

# Physics-Informed Gaussian Process Regression Generalizes Linear PDE Solvers

Marvin Pförtner<sup>1</sup>  
Ingo Steinwart<sup>2</sup>  
Philipp Hennig<sup>1</sup>  
Jonathan Wenger<sup>1</sup>

MARVIN.PFOERTNER@UNI-TUEBINGEN.DE  
INGO.STEINWART@MATHEMATIK.UNI-STUTTGAERT.DE  
PHILIPP.HENNIG@UNI-TUEBINGEN.DE  
JONATHAN.WENGER@UNI-TUEBINGEN.DE

<sup>1</sup>University of Tübingen, Tübingen AI Center

<sup>2</sup>University of Stuttgart

## Abstract

Linear partial differential equations (PDEs) are an important, widely applied class of mechanistic models, describing physical processes such as heat transfer, electromagnetism, and wave propagation. In practice, specialized numerical methods based on discretization are used to solve PDEs. They generally use an estimate of the unknown model parameters and, if available, physical measurements for initialization. Such solvers are often embedded into larger scientific models with a downstream application and thus error quantification plays a key role. However, by ignoring parameter and measurement uncertainty, classical PDE solvers may fail to produce consistent estimates of their inherent approximation error. In this work, we approach this problem in a principled fashion by interpreting solving linear PDEs as physics-informed Gaussian process (GP) regression. Our framework is based on a key generalization of the Gaussian process inference theorem to observations made via an arbitrary bounded linear operator. Crucially, this probabilistic viewpoint allows to (1) *quantify the inherent discretization error*; (2) *propagate uncertainty about the model parameters to the solution*; and (3) *condition on noisy measurements*. Demonstrating the strength of this formulation, we prove that it strictly generalizes methods of weighted residuals, a central class of PDE solvers including collocation, finite volume, pseudospectral, and (generalized) Galerkin methods such as finite element and spectral methods. This class can thus be directly equipped with a structured error estimate. In summary, our results enable the seamless integration of mechanistic models as modular building blocks into probabilistic models by blurring the boundaries between numerical analysis and Bayesian inference.

**Keywords:** physics-informed machine learning, probabilistic numerics, partial differential equations, method of weighted residuals, Galerkin methods, Gaussian processes, bounded linear operators

## 1. Introduction

Partial differential equations (PDEs) are powerful mechanistic models of static and dynamic systems with continuous spatial interactions (Borthwick, 2018). They are widely used in the natural sciences, especially in physics, and in applied fields like engineering, medicine and finance. Linear PDEs form a subclass describing physical phenomena such as heat diffusion (Fourier, 1822), electromagnetism (Maxwell, 1865), and continuum mechanics (Lautrup, 2005). Additionally, they are used in applications as diverse as computer graphics (Kazhdan et al., 2006), medical imaging (Holder, 2005), or option pricing (Black and Scholes, 1973).

**Scientific inference with PDEs** Given a mechanistic model of a (physical) system in the form of a linear PDE  $\mathcal{D}[\mathbf{u}] = f$ , where  $\mathcal{D}$  is a linear differential operator mapping between vector spaces of functions, the system can be simulated by solving the PDE subject to a set of linear boundary conditions (BC), given by a linear operator  $\mathcal{B}$  and a function  $g$  defined on the boundary of the domain, s.t.  $\mathcal{B}[\mathbf{u}] = g$  (Evans, 2010). For instance, given all material parameters and heat sources involved, a PDE can describe the temperature distribution in an electronic component, while the boundary conditions describe the heat flux out of the component at the surface. Since hardly any practically relevant PDE can be solved analytically (Borthwick, 2018), in practice, specialized numerical methods relying on discretization are employed. Often such solvers are embedded into larger scientific models, where model parameters are inferred from measurement and downstream analyses depend on the resulting simulation. For example, we would like to model whether said electronic component hits critical temperature thresholds during operation to assess its longevity.

**Challenges when solving PDEs** When performing scientific inference with PDEs via numerical simulation, one is faced with three fundamental challenges.

