

Active Learning for a Recursive Non-Additive Emulator for Multi-Fidelity Computer Experiments

Junoh Heo and Chih-Li Sung*

Michigan State University

Abstract

Computer simulations have become essential for analyzing complex systems, but high-fidelity simulations often come with significant computational costs. To tackle this challenge, multi-fidelity computer experiments have emerged as a promising approach that leverages both low-fidelity and high-fidelity simulations, enhancing both the accuracy and efficiency of the analysis. In this paper, we introduce a new and flexible statistical model, the *Recursive Non-Additive (RNA) emulator*, that integrates the data from multi-fidelity computer experiments. Unlike conventional multi-fidelity emulation approaches that rely on an additive auto-regressive structure, the proposed RNA emulator recursively captures the relationships between multi-fidelity data using Gaussian process priors without making the additive assumption, allowing the model to accommodate more complex data patterns. Importantly, we derive the posterior predictive mean and variance of the emulator, which can be efficiently computed in a closed-form manner, leading to significant improvements in computational efficiency. Additionally, based on this emulator, we introduce four active learning strategies that optimize the balance between accuracy and simulation costs to guide the selection of the fidelity level and input locations for the next simulation run. We demonstrate the effectiveness of the proposed approach in a suite of synthetic examples and a real-world problem. An R package `RNAmf` for the proposed methodology is provided on CRAN.

Keywords: Surrogate model; Sequential design; Uncertainty quantification; Gaussian process; Auto-regressive model.

*These authors gratefully acknowledge funding from NSF DMS 2113407 and 2338018.

1 Introduction

Computer simulations play a crucial role in engineering and scientific research, serving as valuable tools for predicting the performance of complex systems across diverse fields such as aerospace engineering (Mak et al., 2018), natural disaster prediction (Ma et al., 2022), and cell biology (Sung et al., 2020). However, conducting high-fidelity simulations for parameter space exploration can be demanding due to prohibitive costs. To address this challenge, *multi-fidelity emulation* has emerged as a promising alternative. It leverages computationally expensive yet accurate high-fidelity simulations alongside computationally inexpensive but potentially less accurate low-fidelity simulations to create an efficient predictive model, *emulating* the expensive computer code. By strategically integrating these simulations and designing multi-fidelity experiments, we can potentially improve accuracy without excessive computational resources.

The usefulness of the multi-fidelity emulation framework has driven extensive research in recent years. One popular approach is the Kennedy-O'Hagan (KO) model (Kennedy and O'Hagan, 2000), which models a sequence of computer simulations from lowest to highest fidelity using a sequence of Gaussian process (GP) models (Gramacy, 2020; Rasmussen and Williams, 2006), linked by a linear auto-regressive framework. This model has made significant contributions across various fields employing multi-fidelity computer experiments (see, e.g., Patra et al., 2020; Kuya et al., 2011; Demeyer et al., 2017), and several recent developments, including Qian et al. (2006), Le Gratiet (2013), Le Gratiet and Garnier (2014), Qian and Wu (2008), Perdikaris et al. (2017), and Ji et al. (2024) (among many others), have investigated modeling strategies for efficient posterior prediction and Bayesian uncertainty quantification.

Despite this body of work, most of these approaches rely on the assumption of linear correlation between low-fidelity and high-fidelity data, resulting in an *additive* GP structure. With the growing complexity of modern data, such models face challenges in capturing complex relationships between data with different fidelity levels. As shown in the left panel of Figure 1, where the relationship between high-fidelity data and low-fidelity data is nonlinear, the KO model falls short in providing accurate predictions due to its limited

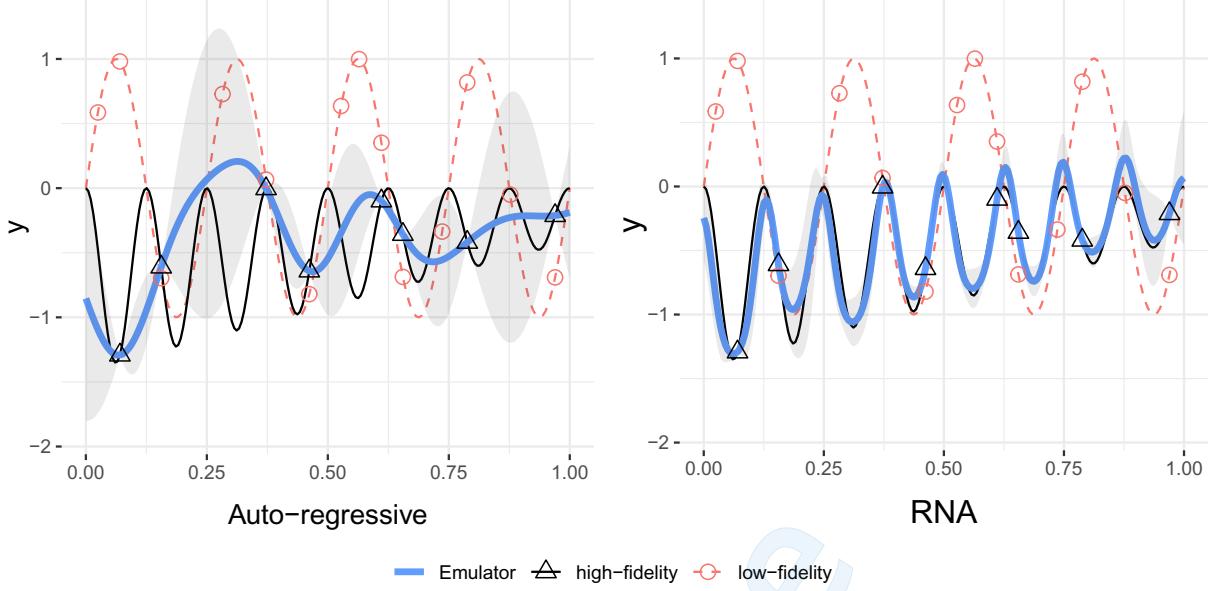


Figure 1: An example adapted from Perdikaris et al. (2017), where $n_1 = 13$ samples (red dots) are collected from the low-fidelity simulator $f_1(x) = \sin(8\pi x)$ (red dashed line), and $n_2 = 8$ samples (black triangles) are collected from the high-fidelity simulator $f_2(x) = (x - \sqrt{2})f_1(x)^2$ (black solid line). The KO emulator (left panel) and the RNA emulator (right panel) are shown as blue lines. Gray shaded regions represent the 95% pointwise confidence intervals.

flexibility.

In this paper, we propose a new and flexible model that captures the *nonlinear* relationships between multi-fidelity data in a *recursive* manner. This flexible nonlinear functional form can encompass many existing models, including the KO model, as a special case. Specifically, we compose GP priors to model multi-fidelity data non-additively. Hence, we refer to this proposed method as the *Recursive Non-Additive (RNA) emulator*. As shown in the right panel of Figure 1, the RNA emulator demonstrates superiority over the KO model by emulating the high-fidelity simulator with high accuracy and low uncertainty.

The RNA emulator belongs to the emerging field of linked/deep GP models (see, e.g., Kyzyurova et al., 2018; Ming and Guillas, 2021; Sauer et al., 2023; Ming et al., 2023), where different GPs are connected in a coupled manner. To the best of our knowledge, there has been limited research on extending such results for the analysis of multi-fidelity computer experiments, and we aim to address this gap in our work. Notably, recent work by Perdikaris et al. (2017) has made progress in this direction, but their approach assumes

an additive structure for the kernel function, and employs the Monte Carlo integration to handle intractable posterior distributions. Recent advancement by Ko and Kim (2022) extends deep GP models for multi-fidelity computer experiments, but still relies on the additive structure of the KO model. Similarly, Cutajar et al. (2019) employ an additive kernel akin to Perdikaris et al. (2017) and rely on the sparse variational approximation for inference. In a similar vein, Meng and Karniadakis (2020), Li et al. (2020), Meng et al. (2021), and Kerleguer et al. (2024) establish connections between different fidelities using (Bayesian) neural networks. In contrast, our proposed model not only provides great flexibility using GP priors with commonly used kernel structures to connect multi-fidelity data, but also provides analytical expressions for both the posterior mean and variance. This computational improvement allows for more efficient calculations, facilitating efficient uncertainty quantification.

Leveraging this newly developed RNA emulator, we introduce four *active learning* strategies to achieve enhanced accuracy while carefully managing the limited simulation resources, which is particularly crucial for computationally expensive simulations. Active learning, also known as *sequential design*, involves sequentially searching for and acquiring new data points at optimal locations based on a given sampling criterion, to construct an accurate surrogate model/emulator. While active learning has been well-established for single-fidelity GP emulators (Gramacy, 2020; Rasmussen and Williams, 2006; Santner et al., 2018), research in the context of multi-fidelity computer experiments is scarce and more challenging. This is because it requires simultaneous selection of optimal input locations and fidelity levels, accounting for their respective simulation costs. Although some recent works have considered cost in specific cases like single-objective unconstrained optimization (Huang et al., 2006; Swersky et al., 2013; He et al., 2017) and global approximation (Xiong et al., 2013; Le Gratiet and Cannamela, 2015; Stroh et al., 2022; Ehara and Guillas, 2023; Sung et al., 2024a), most of these works were developed based on the KO model. Active learning for non-additive GP models has not been fully explored in the literature. In addition, popular sampling criteria for global approximation, such as “Active Learning MacKay” (McKay et al., 2000, ALM) and “Active Learning Cohn” (Cohn, 1993, ALC),

remain largely unexplored in the context of multi-fidelity computer experiments. Recent successful applications of these sampling criteria to other learning problems can be found in Binois et al. (2019), Park et al. (2023), Koermer et al. (2024), and Sauer et al. (2023).

Our main contribution lies in advancing active learning with these popular sampling criteria, based on this newly developed RNA emulator. It is important to note that few existing works in the multi-fidelity deep GP literature (Perdikaris et al., 2017; Cutajar et al., 2019; Ko and Kim, 2022) delve into active learning, mainly due to computational complexities associated with computing acquisition functions. In contrast, our closed-form posterior mean and variance of the RNA emulator not only facilitate efficient computation of these sampling criteria, but also provide valuable mathematical insights into the active learning. To facilitate broader usage, we implement an R (R Core Team, 2018) package called **RNAmf** (Heo and Sung, 2024), which is available on CRAN.

The structure of this article is as follows. In Section 2, we provide a brief review of the KO model. Our proposed RNA emulator is introduced in Section 3. Section 4 outlines our active learning strategies based on the RNA emulator. Numerical and real data studies are presented in Sections 5 and 6, respectively. Lastly, we conclude the paper in Section 7.

2 Background

2.1 Problem Setup

Let $f_l(\mathbf{x})$ represent the scalar simulation output of the computer code with input parameters $\mathbf{x} \in \Omega \subseteq \mathbb{R}^d$ at fidelity level $l = 1, \dots, L$. We assume that L distinct fidelity levels of simulations are conducted for training an emulator, where a higher fidelity level corresponds to a simulator with more accurate outputs but also higher simulation costs per run.

Our primary objective is to construct an efficient emulator for the highest-fidelity simulation code, $f_L(\mathbf{x})$. For each fidelity level l , we perform simulations at n_l design points denoted by $\mathcal{X}_l = \{\mathbf{x}_i^{[l]}\}_{i=1}^{n_l}$. These simulations yield the corresponding simulation outputs $\mathbf{y}_l := (f_l(\mathbf{x}))_{\mathbf{x} \in \mathcal{X}_l}$, representing the vector of outputs for $f_l(\mathbf{x})$ at design points $\mathbf{x} \in \mathcal{X}_l$, and each element of \mathbf{y}_l is denoted by $y_i^{[l]} = f_l(\mathbf{x}_i^{[l]})$. We assume that the designs \mathcal{X}_l are

sequentially nested, i.e.,

$$\mathcal{X}_L \subseteq \mathcal{X}_{L-1} \subseteq \cdots \subseteq \mathcal{X}_1 \subseteq \Omega, \quad (1)$$

and $\mathbf{x}_i^{[l]} = \mathbf{x}_i^{[l-1]}$ for $i = 1, \dots, n_l$. In other words, design points run for a higher-fidelity simulator are contained within the design points run for a lower-fidelity simulator. This property has been shown to lead to more efficient inference in various multi-fidelity emulation approaches (Qian, 2009; Qian et al., 2009; Haaland and Qian, 2010).

