a structured kernel as mentioned in equation 18. We named the resulting method Deep Gaussian Process with Delay Fusion and Composite kernel (DGPDPFC).

Methods discussed in Section 3.1.2 and Section 3.1.3 have similar ideas but different approaches to approximating the posterior distribution. In future discussions, we will sometimes mention the methods in Section 3.1.2 as full GP to differentiate it from DGP because they do not involve sparse variational approximation of the posterior.

3.2. Adaptive model improvement

In many cases, the available training data may not provide sufficient information to accurately model predictions, resulting in poor performance. One approach to enhance the model is adding more data points to the training set. This comes with the added cost of evaluating the corresponding function at the newly added training data locations. Therefore, judiciously choosing where to add points is crucial. The goal is to ensure that we capture the essential features of the target functions using as few training data points as possible.

In this section, we discuss an adaptive algorithm [48]. It is designed to enhance training datasets applicable to any of the methods mentioned earlier, all of which utilize combinations of Gaussian processes. This adaptive algorithm improves the overall surrogate prediction by focusing on one Gaussian process at a time. Suppose we are training a GP using n data points $\mathcal{D}_n = \{X_1, X_2, \dots, X_n\}$. Let y^{n+1} be the value of a sample from GP at X^{n+1} . We select the next point X_{n+1} where we gain maximum information which is defined as follows

$$\mathcal{I}(X_{n+1}|\mathcal{D}_n) = H(y^*|\mathcal{D}_n) - H(y^*|\mathcal{D}_n \cup X_{n+1}), \quad (26)$$

where $H(y^{n+1}|\mathcal{D}_n)$ is the conditional entropy. We suggest a greedy algorithm to selecting the next point X_{n+1} by solving the optimization problem

$$X_{n+1} = \underset{X_{n+1}}{\operatorname{argmax}} \mathcal{I}(X_{n+1}|\mathcal{D}_n). \quad (27)$$

Upon simplification, it can be shown that for GPs, the gain in information is directly proportional to the logarithm of the posterior predicted standard deviation. Thus, the optimization problem stated in equation 27 can be converted to

$$X_{n+1} = \operatorname{argmax}_{X^*} \sigma_p(X^* | \mathcal{D}_n), \quad (28)$$

where σ_p is the predicted standard deviation. We can keep adding new points for a particular fidelity level until the maximum posterior standard deviation is below a certain specified threshold level. After that, we can move to the next level until we cover all the fidelity levels.

There are certain drawbacks to this method. Firstly, this method adds one point at a time, limiting its overall speed. Secondly, the suggested points are always towards the boundary in high-dimensional space. Moreover, there are many other methods suggested in the literature that rely on other metrics to sequentially improve the Gaussian Process surrogate [49, 50]. We leave the discussion and effects of those algorithms on multi-fidelity GP as future works.

4. Numerical Examples

4.1. Academic Examples

Before applying the problem to real-world scenarios, we test different multi-fidelity surrogate modeling methods on some academic problems shown in Table 1, each posing a different modeling challenge. We deal with two fidelity cases and train the multi-fidelity GP surrogate using fifty low-fidelity and eight high-fidelity function evaluations. The high number of low-fidelity train data ensures low predictive variance for the low-fidelity Gaussian process. We use twenty-five and eight induction points in each level for all the cases involving deep Gaussian processes. Additionally, we also train a GP on high-fidelity data without multi-fidelity features. We then compare the multi-fidelity methods to the single-fidelity GP surrogate.

After this initial training, we run the adaptivity algorithm for ten steps to add ten additional high-fidelity data points. To ensure the robustness of the adaptivity algorithm, we perform multiple runs with distinct seed values. We take the average of the mean square error (MSE) and plot it against the corresponding number of high-fidelity function evaluations shown in Figure 6. The resulting plot visualizes the rate at which the respective algorithm convergences to the target high-fidelity curve. We also summarize the average MSE in Table 2.

We can observe from Table 2 that across all scenarios, the multi-fidelity methods outperform the single-fidelity GP by a considerable margin. This is because the multi-fidelity methods take advantage of the additional information the low-fidelity models provide. This shows the advantage of using multi-fidelity GP over single-fidelity GP. One can also conclude from Table 2 that methods involving the DGP have higher MSE as compared to the corresponding non-linear autoregressive methods with full GP, specifically when the number of high-fidelity function evaluations is increased. This

Table 1: Benchmark problems for the academic examples.

Benchmark name	Low-fidelity function (f_l)	High-fidelity function (f_h)
Linear transformation	$\sin(8\pi x)$	$0.8 \sin(8\pi x) + 0.3 \sin(2\pi x)$
Non-linear transformation		$\sin^2(8\pi x)$
Phase-shifted oscillation		$\sin^2(8\pi x + \frac{\pi}{10}) + \cos(4\pi x)$

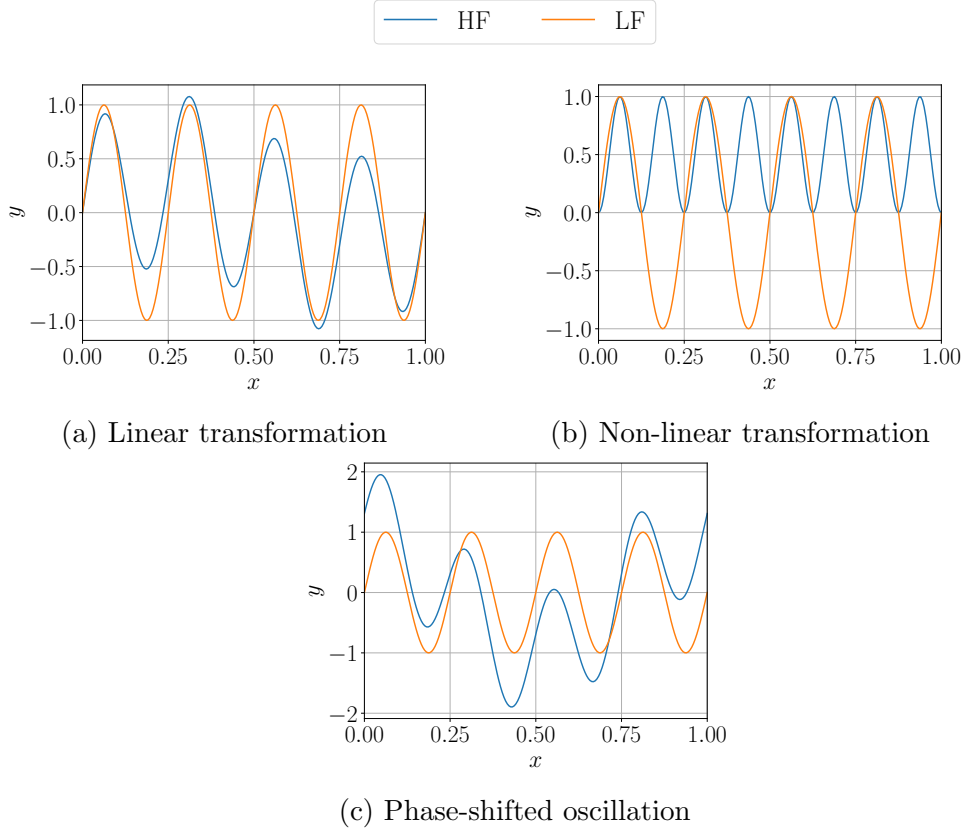


Figure 5: Visualization of different academic problems. The low and high-fidelity functions are visualized in each sub-plot by orange and blue curves, respectively.

is because DGP involves the approximation of the posterior distribution using some finite induction points.

We observe from Table 2 that the AR1 surrogate demonstrates the lowest MSE compared to the other methods in a linear transformation problem. The high-fidelity function of the linear scaling problem is written as constant scaling of the low-fidelity function with a non-linear additive term. This aligns with the assumed formulation of AR1 as described in Equation 10, giving AR1 an edge over the other methods. Every other non-linear method except DGPf also has a low MSE compared to the single-fidelity GP. The non-linear transformation methods involve expanding the surrogate dimensionality by introducing additional dimensions for low-fidelity function evaluations and/or delay terms. The assumed non-linear transformation of the low-fidelity term and

Table 2: Average value of MSE of each method in three different transformation scenarios, namely linear scaling, non-linear, and phase-shifted transformation as shown in Table 1, w.r.t. number of high-fidelity data points used in training.

MF method	Type of transformation					
	Linear scaling		Non-linear transformation		Phase-shifted oscillation	
	8 HF points	18 HF points	8 HF points	18 HF points	8 HF points	18 HF points
GP	2.36×10^{-1}	5.85×10^{-6}	2.45×10^{-1}	4.79×10^{-2}	7.32×10^{-1}	1.37×10^{-1}
AR1	4.46×10^{-5}	5.24×10^{-8}	3.73×10^{-1}	2.71×10^{-1}	2.25×10^{-1}	1.14×10^{-6}
NARGP	2.57×10^{-3}	3.02×10^{-6}	6.57×10^{-6}	2.00×10^{-6}	3.62×10^{-1}	2.19×10^{-1}
GPDPF	9.15×10^{-3}	3.11×10^{-6}	2.04×10^{-2}	1.68×10^{-6}	2.50×10^{-1}	6.99×10^{-6}
GPDFC	9.15×10^{-3}	3.11×10^{-6}	2.04×10^{-2}	1.68×10^{-6}	2.50×10^{-1}	6.99×10^{-6}
NARDGP	7.63×10^{-4}	2.88×10^{-5}	3.34×10^{-2}	1.14×10^{-4}	4.84×10^{-2}	7.16×10^{-2}
DGPDPF	1.09×10^{-1}	2.41×10^{-2}	8.81×10^{-4}	1.01×10^{-4}	17.0×10^{-1}	8.13×10^{-1}
DGPDFC	9.99×10^{-4}	2.52×10^{-5}	1.00×10^{-1}	1.37×10^{-4}	3.53×10^{-2}	2.52×10^{-4}

