

SPARSE GAUSSIAN PROCESSES FOR SOLVING NONLINEAR PDES

RUI MENG¹, XIANJIN YANG^{2,3,*}

¹*Lawrence Berkeley National Laboratory, Berkeley, California, USA.*

²*Yau Mathematical Sciences Center, Tsinghua University, Haidian District, Beijing, 100084, China.*

³*Beijing Institute of Mathematical Sciences and Applications, Huairou District, Beijing, 101408, China.*

ABSTRACT. This article proposes an efficient numerical method for solving nonlinear partial differential equations (PDEs) based on sparse Gaussian processes (SGPs). Gaussian processes (GPs) have been extensively studied for solving PDEs by formulating the problem of finding a reproducing kernel Hilbert space (RKHS) to approximate a PDE solution. The approximated solution lies in the span of base functions generated by evaluating derivatives of different orders of kernels at sample points. However, the RKHS specified by GPs can result in an expensive computational burden due to the cubic computation order of the matrix inverse. Therefore, we conjecture that a solution exists on a “condensed” subspace that can achieve similar approximation performance, and we propose a SGP-based method to reformulate the optimization problem in the “condensed” subspace. This significantly reduces the computation burden while retaining desirable accuracy. The paper rigorously formulates this problem and provides error analysis and numerical experiments to demonstrate the effectiveness of this method. The numerical experiments show that the SGP method uses fewer than half the uniform samples as inducing points and achieves comparable accuracy to the GP method using the same number of uniform samples, resulting in a significant reduction in computational cost.

Our contributions include formulating the nonlinear PDE problem as an optimization problem on a “condensed” subspace of RKHS using SGP, as well as providing an existence proof and rigorous error analysis. Furthermore, our method can be viewed as an extension of the GP method to account for general positive semi-definite kernels.

1. INTRODUCTION

This work develops a numerical method based on sparse Gaussian processes (SGPs) [19, 37] to solve nonlinear partial differential equations (PDEs). PDEs have been widely used to model applications in science, economics, and biology [25, 36]. However, very few PDEs admit explicit solutions. Standard numerical approaches including finite difference [36] and finite element methods [13] for solving PDEs are prone to the curse of dimensions. Recently, to keep pace with increasing problem sizes, machine learning methods have attracted a lot of attention [4, 20, 21, 29, 42]. Probabilistic algorithms to solve linear problems are presented in [6, 22, 23]. The authors of [4] extend these ideas and propose a Gaussian process (GP) framework for solving general nonlinear PDEs. In particular, for time-dependent PDEs, GP methods based on time-discretization are considered in [15, 28, 38]. Most of the algorithms mentioned above seek approximations of a solution to a PDE in a normed vector space with base functions chosen beforehand. For instance, for the GP method in [4], by the representer theorem [24, Section 17.8], a solution of a PDE is approximated in the span of base functions, each of which is associated with a sample. If the samples are not chosen properly,

E-mail address: mengrui6351@gmail.com, yxjmath@gmail.com.

Date: July 4, 2023.

2010 Mathematics Subject Classification. 65L30, 65N75, 82C80, 35A01.

Key words and phrases. Partial Differential Equations; Sparse Gaussian Process.

*Corresponding author.




FIG. 1. The illustration of Problem 1. u^* is the solution of a PDE in the space \mathcal{U} . Given an approximation $u_{\tilde{\mathcal{U}}}^\dagger$ of u^* in the subspace $\tilde{\mathcal{U}}$ of \mathcal{U} such that $\|u_{\tilde{\mathcal{U}}}^\dagger - u^*\|_{\mathcal{U}} \leq \epsilon$ for some $\epsilon > 0$, and given a positive number δ such that $\epsilon \leq \delta$, we “compress” $\tilde{\mathcal{U}}$ to find a minimum subspace \mathcal{U}_c of $\tilde{\mathcal{U}}$ such that the approximation $u_{\mathcal{U}_c}^\dagger$ of u^* in \mathcal{U}_c satisfies $\|u^* - u_{\mathcal{U}_c}^\dagger\|_{\mathcal{U}} \leq \delta$.

the base functions may contain redundant information about the approximated solution. Therefore, a set of random samples may result in consuming unnecessary extra storage and in a waste of computational time. The computational cost of the GP method [4] grows cubically with respect to the number of samples due to the inversion of covariance matrices of samples. To improve the performance of [4], [21] proposes to approximate the solution of a PDE in the space generated by random Fourier features [27, 41], where the corresponding covariance matrix is a low-rank approximation to that of the GP method in [4]. The numerical experiments in [21] show that approximating solutions of PDEs in the spaces generated by low-rank kernels reduces the precomputation time, saves the memory, and achieves comparable accuracy to the GP method. This motivates us to question whether the space we use to approximate a solution of a PDE contains redundant information for general numerical methods of solving nonlinear PDEs and whether we can “compress” the space until it reaches a minimum subspace but contains sufficient information to express a solution of the PDE. In other words, we are interested in the following problem:

Problem 1. Let \mathcal{U} be a normed vector space equipped with the norm $\|\cdot\|_{\mathcal{U}}$ and let $u^* \in \mathcal{U}$ be the solution of a nonlinear PDE. For any $\mathcal{W} \subset \mathcal{U}$, denote by $u_{\mathcal{W}}^\dagger$ a best approximation of u^* in \mathcal{W} in terms of a specific numerical method and by $|\mathcal{W}|$ the cardinality of \mathcal{W} . Let $u_{\tilde{\mathcal{U}}}^\dagger$ be the approximation of u^* in a subspace $\tilde{\mathcal{U}}$ of \mathcal{U} , and $\|u_{\tilde{\mathcal{U}}}^\dagger - u^*\|_{\mathcal{U}} \leq \epsilon$ for some $\epsilon > 0$. Given a positive number δ such that $\epsilon \leq \delta$, find a subspace \mathcal{U}_c , with the minimum cardinality, of $\tilde{\mathcal{U}}$ such that the approximation $u_{\mathcal{U}_c}^\dagger$ of u^* in \mathcal{U}_c satisfies $\|u^* - u_{\mathcal{U}_c}^\dagger\|_{\mathcal{U}} \leq \delta$. That is, we find \mathcal{U}_c such that

$$\mathcal{U}_c = \operatorname{argmin}_{\mathcal{V} \subset \tilde{\mathcal{U}}, \|u^* - u_{\mathcal{V}}^\dagger\|_{\mathcal{U}} \leq \delta} |\mathcal{V}|.$$

Problem 1 is concerned about the efficient information in the space $\tilde{\mathcal{U}}$ used to approximate the true solution u^* , which is illustrated in Figure 1. Using a “condensed” space not only saves the storage but reduces the computational costs. Merely decreasing the number of base functions of \mathcal{U} may not sufficiently handle Problem 1, which is implied by the numerical experiments in Section 4 (see discussions in Subsection 4.1). A thorough general discussion of Problem 1 for all the numerical methods of solving nonlinear PDEs is out of the scope of this paper. Here, we propose to study Problem 1 in terms of the framework in [4], where a numerical approximation of the solution to a nonlinear PDE is viewed as the maximum *a posteriori* (MAP) estimator of a GP conditioned on solving the PDE at a finite set of sample points. By using the representer theorem [24, Section 17.8], the approximated solution lies in the span of base functions which are obtained by evaluating kernel functions or derivatives of kernels at sample points. Thus, base functions generated by random samples may provide redundant information about the approximated solution. To address Problem 1 with respect to the GP method in [4], we propose a low-rank approximation method based on SGPs, which leverages a few inducing points to summarize the information from observations. In the SGP method, we

replace the Gaussian prior by an alternative with a zero mean and a low-rank kernel induced from inducing points. The approximation to the solution of a PDE can be viewed as a MAP estimator of a SGP conditioned on a noisy observation about the values of linear operators of the true solution, where the observation satisfies the PDE at the samples. Meanwhile, the numerical approximation can also be viewed as a function with the minimum norm in the reproducing kernel Hilbert space (RKHS) associated with the low-rank kernel. By Theorem 3.4 in Section 3, a solution given by the SGP method lies in the span of base functions generated by evaluating kernels or derivatives of kernels at inducing points. Thus, if a small number of inducing points are sufficient to capture the information in the sample points, the RKHS generated by the low-rank kernel is sufficient to serve as a candidate of the minimum subspace in Problem 1 for the GP method [4]. By a stable Woodbury matrix identity stated in Subsection 2.2, the dimension of the matrix to be inverted in the SGP method depends only on the number of inducing points selected. In the numerical experiments in Section 4, the SGP method reduces the computational cost by sacrificing only negligible amounts of accuracy by using half of the samples as inducing points.

This inducing-point approach is initially proposed in the deterministic training conditional (DTC) approximation [32], and then it has been fully studied in several works. [26] provides a unifying view for the approximation of GP regressions. This framework includes the DTC, the fully independent training conditional (FITC) approximation [35], and the partial independent training conditional (PITC) approximation. All the above methods modify the joint prior using a conditional independence assumption between training and test data given inducing variables. Alternatively, [12, 37] leverage the inducing-point approach into the computing of the evidence lower bound of the log marginal likelihood. They retain exact priors but approximate the posteriors via variational inference. On the other hand, the estimation of inducing inputs has been generalized into an augmented feature space in [16]. In particular, SGPs are extended in the spectral domain with and without variational inference [11, 17]. Moreover, some works directly speed up the computation of GP regressions through fast matrix-vector multiplication and pre-conditional conjugate gradients [9], and through structured kernel interpolation [40].

Our contributions are summarized as follows:

1. We introduce a general framework based on SGPs to solve nonlinear PDEs. Our method seeks approximated solutions in RKHSs associated with positive semi-definite kernels generated by inducing points. The SGP method provides a way of “compressing” function spaces such that “condensed” subspaces provide sufficient information about approximated solutions to PDEs. Meanwhile, the framework in this paper can be viewed as an extension of the one in [4] to take into consideration of more general positive semi-definite kernels;
2. In the SGP method, the inversion of the covariance matrix requires only the inversion of a matrix whose size is proportional to the number of inducing points which is much less than the number of samples. The numerical experiments in Section 4 show that the SGP method, which leverages less than half of the uniform samples as inducing points, reduces the computational cost by sacrificing only negligible amounts of accuracy compared to the GP method in [4] using the same number of uniform samples;
3. We give a rigorous existence proof for the numerical approximation to the solution of a general PDE and analyze error bounds in Section 3. Our analysis bases on the arguments in [4], learning theory [33, 34], and Nyström approximations [14]. The error bound given in Theorem 3.7 provides qualitative hints to improve the accuracy of our method;
4. The numerical experiments show that the choice of the positions of inducing points and parameters of kernels has a profound impact on the accuracy of the numerical approximations. In Section 2, we give a probabilistic interpretation for the SGP method in the setting of [37], which motivates future work for hyperparameter learning.

The paper is organized as follows. Section 2 summarizes our SGP method by solving a nonlinear elliptic PDE. In Section 3, we provide a framework of the SGP method to solve general PDEs. Numerical experiments are presented in Section 4. Further discussion and future work appear in Section 5.

Notations. Vectors are in bold font. For a real-valued vector \mathbf{v} , we represent by $|\mathbf{v}|$ the Euclidean norm of \mathbf{v} and by \mathbf{v}^T its transpose. For a matrix A , we denote by $\|A\|$ the spectral norm of A . Let $\langle \mathbf{u}, \mathbf{v} \rangle$ or $\mathbf{u}^T \mathbf{v}$ be the inner product of vectors \mathbf{u} and \mathbf{v} . For a normed vector space V , let $\|\cdot\|_V$ be the norm of V . Let \mathcal{U} be a Banach space associated with a quadratic norm $\|\cdot\|_{\mathcal{U}}$ and let \mathcal{U}^* be the dual of \mathcal{U} . Denote by $[\cdot, \cdot]$ the duality pairing between \mathcal{U}^* and \mathcal{U} . We assume that there exists a linear, bijective, symmetric ($[\mathcal{K}_{\mathcal{U}}\phi, \psi] = [\mathcal{K}_{\mathcal{U}}\psi, \phi]$), and positive ($[\mathcal{K}_{\mathcal{U}}\phi, \phi] > 0$ for $\phi \neq 0$) covariance operator $\mathcal{K}_{\mathcal{U}} : \mathcal{U}^* \mapsto \mathcal{U}$, such that

$$\|u\|_{\mathcal{U}}^2 = [\mathcal{K}^{-1}u, u], \forall u \in \mathcal{U}.$$

Let $\{\phi_i\}_{i=1}^P$ be a set of $P \in \mathbb{N}$ elements in \mathcal{U}^* and let $\phi := (\phi_1, \dots, \phi_P)$ be in the product space $(\mathcal{U}^*)^{\otimes P}$. Then, for $u \in \mathcal{U}$, we denote the pairing $[\phi, u]$ by

$$[\phi, u] := ([\phi_1, u], \dots, [\phi_P, u]).$$

Furthermore, for $\mathbf{u} := (u_1, \dots, u_S) \in \mathcal{U}^{\otimes S}$, $S \in \mathbb{N}$, we represent by $[\phi, \mathbf{u}] \in \mathbb{R}^{P \times S}$ the matrix with entries $[\phi_i, u_j]$. We write $\|\cdot\|$ as the induced norm for linear operators on \mathcal{U} , i.e., for any $\mathcal{F} : \mathcal{U} \mapsto \mathcal{U}$, we define $\|\mathcal{F}\| = \max_{\|u\|_{\mathcal{U}}=1} \|\mathcal{F}(u)\|_{\mathcal{U}}$. Given any $R \in \mathbb{N}$, we denote by $[R]$ the set of nonnegative integers less equal to R . Finally, we represent by C a positive real number whose value may change line by line.

2. A RECAPITULATION OF THE SGP METHOD

To briefly summarize the SGP method and compare it to the GP method proposed in [4], in Subsection 2.1, we demonstrate the SGP method by solving a nonlinear elliptic PDE (2.1). Both the GP and the SGP methods can be applied to more general forms of PDEs, which is discussed in Section 3. Furthermore, we can naturally extend what we discuss below to solve PDE systems. We show a numerical example about solving a PDE system in Subsection 4.2. The SGP method approximates the solution of the PDE by a minimizer of an optimal recovery problem, which searches a solution with the minimum norm in the RKHS associated with a low-rank kernel generated by inducing points and where the values of linear operators of the solution are close to those satisfying the PDE at finite samples. From a probabilistic perspective, the optimal recovery problem can also be interpreted as the MAP estimation for a SGP constrained by a noisy observation of linear operator values of the true solution. The SGP method also requires the inversion of a covariance matrix. In Subsection 2.2, we describe a robust way to compute the inverse of the covariance matrix, which is the cornerstone of the SGP method to reduce computational complexity. Finally, in Subsection 2.3, we give a probabilistic explanation for the SGP method, which motivates a way to choose hyperparameters in future work.

2.1. The SGP Method. Let Ω be a bounded open subset of \mathbb{R}^2 with a Lipschitz boundary $\partial\Omega$. Given continuous functions $f : \Omega \mapsto \mathbb{R}$ and $g : \partial\Omega \mapsto \mathbb{R}$, we seek a function u^* solving

$$\begin{cases} \Delta u^*(x) = u^*(x)(\partial_{x_1} u^*(x) + \partial_{x_2} u^*(x)) + f(x), \forall x := (x_1, x_2) \in \Omega, \\ u^*(x) = g(x), \forall x := (x_1, x_2) \in \partial\Omega, \end{cases} \quad (2.1)$$

where we assume that (2.1) admits a unique strong solution. We propose to approximate the solution to (2.1) by functions in a RKHS generated by inducing points. First, we show the construction of such RKHS. We assume that the solution of (2.1) is in the space \mathcal{U} , where \mathcal{U} is a RKHS associated with a nondegenerate, symmetric, and positive definite kernel $K : \bar{\Omega} \times \bar{\Omega} \mapsto \mathbb{R}$. We view \mathcal{U} as the ambient space. Let \mathcal{U}^* be the dual space of \mathcal{U} and let $[\cdot, \cdot] : \mathcal{U}^* \times \mathcal{U} \mapsto \mathbb{R}$ be the duality pairing. Next, we sample M inducing points $\{\hat{\mathbf{x}}_j\}_{j=1}^M$ in $\bar{\Omega}$ such that $\{\hat{\mathbf{x}}_j\}_{j=1}^{M_{\Omega}} \subset \Omega$ and $\{\hat{\mathbf{x}}_j\}_{j=M_{\Omega}+1}^M \subset \partial\Omega$ for $1 \leq M_{\Omega} \leq M$. Then, we define the functionals $\phi_j^{(1)} = \delta_{\hat{\mathbf{x}}_j}$ for $1 \leq j \leq M$, $\phi_k^{(2)} = \delta_{\hat{\mathbf{x}}_k} \circ (\partial_{x_1} + \partial_{x_2})$ and $\phi_k^{(3)} = \delta_{\hat{\mathbf{x}}_k} \circ \Delta$ for $1 \leq k \leq M_{\Omega}$. Denote by $\phi^{(1)}$, $\phi^{(2)}$, $\phi^{(3)}$ the vectors with entries $\phi_j^{(1)}$, $\phi_k^{(2)}$, $\phi_k^{(3)}$, respectively. Let ϕ be the vector concatenating $\phi^{(1)}$, $\phi^{(2)}$, and $\phi^{(3)}$. For simplicity, we denote by ϕ_j the j^{th} component of ϕ . Define $K(\cdot, \phi)$ as the vector with entries $\int K(\cdot, \mathbf{x}') \phi_j(\mathbf{x}') d\mathbf{x}'$ and define $K(\phi, \phi)$ as the matrix with elements $\int \int K(\mathbf{x}, \mathbf{x}') \phi_i(\mathbf{x}) \phi_j(\mathbf{x}') d\mathbf{x} d\mathbf{x}'$.

Suppose that the inducing points are chosen such that $K(\phi, \phi)$ is positive definite. Then, we define a new kernel $Q : \bar{\Omega} \times \bar{\Omega} \mapsto \mathbb{R}$ as

$$Q(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}, \phi)(K(\phi, \phi))^{-1}K(\phi, \mathbf{x}').$$

Using the positivity of the kernel K , we see that Q is positive semi-definite (see Lemma 3.1 in Section 3). Thus, by the Moore–Aronszajn theorem, Q induces a unique RKHS, denoted by \mathcal{U}_Q . We call \mathcal{U}_Q the RKHS generated by inducing points.

Next, we approximate the solution u in \mathcal{U}_Q . We take N samples $\{\mathbf{x}_i\}_{i=1}^N$ such that $\{\mathbf{x}_i\}_{i=1}^{N_\Omega} \subset \Omega$ and $\{\mathbf{x}_i\}_{i=N_\Omega+1}^N \subset \partial\Omega$ for $1 \leq N_\Omega \leq N$. Let $\|\cdot\|_{\mathcal{U}_Q}$ be the norm of \mathcal{U}_Q . To find an approximated solution in \mathcal{U}_Q , for the first attempt, we consider a similar formulation to that of the GP method in [4] and solve the optimal recovery problem

$$\begin{cases} \min_{u \in \mathcal{U}_Q} \|u\|_{\mathcal{U}_Q}^2 \\ \text{s.t. } \Delta u(\mathbf{x}_j) = u(\mathbf{x}_j)(\partial_{x_1} u(\mathbf{x}_j) + \partial_{x_2} u(\mathbf{x}_j)) + f(\mathbf{x}_j), \forall j = 1, \dots, N_\Omega, \\ u(\mathbf{x}_j) = g(\mathbf{x}_j), \forall j = N_\Omega + 1, \dots, N. \end{cases} \quad (2.2)$$

However, since Q is only positive semi-definite, the representer theorem [24, Section 17.8] does not apply. Furthermore, since \mathcal{U}_Q is a subspace of \mathcal{U} , the solution u^* to (2.1) may not lie in \mathcal{U}_Q . Hence, we may not find a function in \mathcal{U}_Q satisfying the constraints in (2.2). Thus, instead of seeking a function in \mathcal{U}_Q satisfying the constraints of (2.2) exactly, we consider a relaxed version of (2.2).

