MIT Open Access Articles

*Gaussian functional regression for
linear partial differential equations*

The MIT Faculty has made this article openly available. **Please share**
how this access benefits you. Your story matters.

Citation: Nguyen, N.C., and J. Peraire. "Gaussian Functional Regression for Linear Partial Differential Equations." Computer Methods in Applied Mechanics and Engineering 287 (2015): 69–89.

As Published: <http://dx.doi.org/10.1016/j.cma.2015.01.008>

Publisher: Elsevier

Persistent URL: <http://hdl.handle.net/1721.1/107776>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License

Gaussian functional regression for state estimation with linear partial differential equations

N. C. Nguyen^{*} and J. Peraire[†]

Abstract

In this paper, we present a new statistical approach to the problem of incorporating experimental observations into a mathematical model described by linear partial differential equations (PDEs) to improve the prediction of the state of a physical system. We augment the linear PDE with a functional that accounts for the uncertainty in the mathematical model and is modeled as a *Gaussian process*. This gives rise to a stochastic PDE which is characterized by the Gaussian functional. We develop a *Gaussian functional regression* method to determine the posterior mean and covariance of the Gaussian functional, thereby solving the stochastic PDE to obtain the posterior distribution for our prediction of the physical state. Our method has the following features which distinguish itself from other regression methods. First, it incorporates both the mathematical model and the observations into the regression procedure. Second, it can handle the observations given in the form of linear functionals of the field variable. Third, the method is non-parametric in the sense that it provides a systematic way to optimally determine the prior covariance operator of the Gaussian functional based on the observations. Fourth, it provides the posterior distribution quantifying the magnitude of uncertainty in our prediction of the physical state. We present numerical results to illustrate these features of the method and compare its performance to that of the standard Gaussian process regression.

^{*}MIT Department of Aeronautics and Astronautics, 77 Massachusetts Ave., Cambridge, MA 02139, USA,
email: cuongng@mit.edu

[†]MIT Department of Aeronautics and Astronautics, 77 Massachusetts Ave., Cambridge, MA 02139, USA,
email: peraire@mit.edu

1 Introduction

Partial differential equations (PDEs) are used to mathematically model a wide variety of physical phenomena such as heat transfer, fluid flows, electromagnetism, and structural deformations. A PDE model of a physical system is typically described by conservation laws, constitutive laws, material properties, boundary conditions, boundary data, and geometry. In practical applications, the mathematical model described by the PDEs is only an approximation to the real physical system due to *(i)* the deliberate simplification of the mathematical model to keep it tractable (by ignoring certain physics or certain boundary conditions that pose computational difficulties), and *(ii)* the uncertainty of the available data (by using geometry, material property and boundary data that are not exactly the same as those of the physical system). We refer to the PDE model (available to us) as the *best knowledge PDE model* [24] and to its solution as the *best knowledge state*. To assess the accuracy of the best knowledge model in predicting the physical system, the best knowledge state needs to be compared against experimental data, which typically will have some level of noise.

In cases where the discrepancy between the PDE model and the experimental data is beyond an acceptable level of accuracy, we need to improve the current PDE model. There are several approaches to defining a new improved model. *Parameter estimation* [20, 25] involves calibrating some parameters in the model to match the data. An alternative approach to obtain an improved model is *data assimilation* [8, 9, 10, 11, 19]. Broadly speaking, data assimilation is a numerical procedure by which we incorporate observations into a mathematical model to reflect the errors inherent in our mathematical modeling of the physical system. Although data assimilation shares the same objective as parameter estimation, it differs from the latter in methodology. More specifically, data assimilation does not assume any parameters to be calibrated; instead, data assimilation defines a new model that matches the observations as well as possible, while being as close as possible to the best knowledge model. Another approach is *data interpolation* [2, 6, 12, 13, 15, 22] which involves computing a collection of solutions (snapshots) of a parametrized or time-varying mathematical model and reconstructing the physical state by fitting the experimental data to the snapshots.

A widely used technique for obtaining an improved model in parameter estimation and data assimilation is *least squares regression* [4, 9, 23]. Least squares is a deterministic regression approach that provides an estimate for the physical state which is optimal in least squares sense. However, it does not provide a means to quantify the prediction uncertainty. A recent work [24] poses the least-square regression as a regularized saddle point Galerkin formulation which admits interpretation from a variational framework and permits its extension to *Petrov-Galerkin formulation*. While the Petrov-Galerkin formulation provides more flexibility than the Galerkin formulation, it does not quantify the uncertainty in the prediction either. A popular statistical approach in parameter estimation and data assimilation is *Bayesian inference* [5, 7, 18, 21]. In Bayesian inference an estimate of the physical state is described by random variables and the posterior probability distribution of the estimate is determined by the data according to Bayes' rule [7, 18]. Therefore, Bayesian inference provides a powerful framework to quantify the prediction uncertainties.

In this paper, we introduce a new statistical approach to the problem of incorporating

observations into the best knowledge model to predict the state of a physical system. Our approach has its root in Gaussian process (GP) regression [14, 16, 17]. We augment the linear PDE with a functional that accounts for the uncertainty in the mathematical model and is modeled as a *Gaussian process*.¹ This gives rise to a stochastic PDE whose solutions are characterized by the Gaussian functional. By extending the standard GP regression for *functions of vectors* to *functionals of functions*, we develop a *Gaussian functional regression* method to determine the posterior distribution of the Gaussian functional, thereby solving the stochastic PDE for our prediction of the physical state. Our method is devised as follows. We first derive a *functional regression problem* by making use of the adjoint states and the observations. We next solve the functional regression problem by an application of the principle of Gaussian processes to obtain the posterior mean and covariance of the Gaussian functional. Finally, we compute the posterior distribution for our estimate of the physical state. A crucial ingredient in our method is the *covariance operator* representing the *prior* of the Gaussian functional. The bilinear covariance operators considered incorporate a number of free parameters (the so-called *hyperparameters*) that can be optimally determined from the measured data by maximizing a marginal likelihood.

Our Gaussian functional regression method can be viewed as a generalization of the standard GP regression from a finite dimensional vector (input) space to an infinite dimensional function (input) space. GP regression is a well-established technique to construct maps between inputs and outputs based on a set of sample, or training, input and output pairs, but does not offer a direct method to incorporate prior knowledge, albeit approximate, from an existing model relating the inputs and outputs. By combining the best knowledge model with the data, our method can greatly improve the prediction of the physical system.

Furthermore, we introduce a *nonparametric Bayesian inference* method for linear functional regression with Gaussian noise. It turns out that nonparametric Bayesian inference and Gaussian functional regression represent two different views of the same procedure. Specifically, we can think of Gaussian functional regression as defining a distribution over functionals and doing inference in the space of functionals — the *functional-space view*. We can think of nonparametric Bayesian inference as defining a distribution over weights and doing inference in the space of weights — the *weight-space view*. Theoretically, Gaussian functional regression can be interpreted as an application of the *kernel trick* [16] to nonparametric Bayesian inference, thereby avoiding an explicit construction of the feature map.

The paper is organized as follows. In Section 2, we present a description of the problem considered. In Section 3, we give an overview of Gaussian processes for regression problems. In Section 4, we introduce Gaussian functional regression for state estimation problems via linear PDE models. In Section 5 and Section 6, we present numerical results to demonstrate our method and compare its performance to that of function Gaussian process regression. In Section 7, we provide some concluding remarks on future research. Finally, in the Appendix, we describe our nonparametric Bayesian inference method.

¹In the cases considered, the physical system is not stochastic but deterministic. The introduction of the Gaussian functional serves to represent uncertainties in the best knowledge model and in the data, *not* in the physical system *per se*.

2 Motivation and Problem Statement

Let $\Omega^{\text{true}} \in \mathbb{R}^n$ denote a bounded open domain with Lipschitz boundary. Let $V^{\text{true}}(\Omega^{\text{true}})$ be an appropriate real-valued function space in which the true state u^{true} resides. A weak formulation of the true PDE model can be stated as: Find $u^{\text{true}} \in V^{\text{true}}$ and $\mathbf{s}^{\text{true}} \in \mathbb{R}^M$ such that

$$a^{\text{true}}(u^{\text{true}}, v) = \ell^{\text{true}}(v), \quad \forall v \in V^{\text{true}}, \quad (1a)$$

$$s_i^{\text{true}} = c_i^{\text{true}}(u^{\text{true}}), \quad i = 1, \dots, M, \quad (1b)$$

where $a^{\text{true}} : V^{\text{true}} \times V^{\text{true}} \rightarrow \mathbb{R}$ is a bilinear form, $\ell^{\text{true}} : V^{\text{true}} \rightarrow \mathbb{R}$ is a linear functional, and $c_i^{\text{true}} : V^{\text{true}} \rightarrow \mathbb{R}, i = 1, \dots, M$ are observation functionals. We assume that the true PDE model (1) is well defined and accurately describes the physical system of interest.

In actual practice, we do not have access to $a^{\text{true}}, \ell^{\text{true}}, c_i^{\text{true}}$, and $V^{\text{true}}(\Omega^{\text{true}})$. Hence, we can not compute u^{true} and $\mathbf{s}^{\text{true}}$. However, we assume that we have access to the “best knowledge” of $a^{\text{true}}, \ell^{\text{true}}, c_i^{\text{true}}, i = 1, \dots, M$, and $V^{\text{true}}(\Omega^{\text{true}})$, which shall be denoted by $a^{\text{bk}}, \ell^{\text{bk}}, c_i^{\text{bk}}, i = 1, \dots, M$, and $V^{\text{bk}}(\Omega^{\text{bk}})$, respectively. We then define the best knowledge PDE model: Find $u^{\text{bk}} \in V^{\text{bk}}$ and $\mathbf{s}^{\text{bk}} \in \mathbb{R}^M$ such that

$$a^{\text{bk}}(u^{\text{bk}}, v) = \ell^{\text{bk}}(v), \quad \forall v \in V^{\text{bk}}, \quad (2a)$$

$$s_i^{\text{bk}} = c_i^{\text{bk}}(u^{\text{bk}}), \quad i = 1, \dots, M. \quad (2b)$$

In the remainder of this paper, we shall drop the superscript “bk” for the quantities associated with the best knowledge model to simplify the notation. (In practice, we replace the continuous function space $V(\Omega)$ with a finite approximation space, which is assumed to be large enough that the numerical solution is indistinguishable from the continuous one.)