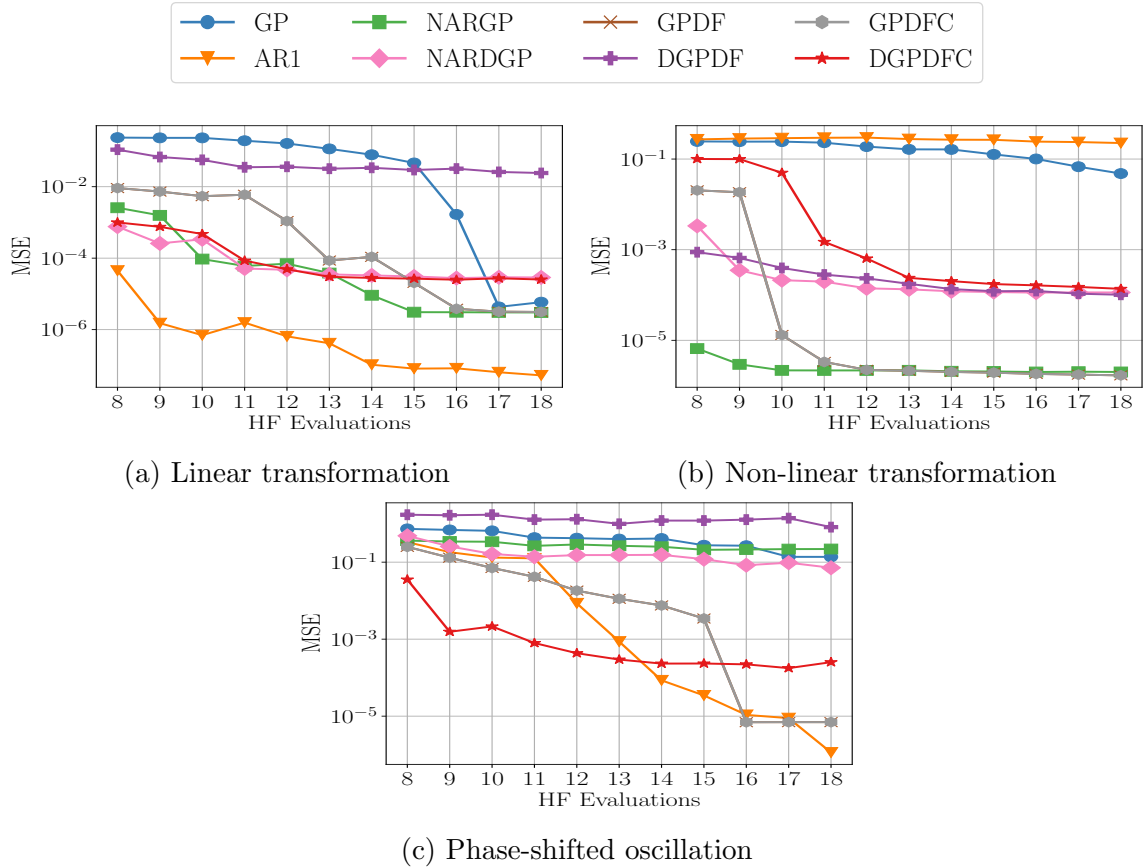


Figure 6: Average MSE evolution with respect to the number of high-fidelity function evaluations for different academic problems discussed in Section 4.1

the additional dimension of the delay term do not contribute meaningful information to the surrogate model in the case of the linear scaling problem. The additional kernel hyperparameters and the increase in dimension might require more evaluation points to represent the function. Therefore, non-linear auto-regressive methods underperform when compared to the AR1 method.

Table 2 shows that the AR1 surrogate cannot accurately capture non-linear transformations as expected because of the underlying linear formulation. In contrast, other methods tailored for non-linear transformations exhibit effective predictions with low MSE.

In the next example, we discuss the case where the high-fidelity function (f_h) and the low-fidelity function (f_l) are oscillating functions with phase differences. We observe such cases in our brain where neurons oscillate with certain frequencies where one uses the Hodgekin-Huxley model of them as a surrogate for real data [51].

The high-fidelity function can be expressed as a combination of sine and cosine terms, with the cosine term representing the derivative of the low-fidelity function. Consequently, methods incorporating delay terms are anticipated to outperform others.

$$\begin{aligned} f_h(x) &= \sin(8\pi x) \cos\left(\frac{\pi}{10}\right) + \cos(8\pi x) \sin\left(\frac{\pi}{10}\right) + \cos(4\pi x) \\ &= \cos\left(\frac{\pi}{10}\right) f_l(x) + \sin\left(\frac{\pi}{10}\right) \frac{df_l}{dx} + \cos(4\pi x). \end{aligned}$$

NARGP and NARDGP exhibited less accurate fitting to the target high-fidelity curve, as their formulations lack derivatives directly applicable to this example. Conversely, methods incorporating delay terms in surrogate modeling demonstrated accurate predictions by mimicking the derivative of the low-fidelity function. Surprisingly, AR1 also predicts accurately because the phase-shifted oscillation gets simplified into a linear scaling problem which is suitable for AR1.

DGPDF specifically underperforms for linear-scaling problems and phase-shifted oscillations even when other non-linear autoregressive methods had low MSE. Only one kernel is responsible for capturing all the features of the target high-fidelity function. In other non-linear autoregressive methods, the kernel has a well-defined structure. This underscores the significance of a well-defined kernel structure specifically for DGP.

The three presented cases with different transformations underscore the variability in the performance of different methods across diverse scenarios. This observation emphasizes the necessity of a judicious selection of methods based on the specific characteristics of the modeled system. Prior knowledge about the system can significantly contribute to choosing the most suitable methodology.

4.2. Terramechanical example

Terramechanics explores the interaction of wheels with the underlying soil [52]. There are no exact physical formulas that describe the process of the wheel-soil interaction because of the non-linearity of the pressure-sinkage relations in the soft soil [53]. Therefore, to simulate wheel movement on soft sands and to estimate forces and

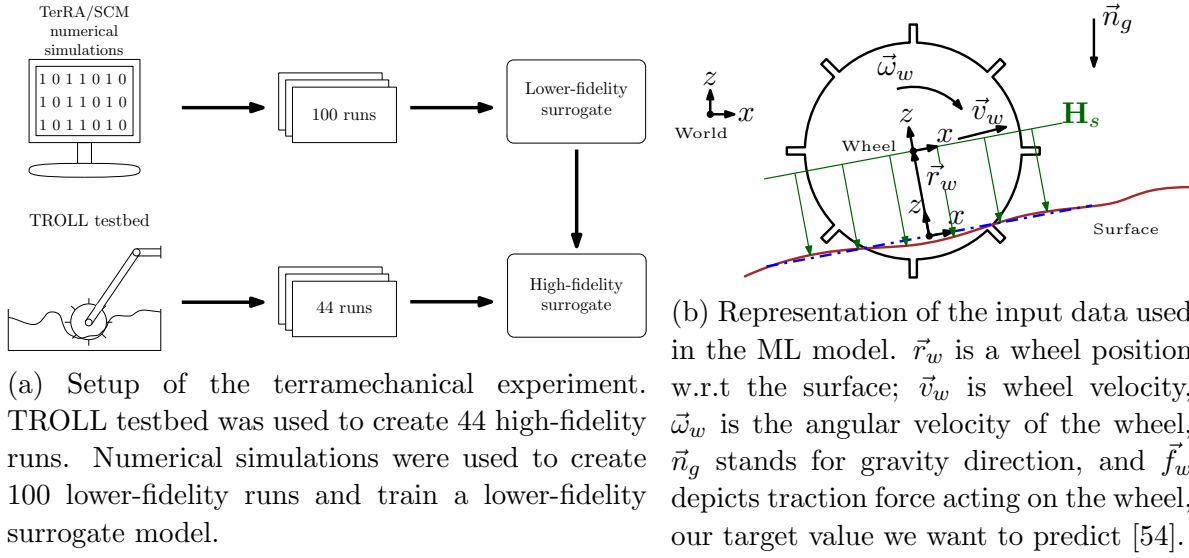


Figure 7: Different data generators were simulating the wheel-soil interaction for the same scenarios, thus creating an ideal case for multi-fidelity modeling.

Variable	Specification	Description
\vec{r}_w	3D Input	Wheel position w.r.t the surface
\vec{v}_w	3D Input	Translational velocity of the wheel
$\vec{\omega}_w$	3D Input	Angular velocity of the wheel
\vec{n}_g	3D Input	Gravity acting on the wheel
f	1D Output	Traction force acting on the wheel

Table 3: Description of the variables from the terramechanical dataset.

torques acting on that wheel, we exploit a wide range of numerical simulations, each one with a different level of discretization, assumptions, accuracy, and computational time. This makes terramechanics a good example of multi-fidelity modeling because different numerical models are dedicated to simulating the same physical process, but with different levels of fidelity.

As data sources for lower-fidelity models for this experiment, we used two types of models, both of which were developed at the Terramechanical lab at German Aerospace Center (DLR): TerRA [55] and SCM [53]. As a high-fidelity data source, DLR’s testbed TROLL was used [56]. This testbed was initially created to test and validate simulations for the wheel-soil interactions for extraterrestrial rover projects like MMX [57] and Scout [58]. The 44 high-fidelity runs from TROLL were also replicated in the simulation environment to achieve both higher- and lower-fidelity observations on the same set of inputs. These runs included a variety of different steering and titling angles, different rotation and horizontal velocities, and different movement trajectories and accelerations. Movement scenarios for dataset creation were described and justified in the previous

paper [54]. The scheme of the overall setup of experiments is depicted in Figure 7a.