- (C1) *Limited computation.* Any numerically computed solution  $\hat{\mathbf{u}} \approx \mathbf{u}$  suffers from approximation error. In practice, a sufficiently accurate simulation often requires vast amounts of computational resources.
- (C2) *Partially-known physics.* While the underlying physical mechanism is encoded in the formulation of the PDE, in practice, its exact parameters and boundary conditions are often unknown. For example, the position and strength of heat sources  $f$  within the aforementioned electric component are only approximately known. Similarly, material parameters like thermal conductivity, which define  $\mathcal{D}$ , can often only be estimated. Finally, the initial or boundary conditions  $\mathcal{B}[\mathbf{u}] = g$  are also only partially known. For example, how much heat an electrical component dissipates via its surface.
- (C3) *Error propagation.* Limited computation and partially-known physics inevitably introduce error into the simulation. This resulting bias can fundamentally alter conclusions drawn from downstream analysis steps, in particular if these are sensitive to input variability. For example, an electronic component may be deemed safe based on the simulation, although its true internal temperature hits safety-critical levels repeatedly.

**Solving PDEs as a learning problem** The challenges of scientific inference with PDEs are fundamentally issues of partial information. Here, we interpret solving a PDE as a *learning problem*, specifically as physics-informed regression, in the spirit of probabilistic numerics (Hennig et al., 2015; Cockayne et al., 2019b; Oates and Sullivan, 2019; Owhadi et al., 2019; Hennig et al., 2022). By leveraging the tools of Bayesian inference, we can tackle the challenges (C1) to (C3). As illustrated in Figure 1(a), we model the solution of the PDE with a Gaussian process, which we condition on observations of the boundary conditions, the PDE itself and any physical measurements:

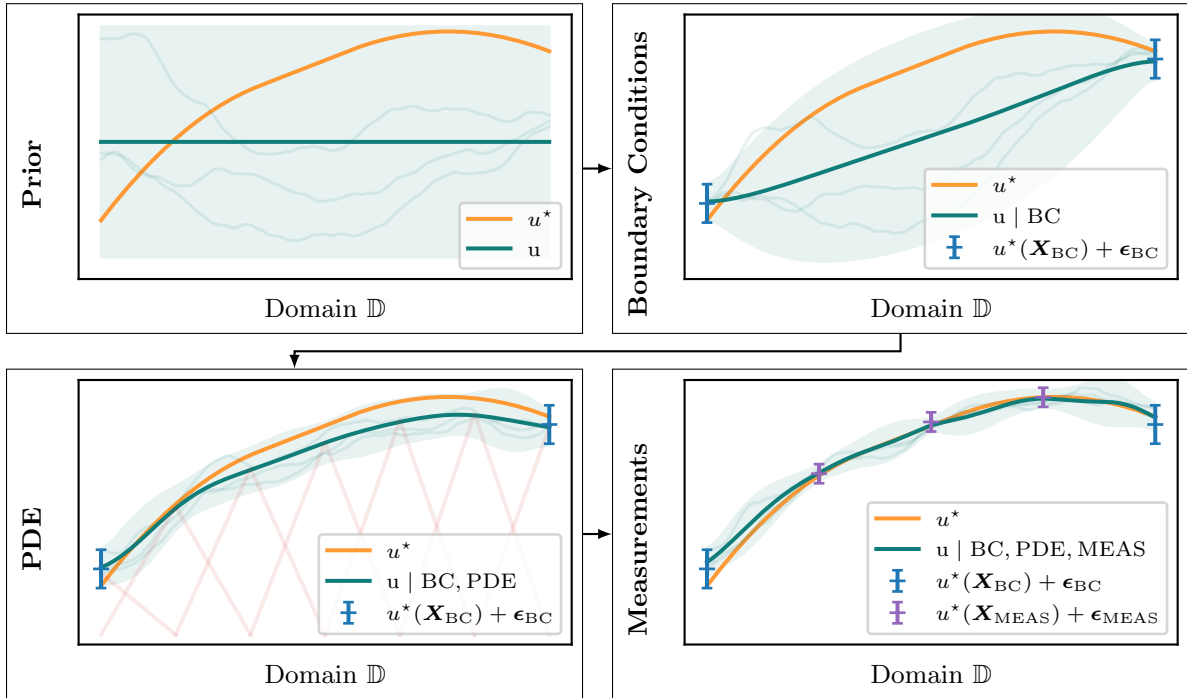
- *Encoding prior knowledge.* We can efficiently leverage any available computation by encoding inductive bias about the solution of the PDE. For example, we can identify the

solution space by “partial derivative counting”. Moreover, since PDEs typically model physical systems, expert knowledge is often available. This includes known physical properties of the system such as symmetries, as well as more subjective estimates from previous experience with similar systems or computationally cheap approximations.