We now assume that we are given the observed data $\mathbf{d} \in \mathbb{R}^M$, which are the M measurements of the true output vector $\mathbf{s}^{\text{true}}$. We further assume that the measurements differ from the true outputs $\mathbf{s}^{\text{true}}$ by additive Gaussian noise $\boldsymbol{\varepsilon}$, namely,

$$\mathbf{d} = \mathbf{s}^{\text{true}} + \boldsymbol{\varepsilon}, \quad (3)$$

where $\varepsilon_i, i = 1, \dots, M$ are independent, identically distributed Gaussian distributions with zero mean and variance σ^2 . If σ is sufficiently small within the acceptable accuracy then we can use the observed data $\mathbf{d}$ to validate the best knowledge model (2). If the best knowledge outputs $\mathbf{s}$ are close enough to $\mathbf{d}$ within the noise level then we may trust the best knowledge model to predict the behavior of the true model. In many cases, the best knowledge outputs do not match the observed data due to various sources of uncertainty from physical modeling, constitutive laws, boundary conditions, boundary data, material properties, and geometry.

We are interested in improving the best knowledge model when it does not produce a good estimate of the true state. In particular, we propose a method to compute a better estimate for the true state by combining the best knowledge model with the observed data. Our method has its root in Gaussian process regression. Before proceeding to describe the proposed method we review the ideas behind Gaussian processes.

3 Gaussian Process Regression

We begin by assuming that we are given a training set of M observations

$$S = \{(\mathbf{x}_i, y_i), i = 1, \dots, M\}, \quad (4)$$

where $\mathbf{x}_i \in \mathbb{R}^N$ denotes an input vector of dimension N and $y_i \in \mathbb{R}$ denotes a scalar real-valued output. The training input vectors $\mathbf{x}_i, i = 1, \dots, M$ are aggregated in the $N \times M$ real-valued matrix $\mathbf{X}$, and the outputs are collected in the real-valued vector $\mathbf{y}$, so we can write $S = (\mathbf{X}, \mathbf{y})$. We assume that

$$y_i = h(\mathbf{x}_i) + \mathcal{N}(0, \sigma^2), \quad (5)$$

where $h(\mathbf{x})$ is the true but unknown function which we want to infer. The unknown function $h(\mathbf{x})$ is modeled as a Gaussian process² with zero mean³, for simplicity, and covariance function $\kappa(\mathbf{x}, \mathbf{x}')$, namely,

$$h(\mathbf{x}) \sim \mathcal{GP}(0, \kappa(\mathbf{x}, \mathbf{x}')) . \quad (6)$$

From a Bayesian perspective, we encode our belief that instances of $h(\mathbf{x})$ are drawn from a Gaussian process with zero mean and covariance function $\kappa(\mathbf{x}, \mathbf{x}')$ *prior to taking into account observations*.⁴ Mathematically speaking, $h(\mathbf{x})$ is assumed to reside in a reproducing kernel Hilbert space [1] spanned by the eigenfunctions of the covariance function $\kappa(\mathbf{x}, \mathbf{x}')$.

Let $\mathbf{X}^*$ be the $N \times M^*$ matrix that contains M^* test input vectors $\mathbf{x}_j^*, j = 1, \dots, M^*$ as its columns. Since $h(\mathbf{x}) \sim \mathcal{GP}(0, \kappa(\mathbf{x}, \mathbf{x}'))$, the joint distribution of the observed outputs and the function values at the test input vectors is

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{h}^* \end{bmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} \mathcal{K}(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I} & \mathcal{K}(\mathbf{X}, \mathbf{X}^*) \\ \mathcal{K}(\mathbf{X}^*, \mathbf{X}) & \mathcal{K}(\mathbf{X}^*, \mathbf{X}^*) \end{bmatrix} \right), \quad (7)$$

where $\mathbf{h}^* \in \mathbb{R}^{M^*}$, $\mathcal{K}(\mathbf{X}, \mathbf{X}) \in \mathbb{R}^{M \times M}$, $\mathcal{K}(\mathbf{X}, \mathbf{X}^*) \in \mathbb{R}^{M \times M^*}$, $\mathcal{K}(\mathbf{X}^*, \mathbf{X}) \in \mathbb{R}^{M^* \times M}$, and $\mathcal{K}(\mathbf{X}^*, \mathbf{X}^*) \in \mathbb{R}^{M^* \times M^*}$ have entries

$$\begin{aligned} h_i^* &= h(\mathbf{x}_i^*), & i &= 1, \dots, M^*, \\ \mathcal{K}_{ij}(\mathbf{X}, \mathbf{X}) &= \kappa(\mathbf{x}_i, \mathbf{x}_j), & i &= 1, \dots, M, j = 1, \dots, M, \\ \mathcal{K}_{ij}(\mathbf{X}, \mathbf{X}^*) &= \kappa(\mathbf{x}_i, \mathbf{x}_j^*), & i &= 1, \dots, M, j = 1, \dots, M^*, \\ \mathcal{K}_{ij}(\mathbf{X}^*, \mathbf{X}) &= \kappa(\mathbf{x}_i^*, \mathbf{x}_j), & i &= 1, \dots, M^*, j = 1, \dots, M, \\ \mathcal{K}_{ij}(\mathbf{X}^*, \mathbf{X}^*) &= \kappa(\mathbf{x}_i^*, \mathbf{x}_j^*), & i &= 1, \dots, M^*, j = 1, \dots, M^*, \end{aligned} \quad (8)$$

²A Gaussian process is a generalization of the Gaussian probability distribution. A Gaussian process governs the properties of Gaussian random functions, whereas a Gaussian probability distribution describes the properties of Gaussian random variables (scalars or vectors).

³The Gaussian process is assumed to have zero mean because we can always subtract the original outputs $\mathbf{y}$ from its average $\bar{y} = \frac{1}{M} \sum_{i=1}^M y_i$ to obtain new outputs with zero average. We then work with the new outputs and add the average $\bar{y}$ to our Gaussian process estimator.

⁴At this point, one may ask what if our belief is wrong, that is, what if the covariance function κ is not correctly chosen. Of course, choosing a wrong covariance function will result in very poor prediction. Hence, the covariance function should not be chosen arbitrarily. As discussed later, Gaussian processes provide a framework for optimal selection of a covariance function based on the observed data.

respectively. We next apply the conditional distribution formula (see [16]) to the joint distribution (7) to obtain the *predictive distribution* for $\mathbf{h}^*$ as

$$\mathbf{h}^*|\mathbf{y}, \mathbf{X}, \mathbf{X}^* \sim \mathcal{N}(\bar{\mathbf{h}}^*, \text{cov}(\mathbf{h}^*)), \quad (9)$$

where

$$\begin{aligned} \bar{\mathbf{h}}^* &= \mathcal{K}(\mathbf{X}^*, \mathbf{X})\boldsymbol{\alpha}, \\ \text{cov}(\mathbf{h}^*) &= \mathcal{K}(\mathbf{X}^*, \mathbf{X}^*) - \mathcal{K}(\mathbf{X}^*, \mathbf{X})\mathbf{C}^{-1}\mathcal{K}(\mathbf{X}, \mathbf{X}^*), \end{aligned} \quad (10)$$

and $\boldsymbol{\alpha} \in \mathbb{R}^M$ and $\mathbf{C} \in \mathbb{R}^{M \times M}$ are given by

$$\mathbf{C}\boldsymbol{\alpha} = \mathbf{y}, \quad \mathbf{C} = \mathcal{K}(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I}. \quad (11)$$

Note that the predictive mean $\bar{\mathbf{h}}^*$ is a linear combination of M kernel functions, each one centered on a training input vector. Note also that the predictive covariance $\text{cov}(\mathbf{h}^*)$ does not explicitly depend on the observed data $\mathbf{y}$, but only on the training input vectors $\mathbf{X}$ and the covariance function κ . This is a property of the Gaussian distribution. However, as discussed below, the covariance function κ can be determined by using the observed data. As a result, the predictive covariance implicitly depends on the observed data.

Typically, the covariance function κ has some free parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_Q)$, so that the matrix $\mathbf{C}$ depends on $\boldsymbol{\theta}$. These free parameters are called *hyperparameters*. The hyperparameters have a significant impact on the predictive mean and covariance. They are determined by maximizing the log marginal likelihood (see [16]):

$$\log p(\mathbf{y}|\mathbf{X}, \boldsymbol{\theta}) = -\frac{1}{2}\mathbf{y}^T \mathbf{C}(\boldsymbol{\theta})^{-1}\mathbf{y} - \frac{1}{2}\log(\det(\mathbf{C}(\boldsymbol{\theta}))) - \frac{M}{2}\log(2\pi). \quad (12)$$

Once we choose a specific form for κ and determine its hyperparameters, we can compute $\bar{\mathbf{h}}^*$ and $\text{cov}(\mathbf{h}^*)$ for any given $\mathbf{X}^*$. Gaussian processes also provide us a mean to choose an appropriate family among many possible families of covariance functions. Choosing a covariance function for a particular application involves both determining hyperparameters within a family and comparing across different families. This step is termed as *model selection* [16].

We see that the standard GP regression provides us not only a posterior mean, but also a posterior covariance which characterizes uncertainty in our prediction of the true function. Moreover, it allows us to determine the optimal covariance function and thus the optimal reproducing Kernel Hilbert space in which the true function is believed to reside. These features differentiate GP regression from parametric regression methods such as least-squares regression, which typically provides the maximum likelihood estimate only. However, GP regression tends to require larger sample sizes than parametric regression methods because the data must supply enough information to yield a good covariance function by using model selection.