Lower fidelity simulations are iterative simulations with each next step dependent on the state of the previous step, and the output of the simulation depends on several technical variables, which do not have a clear physical meaning and were not included as the input features in the dataset. The Gaussian process surrogate takes only a subset of the input variables of the simulation into account. Surrogate GPs were trained to approximate the outputs of TerRA and SCM simulations, and 100 runs were used in both cases. Each run was 20 seconds long and with a sampling rate of 3 points per second. Runs were conducted with randomly selected surfaces, random sequences of acceleration and deceleration and random sequences of steering commands. The randomness of a surface is introduced by varying its profile. The profile for each run has 6 breaking points, where after each point the change is introduced. Changes are normally distributed with a mean of 0.1 m and a variance of 0.6 m and are gradually introduced at each breakpoint. Velocities of the wheel simulation are also distributed normally, with a mean of 1 m/s and variance of 2 m/s. Steering angle is distributed normally, with 0 mean and 0.5 radians variance. Both velocities and steering angle were changing each second of the simulation.

We split high-fidelity runs into 29 runs used for training purposes and 15 reserved for the test dataset. We varied sampling frequency from the high-fidelity data, to verify how performance changes with changes in the amount of high-fidelity data. Harvesting high-fidelity data is challenging and we want to minimize the damage to the model’s performance, therefore we want to analyze the degradation of the prediction accuracy with the decrease in the number of training points. Each next subset is nested into the subset with a bigger number of training points. This logic assures the consistency of training points in all down-sampling examples.

All signals were smoothed with a low-pass filter, as in reality spikes of traction force are considered to be noise because only constant application of a force can change the wheel’s movement. The data consists of 12-dimensional input and the objective function is a traction force acting on the wheel. A description of the variables is detailed in Table 3. A visualization of the wheel and features describing it is depicted in Figure 7b.

We tested all MF models discussed earlier on the terramechanical data. MF models shown here were trained with both SCM and TerRA surrogates being the lower-fidelity function and compared with each other. Table 4 shows the performance of different MF methods w.r.t. the number of high-fidelity data points used during the modeling and numerical simulation which generated the data for the lower-fidelity level. The performance of the MF models with the TerRA as a lower-fidelity data source shows marginally worse results than with SCM, however still very compatible. This is a positive result, given that the TerRA model is much simpler in implementation and less expensive in exploitation. Empirically, there is no advantage in building more than two layers of fidelity for regression tasks, if we can already approximate the medium-fidelity surrogate with high accuracy. And given an abundance of data on medium-fidelity levels, we can skip this step. Changing the number of high-fidelity points can give a hint of the lower

Table 4: Normalized MSE of each MF method for wheel-soil locomotion simulation, w.r.t. number of high-fidelity data points used in training. Simulations using SCM were used as a lower-fidelity layer. The normalized MSE of conventional numerical simulation models is presented for comparison. GP trained only on the high-fidelity points shows good results when we have a lot of data points, but its performance deteriorates with a decrease in the training points. While the MF method, especially NARDGP, can show good results with fewer high-fidelity points and has more consistent performance.

Method Name	Number of HF points									
	2085		1042		521		208		104	
	Numerical simulation used as a lower fidelity layer									
	TerRA	SCM	TerRA	SCM	TerRA	SCM	TerRA	SCM	TerRA	SCM
Single fidelity GP	0.149		0.239		0.331		0.32		0.756	
AR1	0.235	0.235	0.235	0.235	0.235	0.235	0.235	0.235	0.235	0.235
NARGP	0.184	0.211	0.2	0.23	0.218	0.726	0.251	0.256	0.724	0.726
GPDPF	0.41	0.207	0.327	0.231	0.3	0.218	0.25	0.233	0.724	0.724
GPDFC	0.41	0.207	0.327	0.231	0.3	0.218	0.25	0.233	0.724	0.724
NARDGP	0.116	0.122	0.12	0.118	0.123	0.12	0.118	0.117	0.121	0.125
DGPDF	0.726	0.726	0.726	0.726	0.726	0.726	0.726	0.726	0.726	0.726
DGPDFC	0.217	0.22	0.249	0.233	0.224	0.223	0.226	0.23	0.22	0.235
Conventional numerical simulations										
TerRA						799.92				
SCM						0.209				

acceptable bar when we construct the MF model.

Models comparison with different subsets is illustrated in Table 4. As we can see there, the multi-fidelity approach can handle the drastic decrease in the number of used high-fidelity points, while the single-fidelity model’s performance decreases, as expected. Several examples of best-performing methods, NARDGP, are shown in Figure 8.

One of the main advantages of MF for complex systems modeling is that it can decrease computational and time costs without decreasing prediction accuracy. MF ML model outperforms both conventional numerical simulations TerRA and SCM simultaneously by MSE and by the runtime. Figure 9 shows 44 runs for both the MF ML model and two baseline numerical simulations (TerRA and SCM), distributed by their MSE and runtime. As we can see, the performance of the MF ML model is indeed faster and more accurate than the conventional terramechanical models.

Another advantage of the Gaussian process is the implicit handling of uncertainties. Due to its Bayesian nature, we can propagate uncertainties easily using the equation 5. It is an important part of predicting complex simulation systems, both for operational and research purposes. This adds more depth to the understanding and explainability of predictions and makes ML black boxes more transparent.

As we can see in Figure 8, the main disadvantage of the NARDGP multi-fidelity prediction is that confidence intervals are too narrow, as discussed in Section 2.3. This is a critical flaw when it comes to deploying ML models in practice, especially for highly sensitive operations like extraterrestrial rovers, where the cost of error is very high due

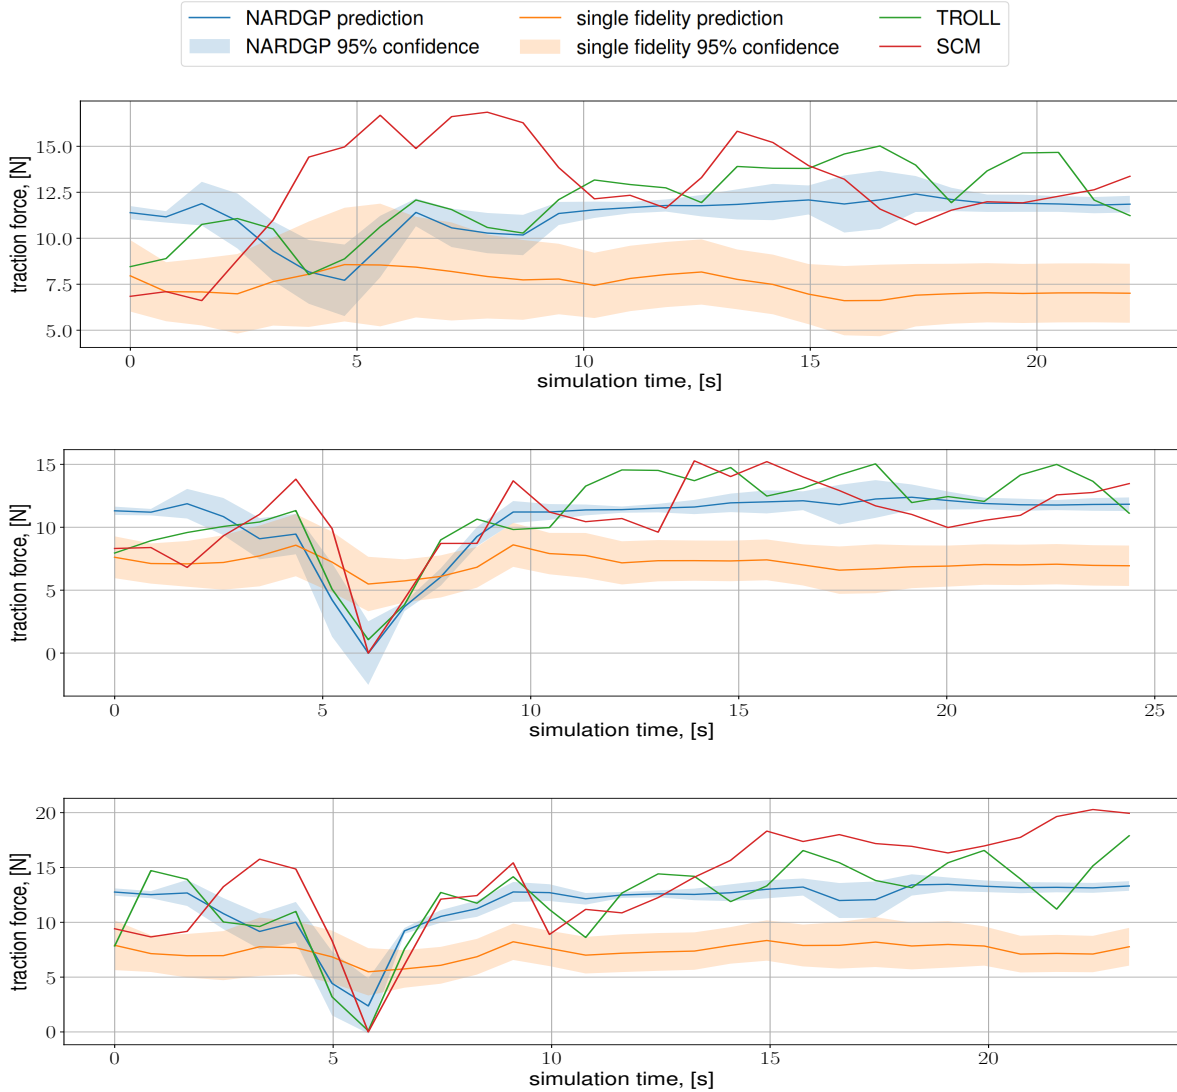


Figure 8: Visual comparison of the best performing MF ML model, NARDGP, on several test runs. It is compared against the ground truth from the TROLL sensors, single-fidelity model and the most popular numerical simulation SCM. The overall performance is reasonably better compared to other methods, but the model is visibly over-confident in its uncertainty estimation.

to the inability to quickly compensate for the wrong move or repair the rover. To solve this problem, we applied several calibration methods, mentioned in Section 2.3. Isotonic regression is an algorithm enabling quantile calibration, inspired by a Platt scaling algorithm from classification [38]. Beta calibration was chosen as an algorithm for achieving distribution calibration [37]. Normal calibration [36] introduces a scaling parameter for the predicted variance, preserving a Gaussian nature of the posterior prediction. A comparison of all calibration techniques on all of the test runs can be seen in Table 5. As an example, we took NARDGP trained with 385 high-fidelity data




Figure 9: Performance comparison of predictions from the conventional numerical simulation methods and MF ML model. The multi-fidelity approach is computationally efficient while maintaining a high level of precision.

points and calibrated it with the calibration mentioned above techniques. From this table, we can see that isotonic regression performs the best calibration, despite that it is only a global calibration of the marginal distribution. This means that the NARDGP model was constantly making overconfident predictions and global scaling of the variance solves the problem. These findings are consistent with the previous study, researching Gaussian process regression calibration for terramechanical data, but conducted only in the context of single-level medium-fidelity SCM data [61]. An example of a calibration on one test run could be seen in Figure 10, where the increased reliability of the predicted uncertainties is noticeable.