More precisely, our SGP method introduces a slack vector \mathbf{z} and approximates the solution u^* of (2.1) with a minimizer of the following regularized optimal recovery problem

$$\begin{cases} \min_{\mathbf{z} \in \mathbb{R}^{2N_\Omega+N}, u \in \mathcal{U}_Q} \sum_{j=1}^N |z_j^{(1)} - u(\mathbf{x}_j)|^2 + \sum_{j=1}^{N_\Omega} |z_j^{(2)} - (\partial_{x_1} + \partial_{x_2})u(\mathbf{x}_j)|^2 + \sum_{j=1}^{N_\Omega} |z_j^{(3)} - \Delta u(\mathbf{x}_j)|^2 + \gamma \|u\|_{\mathcal{U}_Q}^2 \\ \text{s.t. } z_j^{(3)} = z_j^{(1)} z_j^{(2)} + f(\mathbf{x}_j), \forall j = 1, \dots, N_\Omega, \\ z_j^{(1)} = g(\mathbf{x}_j), \forall j = N_\Omega + 1, \dots, N, \end{cases} \quad (2.3)$$

where $\gamma > 0$ is a given small regularization parameter and

$$\mathbf{z} := (z_1^{(1)}, \dots, z_{N_\Omega}^{(1)}, z_{N_\Omega+1}^{(1)}, \dots, z_N^{(1)}, z_1^{(2)}, \dots, z_{N_\Omega}^{(2)}, z_1^{(3)}, \dots, z_{N_\Omega}^{(3)}).$$

Let $(u^\dagger, \mathbf{z}^\dagger)$ be a minimizer of (2.3). A minimizer u^\dagger of (2.3) can be viewed as a MAP estimator of a SGP $\xi \sim \mathcal{N}(0, Q)$ conditioned on a noisy observation \mathbf{z}^\dagger about values of the linear operators of u^\dagger , where \mathbf{z}^\dagger satisfies the PDE at the sample points. A probabilistic interpretation of (2.3) is given in Subsection 2.3.

To solve the infinite dimensional minimization problem (2.3), we derive a representer formula for a minimizer u . Define $\boldsymbol{\psi}$ as the vector consisting of $\delta_{\mathbf{x}_j}$, $\delta_{\mathbf{x}_k} \circ (\partial_{x_1} + \partial_{x_2})$, and $\delta_{\mathbf{x}_k} \circ \Delta$, for $1 \leq j \leq N$ and $1 \leq k \leq N_\Omega$. We represent by ψ_j the j^{th} component of $\boldsymbol{\psi}$. Let $Q(\cdot, \boldsymbol{\psi})$ be the vector containing entries $\int Q(\cdot, \mathbf{x}') \psi_j(\mathbf{x}') d\mathbf{x}'$ and let $Q(\boldsymbol{\psi}, \boldsymbol{\psi})$ be the matrix consisting of elements $\int \int Q(\mathbf{x}, \mathbf{x}') \psi_i(\mathbf{x}) \psi_j(\mathbf{x}') d\mathbf{x} d\mathbf{x}'$. Theorem 3.4 in Section 3 shows that (2.3) admits a minimizer u^\dagger such that

$$u^\dagger(x) = Q(x, \boldsymbol{\psi})^T (\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathbf{z}^\dagger, \quad (2.4)$$

where I is the identity matrix and \mathbf{z}^\dagger is a minimizer of

$$\begin{cases} \min_{\mathbf{z} \in \mathbb{R}^{2N_\Omega+N}} \mathbf{z}^T (\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathbf{z} \\ \text{s.t. } z_j^{(3)} = z_j^{(1)} z_j^{(2)} + f(\mathbf{x}_j), \forall j = 1, \dots, N_\Omega, \\ z_j^{(1)} = g(\mathbf{x}_j), \forall j = N_\Omega + 1, \dots, N. \end{cases} \quad (2.5)$$

Next, we use the technique of eliminating variables in [4] to remove the constraints of (2.5). More precisely, observing $z_j^{(1)} = g(\mathbf{x}_j)$ and $z_j^{(3)} = z_j^{(1)}z_j^{(2)} - f(\mathbf{x}_j)$, we rewrite (2.5) as

$$\min_{(\mathbf{z}^{(1)}, \mathbf{z}^{(2)}) \in \mathbb{R}^{2N_\Omega}} (\mathbf{z}^{(1)}, g(\mathbf{x}_{\partial\Omega}), \mathbf{z}^{(2)}, \mathbf{z}^{(1)}\mathbf{z}^{(2)} - f(\mathbf{x}_\Omega))(\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \begin{pmatrix} \mathbf{z}^{(1)} \\ g(\mathbf{x}_{\partial\Omega}) \\ \mathbf{z}^{(2)} \\ \mathbf{z}^{(1)}\mathbf{z}^{(2)} - f(\mathbf{x}_\Omega) \end{pmatrix}, \quad (2.6)$$

where $\mathbf{z}^{(1)} = (z_j^{(1)})_{j=1}^{N_\Omega}$, $\mathbf{z}^{(2)} = (z_j^{(2)})_{j=1}^{N_\Omega}$, $g(\mathbf{x}_{\partial\Omega}) = (g(\mathbf{x}_j))_{j=N_\Omega+1}^N$, $f(\mathbf{x}_\Omega) = (f(\mathbf{x}_j))_{j=1}^{N_\Omega}$, and $\mathbf{z}^{(1)}\mathbf{z}^{(2)} = (z_j^{(1)}z_j^{(2)})_{j=1}^{N_\Omega}$. Then, we obtain \mathbf{z}^\dagger by solving (2.6) using the Gauss–Newton method and get u^\dagger by (2.4).

2.2. Computing $(\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}$. The bottleneck of solving (2.6) is to calculate the inverse of $\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi})$, whose size is proportional to the number of sample points. The computational complexity of standard Cholesky decomposition algorithms is $O(N^3)$ in general. Some fast algorithms [3, 31] are developed by considering sparse inverse Cholesky factorizations. However, their works only consider the case where the covariance matrix contains pointwise evaluations but not pointwise values of differential operators, which are required in the SGP method. In a recent study by [5], a sparse Cholesky factorization algorithm was proposed for kernel matrices, leveraging the sparsity of the Cholesky factor through a novel ordering of Diracs and derivative measurements. Our approach, along with the algorithm described in [5], focuses on effectively solving nonlinear PDEs. Our method accomplishes this objective by approximating representation functions using a smaller set of bases. In contrast, their approach achieves the same goal by efficiently performing computations with dense kernel matrices through a novel sparse Cholesky decomposition technique.

In this paper, by using the structure of $Q(\boldsymbol{\psi}, \boldsymbol{\psi})$, we show that the computational complexity can be reduced and depends only on the number of inducing points. The technique is from Section 3.4.3 of [39] and is well explained in the documents of GPflows*. We provide the full details below for the sake of completeness.

Using the Woodbury identity, we get

$$\begin{aligned} (\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} &= (\gamma I + K(\boldsymbol{\psi}, \boldsymbol{\phi})(K(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi}))^{-1} \\ &= \gamma^{-1}I - \gamma^{-2}K(\boldsymbol{\psi}, \boldsymbol{\phi})(K(\boldsymbol{\phi}, \boldsymbol{\phi}) + \gamma^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi})K(\boldsymbol{\psi}, \boldsymbol{\phi}))^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi}). \end{aligned} \quad (2.7)$$

The equation (2.7) is not numerically stable when γ is small. To get a better conditioned matrix, we perform the Cholesky decomposition on $K(\boldsymbol{\phi}, \boldsymbol{\phi})$ and get $K(\boldsymbol{\phi}, \boldsymbol{\phi}) = LL^T$ (see Remark 2.1 below). Then, we consider

$$\begin{aligned} (\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} &= \gamma^{-1}I - \gamma^{-2}K(\boldsymbol{\psi}, \boldsymbol{\phi})(K(\boldsymbol{\phi}, \boldsymbol{\phi}) + \gamma^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi})K(\boldsymbol{\psi}, \boldsymbol{\phi}))^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi}) \\ &= \gamma^{-1}I - \gamma^{-2}K(\boldsymbol{\psi}, \boldsymbol{\phi})L^{-T}L^T(K(\boldsymbol{\phi}, \boldsymbol{\phi}) + \gamma^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi})K(\boldsymbol{\psi}, \boldsymbol{\phi}))^{-1}LL^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi}) \\ &= \gamma^{-1}I - \gamma^{-2}K(\boldsymbol{\psi}, \boldsymbol{\phi})L^{-T}(L^{-1}(K(\boldsymbol{\phi}, \boldsymbol{\phi}) + \gamma^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi})K(\boldsymbol{\psi}, \boldsymbol{\phi}))L^{-T})^{-1}L^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi}) \\ &= \gamma^{-1}I - \gamma^{-2}K(\boldsymbol{\psi}, \boldsymbol{\phi})L^{-T}(I + \gamma^{-1}L^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi})K(\boldsymbol{\psi}, \boldsymbol{\phi})L^{-T})^{-1}L^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi}). \end{aligned} \quad (2.8)$$

Let $A := \gamma^{-1/2}L^{-1}K(\boldsymbol{\phi}, \boldsymbol{\psi})$. Then, (2.8) gives

$$(\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} = \gamma^{-1}I - \gamma^{-1}A^T(I + AA^T)^{-1}A. \quad (2.9)$$

The identity (2.9) is the cornerstone for the SGP method to reduce the computational complexity. We note that the dimension of the matrix needed to be inverted at the right-hand side of (2.9) depends only on the length of $\boldsymbol{\phi}$, which is proportional to the number of inducing points. Hence, if $M \ll N$, using (2.9), we are able to save computational time and memory storage. The numerical results in Section 4 show that a small number of inducing points is sufficient for the SGP approach to achieve comparable accuracy to the GP method in [4].

*https://gpflow.github.io/GPflow/2.4.0/notebooks/theory/SGPR_notes.html

Remark 2.1. In general, $K(\phi, \phi)$ is ill-conditioned. Thus, to deal with the inverse of $K(\phi, \phi)$, we perform the Cholesky decomposition on $K(\phi, \phi) + \eta\mathcal{R}$, where \mathcal{R} is a block diagonal nugget built using the approach in [4] and $\eta > 0$ is a small regularization parameter.

2.3. A Probabilistic Perspective of the SGP Method. Here, we give a probabilistic interpretation for the SGP method. First, we put a Gaussian prior on the solution u^* to (2.1), i.e., let $u^* \sim \mathcal{N}(0, K)$, where K is the kernel associated with the RKHS \mathcal{U} in Subsection 2.1. Let $\{\mathbf{x}_i\}_{i=1}^N$ be the set of sample points in Subsection 2.1. Define

$$\mathbf{u} = (u^*(\mathbf{x}_1), \dots, u^*(\mathbf{x}_N), (\partial_{x_1} + \partial_{x_2})u^*(\mathbf{x}_1), \dots, (\partial_{x_1} + \partial_{x_2})u^*(\mathbf{x}_{N_\Omega}), \Delta u^*(\mathbf{x}_1), \dots, \Delta u^*(\mathbf{x}_{N_\Omega})).$$

Then, \mathbf{u} is a multivariate Gaussian random variable following the distribution $\mathcal{N}(0, K(\boldsymbol{\psi}, \boldsymbol{\psi}))$. Let \mathbf{z} be a noisy observation of \mathbf{u} , i.e.,

$$\mathbf{z} = \mathbf{u} + \boldsymbol{\epsilon},$$

where $\boldsymbol{\epsilon}$ is a noise following the distribution $\mathcal{N}(0, \gamma I)$, which is independent with u^* . Thus, $\mathbf{z} \sim \mathcal{N}(0, K(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I)$ and the probability density of \mathbf{z} is

$$p(\mathbf{z}) = \frac{1}{\sqrt{(2\pi)^{2(N_\Omega+N)} \det(K(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I)}} e^{-\frac{1}{2}\mathbf{z}^T(K(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I)^{-1}\mathbf{z}}.$$

Then, solving (2.1) by the GP method in [4] is equivalent to finding \mathbf{z} with the maximum likelihood while satisfying the PDE system. More precisely, the GP method seeks to find \mathbf{z} solving

$$\begin{cases} \max_{\mathbf{z}} \ln p(\mathbf{z}) \\ \text{s.t. } z_j^{(3)} = z_j^{(1)} z_j^{(2)} + f(\mathbf{x}_j), \forall j = 1, \dots, N_\Omega, \\ z_j^{(1)} = g(\mathbf{x}_j), \forall j, N_\Omega + 1 \leq j \leq N. \end{cases} \quad (2.10)$$

In our SGP method, we replace the log-likelihood in (2.10) with its lower bound. The derivation follows [37]. More precisely, let $\{\widehat{\mathbf{x}}_i\}_{i=1}^M$ be the inducing points in Subsection 2.1 and define $\mathbf{v}^{(1)} = (u(\widehat{\mathbf{x}}_1), \dots, u(\widehat{\mathbf{x}}_M))$, $\mathbf{v}^{(2)} = ((\partial_{x_1} + \partial_{x_2})u(\widehat{\mathbf{x}}_1), \dots, (\partial_{x_1} + \partial_{x_2})u(\widehat{\mathbf{x}}_{M_\Omega}))$, and $\mathbf{v}^{(3)} = (\Delta u(\widehat{\mathbf{x}}_1), \dots, \Delta u(\widehat{\mathbf{x}}_{M_\Omega}))$. Let $\mathbf{v} = (\mathbf{v}^{(1)}, \mathbf{v}^{(2)}, \mathbf{v}^{(3)})$. Then, \mathbf{v} and \mathbf{z} are joint Gaussian random variables. Let $q(\mathbf{v})$ be the variational distribution of \mathbf{v} and let $q(\mathbf{u}, \mathbf{v}) = p(\mathbf{u}|\mathbf{v})q(\mathbf{v})$ be the probability density function of the variational distribution of (\mathbf{u}, \mathbf{v}) . Then, by Jensen's inequality, we have

$$\begin{aligned} \ln p(\mathbf{z}) &= \ln \int \int p(\mathbf{z}, \mathbf{u}, \mathbf{v}) \, d\mathbf{u} \, d\mathbf{v} = \ln \int \int q(\mathbf{u}, \mathbf{v}) \frac{p(\mathbf{z}, \mathbf{u}, \mathbf{v})}{q(\mathbf{u}, \mathbf{v})} \, d\mathbf{u} \, d\mathbf{v} \\ &\geq \int \int q(\mathbf{u}, \mathbf{v}) \ln \frac{p(\mathbf{z}, \mathbf{u}, \mathbf{v})}{q(\mathbf{u}, \mathbf{v})} \, d\mathbf{u} \, d\mathbf{v} = \int \int p(\mathbf{u}|\mathbf{v})q(\mathbf{v}) \ln \frac{p(\mathbf{z}|\mathbf{u})p(\mathbf{v})}{q(\mathbf{v})} \, d\mathbf{u} \, d\mathbf{v} \\ &=: \mathcal{F}(\mathbf{z}, q). \end{aligned} \quad (2.11)$$

Thus, $\mathcal{F}(\mathbf{z}, q)$ provides a lower bound for $\ln p(\mathbf{z})$. To get an optimal lower bound, we take the derivative of \mathcal{F} w.r.t. q and get

$$\frac{\delta \mathcal{F}(\mathbf{z}, q(\mathbf{v}))}{\delta q(\mathbf{v})} = \int p(\mathbf{u}|\mathbf{v}) \left(\ln \frac{p(\mathbf{z}|\mathbf{u})p(\mathbf{v})}{q(\mathbf{v})} - 1 \right) \, d\mathbf{u}.$$

Letting $\delta \mathcal{F}(q(\mathbf{v}))/\delta q(\mathbf{v}) = 0$, we get the optimal variational distribution

$$q^*(\mathbf{v}) = \frac{p(\mathbf{v})}{Z} \exp \left(\int p(\mathbf{u}|\mathbf{v}) \ln p(\mathbf{z}|\mathbf{u}) \, d\mathbf{u} \right), \quad (2.12)$$

where Z is a real number guaranteeing the unit mass of q^* . Then, plugging (2.12) into (2.11), we get

$$\mathcal{F}(\mathbf{z}, q) = -\frac{N + N_\Omega}{2} \ln 2\pi - \frac{1}{2} \ln \det(Q(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I) - \mathbf{z}^T (Q(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I)^{-1} \mathbf{z} - \frac{1}{2\gamma} \text{Tr}(K(\boldsymbol{\psi}, \boldsymbol{\psi}) - Q(\boldsymbol{\psi}, \boldsymbol{\psi})).$$

Hence, if we replace $\ln p(\mathbf{z})$ by $\mathcal{F}(\mathbf{z}, q^*)$ in (2.10), we note that (2.5) is equivalent to

$$\begin{cases} \max_{\mathbf{z}} \mathcal{F}(\mathbf{z}, q^*) \\ \text{s.t. } z_j^{(3)} = z_j^{(1)} z_j^{(2)} + f(\mathbf{x}_j), \forall j = 1, \dots, N_\Omega, \\ z_j^{(1)} = g(\mathbf{x}_j), \forall j = N_\Omega + 1, \dots, N. \end{cases} \quad (2.13)$$

Remark 2.2. The probabilistic interpretation of the SGP method above suggests a way for hyperparameter learning. Let $\boldsymbol{\theta}$ be the parameter vector consisting of the positions of the inducing points and the kernel's parameters. Write $q_{\boldsymbol{\theta}}^*$ to emphasize the dependence of q^* in (2.12) on $\boldsymbol{\theta}$. Then, the SGP method with hyperparameter learning for solving (2.1) considers the maximization problem

$$\begin{cases} \max_{\mathbf{z}, \boldsymbol{\theta}} \mathcal{F}(\mathbf{z}, q_{\boldsymbol{\theta}}^*) \\ \text{s.t. } z_j^{(3)} = z_j^{(1)} z_j^{(2)} + f(\mathbf{x}_j), \forall j = 1, \dots, N_\Omega, \\ z_j^{(1)} = g(\mathbf{x}_j), \forall j = N_\Omega + 1, \dots, N. \end{cases} \quad (2.14)$$

However, the objective function in (2.14) is highly nonlinear. Thus, the computation of (2.14) is very costly. In the numerical experiments of Section 4, we fix the positions of inducing points. Then, we use the Gauss–Newton method to solve (2.14) for different values of the parameters of the kernel on the uniform grid in the domain of parameters and accept the result that achieves the largest value of the objective function in (2.14). If we choose Gaussian kernels for solving PDEs, which contain only one or two lengthscale parameters, the above grid search method for choosing hyperparameters is efficient. When the initial search interval is appropriately chosen, the combination of the maximum likelihood approach and grid search typically produces favorable outcomes. The optimal choice of hyperparameters depends on the underlying PDE's regularity, as a larger lengthscale tends to result in a smoother solution, while a smaller lengthscale tends to favor solutions with lower regularity. Given our understanding of the regularity of PDEs in our experiments, we search a lengthscale within the range of 0.01 to 1 using a step size of 0.01. We leave the full optimization over \mathbf{z} and $\boldsymbol{\theta}$ in (2.14) to future work.