Furthermore, we let C_l denote the simulation cost (e.g., in CPU hours) for a single run of the simulator at fidelity level l . Since higher-fidelity simulators are more computationally demanding, this implies that $0 < C_1 < C_2 < \dots < C_L$.

2.2 Auto-regressive model

One of the prominent approaches for modeling $f_L(\mathbf{x})$ is the KO model (also known as co-kriging model) proposed by Kennedy and O'Hagan (2000), which can be expressed in an auto-regressive form as follows:

$$\begin{cases} f_1(\mathbf{x}) = Z_1(\mathbf{x}), \\ f_l(\mathbf{x}) = \rho_{l-1} f_{l-1}(\mathbf{x}) + Z_l(\mathbf{x}), \quad \text{for } 2 \leq l \leq L, \end{cases} \quad (2)$$

where ρ_{l-1} is an unknown auto-regressive parameter, and $Z_l = (f_l - \rho_{l-1} f_{l-1})$ represents the *discrepancy* between the $(l-1)$ -th and l -th code. The KO model considers a probabilistic surrogate model by assuming that $\{Z_l\}_{l=1}^L$ follow independent zero-mean GP models:

$$Z_l(\mathbf{x}) \stackrel{\text{indep.}}{\sim} \mathcal{GP}\{\alpha_l(\mathbf{x}), \tau_l^2 K_l(\mathbf{x}, \mathbf{x}')\}, \quad l = 1, \dots, L, \quad (3)$$

where $\alpha_l(\mathbf{x})$ is a mean function, τ_l^2 is a variance parameter, and $K_l(\mathbf{x}, \mathbf{x}')$ is a positive-definite kernel function defined on $\Omega \times \Omega$. In the original KO paper, $\alpha_l(\mathbf{x})$ is assumed to be $h(\mathbf{x})\beta_l$, where $h(\mathbf{x})$ is a vector of d regression functions. Other common choices for $\alpha_l(\mathbf{x})$ include $\alpha_l(\mathbf{x}) \equiv 0$ or $\alpha_l(\mathbf{x}) \equiv \mu_l$. As for the kernel function, popular choices include the squared exponential kernel and Matérn kernel (Stein, 1999). Specifically, the *anisotropic* squared

exponential kernel takes the form:

$$K_l(\mathbf{x}, \mathbf{x}') = \prod_{j=1}^d \psi(x_j, x'_j; \theta_{lj}) = \prod_{j=1}^d \exp\left(-\frac{(x_j - x'_j)^2}{\theta_{lj}}\right), \quad (4)$$

where $(\theta_{l1}, \dots, \theta_{ld})$ is the *lengthscale* hyperparameter, indicating that the correlation decays exponentially fast in the squared distance between \mathbf{x} and \mathbf{x}' . The GP model (3), combined with the auto-regressive model (2), implies that, conditional on the parameters τ_l^2, θ_{lj} , and $\mu(\cdot)$, the joint distribution of $(\mathbf{y}_1, \dots, \mathbf{y}_L)$ follows a multivariate normal distribution, and these unknown parameters can be estimated via maximum likelihood estimation or Bayesian inference. Given the data $(\mathbf{y}_1, \dots, \mathbf{y}_L)$, it can be shown that the posterior distribution of $f_L(\mathbf{x})$ is also a GP. The posterior mean function can then be used to emulate the expensive simulator, while the posterior variance function can be employed to quantify the emulation uncertainty. Refer to Kennedy and O'Hagan (2000) for further details.

The auto-regressive framework of the KO model has led to the development of several variants. For instance, Le Gratiet and Garnier (2014) introduce a faster algorithm based on a *recursive* formulation for computing the posterior of $f_l(\mathbf{x})$ more efficiently. To enhance the model's flexibility, Qian et al. (2006), Le Gratiet and Garnier (2014), and Qian and Wu (2008) allow the auto-regressive parameter ρ_l to depend on the input \mathbf{x} , that is, $f_l(\mathbf{x}) = \rho_{l-1}(\mathbf{x})f_{l-1}(\mathbf{x}) + Z_l(\mathbf{x})$ for $2 \leq l \leq L$, where the first two assume $\rho_{l-1}(\mathbf{x})$ to be a linear function, while the last one assumes $\rho_{l-1}(\mathbf{x})$ to be a GP.

3 Recursive Non-Additive (RNA) emulator

Despite the advantages of the KO model, it results in an *additive* GP model based on (2) and (3), which may not adequately capture the *nonlinear* relationships between data at different fidelity levels. To overcome this limitation and achieve a more flexible representation, we

propose a novel Recursive Non-Additive (RNA) emulator:

$$\begin{cases} f_1(\mathbf{x}) = W_1(\mathbf{x}), \\ f_l(\mathbf{x}) = W_l(\mathbf{x}, f_{l-1}(\mathbf{x})), \quad \text{for } l = 2, \dots, L. \end{cases} \quad (5)$$

The model structure is illustrated in Figure 2. This RNA model offers greater flexibility and can encompass many existing models as special cases. For instance, the auto-regressive KO model can be represented in the form of (5) by setting $W_l(\mathbf{x}, f_{l-1}(\mathbf{x})) = \rho_{l-1} f_{l-1}(\mathbf{x}) + \delta(\mathbf{x})$. Similarly, the model in Qian et al. (2006), Le Gratiet and Garnier (2014), and Qian and Wu (2008) can be considered a special case by setting $W_l(\mathbf{x}, f_{l-1}(\mathbf{x})) = \rho_{l-1}(\mathbf{x}) f_{l-1}(\mathbf{x}) + \delta(\mathbf{x})$.

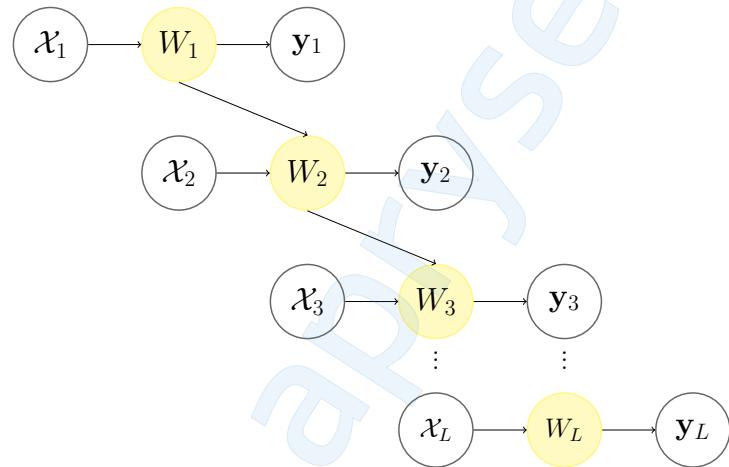


Figure 2: An illustration of the recursive structure of the RNA model.

We model the relationship W_l using a GP prior: $W_1(\mathbf{x}) \sim \mathcal{GP}\{\alpha_1(\mathbf{x}), \tau_1^2 K_1(\mathbf{x}, \mathbf{x}')\}$ and

$$W_l(\mathbf{z}) \sim \mathcal{GP}\{\alpha_l(\mathbf{z}), \tau_l^2 K_l(\mathbf{z}, \mathbf{z}')\}, \quad l = 2, \dots, L, \quad (6)$$

where $\mathbf{z} = (\mathbf{x}, y)$, forming a vector of size $(d + 1)$, and K_l is a positive-definite kernel. We consider a constant mean, i.e., $\alpha_1(\mathbf{x}) = \alpha_1$ and $\alpha_l(\mathbf{z}) = \alpha_l$ for $l \geq 2$. We adopt popular kernel choices for K_l , such as the squared exponential kernel and Matérn kernel. In particular, the

squared exponential kernel K_l for $l \geq 2$, following (4), can be expressed as:

$$\begin{aligned} K_l(\mathbf{z}, \mathbf{z}') &= \psi(y_{l-1}, y'_{l-1}; \theta_{ly}) \prod_{j=1}^d \psi(x_j, x'_j; \theta_{lj}) \\ &= \exp\left(-\frac{(y_{l-1} - y'_{l-1})^2}{\theta_{ly}}\right) \prod_{j=1}^d \exp\left(-\frac{(x_j - x'_j)^2}{\theta_{lj}}\right), \end{aligned} \quad (7)$$

where the lengthscale hyperparameter, $\boldsymbol{\theta}_l = (\theta_{l1}, \dots, \theta_{ld}, \theta_{ly})$, represents a vector of size $(d+1)$, and $K_1(\mathbf{x}, \mathbf{x}')$ takes the form of (4). The Matérn kernel can be similarly constructed, which is given in the Supplementary Materials S1.

Combining the GP model (6) with the recursive formulation (5) and assuming the nested design as in (1), the observed simulations \mathbf{y}_l follow a multivariate normal distribution:

$$\mathbf{y}_l \sim \mathcal{N}_{n_l}(\alpha_l \mathbf{1}_{n_l}, \tau_l^2 \mathbf{K}_l) \quad \text{for } l = 1, \dots, L,$$

where $\mathbf{1}_{n_l}$ is a unit vector of size n_l , \mathbf{K}_l is an $n_l \times n_l$ matrix with each element $(\mathbf{K}_l)_{ij} = K_1(\mathbf{x}_i^{[1]}, \mathbf{x}_j^{[1]})$, and $(\mathbf{K}_l)_{ij} = K_l(\mathbf{z}_i^{[l]}, \mathbf{z}_j^{[l]})$ for $l \geq 2$, where $\mathbf{z}_i^{[l]} = (\mathbf{x}_i^{[l]}, f_{l-1}(\mathbf{x}_i^{[l]}))$. Note that $f_{l-1}(\mathbf{x}_i^{[l]}) = y_i^{[l-1]}$ due to the nested design assumption, which is the i -th simulation output at level $l-1$. The parameters $\{\alpha_l, \tau_l^2, \boldsymbol{\theta}_l\}_{l=1}^L$ can be estimated by maximum likelihood estimation. Specifically, the log-likelihood (up to an additive constant) is

$$-\frac{1}{2} \left(n_l \log(\tau_l^2) + \log(\det(\mathbf{K}_l)) + \frac{1}{\tau_l^2} (\mathbf{y}_l - \alpha_l \mathbf{1}_{n_l})^T \mathbf{K}_l^{-1} (\mathbf{y}_l - \alpha_l \mathbf{1}_{n_l}) \right).$$

The parameters can then be efficiently estimated by maximizing the log-likelihood via an optimization algorithm, such as quasi-Newton optimization method of Byrd et al. (1995).

It is important to acknowledge that the idea of a recursive GP was previously proposed by Perdikaris et al. (2017), referred to as a *nonlinear auto-regressive model* therein. However, there are two key distinctions between their model and ours. The first distinction lies in the kernel assumption, where they assume an additive form of the kernel to better capture the auto-regressive nature. Specifically, they use $K_l(\mathbf{z}, \mathbf{z}') = K_{l1}(\mathbf{x}, \mathbf{x}')K_{l2}(f_{l-1}(\mathbf{x}), f_{l-1}(\mathbf{x}')) + K_{l3}(\mathbf{x}, \mathbf{x}')$ with valid kernel functions K_{l1} , K_{l2} , and K_{l3} . While our kernel function shares

some similarities, particularly the first component $K_{l1}(\mathbf{x}, \mathbf{x}')K_{l2}(f_{l-1}(\mathbf{x}), f_{l-1}(\mathbf{x}'))$, the role of the second component K_{l3} in predictions remains unclear. The inclusion of this component introduces d hyperparameters for an anisotropic kernel, making the estimation more challenging, especially for high-dimensional problems. In contrast, we adopt the natural form of popular kernel choices, such as the squared exponential kernel in (7) and the Matérn kernel, placing our model within the emerging field of linked/deep GP models, which has shown promising results in the computer experiment literature (Kyzyurova et al., 2018; Ming and Guillas, 2021; Sauer et al., 2023). The second distinction is in the computation for the posterior of $f_L(\mathbf{x})$. Specifically, their model relies on Monte Carlo (MC) integration for their computation, which can be quite computationally demanding, especially in this recursive formulation. In contrast, with these popular kernel choices, we can derive the posterior mean and variance of $f_L(\mathbf{x})$ in a closed form, which is presented in the following proposition, enabling more efficient predictions and uncertainty quantification.