4.3. Plasma-physics example

Harnessing energy from plasma fusion promises to offer a clean alternative energy source. To achieve this goal, creating a self-sustained burning plasma is essential. Achieving a sustained burning plasma remains elusive due to physical and technological challenges.

Table 5: Mean pinball loss measures the errors of over- and under-confidence for each quantile individually [59]. Negative loglikelihood (NLL) measures how likely observed data was generated by the model [37]. Expected normalized calibration error (ENCE) orders and splits instances by variance into bins and calculates the scaled difference between the predicted error and variance in each bin [39]. Mean predictive interval width (MPIW) measures the sharpness of the prediction and is defined for each confidence interval α as a width of this interval [60]. Normal and isotonic calibration have similar results almost by every metric, but normal calibration has a significantly worse sharpness metric, which leaves us only with Isotonic regression.

Calibration method	Calibration metrics			
	Pinball	NLL	ENCE	MPIW
Uncalibrated	0.696	6.648	2.232	0.093
Isotonic regression	0.424	1.932	0.237	0.422
Normal calibration	0.424	2.229	0.21	175.675
Beta calibration	0.484	2.74	0.62	0.187

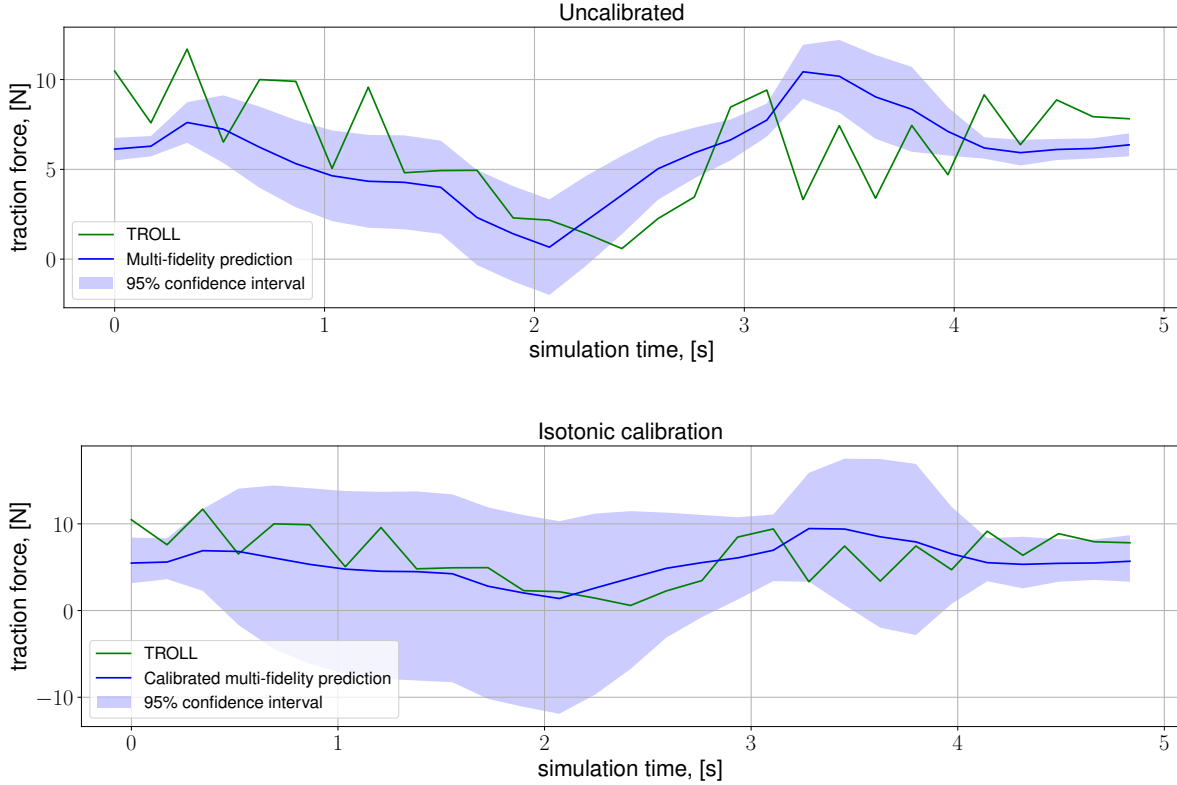


Figure 10: Original uncalibrated prediction is visibly overconfident because the ground truth traction force from TROLL quite often finds itself outside even the 95% confidence interval. By applying Isotonic regression, we achieve quantile calibrated prediction and here predicted confidence intervals capturing ground truth much better.

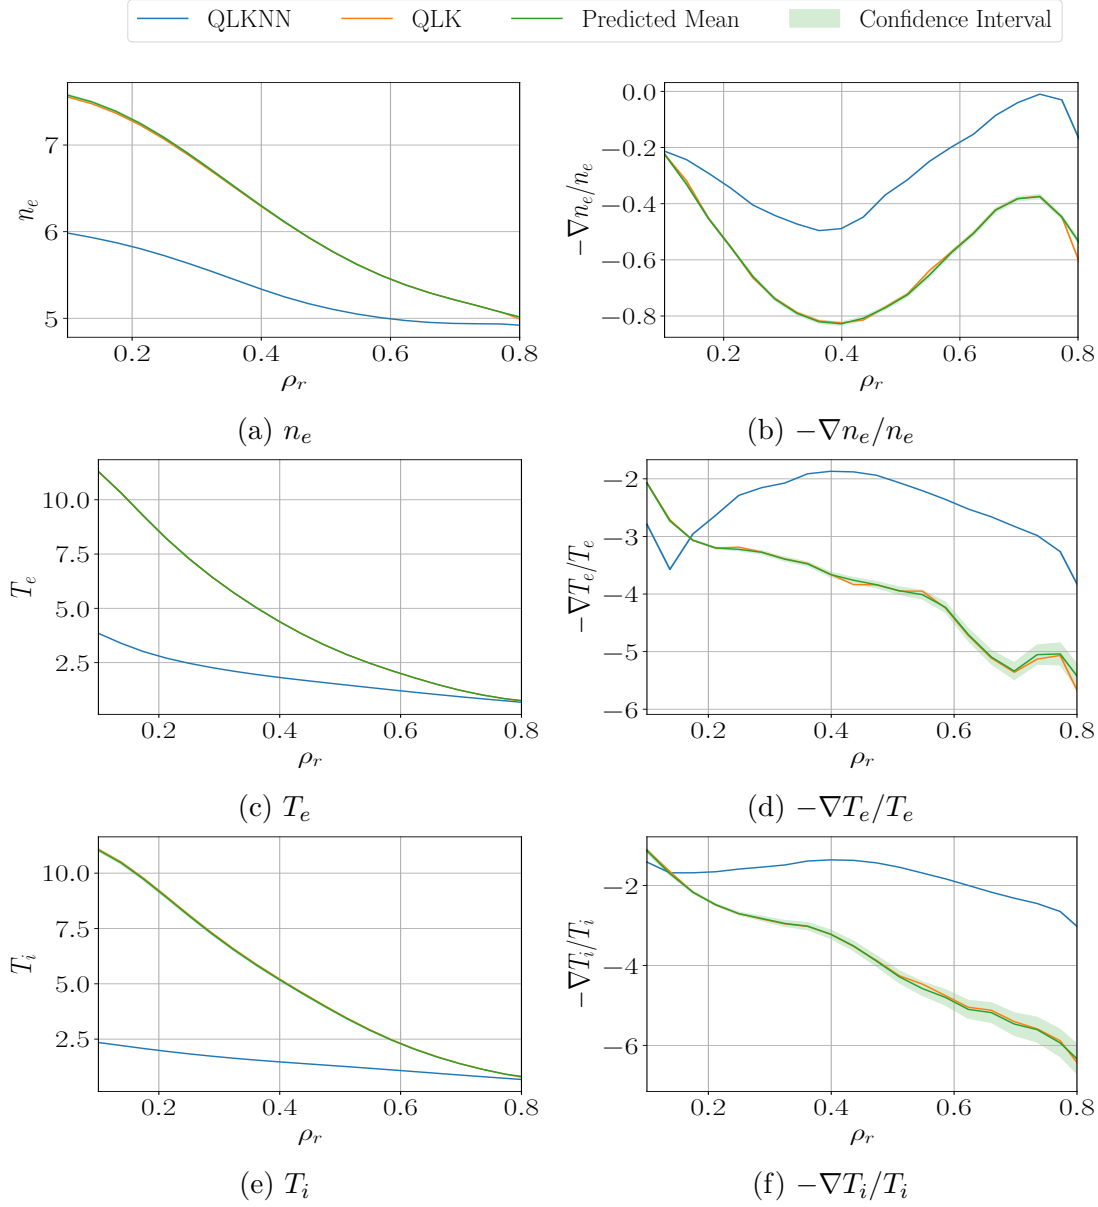


Figure 11: Steady state solution for different quantities simulated for the shot #34954. Parameters other than radial distance (ρ_r) are set as 1. The blue and orange curves show the simulation results for QLKNN (low-fidelity) and QLK (high-fidelity), respectively. Green curves show the predicted mean and 95% confidence interval for the NARGP surrogate.