Remark 2.3. The log-determinant term in \mathcal{F} will not introduce extra computational costs if the Cholesky decomposition of $I + AA^T$ in (2.9) is computed beforehand. To see this, we first recall the Weinstein–Aronszajn identity, which states that for n -by- m matrices U and V , and for an invertible m -by- m matrix W , we have

$$\det(I + U W V^T) = \det(W^{-1} + V^T U) \det(W),$$

where I is the identity matrix. In the following presentation of this remark, for ease of presentation, we denote by n and by m the dimensions of $Q(\boldsymbol{\psi}, \boldsymbol{\psi})$ and AA^T , separately, where A appears in (2.9). Thus, using (2.9) and the Weinstein–Aronszajn identity, we have

$$\begin{aligned} \det((\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}) &= \det(\gamma^{-1} I - \gamma^{-1} A^T (I + AA^T)^{-1} A) \\ &= \gamma^{-n} \det(-(I + AA^T) + AA^T) \det(-(I + AA^T)) \\ &= \gamma^{-n} \det(I + AA^T). \end{aligned} \quad (2.15)$$

Hence, if the Cholesky decomposition of $I + AA^T$ is computed in advance, i.e. $I + AA^T = JJ^T$, we conclude from (2.15) that

$$\begin{aligned} \log \det(\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi})) &= -\log \det((\gamma I + Q(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}) \\ &= n \log \gamma - \log \det(JJ^T) = n \log \gamma - 2 \log \det(J) \\ &= n \log \gamma - 2 \sum_{i=1}^m \log J_{ii}. \end{aligned} \quad (2.16)$$

Therefore, (2.16) implies that the computational cost of the log-determinant in \mathcal{F} is not high once we get the Cholesky decomposition of $I + AA^T$, which is required for solving (2.5) in our current SGP method.

3. A GENERAL FRAMEWORK FOR THE SGP METHOD

In this section, we provide a framework of the SGP method for solving a general PDE. Then, we give the existence argument for a minimizer of the optimal recovery problem used to approximate the solution of the PDE. Finally, we analyze the approximation errors.

3.1. A RKHS Generated by Linear Operators. In this subsection, we show the construction of a RKHS generated by linear operators based on the abstract theory of RKHSs (see [24]). The RKHS forms a space where we seek approximations for solutions to general PDEs.

Let \mathcal{U} be a Banach space, let \mathcal{U}^* be its dual, and let $[\cdot, \cdot]$ be their duality pairing. We assume further that there exists a covariance operator $\mathcal{K} : \mathcal{U}^* \mapsto \mathcal{U}$ that is linear, bijective, symmetric ($[\mathcal{K}\phi, \varphi] = [\mathcal{K}\varphi, \phi]$), and positive ($[\mathcal{K}\phi, \phi] > 0$ for $\phi \neq 0$), and suppose that the norm of \mathcal{U} is given by

$$\|u\|_{\mathcal{U}} = \sqrt{[\mathcal{K}^{-1}u, u]}, \forall u \in \mathcal{U}.$$

Meanwhile, the inner products of \mathcal{U} and \mathcal{U}^* are given by

$$\begin{aligned} \langle u, v \rangle &:= [\mathcal{K}^{-1}u, v], \forall u, v \in \mathcal{U}, \\ \langle \phi, \varphi \rangle_* &:= [\phi, \mathcal{K}\varphi], \forall \phi, \varphi \in \mathcal{U}^*. \end{aligned}$$

Then, \mathcal{U} coincides with the RKHS space of the kernel K defined by

$$K(\mathbf{x}, \mathbf{x}') = [\delta_{\mathbf{x}}, \mathcal{K}\delta_{\mathbf{x}'}], \forall \mathbf{x}, \mathbf{x}' \in \bar{\Omega}, \quad (3.1)$$

where $\delta_{\mathbf{x}}$ is the Dirac delta function centered at \mathbf{x} . Let $\xi \sim \mathcal{N}(0, \mathcal{K})$ be the *canonical* GP on \mathcal{U} [24, Chapter 17.6] and let $[\phi, \xi]$ be the image of $\phi \in \mathcal{U}^*$ under ξ . Then, we have

$$E[\phi, \xi] = 0 \text{ and } E[\phi, \xi][\varphi, \xi] = [\phi, \mathcal{K}\varphi], \forall \phi, \varphi \in \mathcal{U}^*.$$

The GP method proposed in [4] approximates a solution of a nonlinear PDE by a MAP point for the GP $\xi \sim \mathcal{N}(0, \mathcal{K})$ conditioned on PDE constraints at the collocation points. To reduce the computation time, we consider a Gaussian prior with a different covariance operator generated by selected linear operators.

Let $\phi := \{\phi_i\}_{i=1}^R$ be the collection of R non-trivial elements of \mathcal{U}^* . For the *canonical* GP $\xi \sim \mathcal{N}(0, \mathcal{K})$, $[\phi, \xi]$ is a \mathbb{R}^R -valued Gaussian vector and $[\phi, \xi] \sim \mathcal{N}(0, \Theta)$, where

$$\Theta \in \mathbb{R}^{R \times R}, \Theta_{i,n} = [\phi_i, \mathcal{K}\phi_n], \forall 1 \leq i, n \leq R. \quad (3.2)$$

Suppose that Θ is invertible. We define the induced covariance operator \mathcal{Q} as

$$[\varphi, \mathcal{Q}\varphi] = \mathcal{K}(\varphi, \phi)\Theta^{-1}\mathcal{K}(\phi, \varphi), \forall \varphi \in \mathcal{U}^*. \quad (3.3)$$

The next lemma guarantees that \mathcal{Q} is a covariance operator and associates with a RKHS.

Lemma 3.1. *Let $\phi := \{\phi_i\}_{i=1}^R$ be the collection of R linearly independent non-trivial elements of \mathcal{U}^* . Then, the operator \mathcal{Q} defined in (3.3) is linear, symmetric, and nonnegative. Hence, there exists a unique RKHS $\mathcal{U}_{\mathcal{Q}}$ associated with the kernel Q defined by*

$$Q(\mathbf{x}, \mathbf{x}') = [\delta_{\mathbf{x}}, \mathcal{Q}\delta_{\mathbf{x}'}], \forall \mathbf{x}, \mathbf{x}' \in \bar{\Omega}. \quad (3.4)$$

Furthermore, $\mathcal{U}_{\mathcal{Q}} \subset \mathcal{U}$ and

$$\|u\|_{\mathcal{U}_{\mathcal{Q}}} = \|u\|_{\mathcal{U}}, \forall u \in \mathcal{U}_{\mathcal{Q}}. \quad (3.5)$$

Proof. It is easy to see that \mathcal{Q} is linear and symmetric. Since \mathcal{K} is positive, Θ is positive definite. Then, $[\varphi, \mathcal{Q}\varphi] = \mathcal{K}(\varphi, \phi)\Theta^{-1}\mathcal{K}(\phi, \varphi) \geq 0, \forall \varphi \in \mathcal{U}^*$. The equality holds if and only if φ is perpendicular to ϕ w.r.t. \mathcal{K} . Thus, \mathcal{Q} is nonnegative. Hence, the kernel Q in (3.4) is symmetric and positive semi-definite. According to the Moore–Aronszajn theorem, there exists a unique RKHS \mathcal{U}_Q associated with the kernel Q . Next, we show the relation between \mathcal{U}_Q and \mathcal{U} .

Since $Q(\mathbf{x}, \cdot) \in \mathcal{U}$ for any $\mathbf{x} \in \bar{\Omega}$, the linear span of $\{Q(\mathbf{x}, \cdot), \mathbf{x} \in \bar{\Omega}\}$, denoted by \mathcal{U}_{Q_0} , is a subset of \mathcal{U} . By the Moore–Aronszajn theorem, \mathcal{U}_Q is the completion of \mathcal{U}_{Q_0} . Thus, $\mathcal{U}_Q \subset \mathcal{U}$. For any $u \in \mathcal{U}_{Q_0}$, there exist sequences $\{\mathbf{x}_i\}_{i=1}^m$ and $\{\alpha_i\}_{i=1}^m$, $m \in \mathbb{N}$, such that

$$\|u\|_{\mathcal{U}_Q}^2 = \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j Q(\mathbf{x}_i, \mathbf{x}_j) = \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j \mathcal{K}(\delta_{\mathbf{x}_i}, \phi) \Theta^{-1} \mathcal{K}(\phi, \delta_{\mathbf{x}_j}) = \|u\|_{\mathcal{U}}^2,$$

where the last two equalities follow by (3.3) and (3.4). Thus, we conclude (3.5). \square

3.2. The SGP method. In this subsection, we present a general framework of the SGP method to solve nonlinear PDEs. Let $\Omega \subset \mathbb{R}^d$ be a bounded open domain with the boundary $\partial\Omega$ for $d \geq 1$. Let $L_1, \dots, L_{D_b} \in \mathcal{L}(\mathcal{U}; C(\partial\Omega))$ and $L_{D_b+1}, \dots, L_D \in \mathcal{L}(\mathcal{U}; C(\Omega))$ be bounded linear operators for $1 \leq D_b \leq D$. We seek to find a function u^* solving the nonlinear PDE

$$\begin{cases} P(L_{D_b+1}(u^*)(\mathbf{x}), \dots, L_D(u^*)(\mathbf{x})) = f(\mathbf{x}), \forall \mathbf{x} \in \Omega, \\ B(L_1(u^*)(\mathbf{x}), \dots, L_{D_b}(u^*)(\mathbf{x})) = g(\mathbf{x}), \forall \mathbf{x} \in \partial\Omega, \end{cases} \quad (3.6)$$

where P, B represent nonlinear operators, and f, g are given data. Throughout this section, we assume that (3.6) admits a unique strong solution in a quadratic Banach space \mathcal{U} associated with the covariance operator \mathcal{K} , which means that u^* has enough regularity for the linear operators to be well defined pointwisely in (3.6).

We propose to approximate the solution to (3.6) in a RKHS generated by linear operators associated with inducing points. To do that, we take M inducing points $\{\hat{\mathbf{x}}_j\}_{j=1}^M$ in $\bar{\Omega}$ such that $\{\hat{\mathbf{x}}_j\}_{j=1}^{M_\Omega} \subset \Omega$ and $\{\hat{\mathbf{x}}_j\}_{j=M_\Omega+1}^M \subset \partial\Omega$. Then, we define the functionals $\phi_j^{(i)} \in \mathcal{U}^*$ as

$$\phi_j^{(i)} := \delta_{\hat{\mathbf{x}}_j} \circ L_i, \text{ and } \begin{cases} M_\Omega + 1 \leq j \leq M, & \text{if } 1 \leq i \leq D_b, \\ 1 \leq j \leq M_\Omega, & \text{if } D_b + 1 \leq i \leq D. \end{cases}$$

Let $\phi^{(i)}$ be the vector concatenating the functionals $\phi_j^{(i)}$ for fixed i and define

$$\phi := (\phi^{(1)}, \dots, \phi^{(D)}) \in (\mathcal{U}^*)^{\otimes R}, \text{ where } R = (M - M_\Omega)D_b + M_\Omega(D - D_b). \quad (3.7)$$

The next corollary gives a RKHS generated by ϕ .

Corollary 3.2. *Let ϕ be as in (3.7). Assume that the elements of ϕ are linearly independent in \mathcal{U}^* . Define Θ as in (3.2). Then, there exists a RKHS \mathcal{U}_Q associated with the kernel Q defined by*

$$Q(\mathbf{x}, \mathbf{x}') = K(\mathbf{x}, \phi) \Theta^{-1} K(\phi, \mathbf{x}'), \forall \mathbf{x}, \mathbf{x}' \in \bar{\Omega},$$

such that $\mathcal{U}_Q \subset \mathcal{U}$. Denote by $\|\cdot\|_{\mathcal{U}_Q}$ the norm of \mathcal{U}_Q . Then,

$$\|u\|_{\mathcal{U}_Q} = \|u\|_{\mathcal{U}}, \forall u \in \mathcal{U}_Q.$$

Proof. The claim is a direct result of Lemma 3.1. \square

Corollary 3.2 implies that \mathcal{U}_Q is a ‘‘condensed’’ subspace of \mathcal{U} and can be treated as a candidate solution for Problem 1. Next, we propose to approximate the solution u^* of (3.6) in the space \mathcal{U}_Q . More precisely, we take N samples $\bar{\mathbf{x}} := \{\mathbf{x}_i\}_{i=1}^N$ in $\bar{\Omega}$ such that $\{\mathbf{x}_i\}_{i=1}^{N_\Omega} \subset \Omega$ and $\{\mathbf{x}_i\}_{i=N_\Omega+1}^N \subset \partial\Omega$. Given a regularization

parameter $\gamma > 0$, we approximate u^* by the minimizer of the following optimization problem

$$\begin{cases} \min_{\mathbf{z} \in \mathbb{R}^{\bar{R}}, u \in \mathcal{U}_Q} & \sum_{i=1}^{D_b} \sum_{j=N_\Omega+1}^N |L_i(u)(\mathbf{x}_j) - z_j^{(i)}|^2 + \sum_{i=D_b+1}^D \sum_{j=1}^{N_\Omega} |L_i(u)(\mathbf{x}_j) - z_j^{(i)}|^2 + \gamma \|u\|_{\mathcal{U}_Q}^2 \\ \text{s. t.} & P(z_j^{(D_b+1)}, \dots, z_j^{(D)}) = f(\mathbf{x}_j), \text{ for } j = 1, \dots, N_\Omega, \\ & B(z_j^{(1)}, \dots, z_j^{(D_b)}) = g(\mathbf{x}_j), \text{ for } j = N_\Omega + 1, \dots, N, \end{cases} \quad (3.8)$$

where

$$\bar{R} = (N - N_\Omega)D_b + N_\Omega(D - D_b) \quad (3.9)$$

and

$$\mathbf{z} := (z_{N_\Omega+1}^{(1)}, \dots, z_N^{(1)}, \dots, z_{N_\Omega+1}^{(D_b)}, \dots, z_N^{(D_b)}, z_1^{(D_b+1)}, \dots, z_{N_\Omega}^{(D_b+1)}, \dots, z_1^{(D)}, \dots, z_{N_\Omega}^{(D)}).$$

For ease of presentation, we rewrite (3.8) into a compact form. To do that, we define the functionals $\psi_j^{(i)} \in \mathcal{U}^*$ as

$$\psi_j^{(i)} := \delta_{\mathbf{x}_j} \circ L_i, \text{ and } \begin{cases} N_\Omega + 1 \leq j \leq N, \text{ if } 1 \leq i \leq D_b, \\ 1 \leq j \leq N_\Omega, \text{ if } D_b + 1 \leq i \leq D. \end{cases}$$

Meanwhile, let $\boldsymbol{\psi}^{(i)}$ be the vector consisting of $\psi_j^{(i)}$ for fixed i and define

$$\boldsymbol{\psi} := (\boldsymbol{\psi}^{(1)}, \dots, \boldsymbol{\psi}^{(D)}) \in (\mathcal{U}^*)^{\otimes \bar{R}}. \quad (3.10)$$

Next, we define the data vector $\mathbf{y} \in \mathbb{R}^N$ by

$$y_i = \begin{cases} f(\mathbf{x}_i), \text{ if } i \in \{1, \dots, N_\Omega\}, \\ g(\mathbf{x}_i), \text{ if } i \in \{N_\Omega + 1, \dots, N\}. \end{cases}$$

Furthermore, we define the nonlinear map

$$(F(\mathbf{z}))_j := \begin{cases} P(z_j^{(D_b+1)}, \dots, z_j^{(D)}), \text{ for } j = 1, \dots, N_\Omega, \\ B(z_j^{(1)}, \dots, z_j^{(D_b)}), \text{ for } j = N_\Omega + 1, \dots, N. \end{cases}$$

Thus, we reformulate (3.8) as

$$\begin{cases} \min_{\mathbf{z} \in \mathbb{R}^{\bar{R}}, u \in \mathcal{U}_Q} & |[\boldsymbol{\psi}, u] - \mathbf{z}|^2 + \gamma \|u\|_{\mathcal{U}_Q}^2 \\ \text{s. t.} & F(\mathbf{z}) = \mathbf{y}. \end{cases} \quad (3.11)$$

We observe that the minimization on u in (3.11) is over the infinite dimensional function space \mathcal{U}_Q . Fortunately, we can derive a representer formula for u and transform (3.11) into a finite dimensional minimization problem. However, we are not able to use the representer theorem in [24, Section 17.8] directly, since the inducing kernel Q in Corollary 3.2 is not positive definite. Instead, the arguments are based on the framework in [33]. The basic idea is to define a sampling operator $S_{\bar{\mathbf{x}}}$ (explained in more details below) such that $S_{\bar{\mathbf{x}}}u = [\boldsymbol{\psi}, u], \forall u \in \mathcal{U}_Q$. Then, the objective function in (3.11) is reformulated as $\langle u, S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}}u - S_{\bar{\mathbf{x}}}^* \mathbf{z} \rangle + \gamma \|u\|_{\mathcal{U}_Q}^2$, where $S_{\bar{\mathbf{x}}}^*$ is the adjoint operator of $S_{\bar{\mathbf{x}}}$ and will be defined later. Then, by taking a functional derivative of the objective function and setting the derivative to be zero, we obtain $u = (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} S_{\bar{\mathbf{x}}}^* \mathbf{z}$ for fixed \mathbf{z} . Then, we plug $u = (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} S_{\bar{\mathbf{x}}}^* \mathbf{z}$ back into (3.11) and get a finite dimensional minimization problem over \mathbf{z} . Next, we explain rigorously the definitions and the properties of $S_{\bar{\mathbf{x}}}$ and its adjoint in more details.

Let $\ell^2(\bar{\mathbf{x}})$ be the set of sequences $\mathbf{a} = (a_{\mathbf{x}})_{\mathbf{x} \in \bar{\mathbf{x}}}$ with $\langle \mathbf{a}, \mathbf{b} \rangle_{\ell^2(\bar{\mathbf{x}})} = \sum_{\mathbf{x} \in \bar{\mathbf{x}}} a_{\mathbf{x}} b_{\mathbf{x}}$ defining the inner product. Define the sampling operator $S_{\bar{\mathbf{x}}} : \mathcal{U}_Q \mapsto \ell^2(\bar{\mathbf{x}})$ by

$$S_{\bar{\mathbf{x}}}(u) = [\boldsymbol{\psi}, u], \forall u \in \mathcal{U}_Q. \quad (3.12)$$

One can view $S_{\bar{x}}$ as a generalization of the sampling operator in [33], which only evaluates a function at the sample points. Denote by $\langle \cdot, \cdot \rangle_{\mathcal{U}_Q}$ the inner product of \mathcal{U}_Q and by $S_{\bar{x}}^*$ the adjoint of $S_{\bar{x}}$. Then, for each $c \in \ell^2(\bar{x})$ and $u \in \mathcal{U}_Q$, we have

$$\langle S_{\bar{x}}^* c, u \rangle_{\mathcal{U}_Q} = \langle S_{\bar{x}} u, c \rangle_{\ell^2(\bar{x})} = c^T [\boldsymbol{\psi}, u] = \langle c^T \mathcal{Q} \boldsymbol{\psi}, u \rangle_{\mathcal{U}_Q}.$$

Thus, we obtain

$$S_{\bar{x}}^* c = c^T \mathcal{Q} \boldsymbol{\psi}, \forall c \in \ell^2(\bar{x}). \quad (3.13)$$

The next lemma gives some basic properties of $S_{\bar{x}}$ and $S_{\bar{x}}^*$. For a little abuse of notations, in the rest of this section, we denote by I the identity map or the identity matrix, whose meaning is easy to recognize from the context.

Lemma 3.3. *Let $S_{\bar{x}}$ and $S_{\bar{x}}^*$ be as in (3.12) and (3.13). Then, for any $\gamma > 0$, $S_{\bar{x}}^* S_{\bar{x}} + \gamma I$ and $S_{\bar{x}} S_{\bar{x}}^* + \gamma I$ are bijective. Meanwhile,*

$$S_{\bar{x}}^* (S_{\bar{x}} S_{\bar{x}}^* + \gamma I)^{-1} = (S_{\bar{x}}^* S_{\bar{x}} + \gamma I)^{-1} S_{\bar{x}}^*. \quad (3.14)$$

For any $c \in \ell^2(\bar{x})$, we have

$$(S_{\bar{x}}^* S_{\bar{x}} + \gamma I)^{-1} S_{\bar{x}}^* c = (\mathcal{Q} \boldsymbol{\psi})^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} c. \quad (3.15)$$

Furthermore,

$$I - S_{\bar{x}} (S_{\bar{x}}^* S_{\bar{x}} + \gamma I)^{-1} S_{\bar{x}}^* = \gamma (\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I)^{-1}. \quad (3.16)$$

We refer the reader to Appendix A for the proof of Lemma 3.3. We note that (3.15) and (3.16) establishes the relations between the operations of sampling operators $S_{\bar{x}}$ and the covariance matrix $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})$. With the properties of the sampling operator at hand, we are able to prove the existence of minimizers to (3.11). The following theorem is similar to Proposition 2.3 in [4]. A key difference is that the arguments of Proposition 2.3 in [4] rely on the representer theorem [24, Section 17.8], which requires the covariance operator to be positive definite, but our proof here holds for general positive semi-definite covariance operators.