There are a number of obstacles that prevent us from applying the standard GP regression to our problem of interest described in the previous section. First, our outputs are in general

not the evaluations of the state at spatial coordinates. Instead, they are linear functionals of the state. Second, the standard GP regression described here does not allow us to make use of the best knowledge model. The best knowledge model plays an important role because it carries crucial prior information about the true model. By taking advantage of the best knowledge model, we may be able to use far less observations to obtain a good prediction and thus address the main disadvantage of the standard GP regression. We propose a new approach to overcome these obstacles.

4 Gaussian Functional Regression

4.1 A stochastic PDE model

Let $g : V \rightarrow \mathbb{R}$ be a linear functional. We introduce a new mathematical model: Find $u^* \in V$ and $\mathbf{s}^* \in \mathbb{R}^M$ such that

$$a(u^*, v) + g(v) = \ell(v), \quad \forall v \in V, \quad (13a)$$

$$s_i^* = c_i(u^*), \quad i = 1, \dots, M. \quad (13b)$$

Notice that the new model (13) differs from the best knowledge model (2) by the functional g . We can determine u^* and $\mathbf{s}^*$ only if g is known. The functional g thus characterizes the solution u^* and the output vector $\mathbf{s}^*$ of the model (13). We note that if $g(v) = \ell(v) - a(u^{\text{true}}, v)$ then we $u^* = u^{\text{true}}$. Unfortunately, this particular choice of g requires the true state u^{true} which we do not know and thus want to infer.

In order to capture various sources of uncertainty in the best knowledge model, we represent g as a *Gaussian functional*⁵ with zero mean and covariance operator k , namely,

$$g(v) \sim \mathcal{GF}(0, k(v, v')), \quad \forall v, v' \in V. \quad (14)$$

Notice that there are three main differences between the Gaussian functional (14) and the Gaussian process (6). First, g is a functional, whereas h is a function. Second, v is a function, whereas $\mathbf{x}$ is a vector. And third, $k(v, v')$ is generally a differential and integral operator of v and v' , whereas $\kappa(\mathbf{x}, \mathbf{x}')$ is a function of $\mathbf{x}$ and $\mathbf{x}'$. We will require that the covariance operator $k : V \times V \rightarrow \mathbb{R}$ is symmetric positive-definite. That is,

$$k(v, v') = k(v', v), \quad \text{and} \quad k(v, v) > 0, \quad \forall v, v' \in V. \quad (15)$$

As the covariance operator k characterizes the space of all possible functionals prior to taking into account the observations, it plays an important role in our method. The selection of a covariance operator will be discussed later.

Since g is a Gaussian functional the model (13) becomes a stochastic PDE. In order to solve the stochastic PDE (13), we need to compute the *posterior mean* and *posterior covariance* of g after accounting for the observed data $\mathbf{d}$. To this end, we formulate a functional regression problem and describe a procedure for solving it as follows.

⁵A Gaussian functional can be thought of as a generalization of the Gaussian process from a finite dimensional vector space to an infinite dimensional function space.

4.2 Functional regression problem

We first introduce the adjoint problems: for $i = 1, \dots, M$ we find $\phi_i \in V$ such that

$$a(v, \phi_i) = -c_i(v), \quad \forall v \in V. \quad (16)$$

We note that the adjoint states ϕ_i depend only on the output functionals c_i and the bilinear form a . It follows from (2), (13), and (16) that

$$g(\phi_i) = \ell(\phi_i) - a(u^*, \phi_i) = a(u, \phi_i) + c_i(u^*) = c_i(u^*) - c_i(u) = s_i^* - s_i, \quad (17)$$

for $i = 1, \dots, M$. Moreover, we would like our stochastic PDE to produce the outputs $\mathbf{s}^*$ that are consistent with the observed data $\mathbf{d}$ in such a way that

$$d_i = s_i^* + \mathcal{N}(0, \sigma^2), \quad i = 1, \dots, M. \quad (18)$$

This equation is analogous to (3) which relates the observed data $\mathbf{d}$ to the true outputs $\mathbf{s}^{\text{true}}$. We substitute $s_i^* = d_i - \mathcal{N}(0, \sigma^2)$ into (17) to obtain

$$d_i - s_i = g(\phi_i) + \mathcal{N}(0, \sigma^2), \quad i = 1, \dots, M. \quad (19)$$

Notice that this expression characterizes the relationship between $g(\phi_i)$ and $d_i - s_i$ in the same way (5) characterizes the relationship between $h(\mathbf{x}_i)$ and y_i .

We now introduce a training set of M observations

$$T = \{(\phi_i, d_i - s_i), \quad i = 1, \dots, M\}, \quad (20)$$

and use this training set to learn about g . More specifically, we wish to determine $g(\phi^*)$ for any given $\phi^* \in V$ based on the training set T . This problem is similar to the regression problem described in the previous section and is named the *functional regression problem* to emphasize that the object of interest g is a functional. We next describe the solution of the functional regression problem.

4.3 Regression procedure

Let $\Phi = [\phi_1, \dots, \phi_M]$ be a collection of M adjoint states as determined by (16). Let $\Phi^* = [\phi_1^* \in V, \dots, \phi_M^* \in V]$ be a collection of M^* test functions. The joint distribution of the observed outputs and the functional values for the test functions according to the prior (14) is given by

$$\begin{bmatrix} \mathbf{d} - \mathbf{s} \\ \mathbf{g}^* \end{bmatrix} \sim \mathcal{N} \left(0, \begin{bmatrix} \mathbf{K}(\Phi, \Phi) + \sigma^2 \mathbf{I} & \mathbf{K}(\Phi, \Phi^*) \\ \mathbf{K}(\Phi^*, \Phi) & \mathbf{K}(\Phi^*, \Phi^*) \end{bmatrix} \right), \quad (21)$$

where $\mathbf{g}^* \in \mathbb{R}^{M^*}$, $\mathbf{K}(\Phi, \Phi) \in \mathbb{R}^{M \times M}$, $\mathbf{K}(\Phi, \Phi^*) \in \mathbb{R}^{M \times M^*}$, $\mathbf{K}(\Phi^*, \Phi) \in \mathbb{R}^{M^* \times M}$, and $\mathbf{K}(\Phi^*, \Phi^*) \in \mathbb{R}^{M^* \times M^*}$ have entries

$$\begin{aligned} g_i^* &= g(\phi_i^*), & i &= 1, \dots, M^*, \\ K_{ij}(\Phi, \Phi) &= k(\phi_i, \phi_j), & i &= 1, \dots, M, j = 1, \dots, M, \\ K_{ij}(\Phi, \Phi^*) &= k(\phi_i, \phi_j^*), & i &= 1, \dots, M, j = 1, \dots, M^*, \\ K_{ij}(\Phi^*, \Phi) &= k(\phi_i^*, \phi_j), & i &= 1, \dots, M^*, j = 1, \dots, M, \\ K_{ij}(\Phi^*, \Phi^*) &= k(\phi_i^*, \phi_j^*), & i &= 1, \dots, M^*, j = 1, \dots, M^*, \end{aligned} \quad (22)$$

respectively. It thus follows that the *predictive distribution* for $\mathbf{g}^*$ is

$$\mathbf{g}^* | (\mathbf{d} - \mathbf{s}), \Phi, \Phi^* \sim \mathcal{N}(\bar{\mathbf{g}}^*, \text{cov}(\mathbf{g}^*)), \quad (23)$$

where

$$\bar{\mathbf{g}}^* = \mathbf{K}(\Phi^*, \Phi)\boldsymbol{\beta}, \quad \text{cov}(\mathbf{g}^*) = \mathbf{K}(\Phi^*, \Phi^*) - \mathbf{K}(\Phi^*, \Phi)\mathbf{D}^{-1}\mathbf{K}(\Phi, \Phi^*), \quad (24)$$

and $\boldsymbol{\beta} \in \mathbb{R}^M$ and $\mathbf{D} \in \mathbb{R}^{M \times M}$ are given by

$$\mathbf{D}\boldsymbol{\beta} = \mathbf{d} - \mathbf{s}, \quad \mathbf{D} = \mathbf{K}(\Phi, \Phi) + \sigma^2 \mathbf{I}. \quad (25)$$

Notice that we have correspondence with function Gaussian process regression described in the previous section, when identifying $(\Phi, \mathbf{d} - \mathbf{s})$ with $(\mathbf{X}, \mathbf{y})$, $(\Phi^*, \mathbf{g}^*)$ with $(\mathbf{X}^*, \mathbf{h}^*)$, and $k(\cdot, \cdot)$ with $\kappa(\cdot, \cdot)$.

While our approach share similarities with Gaussian process regression, it differs from the latter in many important ways. We summarize in Table 1 the differences between Gaussian process regression and Gaussian functional regression.

Quantities	Gaussian Process Regression	Gaussian Functional Regression
Input	vector $\mathbf{x} \in \mathbb{R}^N$	function $v \in V$
Output	function $h(\mathbf{x})$	functional $g(v)$
Prior	$h \sim \mathcal{N}(0, \kappa(\cdot, \cdot))$	$g \sim \mathcal{N}(0, k(\cdot, \cdot))$
Kernel	function $\kappa : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$	operator $k : V \times V \rightarrow \mathbb{R}$
Training inputs	$\mathbf{X} \in \mathbb{R}^{N \times M}$	$\Phi \in V^M$ adjoint states
Observations	$\mathbf{y} = h(\mathbf{X}) + \mathcal{N}(0, \sigma)$	$\mathbf{d} - \mathbf{s} = g(\Phi) + \mathcal{N}(0, \sigma)$
Coefficients	$\mathbf{C}\boldsymbol{\alpha} = \mathbf{y}, \mathbf{C} = [\mathbf{K}(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I}]$	$\mathbf{D}\boldsymbol{\beta} = \mathbf{d} - \mathbf{s}, \mathbf{D} = [\mathbf{K}(\Phi, \Phi) + \sigma^2 \mathbf{I}]$
Test inputs	$\mathbf{X}^* \in \mathbb{R}^{N \times M^*}$	$\Phi^* \in V^{M^*}$
Mean	$\mathbf{K}(\mathbf{X}^*, \mathbf{X})\boldsymbol{\alpha}$	$\mathbf{K}(\Phi^*, \Phi)\boldsymbol{\beta}$
Covariance	$\mathbf{K}(\mathbf{X}^*, \mathbf{X}^*) - \mathbf{K}(\mathbf{X}^*, \mathbf{X})\mathbf{C}^{-1}\mathbf{K}(\mathbf{X}, \mathbf{X}^*)$	$\mathbf{K}(\Phi^*, \Phi^*) - \mathbf{K}(\Phi^*, \Phi)\mathbf{D}^{-1}\mathbf{K}(\Phi, \Phi^*)$

Table 1: Comparison between Gaussian process regression and Gaussian functional regression. Note that the best knowledge model enters in the Gaussian functional regression through the adjoint states Φ and the best knowledge outputs $\mathbf{s}$.