One prominent physical hurdle is the occurrence of small-scale fluctuations in confined plasma, which lead to energy loss. These small-scale fluctuations are referred to as microturbulence. Mitigating energy losses caused by microturbulence is a significant challenge within the plasma physics community.

This work simulates plasma fusion using the Automated System for TRANsport Analysis (ASTRA) [62, 63] modeling suite. ASTRA solves four 1-dimensional transport

Table 6: List of parameters for plasma microturbulence surrogate modeling.

Parameter Name	Range
Radial distance (ρ_r)	[0.1, 0.8]
Initial T_i scaling	[0.9, 1.1]
Initial T_e scaling	[0.9, 1.1]
Initial n_e scaling	[0.9, 1.1]
Toroidal velocity scaling	[0.9, 1.1]
Safety factor scaling	[0.9, 1.1]

equations for electron density (n_e), electron temperature (T_i), ion temperature (T_i) and poloidal flux (Ψ). Interested readers can refer to [64] for details. It uses different turbulent solvers to predict the density and temperature profiles given some initial conditions (initial profiles of T_i , T_e , and n_e) and the last-closed-flux surface. The calculation of the turbulent flux is one of the computationally expensive parts of the code. One can couple ASTRA with different turbulent solvers. In this work, we will use QuaLiKiz (QLK) [65] and QuaLiKiz Neural Network (QLKNN) [66, 67] as the two subroutines. QLK is a quasi-linear turbulence solver for the linearized gyrokinetic Vlasov equation. QLKNN is a neural network surrogate trained on a 10-dimensional Latin hypercube for a quick flux evaluation [67]. ASTRA simulation with QLK is the high-fidelity model, whereas ASTRA with QLKNN is the low-fidelity model.

We simulate a plasma discharge until it reaches a steady state. At the steady state, the heat and particle fluxes match the integrated sources at all the radial locations. In this work we simulate shot #34954 until it reaches a steady state. We train the surrogate on a six-dimensional space mentioned in Table 6. We create six surrogates for six different quantities, namely, steady-state T_i , T_e , and n_e , and the corresponding negative derivative of the logarithm of each quantity ($-\nabla T_i/T_i$, $-\nabla T_e/T_e$, and $-\nabla n_e/n_e$). We normalize each quantity before training the surrogate. We use 40 high-fidelity and 400 low-fidelity training points.

Figure 11 shows the simulation results for QLK (high-fidelity model) and QLKNN (low-fidelity model) when parameters other than radial distance (ρ_r) are set as 1. The figure also shows the predicted results for multi-fidelity surrogates using NARGP. Table 7 shows the spatial average of mean square error for each surrogate.

We observe that non-linear auto-regressive multi-fidelity surrogates perform better than the single-fidelity GP. However, the difference between single-fidelity GP and multi-fidelity surrogates is not too significant. This may be due to the underlying simple structure that single-fidelity GP was able to learn relatively easily. AR1 struggles to model the quantity $-\nabla n_e/n_e$, which suggests that the relation between low-fidelity and high-fidelity could be non-linear for the case of $-\nabla n_e/n_e$. Moreover, DGPDPF also has a high MSE for $-\nabla n_e/n_e$, whereas DGPDPFC fits the quantity accurately. This re-iterates the significance of structured kernels, especially for deep Gaussian processes.

Table 7: Spatial average value of MSE for different surrogates across different quantities for the plasma microturbulence surrogate modelling.

Method Name	Quantity of Interest					
	T_i	T_e	n_e	$-\nabla T_i/T_i$	$-\nabla T_e/T_e$	$-\nabla n_e/n_e$
Single-fidelity GP	8.57×10^{-3}	8.68×10^{-4}	1.07×10^{-2}	2.26×10^{-1}	1.04×10^{-1}	7.83×10^{-2}
AR1	1.73×10^{-2}	2.26×10^{-3}	3.31×10^{-2}	4.89×10^{-1}	8.21×10^{-2}	1.33
NARGP	4.46×10^{-3}	8.13×10^{-4}	7.12×10^{-3}	1.39×10^{-1}	6.9×10^{-2}	9.57×10^{-3}
GPDPF	3.07×10^{-3}	8.13×10^{-4}	6.3×10^{-3}	1.94×10^{-1}	6.9×10^{-2}	1.44×10^{-2}
GPDFC	3.07×10^{-3}	8.13×10^{-4}	6.3×10^{-3}	1.94×10^{-1}	6.9×10^{-2}	1.44×10^{-2}
NARDGP	2.89×10^{-3}	8.56×10^{-4}	6.15×10^{-3}	4.41×10^{-1}	1.65×10^{-1}	1.72×10^{-1}
DGPDPF	4.52×10^{-3}	8.58×10^{-4}	5.42×10^{-3}	1.83×10^{-1}	8.06×10^{-2}	3.57
DGPDPFC	2.88×10^{-3}	8.58×10^{-4}	5.42×10^{-3}	1.38×10^{-1}	1.82×10^{-1}	4.37×10^{-2}

5. Conclusions

This paper describes different multi-fidelity Gaussian process surrogate modeling methods. We extend non-linear autoregressive models for full GP to accommodate cases that involve more than two fidelities. Using a structured kernel with delay terms, we also suggest a new family of multi-fidelity GP models (GPDPF and DGPDPF). Finally, we test all the modeling methods on different academic and real-world problems.

We observe that not all the multi-fidelity methods performed well under all the scenarios. The quality of prediction depends upon the underlying relationship between the low and high-fidelity models. Prior knowledge about that would help us choose the correct modeling method. In many scenarios, that relationship is not known. Therefore, after constructing the surrogate, one must carefully validate its predictions. We observe that the structured kernel significantly improves models involving deep Gaussian processes.

We also conclude that the multi-fidelity simulations can preserve the performance of higher-fidelity simulations but reach the speed and cost of low-fidelity ones. This can aid different research fields in analyzing the underlying system or improving the algorithm using outer loop methods [7].

Acknowledgments

Research by Vladyslav Fediukov, Michael Bergmann and Kislaya Ravi is funded by the Helmholtz Association under the "Munich School for Data Science - MuDS". Felix Dietrich acknowledges funding by the DFG project no. 468830823, and also association to DFG-SPP-229.

Reference

- [1] Ping Jiang, Qi Zhou, Xinyu Shao, Ping Jiang, Qi Zhou, and Xinyu Shao. *Surrogate-model-based design and optimization*. Springer, 2020.
- [2] Yolanda Mack, Tushar Goel, Wei Shyy, and Raphael Haftka. Surrogate model-based optimization framework: a case study in aerospace design. *Evolutionary computation in dynamic and uncertain environments*, pages 323–342, 2007.
- [3] Chong Wang, Xin Qiang, Menghui Xu, and Tao Wu. Recent advances in surrogate modeling methods for uncertainty quantification and propagation. *Symmetry*, 14(6):1219, 2022.
- [4] Ionuț-Gabriel Farcaș, Benjamin Uekermann, Tobias Neckel, and Hans-Joachim Bungartz. Nonintrusive uncertainty analysis of fluid-structure interaction with spatially adaptive sparse grids and polynomial chaos expansion. *SIAM Journal on Scientific Computing*, 40(2):B457–B482, 2018.
- [5] Raul Yondo, Esther Andrés, and Eusebio Valero. A review on design of experiments and surrogate models in aircraft real-time and many-query aerodynamic analyses. *Progress in aerospace sciences*, 96:23–61, 2018.
- [6] Aditya Balu, Soumik Sarkar, Baskar Ganapathysubramanian, and Adarsh Krishnamurthy. Physics-aware machine learning surrogates for real-time manufacturing digital twin. *Manufacturing Letters*, 34:71–74, 2022.
- [7] Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *Siam Review*, 60(3):550–591, 2018.
- [8] Michael B Giles. Multilevel monte carlo methods. *Acta numerica*, 24:259–328, 2015.
- [9] Leo Wai-Tsun Ng and Michael Eldred. Multifidelity uncertainty quantification using non-intrusive polynomial chaos and stochastic collocation. In *53rd aiaa/asme/asce/ahs/asc structures, structural dynamics and materials conference 20th aiaa/asme/ahs adaptive structures conference 14th aiaa*, page 1852, 2012.
- [10] Ionuț-Gabriel Farcaș, Benjamin Peherstorfer, Tobias Neckel, Frank Jenko, and Hans-Joachim Bungartz. Context-aware learning of hierarchies of low-fidelity models for multi-fidelity uncertainty quantification. *Computer Methods in Applied Mechanics and Engineering*, 406:115908, 2023.
- [11] Vivin Vinod, Ulrich Kleinekathöfer, and Peter Zaspel. Optimized multifidelity machine learning for quantum chemistry. *Machine Learning: Science and Technology*, 2023.
- [12] Kislaya Ravi, Tobias Neckel, and Hans-Joachim Bungartz. Multi-fidelity no-u-turn sampling. *arXiv preprint arXiv:2310.02703*, 2023.
- [13] Mikkel Bue Lykkegaard, Tim J Dodwell, Colin Fox, Grigorios Mingas, and Robert Scheichl. Multilevel delayed acceptance mcmc. *SIAM/ASA Journal on Uncertainty Quantification*, 11(1):1–30, 2023.
- [14] Thomas P Prescott and Ruth E Baker. Multifidelity approximate bayesian computation. *SIAM/ASA Journal on Uncertainty Quantification*, 8(1):114–138, 2020.
- [15] Atul Agrawal, Kislaya Ravi, Phaedon-Stelios Koutsourelakis, and Hans-Joachim Bungartz. Multi-fidelity constrained optimization for stochastic black box simulators. *arXiv preprint arXiv:2311.15137*, 2023.
- [16] Faran Irshad, Stefan Karsch, and Andreas Döpp. Leveraging trust for joint multi-objective and multi-fidelity optimization. Technical report.
- [17] Marjuka F Lazin, Christian R Shelton, Simon N Sandhofer, and Bryan M Wong. High-dimensional multi-fidelity bayesian optimization for quantum control. *Machine Learning: Science and Technology*, 4(4):045014, 2023.
- [18] Stefan Görtz, Mohammad Abu-Zurayk, Caslav Ilic, Tobias Wunderlich, Stefan Keye, Matthias Schulze, Thomas Klimmek, Christoph Kaiser, Özge Süelözgen, Thiemo Kier, Andreas Schuster, Sascha Dähne, Michael Petsch, Dieter Kohlgrüber, Jannik Häßy, Robert Mischke, Alexander Weinert, Philipp Knechtges, Sebastian Gottfried, Johannes Hartmann, and Benjamin Fröhler.