- *Conditioning on the boundary conditions.* The linear boundary conditions can be interpreted as measurements of the solution of the PDE on the boundary. By conditioning on (some of) these measurements, we are not limited to satisfying the boundary conditions exactly, but can directly model uncertain constraints without having to resort to point estimates. Instead, we propagate the uncertainty to the solution estimate. This also allows us to handle cases where we do not have a functional form  $g$  of the constraints, but only a discrete set of constraints at boundary points.
- *Conditioning on the PDE.* Conditioning a probability measure over the solution on the analytic “observation” that the PDE holds is generally intractable. In the spirit of classic approaches for solving PDEs, we relax the PDE-constraint by requiring only a finite number of projections of the associated PDE residual onto carefully chosen test functions to be zero. This choice of projections defines the discretization and allows for control over the amount of expended computation. The resulting posterior quantifies the algorithm’s uncertainty within a whole set of solution candidates.
- *Conditioning on measurements.* Finally, we can also condition on direct measurements of the solution itself. This is especially useful if parameters of the differential operator or boundary conditions are uncertain, or if the computational budget is restrictive.

The resulting posterior belief quantifies the uncertainty about the true solution induced by limited computation and partially-known physics (see Figure 1(b)). By quantifying this error probabilistically, we can propagate it to any downstream analysis or decision. For example, to project the longevity of a newly designed electrical component, we want to simulate how likely it will hit a critical temperature threshold during operation. Given our posterior belief, we can simply compute the marginal probability instead of performing Monte-Carlo sampling, which would require repeated PDE solves at significant computational expense.


**Contribution** We introduce a *probabilistic learning framework for the solution of (systems of) linear PDEs*. Our framework can be viewed as physics-informed Gaussian process regression. It is based on a crucial generalization of a popular result on conditioning GPs on linear observations to observations made via an arbitrary bounded linear operator with values in  $\mathbb{R}^n$  (Theorem 1). This enables *combined quantification of uncertainty* from the inherent discretization error, uncertain initial or boundary conditions, as well as noisy measurements of the solution. While connections between GP inference and the solution of PDEs were made in the past (see Section 3.5), corresponding methods have largely focused on estimating strong solutions by leveraging finite difference or collocation schemes. In contrast, our framework applies to both weak and strong formulations and generalizes a significantly broader class of existing numerical methods. Our approach is a strict probabilistic *generalization of methods of weighted residuals* (Corollary 3.3), including collocation, finite volume, (pseudo)spectral, and (generalized) Galerkin methods such as finite element methods. The resulting probabilistic methods thus have the same convergence properties as



(a) *Learning to solve the Poisson equation.* A problem-specific Gaussian process prior  $u$  is conditioned on partially-known physics, given by uncertain boundary conditions (BC) and a linear PDE, as well as on noisy physical measurements from experiment. The boundary conditions and the right-hand side of the PDE are not known but inferred from a small set of noise-corrupted measurements. The plots juxtapose the belief  $u | \dots$  with the true solution  $u^*$  of the latent boundary value problem.



(b) *Uncertainty quantification.* Marginal posterior standard deviation after conditioning on uncertain boundary conditions, a linear PDE, and noisy (physical) measurements.



(c) *Generalization of Classical Solvers.* For certain priors our framework reproduces any method of weighted residuals, e.g. the finite element method, in its posterior mean.

Figure 1: A physics-informed Gaussian process framework for the solution of linear PDEs.

their classic counterparts, while providing a structured error estimate. Moreover, the probabilistic viewpoint allows to incorporate partially-known physics and (noisy) experimental measurements.

## 2. Background

### 2.1 Linear Partial Differential Equations

A linear *partial differential equation* (PDE) is an equation of the form

$$\mathcal{D}[\mathbf{u}] = f, \quad (2.1)$$

where  $\mathcal{D}: \mathbb{U} \rightarrow \mathbb{V}$  is a linear differential operator (see Definition C.2) between a Banach space  $\mathbb{U}$  of  $\mathbb{R}^{d'}$ -valued functions and a Banach space  $\mathbb{V}$  of real-valued functions on a common open and bounded domain  $\mathbb{D} \subset \mathbb{R}^d$ , and  $f \in \mathbb{V}$  is the right-hand side function. For simplicity of exposition, we will often focus on the case  $d' = 1$ , in which case we write  $u$  instead of  $\mathbf{u}$ . Systems modeled by linear PDEs are often further constrained by linear *boundary conditions* (BCs)  $\mathcal{B}[\mathbf{u}] = g$  describing the behavior of the system on the boundary  $\partial\mathbb{D}$  of the domain, where  $\mathcal{B}$  is a linear operator mapping functions  $\mathbf{u} \in \mathbb{U}$  onto functions  $\mathcal{B}[\mathbf{u}]: \partial\mathbb{D} \rightarrow \mathbb{R}$  defined on the boundary and  $g: \partial\mathbb{D} \rightarrow \mathbb{R}$ . Common types of boundary conditions for  $d' = 1$  are:

- *Dirichlet*: Specify the values of the solution on the boundary, i.e.  $\mathcal{B}[u] = u|_{\partial\mathbb{D}}$ .
- *Neumann*: Specify the exterior normal derivative on the boundary, i.e.  $\mathcal{B}[u](\mathbf{x}) := \partial_{\boldsymbol{\eta}(\mathbf{x})}u(\mathbf{x})$ , where  $\boldsymbol{\eta}(\mathbf{x})$  is the exterior normal vector at each point of the boundary.

A PDE and a set of boundary conditions is referred to as a *boundary value problem* (BVP). A prototypical example of a linear PDE, used in thermodynamics, electrostatics and Newtonian gravity, is the *Poisson equation*  $-\Delta u = f$ , where  $\Delta u = \sum_{i=1}^d \frac{\partial^2 u}{\partial x_i^2}$  is the Laplacian.

#### 2.1.1 WEAK FORMULATION

Many models of physical phenomena are expressed as functions  $\mathbf{u}$ , which are not (continuously) differentiable or even continuous (Evans, 2010; Borthwick, 2018; von Harrach, 2021). In other words, they are not *strong solutions* to any PDE. There are also PDEs derived from established physical principles, which do not admit strong solutions at all. To address this, one can weaken the notion of differentiability leading to the concept of weak solutions. Many of the aforementioned physical phenomena are in fact weak solutions. As an example<sup>1</sup>, consider the weak formulation of the stationary heat equation for non-homogeneous media

$$-\operatorname{div}(\kappa \nabla u) = \dot{q}_V. \quad (2.2)$$

Assume that  $u \in C^2(\mathbb{D})$ ,  $\kappa \in C^1(\mathbb{D})$ , and  $\dot{q}_V \in C^0(\mathbb{D})$ . If  $u$  is a solution to Equation (2.2), then we can integrate both sides of the equation against a *test function*  $v \in C_c^\infty(\mathbb{D})$ , i.e. an infinitely smooth function with compact support (see Definition C.3), which results in

$$-\int_{\mathbb{D}} \operatorname{div}(\kappa \nabla u)(\mathbf{x}) v(\mathbf{x}) \, d\mathbf{x} = \int_{\mathbb{D}} \dot{q}_V(\mathbf{x}) v(\mathbf{x}) \, d\mathbf{x}.$$

---

1. Our exposition is a strongly abbreviated version of Evans (2010, Section 6.1.2).

Since both  $u$  and  $v$  are sufficiently differentiable, we can apply integration by parts (Green's first identity) to the first integral to obtain

$$\underbrace{\int_{\mathbb{D}} \langle \kappa(\mathbf{x}) \nabla u(\mathbf{x}), \nabla v(\mathbf{x}) \rangle d\mathbf{x}}_{=: B[u, v]} = \int_{\mathbb{D}} \dot{q}_V(\mathbf{x}) v(\mathbf{x}) d\mathbf{x}, \quad (2.3)$$

since  $v|_{\partial\mathbb{D}} = 0$ . This expression does not require  $u$  to be twice differentiable. Rather,  $u$  only needs to be once *weakly differentiable* (see Evans 2010, Section 5.2.1) with  $(\nabla u)_i \in L_2(\mathbb{D})$ . Intuitively speaking, a weak derivative of a (classically non-differentiable) function “behaves like a derivative” when integrated against a smooth test function. These relaxed requirements on  $u$  are exactly the defining properties of the Sobolev space  $H^1(\mathbb{D})$ , i.e. it suffices that  $u \in H^1(\mathbb{D})$ . Similarly, we can weaken all other assumptions to  $v \in H_0^1(\mathbb{D})$ ,  $\dot{q}_V \in L_2(\mathbb{D})$  and  $\kappa \in L_\infty(\mathbb{D})$ . Then, for  $u \in H^1(\mathbb{D})$  and  $v \in H_0^1(\mathbb{D})$ , Equation (2.3) is equivalent to