Theorem 3.4. *Let $S_{\bar{x}}$ and $S_{\bar{x}}^*$ be defined in (3.12) and (3.13). Let \bar{R} be as in (3.9). Given $\gamma > 0$, the system (3.11) admits a solution*

$$u^\dagger = (S_{\bar{x}}^* S_{\bar{x}} + \gamma I)^{-1} S_{\bar{x}}^* \mathbf{z}^\dagger, \quad (3.17)$$

where \mathbf{z}^\dagger is a solution to

$$\begin{cases} \min_{\mathbf{z} \in \mathbb{R}^{\bar{R}}} & \mathbf{z}^T (\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I)^{-1} \mathbf{z} \\ \text{s. t.} & F(\mathbf{z}) = \mathbf{y}. \end{cases} \quad (3.18)$$

Proof. In (3.11), given $\mathbf{z} \in \mathbb{R}^{\bar{R}}$, we first consider the minimization problem over u

$$\min_{u \in \mathcal{U}_Q} \|\boldsymbol{\psi}, u\| - \mathbf{z} + \gamma \|u\|_{\mathcal{U}_Q}^2. \quad (3.19)$$

By the definition of $S_{\bar{x}}$ in (3.12), we get

$$\begin{aligned} \|\boldsymbol{\psi}, u\| - \mathbf{z} + \gamma \|u\|_{\mathcal{U}_Q}^2 &= \langle S_{\bar{x}} u - \mathbf{z}, S_{\bar{x}} u - \mathbf{z} \rangle_{\ell^2(\bar{x})} + \gamma \langle u, u \rangle_{\mathcal{U}_Q} \\ &= \langle S_{\bar{x}}^* S_{\bar{x}} u - 2S_{\bar{x}}^* \mathbf{z} + \gamma u, u \rangle_{\mathcal{U}_Q} + |\mathbf{z}|^2. \end{aligned} \quad (3.20)$$

Taking the function derivative w.r.t. u , we obtain from (3.20) that

$$S_{\bar{x}}^* S_{\bar{x}} u - S_{\bar{x}}^* \mathbf{z} + \gamma u = 0,$$

which implies that the minimizer u^\dagger of (3.19) satisfies

$$u^\dagger = (S_{\bar{x}}^* S_{\bar{x}} + \gamma I)^{-1} S_{\bar{x}}^* \mathbf{z}. \quad (3.21)$$

Thus, using (3.20) and (3.21), we have

$$\begin{aligned} \|[\boldsymbol{\psi}, u^\dagger] - \mathbf{z}\|^2 + \gamma \|u^\dagger\|_{\mathcal{U}_Q}^2 &= \langle (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} - 2(S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I) + \gamma I)(S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} S_{\bar{\mathbf{x}}}^* \mathbf{z}, u^\dagger \rangle_{\mathcal{U}_Q} + |\mathbf{z}|^2 \\ &= -\langle S_{\bar{\mathbf{x}}}^* \mathbf{z}, u^\dagger \rangle_{\mathcal{U}_Q} + |\mathbf{z}|^2 = -\langle \mathbf{z}, S_{\bar{\mathbf{x}}} u^\dagger \rangle_{\ell^2(\bar{\mathbf{x}})} + |\mathbf{z}|^2 \\ &= \langle (I - S_{\bar{\mathbf{x}}}(S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} S_{\bar{\mathbf{x}}}^*) \mathbf{z}, \mathbf{z} \rangle_{\ell^2(\bar{\mathbf{x}})} = \gamma \mathbf{z}^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathbf{z}, \end{aligned} \quad (3.22)$$

where the last equality results from (3.16). Thus, (3.22) implies that (3.11) is equivalent to (3.18). Using a similar compactness argument as in the proof of Theorem 1.1 in [4], one can show that (3.18) admits a minimizer \mathbf{z}^\dagger . Therefore, we conclude (3.17) by (3.21). \square

Remark 3.5. Theorem 3.4 implies that in order to find the solution u^* of the PDE system (3.6), we need to solve (3.18). To handle the constraints of (3.18), we use the methods of eliminating variables or relaxation described in Subsection 3.3 of [4] when necessary.

Remark 3.6. To compute the inversion of $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I$, we use the techniques introduced in Subsection 2.2 and Remark 2.1.

Next, we estimate the error between the approximation solution u^\dagger given in (3.17) and the strong solution u^* of (3.6). The proof appears in Appendix A.

Theorem 3.7. *Let u^* be the strong solution of (3.6). Given $\gamma > 0$, let $(u^\dagger, \mathbf{z}^\dagger)$ be a solution of (3.11), where u^\dagger is given in (3.17), and \mathbf{z}^\dagger solves (3.18). Then,*

$$\begin{aligned} \|u^\dagger - u^*\|_{\mathcal{U}} &\leq \inf_{\psi \in \text{span}(\boldsymbol{\psi})} \|u^* - \mathcal{K}\psi\|_{\mathcal{U}} + \gamma \sqrt{[\boldsymbol{\psi}, u^*]^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-3} [\boldsymbol{\psi}, u^*]} \\ &\quad + \sqrt{3} |\mathbf{z}^\dagger| \sqrt{\|(\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} - (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}\|} \\ &\quad + \sqrt{(\mathbf{z}^\dagger - [\boldsymbol{\psi}, u^*]^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (\mathbf{z}^\dagger - [\boldsymbol{\psi}, u^*]))}. \end{aligned} \quad (3.23)$$

The first term at the right-hand side (RHS) of (3.23) measures the ability of $\text{span}(\mathcal{K}\boldsymbol{\psi})$ to approximate u^* . Since $u^* \in \mathcal{U}$, we conclude that as the samples become dense enough, $\inf_{\psi \in \text{span}(\boldsymbol{\psi})} \|u^* - \mathcal{K}\psi\|_{\mathcal{U}}$ is small. The second term at the RHS of (3.23) measures the effect of the nugget term γI , which converges to 0 as $\gamma \rightarrow 0$ for a fixed number of samples. The magnitudes of the last two terms of (3.23) depend on how well $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})$ approximates $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$. Indeed, when $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) = \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$, the SGP method recovers the GP algorithm in [4]. Hence, $|\mathbf{z}^\dagger|$ is bounded in this case and the third term at the RHS of (3.23) equals 0. Besides, \mathbf{z}^\dagger converges to $[\boldsymbol{\psi}, u^*]$ as the number of samples, N , goes to infinity and $\gamma \rightarrow 0$ by Theorem 3.3 of [4]. Thus, when $\mathcal{Q} = \mathcal{K}$, as the samples gets dense and $\gamma \rightarrow 0$, $\mathbf{z}^\dagger \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} \mathbf{z}^\dagger$ and $\mathbf{z}^\dagger \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} [\boldsymbol{\psi}, u^*]$ converge to $\|u^*\|_{\mathcal{U}}^2$. Hence, the last term of (3.23) converges to 0.

3.3. The Nyström Approximation Error. As mentioned above, the upper bound in (3.23) decreases when $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})$ approximates $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$ well. In this subsection, we study the spectral norm of $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})$, denoted by $\|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})\|$.

We recall that $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) = \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1} \mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})$. Thus, we can view $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})$ as the Nyström approximation to $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$. Hence, we call $\|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})\|$ the Nyström approximation error in the SGP method. A typical Nyström approximation approximates a symmetric positive semi-definite matrix $G \in \mathbb{R}^{n \times n}$, $n \in \mathbb{N}$, by sampling m columns from G and does not assume that columns of G have close relations with each other. The authors in [7] show that for any m uniformly sampled columns, with a high probability, the approximation error is $O(n/\sqrt{m})$. Later, [14] improves the bound from $O(n/\sqrt{m})$ to $O(n/m^{1-\rho})$ for $\rho \in (0, 1/2)$ under the assumption that the eigenvalues of G have a big eigengap. In our setting, since $\boldsymbol{\psi}$ has linear operators corresponding to the same sample points, the columns of $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$ are highly correlated. Thus, we approximate $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$ using a small number of inducing points instead of using

columns from $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$. In general, instead of being taken from the samples, inducing points can be selected anywhere in the input space [8]. The error analysis has been thoroughly studied in [43], which indicates that the approximation error in the matrix Frobenious norm has an upper bound that is influenced by the quantization error defined in [43]. Since that quantization error is the objective function in the k-means clustering [10], the k-means algorithm is typically used for the initialization of inducing inputs. In our settings, we assume that the inducing points are sampled from the set of samples. We adapt the framework in [14] to our settings and show that the Nyström approximation error has an upper bound that depends only on the numbers of samples and inducing points. In other words, though the dimension of $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})$ is proportional to the product of the size of samples and the number of linear operators, the error bound is not influenced by the size of linear operators.

Our arguments are an adaptation to those in [14] to take into consideration of covariance matrices with values of differential operators. The strategy is to turn $\|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})\|$ into a functional approximation problem, which is easier to bound. Define the sets

$$\mathcal{H}_a = \text{span}\{\mathcal{K}\boldsymbol{\phi}\} \text{ and } \mathcal{H}_b = \{\boldsymbol{\beta}^T \mathcal{K}\boldsymbol{\psi}, |\boldsymbol{\beta}| \leq 1\}. \quad (3.24)$$

For $h \in \mathcal{H}_b$, we define

$$\mathcal{E}(h, \mathcal{H}_a) = \min_{v \in \mathcal{H}_a} \|v - h\|_{\mathcal{U}}^2.$$

Then, $\mathcal{E}(h, \mathcal{H}_a)$ is the minimum error for approximating a function $h \in \mathcal{H}_b$ by functions in \mathcal{H}_a . Meanwhile, define $\mathcal{E}(\mathcal{H}_a)$ as the worst error in approximating any function $h \in \mathcal{H}_b$ by functions in \mathcal{H}_a , i.e.

$$\mathcal{E}(\mathcal{H}_a) = \max_{h \in \mathcal{H}_b} \mathcal{E}(h, \mathcal{H}_a). \quad (3.25)$$

The next proposition shows the equivalence of $\|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})\|$ and $\mathcal{E}(\mathcal{H}_a)$.

Proposition 3.8. Let $\mathcal{E}(\mathcal{H}_a)$, $\boldsymbol{\psi}$, and $\boldsymbol{\phi}$ be as in (3.25), (3.10), and (3.7), respectively. Then,

$$\|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})\| = \mathcal{E}(\mathcal{H}_a). \quad (3.26)$$

Proof. Let $h \in \mathcal{H}_b$ and $v \in \mathcal{H}_a$. Then, there exists $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that $v = \boldsymbol{\alpha}^T \mathcal{K}\boldsymbol{\phi}$ and $h = \boldsymbol{\beta}^T \mathcal{K}\boldsymbol{\psi}$. We have

$$\begin{aligned} \mathcal{E}(h, \mathcal{H}_a) &= \min_{\boldsymbol{\alpha}} \boldsymbol{\alpha}^T \mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}) \boldsymbol{\alpha} - 2\boldsymbol{\alpha}^T \mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi}) \boldsymbol{\beta} + \boldsymbol{\beta}^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) \boldsymbol{\beta} \\ &= \boldsymbol{\beta}^T (\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})) \boldsymbol{\beta}. \end{aligned}$$

Hence, we obtain

$$\begin{aligned} \mathcal{E}(\mathcal{H}_a) &= \max_{h \in \mathcal{H}_b} \mathcal{E}(h, \mathcal{H}_a) \\ &= \max_{|\boldsymbol{\beta}| \leq 1} \boldsymbol{\beta}^T (\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})) \boldsymbol{\beta} \\ &= \|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})\|, \end{aligned}$$

which concludes (3.26). \square

Thus, Proposition 3.8 implies that estimating an upper bound for $\|\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\phi})(\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\phi}))^{-1}\mathcal{K}(\boldsymbol{\phi}, \boldsymbol{\psi})\|$ is equivalent to bounding $\mathcal{E}(\mathcal{H}_a)$. To proceed with the argument, we make the following assumptions about the inducing points and the kernel associated with the RKHS \mathcal{U} .

The first assumption assumes that the inducing points are uniformly sampled from the collection of sample points.

Assumption 1. Let $\bar{\boldsymbol{x}} := (\bar{\boldsymbol{x}}_1, \bar{\boldsymbol{x}}_2)$, where $\bar{\boldsymbol{x}}_1$ is the set of interior sample points in Ω and $\bar{\boldsymbol{x}}_2$ is the collection of boundary samples, and let $\hat{\boldsymbol{x}} := (\hat{\boldsymbol{x}}_1, \hat{\boldsymbol{x}}_2)$ be the set of inducing points. Assume that $\hat{\boldsymbol{x}}_1$ and $\hat{\boldsymbol{x}}_2$ are uniformly sampled from $\bar{\boldsymbol{x}}_1$ and $\bar{\boldsymbol{x}}_2$, separately.

The next assumption supposes that the covariance operator \mathcal{K} admits enough regularity.

Assumption 2. Let ψ be as in (3.10). Let $\psi^{\mathbf{x}} \subset (\delta_{\mathbf{x}} \circ L_1, \dots, \delta_{\mathbf{x}} \circ L_D)$ be the vector of linear operators in ψ that is associated with the sample point at $\mathbf{x} \in \bar{\mathbf{x}}$. Assume that there exists a constant $C > 0$ such that

$$\text{trace}(\mathcal{K}(\psi^{\mathbf{x}}, \psi^{\mathbf{x}})) \leq C.$$

Remark 3.9. We note that if the kernel K in (3.1) generated by \mathcal{K} is smooth enough and the domain $\bar{\Omega}$ is compact, Assumption 2 holds.

We provide an upper bound for $\mathcal{E}(\mathcal{H}_a)$ using arguments based on the eigenvalue decomposition. To do this, we define some notation. We denote $\psi := (\psi_{\Omega}, \psi_{\partial\Omega})$, where ψ_{Ω} represents linear operators for interior points, and $\psi_{\partial\Omega}$ contains linear operators for boundary samples. We recall that $\bar{\mathbf{x}} := (\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2)$ contains all sample points, with $\bar{\mathbf{x}}_1$ being the set of interior points in Ω and $\bar{\mathbf{x}}_2$ representing the collection of boundary points. We define N_{Ω} as the size of $\bar{\mathbf{x}}_1$ and $N_{\partial\Omega}$ as the length of $\bar{\mathbf{x}}_2$. We also define R_1 and R_2 as $R_1 = N_{\Omega}(D - D_b)$ and $R_2 = (N - N_{\Omega})D_b$, respectively, where D and D_b are given in (3.6). We use $\{\lambda_j\}_{j=1}^{R_1}$ and $\{\tau_j\}_{j=1}^{R_2}$ to represent the set of eigenvalues of $\mathcal{K}(\psi_{\Omega}, \psi_{\Omega})$ and $\mathcal{K}(\psi_{\partial\Omega}, \psi_{\partial\Omega})$, respectively, sorted in descending order. We denote $[i]$ as the set of all non-negative integers less than or equal to i . Using these notations, we establish an upper bound for $\mathcal{E}(\mathcal{H}_a)$ in the following theorem. The proof is presented in Appendix B.

Theorem 3.10. Suppose that Assumptions 1 and 2 hold. Let $\mathcal{E}(\mathcal{H}_a)$ be as in (3.25). With a probability at least $1 - \delta$, $\delta \in (0, 1)$, for any $\mathbf{r} := (r_1, r_2) \in [R_1] \times [R_2]$, there exists a constant C such that

$$\mathcal{E}(\mathcal{H}_a) \leq \max \left\{ \frac{CN_{\Omega}^2 \ln^2(2/\delta)}{\lambda_{r_1} M_{\Omega}}, \frac{CN_{\partial\Omega}^2 \ln^2(2/\delta)}{\tau_{r_2} M_{\partial\Omega}} \right\} + 4\lambda_{r_1+1} + 4\tau_{r_2+1}, \quad (3.27)$$

with the convention that $\lambda_{R_1+1} = 0$ and $\tau_{R_2+1} = 0$.

Proposition 3.8 and Theorem 3.10 imply that $\|\mathcal{K}(\psi, \psi) - \mathcal{K}(\psi, \phi)(\mathcal{K}(\phi, \phi))^{-1}\mathcal{K}(\phi, \psi)\|$ admits an upper bound that depends only on the size of sample points and is not influenced by the number of linear operators in the PDE.

4. NUMERICAL RESULTS

In this section, we study numerically how much information the SGP method preserves for different choices of numbers of inducing points. We demonstrate the efficacy of the SGP method by solving a nonlinear elliptic equation in Subsection 4.1, a mean-field game system in Subsection 4.2, Burger's equation in Subsection 4.3, and a parabolic equation in Subsection 4.4. For all the experiments, we take N samples in the domain in a way such that N_{Ω} samples are in the interior. Besides, we randomly choose M points from the samples and treat them as inducing points. To show the performance of the SGP method, we first plot samples, inducing points, loss histories of the Gauss–Newton iterations, and contours of pointwise solution errors for fixed values of N and M . Though we solve at a collection of uniform random samples, we compute the solution errors by reevaluating the numerical approximation function at a 60×60 equally spaced grid using (3.17) and by comparing the results with the values of reference solutions. Then, we record the L^{∞} errors of the GP method in [4] and the SGP algorithm for different values of N and M . Results are averaged over 10 realizations of random sample points.

Our implementation is based on the code[†] of [4], which leverages Python with the JAX package for automatic differentiation. Our experiments are conducted only on CPUs. Additional performance speedups can be obtained by considering accelerated hardware such as Graphics Processing Units (GPU).

[†]<https://github.com/yifanc96/NonlinearPDEs-GPsolver.git>

4.1. A Nonlinear Elliptic Equation. In this example, we reconsider the nonlinear elliptic equation in Subsection 2.1. More precisely, we consider (2.1) with $\Omega = (0, 3)^2$ and $g(\mathbf{x}) = 0$. We prescribe the solution u to be $\sin(\pi x_1) \sin(\pi x_2) + 4 \sin(4\pi x_1) \sin(4\pi x_2)$ for $\mathbf{x} := (x_1, x_2) \in \bar{\Omega}$ and compute f accordingly. We use the Gaussian kernel $K(\mathbf{x}, \mathbf{y}) = \exp(-\frac{|\mathbf{x}-\mathbf{y}|^2}{2\sigma^2})$ with the lengthscale parameter $\sigma = 0.2$. Given N and M , we set $N_\Omega = 0.75 \times N$ and $M_\Omega = 0.75 \times M$. In this example, we choose $\gamma = \eta = 10^{-12}$ (see Subsection 2.2 and Remark 2.1). The algorithm stops once the error between two successive steps is less than 10^{-5} .

Figure 2 shows the numerical results for $N = 9600$ and $M = 1200$. The uniform sample points used are plotted in Figure 2a, from which we randomly choose M including points shown in Figure 2b. We solve the unconstrained minimization problem in (2.6). The convergence history of the Gauss–Newton iteration in Figure 2c shows that the SGP method converges in 5 iterations. In Figure 2d, we plot the contour of pointwise errors between the numerical approximation and the true solution on a 60×60 equally spaced grid.