The derivation of the posterior follows these steps. Starting with the GP assumption (6), and utilizing the properties of conditional multivariate normal distribution, the posterior distribution of f_l given \mathbf{y}_l and f_{l-1} at a new input location \mathbf{x} is normally distributed, namely: $f_1(\mathbf{x})|\mathbf{y}_1 \sim \mathcal{N}(\mu_1(\mathbf{x}), \sigma_1^2(\mathbf{x}))$ with

$$\mu_1(\mathbf{x}) = \alpha_1 \mathbf{1}_{n_1} + \mathbf{k}_1(\mathbf{x})^T \mathbf{K}_1^{-1} (\mathbf{y}_1 - \alpha_1 \mathbf{1}_{n_1}), \quad \text{and} \quad (8)$$

$$\sigma_1^2(\mathbf{x}) = \tau_1^2 (1 - \mathbf{k}_1(\mathbf{x})^T \mathbf{K}_1^{-1} \mathbf{k}_1(\mathbf{x})), \quad (9)$$

and $f_l(\mathbf{x})|\mathbf{y}_l, f_{l-1}(\mathbf{x}) \sim \mathcal{N}(\mu_l(\mathbf{x}, f_{l-1}(\mathbf{x})), \sigma_l^2(\mathbf{x}, f_{l-1}(\mathbf{x})))$ for $l = 2, \dots, L$ with

$$\mu_l(\mathbf{x}, f_{l-1}(\mathbf{x})) = \alpha_l \mathbf{1}_{n_l} + \mathbf{k}_l(\mathbf{x}, f_{l-1}(\mathbf{x}))^T \mathbf{K}_l^{-1} (\mathbf{y}_l - \alpha_l \mathbf{1}_{n_l}), \quad \text{and} \quad (10)$$

$$\sigma_l^2(\mathbf{x}, f_{l-1}(\mathbf{x})) = \tau_l^2 (1 - \mathbf{k}_l(\mathbf{x}, f_{l-1}(\mathbf{x}))^T \mathbf{K}_l^{-1} \mathbf{k}_l(\mathbf{x}, f_{l-1}(\mathbf{x}))), \quad (11)$$

where $\mathbf{k}_1(\mathbf{x})$ and $\mathbf{k}_l(\mathbf{x}, f_{l-1}(\mathbf{x}))$ are an $n_l \times 1$ matrix with each element $(\mathbf{k}_1(\mathbf{x}))_{i,1} = K_1(\mathbf{x}, \mathbf{x}_i^{[1]})$ and $(\mathbf{k}_l(\mathbf{x}, f_{l-1}(\mathbf{x})))_{i,1} = K_l((\mathbf{x}, f_{l-1}(\mathbf{x})), (\mathbf{x}_i^{[l]}, y_i^{[l-1]}))$ for $l \geq 2$, respectively. The posterior

distribution of f_l can then be obtained by

$$\begin{aligned} & p(f_l(\mathbf{x})|\mathbf{y}_1, \dots, \mathbf{y}_l) \\ &= \int \cdots \int p(f_l(\mathbf{x})|\mathbf{y}_l, f_{l-1}(\mathbf{x})) p(f_{l-1}(\mathbf{x})|\mathbf{y}_{l-1}, f_{l-2}(\mathbf{x})) \cdots p(f_1(\mathbf{x})|\mathbf{y}_1) d(f_{l-1}(\mathbf{x})) \dots d(f_1(\mathbf{x})). \end{aligned}$$

This posterior is analytically intractable but can be numerically approximated using MC integration, as done in Perdikaris et al. (2017), which involves sequential sampling from the normal distribution $p(f_l(\mathbf{x})|\mathbf{y}_l, f_{l-1}(\mathbf{x}))$ from $l = 1$ to $l = L$. However, this method can be computationally demanding, especially when the dimension of \mathbf{x} and the number of fidelity levels increase. To address this, we derive recursive closed-form expressions for the posterior mean and variance under popular kernel choices as follows.

Proposition 3.1. *Under the squared exponential kernel function (7), the posterior mean and variance of $f_l(\mathbf{x})$ given the data $\{\mathbf{y}_l\}_{l=1}^L$ for $l \geq 2$ can be expressed in a recursive fashion:*

$$\begin{aligned} \mu_l^*(\mathbf{x}) &:= \mathbb{E}[f_l(\mathbf{x})|\mathbf{y}_1, \dots, \mathbf{y}_l] \\ &= \alpha_l + \sum_{i=1}^{n_l} r_i \prod_{j=1}^d \exp\left(-\frac{(x_j - x_{ij}^{[l]})^2}{\theta_{lj}}\right) \frac{1}{\sqrt{1 + 2\frac{\sigma_{l-1}^{*2}(\mathbf{x})}{\theta_{ly}}}} \exp\left(-\frac{(y_i^{[l-1]} - \mu_{l-1}^*(\mathbf{x}))^2}{\theta_{ly} + 2\sigma_{l-1}^{*2}(\mathbf{x})}\right), \end{aligned}$$

and

$$\begin{aligned} \sigma_l^{*2}(\mathbf{x}) &:= \mathbb{V}[f_l(\mathbf{x})|\mathbf{y}_1, \dots, \mathbf{y}_l] = \tau_l^2 - (\mu_l^*(\mathbf{x}) - \alpha_l)^2 + \\ & \left(\sum_{i,k=1}^{n_l} \zeta_{ik} (r_i r_k - \tau_l^2 (\mathbf{K}_l^{-1})_{ik}) \prod_{j=1}^d \exp\left(-\frac{(x_j - x_{ij}^{[l]})^2 + (x_j - x_{kj}^{[l]})^2}{\theta_{lj}}\right) \right), \quad (12) \end{aligned}$$

where $r_i = (\mathbf{K}_l^{-1}(\mathbf{y}_l - \alpha_l \mathbf{1}_{n_l}))_i$, and

$$\zeta_{ik} = \frac{1}{\sqrt{1 + 4\frac{\sigma_{l-1}^{*2}(\mathbf{x})}{\theta_{ly}}}} \exp\left(-\frac{\frac{(y_i^{[l-1]} + y_k^{[l-1]})^2}{2} - \mu_{l-1}^*(\mathbf{x})^2}{\frac{\theta_{ly}}{2} + 2\sigma_{l-1}^{*2}(\mathbf{x})} - \frac{(y_i^{[l-1]} - y_k^{[l-1]})^2}{2\theta_{ly}}\right). \quad (13)$$

For $l = 1$, it follows that $\mu_1^*(\mathbf{x}) = \mu_1(\mathbf{x})$ and $\sigma_1^{*2}(\mathbf{x}) = \sigma_1^2(\mathbf{x})$ as in (8) and (9), respectively.

The posterior mean and variance under a Matérn kernel with the smoothness parameter $\nu = 1.5$ and $\nu = 2.5$ are provided in the Supplementary Materials S3, and the detailed derivations for Proposition 3.1 are provided in Supplementary Materials S2, which follow the proof of Kyzyurova et al. (2018) and Ming and Guillas (2021). With this proposition, the posterior mean and variance can be efficiently computed in a recursive fashion. Similar to Kyzyurova et al. (2018) and Ming and Guillas (2021), we adopt the *moment matching* method, using a Gaussian distribution to approximate the posterior distribution with the mean and variance presented in the proposition. The parameters in the posterior distribution, including $\{\alpha_l, \tau_l^2, \theta_l\}_{l=1}^L$, can be plugged in by their estimates.

Notably, Proposition 3.1 can be viewed as a simplified representation of Theorem 3.3 from Ming and Guillas (2021) for constructing a linked GP surrogate. However, it is important to highlight the distinctions and contributions of our work, particularly in the context of multi-fidelity computer experiments. Firstly, there are currently no existing closed-form expressions for the posterior mean and variance in the multi-fidelity deep GP literature. By providing such expressions, our work fills this gap, offering valuable mathematical insights and enhancing computational efficiency for active learning strategies, which will be discussed in Section 3.1 and Section 4. Additionally, while the linked GP model provides a general framework, much of the discussion in their work focuses on *sequential* GPs, where the output of the high-layer emulator depends solely on the output of the low-layer emulator, i.e., $W_2(W_1(\mathbf{x}))$. Our setup differs slightly, as the high-fidelity emulator in our RNA framework depends not only on the output of the low-fidelity emulator but also on the input variables directly, i.e., $W_2(\mathbf{x}, W_1(\mathbf{x}))$. This difference in formulation is important and impacts the design of active learning strategies in our framework.

Similar to conventional GP emulators for single-fidelity deterministic computer models, the proposed RNA emulator also exhibits the interpolation property, which is described in the following proposition. The proof is provided in Supplementary Materials S4.

Proposition 3.2. *The RNA emulator satisfies interpolation property, that is, $\mu_l^*(\mathbf{x}_i^{[l]}) = y_i^{[l]}$, and $\sigma_l^{*2}(\mathbf{x}_i^{[l]}) = 0$, where $\{(\mathbf{x}_i^{[l]}, y_i^{[l]})\}_{i=1, \dots, n_l}$ are the training samples.*

An example of this posterior distribution is presented in the right panel of Figure 1,

illustrating that the posterior mean closely aligns with the true function, and the confidence intervals constructed by the posterior variance cover the true function. For further insights into how this nonlinear relationship modeling can effectively reconstruct the high-fidelity function $f_2(x)$ for this example, we refer to Perdikaris et al. (2017).

Our R package, **RNAmf**, implements the parameter estimation and computations for the closed-form posterior mean and variance using a squared exponential kernel and a Matérn kernel with smoothness parameters of 1.5 and 2.5.

3.1 Insights into the RNA emulator

We delve into the RNA emulator, exploring its mathematical insights and investigating scenarios where this method may succeed or encounter challenges.

For the sake of simplicity in explanation, we consider two fidelity levels ($L = 2$) and assume the input x is one-dimensional. According to Proposition 3.1, under a squared exponential kernel function, the RNA emulator yields the following posterior mean:

$$\mu_2^*(x) = \alpha_2 + \sqrt{\frac{\theta_{2y}}{\theta_{2y} + 2\sigma_1^{*2}(x)}} \sum_{i=1}^{n_2} r_i \exp\left(-\frac{(x - x_i^{[2]})^2}{\theta_2} - \frac{(y_i^{[1]} - \mu_1^*(x))^2}{\theta_{2y} + 2\sigma_1^{*2}(x)}\right),$$

where $r_i = (\mathbf{K}_2^{-1}(\mathbf{y}_2 - \alpha_2 \mathbf{1}_{n_2}))_i$.

The mathematical expression reveals several insights into the behavior of the RNA emulator. Firstly, it reveals the impact of the uncertainty in the low-fidelity model, $\sigma_1^{*2}(x)$, on the posterior mean $\mu_2^*(x)$. In scenarios where $\sigma_1^{*2}(x) = 0$ for all $x \in \Omega$, $\mu_2^*(x)$ mirrors the posterior mean when $\mu_1^*(x)$ is replaced with the true low-fidelity function $f_1(x)$. Consequently, the term $\sqrt{\frac{\theta_{2y}}{\theta_{2y} + 2\sigma_1^{*2}(x)}}$ acts as a scaling factor for the posterior mean, adjusting the influence of the uncertainty $\sigma_1^{*2}(x)$ on the overall prediction to account for the approximation error between $\mu_1^*(x)$ and $f_1(x)$. Additionally, the inflated denominator of $\frac{(y_i^{[1]} - \mu_1^*(x))^2}{\theta_{2y} + 2\sigma_1^{*2}(x)}$ by the low-fidelity model uncertainty also aids in mitigating the approximation error, indicating a slower decay in correlation with the squared distance between the low-fidelity observations $y_i^{[1]}$ and $\mu_1^*(x)$. Both aspects ensure a balanced integration of high and low-fidelity information, which is particularly crucial when dealing with limited samples from low-fidelity data.