$$\mathcal{D}^w[u] = f^w, \quad (2.4)$$

where  $\mathcal{D}^w : H^1(\mathbb{D}) \rightarrow H_0^1(\mathbb{D})'$ ,  $u \mapsto B[u, \cdot]$  and  $f^w = \langle \dot{q}_V, \cdot \rangle_{L_2(\mathbb{D})} \in H_0^1(\mathbb{D})'$ . Here,  $H_0^1(\mathbb{D})'$  denotes the continuous dual space of  $H_0^1(\mathbb{D})$ . We define a *weak solution* of Equation (2.2) as  $u \in H^1(\mathbb{D})$  such that Equation (2.4), known as the *weak* or *variational formulation*, holds.

**Definition 2.1.** A weak formulation of a linear PDE  $\mathcal{D}[\mathbf{u}] = f$  is an equation of the form

$$\mathcal{D}^w[\mathbf{u}] = f^w, \quad (2.5)$$

where  $\mathcal{D}^w : \mathbb{U} \rightarrow \mathbb{V}'$  is a linear operator induced by the differential operator  $\mathcal{D}$  and  $f^w \in \mathbb{V}'$  is a linear functional induced by the right-hand side  $f$ . A solution to Equation (2.5) is called a weak solution of the PDE. In this context,  $\mathcal{D}[\mathbf{u}] = f$  is called the strong formulation of the PDE and any solution to it is called a strong or classical solution.

### 2.1.2 METHODS OF WEIGHTED RESIDUALS

Unfortunately, linear PDEs both in weak and strong formulation are in general not analytically solvable, so approximate solutions are sought instead. *Methods of weighted residuals* (MWR) constitute a large family of popular numerical approximation schemes for linear PDEs, including *collocation*, *finite volume*, *(pseudo)spectral*, and *(generalized) Galerkin methods* such as *finite-element methods* (Fletcher, 1984). Loosely speaking, MWRs interpret a linear PDE as a root-finding problem for the associated PDE *residual*, i.e.  $\mathcal{D}[\mathbf{u}] - f = 0$ . Finding the solution of such a system of an uncountably infinite number of equations with infinitely many unknowns is generally intractable. To render the problem tractable, we reduce the number of equations by “projecting” onto  $\mathbb{R}^n$  using a finite number of continuous linear *test functionals*  $l^{(1)}, \dots, l^{(n)} \in \mathbb{V}'$ , i.e. we use that the residual being zero implies

$$l^{(i)}[\mathcal{D}[\mathbf{u}] - f] = (l^{(i)} \circ \mathcal{D})[\mathbf{u}] - l^{(i)}[f] = 0 \quad (2.6)$$

for all  $i = 1, \dots, n$ . This is a relaxation of the original problem, since the above is not an equivalence but only an implication.<sup>2</sup> A common choice for the test functionals appearing

---

2. This means that Equation (2.6) will generally have infinitely many solutions and needs regularization to have a unique solution.

in a large class of MWRs is the integral  $l^{(i)}[v] := \int_{\mathbb{D}} \psi^{(i)}(x)v(x) dx$ , where  $\psi^{(i)} \in \mathbb{V}$  is a so-called *test function*. In this case, the test functionals define a weighted average of the current residual, giving rise to the name of the method.

To reduce the number of unknowns, MWRs also often approximate the unknown solution function  $\mathbf{u}$  via finite linear combinations of *trial functions*  $\phi^{(1)}, \dots, \phi^{(m)} \in \mathbb{U}$ , i.e.

$$\mathbf{u} \approx \hat{\mathbf{u}} := \sum_{i=1}^m c_i \phi^{(i)}, \quad (2.7)$$

where  $\mathbf{c} \in \mathbb{R}^m$  is the coordinate vector of  $\hat{\mathbf{u}}$  in the finite-dimensional subspace  $\hat{\mathbb{U}} := \text{span}(\phi^{(1)}, \dots, \phi^{(m)}) \subset \mathbb{U}$ . By substituting Equation (2.7) into Equation (2.6), we arrive at a linear system  $\hat{\mathbf{D}}\mathbf{c} = \hat{\mathbf{f}}$ , where  $\hat{D}_{ij} := l^{(i)}[\mathcal{D}[\phi^{(j)}]]$  and  $\hat{f}_i := l^{(i)}[f]$ . Hence, the approximate solution function obtained from this method is given by

$$\mathbf{u}^{\text{MWR}} = \sum_{i=1}^m c_i^{\text{MWR}} \phi^{(i)}, \quad \text{where} \quad \mathbf{c}^{\text{MWR}} = \hat{\mathbf{D}}^{-1} \hat{\mathbf{f}} \quad (2.8)$$

assuming that  $\hat{\mathbf{D}}$  is invertible. Above, we implicitly assume that the trial functions  $\phi^{(i)}$  satisfy the boundary conditions, i.e. we describe so-called *interior methods*.<sup>3</sup>