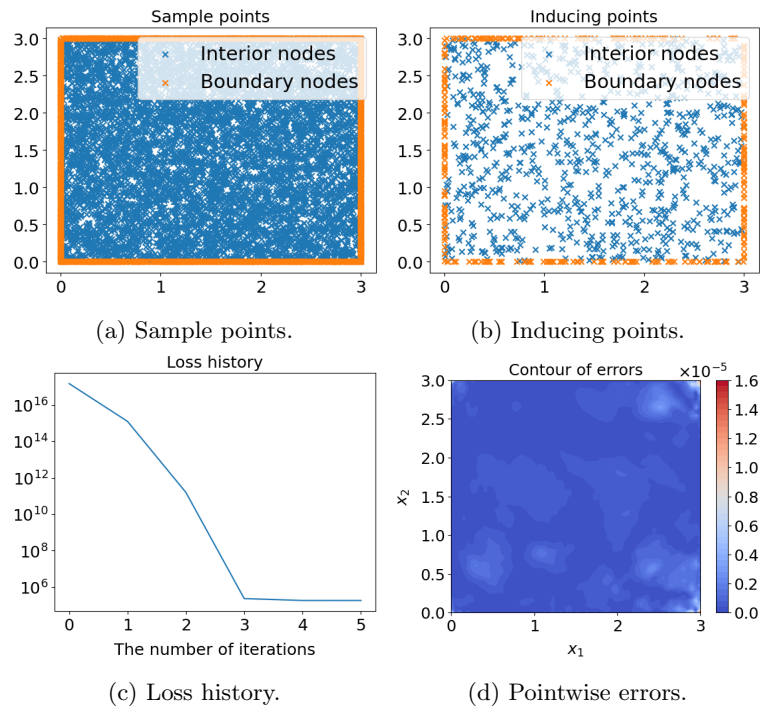


FIG. 2. Numerical results for the nonlinear elliptic equation: (a) a set of sample points; (b) a collection of inducing points; (c) the convergence history of the Gauss–Newton iterations; (d) contours of pointwise errors. The regularization parameters $\gamma = \eta = 10^{-12}$. $N = 9600$, $N_\Omega = 0.75 \times N$, $M = 1200$, $M_\Omega = 0.75 \times M$.

To further explore the efficacy of our algorithm. We compare the L^∞ errors of the SGP method and the GP algorithm in [4] for different numbers of sample points and inducing points in Table 1. We see that when $N = M$, both algorithms achieve accuracy of the same magnitude. Meanwhile, when $N \gg M$, the SGP method with N sample points and M inducing points is more accurate than the GP method with M sample points. This phenomenon implies that for answering Problem 1, purely reducing the number of base functions in a function space may not result in an efficient subspace. Furthermore, Table 1 implies that the SGP method with N sample points, half of which are used as inducing points, achieves comparable accuracy to the GP method with the same number of sample points. This is also confirmed in Figure 3, where we plot the means and the standard variances of the maximum pointwise errors over 10 different realizations against

different numbers of inducing points when the number of samples is fixed to be 4800. Figure 3 shows that the accuracy is not deteriorated if the number of inducing points is larger than 2000, which is smaller than half the number of samples. Hence, by using (2.9), the size of the matrix needed to be inverted in the SGP method can be significantly reduced. In other words, the SGP approach consumes much less computational time while retaining desirable accuracy compared to the GP method. Figure 3 also implies that if the number of inducing points is too small, the induced space \mathcal{U}_Q cannot maintain sufficient information about the approximated solution. Therefore, there is a trade-off between accuracy and efficiency. If the number of inducing points is chosen properly, the subspace generated by the SGP method serves as a candidate answer to Problem 1 in terms of the GP method.

N	1200	2400	4800	9600
GP, L^∞ error	1.34×10^{-1}	1.01×10^{-3}	6.78×10^{-5}	4.46×10^{-6}
SGP($M = 600$), L^∞ error	1.46×10^{-1}	9.34×10^{-3}	2.55×10^{-3}	2.91×10^{-3}
SGP($M = 1200$), L^∞ error	1.40×10^{-1}	1.37×10^{-3}	7.58×10^{-5}	2.16×10^{-5}
SGP($M = 2400$), L^∞ error	N/A	1.92×10^{-3}	2.62×10^{-5}	6.20×10^{-6}
SGP($M = 4800$), L^∞ error	N/A	N/A	3.15×10^{-5}	5.13×10^{-6}

TABLE 1. The nonlinear elliptic equation: L^∞ errors of the SGP method and the GP algorithm for different numbers of sample points and inducing points. The N/A means not available. We set $\gamma = \eta = 10^{-12}$.




FIG. 3. The nonlinear elliptic equation: L^∞ errors for different numbers of inducing points when the number of samples N equals to 4800.

4.2. A Stationary Mean Field Game. In this subsection, we explore the behavior of the SGP method by solving a stationary mean field game (MFG) system, which consists of two coupled PDEs, a Hamilton–Jacobi equation and a Fokker–Planck equation. MFGs model the behaviors of agents in a large population. To learn more about the use of finite differences and machine learning methods for solving MFGs, we suggest referring to the references [1, 2, 18, 30]. The GP method and a random Fourier feature algorithm for solving MFG are presented in [21]. Following the same framework as in [21], our SGP method can be naturally extended to solve PDE systems. In this example, we seek $(u, m, \lambda) \in C^\infty(\mathbb{T}^2) \times C^\infty(\mathbb{T}^2) \times \mathbb{R}$ solving

$$\begin{cases} -\nu\Delta u + \frac{|\nabla u|^2}{2} + \frac{\sin(4\pi x_1) + \cos(4\pi x_1) + \sin(4\pi x_2)}{2} = m^2 + \lambda, \forall x := (x_1, x_2) \in \mathbb{T}^2, \\ -\nu\Delta m - \operatorname{div}(m\nabla u) = 0, \forall x := (x_1, x_2) \in \mathbb{T}^2, \\ \int_{\mathbb{T}^2} m \, dx = 1, \int_{\mathbb{T}^2} u \, dx = 0, \end{cases} \quad (4.1)$$

where $\nu = 0.1$ and \mathbb{T}^2 is the two-dimensional torus. In the computations, we identify \mathbb{T}^2 with $[-0.5, 0.5]^2$. Meanwhile, to deal with the periodic condition, we choose the kernel

$$K((x_1, x_2), (x'_1, x'_2)) = \exp(\cos(2\pi(x_1 - x'_1)) + \cos(2\pi(x_2 - x'_2)) - 2). \quad (4.2)$$

Then, we randomly select M inducing points from N uniform random samples. For simplicity, we approximate both u and m in the same RKHS \mathcal{U}_Q generated by inducing points as described in Subsection 3.1. More precisely, we consider the following minimization problem

$$\left\{ \begin{array}{l} \min_{u, m \in \mathcal{U}_Q, \lambda \in \mathbb{R}, \mathbf{z}, \rho \in \mathbb{R}^N} \gamma \|u\|_{\mathcal{U}_Q}^2 + \gamma \|m\|_{\mathcal{U}_Q}^2 + \gamma \lambda^2 + \sum_{j=1}^N |z_j^{(1)} - u(\mathbf{x}_j)|^2 + \sum_{j=1}^N |z_j^{(2)} - \partial_{x_1} u(\mathbf{x}_j)|^2 \\ \quad + \sum_{j=1}^N |z_j^{(3)} - \partial_{x_2} u(\mathbf{x}_j)|^2 + \sum_{j=1}^N |z_j^{(4)} - \Delta u(\mathbf{x}_j)|^2 + \sum_{j=1}^N |\rho_j^{(1)} - m(\mathbf{x}_j)|^2 \\ \quad + \sum_{j=1}^N |\rho_j^{(2)} - \partial_{x_1} m(\mathbf{x}_j)|^2 + \sum_{j=1}^N |\rho_j^{(3)} - \partial_{x_2} m(\mathbf{x}_j)|^2 + \sum_{j=1}^N |\rho_j^{(4)} - \Delta m(\mathbf{x}_j)|^2 \\ \quad + \sum_{j=1}^N |\nu z_j^{(4)} - \frac{(z_j^{(2)} + z_j^{(3)})^2}{2} - \frac{\sin(4\pi x_{1,j}) + \cos(4\pi x_{1,j}) + \sin(4\pi x_{2,j})}{2} + (\rho_j^{(1)})^2 + \lambda|^2 \\ \quad + \sum_{j=1}^N |-\nu \rho_j^{(4)} - \rho_j^{(2)} z_j^{(2)} - \rho_j^{(3)} z_j^{(3)} - \rho_j^{(1)} z_j^{(4)}|^2, \\ \text{s.t. } \sum_{j=1}^N z_j^{(1)} = 0, \frac{1}{N} \sum_{j=1}^N \rho_j^{(1)} = 1. \end{array} \right. \quad (4.3)$$

Then, we derive the representer formulas for u and m as what we did in Section 3. To study the efficacy of our SGP method, for clarification, we only compare the pointwise errors between a minimizer m^\dagger of (4.3) and a reference solution m^* on a 60×60 uniform grid. To get m^* , we use the proximal method proposed in [2] to compute (4.1) on a 120×120 uniform grid and linearly interpolate the solution on the coarser 60×60 grid.

Figure 4 illustrates the numerical results when we take $N = 800$ samples and $M = 50$ inducing points. Here, we set $\gamma = 10^{-10}$ and $\eta = 10^{-4}$. Figures 4a and 4b plot the samples and inducing points, respectively. Figure 4c shows the loss history of the Gauss–Newton iterations and Figure 4d presents the contour of pointwise errors. Table 2 compares the L^∞ errors of m^* and m^\dagger for different values of N and M . We see that for solving the MFG system (4.1), the SGP method exhibits similar efficacy as in the previous experiments for solving a single nonlinear elliptic equation. That is, the SGP method which takes half of the samples as inducing points, achieves comparable accuracy to the GP method using the same number of samples.

N	100	200	400	800
GP, L^∞ error	1.09×10^{-1}	2.31×10^{-2}	1.05×10^{-2}	4.36×10^{-3}
SGP($M = 50$), L^∞ error	1.16×10^{-1}	2.61×10^{-2}	1.10×10^{-2}	6.20×10^{-3}
SGP($M = 100$), L^∞ error	1.18×10^{-1}	2.31×10^{-2}	1.25×10^{-2}	5.32×10^{-3}
SGP($M = 200$), L^∞ error	N/A	2.26×10^{-2}	1.15×10^{-2}	6.97×10^{-3}
SGP($M = 400$), L^∞ error	N/A	N/A	1.12×10^{-2}	4.04×10^{-3}

TABLE 2. The stationary MFG: The L^∞ errors of the SGP method and the GP algorithm are analyzed under varying sample point and inducing point configurations. The N/A means not available. We set $\gamma = 10^{-10}$ and $\eta = 10^{-4}$.

4.3. Burgers' equation. In this example, we consider the same Burger's equation as in [4]. More precisely, given $\nu = 0.02$, we seek to find u solving

$$\begin{cases} \partial_t u + u \partial_x u + \nu \partial_x^2 u = 0, \forall (t, x) \in (0, 1] \times (-1, 1), \\ u(0, x) = -\sin(\pi x), \forall x \in [-1, 1], \\ u(t, -1) = u(t, 1) = 0, \forall t \in [0, 1]. \end{cases}$$

We uniformly sample N points in the space-time domain, from which M samples are randomly chosen as inducing points. All experiments in this example use a fixed ratio of interior points, $N_\Omega/N = M_\Omega/M = 5/6$. To take into consideration the space and time variability of the solution to Burger's equation, as in [4], we choose the anisotropic kernel

$$K((t, x), (t', x')) = \exp(-(t - t')^2/\sigma_1^2 - (x - x')^2/\sigma_2^2), \quad (4.4)$$

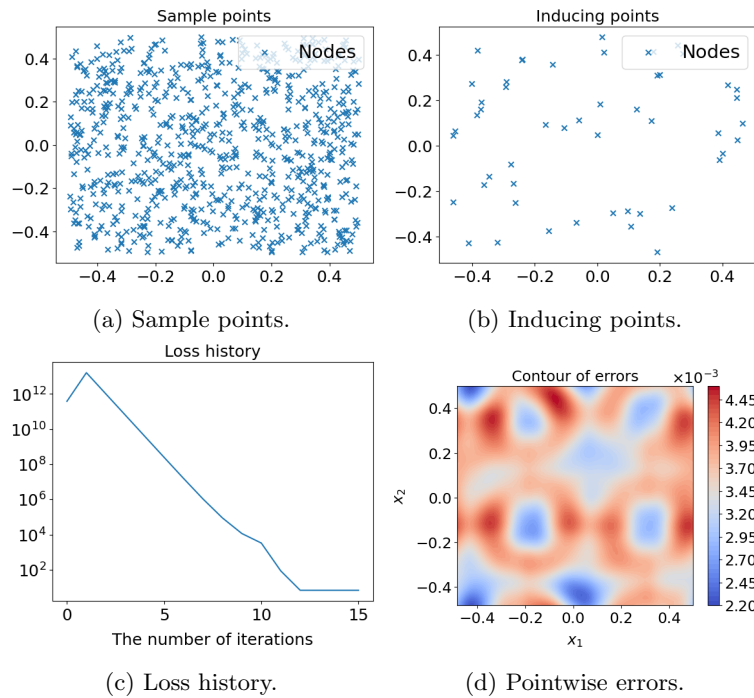


FIG. 4. Numerical results for the stationary MFG: (a) a set of sample points; (b) a collection of inducing points; (c) the convergence history of the Gauss–Newton iterations; (d) contours of pointwise errors. The regularization parameters $\gamma = 10^{-10}$ and $\eta = 10^{-4}$. $N = 800$ and $M = 50$.

with $(\sigma_1, \sigma_2) = (0.3, 0.05)$. For comparison, as in [4], the true solution is calculated from the Cole–Hopf transformation and the numerical quadrature. We use the technique of eliminating variables to deal with the constraints in (3.18) (see Subsection 3.3 in [4]) and use the Gauss–Newton method to solve the resulting unconstrained minimization problem. All the experiments in this example stop once the errors between two successive steps are less than 10^{-5} .

Figure 5 shows the numerical results for $N = 2400$ and $M = 600$ with parameters $\gamma = \eta = 10^{-6}$. The samples and the inducing points are given in Figures 5a and 5b. The history of Gauss–Newton iterations is plotted in Figure 5c, which implies the fast convergence of the SGP method. The absolute values of the pointwise errors between the numerical solution and the true solution are shown in Figure 5d. Similar to the results of the GP method given in [4], the maximum errors occurred close to the shock. In Figures 5e–5g, we compare the numerical and true solutions at times $t = 0.25, 0.5, 0.75$ to highlight the accuracy of our SGP method.

In Table 3, we present the L^∞ errors of the SGP method and the GP algorithm in [4] for different choices of N and M . Similar to the nonlinear elliptic example, the SGP algorithm yields comparable accurate numerical solutions to the GP method by using half of the samples as inducing points.

4.4. A Nonlinear Parabolic Equation. We consider the numerical solution of the parabolic equation

$$\begin{cases} \partial_t u - \partial_x^2 u + \frac{1}{2} |\partial_x u|^2 + u + x \partial_x u = f, \forall (t, x) \in (0, 1] \times (0, 3/2), \\ u(0, x) = g(x), \forall x \in (0, 3/2), \\ u(t, -1) = h_1(x), u(t, 1) = h_2(x), \forall t \in (0, 1). \end{cases}$$

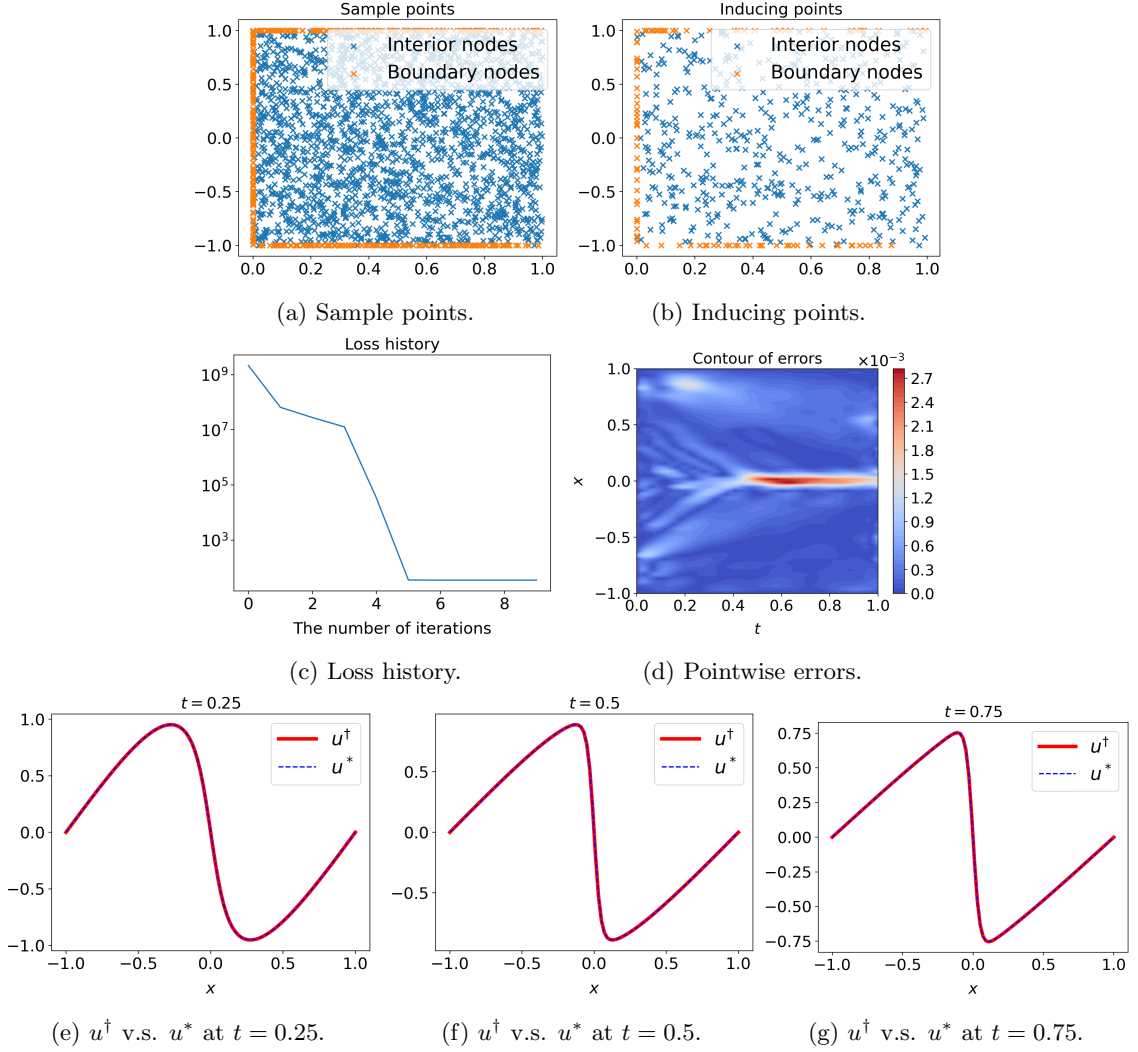


FIG. 5. Numerical results for the Burger's equation: (a) a set of sample points; (b) a collection of inducing points; (c) the convergence history of the Gauss–Newton iterations; (d) contours of pointwise errors; (e)–(g) time slices of the numerical approximation u^\dagger and the true solution u^* at $t = 0.25, 0.5, 0.75$. The parameters $\gamma = \eta = 10^{-6}$, $N = 2400$, $N_\Omega = 5N/6$, $M = 600$, and $M_\Omega = 5M/6$.

We prescribe the true solution $u(t, x) = (\sin(\pi x) + 2 \cos(2\pi x))e^{-t}$ and compute f , g , h_1 , and h_2 accordingly. We fix the ratios $N_\Omega/N = M_\Omega/M = 6/7$. Meanwhile, we use the anisotropic kernel in (4.4) with the same lengthscales, i.e., $(\sigma_1, \sigma_2) = (0.3, 0.05)$. We use the technique of eliminating variables to handle the constraints in (3.18) (see Subsection 3.3 of [4]). All the experiments in this example stop once the errors between two successive steps are less than 10^{-5} or the iteration numbers exceed 20.