In the Appendix A, we introduce a nonparametric Bayesian framework for linear functional regression with Gaussian noise. It turns out that this nonparametric Bayesian framework is equivalent to the Gaussian functional regression described here. In fact, Gaussian functional regression can be viewed as an application of the *kernel trick* to nonparametric Bayesian inference for linear functional regression, thereby avoiding the computation of the eigenfunctions of the covariance operator k . We next introduce a family of bilinear covariance operators and then describe a method for determining the hyperparameters.

4.4 Covariance operators

The covariance operator k is a crucial ingredient in our approach. Here, we consider a class of bilinear covariance operators parametrized by $\boldsymbol{\theta} = (\theta_1, \theta_2)$ of the form:

$$k(v, v'; \boldsymbol{\theta}) = \theta_1 \int_{\Omega} vv' d\mathbf{x} + \theta_2 \int_{\Omega} \nabla v \cdot \nabla v' d\mathbf{x} . \quad (26)$$

More general forms of the covariance operator are possible provided that they are symmetric and positive definite.

In order for a covariance operator to be used in our method, we need to specify its hyperparameters $\boldsymbol{\theta}$. Fortunately, Gaussian processes allow us to determine the hyperparameters by using the observed data. In order to do this, we first calculate the probability of the observed data given the hyperparameters, or *marginal likelihood* and choose $\boldsymbol{\theta}$ so that this likelihood is maximized. We note from (21) that

$$p(\mathbf{d} - \mathbf{s} | \Phi, \boldsymbol{\theta}) = \mathcal{N}(0, \mathbf{D}(\boldsymbol{\theta})), \quad (27)$$

where the matrix $\mathbf{D}(\boldsymbol{\theta})$ as defined in (25) depends on $k(\cdot, \cdot; \boldsymbol{\theta})$ and thus on $\boldsymbol{\theta}$ as well. Rather than maximizing (27), it is more convenient to maximize the log marginal likelihood which is given by,

$$\log p(\mathbf{d} - \mathbf{s} | \Phi, \boldsymbol{\theta}) = -\frac{1}{2}(\mathbf{d} - \mathbf{s})^T \mathbf{D}(\boldsymbol{\theta})^{-1}(\mathbf{d} - \mathbf{s}) - \frac{1}{2} \log(\det(\mathbf{D}(\boldsymbol{\theta}))) - \frac{M}{2} \log(2\pi). \quad (28)$$

Thus, we find $\boldsymbol{\theta}$ by solving the maximization problem

$$\boldsymbol{\theta} = \arg \max_{\boldsymbol{\theta}' \in \mathbb{R}^2} \log p(\mathbf{d} - \mathbf{s} | \Phi, \boldsymbol{\theta}'). \quad (29)$$

Hence, the hyperparameters $\boldsymbol{\theta}$ are chosen as the maximizer of the log marginal likelihood.

Once we determine the covariance operator, we can compute $(\bar{\mathbf{g}}^*, \text{cov}(\mathbf{g}^*))$ for any given set Φ^* of test functions as described in Subsection 4.3. Therefore, Gaussian functional regression is *non-parametric* in the sense that the hyperparameters are chosen in light of the observed data. In other words, the data are used to define both the prior covariance and the posterior covariance. In contrast, parametric regression methods use a number of parameters to define the prior and combine this prior with the data to determine the posterior prediction. It remains to describe how to compute the posterior mean and covariance of the solution u^* of the stochastic PDE.

4.5 Computation of the mean state and covariance

We recall that our stochastic PDE model consists of finding $u^* \in V$ such that

$$a(u^*, v) = \ell(v) - g(v), \quad \forall v \in V. \quad (30)$$

Let $\{v_j(\mathbf{x})\}_{j=1}^J$ be a “suitable” basis set of the function space $V(\Omega)$, where J is the dimension of $V(\Omega)$. Since the functional g is Gaussian and the best knowledge model is linear, we can express the solution of the stochastic PDE (30) as

$$u^*(\mathbf{x}) = \sum_{j=1}^J \gamma_j^* v_j(\mathbf{x}), \quad \boldsymbol{\gamma}^* \sim \mathcal{N}(\bar{\boldsymbol{\gamma}}^*, \text{cov}(\boldsymbol{\gamma}^*)). \quad (31)$$

In order to determine $\bar{\boldsymbol{\gamma}}^*$ and $\text{cov}(\boldsymbol{\gamma}^*)$, we choose $v = v_i, i = 1, \dots, J$ in (30) to arrive at the stochastic linear system:

$$\mathbf{A}\boldsymbol{\gamma}^* = \mathbf{l} - \mathbf{g}^*, \quad (32)$$

where $A_{ij} = a(v_i, v_j)$, $l_i = \ell(v_i)$ for $i, j = 1, \dots, J$, and $\mathbf{g}^* \sim \mathcal{N}(\bar{\mathbf{g}}^*, \text{cov}(\mathbf{g}^*))$ with

$$\bar{\mathbf{g}}^* = \mathbf{K}(\Phi^*, \Phi)\boldsymbol{\beta}, \quad \text{cov}(\mathbf{g}^*) = \mathbf{K}(\Phi^*, \Phi^*) - \mathbf{K}(\Phi^*, \Phi)\mathbf{D}^{-1}\mathbf{K}(\Phi, \Phi^*), \quad (33)$$

for $\Phi^* \equiv [v_1, v_2, \dots, v_J]$. It thus follows from (32) that

$$\bar{\boldsymbol{\gamma}}^* = \mathbf{A}^{-1}(\mathbf{l} - \bar{\mathbf{g}}^*), \quad \text{cov}(\boldsymbol{\gamma}^*) = \mathbf{A}^{-1}\text{cov}(\mathbf{g}^*)\mathbf{A}^{-T}, \quad (34)$$

as $\mathbf{g}^*$ is Gaussian and $\mathbf{A}$ is invertible.

Now let $\mathbf{x}_i \in \Omega, i = 1, \dots, N$ be spatial points at which we would like to evaluate the predictive mean and covariance of u^* . Let $\mathbf{V} \in \mathbb{R}^{N \times J}$ be a matrix with entries $V_{ij} = v_j(\mathbf{x}_i), i = 1, \dots, N, j = 1, \dots, J$. It then follows from (31) that

$$\mathbf{u}^* = \mathbf{V}\boldsymbol{\gamma}^*, \quad (35)$$

where $u_i^* = u^*(\mathbf{x}_i), i = 1, \dots, N$. It follows from (34) and (35) that

$$\mathbf{u}^* \sim \mathcal{N}(\bar{\mathbf{u}}^*, \text{cov}(\mathbf{u}^*)), \quad (36)$$

where

$$\bar{\mathbf{u}}^* = \mathbf{U}(\mathbf{l} - \bar{\mathbf{g}}^*), \quad \text{cov}(\mathbf{u}^*) = \mathbf{U}\text{cov}(\mathbf{g}^*)\mathbf{U}^T, \quad (37)$$

with $\mathbf{U} = \mathbf{V}\mathbf{A}^{-1}$. We examine the posterior distribution as given by (37). Note first that the posterior mean $\bar{\mathbf{u}}^*$ is the difference between two terms: the first term $\mathbf{u} = \mathbf{U}\mathbf{l}$ is simply the best knowledge state u evaluated at $\mathbf{x}_i, i = 1, \dots, N$; the second term $\mathbf{U}\bar{\mathbf{g}}^*$ is a correction term to the best knowledge state and is obtained by using our Gaussian functional regression. Note also that the posterior covariance is a quadratic form of $\mathbf{U}$ with the posterior covariance matrix $\text{cov}(\mathbf{g}^*)$, showing that the predictive uncertainty grows with the magnitude of $\mathbf{U}$. Hence, the predictive uncertainty depends on the inverse matrix $\mathbf{A}^{-1}$. The implementation of our method for computing the posterior distribution (37) is shown in Figure 1.

<p>Input: $a, \ell, \{c_i\}_{i=1}^M$ (best knowledge), $\{v_j\}_{j=1}^J$ (basis functions), $\mathbf{d}$ (observed data), σ (noise level), k (covariance operator), $\{\mathbf{x}_i \in \Omega\}_{i=1}^N$ (spatial coordinates)</p> <ol style="list-style-type: none"> 1. Compute $\mathbf{s}$ by solving the best knowledge model (2) 2. Compute the adjoint states $\{\phi_i\}_{i=1}^M$ by solving (16) 3. Compute $(\bar{\mathbf{g}}^*, \text{cov}(\mathbf{g}^*))$ in (33) by using Gaussian functional regression 4. Form $\mathbf{A}, \mathbf{l}, \mathbf{V}$ and solve $\mathbf{A}^T \mathbf{U}^T = \mathbf{V}^T$ to obtain $\mathbf{U}$ 5. Compute $\bar{\mathbf{u}}^* = \mathbf{U}(\mathbf{l} - \bar{\mathbf{g}}^*)$, $\text{cov}(\mathbf{u}^*) = \mathbf{U} \text{cov}(\mathbf{g}^*) \mathbf{U}^T$ <p>Return: $\bar{\mathbf{u}}^*$ (posterior mean), $\text{cov}(\mathbf{u}^*)$ (posterior covariance)</p>

Figure 1: Main algorithm for computing the posterior distribution of $\mathbf{u}^*$.