The procedure outlined above can also be applied to approximate weak solutions to linear PDEs by simply substituting  $\mathcal{D} \leftarrow \mathcal{D}^w$ ,  $f \leftarrow f^w$ , and  $\mathbb{V} \leftarrow \mathbb{V}'$ . In this case, it is customary to employ test functionals  $l^{(i)} \in \mathbb{V}''$  induced by test functions  $\psi^{(i)} \in \mathbb{V}$  such that  $l^{(i)}[\mathcal{D}^w[\mathbf{u}]] = \mathcal{D}^w[\mathbf{u}][\psi^{(i)}]$  and  $l^{(i)}[f^w] = f^w[\psi^{(i)}]$ .<sup>4</sup> In particular, in the example from Section 2.1.1, this implies  $l^{(i)}[\mathcal{D}^w[\mathbf{u}]] = B[\mathbf{u}, \psi^{(i)}]$  and  $l^{(i)}[f^w] = \langle \dot{q}_V, \psi^{(i)} \rangle_{L_2(\mathbb{D})}$ . Following Fletcher (1984), we will also refer to these methods as methods of weighted residuals.

Table 1 lists the aforementioned examples of MWRs together with the corresponding trial and test function(al)s that induce the method.

## 2.2 Gaussian Processes

A *Gaussian process* (GP)  $f$  with index set  $\mathbb{X}$  is a family  $\{f_{\mathbf{x}}\}_{\mathbf{x} \in \mathbb{X}}$  of real-valued random variables on a common probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , such that, for each finite set of indices  $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{X}$ , the joint distribution of  $f_{\mathbf{x}_1}, \dots, f_{\mathbf{x}_n}$  is Gaussian. We also write  $f(\mathbf{x}) := f_{\mathbf{x}}$  and  $f(\mathbf{x}, \omega) := f_{\mathbf{x}}(\omega)$ . The function  $\mathbf{x} \mapsto \mathbb{E}_{\mathbb{P}}[f(\mathbf{x})]$  is called the *mean (function)* of  $f$  and the function  $(\mathbf{x}_1, \mathbf{x}_2) \mapsto \text{Cov}_{\mathbb{P}}[f(\mathbf{x}_1), f(\mathbf{x}_2)]$  is called the *covariance function* or *kernel* of  $f$ . We write  $f \sim \mathcal{GP}(m, k)$  to indicate that  $f$  is a Gaussian process with mean function  $m$  and covariance function  $k$ . For each  $\omega \in \Omega$ , the function  $f(\cdot, \omega): \mathbb{X} \rightarrow \mathbb{R}$ ,  $\mathbf{x} \mapsto f(\mathbf{x}, \omega)$  is called a *sample* or *(sample) path* of the Gaussian process. We denote the set of all sample paths of  $f$  by  $\text{paths}(f) := \{f(\cdot, \omega): \omega \in \Omega\} \subset \mathbb{R}^{\mathbb{X}}$ .

The sample paths of Gaussian processes are always real-valued. However, especially in the context of PDEs, vector-valued functions are ubiquitous, e.g. when dealing with vector

3. By stacking the residuals corresponding to the PDE and the boundary conditions, the approach outlined here can be used to realized *mixed methods*, which solve the boundary value problem without requiring that  $\hat{\mathbf{u}}$  fulfills the boundary conditions by construction.