- Overview of collaborative multi-fidelity multidisciplinary design optimization activities in the dlr project victoria. In *AIAA Aviation Forum 2020*, editor, *AIAA Aviation 2020 Forum*, June 2020.
- [19] Mohammad Abu-Zurayk, Andrei Merle, Caslav Ilic, Stefan Keye, Stefan Görtz, Matthias Schulze, Thomas Klimmek, Christoph Kaiser, David Quero-Martin, Jannik Häßy, Richard-Gregor Becker, Benjamin Fröhler, and Johannes Hartmann. Sensitivity-based multifidelity multidisciplinary optimization of a powered aircraft subject to a comprehensive set of loads. In *AIAA Aviation 2020 Forum*, June 2020.
- [20] Alexander Sebastian Zakrzewski, Fabian Lange, and Rene Walter Hollmann. Multi-fidelity aerodynamic design process for moveables at dlr virtual product house. In *AIAA Aviation 2022 Forum*. American Institute of Aeronautics and Astronautics, Inc., June 2022. Video Presentation: <https://doi.org/10.2514/6.2022-3938>.vid.
- [21] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [22] Marc C Kennedy and Anthony O’Hagan. Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13, 2000.
- [23] Loic Le Gratiet. *Multi-fidelity Gaussian process regression for computer experiments*. PhD thesis, Université Paris-Diderot-Paris VII, 2013.
- [24] Seungjoon Lee, Felix Dietrich, George E Karniadakis, and Ioannis G Kevrekidis. Linking gaussian process regression with data-driven manifold embeddings for nonlinear data fusion. *Interface focus*, 9(3):20180083, 2019.
- [25] Paris Perdikaris, Maziar Raissi, Andreas Damianou, Neil D Lawrence, and George Em Karniadakis. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2198):20160751, 2017.
- [26] Kurt Cutajar, Mark Pullin, Andreas Damianou, Neil Lawrence, and Javier González. Deep gaussian processes for multi-fidelity modeling. *arXiv preprint arXiv:1903.07320*, 2019.
- [27] Daniel G Krige. A statistical approach to some basic mine valuation problems on the witwatersrand. *Journal of the Southern African Institute of Mining and Metallurgy*, 52(6):119–139, 1951.
- [28] Klaus-Robert Müller, Sebastian Mika, Koji Tsuda, and Koji Schölkopf. An introduction to kernel-based learning algorithms. In *Handbook of neural network signal processing*, pages 4–1. CRC Press, 2018.
- [29] Houman Owhadi and Gene Ryan Yoo. Kernel flows: From learning kernels from data into the abyss. *Journal of Computational Physics*, 389:22–47, 2019.
- [30] Jan David Hüwel, Fabian Berns, and Christian Beecks. Automated kernel search for gaussian processes on data streams. In *2021 IEEE International Conference on Big Data (Big Data)*, pages 3584–3588. IEEE, 2021.
- [31] David Duvenaud, James Lloyd, Roger Grosse, Joshua Tenenbaum, and Ghahramani Zoubin. Structure discovery in nonparametric regression through compositional kernel search. In *International Conference on Machine Learning*, pages 1166–1174. PMLR, 2013.
- [32] Daniel Horn, Jörg Stork, Nils-Jannik Schüßler, and Martin Zaefferer. Surrogates for hierarchical search spaces: The wedge-kernel and an automated analysis. In *Proceedings of the genetic and evolutionary computation conference*, pages 916–924, 2019.
- [33] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989.
- [34] Alexandre Capone, Armin Lederer, and Sandra Hirche. Gaussian process uniform error bounds with unknown hyperparameters for safety-critical applications. In *International Conference on Machine Learning*, pages 2609–2624. PMLR, 2022.
- [35] Alexandre Capone, Geoff Pleiss, and Sandra Hirche. Sharp calibrated gaussian processes. *arXiv preprint arXiv:2302.11961*, 2023.

- [36] Fabian Küppers, Jonas Schneider, and Anselm Haselhoff. Parametric and multivariate uncertainty calibration for regression and object detection. In *European Conference on Computer Vision*, pages 426–442. Springer, 2022.
- [37] Hao Song, Tom Diethe, Meelis Kull, and Peter Flach. Distribution calibration for regression. In *International Conference on Machine Learning*, pages 5897–5906. PMLR, 2019.
- [38] Volodymyr Kuleshov, Nathan Fenner, and Stefano Ermon. Accurate uncertainties for deep learning using calibrated regression. In *International conference on machine learning*, pages 2796–2804. PMLR, 2018.
- [39] Dan Levi, Liran Gispán, Niv Giladi, and Ethan Fetaya. Evaluating and calibrating uncertainty prediction in regression tasks. *Sensors*, 22(15):5540, 2022.
- [40] Kyle Cranmer, Johann Brehmer, and Gilles Louppe. The frontier of simulation-based inference. *Proceedings of the National Academy of Sciences*, 117(48):30055–30062, 2020.
- [41] Andreas Damianou and Neil D Lawrence. Deep gaussian processes. In *Artificial intelligence and statistics*, pages 207–215. PMLR, 2013.
- [42] Michalis Titsias and Neil D Lawrence. Bayesian gaussian process latent variable model. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 844–851. JMLR Workshop and Conference Proceedings, 2010.
- [43] Francis R Bach and Michael I Jordan. Predictive low-rank decomposition for kernel methods. In *Proceedings of the 22nd international conference on Machine learning*, pages 33–40, 2005.
- [44] Matthias Bauer, Mark van der Wilk, and Carl Edward Rasmussen. Understanding probabilistic sparse gaussian process approximations. *Advances in neural information processing systems*, 29, 2016.
- [45] Joaquin Quinero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.
- [46] David Burt, Carl Edward Rasmussen, and Mark Van Der Wilk. Rates of convergence for sparse variational gaussian process regression. In *International Conference on Machine Learning*, pages 862–871. PMLR, 2019.
- [47] Hugh Salimbeni and Marc Deisenroth. Doubly stochastic variational inference for deep gaussian processes. *Advances in neural information processing systems*, 30, 2017.
- [48] Julien Villemonteix, Emmanuel Vazquez, and Eric Walter. An informational approach to the global optimization of expensive-to-evaluate functions. *Journal of Global Optimization*, 44:509–534, 2009.
- [49] Daniel Busby. Hierarchical adaptive experimental design for gaussian process emulators. *Reliability Engineering & System Safety*, 94(7):1183–1193, 2009.
- [50] Zhihang Xu and Qifeng Liao. Gaussian process based expected information gain computation for bayesian optimal design. *Entropy*, 22(2):258, 2020.
- [51] Alan L Hodgkin and Andrew F Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500, 1952.
- [52] Mieczyslaw Gregory Bekker. Introduction to terrain-vehicle systems. part i: The terrain. part ii: The vehicle. 1969.
- [53] Fabian Buse. *Development and Validation of a Deformable Soft Soil Contact Model for Dynamic Rover Simulations*. PhD thesis, Tohoku University, 2022.
- [54] Vladyslav Fediukov, Felix Dietrich, and Fabian Buse. Multi-fidelity machine learning modeling for wheel locomotion. In *11th Asia-Pacific Regional Conference of the International society for terrain-vehicle systems, ISTVS 2022*. ISTVS, 2022.
- [55] Stefan Barthelmes. Terra: Terramechanics for real-time application. In *5th Joint International Conference on Multibody System Dynamics*, June 2018.
- [56] Fabian Buse, Tobias Bellmann, Roy Lichtenheldt, and Rainer Krenn. The dlr terramechanics robotics locomotion lab. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space*, June 2018.
- [57] Fabian Buse, Antoine Pignède, Jean Bertrand, Sébastien Goulet, and Sandra Lagabarre. Mmx