Figure 6 plots the numerical results when $N = 2800$, $M = 700$, and $\gamma = \eta = 10^{-10}$. Figures 6a and 6b depict the samples and the inducing points. The history of Gauss–Newton iterations is plotted in Figure 6c, which shows that the algorithm converges within 6 iterations. The absolute values of the pointwise errors between the numerical solution and the true solution are given in Figure 6d. Table 4 records the L^∞ errors

N	600	1200	2400	4800
GP, L^∞ error	6.29×10^{-1}	4.66×10^{-2}	4.92×10^{-3}	4.60×10^{-4}
SGP($M = 300$), L^∞ error	6.80×10^{-1}	7.10×10^{-2}	1.72×10^{-2}	1.15×10^{-2}
SGP($M = 600$), L^∞ error	7.15×10^{-1}	7.56×10^{-2}	4.60×10^{-3}	1.30×10^{-3}
SGP($M = 1200$), L^∞ error	N/A	5.83×10^{-2}	4.24×10^{-3}	1.13×10^{-3}
SGP($M = 2400$), L^∞ error	N/A	N/A	3.53×10^{-3}	4.47×10^{-4}

TABLE 3. Burger’s equation: L^∞ errors of the SGP method and the GP algorithm for different numbers of sample points and inducing points. The N/A means not available. We set $\gamma = \eta = 10^{-6}$ for $N \leq 1200$ and choose $\gamma = \eta = 10^{-8}$ for $N = 2400$.

of the SGP method and the GP algorithm in [4] for different values of N and M , which justifies the efficacy of the SGP method.

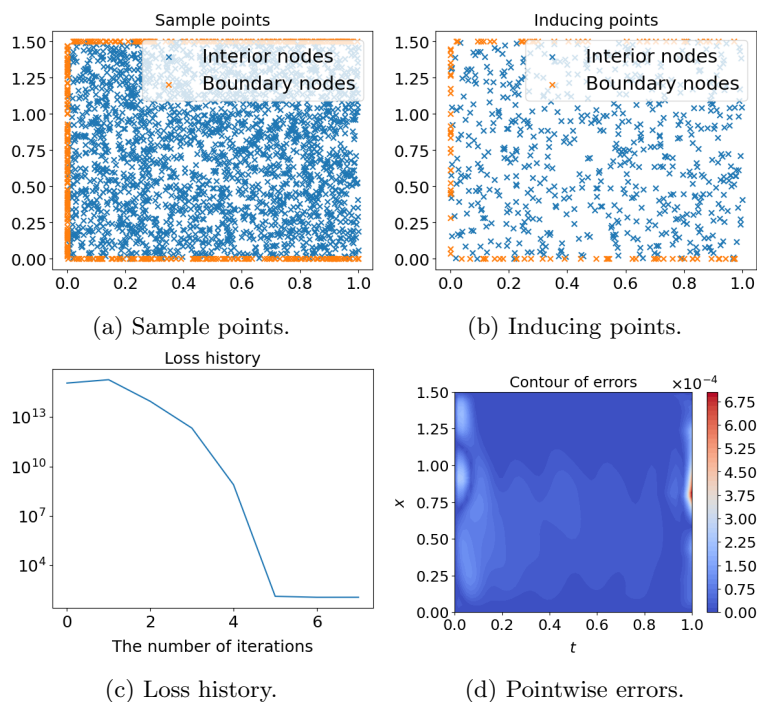


FIG. 6. Numerical results for the parabolic equation: (a) a set of sample points; (b) a collection of inducing points; (c) the convergence history of the Gauss–Newton iterations; (d) contours of pointwise errors. The regularization parameters $\gamma = \eta = 10^{-10}$. $N = 2800$, $N_\Omega = 6/7 \times N$, $M = 700$, $M_\Omega = 6/7 \times M$.

5. CONCLUSION AND FUTURE WORK

In general, numerical techniques used to solve nonlinear PDEs provide an estimation of the solution within a specific function space. The function space may contain redundant information in terms of PDE solutions if it is not chosen properly. In this paper, we formulate the problem of finding a “compressed” subspace, in which we seek an approximated solution to a nonlinear PDE with a negligible loss in accuracy compared to the result of using the original function space. In particular, we examine the GP method that approximates functions in the span of base functions formulated by evaluating derivatives of different orders of kernels

N	700	1400	2800	5600
GP, L^∞ error	2.75×10^{-1}	6.42×10^{-3}	5.13×10^{-4}	7.25×10^{-5}
SGP($M = 350$), L^∞ error	3.09×10^{-1}	2.92×10^{-2}	2.67×10^{-2}	3.63×10^{-3}
SGP($M = 700$), L^∞ error	2.79×10^{-1}	1.50×10^{-2}	6.93×10^{-4}	1.00×10^{-4}
SGP($M = 1400$), L^∞ error	N/A	9.47×10^{-3}	5.22×10^{-4}	4.70×10^{-5}
SGP($M = 2800$), L^∞ error	N/A	N/A	3.36×10^{-4}	4.32×10^{-5}

TABLE 4. The parabolic PDE: L^∞ errors of the SGP method and the GP algorithm for different numbers of sample points and inducing points. The N/A means not available. We set $\gamma = \eta = 10^{-10}$.

at random samples. We present a SGP method which seeks solutions for nonlinear PDEs in “condensed” RKHSs. Specifically, the SGP method finds a numerical solution to a PDE in the RKHS associated with a low-rank kernel generated by inducing points. Meanwhile, the approximation solution can be viewed as a maximum *a posteriori* probability estimator of a SGP conditioned a noisy observation of values of linear operators satisfying the PDE at a finite set of samples. Our numerical experiments imply that if the number of inducing points is chosen properly, the SGP algorithm produces a subspace that holds sufficient enough information about the approximated solution. Hence, the SGP method consumes less computational time than the GP method in [4] and achieves comparable accuracy when both methods use uniform samples. We notice that the positions of inducing points greatly influence the performance of our algorithm. Hence, a potential future work is to investigate a better way to put inducing points. Meanwhile, the choice of hyperparameters has a profound impact on the performance of our method. The probabilistic interpretation of the SGP approach in Subsection 2.3 provides a way for hyperparameter learning in future work.

REFERENCES

- [1] Y. Achdou and I. Capuzzo-Dolcetta. Mean field games: numerical methods. *SIAM Journal on Numerical Analysis*, 48(3):1136–1162, 2010.
- [2] L. M. Briceno-Arias, D. Kalise, and F. Silva. Proximal methods for stationary mean field games with local couplings. *SIAM Journal on Control and Optimization*, 56(2):801–836, 2018.
- [3] J. Cao, M. Kang, F. Jimenez, H. Sang, F. Schafer, and M. Katzfuss. Variational sparse inverse Cholesky approximation for latent Gaussian processes via double Kullback-Leibler minimization. *arXiv preprint arXiv:2301.13303*, 2023.
- [4] Y. Chen, B. Hosseini, H. Owhadi, and A. M. Stuart. Solving and learning nonlinear PDEs with Gaussian processes. *Journal of Computational Physics*, 2021.
- [5] Y. Chen, H. Owhadi, and F. Schäfer. Sparse Cholesky factorization for solving nonlinear PDEs via Gaussian processes. *arXiv preprint arXiv:2304.01294*, 2023.
- [6] J. Cockayne, C. Oates, T. Sullivan, and M. Girolami. Probabilistic numerical methods for PDE-constrained Bayesian inverse problems. In *AIP Conference Proceedings*, volume 1853, page 060001. AIP Publishing LLC, 2017.
- [7] P. Drineas, M. W. Mahoney, and N. Cristianini. On the Nyström method for approximating a gram matrix for improved kernel-based learning. *journal of machine learning research*, 6(12), 2005.
- [8] C. Fowlkes, S. Belongie, F. Chung, and J. Malik. Spectral grouping using the Nyström method. *IEEE transactions on pattern analysis and machine intelligence*, 26(2):214–225, 2004.
- [9] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson. GPytorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration. *Advances in neural information processing systems*, 31, 2018.
- [10] A. Gersho and R. M. Gray. *Vector quantization and signal compression*, volume 159. Springer Science & Business Media, 2012.
- [11] J. Hensman, N. Durrande, and A. Solin. Variational Fourier features for Gaussian processes. *J. Mach. Learn. Res.*, 18(1):5537–5588, 2017.
- [12] J. Hensman, N. Fusi, and N. D. Lawrence. Gaussian processes for big data. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, UAI’13, pages 282–290, Arlington, Virginia, United States, 2013. AUAI Press.
- [13] T. J. Hughes. *The finite element method: linear static and dynamic finite element analysis*. Courier Corporation, 2012.
- [14] R. Jin, T. Yang, M. Mahdavi, Y. Li, and Z. Zhou. Improved bounds for the Nyström method with application to kernel classification. *IEEE Transactions on Information Theory*, 59(10):6939–6949, 2013.
- [15] N. Krämer, J. Schmidt, and P. Hennig. Probabilistic numerical method of lines for time-dependent partial differential equations. *arXiv preprint arXiv:2110.11847*, 2021.

- [16] M. Lázaro-Gredilla and A. Figueiras-Vidal. Inter-domain Gaussian processes for sparse inference using inducing features. *Advances in Neural Information Processing Systems*, 22, 2009.
- [17] M. Lázaro-Gredilla, J. Quinonero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal. Sparse spectrum Gaussian process regression. *The Journal of Machine Learning Research*, 11:1865–1881, 2010.
- [18] A. T. Lin, S. W. Fung, W. Li, L. Nurbekyan, and S. J. Osher. Alternating the population and control neural networks to solve high-dimensional stochastic mean-field games. *Proceedings of the National Academy of Sciences*, 118(31), 2021.
- [19] H. Liu, Y. S. Ong, X. Shen, and J. Cai. When gaussian process meets big data: A review of scalable GPs. *IEEE transactions on neural networks and learning systems*, 31(11):4405–4423, 2020.
- [20] L. Lu, P. Jin, G. Pang, Z. Zhang, and G. E. Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- [21] C. Mou, X. Yang, and C. Zhou. Numerical methods for mean field games based on Gaussian processes and Fourier features. *Journal of Computational Physics*, 2022.
- [22] H. Owhadi. Bayesian numerical homogenization. *Multiscale Modeling and Simulation*, 13(3):812–828, 2015.
- [23] H. Owhadi. Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games. *SIAM Review*, 59(1):99–149, 2017.
- [24] H. Owhadi and C. Scovel. *Operator-Adapted Wavelets, Fast Solvers, and Numerical Homogenization: From a Game Theoretic Approach to Numerical Approximation and Algorithm Design*, volume 35. Cambridge University Press, 2019.
- [25] A. Quarteroni and A. Valli. *Numerical approximation of partial differential equations*, volume 23. Springer Science and Business Media, 2008.
- [26] J. Quinonero-Candela and C. E. Rasmussen. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.
- [27] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, volume 3, page 5. Citeseer, 2007.
- [28] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Numerical Gaussian processes for time-dependent and nonlinear partial differential equations. *SIAM Journal on Scientific Computing*, 40(1):A172–A198, 2018.
- [29] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- [30] L. Ruthotto, S. J. Osher, W. Li, L. Nurbekyan, and S. W. Fung. A machine learning framework for solving high-dimensional mean field game and mean field control problems. *Proceedings of the National Academy of Sciences*, 117(17):9183–9193, 2020.
- [31] F. Schäfer, M. Katzfuss, and H. Owhadi. Sparse Cholesky factorization by Kullback–Leibler minimization. *SIAM Journal on scientific computing*, 43(3):A2019–A2046, 2021.
- [32] M. Seeger. PAC-Bayesian generalisation error bounds for Gaussian process classification. *J. Mach. Learn. Res.*, 3:233–269, March 2003.
- [33] S. Smale and D. Zhou. Shannon sampling II: Connections to learning theory. *Applied and Computational Harmonic Analysis*, 19(3):285–302, 2005.
- [34] S. Smale and D. Zhou. Geometry on probability spaces. *Constructive Approximation*, 30(3):311–323, 2009.
- [35] E. Snelson and Z. Ghahramani. Sparse Gaussian processes using pseudo-inputs. *Advances in neural information processing systems*, 18, 2005.
- [36] J. W. Thomas. *Numerical partial differential equations: finite difference methods*, volume 22. Springer Science and Business Media, 2013.
- [37] M. Titsias. Variational learning of inducing variables in sparse Gaussian processes. In *Artificial intelligence and statistics*, pages 567–574. PMLR, 2009.
- [38] J. Wang, J. Cockayne, O. Chkrebtii, T. J. Sullivan, and C. J. Oates. Bayesian numerical methods for nonlinear partial differential equations. *Statistics and Computing*, 31(5):1–20, 2021.
- [39] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*. MIT press Cambridge, MA., 2006.
- [40] A. Wilson and H. Nickisch. Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *International conference on machine learning*, pages 1775–1784. PMLR, 2015.
- [41] F. X. Yu, A. T. Suresh, K. M. Choromanski, D. N. Holtmann-Rice, and S. Kumar. Orthogonal random features. *Advances in neural information processing systems*, 29:1975–1983, 2016.
- [42] Y. Zang, G. Bao, X. Ye, and H. Zhou. Weak adversarial networks for high-dimensional partial differential equations. *Journal of Computational Physics*, 411, 2020.
- [43] K. Zhang, I. W. Tsang, and J. T. Kwok. Improved Nyström low-rank approximation and error analysis. In *Proceedings of the 25th international conference on Machine learning*, pages 1232–1239, 2008.

A. PROOF FOR THE ERROR OF THE SGP METHOD

We first give a proof for Lemma 3.3.

Proof of Lemma 3.3. We first show that $S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I$ is invertible for any $\gamma > 0$. Let $u, v \in \mathcal{U}_Q$ and $S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} u + \gamma u = v$. Then, according to (3.12) and (3.13), we have

$$[\boldsymbol{\psi}, u]^T \mathcal{Q} \boldsymbol{\psi} + \gamma u = v. \quad (\text{A.1})$$

Acting $\boldsymbol{\psi}$ on both sides of (A.1), we get

$$\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})[\boldsymbol{\psi}, u] + \gamma[\boldsymbol{\psi}, u] = [\boldsymbol{\psi}, v].$$

Since $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi})$ is positive semi-definite, $\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I$ is invertible. Thus,

$$[\boldsymbol{\psi}, u] = (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}[\boldsymbol{\psi}, v]. \quad (\text{A.2})$$

We get from (A.1) and (A.2) that

$$u = \frac{1}{\gamma} v - \frac{1}{\gamma} [\boldsymbol{\psi}, v]^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{Q} \boldsymbol{\psi},$$

which implies that $S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I$ is bijective. Furthermore,

$$(S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} v = \frac{1}{\gamma} v - \frac{1}{\gamma} [\boldsymbol{\psi}, v]^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{Q} \boldsymbol{\psi}, \forall v \in \mathcal{U}_Q. \quad (\text{A.3})$$

Similarly, one can show that $S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I$ is bijective.

To prove (3.14), we see that

$$\begin{aligned} (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I) S_{\bar{\mathbf{x}}}^* (S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} &= (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma S_{\bar{\mathbf{x}}}^*) (S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} \\ &= S_{\bar{\mathbf{x}}}^* (S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I) (S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} = S_{\bar{\mathbf{x}}}^*, \end{aligned}$$

which yields (3.14).

For any $c \in \ell^2(\bar{\mathbf{x}})$, using (A.3) and (3.13), we get

$$\begin{aligned} (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} S_{\bar{\mathbf{x}}}^* c &= \frac{1}{\gamma} c^T \mathcal{Q} \boldsymbol{\psi} - \frac{1}{\gamma} [\boldsymbol{\psi}, c^T \mathcal{Q} \boldsymbol{\psi}] (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{Q} \boldsymbol{\psi} \\ &= \frac{1}{\gamma} c^T \mathcal{Q} \boldsymbol{\psi} - \frac{1}{\gamma} c^T \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{Q} \boldsymbol{\psi} \\ &= \frac{1}{\gamma} c^T \mathcal{Q} \boldsymbol{\psi} - \frac{1}{\gamma} c^T (\mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) + \gamma I - \gamma I) (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{Q} \boldsymbol{\psi} \\ &= c^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{Q} \boldsymbol{\psi}, \end{aligned}$$

which concludes (3.15).

For any $c \in \ell^2(\bar{\mathbf{x}})$, we use (3.15) and obtain

$$\begin{aligned} (I - S_{\bar{\mathbf{x}}} (S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} + \gamma I)^{-1} S_{\bar{\mathbf{x}}}^*) c &= c - S_{\bar{\mathbf{x}}} ((\mathcal{Q} \boldsymbol{\psi})^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} c) \\ &= c - \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} c \\ &= c - (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}) - \gamma I) (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} c \\ &= \gamma (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} c. \end{aligned} \quad (\text{A.4})$$

Since (A.4) holds for any $c \in \ell^2(\bar{\mathbf{x}})$, we conclude (3.16). \square

The next lemma is similar to Lemma 3.3, in which we establish the properties of a sampling operator leveraging the information of all the samples.

Lemma A.1. Define the sampling operator $\widehat{S}_{\bar{\mathbf{x}}} : \mathcal{U} \mapsto \ell^2(\bar{\mathbf{x}})$ by

$$\widehat{S}_{\bar{\mathbf{x}}} = [\boldsymbol{\psi}, u].$$

Denote $\widehat{S}_{\bar{\mathbf{x}}}^*$ as the adjoint of $\widehat{S}_{\bar{\mathbf{x}}}$ in the sense that for any $c \in \ell^2(\bar{\mathbf{x}})$, $\langle \widehat{S}_{\bar{\mathbf{x}}}^* c, u \rangle = \langle \widehat{S}_{\bar{\mathbf{x}}} u, c \rangle_{\ell^2(\bar{\mathbf{x}})}$. Then,

$$\widehat{S}_{\bar{\mathbf{x}}}^* c = c^T \mathcal{K} \boldsymbol{\psi}, \forall c \in \ell^2(\bar{\mathbf{x}}). \quad (\text{A.5})$$

Besides,

$$\widehat{S}_{\bar{\mathbf{x}}}^* (\widehat{S}_{\bar{\mathbf{x}}} \widehat{S}_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} = (\widehat{S}_{\bar{\mathbf{x}}}^* \widehat{S}_{\bar{\mathbf{x}}} + \gamma I)^{-1} \widehat{S}_{\bar{\mathbf{x}}}^*, \quad (\text{A.6})$$

and

$$(\widehat{S}_{\bar{\mathbf{x}}}^* \widehat{S}_{\bar{\mathbf{x}}} + \gamma I)^{-1} \widehat{S}_{\bar{\mathbf{x}}}^* c = (\mathcal{K} \boldsymbol{\psi})^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} c, \forall c \in \ell^2(\bar{\mathbf{x}}). \quad (\text{A.7})$$

Proof. Since $\widehat{S}_{\bar{\mathbf{x}}}^*$ is the adjoint of $\widehat{S}_{\bar{\mathbf{x}}}$, for each $c \in \ell^2(\bar{\mathbf{x}})$, we have

$$\langle \widehat{S}_{\bar{\mathbf{x}}}^* c, u \rangle = \langle \widehat{S}_{\bar{\mathbf{x}}} u, c \rangle_{\ell^2(\bar{\mathbf{x}})} = c^T [\boldsymbol{\psi}, u] = \langle c^T \mathcal{K}(\boldsymbol{\psi}), u \rangle, \forall u \in \mathcal{U}.$$

Thus, we conclude (A.5). Meanwhile, for $\gamma > 0$, by the same arguments as the proof of Lemma 3.3, $\widehat{S}_{\bar{\mathbf{x}}}$ admits similar properties to those of $S_{\bar{\mathbf{x}}}$, i.e., (A.6) and (A.7) holds. \square

Next, we give arguments for our main result, Theorem 3.7.