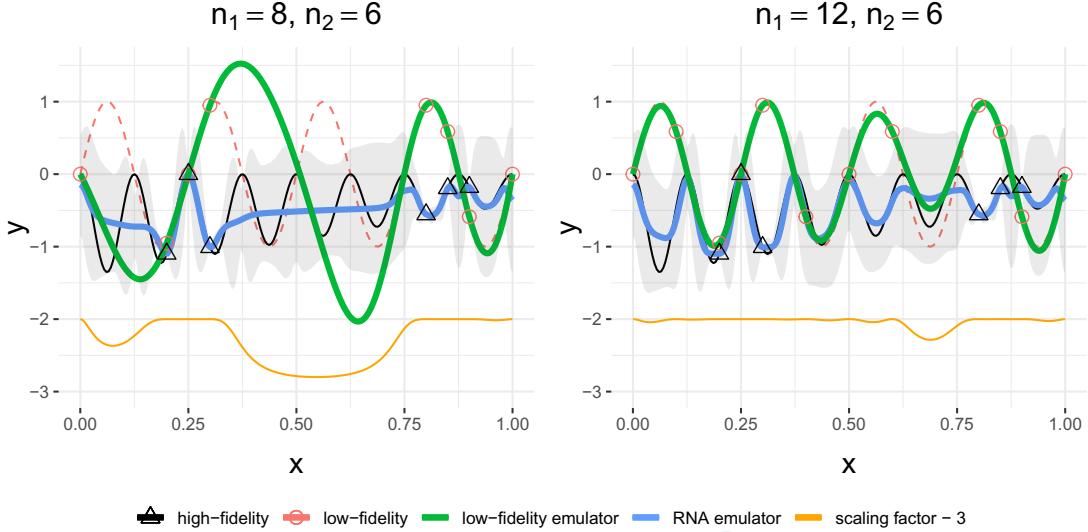


Figure 3: Illustration of RNA emulator insights using the Perdikaris example. The left panel and right panel depict results obtained with different sample sizes of low-fidelity data (red dots), $n_1 = 8$ (left) and $n_1 = 12$ (right), alongside the same high-fidelity data (black triangles) of size $n_2 = 6$. The scaling factor is the orange solid line, with values shifted by subtracting 3.

Figure 3 demonstrates an example of how the low-fidelity emulator impacts RNA emulation performance. The left panel illustrates that with limited low-fidelity data ($n_1 = 8$), especially in the absence of data at $x \in (0.3, 0.8)$, the posterior mean of the low-fidelity emulator, $\mu_1^*(x)$ (represented by the green line), inaccurately predicts the true low-fidelity simulator $f_1(x)$ (red dashed line). In this scenario, the scaling factor (orange line), $\sqrt{\frac{\theta_{2y}}{\theta_{2y} + 2\sigma_1^2(x)}}$, is very small for those poor predictions of $\mu_1^*(x)$, particularly for $x \in (0.3, 0.8)$. This results in $\mu_1^*(x)$ being close to the mean estimate $\hat{\alpha}_2$. This is not surprising because there is no data available from both low-fidelity and high-fidelity simulators in this region, leading to the posterior mean reverting back to the mean estimate. With an increase in low-fidelity data ($n_1 = 12$), which makes $\mu_1^*(x)$ much closer to the true $f_1(x)$, the scaling factor is close to one everywhere, significantly enhancing the accuracy of the RNA emulator.

The posterior variance can be written as (see Supplementary Materials S2)

$$\sigma_2^{*2}(x) = \mathbb{V} [\mathbb{E}[f_2(x)|f_1(x), \mathbf{y}_1, \mathbf{y}_2]] + \mathbb{E} [\mathbb{V}[f_2(x)|f_1(x), \mathbf{y}_1, \mathbf{y}_2]], \quad (14)$$

where both terms can be expressed in a closed form as in (S5.1) and (S5.2), respectively.

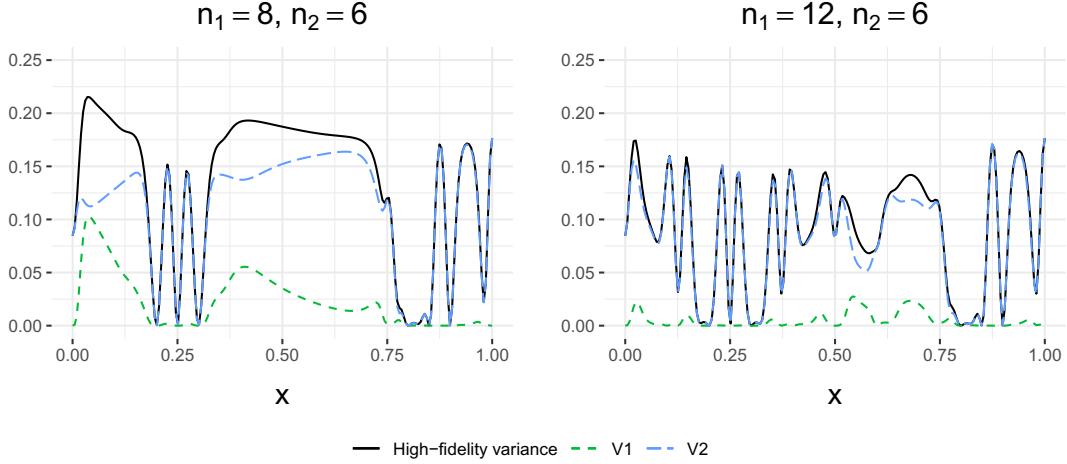


Figure 4: Illustration of decomposition of $\sigma_2^{*2}(x)$ (black solid line) for the examples of Figure 3, where V_1 is the blue dashed line and $V_2(x)$ is the green dashed line.

Define $V_1(x) = \mathbb{V}[\mathbb{E}[f_2(x)|f_1(x), \mathbf{y}_1, \mathbf{y}_2]]$ and $V_2(x) = \mathbb{E}[\mathbb{V}[f_2(x)|f_1(x), \mathbf{y}_1, \mathbf{y}_2]]$, then $V_1(x)$ represents the overall contribution of the GP emulator W_1 to $\sigma_2^{*2}(x)$ and $V_2(x)$ represents the contribution of the GP emulator W_2 to $\sigma_2^{*2}(x)$. This decomposition mirrors that of Ming and Guillas (2021) within the context of linked GPs. Figure 4 illustrates this decomposition for the examples in Figure 3. It can be seen that for both scenarios, V_2 appears to dominate V_1 , indicating that W_2 contributes more uncertainty than W_1 . However, when we have limited low-fidelity data (left panel), V_1 exhibits a very high peak at $x \approx 0.04$ with a value close to 0.10, even very close to the maximum value of V_2 . From an active learning perspective, if the cost of evaluating $f_1(x)$ is cheaper than $f_2(x)$, then it's sensible to select the next sample from the cheaper $f_1(x)$ to reduce $\sigma_2^{*2}(x)$. On the other hand, when we have more low-fidelity data (right panel), V_1 remains very small everywhere compared to V_2 , indicating that selecting the next sample from $f_2(x)$ would be more effective in reducing the predictive uncertainty. More details of active learning strategies will be introduced in the next section.

4 Active learning for RNA emulator

We present four active learning (AL) strategies aimed at enhancing the predictive capabilities of the proposed model through the careful design of computer experiments. These strategies

encompass the dual task of not only identifying the optimal input locations but also determining the most appropriate fidelity level.

We suppose that an initial experiment of sample size n_l for each fidelity level l , following a nested design $\mathcal{X}_L \subseteq \dots \subseteq \mathcal{X}_1$, is conducted, for which a space-filling design is often considered, such as the nested Latin hypercube design (Qian, 2009). AL seeks to optimize a selection criterion for choosing the next point $\mathbf{x}_{n_l+1}^{[l]}$ at fidelity level l , carrying out its corresponding simulation $y_{n_l+1}^{[l]} = f_l(\mathbf{x}_{n_l+1}^{[l]})$, and thus augmenting the dataset.

4.1 Active Learning Decomposition (ALD)

We first introduce an active learning criterion inspired by Section 3.1 and the variance-based adaptive design for linked GPs outlined in Ming and Guillas (2021). Specifically, we extend the decomposition of (14) to encompass L fidelity levels:

$$\sigma_L^{*2}(\mathbf{x}) = \sum_{l=1}^L V_l(\mathbf{x}), \quad (15)$$

where $V_l(\mathbf{x})$ represents the contribution of each GP emulator W_l at fidelity level l to $\sigma_L^{*2}(\mathbf{x})$:

$$V_l(\mathbf{x}) = \mathbb{E} \cdots \mathbb{E} \mathbb{V} \mathbb{E} \cdots \mathbb{E} [f_l(\mathbf{x}) | f_{L-1}(\mathbf{x}), \dots, f_1(\mathbf{x}), \mathbf{y}_L, \dots, \mathbf{y}_1],$$

with \mathbb{V} being at the l -th term. The expectation or variance in the l -th term is taken with respect to the variable $f_l(\mathbf{x})$. When $L = 2$, the closed-form expression for $V_l(\mathbf{x})$ is available, as shown in (14). For $L = 3$, each $V_l(\mathbf{x})$ can be easily approximated using MC methods. We detail the calculation of $V_l(\mathbf{x})$ for the settings of $L = 2$ and $L = 3$ in Supplementary Materials S5. However, the calculation becomes more cumbersome for $L \geq 4$, which we leave as a topic for future development.

Considering the simulation cost C_l , our approach guides the selection of the next point $\mathbf{x}_{n_l+1}^{[l]}$ at fidelity level l by maximizing the criterion, which we refer to as Active Learning Decomposition (ALD):

$$(l^*, \mathbf{x}_{n_{l^*}+1}^{[l^*]}) = \underset{l \in \{1, \dots, L\}; \mathbf{x} \in \Omega}{\operatorname{argmax}} \frac{V_l(\mathbf{x})}{\sum_{j=1}^l C_j},$$

which aims to maximize the ratio between each contribution $V_l(\mathbf{x})$ to $\sigma_L^{*2}(\mathbf{x})$ and the simulation cost $\sum_{j=1}^l C_j$ at each fidelity level l .

Simulation costs are incorporated to account for the nested structure. That is, to run the simulation $f_{l^*}(\mathbf{x}_{n_{l^*}+1}^{[l^*]})$, we also need to run $f_l(\mathbf{x}_{n_l+1}^{[l]})$ with $\mathbf{x}_{n_l+1}^{[l]} = \mathbf{x}_{n_{l^*}+1}^{[l^*]}$ for all $1 \leq l < l^*$. It is also worth mentioning that the cost can be tailored to depend on the input \mathbf{x} , as done in He et al. (2017) and Stroh et al. (2022).

4.2 Active Learning MacKay (ALM)

A straightforward but commonly used sampling criterion in AL is to select the next point that maximizes the posterior predictive variance (MacKay, 1992). Extending this concept to our scenario, we choose the next point by maximizing the ALM criterion:

$$(l^*, \mathbf{x}_{n_{l^*}+1}^{[l^*]}) = \underset{l \in \{1, \dots, L\}; \mathbf{x} \in \Omega}{\operatorname{argmax}} \frac{\sigma_l^{*2}(\mathbf{x})}{\sum_{j=1}^l C_j}. \quad (16)$$

Note that after running the simulation at the optimal input location $\mathbf{x}_{n_{l^*}+1}^{[l^*]}$ at level l^* , the posterior predictive variance $\sigma_{l^*}^{*2}(\mathbf{x}_{n_{l^*}+1}^{[l^*]})$ becomes zero (see Proposition 3.2). In other words, our selection of the optimal level hinges on achieving the highest ratio of *uncertainty reduction at $\mathbf{x}_{n_{l^*}+1}^{[l^*]}$* to the simulation cost.

The computation of ALM criterion is facilitated by the availability of the closed-form expression of the posterior predictive variance as in (12), which in turn simplifies the optimization process of (16). In particular, the optimal input location $\mathbf{x}_{n_l+1}^{[l]}$ for each l can be efficiently obtained through the `optim` library in R, using the `method=L-BFGS-B` option, which performs a quasi-Newton optimization approach of Byrd et al. (1995).