4.6 Relationship with least-squares regression

Here, we show an alternative approach to computing the posterior mean in (37) by solving a deterministic least-squares problem. We note that the posterior mean $\bar{\mathbf{u}}^* \in V$ satisfies

$$a(\bar{\mathbf{u}}^*, v) = \ell(v) - \bar{g}^*(v), \quad \forall v \in V. \quad (38)$$

Here, for any $v \in V$, $\bar{g}^*(v)$ is the posterior mean of $g(v)$ and given by

$$\bar{g}^*(v) = \sum_{j=1}^M \beta_j k(\phi_j, v), \quad \forall v \in V, \quad (39)$$

where the adjoint states $\phi_j, j = 1, \dots, M$ satisfy (16) and the coefficient vector β is the solution of (25). It thus follows that the mean state $\bar{\mathbf{u}}^* \in V$ satisfies

$$a(\bar{\mathbf{u}}^*, v) = \ell(v) - k(\bar{\mathbf{q}}^*, v), \quad \forall v \in V, \quad (40)$$

where $\bar{\mathbf{q}}^* = \sum_{i=1}^M \beta_i \phi_i$ is the weighted sum of the adjoint states. We then evaluate the mean outputs of the stochastic PDE model as

$$\bar{s}_i^* = c_i(\bar{\mathbf{u}}^*), \quad i = 1, \dots, M. \quad (41)$$

The following lemma sheds light on the relationship between the mean outputs and the observed data.

Lemma 1. *Assume that the covariance operator $k(\cdot, \cdot)$ is a bilinear form. We have that $\bar{\mathbf{s}}^* = \mathbf{d} - \sigma \beta$.*

Proof. We first note from the adjoint equation (16) and (40) that

$$\bar{s}_i^* = c_i(\bar{\mathbf{u}}^*) = -a(\bar{\mathbf{u}}^*, \phi_i) = \sum_{j=1}^M k(\phi_j, \phi_i) \beta_j - \ell(\phi_i), \quad i = 1, \dots, M. \quad (42)$$

We next recall that β satisfies

$$\sum_{j=1}^M (k(\phi_j, \phi_i) + \sigma \delta_{ij}) \beta_j = d_i - s_i, \quad i = 1, \dots, M, \quad (43)$$

where δ_{ij} is the Kronecker delta. Moreover, we obtain from the best knowledge model (2) and the adjoint equation (16) that

$$\ell(\phi_i) = a(u, \phi_i) = -c_i(u) = -s_i, \quad i = 1, \dots, M. \quad (44)$$

The desired result immediately follows from the above three equations. This completes the proof. $\square$

This lemma shows that the mean outputs differs from the observed data by the product of the noise level σ and the coefficient vector β . When the observed data is noise-free (namely, $\sigma = 0$) we have that the mean output vector is exactly equal to the observed data. Henceforth, whenever M is sufficiently large and σ is relatively small, we expect that our method will yield a much better estimate of the true state than the best knowledge model. The following theorem shows the optimality of the mean state.

Theorem 2. *Assume that the covariance operator $k(\cdot, \cdot)$ is a bilinear form. Then we have $(\bar{u}^*, \bar{q}^*, \beta) = (u^o, q^o, \gamma^o)$, where*

$$\begin{aligned} (u^o, q^o, \gamma^o) = \arg \min_{z \in V, q \in V, \gamma \in \mathbb{R}^M} & \frac{1}{2} k(q, q) + \frac{1}{2} \sigma^2 \gamma^T \gamma \\ \text{s.t. } & a(z, v) + k(q, v) = \ell(v), \quad \forall v \in V, \\ & c_i(z) + \sigma^2 \gamma_i = d_i, \quad i = 1, \dots, M. \end{aligned} \quad (45)$$

Proof. We introduce the Lagrangian

$$\mathcal{L}(q, z, \gamma, p, \varrho) = \frac{1}{2} k(q, q) + \frac{1}{2} \sigma^2 \gamma^T \gamma - a(z, p) - k(q, p) + \ell(p) - \sum_{i=1}^M \varrho_i (c_i(z) + \sigma^2 \gamma_i - d_i), \quad (46)$$

where $p \in V$ and $\varrho \in \mathbb{R}^M$ are the Lagrange multipliers of the constraints. The optimal solution $(q^o, u^o, \gamma^o, p^o, \varrho^o)$ satisfies

$$\frac{\partial \mathcal{L}}{\partial q} = \frac{\partial \mathcal{L}}{\partial z} = \frac{\partial \mathcal{L}}{\partial \gamma} = \frac{\partial \mathcal{L}}{\partial p} = \frac{\partial \mathcal{L}}{\partial \varrho} = 0, \quad (47)$$

which yields

$$k(q^o, v) - k(p^o, v) = 0, \quad \forall v \in V, \quad (48a)$$

$$a(v, p^o) + \sum_{i=1}^M \varrho_i^o c_i(v) = 0, \quad \forall v \in V, \quad (48b)$$

$$\gamma_i^o - \varrho_i^o = 0, \quad i = 1, \dots, M, \quad (48c)$$

$$a(u^o, v) + k(q^o, v) - \ell(v) = 0, \quad \forall v \in V, \quad (48d)$$

$$c_i(u^o) + \sigma^2 \gamma_i^o = d_i, \quad i = 1, \dots, M. \quad (48e)$$

Note that when taking the partial derivatives we have used the assumption that k is bilinear. The first two equations (48a) and (48b) yield that

$$q^o = p^o = \sum_{i=1}^M \varrho_i^o \phi_i, \quad (49)$$

where $\phi_i, i = 1, \dots, M$ are the adjoint states. Therefore, if we can show that $\varrho^o = \beta$ then (48c), (48d) and (49) imply that $(\bar{u}^*, \bar{q}^*, \beta) = (u^o, q^o, \gamma^o)$. To this end, we note from (48d) and the adjoint equation (16) that

$$c_i(u^o) = -a(u^o, \phi_i) = k(q^o, \phi_i) - \ell(\phi_i), \quad i = 1, \dots, M. \quad (50)$$

Moreover, we obtain from the best knowledge model (2) and the adjoint equation (16) that

$$\ell(\phi_i) = a(u, \phi_i) = -c_i(u) = -s_i, \quad i = 1, \dots, M. \quad (51)$$

Finally, it follows from (48c), (48e), (49), (50), and (51) that

$$\sum_{j=1}^m (k(\phi_i, \phi_j) + \sigma^2 \delta_{ij}) \varrho_j^o = d_i - s_i, \quad i = 1, \dots, M, \quad (52)$$

which implies that $\varrho^o = \beta$. This completes the proof. $\square$

This theorem establishes a connection between Gaussian functional regression and traditional least-squares regression when the covariance operator is bilinear. In particular, the posterior mean state $\bar{u}^*$ is the optimal solution of a least-squares minimization. This is hardly a surprise as the posterior mean state is also the maximum a posteriori (MAP) estimate of linear functional regression model in the Bayesian framework discussed in the Appendix A. It is well known that the MAP estimate coincides with the least-squares solution. The main advantage of our approach over least-squares regression is that we can compute not only the posterior mean state but also the posterior covariance. Another advantage of our approach is that it allows us to choose a covariance operator based on the observed data by exploiting the marginal likelihood function, whereas least-squares regression does not provide a mechanism to optimally set the prior covariance operator.

5 A Simple Heat Conduction Example

5.1 Problem description

For the true PDE model we consider a one-dimensional heat equation:

$$-\frac{\partial^2 u^{\text{true}}}{\partial x^2} = f^{\text{true}}, \quad \text{in } \Omega^{\text{true}} \equiv (-1, 1), \quad (53)$$

with Dirichlet boundary conditions $u^{\text{true}}(-1) = u^{\text{true}}(1) = 0$. The function space $V^{\text{true}}(\Omega^{\text{true}})$ is then given by

$$V^{\text{true}}(\Omega^{\text{true}}) = \left\{ v : \int_{\Omega^{\text{true}}} \left(v^2 + \frac{\partial v}{\partial x} \frac{\partial v}{\partial x} \right) dx < \infty \text{ and } v(-1) = v(1) = 0 \right\}. \quad (54)$$

The true state $u^{\text{true}} \in V^{\text{true}}(\Omega^{\text{true}})$ satisfies

$$a^{\text{true}}(u^{\text{true}}, v) = \ell^{\text{true}}(v), \quad \forall v \in V^{\text{true}}(\Omega^{\text{true}}), \quad (55)$$

where

$$a^{\text{true}}(w, v) = \int_{\Omega^{\text{true}}} \frac{\partial w}{\partial x} \frac{\partial v}{\partial x} dx, \quad \ell^{\text{true}}(v) = \int_{\Omega^{\text{true}}} f^{\text{true}} v dx, \quad \forall w, v \in V^{\text{true}}(\Omega^{\text{true}}). \quad (56)$$

We prescribe a *synthetic source term* as $f^{\text{true}} = \sin(\pi x) + 4 \sin(4\pi x)$. It is easy to see that

$$u^{\text{true}} = \frac{\sin(\pi x)}{\pi^2} + \frac{\sin(4\pi x)}{4\pi^2}. \quad (57)$$

The true state is not known to us and will serve to assess the performance of our method.

We next assume that we know almost everything about the true model except for the source term f^{true} and the boundary data. We introduce a function space $V(\Omega)$ with $\Omega = (-1, 1)$ as

$$V(\Omega) = \left\{ v : \int_{\Omega} \left(v^2 + \frac{\partial v}{\partial x} \frac{\partial v}{\partial x} \right) dx < \infty \text{ and } v(-1) = b_1, v(1) = b_2 \right\}, \quad (58)$$

where the boundary data b_1 and b_2 will be determined from the observed data. We then define our best knowledge model: find $u \in V(\Omega)$ such that

$$a(u, v) = \ell(v), \quad \forall v \in V(\Omega), \quad (59)$$

where

$$a(w, v) = \int_{\Omega} \frac{\partial w}{\partial x} \frac{\partial v}{\partial x} dx, \quad \ell(v) = \int_{\Omega} f v dx, \quad \forall w, v \in V(\Omega), \quad (60)$$

In practice, we replace the continuous space $V(\Omega)$ with a discrete counterpart, for this problem, a 2000-element linear finite element space.