Proof of Theorem 3.7. Let $\widehat{S}_{\bar{\mathbf{x}}}$ and $\widehat{S}_{\bar{\mathbf{x}}}^*$ be defined as in Lemma A.1. To estimate the norm of $u^\dagger - u^*$, we introduce intermediate functions,

$$v := (\widehat{S}_{\bar{\mathbf{x}}}^* \widehat{S}_{\bar{\mathbf{x}}} + \gamma I)^{-1} \widehat{S}_{\bar{\mathbf{x}}}^* z^\dagger, \quad (\text{A.8})$$

$$\widehat{u} := (\widehat{S}_{\bar{\mathbf{x}}}^* \widehat{S}_{\bar{\mathbf{x}}} + \gamma I)^{-1} \widehat{S}_{\bar{\mathbf{x}}}^* \widehat{S}_{\bar{\mathbf{x}}} u^*. \quad (\text{A.9})$$

and

$$\bar{u} := (\mathcal{K} \boldsymbol{\psi})^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} [\boldsymbol{\psi}, u^*]. \quad (\text{A.10})$$

Next, we use the triangle inequality and obtain

$$\|u^\dagger - u^*\|_{\mathcal{U}} \leq \|u^\dagger - v\|_{\mathcal{U}} + \|v - \widehat{u}\|_{\mathcal{U}} + \|\widehat{u} - \bar{u}\|_{\mathcal{U}} + \|\bar{u} - u^*\|_{\mathcal{U}}. \quad (\text{A.11})$$

The accuracy of the approximation of \mathcal{Q} to \mathcal{K} is measured by $\|u^\dagger - v\|_{\mathcal{U}}$, while $\|v - \widehat{u}\|_{\mathcal{U}}$ represents the error between the values of linear operators of u^* and z^\dagger . The effect of the nugget term γI is captured by $\|\widehat{u} - \bar{u}\|_{\mathcal{U}}$, and the extent to which the span of $\mathcal{K} \boldsymbol{\psi}$ can approximate the true solution u^* is reflected in $\|\bar{u} - u^*\|_{\mathcal{U}}$. Next, we split the arguments for bounding terms at the right-hand side of (A.11) into different steps.

Step 1. Here, we give an upper bound for $\|u^\dagger - v\|_{\mathcal{U}}$, where v is defined in (A.8). Given $\gamma > 0$, let u^\dagger be as in (3.17) and let z^\dagger be a solution to (3.18). Then, we have

$$\begin{cases} S_{\bar{\mathbf{x}}}^* S_{\bar{\mathbf{x}}} u^\dagger + \gamma u^\dagger = S_{\bar{\mathbf{x}}}^* z^\dagger, \\ \widehat{S}_{\bar{\mathbf{x}}}^* \widehat{S}_{\bar{\mathbf{x}}} v + \gamma v = \widehat{S}_{\bar{\mathbf{x}}}^* z^\dagger. \end{cases} \quad (\text{A.12})$$

which yields

$$\begin{cases} S_{\bar{\mathbf{x}}} u^\dagger = (S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* z^\dagger, \\ \widehat{S}_{\bar{\mathbf{x}}} v = (\widehat{S}_{\bar{\mathbf{x}}} \widehat{S}_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} \widehat{S}_{\bar{\mathbf{x}}} \widehat{S}_{\bar{\mathbf{x}}}^* z^\dagger. \end{cases} \quad (\text{A.13})$$

Taking the difference of the two equations in (A.12) and using (A.13), we get

$$\begin{aligned} \gamma(u^\dagger - v) &= S_{\bar{\mathbf{x}}}^*(z^\dagger - S_{\bar{\mathbf{x}}} u^\dagger) - \widehat{S}_{\bar{\mathbf{x}}}^*(z^\dagger - \widehat{S}_{\bar{\mathbf{x}}} v) \\ &= \gamma S_{\bar{\mathbf{x}}}^* (S_{\bar{\mathbf{x}}} S_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} z^\dagger - \gamma \widehat{S}_{\bar{\mathbf{x}}}^* (\widehat{S}_{\bar{\mathbf{x}}} \widehat{S}_{\bar{\mathbf{x}}}^* + \gamma I)^{-1} z^\dagger. \end{aligned} \quad (\text{A.14})$$

By (3.14), (3.15), (A.6), (A.7), and (A.14), we have

$$\begin{aligned} u^\dagger - v &= (S_{\bar{x}}^* S_{\bar{x}} + \gamma I)^{-1} S_{\bar{x}}^* z^\dagger - (\widehat{S}_{\bar{x}}^* \widehat{S}_{\bar{x}} + \gamma I)^{-1} \widehat{S}_{\bar{x}}^* z^\dagger \\ &= (\mathcal{Q}\boldsymbol{\psi})^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} z^\dagger - (\mathcal{K}\boldsymbol{\psi})^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} z^\dagger. \end{aligned}$$

Thus, we obtain

$$\begin{aligned} \|u^\dagger - v\|_{\mathcal{U}}^2 &= (z^\dagger)^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} z^\dagger - (z^\dagger)^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} z^\dagger \\ &\quad + 2\gamma (z^\dagger)^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} z^\dagger \\ &\quad - \gamma (z^\dagger)^T (\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-2} z^\dagger - \gamma (z^\dagger)^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-2} z^\dagger \\ &\leq 3 \|(\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} - (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}\| \|z^\dagger\|^2, \end{aligned} \tag{A.15}$$

where we use $\|(\gamma I + \mathcal{Q}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}\| \leq 1/\gamma$ and $\|(\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}\| \leq 1/\gamma$ in the last inequality.

Step 2. Next, we bound $v - \widehat{u}$, where \widehat{u} is given in (A.9). Using (A.7) and the definitions (A.8)-(A.9), we get

$$v = (\mathcal{K}\boldsymbol{\psi})^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} z^\dagger, \tag{A.16}$$

$$\widehat{u} = (\mathcal{K}\boldsymbol{\psi})^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} [\boldsymbol{\psi}, u^*]. \tag{A.17}$$

Thus, we have

$$\begin{aligned} \|v - \widehat{u}\|_{\mathcal{U}}^2 &= (z^\dagger - [\boldsymbol{\psi}, u^*])^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}) (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} (z^\dagger - [\boldsymbol{\psi}, u^*]) \\ &= (z^\dagger - [\boldsymbol{\psi}, u^*])^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} (z^\dagger - [\boldsymbol{\psi}, u^*]) \\ &\quad - \gamma (z^\dagger - [\boldsymbol{\psi}, u^*])^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-2} (z^\dagger - [\boldsymbol{\psi}, u^*]) \\ &\leq (z^\dagger - [\boldsymbol{\psi}, u^*])^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (z^\dagger - [\boldsymbol{\psi}, u^*]), \end{aligned}$$

where the last inequality results from the fact that $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} - (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}$ and $(\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-2}$ are positive definite. Hence, we obtain

$$\|v - \widehat{u}\|_{\mathcal{U}} \leq \sqrt{(z^\dagger - [\boldsymbol{\psi}, u^*])^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (z^\dagger - [\boldsymbol{\psi}, u^*])}. \tag{A.18}$$

Step 3. Then, we estimate the norm of $\widehat{u} - \bar{u}$. Using (A.17), we get

$$\begin{aligned} \bar{u} - \widehat{u} &= (\mathcal{K}\boldsymbol{\psi})^T (\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} - (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}) [\boldsymbol{\psi}, u^*] \\ &= \gamma (\mathcal{K}\boldsymbol{\psi})^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} [\boldsymbol{\psi}, u^*] \end{aligned} \tag{A.19}$$

where we use $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} - (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} = \gamma \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}$ in the last equality. Thus, using (A.19) and the fact that $\mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-3} - (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1}$ is positive definite, we have

$$\begin{aligned} \|\widehat{u} - \bar{u}\|_{\mathcal{U}}^2 &= \gamma^2 [\boldsymbol{\psi}, u^*]^T (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-1} (\gamma I + \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi}))^{-1} [\boldsymbol{\psi}, u^*] \\ &\leq \gamma^2 [\boldsymbol{\psi}, u^*]^T \mathcal{K}(\boldsymbol{\psi}, \boldsymbol{\psi})^{-3} [\boldsymbol{\psi}, u^*]. \end{aligned} \tag{A.20}$$

Step 4. In the last step, we bound $\bar{u} - u^*$, where \bar{u} is given in (A.10). We observe that \bar{u} is the projection of u^* onto the span of $\mathcal{K}\boldsymbol{\psi}$ [24, Theorem 12.3]. Thus,

$$\|u^* - \bar{u}\|_{\mathcal{U}} = \inf_{\boldsymbol{\psi} \in \text{span}(\boldsymbol{\psi})} \|u^* - \mathcal{K}\boldsymbol{\psi}\|_{\mathcal{U}} \tag{A.21}$$

Therefore, (A.11), (A.15), (A.18), (A.20), and (A.21) imply that (3.23) holds. \square

B. PROOFS FOR THE NYSTRÖM APPROXIMATION ERROR

In this section, we give a detailed proof for estimating the upper bound for $\mathcal{E}(\mathcal{H}_a)$. For simplification, we write $\boldsymbol{\psi} := (\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_{\partial\Omega})$, where $\boldsymbol{\psi}_\Omega$ represents the collection of linear operators for interior points and $\boldsymbol{\psi}_{\partial\Omega}$ consists of linear operators for boundary samples. We write $\bar{\boldsymbol{x}} := (\bar{\boldsymbol{x}}_1, \bar{\boldsymbol{x}}_2)$, where we recall that $\bar{\boldsymbol{x}}$ contains all the sample points, $\bar{\boldsymbol{x}}_1$ is the set of interior points in Ω , and $\bar{\boldsymbol{x}}_2$ stands for the collection of boundary points. We define the operators $L_{\bar{\boldsymbol{x}}_1}$ and $L_{\bar{\boldsymbol{x}}_2}$ such that for any function $v \in \mathcal{U}$,

$$L_{\bar{\boldsymbol{x}}_1}[v] = \frac{1}{N_\Omega} (\mathcal{K}\boldsymbol{\psi}_\Omega)^T [\boldsymbol{\psi}_\Omega, v] \text{ and } L_{\bar{\boldsymbol{x}}_2}[v] = \frac{1}{N_{\partial\Omega}} (\mathcal{K}\boldsymbol{\psi}_{\partial\Omega})^T [\boldsymbol{\psi}_{\partial\Omega}, v], \quad (\text{B.1})$$

where N_Ω is the size of $\bar{\boldsymbol{x}}_1$ and $N_{\partial\Omega}$ is the length of $\bar{\boldsymbol{x}}_2$. The next proposition connects the eigenvalues and the eigenfunctions of $L_{\bar{\boldsymbol{x}}_1}$ and $L_{\bar{\boldsymbol{x}}_2}$ to the eigenvalues and the eigenvectors of $\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)$ and $\mathcal{K}(\boldsymbol{\psi}_{\partial\Omega}, \boldsymbol{\psi}_{\partial\Omega})$.

Proposition B.1. Let $R_1 = N_\Omega(D - D_b)$ and $R_2 = (N - N_\Omega)D_b$, where D and D_b are given in (3.6). Let $\{\lambda_j\}_{j=1}^{R_1}$ be the set of eigenvalues of $\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)$, which are ranked in the descending order. Define the operator $L_{\bar{\boldsymbol{x}}_1}$ as in (B.1). Then, the eigenvalues of $L_{\bar{\boldsymbol{x}}_1}$ are $\{\lambda_j/N_\Omega\}_{j=1}^{R_1}$. Moreover, let $\varphi_1, \dots, \varphi_{R_1}$ be the corresponding eigenfunctions of $L_{\bar{\boldsymbol{x}}_1}$ that have normalized functional norms, i.e. $\langle \varphi_i, \varphi_j \rangle = \delta_{ij}$, $1 \leq i, j \leq R_1$. Let $V = [\boldsymbol{v}_1, \dots, \boldsymbol{v}_{R_1}]$ be the orthonormal eigenvector matrix of $\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)$. Then, we have

$$\sqrt{\lambda_j} \boldsymbol{v}_j = [\boldsymbol{\psi}_\Omega, \varphi_j], \forall 1 \leq j \leq R_1, \quad (\text{B.2})$$

$$\sqrt{\lambda_j} \varphi_j = \boldsymbol{v}_j^T \mathcal{K}\boldsymbol{\psi}_\Omega, \forall 1 \leq j \leq R_1, \quad (\text{B.3})$$

$$\mathcal{K}\boldsymbol{\psi}_\Omega = V[\sqrt{\lambda_1}\varphi_1, \dots, \sqrt{\lambda_{R_1}}\varphi_{R_1}]^T, \quad (\text{B.4})$$

and

$$\mathcal{K}(\boldsymbol{\psi}_{\Omega,j}, \boldsymbol{\psi}_{\Omega,j}) = \sum_{i=1}^{R_1} [\boldsymbol{\psi}_{\Omega,j}, \varphi_i]^2, \forall 1 \leq j \leq R_1, \quad (\text{B.5})$$

where $\boldsymbol{\psi}_{\Omega,j}$ is the j^{th} component of $\boldsymbol{\psi}_\Omega$. Similarly, let $\{\tau_j\}_{j=1}^{R_2}$ be the set of eigenvalues of $\mathcal{K}(\boldsymbol{\psi}_{\partial\Omega}, \boldsymbol{\psi}_{\partial\Omega})$ sorted in the descending order and let $L_{\bar{\boldsymbol{x}}_2}$ be as in (B.1). Then, the eigenvalues of $L_{\bar{\boldsymbol{x}}_2}$ are $\{\tau_j/N_{\partial\Omega}\}_{j=1}^{R_2}$. Meanwhile, let $\zeta_1, \dots, \zeta_{R_2}$ be the eigenfunctions of $L_{\bar{\boldsymbol{x}}_2}$ that have normalized functional norms. Let $W = [\boldsymbol{w}_1, \dots, \boldsymbol{w}_{R_2}]$ be the orthonormal eigenvector matrix of $\mathcal{K}(\boldsymbol{\psi}_{\partial\Omega}, \boldsymbol{\psi}_{\partial\Omega})$. Then, we obtain

$$\sqrt{\tau_j} \boldsymbol{w}_j = [\boldsymbol{\psi}_{\partial\Omega}, \zeta_j], \forall 1 \leq j \leq R_2, \quad (\text{B.6})$$

$$\sqrt{\tau_j} \zeta_j = \boldsymbol{w}_j^T \mathcal{K}\boldsymbol{\psi}_{\partial\Omega}, \forall 1 \leq j \leq R_2, \quad (\text{B.7})$$

$$\mathcal{K}\boldsymbol{\psi}_{\partial\Omega} = W[\sqrt{\tau_1}\zeta_1, \dots, \sqrt{\tau_{R_2}}\zeta_{R_2}]^T, \quad (\text{B.8})$$

and

$$\mathcal{K}(\boldsymbol{\psi}_{\partial\Omega,j}, \boldsymbol{\psi}_{\partial\Omega,j}) = \sum_{i=1}^{R_2} [\boldsymbol{\psi}_{\partial\Omega,j}, \zeta_i]^2, \forall 1 \leq j \leq R_2, \quad (\text{B.9})$$

where $\boldsymbol{\psi}_{\partial\Omega,j}$ is the j^{th} element of $\boldsymbol{\psi}_{\partial\Omega}$.

Proof. We only give proofs related to $L_{\bar{\boldsymbol{x}}_1}$ since the arguments for $L_{\bar{\boldsymbol{x}}_2}$ are similar. Let $\tilde{\lambda}_j$ be j^{th} eigenvalue of $L_{\bar{\boldsymbol{x}}_1}$ and let φ_j be the corresponding eigenfunction with a normalized norm. Then, we have

$$L_{\bar{\boldsymbol{x}}_1}[\varphi_j] = \frac{1}{N_\Omega} (\mathcal{K}\boldsymbol{\psi}_\Omega)^T [\boldsymbol{\psi}_\Omega, \varphi_j] = \tilde{\lambda}_j \varphi_j. \quad (\text{B.10})$$

Acting $\boldsymbol{\psi}_\Omega$ on both sides of (B.10), we get

$$\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)[\boldsymbol{\psi}_\Omega, \varphi_j] = N_\Omega \tilde{\lambda}_j [\boldsymbol{\psi}_\Omega, \varphi_j], \quad (\text{B.11})$$

which implies that $N_\Omega \tilde{\lambda}_j$ is the j^{th} eigenvalue of $\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)$ and $[\boldsymbol{\psi}_\Omega, \varphi_j]$ is the corresponding eigenvector. Hence, we confirm that if $\{\lambda_j\}_{j=1}^{R_1}$ contains the eigenvalues of $\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)$, $\{\lambda_j/N_\Omega\}_{j=1}^{R_1}$ is the set of the eigenvalues of $L_{\bar{x}_1}$. Furthermore, let $V = [\mathbf{v}_1, \dots, \mathbf{v}_{R_1}]$ be the orthonormal eigenvector matrix of $\mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega)$. Then, (B.11) implies that

$$\mathbf{v}_j = \frac{[\boldsymbol{\psi}_\Omega, \varphi_j]}{\|[\boldsymbol{\psi}_\Omega, \varphi_j]\|}. \quad (\text{B.12})$$

From (B.10), we get

$$(\mathcal{K}\boldsymbol{\psi}_\Omega)^T[\boldsymbol{\psi}_\Omega, \varphi_j] = \lambda_j \varphi_j, \quad (\text{B.13})$$

which yields

$$\langle (\mathcal{K}\boldsymbol{\psi}_\Omega)^T[\boldsymbol{\psi}_\Omega, \varphi_j], (\mathcal{K}\boldsymbol{\psi}_\Omega)^T[\boldsymbol{\psi}_\Omega, \varphi_j] \rangle = \lambda_j^2 \langle \varphi_j, \varphi_j \rangle. \quad (\text{B.14})$$

Since φ_j is normalized, we obtain from (B.13) and (B.14) that

$$\lambda_j \|[\boldsymbol{\psi}_\Omega, \varphi_j]\|^2 = [\boldsymbol{\psi}_\Omega, \varphi_j]^T \mathcal{K}(\boldsymbol{\psi}_\Omega, \boldsymbol{\psi}_\Omega) [\boldsymbol{\psi}_\Omega, \varphi_j] = \lambda_j^2,$$

which gives

$$\|[\boldsymbol{\psi}_\Omega, \varphi_j]\| = \sqrt{\lambda_j}. \quad (\text{B.15})$$

Hence, (B.12) and (B.15) imply that $\sqrt{\lambda_j} \mathbf{v}_j = [\boldsymbol{\psi}_\Omega, \varphi_j]$, which confirms (B.2). Then, the equality (B.3) follows directly from (B.13) and (B.2). Meanwhile, (B.4) results from (B.3) and the fact that V is orthonormal.