4.3 Active Learning Cohn (ALC)

Another widely employed, more aggregate criterion is Active Learning Cohn (ALC) (Cohn, 1993; Seo et al., 2000). In contrast to ALM, ALC selects an input location that maximizes the reduction in posterior variances *across the entire input space* after running this selected simulation. Extending the concept to our scenario, we choose the next point by maximizing

the ALC criterion:

$$(l^*, \mathbf{x}_{n_{l^*}+1}^{[l^*]}) = \operatorname{argmax}_{l \in \{1, \dots, L\}; \mathbf{x} \in \Omega} \frac{\Delta\sigma_L^2(l, \mathbf{x})}{\sum_{j=1}^l C_j}, \quad (17)$$

where $\Delta\sigma_L^2(l, \mathbf{x})$ is the *average reduction in variance* (of the highest-fidelity emulator) from the current design measured through a choice of the fidelity level l and the input location \mathbf{x} , augmenting the design. That is,

$$\Delta\sigma_L^2(l, \mathbf{x}) = \int_{\Omega} \sigma_L^{*2}(\boldsymbol{\xi}) - \tilde{\sigma}_L^{*2}(\boldsymbol{\xi}; l, \mathbf{x}) d\boldsymbol{\xi}, \quad (18)$$

where $\sigma_L^{*2}(\boldsymbol{\xi})$ is the posterior variance of $f_L(\boldsymbol{\xi})$ based on the current design $\{\mathcal{X}_l\}_{l=1}^L$, and $\tilde{\sigma}_L^{*2}(\boldsymbol{\xi}; l, \mathbf{x})$ is the posterior variance based on the *augmented design* combining the current design and a new input location \mathbf{x} at each fidelity level lower than or equal to l , i.e., $\{(\mathcal{X}_1 \cup \mathbf{x}_{n_1+1}^{[1]}), \dots, (\mathcal{X}_l \cup \mathbf{x}_{n_l+1}^{[l]}), \mathcal{X}_{l+1}, \dots, \mathcal{X}_L\}$ with $\mathbf{x}_{n_1+1}^{[1]} = \dots = \mathbf{x}_{n_l+1}^{[l]} = \mathbf{x}$. Once again, the incorporation of the new input location \mathbf{x} at each fidelity level lower than l is due to the nested structure assumption. In other words, our selection of the optimal level involves maximizing the ratio of *average reduction in the variance of the highest-fidelity emulator* to the associated simulation cost. In practice, the integration in (18) can be approximated by numerical methods, such as MC integration.

Unlike ALM where the influence of design augmentation on the variance of the highest-fidelity emulator is unclear, ALC is specifically designed to *maximize the reduction in variance of the highest-fidelity emulator*. However, the ALC strategy involves requiring knowledge of future outputs $y_{n_s+1}^{[s]} = f_s(\mathbf{x}_{n_s+1}^{[s]})$ for all $1 \leq s \leq l$, as they are involved in $\tilde{\sigma}_L^{*2}(\boldsymbol{\xi}; l, \mathbf{x})$ (as seen in (13)), but these outputs are not available prior to conducting the simulations. A possible approach to address this issue is through MC approximation to *impute* the outputs. Specifically, we can impute $y_{n_s+1}^{[s]}$ for each $1 \leq s \leq l$ by drawing samples from the posterior distribution of $f_s(\mathbf{x}_{n_s+1}^{[s]})$ based on the *current* design, which is a normal distribution with the posterior mean and variance presented in Proposition 3.1. This allows us to repeatedly compute $\tilde{\sigma}_L^{*2}(\boldsymbol{\xi}; l, \mathbf{x})$ using the imputations and average the results to approximate the variance. Notably, with the imputed output $y_{n_s+1}^{[s]}$, the variance

$\tilde{\sigma}_L^{*2}(\boldsymbol{\xi}; l, \mathbf{x})$ can be efficiently computed using the Sherman–Morrison formula (Harville, 1998) for updating the covariance matrix’s inverse, \mathbf{K}_l^{-1} , from $\sigma_L^{*2}(\boldsymbol{\xi})$ (Gramacy, 2020).

In contrast to ALM, maximizing the ALC criterion (17) can be quite computationally expensive due to the costly MC approximation to compute (18). To this end, an alternative strategy is proposed to strike a compromise by combining the two criteria.

4.4 Two-step approach: ALMC

Given the distinct advantages and limitations of both ALM and ALC criteria (details of which are referred to Chapter 6 of Gramacy, 2020), for a comprehensive exploration, we can contemplate their combination. Inspired by Le Gratiet and Cannamela (2015), we introduce a hybrid approach, which we refer to as *ALMC*. First, the optimal input location is selected by maximizing the posterior predictive variance of the *highest* fidelity emulator:

$$\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x} \in \Omega} \sigma_L^{*2}(\mathbf{x}).$$

Then, the ALC criterion determines the fidelity level with the identified input location:

$$l^* = \operatorname{argmax}_{l \in \{1, \dots, L\}} \frac{\Delta \sigma_L^2(l, \mathbf{x}^*)}{\sum_{j=1}^l C_j}.$$

Unlike ALM, this hybrid approach focuses on the *direct impact* on the highest-fidelity emulator. It first identifies the sampling location that maximizes $\sigma_L^{*2}(\mathbf{x})$, and then determines which level selection will effectively reduce the overall variance of the highest-fidelity emulator across the input space *after running this location*. This synergistic approach is not only expected to capture the advantages of both ALM and ALC, but also offers computational efficiency advantages compared to the ALC method in the previous subsection. This is due to the fact that the optimization for \mathbf{x}^* by maximizing the closed-form posterior variance is computationally much cheaper, as discussed in Section 4.2.

Figure 5 demonstrates the effectiveness of these four strategies for the example in the right panel of Figure 1. Consider the simulation costs: $C_1 = 1$ and $C_2 = 3$ for the two simulators. It shows that, for all four criteria, the choice is consistently in favor of selecting

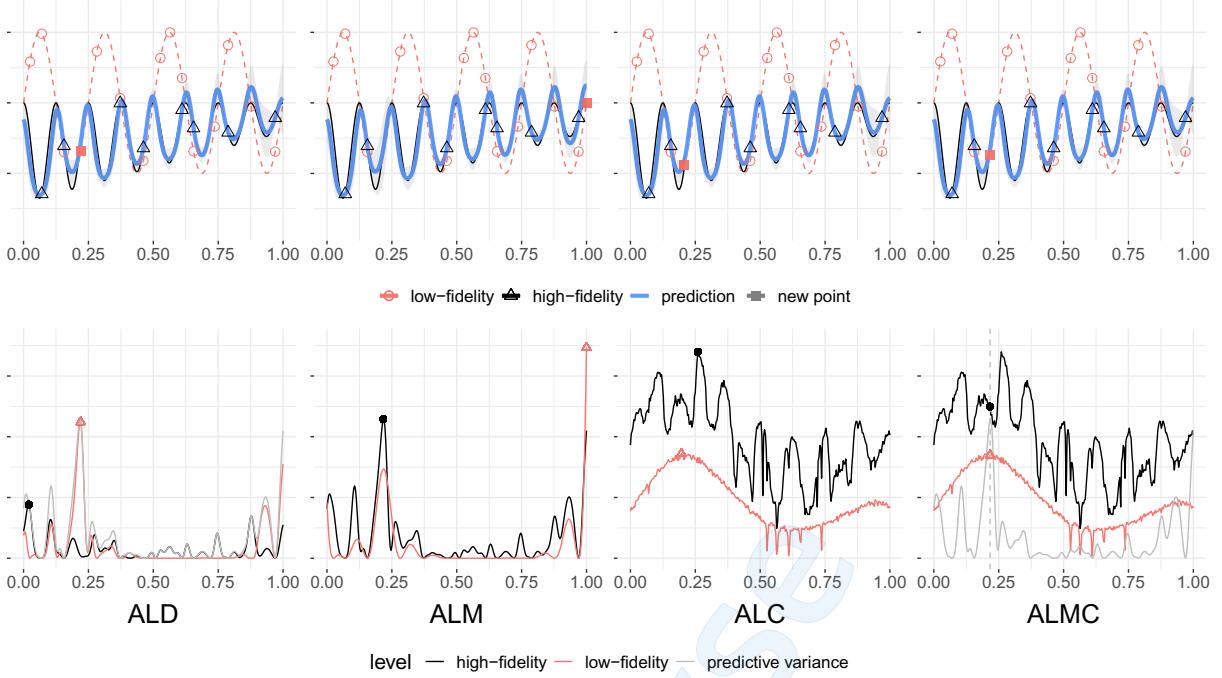


Figure 5: Demonstration of the four active learning strategies using the example in the right panel of Figure 1. The criteria of the four strategies are presented in the bottom panel, where the dots represent the optimal input locations for each of the simulators. Notably, ALD utilizes the gray line to illustrate $\sigma_2^{*2}(x)$, which is decomposed into $V_1(x)$ (depicted in red) and $V_2(x)$ (depicted in black). ALMC, on the other hand, employs the gray line to determine the optimal input location and then utilizes the red and black lines (which are identical to ALC) to decide the fidelity level. The upper panels show the corresponding fits after adding the selected points to the training dataset, where the solid dots represent the chosen samples, all of which select the low-fidelity simulator.

the low-fidelity simulator to augment the dataset. While the selected locations differ, ALD, ALC, and ALMC all fall within the range of $[0.18, 0.25]$, which, as per the current design (prior to running this simulation), holds large uncertainty, as seen in the right panel of Figure 1. ALM selects the sample at the boundary of the input space. All these selection outcomes contribute to an overall improvement in emulation accuracy, while simultaneously reducing global uncertainty, even when opting for low-fidelity data alone.

4.5 Remark on the AL strategies

In this section, we delve deeper into the merits of the AL strategies, with a focus on the conditions favoring each method. To gain deeper insights, we consider a synthetic example

generated from a 2-level Currin function (Xiong et al., 2013; Kerleguer et al., 2024), with the explicit form provided in Supplementary Materials S6. Assuming simulation costs $C_1 = 1$ and $C_2 = 3$, we employ the four AL strategies until reaching a total budget of 15.

Figure 6 showcases the selected sites within the input space $[0, 1]^2$. Similar to discussions on AL for single-fidelity GPs (Seo et al., 2000; Gramacy and Lee, 2009; Bilionis and Zabaras, 2012; Beck and Guillas, 2016), ALM tends to push selected data points towards the boundaries of the input space, whereas ALC avoids boundary locations. ALD and ALMC, inheriting attributes of ALM, exhibit similar behavior to ALM. The choice between them depends on the underlying true function: if the function in the boundary region is flat and exhibits more variability in the interior, then ALC may be preferable. Regarding computational efficiency, ALD, ALM, and ALMC benefit from closed-form expressions of the posterior variance, requiring only a few seconds per acquisition. In contrast, ALC is more computationally demanding due to extensive MC sampling efforts, taking several minutes per acquisition.

It is worth noting that if the scale of low-fidelity outputs significantly exceeds that of high-fidelity outputs, ALM may consistently favor low-fidelity levels in the initial acquisitions, as the maximum of the low-fidelity posterior variance tends to be large. However, it's unclear whether this selection is effective, as maximizing the posterior variance of the low-fidelity emulator doesn't necessarily translate to a reduction in the uncertainty of the high-fidelity emulator. In contrast, the other three methods focus on directly impacting the high-fidelity emulator by selecting points, making them independent of the scale. In summary, considering the discussions above and the findings from our empirical studies in Sections 5 and 6, ALD (for $L \leq 3$) and ALMC generally emerge as favorable choices, offering accurate RNA emulators along with computational efficiency.

5 Numerical Studies

In this section, we conduct a suite of numerical experiments to examine the performance of the proposed approach. The experiments encompass two main aspects. In Section 5.1, we assess the predictive capabilities of the proposed RNA emulator, while Section 5.2 delves

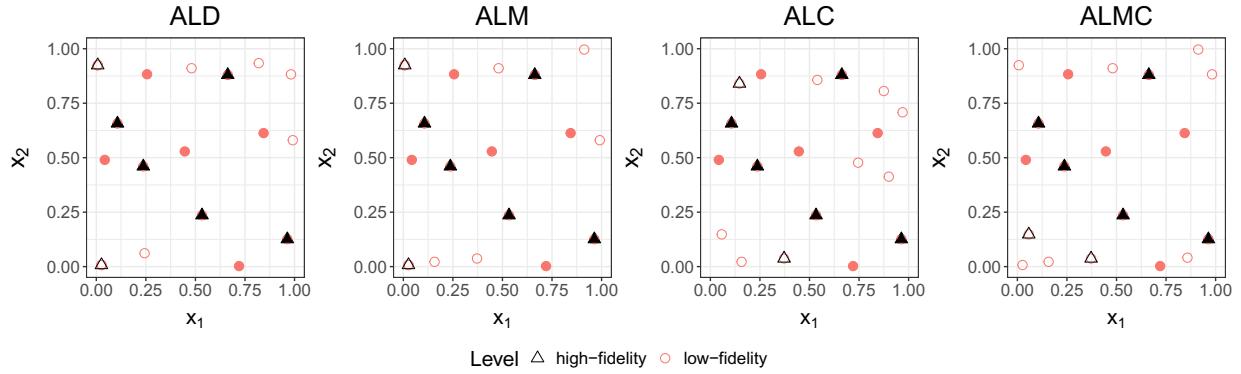


Figure 6: Selected input locations by four proposed strategies with a total budget of 15, where the simulation costs are $C_1 = 1$ and $C_2 = 3$. The initial design points are represented as filled shapes.

into the evaluation of the performance of the proposed AL strategies.