5.2 Model specifications

We now specify the observation functionals $c_i(v) = \int_{\Omega} \delta(x_i) v dx = v(x_i)$, $i = 1, \dots, M$, where $\delta(x)$ is the Dirac delta function and the x_i are the extended Chebyshev nodes [3] in the interval $[-1, 1]$:

$$x_i = -\frac{\cos((2i-1)\pi/(2M))}{\cos(\pi/(2M))}, \quad i = 1, \dots, M. \quad (61)$$

These functionals correspond to pointwise observations taken at the points x_i . Note that the set of measurement points $\{x_i\}_{i=1}^M$ varies with M . Note also that the finite element mesh is designed to include $\{x_i\}_{i=1}^M$ in its grid points.

We shall assume that the observations are noise-free, that is we have $\sigma = 0$ and $d_i = u^{\text{true}}(x_i)$, $i = 1, \dots, M$. Since the observations at x_1 and x_M are used to define the function space $V(\Omega)$ in (58) (that is we set $b_1 = d_1 = 0$ and $b_2 = d_M = 0$), we can only use the remaining $(M - 2)$ observations to construct the training set as

$$T = \{(\phi_i, d_i - s_i), i = 2, \dots, M - 1\}, \quad (62)$$

where $\phi_i \in V(\Omega)$, $i = 2, \dots, M - 1$ satisfies

$$a(v, \phi_i) = -c_i(v), \quad \forall v \in V(\Omega), \quad (63)$$

and $s_i = u(x_i)$, $i = 2, \dots, M - 1$ are the outputs of the best knowledge model. Hence, the training set has only $(M - 2)$ samples. Furthermore, we use a bilinear covariance operator of the form

$$k(w, v; \boldsymbol{\theta}) = \int_{\Omega} \left(\theta_1 wv + \theta_2 \frac{\partial w}{\partial x} \frac{\partial v}{\partial x} \right) dx, \quad (64)$$

The parameters $\boldsymbol{\theta} = (\theta_1, \theta_2)$ are determined by maximizing the log marginal likelihood (28).

We will compare our method to the standard Gaussian process regression described in Section 3 which ignores the best knowledge model and utilizes only the data. To this end, we employ a squared-exponential covariance function of the form

$$\kappa(x, x'; \boldsymbol{\zeta}) = \zeta_1^2 \exp\left(-\frac{1}{2\zeta_2^2}(x - x')^2\right), \quad (65)$$

where ζ_1 represents the signal variance, while ζ_2 represents the length scale. These parameters are set by maximizing the log marginal likelihood (12).

5.3 Results and discussions

We consider $f = 4 \sin(4\pi x)$ for the best knowledge model (59)-(60). This yields the best knowledge state $u = \sin(4\pi x)/(4\pi^2)$. Figure 2 shows the true state u^{true} and the best knowledge state u . We observe that u is considerably different from u^{true} . Therefore, the best knowledge model does not produce a good prediction of the heat equation (53). We now apply Gaussian functional regression to this example and present numerical results to demonstrate the performance of our method relative to the standard GP regression.

We present in Table 2 the optimal hyperparameters, the $L^2(\Omega)$ norm of the prediction error, and the $L^2(\Omega)$ norm of the posterior standard deviation (the square root of the posterior variance) for our method and the standard GP regression. Here the $L^2(\Omega)$ norm of a function v is defined as $\|v\|_{\Omega} = (\int_{\Omega} v^2 dx)^{1/2}$. We observe that while θ_2 is always zero, θ_1 decreases as M increases, indicating that the prediction uncertainty is reduced as the number of observations increases. We also note that the length scale ζ_2 of the squared-exponential

Figure 2: Plots of the true state and the best knowledge state.

M	Gaussian Functional Regression			Standard GP Regression		
	(θ_1, θ_2)	$\ u^{\text{true}} - \bar{u}^*\ _{\Omega}$	$\ \sqrt{\text{var}(u^*)}\ _{\Omega}$	(ζ_1, ζ_2)	$\ u^{\text{true}} - \bar{h}^*\ _{\Omega}$	$\ \sqrt{\text{var}(h^*)}\ _{\Omega}$
4	(0.485, 0)	4.43E-3	4.64E-2	(0.053, 0.064)	9.14E-2	6.84E-2
5	(0.319, 0)	6.02E-3	2.47E-2	(0.076, 0.005)	8.12E-2	9.73E-2
6	(0.247, 0)	1.41E-3	1.60E-2	(0.062, 0.046)	8.68E-2	7.84E-2
7	(0.199, 0)	9.06E-4	1.12E-2	(0.086, 0.621)	3.13E-2	4.40E-4
8	(0.167, 0)	4.16E-4	8.31E-3	(0.074, 0.306)	3.64E-2	8.36E-3
9	(0.143, 0)	2.38E-4	6.42E-3	(0.067, 0.203)	3.25E-2	2.05E-2
10	(0.125, 0)	1.42E-4	5.12E-3	(0.064, 0.186)	9.39E-3	1.85E-2
11	(0.111, 0)	9.10E-5	4.18E-3	(0.062, 0.223)	1.90E-2	7.00E-3
12	(0.100, 0)	6.10E-5	3.48E-3	(0.065, 0.193)	9.24E-4	9.43E-3
13	(0.091, 0)	4.25E-5	2.94E-3	(0.061, 0.196)	4.57E-3	5.96E-3
14	(0.084, 0)	3.06E-5	2.52E-3	(0.064, 0.200)	4.24E-4	3.87E-3
15	(0.077, 0)	2.26E-5	2.18E-3	(0.065, 0.205)	1.58E-3	2.22E-3

Table 2: The optimal hyperparameters, the $L^2(\Omega)$ norm of the prediction error ($u^{\text{true}}(x) - \bar{u}^*(x)$ in our method and $u^{\text{true}}(x) - \bar{h}^*(x)$ in the standard GP regression), and the $L^2(\Omega)$ norm of the standard deviation function ($\sqrt{\text{var}(u^*(x))}$ in our method and $\sqrt{\text{var}(h^*(x))}$ in the standard GP regression) as a function of M for both Gaussian functional regression and standard GP regression. Here $\bar{h}^*(x)$ and $\text{var}(h^*(x))$ are the mean prediction and the posterior variance of $u^{\text{true}}(x)$ for the standard GP regression.

covariance function is relatively small for $M \leq 6$, indicating that the training set may be inadequate for the standard GP regression to produce a good prediction.

We see from Table 2 that the prediction error in our method converges significantly faster than that in the standard GP regression as M increases. Therefore, our method requires fewer observations to achieve the same accuracy. In particular, our method with 4 observations has slightly smaller error than the standard GP regression with 13 observations. This is made possible because our method uses both the best knowledge model and the observations to do regression on the space of functionals, whereas the standard GP regression uses the observations only to do regression on the space of functions. We also observe that the posterior standard deviation (measured in $L^2(\Omega)$ norm) shrinks with increasing M albeit at a slower rate than the prediction error, indicating that our posterior variance of the prediction error is rigorous. In contrast, the standard GP regression has the posterior standard deviation even smaller than the prediction error for small values of M , indicating that the posterior variance of the standard GP regression may not be rigorous when the training set is inadequate. This can be attributed to the fact the standard GP regression requires a large enough set of observations to provide accurate prediction and rigorous error estimation.

Finally, we show in Figure 3 the true state, the mean prediction, and the 95% confidence region (shaded area) for our method (left panels) and the standard GP regression (right panels). Here the 95% confidence region is an area bounded by the mean prediction plus and minus two times the standard deviation function. Note that the prediction error is zero at the measurement points, which is consistent with the theoretical result stated in Lemma 1. Moreover, the standard deviation function is also zero at the measurement points — a consequence of the fact that the prediction error is zero at those points. We see that our method does remarkably well even with just 4 observations when it is compared to the standard GP regression. For our method the true state u^{true} resides in the 95% confidence region which shrinks rapidly with as M increases, whereas for the standard GP regression u^{true} does not always reside in the 95% confidence region. Indeed, as seen in Figure 3(d), the standard GP regression gives poor prediction and erroneous 95% confidence region for $M = 8$. Although the standard GP regression provides more accurate prediction and rigorous 95% confidence region for $M = 12$, it is still not as good as our method. In summary, the numerical results obtained for this simple example show that our method outperforms the standard GP regression.

6 Potential Flow Around a Circular Cylinder

6.1 Problem description

We consider a potential flow around a circular cylinder of radius $a = 1$ in an infinite domain. The stream function for this flow is given by

$$u^{\text{true}} = y \left(1 - \frac{1}{x^2 + y^2} \right). \quad (66)$$

(a) Gaussian functional for $M = 4$

(b) Standard GP for $M = 4$

(c) Gaussian functional for $M = 8$

(d) Standard GP for $M = 8$

(e) Gaussian functional for $M = 12$

(f) Standard GP for $M = 12$

Figure 3: Panels show the true state and the mean prediction obtained using $M = 4, 8, 12$ observations for both Gaussian functional regression (left) and the standard GP regression (right). In these plots the shaded area represents the mean prediction plus and minus two times the standard deviation function (corresponding to the 95% confidence region).

This solution is shown in Figure 5(a). Note that only the upper left quadrant of the domain is considered because of the symmetry.

Since the circular cylinder is a simple geometry the stream function of the potential flow is known analytically. For more complicated geometries, however, the problem is usually solved by using a numerical method. In the case of a finite element method, the infinite domain needs to be truncated to a finite domain. This gives rise to the following PDE model

$$\Delta u = 0, \quad \text{in } \Omega, \quad (67)$$

where the truncated domain Ω and the boundary conditions are shown in Figure 4(a). The weak formulation is given by

$$a(u, v) = \ell(v), \quad \forall v \in V, \quad (68)$$

where, for all $w, v \in V$,

$$a(w, v) = \int_{\Omega} \nabla w \cdot \nabla v dx dy, \quad \ell(v) = \int_{\Gamma_T} v dx dy, \quad (69)$$

and $V \equiv \{v \in H^1(\Omega) : v = 0 \text{ on } \Gamma_B \cup \Gamma_C\}$. (In practice, we replace V with a FE approximation space of piecewise polynomials of degree $p = 3$ defined on a mesh of 690 elements. This FE discretization is fine enough such that the numerical error is insignificant.)