- rover simulation-robotic simulations for phobos operations. In *2022 IEEE Aerospace Conference (AERO)*, pages 1–14. IEEE, 2022.
- [58] Antoine Pignède, Walter Schindler, Roy Lichtenheldt, Bernhard Thiele, Manuel Schütt, and Dennis Franke. Toolchain for a mobile robot applied on the dlr scout rover. In *2022 IEEE Aerospace Conference (AERO)*, pages 1–15. IEEE, 2022.
- [59] Matteo Fasiolo, Simon N Wood, Margaux Zaffran, Raphaël Nedellec, and Yannig Goude. Fast calibrated additive quantile regression. *Journal of the American Statistical Association*, 116(535):1402–1412, 2021.
- [60] Jiayu Yao, Weiwei Pan, Soumya Ghosh, and Finale Doshi-Velez. Quality of uncertainty quantification for bayesian neural network inference. *arXiv preprint arXiv:1906.09686*, 2019.
- [61] Jana Huhne. Uncertainty quantification for gaussian processes. 2023.
- [62] Gregorij V Pereverzev and PN Yushmanov. Astra. automated system for transport analysis in a tokamak. 2002.
- [63] Emiliano Fable, C Angioni, FJ Casson, D Told, AA Ivanov, F Jenko, RM McDermott, S Yu Medvedev, GV Pereverzev, F Ryter, et al. Novel free-boundary equilibrium and transport solver with theory-based models and its validation against asdex upgrade current ramp scenarios. *Plasma Physics and Controlled Fusion*, 55(12):124028, 2013.
- [64] FL Hinton and Richard D Hazeltine. Theory of plasma transport in toroidal confinement systems. *Reviews of Modern Physics*, 48(2):239, 1976.
- [65] Clarisse Bourdelle. *Turbulent transport in tokamak plasmas: bridging theory and experiment*. PhD thesis, Aix Marseille Université, 2015.
- [66] Aaron Ho, Jonathan Citrin, Clarisse Bourdelle, Yann Camenen, Francis J Casson, Karel L van de Plassche, Henri Weisen, and JET Contributors. Neural network surrogate of qualikiz using jet experimental data to populate training space. *Physics of Plasmas*, 28(3), 2021.
- [67] Karel Lucas van de Plassche, Jonathan Citrin, Clarisse Bourdelle, Yann Camenen, Francis J Casson, Victor I Dagnelie, Federico Felici, Aaron Ho, Simon Van Mulders, and JET Contributors. Fast modeling of turbulent transport in fusion plasmas using neural networks. *Physics of Plasmas*, 27(2), 2020.

Appendix

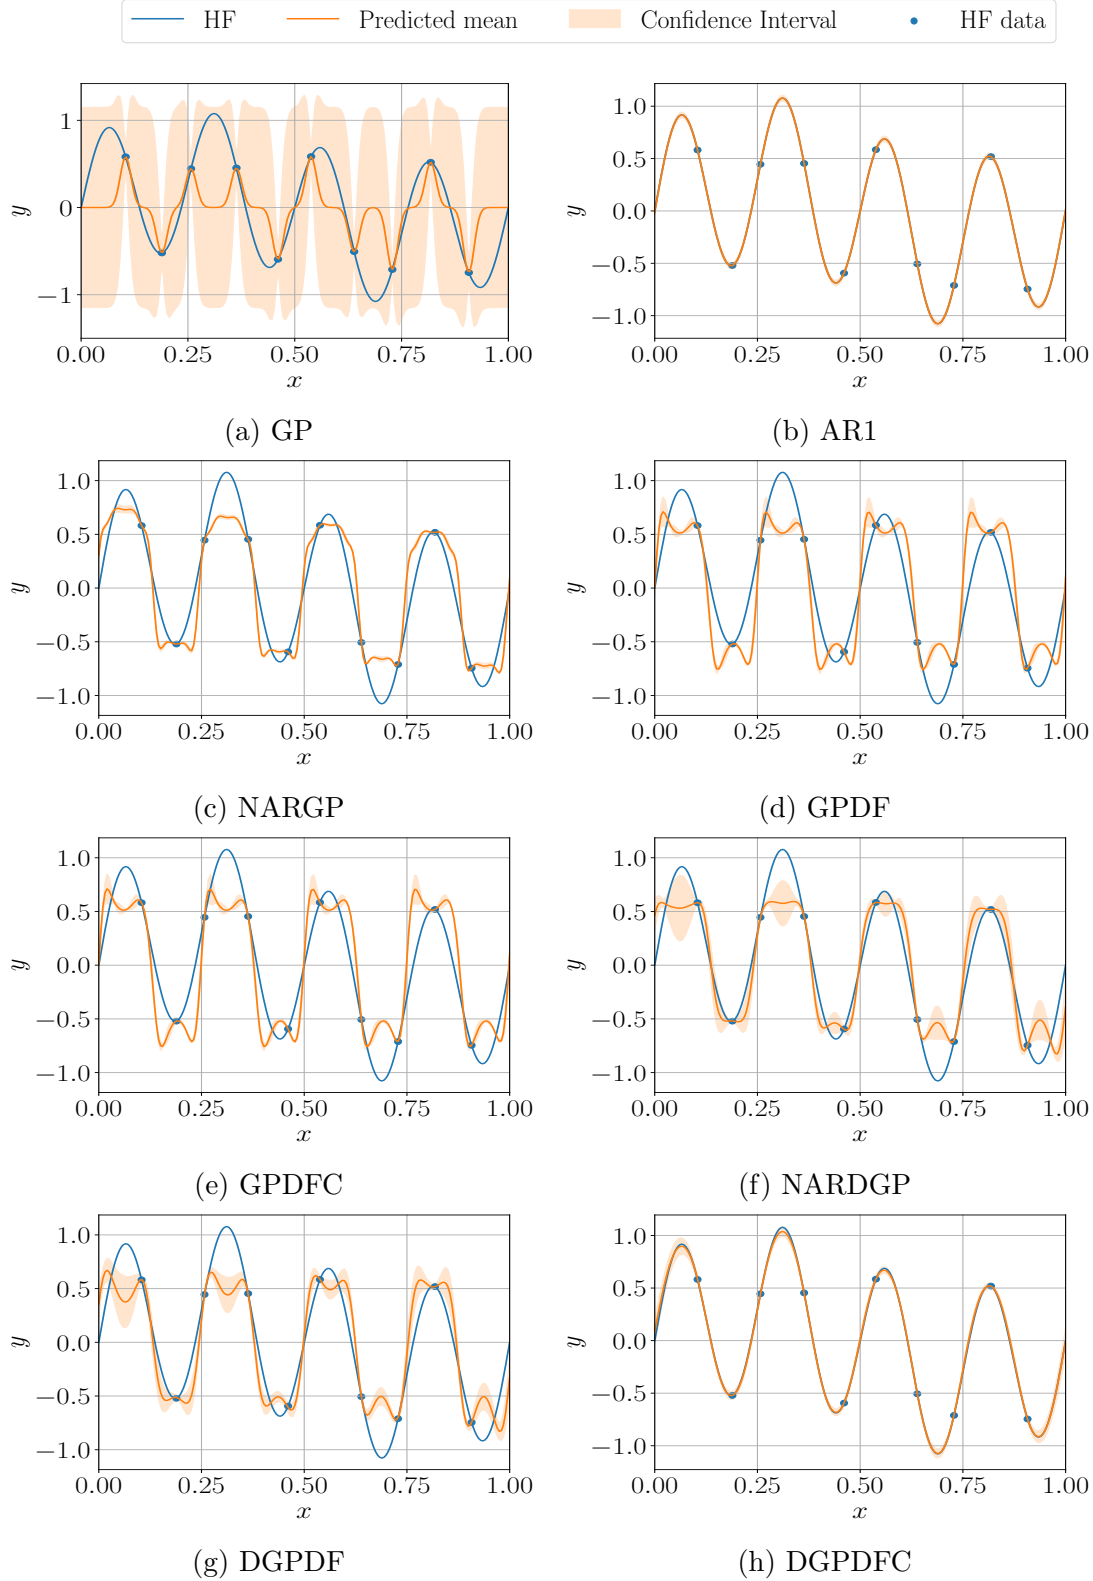


Figure 1: Comparison of different multi-fidelity Gaussian Process surrogate modeling methods on a linear transformation problem stated in Table 1. The plots show the high-fidelity function in the blue curve, the high-fidelity training data in blue dots, the predicted mean in the orange curve, and 95% confidence interval using the orange-shaded region.