We consider the anisotropic squared exponential kernel as in (7) for the proposed model, a choice that is also shared by our competing methods. All experiments are performed on a MacBook Pro laptop with 2.9 GHz 6-Core Intel Core i9 and 16Gb of RAM.

5.1 Emulation performance

We begin by comparing the predictive performance of the proposed RNA emulator (labeled **RNAmf**) with two other methods in the numerical experiments: the co-kriging model (labeled **CoKriging**) by Le Gratiet and Garnier (2014), and the nonlinear auto-regressive multi-fidelity GP (labeled **NARGP**) by Perdikaris et al. (2017). The two methods are readily available through open repositories, specifically the **R** package **MuFiCokriging** (Le Gratiet, 2012) and the **Python** package on the GitHub repository (Perdikaris, 2016), respectively. Note that the multi-fidelity deep GP of Cutajar et al. (2019), which can be implemented using the **Python** package **emukit** (Paley et al., 2019, 2023), is not included in our comparison due to software limitations. We encountered challenges during implementation as the package relies on an outdated package, rendering it incompatible with our current environment.

Five synthetic examples commonly used in the literature to evaluate emulation performance in multi-fidelity simulations are considered, including the two-level Perdikaris function (Perdikaris et al., 2017; Kerleguer et al., 2024), the two-level Park function (Park,

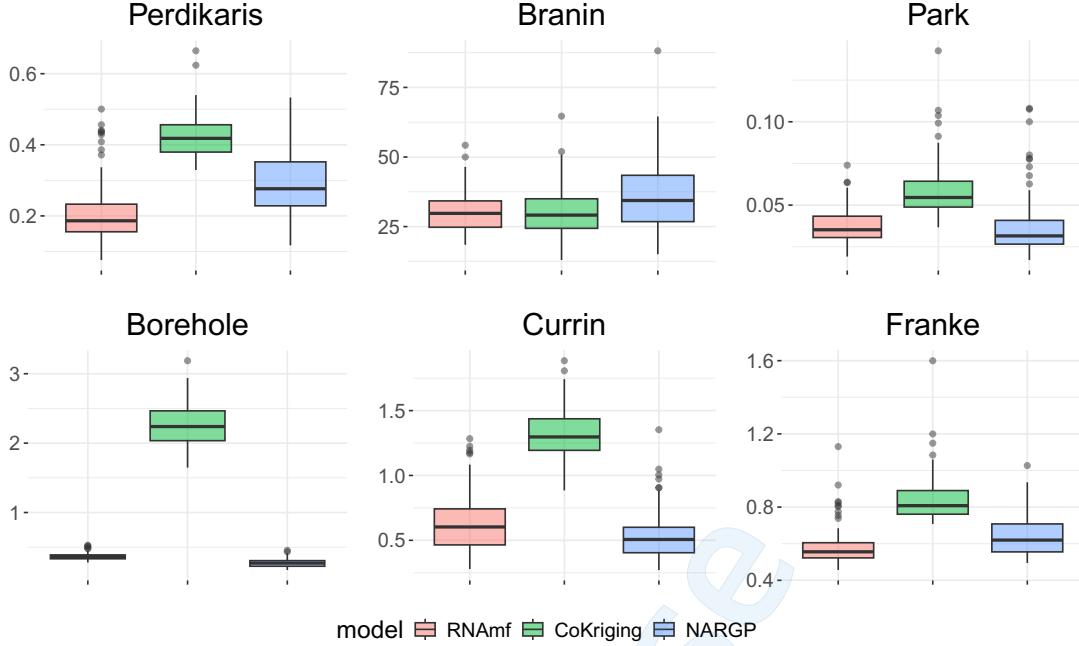


Figure 7: RMSEs of six synthetic examples across 100 repetitions.

1991; Xiong et al., 2013), the three-level Branin function (Sobester et al., 2008), the two-level Borehole function (Morris et al., 1993; Xiong et al., 2013), and the two-level Currin function (Xiong et al., 2013; Kerleguer et al., 2024). Additionally, we introduce a three-level function modified from the Franke function (Franke, 1979). The explicit forms of these functions are available in Supplementary Materials S6.

The data are generated by evaluating these functions at input locations obtained from the nested space-filling design introduced by Le Gratiet and Garnier (2014) with sample sizes $\{n_l\}_{l=1}^L$. The sample sizes and input dimension for each example are outlined in Table S1. To examine the prediction performance, $n_{\text{test}} = 1000$ random test input locations are generated from the same input space. We evaluate the prediction performance based on two criteria: the root-mean-square error (RMSE) and continuous rank probability score (CRPS) (Gneiting and Raftery, 2007), which are defined in Supplementary Materials S7. Note that CRPS serves as a performance metric for the *posterior predictive distribution* of a scalar observation. Lower values for the RMSE and CRPS indicate better model accuracy. Additionally, we assess the computational efficiency by comparing the computation time.

Figures 7 and S14 respectively show the results of RMSE and CRPS metrics across 100 repetitions, each employing a different random nested design for the training input locations. The proposed **RNAmf** consistently outperforms **CoKriging** by both metrics, particularly for examples exhibiting nonlinear relationships between simulators, such as the Perdikaris, Borehole, Currin, and Franke functions. For instances where simulators follow a linear (or nearly linear) auto-regressive model, like the Brainin and Park functions, the proposed **RNAmf** remains competitive with **CoKriging**, which is designed to excel in such scenarios. This highlights the flexibility of our approach, enabled by the GP prior for modeling relationships. On the other hand, **NARGP**, another approach modeling nonlinear relationships, outperforms **CoKriging** in most of the examples and is competitive with **RNAmf**, except in the Perdikaris and Franke examples, where **RNAmf** exhibits superior performance. However, it comes with significantly higher computational costs, as shown in Figure 8, due to its expensive MC approximation, being roughly fifty times slower than both **RNAmf** and **CoKriging** on average. Notably, in scenarios involving three fidelities, including the Brainin and Franke examples, the computational time for **NARGP** exceeds that of **RNAmf** by more than 150 times. This shows that **NARGP** can suffer from intensive computation as the number of fidelity levels increases, while our method remains competitive in this regard. In summary, the performance across these synthetic examples underscores the capability of the proposed method in providing an accurate emulator at a reasonable computational time.

5.2 Active learning performance

With the accurate RNA emulator in place, we now investigate on the performance of AL strategies for the emulator using the proposed criteria. We compare with two existing methods: **CoKriging-CV**, a cokriging-based sequential design utilizing cross-validation techniques (Le Gratiet and Cannamela, 2015), and **MR-SUR**, a sequential design maximizing the rate of stepwise uncertainty reduction using the KO model (Stroh et al., 2022). As for implementing **CoKriging-CV**, we utilized the code provided in the Supplementary Materials of Le Gratiet and Cannamela (2015). Notably, both of these methods employed the (linear) autoregressive model as in (2) in their implementations. To maintain a consistent comparison,

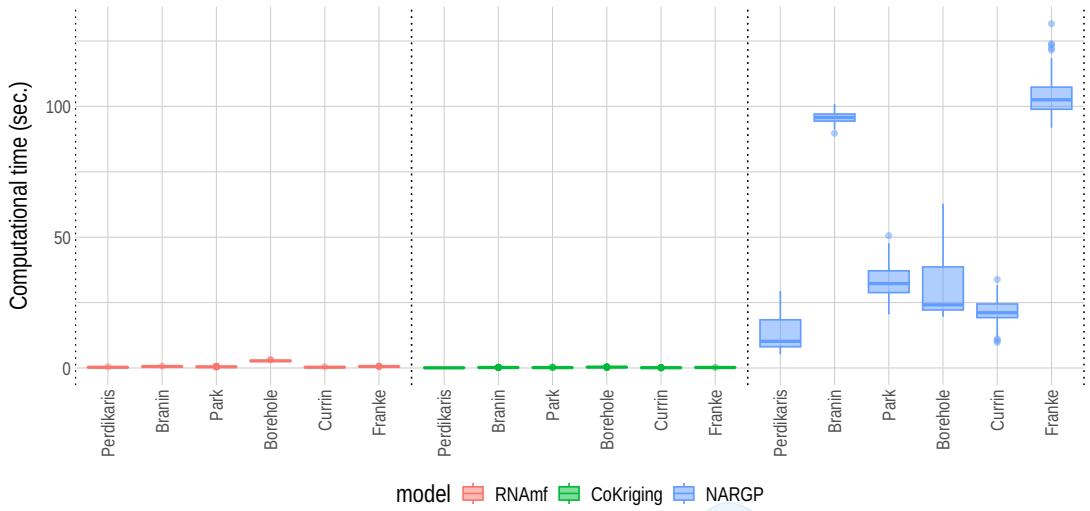


Figure 8: Computational time of six synthetic functions across 100 repetitions.

we use the one-dimensional Perdikaris function (nonlinear) and the 4-dimensional Park function (linear autoregressive) in Section 5.1, to illustrate the performance of these methods.

In this experiment, we suppose that the simulation costs associated with the low- and high-fidelity simulators are $C_1 = 1$ and $C_2 = 3$, respectively. The initial data is established similar to Section 5.1, with sample sizes specified in Table S1. We consider a total simulation budget of $C_{\text{total}} = 80$ for the Perdikaris function and $C_{\text{total}} = 130$ for the Park function. For ALC and ALMC acquisitions which involve the computation of the average reduction in variance as in (18), 1000 and 100 uniform samples are respectively generated from the input space to approximate the integral and impute the future outputs.

Figure 9 shows the results of RMSE and CRPS metrics for the Perdikaris function, with respect to the total simulation costs accrued after each sample selection. The left panel of Figure 10 displays a boxplot depicting the final RMSEs after reaching the total simulation budget across the 10 repetitions. The results show that the proposed AL methods dramatically outperform the two competing methods, CoKriging-CV and MR-SUR, in terms of both accuracy and stability, considering the same costs. As the cost increases, MR-SUR begins to close the gap, while CoKriging-CV lags behind the other methods. Among the four proposed AL strategies, the distinctions are minimal. As noted in Section 4, ALC

acquisitions involve intricate numerical integration approximations and data imputation, taking approximately 400 seconds for each acquisition in this example. In contrast, ALD, ALM and ALMC are significantly more computationally efficient due to the closed-form nature of the criteria, requiring only around 1, 1, and 10 seconds per acquisition, respectively.

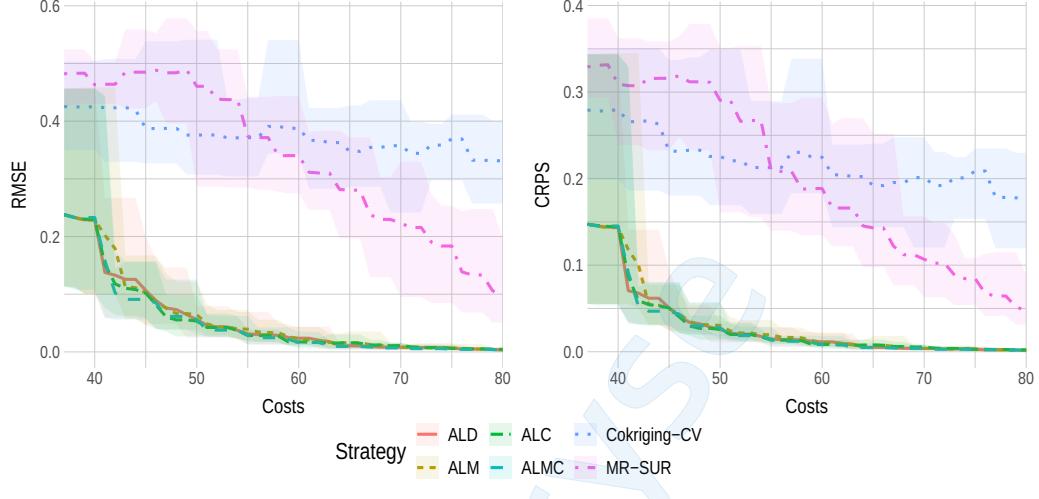


Figure 9: RMSE and CRPS for the Perdikaris function with respect to the simulation cost. Solid lines represent the average over 10 repetitions and shaded regions represent the ranges.