Figure 4: Potential flow around a circular cylinder: (a) Geometry and boundary conditions and (b) FE mesh and measurement points.

This example serves to demonstrate our approach when the uncertainty source comes from the domain truncation and the boundary condition. Due to this uncertainty, the best

Figure 5: Plot of the true state u^{true} , the best knowledge state u , and their difference $u^{\text{true}} - u$.

knowledge state u is significantly different from the true state u^{true} as shown in Figure 5(c). We aim to improve the prediction of the true state u^{true} by combining the best knowledge model with observations of the true state. The observation functionals are specified as

$$c_i(v) = \int_{\Omega} \exp\left(-\frac{(x - x_i)^2 + (y - y_i)^2}{0.25^2}\right) v dx dy, \quad i = 1, \dots, M, \quad (70)$$

where $(x_i, y_i), i = 1, \dots, M$ are selected by using the empirical interpolation method [2, 12] and are shown in Figure 4(b). We shall assume that the observations are noise-free, so that $d_i = c_i(u^{\text{true}}), i = 1, \dots, M$.

6.2 Results and discussions

We consider a bilinear covariance operator of the form

$$k(w, v; \boldsymbol{\theta}) = \int_{\Omega} (\theta_1 w v + \theta_2 \nabla w \cdot \nabla v) dx dy, \quad (71)$$

The parameters $\boldsymbol{\theta} = (\theta_1, \theta_2)$ are determined by maximizing the log marginal likelihood (28). Figure 6(a) shows the hyperparameter θ_1 as a function of the number of measurements. It is interesting to see that our hyperparameter selection method yields $\theta_2 = 0$ for any M . Note further that θ_1 tends to decrease as M increases.

We next show in Figure 6(b) the norm of the standard deviation function $\|\sqrt{\text{var}(u^*)}\|_{\Omega}$ and the error norm $\|u^{\text{true}} - \bar{u}\|_{\Omega}$ as a function of the number of measurements. As expected, both the error norm and the standard deviation norm decrease as M increases. Hence, increasing the number of measurements reduces the uncertainty in our prediction of the true state. This can be seen more clearly in Figure 7 which shows the error function and the standard deviation function for three different values of $M = 5, 10$, and 15 . Moreover, we observe that both the error function the standard deviation function near the top boundary Γ_{T} and the left boundary Γ_{L} of the domain are larger than those in the other region of the domain. This can be attributed to the fact that the uncertainty source comes from the domain truncation and the boundary condition on Γ_{T} and Γ_{L} .

Figure 6: (a) The hyperparameter θ_1 and (b) the error norm and the standard deviation norm as a function of M . Note that $\theta_2 = 0$ for any M .

7 Conclusions

In this paper, we have presented a new statistical approach to the problem of combining the observed data with a mathematical model to improve our prediction of a physical system. A new *Gaussian functional regression* method is presented, which has its root in Gaussian processes for functions. Our approach has the following unique properties. First, it allows for the incorporation of the best knowledge model into the regression procedure. Second, it can handle observations given in the form of linear functionals of the field variable. Third, the method is non-parametric in the sense that it provides a systematic way to optimally determine the prior covariance based on the data. Fourth, our method can compute not only the mean prediction but also the posterior covariance, characterizing the uncertainty in our prediction of the physical state. These features distinguish our method from other regression methods. Numerical results were presented to highlight these features and demonstrate the superior performance of the proposed method relative to the standard GP regression.

We conclude the paper by pointing out several possible extensions and directions for further research. We would like to extend the proposed approach to nonlinear PDE models, which will broaden the application domain of our method. Nonlinear PDEs represent some significant challenges because the adjoint problems will depend on the true state which we do not know and because our stochastic PDE does not preserve the Gaussian property of the functional g due to nonlinearity. We would also like to extend the method to goal-oriented statistical estimation in which we would like to infer new outputs rather the state of the physical system.

Figure 7: Panels show the error function (left) and the standard deviation function (right) obtained using $M = 5, 10, 15$ observations.

Acknowledgments

We would like to thank Professor Robert M. Freund of Sloan School of Management at MIT for fruitful discussions. This work was supported by AFOSR Grant No. FA9550-11-1-0141, AFOSR Grant No. FA9550-12-0357, and the Singapore-MIT Alliance.

A Bayesian linear functional regression

In this section, we develop a nonparametric Bayesian framework for linear functional regression with Gaussian noise. We then show that this framework is equivalent to our Gaussian functional regression method described in Section 3.

We begin by introducing an orthonormal basis set $\{\psi_i \in V\}_{i=1}^J$ such that $\int_{\Omega} \psi_i \psi_j d\mathbf{x} = \delta_{ij}$, where δ_{ij} is the Kronecker delta. We define associated linear functionals $\ell_i : V \rightarrow \mathbb{R}$ as

$$\ell_i(v) = \int_{\Omega} \psi_i v d\mathbf{x} \equiv m(\psi_i, v), \quad i = 1, \dots, J. \quad (72)$$

The linear functional regression model is defined by

$$g(v) = \sum_{i=1}^J w_i \ell_i(v), \quad y = g(v) + \varepsilon, \quad (73)$$

where $v \in V$ is the input function, $\mathbf{w} = [w_1, \dots, w_J]^T$ is a vector of weights (parameters) of the linear model, g is the functional and $y \in \mathbb{R}$ is the observed target value. We assume that the observed value y differs from the functional value $g(v)$ by additive noise

$$\varepsilon \sim \mathcal{N}(0, \sigma^2). \quad (74)$$

This noise assumption together with the model directly gives rise to the following likelihood

$$p(\mathbf{y}|\Phi, \mathbf{w}) = \prod_{j=1}^M \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(y_j - g(\phi_j))^2}{2\sigma^2}\right) = \mathcal{N}(\mathbf{L}^T \mathbf{w}, \sigma^2), \quad (75)$$

where $\mathbf{L}$ has entries $L_{ij} = \ell_i(\phi_j)$, $i = 1, \dots, J$, $j = 1, \dots, M$. Recall that ϕ_i , $i = 1, \dots, M$ are the adjoint states and that $y_i = d_i - s_i$, $i = 1, \dots, M$ are the observed outputs which are the differences between the observed data and the best knowledge outputs.

In the Bayesian formalism we need to specify a *prior distribution* for $\mathbf{w}$, which encodes our belief about the weights prior to taking into account the observations. We consider a zero-mean Gaussian prior distribution with covariance matrix $\mathbf{\Lambda}$ for the weights

$$p(\mathbf{w}) = \mathcal{N}(0, \mathbf{\Lambda}). \quad (76)$$

Without loss of generality we assume that $\mathbf{\Lambda}$ is a diagonal matrix, that is, $\Lambda_{ij} = 0$ if $i \neq j$. If $\mathbf{\Lambda}$ is not a diagonal matrix then we can always introduce a new weight vector $\mathbf{w}' = \mathbf{\Lambda}^{-1/2} \mathbf{w}$,

so that the covariance of $\mathbf{w}'$ is a diagonal matrix. Then we choose to work with $\mathbf{w}'$ instead of $\mathbf{w}$. The posterior distribution over the weights is then obtained by using Bayes' rule as

$$p(\mathbf{w}|\mathbf{y}, \Phi) = \frac{p(\mathbf{y}|\Phi, \mathbf{w})p(\mathbf{w})}{p(\mathbf{y}|\Phi)}, \quad (77)$$

where the normalizing constant is given by

$$p(\mathbf{y}|\Phi) = \int p(\mathbf{y}|\Phi, \mathbf{w})p(\mathbf{w})d\mathbf{w}. \quad (78)$$

Substituting (75) and (76) into (77) and working through some algebraic manipulations, we obtain

$$p(\mathbf{w}|\mathbf{y}, \Phi) \sim \mathcal{N}\left(\frac{1}{\sigma^2}\mathbf{B}^{-1}\mathbf{L}\mathbf{y}, \mathbf{B}^{-1}\right), \quad (79)$$

where $\mathbf{B} = \sigma^{-2}\mathbf{L}\mathbf{L}^T + \mathbf{\Lambda}^{-1}$. The posterior distribution (79) combines the likelihood and the prior distribution, and captures all information about the weights.

To make predictions for a test function $\phi^* \in V$ we average the linear functional regression model (73) over all possible values of the weights under their posterior distribution (79). Hence, the predictive distribution for $g^* = g(\phi^*)$ at ϕ^* is given by

$$g^*|\phi^*, \Phi, \mathbf{y} \sim \mathcal{N}\left(\frac{1}{\sigma^2}\mathbf{l}(\phi^*)^T\mathbf{B}^{-1}\mathbf{L}\mathbf{y}, \mathbf{l}(\phi^*)^T\mathbf{B}^{-1}\mathbf{l}(\phi^*)\right), \quad (80)$$

where $\mathbf{l}(\phi^*)$ has entries $l_i(\phi^*) = \ell_i(\phi^*)$, $i = 1, \dots, J$. The predictive distribution is a Gaussian distribution as expected. However, the predictive distribution (80) requires the matrix inversion $\mathbf{B}^{-1}$ which may be expensive. Fortunately, we can derive an equivalent distribution which is more efficient to compute than the original one (80).