Let $\boldsymbol{\psi}_{\Omega,j}$ be the j^{th} element of $\boldsymbol{\psi}_\Omega$. Denote by V_{ji} the component of V at the j^{th} row and the i^{th} column. From (B.4), we get

$$\mathcal{K}\boldsymbol{\psi}_{\Omega,j} = \sum_{i=1}^{R_1} V_{ji} \sqrt{\lambda_i} \varphi_i. \quad (\text{B.16})$$

Thus, we get from (B.16) that

$$[\boldsymbol{\psi}_{\Omega,j}, \mathcal{K}\boldsymbol{\psi}_{\Omega,j}] = \langle \mathcal{K}\boldsymbol{\psi}_{\Omega,j}, \mathcal{K}\boldsymbol{\psi}_{\Omega,j} \rangle = \sum_{i=1}^{R_1} \lambda_i V_{ji}^2 = \sum_{i=1}^{R_1} [\boldsymbol{\psi}_{\Omega,j}, \varphi_i]^2,$$

where the last equality follows (B.2). Thus, we conclude (B.5). \square

Using (B.4) and (B.8), \mathcal{H}_b in (3.24) can be rewritten in the basis of the eigenfunctions $\{\varphi_i\}_{i=1}^{R_1}$ and $\{\zeta_i\}_{i=1}^{R_2}$, i.e.,

$$\mathcal{H}_b = \left\{ f = \sum_{i=1}^{R_1} b_{1,i} \sqrt{\lambda_i} \varphi_i + \sum_{i=1}^{R_2} b_{2,i} \sqrt{\tau_i} \zeta_i, \sum_{i=1}^{R_1} b_{1,i}^2 + \sum_{i=1}^{R_2} b_{2,i}^2 \leq 1 \right\}. \quad (\text{B.17})$$

For any $\mathbf{r} := (r_1, r_2) \in [R_1] \times [R_2]$, we define

$$\begin{aligned} \mathcal{H}_b^{\mathbf{r}} &= \left\{ f = \sum_{i=1}^{r_1} b_{1,i} \sqrt{\lambda_i} \varphi_i + \sum_{i=1}^{r_2} b_{2,i} \sqrt{\tau_i} \zeta_i, \sum_{i=1}^{r_1} b_{1,i}^2 + \sum_{i=1}^{r_2} b_{2,i}^2 \leq 1 \right\}, \\ \bar{\mathcal{H}}_b^{\mathbf{r}} &= \left\{ f = \sum_{i=1}^{R_1-r_1} b_{1,i} \sqrt{\lambda_{i+r_1}} \varphi_{i+r_1} + \sum_{i=1}^{R_2-r_2} b_{2,i} \sqrt{\tau_{i+r_2}} \zeta_{i+r_2}, \sum_{i=1}^{R_1-r_1} b_{1,i}^2 + \sum_{i=1}^{R_2-r_2} b_{2,i}^2 \leq 1 \right\}. \end{aligned} \quad (\text{B.18})$$

Define

$$\mathcal{E}(\mathcal{H}_a, \mathbf{r}) = \max_{h \in \mathcal{H}_b^{\mathbf{r}}} \mathcal{E}(h, \mathcal{H}_a), \quad (\text{B.19})$$

which is the worst error in approximating any $h \in \mathcal{H}_b^{\mathbf{r}}$ by functions in \mathcal{H}_a . The next proposition bounds $\mathcal{E}(\mathcal{H}_a)$ by $\mathcal{E}(\mathcal{H}_a, \mathbf{r})$.

Proposition B.2. Let $\mathbf{r} := (r_1, r_2) \in [R_1] \times [R_2]$ and $\mathcal{E}(\mathcal{H}_a, \mathbf{r})$ be as in (B.19). Under the conditions of Proposition B.1, for any $\mathbf{r} \in [R_1] \times [R_2]$, we have

$$\mathcal{E}(\mathcal{H}_a) \leq 2\mathcal{E}(\mathcal{H}_a, \mathbf{r}) + 4\lambda_{r_1+1} + 4\tau_{r_2+1}, \quad (\text{B.20})$$

with the convention that $\lambda_{R_1+1} = 0$ and $\tau_{R_2+1} = 0$.

Proof. Let \mathcal{H}_b^r and $\overline{\mathcal{H}}_b^r$ be defined in (B.18). For any $h \in \mathcal{H}_b$, there exist $h_1 \in \mathcal{H}_b^r$ and $h_2 \in \overline{\mathcal{H}}_b^r$ such that $h = h_1 + h_2$. Thus, we have

$$\begin{aligned} \mathcal{E}(\mathcal{H}_a) &= \max_{\substack{h_1 \in \mathcal{H}_b^r \\ h_2 \in \overline{\mathcal{H}}_b^r}} \min_{v \in \mathcal{H}_a} \|v - h_1 - h_2\|_{\mathcal{U}}^2 \\ &\leq 2 \max_{h_1 \in \mathcal{H}_b^r} \min_{v \in \mathcal{H}_a} \|v - h_1\|_{\mathcal{U}}^2 + 2 \max_{h_2 \in \overline{\mathcal{H}}_b^r} \|h_2\|_{\mathcal{U}}^2 \leq 2\mathcal{E}(\mathcal{H}_a, \mathbf{r}) + 4\lambda_{r_1+1} + 4\tau_{r_2+1}, \end{aligned}$$

which concludes (B.20). \square

The above proposition implies that it is sufficient to bound $\mathcal{E}(\mathcal{H}_a, \mathbf{r})$ in order to estimate the Nyström approximation error. Let ϕ be as in (3.7). For clarification, we denote by $\phi := (\phi_\Omega, \phi_{\partial\Omega})$, where ϕ_Ω represents the set of linear operators for interior points and $\phi_{\partial\Omega}$ includes boundary linear operators. Let $\widehat{\mathbf{x}} := (\widehat{\mathbf{x}}_1, \widehat{\mathbf{x}}_2)$ be as in Assumption 1, let M_Ω be the size of $\widehat{\mathbf{x}}_1$, and let $M_{\partial\Omega}$ be the length of $\widehat{\mathbf{x}}_2$. We define the operators $L_{\widehat{\mathbf{x}}_1}$ and $L_{\widehat{\mathbf{x}}_2}$ such that for any function $v \in \mathcal{U}$,

$$L_{\widehat{\mathbf{x}}_1}[v] = \frac{1}{M_\Omega} (\mathcal{K}\phi_\Omega)^T[\phi_\Omega, v] \text{ and } L_{\widehat{\mathbf{x}}_2}[v] = \frac{1}{M_{\partial\Omega}} (\mathcal{K}\phi_{\partial\Omega})^T[\phi_{\partial\Omega}, v]. \quad (\text{B.21})$$

The following proposition provides upper bounds for the Hilbert–Schmidt norms of $L_{\overline{\mathbf{x}}_1} - L_{\widehat{\mathbf{x}}_1}$ and $L_{\overline{\mathbf{x}}_2} - L_{\widehat{\mathbf{x}}_2}$. We recall that for any linear operator $L : \mathcal{H} \mapsto \mathcal{H}$, where \mathcal{H} is a Hilbert space, we denote by $\|L\|_{HS}$ and $\|L\|_2$ the Hilbert–Schmidt norm and the spectral norm, respectively, i.e.

$$\|L\|_{HS} = \sqrt{\sum_{i,j} \langle \mathbf{e}_i, L\mathbf{e}_j \rangle_{\mathcal{H}}^2} \text{ and } \|L\|_2 = \max_{\|v\|_{\mathcal{H}} \leq 1} \|Lv\|_{\mathcal{H}}. \quad (\text{B.22})$$

where $\{\mathbf{e}_i, i = 1, \dots\}$ is a complete orthogonal basis of \mathcal{H} . We note that for any linear operator $L : \mathcal{H} \mapsto \mathcal{H}$, $\|L\|_2 \leq \|L\|_{HS}$.

Proposition B.3. Suppose that Assumptions 1 and 2 hold. Let $L_{\overline{\mathbf{x}}_1}, L_{\overline{\mathbf{x}}_2}$ be as in (B.1) and let $L_{\widehat{\mathbf{x}}_1}, L_{\widehat{\mathbf{x}}_2}$ be given in (B.21). Then, with a probability $1 - \delta$, $0 < \delta < 1$, there exists a constant $C > 0$ such that

$$\|L_{\overline{\mathbf{x}}_1} - L_{\widehat{\mathbf{x}}_1}\|_{HS} \leq \frac{4C \ln(2/\delta)}{\sqrt{M_\Omega}} \text{ and } \|L_{\overline{\mathbf{x}}_2} - L_{\widehat{\mathbf{x}}_2}\|_{HS} \leq \frac{4C \ln(2/\delta)}{\sqrt{M_{\partial\Omega}}}. \quad (\text{B.23})$$

Proof. Here, we adapt the proof of Corollary 8 in [14]. Let $\psi^{\mathbf{x}}$ be given in Assumption 2 for $\mathbf{x} \in \overline{\mathbf{x}}$. For any $v \in \mathcal{U}$, we define

$$\xi(\mathbf{x})[v] = \sum_{i=D_b+1}^D \mathcal{K}\psi_i^{\mathbf{x}}[\psi_i^{\mathbf{x}}, v], \forall \mathbf{x} \in \overline{\mathbf{x}}_1, \text{ and } \widehat{\xi}(\mathbf{x})[v] = \sum_{i=1}^{D_b} \mathcal{K}\psi_i^{\mathbf{x}}[\psi_i^{\mathbf{x}}, v], \forall \mathbf{x} \in \overline{\mathbf{x}}_2.$$

Then, $L_{\widehat{\mathbf{x}}_1} = \frac{1}{M_\Omega} \sum_{i=1}^{M_\Omega} \xi(\widehat{\mathbf{x}}_i)$, and $L_{\widehat{\mathbf{x}}_2} = \frac{1}{M_{\partial\Omega}} \sum_{i=M_\Omega+1}^M \widehat{\xi}(\widehat{\mathbf{x}}_i)$. According to Assumption 1, $E[\xi(\mathbf{x})] = L_{\overline{\mathbf{x}}_1}$ and $E[\widehat{\xi}(\mathbf{x})] = L_{\overline{\mathbf{x}}_2}$. Denote by \mathcal{H}_1 the span of $\{\varphi_i\}_{i=1}^{R_1}$ and by \mathcal{H}_2 the span of $\{\zeta_i\}_{i=1}^{R_2}$, which are endowed with the norm of \mathcal{U} . Under Assumption 1, we observe that $L_{\overline{\mathbf{x}}_1}$ and $L_{\widehat{\mathbf{x}}_1}$ map \mathcal{H}_1 to \mathcal{H}_1 , and that $L_{\overline{\mathbf{x}}_2}$ and $L_{\widehat{\mathbf{x}}_2}$ map \mathcal{H}_2 to \mathcal{H}_2 . Thus, by the definition of the Hilbert–Schmidt norm in (B.22), there exists a constant

C such that

$$\begin{aligned}
\|\xi(\mathbf{x})\|_{HS} &= \sqrt{\sum_{j,k}^{R_1} \left\langle \sum_{i=D_b+1}^D \mathcal{K}\psi_i^{\mathbf{x}}[\psi_i^{\mathbf{x}}, \varphi_j], \varphi_k \right\rangle^2} = \sqrt{\sum_{j,k}^{R_1} \left(\sum_{i=D_b+1}^D [\psi_i^{\mathbf{x}}, \varphi_j][\psi_i^{\mathbf{x}}, \varphi_k] \right)^2} \\
&\leq C \sqrt{\sum_{j,k}^{R_1} \sum_{i=D_b+1}^D [\psi_i^{\mathbf{x}}, \varphi_j]^2 [\psi_i^{\mathbf{x}}, \varphi_k]^2} = C \sqrt{\sum_{i=D_b+1}^D \left(\sum_j^{R_1} [\psi_i^{\mathbf{x}}, \varphi_j]^2 \right)^2} \\
&\leq C \sum_{i=D_b+1}^D \sum_{j=1}^{R_1} [\psi_i^{\mathbf{x}}, \varphi_j]^2 = C \sum_{i=D_b+1}^D [\psi_i^{\mathbf{x}}, \mathcal{K}\psi_i^{\mathbf{x}}] \leq C,
\end{aligned} \tag{B.24}$$

where the last equality follows by (B.5) and we use Assumption 2 in the last inequality. Similarly, we get

$$\|\widehat{\xi}(\mathbf{x})\|_{HS} \leq C \sum_{i=1}^{D_b} [\psi_i^{\mathbf{x}}, \mathcal{K}\psi_i^{\mathbf{x}}] \leq C. \tag{B.25}$$

Hence, using (B.24), (B.25), and Proposition 1 of [34] (see also Proposition 6 of [14]), we conclude (B.23). \square

The following proposition gives an upper bound for $\mathcal{E}(\mathcal{H}_a, \mathbf{r})$.

Proposition B.4. Suppose that Assumptions 1 and 2 hold. Let $\mathcal{E}(\mathcal{H}_a, \mathbf{r})$ be as in (B.19). Let $M_{\partial\Omega} = M - M_\Omega$ and let $N_{\partial\Omega} = N - N_\Omega$. Then, for any $\mathbf{r} := (r_1, r_2) \in [R_1] \times [R_2]$, with a probability at least $1 - \delta$, $0 < \delta < 1$, there exists a constant C such that

$$\mathcal{E}(\mathcal{H}_a, \mathbf{r}) \leq \max \left\{ \frac{CN_\Omega^2 \ln^2(2/\delta)}{\lambda_{r_1} M_\Omega}, \frac{CN_{\partial\Omega}^2 \ln^2(2/\delta)}{\tau_{r_2} M_{\partial\Omega}} \right\}. \tag{B.26}$$

Proof. The proof is an adaption of the arguments of Theorem 7 in [14]. Let $\{\lambda_i\}_{i=1}^{R_1}$, $\{\tau_i\}_{i=1}^{R_2}$, $\{\varphi_i\}_{i=1}^{R_1}$, and $\{\zeta_i\}_{i=1}^{R_2}$ be the eigenvalues and the eigenfunctions in Proposition B.1. Given $\delta \in [0, 1]$ and $\mathbf{r} := (r_1, r_2) \in [R_1] \times [R_2]$, we define

$$\begin{aligned}
\mathcal{H}_{c,\delta}^{\mathbf{r}} &= \left\{ h = \sum_{i=1}^{r_1} c_{1,i} \sqrt{\lambda_i} \varphi_i, \frac{1}{N_\Omega^2} \sum_{i=1}^{r_1} c_{1,i}^2 \lambda_i^2 \leq \delta \right\} \text{ and } \overline{\mathcal{H}}_{c,\delta}^{\mathbf{r}} = \left\{ h = \sum_{i=1}^{r_2} c_{2,i} \sqrt{\tau_i} \zeta_i, \frac{1}{N_{\partial\Omega}^2} \sum_{i=1}^{r_2} c_{2,i}^2 \tau_i^2 \leq 1 - \delta \right\}, \\
\mathcal{H}_{d,\delta}^{\mathbf{r}} &= \{v \in \mathcal{U}, \|v\|_{\mathcal{U}}^2 \leq \delta N_\Omega^2 / \lambda_{r_1}\}, \text{ and } \overline{\mathcal{H}}_{d,\delta}^{\mathbf{r}} = \{v \in \mathcal{U}, \|v\|_{\mathcal{U}}^2 \leq (1 - \delta) N_{\partial\Omega}^2 / \tau_{r_2}\}.
\end{aligned}$$

Thus, $\mathcal{H}_{c,\delta}^{\mathbf{r}} \subset \mathcal{H}_{d,\delta}^{\mathbf{r}}$ and $\overline{\mathcal{H}}_{c,\delta}^{\mathbf{r}} \subset \overline{\mathcal{H}}_{d,\delta}^{\mathbf{r}}$. Meanwhile, for $h_1 \in \mathcal{H}_{c,\delta}^{\mathbf{r}}$ and $h_2 \in \overline{\mathcal{H}}_{c,\delta}^{\mathbf{r}}$, using Proposition B.1, we have

$$\begin{aligned}
L_{\overline{\mathbf{x}}_1}[h_1] + L_{\overline{\mathbf{x}}_2}[h_2] &= \frac{1}{N_\Omega} \sum_{i=1}^{r_1} c_{1,i} \sqrt{\lambda_i} (\mathcal{K}\psi_\Omega)^T [\psi_\Omega, \varphi_i] + \frac{1}{N_{\partial\Omega}} \sum_{i=1}^{r_2} c_{2,i} \sqrt{\tau_i} (\mathcal{K}\psi_{\partial\Omega})^T [\psi_{\partial\Omega}, \zeta_i] \\
&= \frac{1}{N_\Omega} \sum_{i=1}^{r_1} c_{1,i} \lambda_i (\mathcal{K}\psi_\Omega)^T \mathbf{v}_i + \frac{1}{N_{\partial\Omega}} \sum_{i=1}^{r_2} c_{2,i} \tau_i (\mathcal{K}\psi_{\partial\Omega})^T \mathbf{w}_i \\
&= \frac{1}{N_\Omega} \sum_{i=1}^{r_1} c_{1,i} \lambda_i \sqrt{\lambda_i} \varphi_i + \frac{1}{N_{\partial\Omega}} \sum_{i=1}^{r_2} c_{2,i} \tau_i \sqrt{\tau_i} \zeta_i.
\end{aligned} \tag{B.27}$$

Thus, (B.27) implies that for any $h \in \mathcal{H}_b^r$, there exist $\delta \in [0, 1]$, $h_1 \in \mathcal{H}_{c,\delta}^r$, and $h_2 \in \overline{\mathcal{H}}_{c,\delta}^r$ such that $h = L_{\overline{\mathbf{x}}_1}[h_1] + L_{\overline{\mathbf{x}}_2}[h_2]$. Hence, we have

$$\begin{aligned} \mathcal{E}(\mathcal{H}_a, r) &= \max_{h \in \mathcal{H}_b^r} \mathcal{E}(g, \mathcal{H}_a) = \max_{\substack{h_1 \in \mathcal{H}_{c,\delta}^r \\ h_2 \in \overline{\mathcal{H}}_{c,\delta}^r}} \min_{v \in \mathcal{H}_a} \|L_{\overline{\mathbf{x}}_1}[h_1] + L_{\overline{\mathbf{x}}_2}[h_2] - v\|_{\mathcal{U}}^2 \\ &\leq \max_{\substack{h_1 \in \mathcal{H}_{d,\delta}^r \\ h_2 \in \overline{\mathcal{H}}_{d,\delta}^r}} \min_{v \in \mathcal{H}_a} \|L_{\overline{\mathbf{x}}_1}[h_1] + L_{\overline{\mathbf{x}}_2}[h_2] - v\|_{\mathcal{U}}^2. \end{aligned}$$

By the definitions of $L_{\widehat{\mathbf{x}}_1}$ and $L_{\widehat{\mathbf{x}}_2}$, we have $L_{\widehat{\mathbf{x}}_1}[h_1] + L_{\widehat{\mathbf{x}}_2}[h_2] \in \mathcal{H}_a$. Thus, we get

$$\begin{aligned} \mathcal{E}(\mathcal{H}_a, r) &\leq \max_{\substack{h_1 \in \mathcal{H}_{d,\delta}^r \\ h_2 \in \overline{\mathcal{H}}_{d,\delta}^r}} \min_{v \in \mathcal{H}_a} \|L_{\overline{\mathbf{x}}_1}[h_1] + L_{\overline{\mathbf{x}}_2}[h_2] - v\|_{\mathcal{U}}^2 \\ &\leq \max_{\substack{h_1 \in \mathcal{H}_{d,\delta}^r \\ h_2 \in \overline{\mathcal{H}}_{d,\delta}^r}} \|L_{\overline{\mathbf{x}}_1}[h_1] + L_{\overline{\mathbf{x}}_2}[h_2] - L_{\widehat{\mathbf{x}}_1}[h_1] - L_{\widehat{\mathbf{x}}_2}[h_2]\|_{\mathcal{U}}^2 \\ &\leq \max_{\delta \in [0,1]} \left\{ \frac{2\delta N_{\Omega}^2}{\lambda_{r_1}} \|L_{\overline{\mathbf{x}}_1} - L_{\widehat{\mathbf{x}}_1}\|_2^2 + \frac{2(1-\delta)N_{\partial\Omega}^2}{\tau_{r_2}} \|L_{\overline{\mathbf{x}}_2} - L_{\widehat{\mathbf{x}}_2}\|_2^2 \right\} \\ &\leq \max \left\{ \frac{2N_{\Omega}^2}{\lambda_{r_1}} \|L_{\overline{\mathbf{x}}_1} - L_{\widehat{\mathbf{x}}_1}\|_{HS}^2, \frac{2N_{\partial\Omega}^2}{\tau_{r_2}} \|L_{\overline{\mathbf{x}}_2} - L_{\widehat{\mathbf{x}}_2}\|_{HS}^2 \right\}. \end{aligned} \tag{B.28}$$

Thus, (B.28) together with Proposition B.3 implies that (B.26) holds. \square

Finally, we are ready to give a bound for $\mathcal{E}(\mathcal{H}_a)$.

Proof of Theorem 3.10. The inequality (3.27) follows directly from Propositions B.2 and B.4. \square