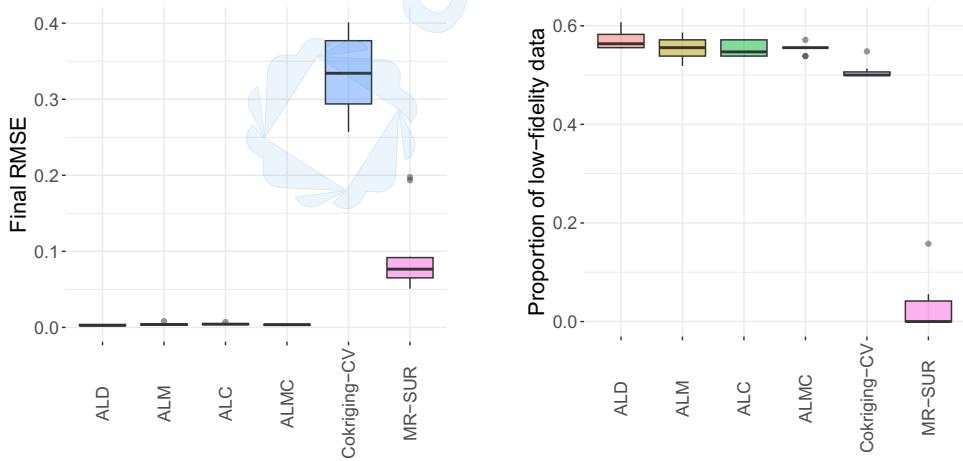


Figure 10: Final RMSE (left) and proportion of AL acquisitions choosing low-fidelity data (right) for the Perdikaris function. Boxplots indicate spread over 10 repetitions.

From the right panel of Figure 10, it can be seen that the proposed AL methods tend

to select low-fidelity simulators more frequently than the other two comparative methods, notably MR-SUR, which consistently chooses samples exclusively from the high-fidelity simulator. This suggests that the proposed RNA model can effectively infer the high-fidelity simulation using primarily low-fidelity data for the nonlinear Perdikaris function, while the other two KO-based methods (CoKriging-CV and MR-SUR) require more high-fidelity data to reduce the uncertainty.

Figures S15 and S16 present the results for the Park function. As expected, the distinctions between these strategies are not as significant because the function aligns more closely with the KO model (linear autoregressive). Nonetheless, our proposed AL strategies still exhibit better average performance. At the final cost budget of $C_{\text{total}} = 130$, ALM and ALMC perform the best, collecting a larger portion of high-fidelity data, as indicated in Figure S16. In contrast, the KO-based strategies collect more low-fidelity data, which is again expected because KO-based models are efficient at leveraging low-fidelity data to infer the high-fidelity simulator. In these scenarios, our strategies efficiently prioritize the selection of high-fidelity data to minimize uncertainty, resulting in superior prediction accuracy at the same cost.

6 Thermal Stress Analysis of Jet Engine Turbine Blade

We leverage our proposed method for a real application involving the analysis of thermal stress in a jet turbine engine blade under steady-state operating conditions. The turbine blade, which forms part of the jet engine, is constructed from nickel alloys capable of withstanding extremely high temperatures. It is crucial for the blade's design to ensure that it can endure stress and deformations while avoiding mechanical failure and friction between the blade tip and the turbine casing. Refer to Carter (2005), Wright and Han (2006), and Sung et al. (2024a,b) for more details.

This problem can be treated as a static structural model and can be solved numerically using finite element methods. There are two input variables denoted as x_1 and x_2 , which represent the pressure load on the pressure and suction sides of the blade, both of which fall within the range of 0.25 to 0.75 MPa, i.e., $\mathbf{x} = (x_1, x_2) \in \Omega = [0.25, 0.75]^2$. The response of

interest is the maximum value over the thermal stress profile, which is a critical parameter used to assess the structural stability of the turbine blade. We perform finite element simulations using the **Partial Differential Equation Toolbox** in MATLAB (MATLAB, 2021).

The simulations are conducted at two fidelity levels, each using different mesh densities for finite element methods. A denser mesh provides higher fidelity and more accurate results but demands greater computational resources. Conversely, a coarser mesh sacrifices some accuracy for reduced computational cost. Figure S17 demonstrates the turbine blade structure and thermal stress profiles obtained at these two fidelity levels for the input location $\mathbf{x} = (0.5, 0.45)$.

We perform the finite element simulations with sample sizes of $n_1 = 20$ and $n_2 = 10$ to examine the emulation performance. Similar to Section 5.1, we use the nested space-filling design of Le Gratiet and Garnier (2014) to generate the input locations of the computer experiments. We record the simulation time of the finite element simulations, which are respectively $C_1 = 2.25$ and $C_2 = 6.85$ (seconds) and will be used later for comparing AL strategies. To examine the performance, we conduct the high-fidelity simulations (i.e. $f_2(\mathbf{x})$) at the test input locations of size $n_{\text{test}} = 100$ generated from a set of Latin hypercube samples from the same design space. The experiment is repeated 10 times, each time considering different nested space-filling designs for the training input locations.

Figure 11 presents a comparison of emulation performance with the other two competing methods, **CoKriging** and **NARGP**. Our proposed method, **RNAmf**, outperforms the other two methods in terms of CRPS and is comparable in terms of RMSE. While **NARGP** delivers competitive prediction performance, it comes at a significantly higher computational cost compared to **RNAmf**.

Figures 12 and 13 present a comparison of the AL strategies with a fixed cost budget of $C_{\text{total}} = 160$ seconds. The right panel of Figure 13 reveals that these strategies collect a similar number of low-fidelity data points. Notably, **CoKriging-CV** exhibits significant variability across the 10 repetitions, so we have removed the shaded region and only show the average, indicating that it yields poorer prediction performance compared to the other strategies. Another KO-based strategy, **MR-SUR**, performs better but still falls short of our

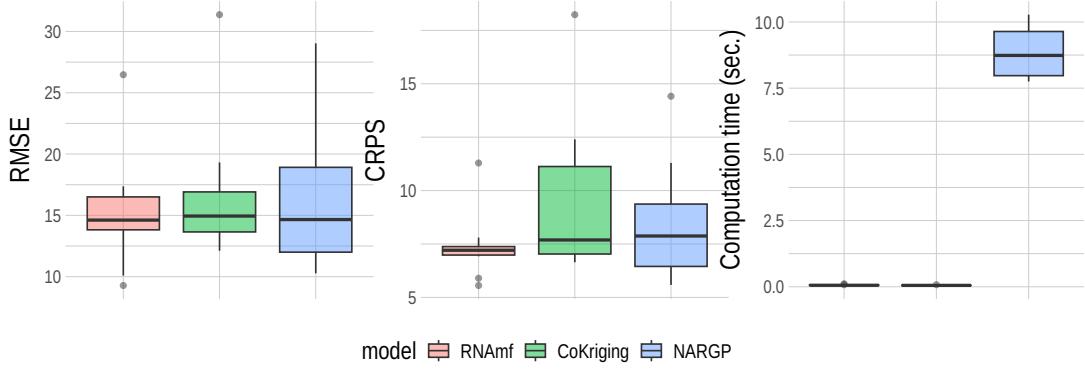


Figure 11: *RMSE, CRPS, and computation time across 10 repetitions in the turbine blade application.*

proposed AL strategies at any given simulation cost. Conversely, our proposed AL strategies demonstrate effective results and outperform the others. This is evident from RMSE and CRPS values exhibiting a leveling-off trend, with final results around 10 and 5, respectively, compared to the initial designs yielding both metrics averaging around 15 and 7. Among the AL strategies, the performance of the four strategies does not show significant differences at the final cost budget.

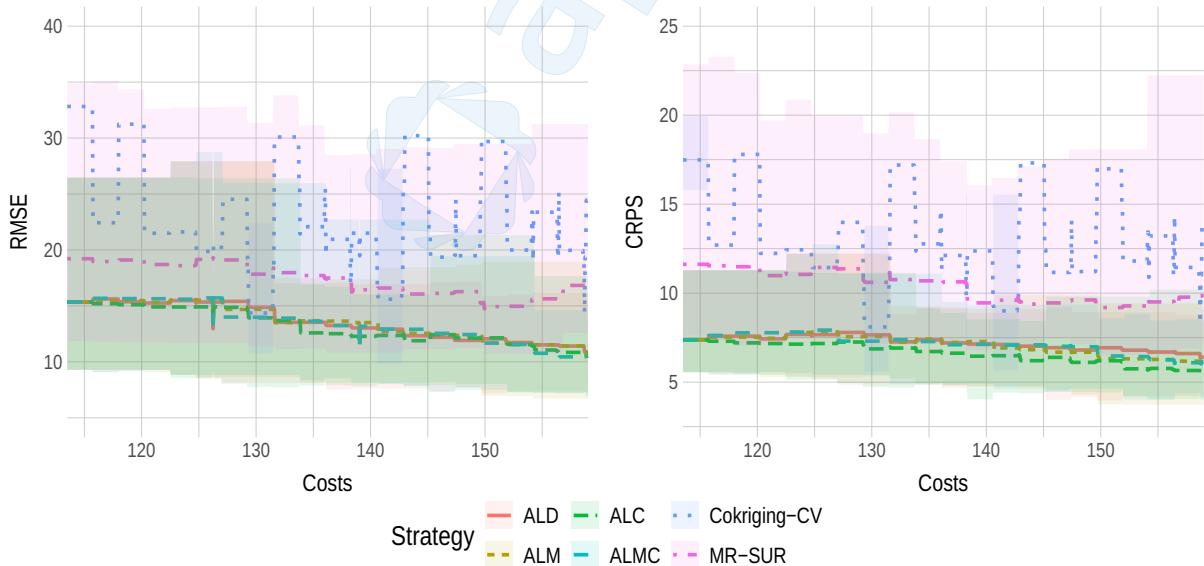


Figure 12: *RMSE and CRPS for the turbine blade application with respect to the cost. Solid lines represent the average over 10 repetitions and shaded regions represent the ranges.*

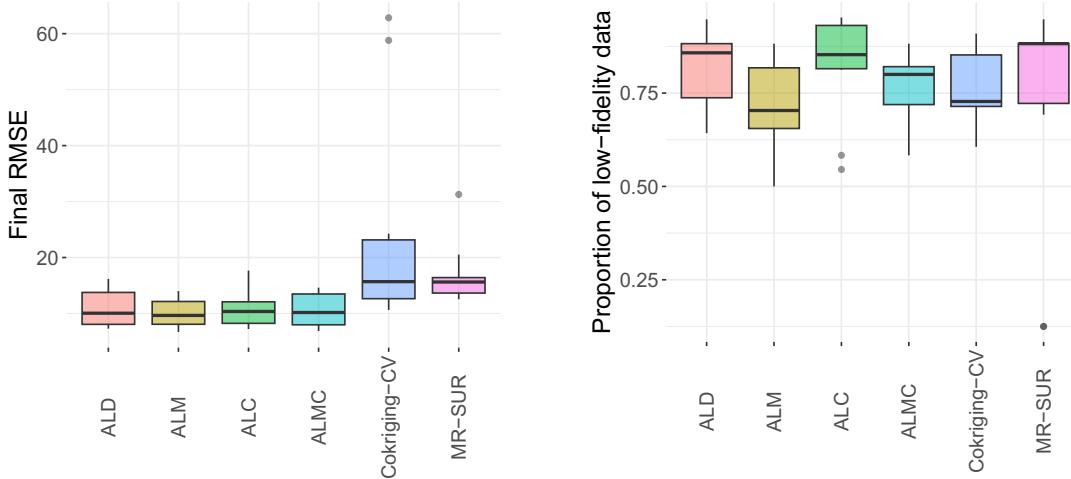


Figure 13: Final RMSE (left) and proportion of AL acquisitions choosing low-fidelity data (right) for the turbine blade application. Boxplots indicate spread over 10 repetitions.

7 Conclusion

Multi-fidelity computer experiments have become an essential tool in simulating complex scientific problems. This paper introduces a new emulator tailored for multi-fidelity simulations, which proves effective in producing accurate, efficient predictions for high-fidelity simulations, especially when dealing with nonlinear relationships between simulators. Building upon this new emulator, we present four AL strategies designed to select optimal input locations and fidelity levels to augment data, thereby enhancing emulation performance.