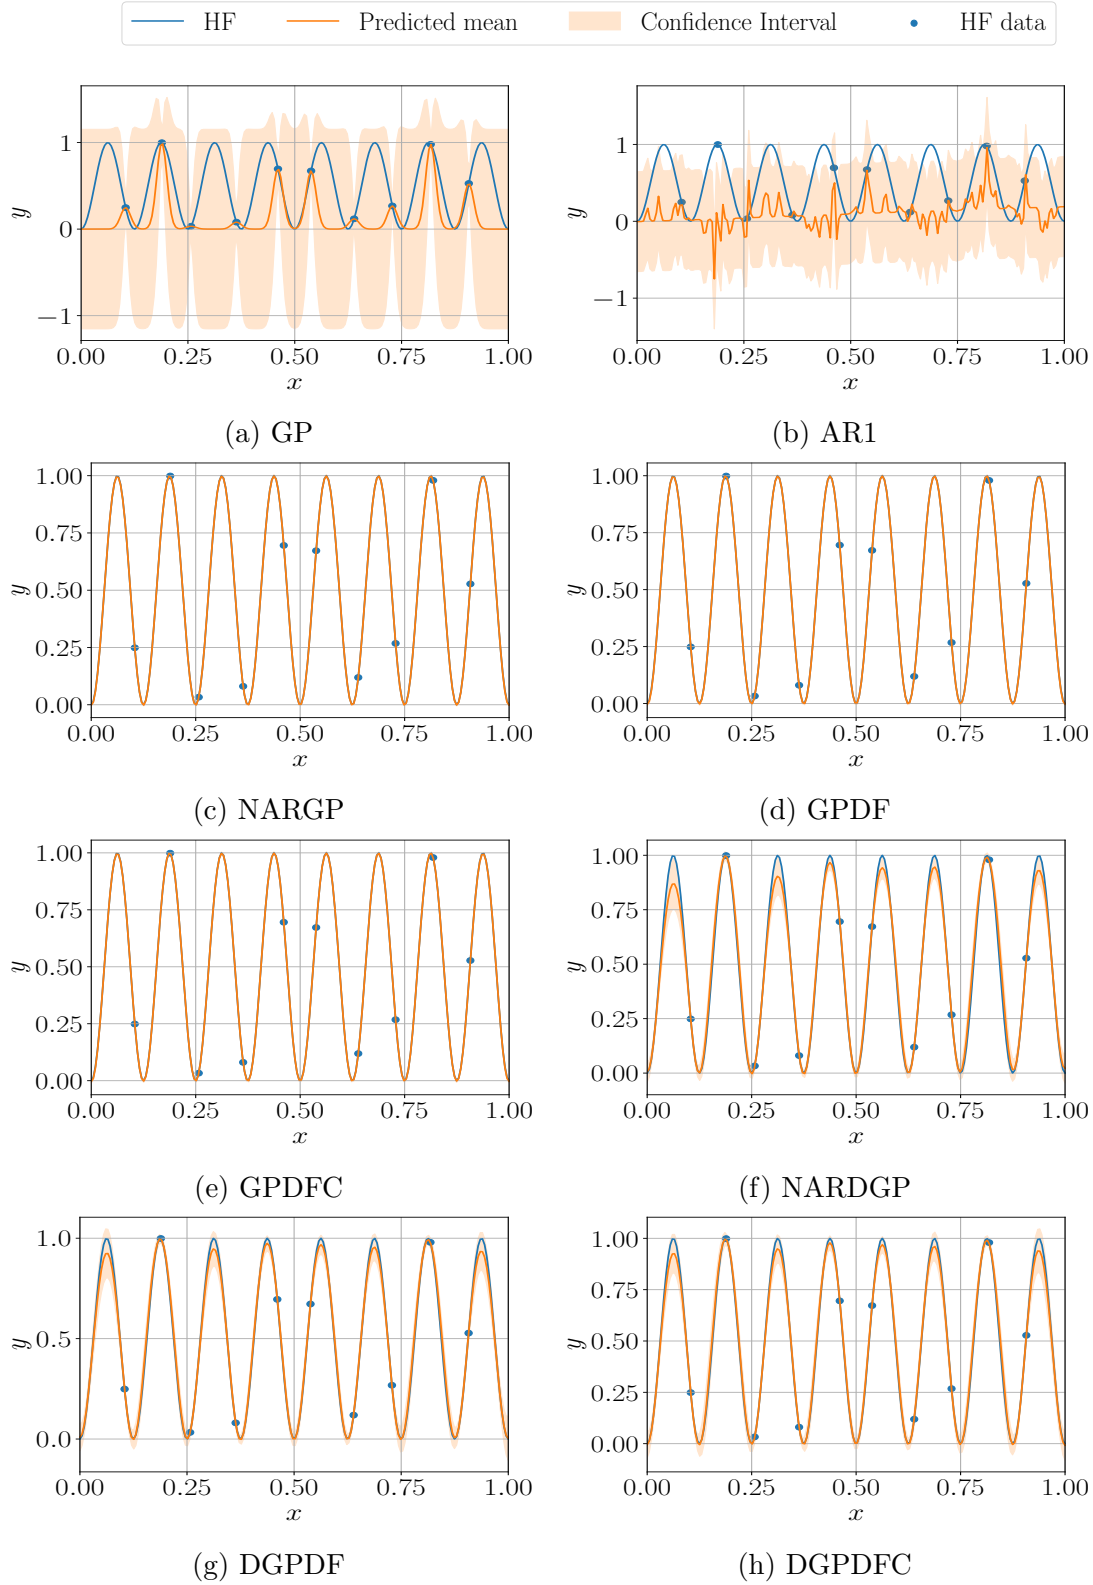


Figure 2: Comparison of different multi-fidelity Gaussian Process surrogate modeling methods on a non-linear transformation problem stated in Table 1. The plots show the high-fidelity function in the blue curve, the high-fidelity training data in blue dots, the predicted mean in the orange curve, and 95% confidence interval using the orange-shaded region.

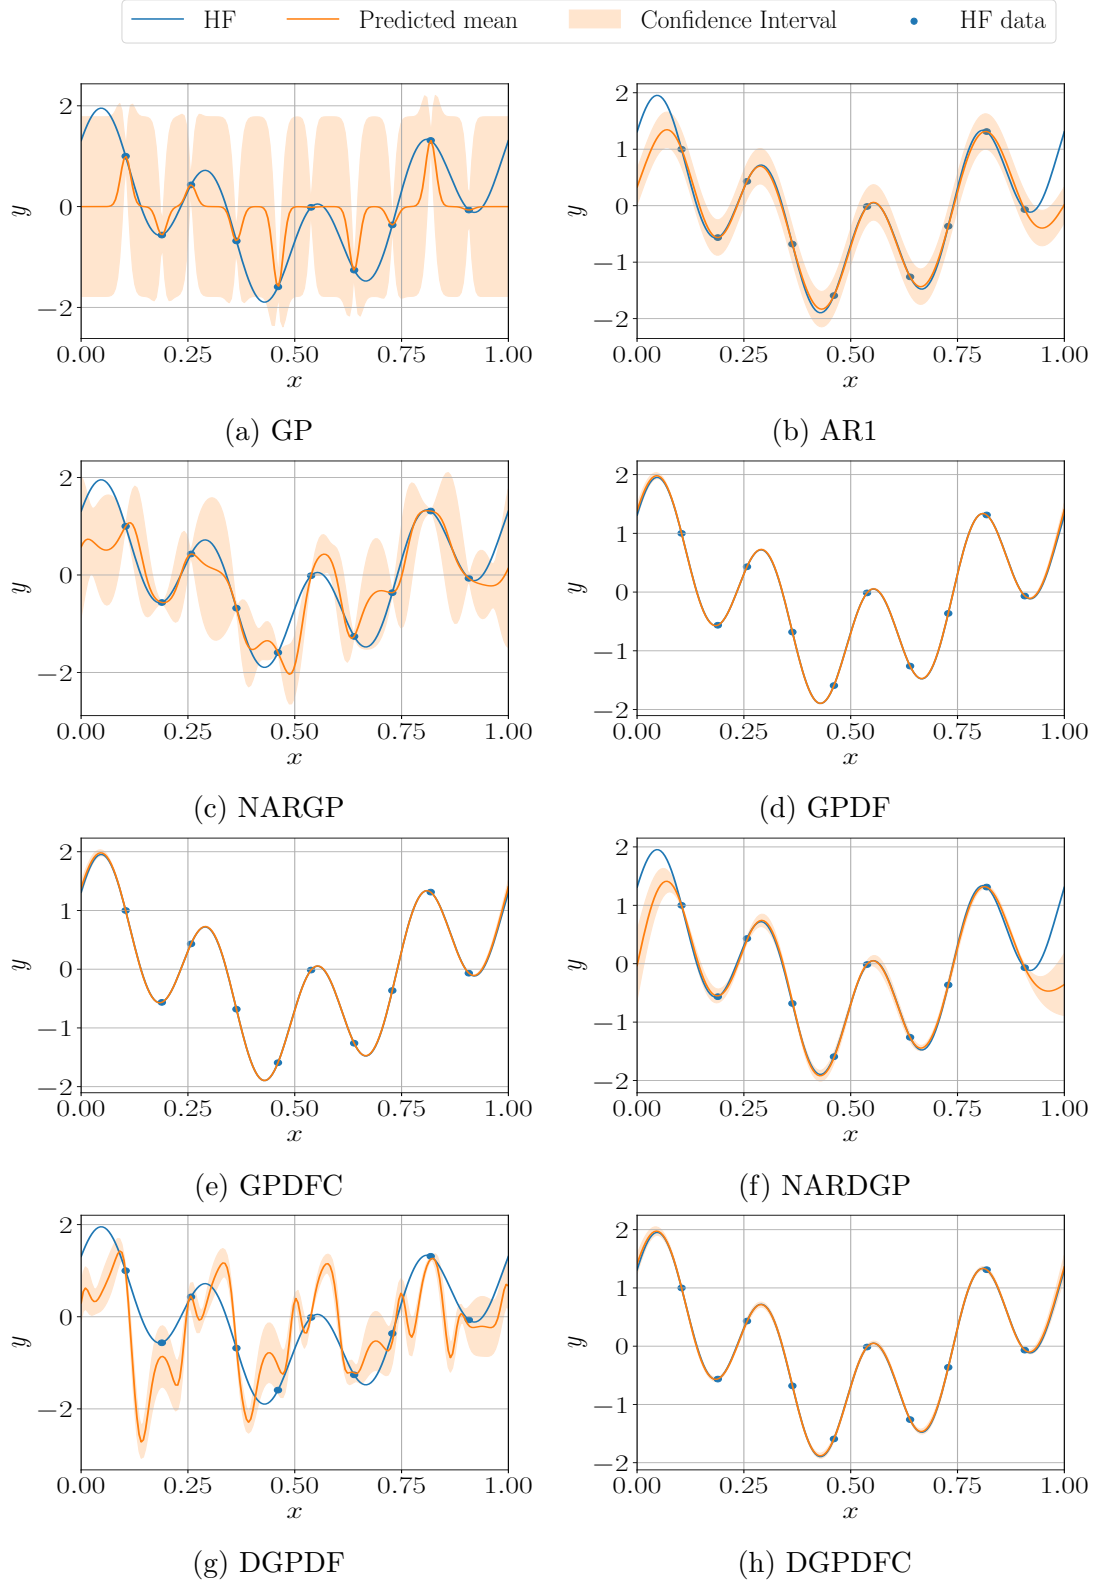


Figure 3: Comparison of different Multi-fidelity Gaussian Process surrogate modeling methods on a phase-shifted oscillation stated in Table 1. The plots show the high-fidelity function in the blue curve, the high-fidelity training data in blue dots, the predicted mean in the orange curve, and 95% confidence interval using the orange-shaded region.