4. This uses the fact that there is an isometric embedding  $\iota: \mathbb{V} \rightarrow \mathbb{V}''$ ,  $v \mapsto (l \mapsto l[v])$ , where  $\mathbb{V}''$  denotes the strong bidual of  $\mathbb{V}$  (Yosida, 1995, Section IV.8).

fields such as the electric field. Fortunately, the index set of a Gaussian process can be chosen freely, which means that we can “emulate” vector-valued GPs. More precisely, a function  $\mathbf{f}: \mathbb{X} \rightarrow \mathbb{R}^{d'}$  can be equivalently viewed as a function  $\tilde{f}: \{1, \dots, d'\} \times \mathbb{X} \rightarrow \mathbb{R}$  with  $\tilde{f}(i, \mathbf{x}) := f_i(\mathbf{x})$ . Applying this construction to a Gaussian process leads to the notion of a multi-output Gaussian process: A  $d'$ -output Gaussian process  $\mathbf{f}$  with index set  $\mathbb{X}$  is a family  $\{\mathbf{f}_{\mathbf{x}}\}_{\mathbf{x} \in \mathbb{X}}$  of  $\mathbb{R}^{d'}$ -valued random variables on  $(\Omega, \mathcal{F}, \mathbb{P})$  such that  $\tilde{\mathbf{f}} := \{(\mathbf{f}_{\mathbf{x}})_i\}_{(i, \mathbf{x}) \in \{1, \dots, d'\} \times \mathbb{X}}$  is a Gaussian process. As before, we define  $\mathbf{f}(\mathbf{x}) := \mathbf{f}_{\mathbf{x}}$  and  $\mathbf{f}(\mathbf{x}, \omega) := \mathbf{f}_{\mathbf{x}}(\omega)$ . The mean function  $\mathbf{m}: \mathbb{X} \rightarrow \mathbb{R}^{d'}$  and covariance function  $\mathbf{k}: \mathbb{X} \times \mathbb{X} \rightarrow \mathbb{R}^{d' \times d'}$  of  $\mathbf{f}$  are defined by

$$\mathbf{m}(\mathbf{x}) = \begin{pmatrix} \tilde{m}(1, \mathbf{x}) \\ \vdots \\ \tilde{m}(d', \mathbf{x}) \end{pmatrix} \quad \text{and} \quad \mathbf{k}(\mathbf{x}_1, \mathbf{x}_2) = \begin{pmatrix} \tilde{k}((1, \mathbf{x}_1), (1, \mathbf{x}_2)) & \dots & \tilde{k}((1, \mathbf{x}_1), (d', \mathbf{x}_2)) \\ \vdots & \ddots & \vdots \\ \tilde{k}((d', \mathbf{x}_1), (1, \mathbf{x}_2)) & \dots & \tilde{k}((d', \mathbf{x}_1), (d', \mathbf{x}_2)) \end{pmatrix},$$

where  $\tilde{\mathbf{f}} \sim \mathcal{GP}(\tilde{\mathbf{m}}, \tilde{\mathbf{k}})$ , and we write  $\mathbf{f} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$ .

### 3. Learning the Solution to a Linear PDE

Consider a linear partial differential equation  $\mathcal{D}[\mathbf{u}] = f$  subject to linear boundary conditions  $\mathcal{B}[\mathbf{u}] = g$  as in Section 2.1. Our goal is to find a solution  $\mathbf{u} \in \mathbb{U}$  satisfying the PDE for (partially) known  $(\mathcal{D}, f)$  and  $(\mathcal{B}, g)$ . In general, one cannot find a closed-form expression for the solution  $\mathbf{u}$  (Borthwick, 2018). Therefore, we aim to compute an accurate approximation  $\hat{\mathbf{u}} \approx \mathbf{u}$  instead. Motivated by the challenges (C1) to (C3) of partial information inherent to numerically solving PDEs, we approach the problem from a statistical inference perspective. In other words, we will *learn* the solution of the PDE from multiple heterogeneous sources of information. This way we can quantify the epistemic uncertainty about the solution at any time during the computation, as Figure 1(a) illustrates.

**Indirectly Observing the Solution of a PDE** Typically, we think of observations as a finite number of direct measurements  $\mathbf{u}(\mathbf{x}_i) = \mathbf{y}_i$  of the latent function  $\mathbf{u}$ . As it turns out, we can generalize this notion of a measurement and even interpret the PDE itself as an (indirect) observation of  $\mathbf{u}$ . As an example, consider the important case where  $\mathbf{u}$  models the state of a physical system. The laws of physics governing such a system are often formulated as *conservation laws* in the language of PDEs. For example, they may require physical quantities like mass, momentum, charge or energy to be conserved over time.