We recall the Woodbury, Sherman and Morrison (WSM) formula for the matrix inversion

$$(\mathbf{Z} + \mathbf{U}\mathbf{W}\mathbf{V}^T)^{-1} = \mathbf{Z}^{-1} - \mathbf{Z}^{-1}\mathbf{U}(\mathbf{W}^{-1} + \mathbf{V}^T\mathbf{Z}^{-1}\mathbf{U})^{-1}\mathbf{V}^T\mathbf{Z}^{-1}. \quad (81)$$

Using the WSM formula with $\mathbf{Z} = \mathbf{\Lambda}$, $\mathbf{W} = \sigma^{-2}\mathbf{I}$, $\mathbf{U} = \mathbf{V} = \mathbf{L}$ we obtain

$$\mathbf{B}^{-1} = (\sigma^{-2}\mathbf{L}\mathbf{L}^T + \mathbf{\Lambda}^{-1})^{-1} = \mathbf{\Lambda} - \mathbf{\Lambda}\mathbf{L}(\sigma^2\mathbf{I} + \mathbf{L}^T\mathbf{\Lambda}\mathbf{L})^{-1}\mathbf{L}^T\mathbf{\Lambda}. \quad (82)$$

It thus follows that

$$\begin{aligned} \mathbf{B}^{-1}\mathbf{L} &= \mathbf{\Lambda}\mathbf{L} - \mathbf{\Lambda}\mathbf{L}(\sigma^2\mathbf{I} + \mathbf{L}^T\mathbf{\Lambda}\mathbf{L})^{-1}\mathbf{L}^T\mathbf{\Lambda}\mathbf{L} \\ &= \mathbf{\Lambda}\mathbf{L}(\sigma^2\mathbf{I} + \mathbf{L}^T\mathbf{\Lambda}\mathbf{L})^{-1}((\sigma^2\mathbf{I} + \mathbf{L}^T\mathbf{\Lambda}\mathbf{L}) - \mathbf{L}^T\mathbf{\Lambda}\mathbf{L}) \\ &= \sigma^2\mathbf{\Lambda}\mathbf{L}(\sigma^2\mathbf{I} + \mathbf{L}^T\mathbf{\Lambda}\mathbf{L})^{-1}. \end{aligned} \quad (83)$$

Inserting (82) and (83) into (80) we get

$$g^*|\phi^*, \Phi, \mathbf{y} \sim \mathcal{N}(\bar{g}^*, \text{var}(g^*)), \quad (84)$$

where

$$\begin{aligned}\bar{g}^* &= \mathbf{l}(\phi^*)^T \mathbf{\Lambda} \mathbf{L} (\sigma^2 \mathbf{I} + \mathbf{L}^T \mathbf{\Lambda} \mathbf{L})^{-1} \mathbf{y}, \\ \text{var}(g^*) &= \mathbf{l}(\phi^*)^T \mathbf{\Lambda} \mathbf{l}(\phi^*) - \mathbf{l}(\phi^*)^T \mathbf{\Lambda} \mathbf{L} (\sigma^2 \mathbf{I} + \mathbf{L}^T \mathbf{\Lambda} \mathbf{L})^{-1} \mathbf{L}^T \mathbf{\Lambda} \mathbf{l}(\phi^*)\end{aligned}\tag{85}$$

In the predictive distribution (84), we need to invert the matrix $(\sigma^2 \mathbf{I} + \mathbf{L}^T \mathbf{\Lambda} \mathbf{L})$ of size $M \times M$. Hence, it is more attractive than the original distribution (80).

Thus far, we have not mentioned how we construct the basis functions $\{\psi_i\}_{i=1}^J$ and the covariance matrix $\mathbf{\Lambda}$. We observe that in the predictive distribution (84) we need to compute $\mathbf{L}^T \mathbf{\Lambda} \mathbf{L}, \mathbf{l}(\phi^*)^T \mathbf{\Lambda} \mathbf{L}, \mathbf{l}(\phi^*)^T \mathbf{\Lambda} \mathbf{l}(\phi^*)$, where both $\mathbf{L}$ and $\mathbf{l}(\phi^*)$ depend on the basis set $\{\psi_i\}_{i=1}^J$ through the definition of the linear functionals $\ell_i(\cdot), i = 1, \dots, J$. The entries of these quantities have the form $\mathbf{l}(v)^T \mathbf{\Lambda} \mathbf{l}(v')$ for $v, v' \in V$. We can thus define

$$k(v', v) = \mathbf{l}(v')^T \mathbf{\Lambda} \mathbf{l}(v), \quad \text{for all } v, v' \in V. \tag{86}$$

It follows that

$$k(v', v) = \sum_{i=1}^J \sum_{j=1}^J \ell_i(v') \Lambda_{ij} \ell_j(v) = \sum_{i=1}^J m(\psi_i, v') \Lambda_{ii} m(\psi_i, v), \quad \forall v', v \in V, \tag{87}$$

since $\ell_i(v) = m(\psi_i, v)$ and $\mathbf{\Lambda}$ is a diagonal matrix. We next choose $v' = \psi_n$ and invoke $m(\psi_i, \psi_n) = \delta_{in}$ to arrive at the following eigenvalue problem:

$$k(\psi_n, v) = \Lambda_{nn} m(\psi_n, v), \quad \forall v \in V. \tag{88}$$

This equation shows that $\{\psi_n\}_{n=1}^J$ and $\{\Lambda_{nn}\}_{n=1}^J$ are the eigenfunctions and eigenvalues of the covariance operator $k(\cdot, \cdot)$. Once the covariance operator $k(\cdot, \cdot)$ is specified, we can construct the basis set $\{\psi_n\}_{n=1}^J$ and the covariance matrix $\mathbf{\Lambda}$ by solving the eigenvalue problem (88).⁶ We can then make predictions using the equations (84) and (85). The predictive distribution (84) is nothing but the same as that of Gaussian functional regression described in Section 3. However, unlike nonparametric Bayesian inference, Gaussian functional regression does not require the computation of eigenfunctions and eigenvalues. This is because Gaussian functional regression needs only to compute the inner products $k(v, v')$ for some pairs $v, v' \in V$, which is known as the *kernel trick* [16].

⁶Of course, we can also define the covariance operator $k(v', v)$ in terms of the basis set $\{\psi_n\}_{n=1}^J$ and the covariance matrix $\mathbf{\Lambda}$ by using the expression (87). However, this series expansion of the covariance operator render the inner products too computationally expensive and should be avoided. Instead, one should work directly with an explicitly analytical form of the covariance operator as listed in Table 2.

References

- [1] N. Aronszajn. Theory of reproducing kernels, 1950.
- [2] M. Barrault, Y. Maday, N. C. Nguyen, and A. T. Patera. An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations. *Comptes Rendus Mathematique*, 339(9):667–672, 2004.
- [3] L. Brutman. On the Lebesgue Function for Polynomial Interpolation. *SIAM Journal on Numerical Analysis*, 15(4):694–704, August 1978.
- [4] A. Cohen, M. A. Davenport, and D. Leviatan. On the Stability and Accuracy of Least Squares Approximations. *Foundations of Computational Mathematics*, 13(5):819–834, March 2013.
- [5] T. A. El Moselhy and Y. M. Marzouk. Bayesian inference with optimal maps. *Journal of Computational Physics*, 231(23):7815–7850, October 2012.
- [6] R. Everson and L. Sirovich. Karhunen-Loeve procedure for gappy data. *Opt. Soc. Am. A*, 12(8):1657–1664, 1995.
- [7] A. Gelman, J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC; 3 edition, 2013.
- [8] R. E. Kalman. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME Journal of Basic Engineering*, 82(1):35 – 45, 1960.
- [9] J. M. Lewis, S. Lakshmivarahan, and S. Dhall. *Dynamic Data Assimilation: A Least Squares Approach*. Cambridge University Press, 2006.
- [10] Z. Li and I. M. Navon. Optimality of variational data assimilation and its relationship with the Kalman filter and smoother. *Quarterly Journal of the Royal Meteorological Society*, 127(572):661–683, January 2001.
- [11] A. C. Lorenc. A Global Three-Dimensional Multivariate Statistical Interpolation Scheme. *Monthly Weather Review*, 109(4):701–721, April 1981.
- [12] Y. Maday, N. C. Nguyen, A. T. Patera, and G. S. H. Pau. A general multipurpose interpolation procedure: the magic points. *Communications on Pure and Applied Analysis*, 8(1):383–404, October 2008.
- [13] N. C. Nguyen, A. T. Patera, and J. Peraire. A best points interpolation method for efficient approximation of parametrized functions. *International Journal for Numerical Methods in Engineering*, 73(4):521–543, January 2008.
- [14] A. O'Hagan and J. F. C. Kingman. Curve Fitting and Optimal Design for Prediction. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40(1):1 – 42, 1978.

- [15] A. T. Patera and E. M. Ronquist. Regression on parametric manifolds: Estimation of spatial fields, functional outputs, and parameters from noisy data. *Comptes Rendus Mathematique*, 350(9-10):543–547, May 2012.
- [16] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [17] J. Sacks, W. J. Welch, T. J. Mitchell, and H. P. Wynn. Design and Analysis of Computer Experiments. *Statistical Science*, 4(4):409–423, November 1989.
- [18] Simo Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.
- [19] T. Sondergaard and P. F. J. Lermusiaux. Data Assimilation with Gaussian Mixture Models Using the Dynamically Orthogonal Field Equations. Part I: Theory and Scheme. *Monthly Weather Review*, 141(6):1737–1760, June 2013.
- [20] Albert Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2005.
- [21] C. K. Wikle and L. M. Berliner. A Bayesian tutorial for data assimilation. *Physica D: Nonlinear Phenomena*, 230(1-2):1–16, June 2007.
- [22] K. Willcox. Unsteady Flow Sensing and Estimation via the Gappy Proper Orthogonal Decomposition. *Computers and Fluids*, 35:208–226, 2006.
- [23] J. Wolberg. *Data Analysis Using the Method of Least Squares: Extracting the Most Information from Experiments*. Springer; 2006 edition, 2005.
- [24] M. Yano, J. D. Penn, and A. T. Patera. A model-data weak formulation for simultaneous estimation of state and model bias. *Comptes Rendus Mathematique*, 351(23-24):937–941, December 2013.
- [25] Z. Zhang. Parameter estimation techniques: a tutorial with application to conic fitting. *Image and Vision Computing*, 15(1):59–76, January 1997.