With the RNA emulator’s success, it is worthwhile to explore emulators and AL strategies built upon similar principles for addressing multi-fidelity problems with *tunable* fidelity parameters, such as mesh density (Picheny et al., 2013; Tuo et al., 2014). Designing experiments for such scenarios presents intriguing challenges, as shown in recent studies (see, e.g., Shaowu Yuchi et al., 2023; Sung et al., 2024a). Furthermore, considering the increasing prevalence of *stochastic computer models* (Baker et al., 2022), extending the proposed RNA emulator to accommodate noisy data would significantly enhance its relevance in real-world applications. While this article assumes noise-free data, introducing noise into the model is a feasible endeavor, a task we leave for our future research.

Supplemental Materials Additional supporting materials can be found in Supplemental Materials, including the closed-form posterior mean and variance under a Matérn kernel, the proof of Proposition 3.1, and the supporting tables and figures for Sections 5 and 6. The R code and package for reproducing the results in Sections 5 and 6 are also provided.

Acknowledgments The authors express sincere gratitude for the conscientious efforts of the associate editor and two anonymous reviewers whose insightful comments greatly strengthen this article.

Disclosure Statement No potential conflict of interest was reported by the authors.

References

Baker, E., Barbillon, P., Fadikar, A., Gramacy, R. B., Herbei, R., Higdon, D., Huang, J., Johnson, L. R., Ma, P., Mondal, A., Pires, B., Sacks, J., and Sokolov, V. (2022). Analyzing stochastic computer models: A review with opportunities. *Statistical Science*, 37(1):64–89.

Beck, J. and Guillas, S. (2016). Sequential design with mutual information for computer experiments (MICE): Emulation of a tsunami model. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):739–766.

Bilionis, I. and Zabaras, N. (2012). Multi-output local gaussian process regression: Applications to uncertainty quantification. *Journal of Computational Physics*, 231(17):5718–5746.

Binois, M., Huang, J., Gramacy, R. B., and Ludkovski, M. (2019). Replication or exploration? sequential design for stochastic simulation experiments. *Technometrics*, 61(1):7–23.

Byrd, R. H., Lu, P., Nocedal, J., and Zhu, C. (1995). A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208.

Carter, T. J. (2005). Common failures in gas turbine blades. *Engineering Failure Analysis*, 12(2):237–247.

Cohn, D. (1993). Neural network exploration using optimal experiment design. *Advances in Neural Information Processing Systems*, 6:1071–1083.

Cutajar, K., Pullin, M., Damianou, A., Lawrence, N., and González, J. (2019). Deep Gaussian processes for multi-fidelity modeling. *arXiv preprint arXiv:1903.07320*.

Demeyer, S., Fischer, N., and Marquis, D. (2017). Surrogate model based sequential sampling estimation of conformance probability for computationally expensive systems: application to fire safety science. *Journal de la société française de statistique*, 158(1):111–138.

Ehara, A. and Guillas, S. (2023). An adaptive strategy for sequential designs of multilevel computer experiments. *International Journal for Uncertainty Quantification*, 13(4):61–98.

Franke, R. (1979). A critical comparison of some methods for interpolation of scattered data. Technical report, Monterey, California: Naval Postgraduate School.

Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.

Gramacy, R. B. (2020). *Surrogates: Gaussian Process Modeling, Design, and Optimization for the Applied Sciences*. CRC press.

Gramacy, R. B. and Lee, H. K. (2009). Adaptive design and analysis of supercomputer experiments. *Technometrics*, 51(2):130–145.

Haaland, B. and Qian, P. Z. G. (2010). An approach to constructing nested space-filling designs for multi-fidelity computer experiments. *Statistica Sinica*, 20(3):1063–1075.

Harville, D. A. (1998). *Matrix Algebra from a Statistician’s Perspective*. New York: Springer-Verlag.

He, X., Tuo, R., and Wu, C. F. J. (2017). Optimization of multi-fidelity computer experiments via the EQIE criterion. *Technometrics*, 59(1):58–68.

Heo, J. and Sung, C.-L. (2024). *RNAmf: Recursive Non-Additive Emulator for Multi-Fidelity Data*. R package version 0.1.2.

Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A. (2006). Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32:369–382.

Ji, Y., Mak, S., Soeder, D., Paquet, J., and Bass, S. A. (2024). A graphical multi-fidelity Gaussian process model, with application to emulation of heavy-ion collisions. *Technometrics*, 66(2):267–281.

Kennedy, M. C. and O’Hagan, A. (2000). Predicting the output from a complex computer code when fast approximations are available. *Biometrika*, 87(1):1–13.

Kerleguer, B., Cannamela, C., and Garnier, J. (2024). A Bayesian neural network approach to multi-fidelity surrogate modelling. *International Journal for Uncertainty Quantification*, 14(1):43–60.

Ko, J. and Kim, H. (2022). Deep gaussian process models for integrating multifidelity experiments with nonstationary relationships. *IISE Transactions*, 54(7):686–698.

Koermer, S., Loda, J., Noble, A., and Gramacy, R. B. (2024). Augmenting a simulation campaign for hybrid computer model and field data experiments. *Technometrics*, to appear.

Kuya, Y., Takeda, K., Zhang, X., and Forrester, A. I. (2011). Multifidelity surrogate modeling of experimental and computational aerodynamic data sets. *AIAA Journal*, 49(2):289–298.

Kyzyurova, K. N., Berger, J. O., and Wolpert, R. L. (2018). Coupling computer models through linking their statistical emulators. *SIAM/ASA Journal on Uncertainty Quantification*, 6(3):1151–1171.

Le Gratiet, L. (2012). *MuFiCokriging: Multi-Fidelity Cokriging models*. R package version 1.2.

Le Gratiet, L. (2013). Bayesian analysis of hierarchical multifidelity codes. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):244–269.

Le Gratiet, L. and Cannamela, C. (2015). Cokriging-based sequential design strategies using fast cross-validation techniques for multi-fidelity computer codes. *Technometrics*, 57(3):418–427.

Le Gratiet, L. and Garnier, J. (2014). Recursive co-kriging model for design of computer experiments with multiple levels of fidelity. *International Journal for Uncertainty Quantification*, 4(5):365–386.

Li, S., Xing, W., Kirby, R., and Zhe, S. (2020). Multi-fidelity Bayesian optimization via deep neural networks. *Advances in Neural Information Processing Systems*, 33:8521–8531.

Ma, P., Karagiannis, G., Konomi, B. A., Asher, T. G., Toro, G. R., and Cox, A. T. (2022). Multifidelity computer model emulation with high-dimensional output: An application to storm surge. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 71(4):861–883.

MacKay, D. J. C. (1992). Information-based objective functions for active data selection. *Neural Computation*, 4(4):590–604.

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.

MATLAB (2021). *MATLAB version 9.11.0.1769968 (R2021b)*. The Mathworks, Inc., Natick, Massachusetts.

McKay, M. D., Beckman, R. J., and Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61.

Meng, X., Babaee, H., and Karniadakis, G. E. (2021). Multi-fidelity Bayesian neural networks: Algorithms and applications. *Journal of Computational Physics*, 438:110361.

Meng, X. and Karniadakis, G. E. (2020). A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems. *Journal of Computational Physics*, 401:109020.

Ming, D. and Guillas, S. (2021). Linked Gaussian process emulation for systems of computer models using Matérn kernels and adaptive design. *SIAM/ASA Journal on Uncertainty Quantification*, 9(4):1615–1642.

Ming, D., Williamson, D., and Guillas, S. (2023). Deep Gaussian process emulation using stochastic imputation. *Technometrics*, 65(2):150–161.

Morris, M. D., Mitchell, T. J., and Ylvisaker, D. (1993). Bayesian design and analysis of computer experiments: use of derivatives in surface prediction. *Technometrics*, 35(3):243–255.

Paleyes, A., Mahsereci, M., and Lawrence, N. D. (2023). Emukit: A Python toolkit for decision making under uncertainty. *Proceedings of the Python in Science Conference*.

Paleyes, A., Pullin, M., Mahsereci, M., McCollum, C., Lawrence, N., and González, J. (2019). Emulation of physical processes with Emukit. In *Second Workshop on Machine Learning and the Physical Sciences, NeurIPS*.

Park, C., Waelder, R., Kang, B., Maruyama, B., Hong, S., and Gramacy, R. B. (2023). Active learning of piecewise Gaussian process surrogates. *arXiv preprint arXiv:2301.08789*.

Park, J. S. (1991). *Tuning Complex Computer Codes to Data and Optimal Designs*. University of Illinois at Urbana-Champaign.

Patra, A., Batra, R., Chandrasekaran, A., Kim, C., Huan, T. D., and Ramprasad, R. (2020). A multi-fidelity information-fusion approach to machine learn and predict polymer bandgap. *Computational Materials Science*, 172:109286.

Perdikaris, P. (2016). Multi-fidelity modeling using Gaussian processes and nonlinear auto-regressive schemes. <https://github.com/paraklas/NARGP>.

Perdikaris, P., Raissi, M., Damianou, A., Lawrence, N. D., and Karniadakis, G. E. (2017). Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2198):20160751.

Picheny, V., Ginsbourger, D., Richet, Y., and Caplin, G. (2013). Quantile-based optimization of noisy computer experiments with tunable precision. *Technometrics*, 55(1):2–13.

Qian, P. Z. G. (2009). Nested Latin hypercube designs. *Biometrika*, 96(4):957–970.

Qian, P. Z. G., Ai, M., and Wu, C. F. J. (2009). Construction of nested space-filling designs. *Annals of Statistics*, 37(6A):3616–3643.

Qian, P. Z. G. and Wu, C. F. J. (2008). Bayesian hierarchical modeling for integrating low-accuracy and high-accuracy experiments. *Technometrics*, 50(2):192–204.

Qian, Z., Seepersad, C. C., Joseph, V. R., Allen, J. K., and Wu, C. F. J. (2006). Building surrogate models based on detailed and approximate simulations. *Journal of Mechanical Design*, 128(4):668–677.

R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

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

Santner, T. J., Williams, B. J., and Notz, W. I. (2018). *The Design and Analysis of Computer Experiments*. Springer New York.

Sauer, A., Gramacy, R. B., and Higdon, D. (2023). Active learning for deep Gaussian process surrogates. *Technometrics*, 65(1):4–18.

Seo, S., Wallat, M., Graepel, T., and Obermayer, K. (2000). Gaussian process regression: Active data selection and test point rejection. In *Mustererkennung 2000: 22. DAGM-Symposium. Kiel, 13.–15. September 2000*, pages 27–34. Springer.

Shaowu Yuchi, H., Roshan Joseph, V., and Wu, C. F. J. (2023). Design and analysis of multifidelity finite element simulations. *Journal of Mechanical Design*, 145(6):061703.

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

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

Stroh, R., Bect, J., Demeyer, S., Fischer, N., Marquis, D., and Vazquez, E. (2022). Sequential design of multi-fidelity computer experiments: maximizing the rate of stepwise uncertainty reduction. *Technometrics*, 64(2):199–209.

Sung, C.-L., Hung, Y., Rittase, W., Zhu, C., and Wu, C. F. J. (2020). Calibration for computer experiments with binary responses and application to cell adhesion study. *Journal of the American Statistical Association*, 115(532):1664–1674.

Sung, C.-L., Ji, Y., Mak, S., Wang, W., and Tang, T. (2024a). Stacking designs: Designing multifidelity computer experiments with target predictive accuracy. *SIAM/ASA Journal on Uncertainty Quantification*, 12(1):157–181.

Sung, C.-L., Wang, W., Ding, L., and Wang, X. (2024b). Mesh-clustered Gaussian process emulator for partial differential equation boundary value problems. *Technometrics*, to appear.

Swersky, K., Snoek, J., and Adams, R. P. (2013). Multi-task Bayesian optimization. *Advances in Neural Information Processing Systems*, 26:2004–2012.

Tuo, R., Wu, C. F. J., and Yu, D. (2014). Surrogate modeling of computer experiments with different mesh densities. *Technometrics*, 56(3):372–380.

Wright, L. M. and Han, J.-C. (2006). Enhanced internal cooling of turbine blades and vanes. *The Gas Turbine Handbook*, 4:1–5.

Xiong, S., Qian, P. Z. G., and Wu, C. F. J. (2013). Sequential design and analysis of high-accuracy and low-accuracy computer codes. *Technometrics*, 55(1):37–46.