**Example 3.1** (Thermal Conduction and the Heat Equation). *Say we want to simulate heat conduction in a solid object with shape  $\mathbb{D} \subset \mathbb{R}^3$ , i.e. we want to find the time-varying temperature distribution  $u: [0, T] \times \mathbb{D} \rightarrow \mathbb{R}$ . Neglecting radiation and convection,  $u(t, \mathbf{x})$  is described by a linear PDE known as the heat equation (Lienhard and Lienhard, 2020). Assuming spatially and temporally uniform material parameters  $c_p, \rho, \kappa \in \mathbb{R}$ , it reduces to*

$$\left( c_p \rho \frac{\partial}{\partial t} - \kappa \Delta \right) u - \dot{q}_V = 0. \quad (3.1)$$

*Thermal conduction is described by  $-\kappa \Delta u$ , while  $\dot{q}_V$  are local heat sources, e.g. from electric currents. Any energy flowing into a region due to conduction or a heat source is balanced by an increase in energy of the material. The net-zero balance shows that energy is conserved.*



Notice how a conservation law is an *observation* of the behavior of the physical system! To formalize this, we begin by rephrasing the classical notion of an observation at a point  $\mathbf{x}_i$  as measuring the result of a specific linear operator applied to the solution  $\mathbf{u}$ :

$$\mathbf{u}(\mathbf{x}_i) = \mathbf{y}_i \iff \delta_{\mathbf{x}_i}[\mathbf{u}] = \mathbf{y}_i$$

where  $\delta_{\mathbf{x}_i}$  is the evaluation functional. Now, the key idea is to generalize the notion of a direct observation to collecting information about the solution via an arbitrary linear operator  $\mathcal{L}$  with values in  $\mathbb{R}^n$  applied to the solution  $\mathbf{u}$ , such that  $\mathcal{L}[\mathbf{u}] = \mathbf{y} \iff \mathcal{L}[\mathbf{u}] - \mathbf{y} = \mathbf{0}$ . The affine operator

$$\mathcal{I}[\mathbf{u}] := \mathcal{L}[\mathbf{u}] - \mathbf{y} \tag{3.2}$$

is a specific kind of *information operator* (Cockayne et al., 2019b). In this setting the information operator may describe a conservation law as in Equation (3.1), a general linear PDE of the form (2.1) or an arbitrary affine operator mapping a function space into  $\mathbb{R}^n$ . This generalized notion of an observation turns out to be very powerful to incorporate different kinds of mathematical, physical, or experimental properties of the solution. Since PDEs and conservation laws are often assumed to hold exactly, we focused on noise-free observations above. However, generally we are not limited to this case and can also model  $\mathbf{y}$  as random variable, in which case the information operator  $\mathcal{I}[(\mathbf{u}, \mathbf{y})]$  is a (jointly) linear functional of the solution  $\mathbf{u}$  and the right-hand side  $\mathbf{y}$ .

### 3.1 Solving PDEs as a Bayesian Inference Problem

One of the main challenges (C1) to (C3) outlined in the beginning is the limited computational budget available to us to approximate the solution. Fortunately, in practice, the solution  $\mathbf{u}$  is not hopelessly unconstrained, but we usually a-priori have information about it. At the very least, we know the space of functions  $\mathbb{U}$  in which to search for the solution. Additionally, we might have expert knowledge about its rough shape and value range, or solutions to related PDEs at our disposal. Now, the question becomes: How do we combine this prior knowledge with indirect observations of the solution through the information operator  $\mathcal{I}$  (3.2)? To do so, we turn to the Bayesian inference framework. This provides a different perspective on the numerical problem of solving a linear PDE as a *learning task*.

**Gaussian Process Inference** We represent our belief about the solution of the linear PDE via a (multi-output) Gaussian process  $\mathbf{u} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$  with mean function  $\mathbf{m}: \mathbb{D} \rightarrow \mathbb{R}^d$  and kernel  $\mathbf{k}: \mathbb{D} \times \mathbb{D} \rightarrow \mathbb{R}^{d \times d}$ . Gaussian processes are well-suited for this purpose since:

- (i) For an appropriate choice of kernel, the Gaussian process defines a probability measure over the function space in which the PDE’s solution is sought.
- (ii) Kernels provide a powerful modeling toolkit to incorporate prior information (e.g. variability, periodicity, multi-scale effects, in- / equivariances, . . .) in a modular fashion.
- (iii) Measurement noise often follows a Gaussian distribution.
- (iv) Conditioning a Gaussian process on observations made via a linear map again results in a Gaussian process.

While the result in (iv) is used ubiquitously in the literature, its general form where observations are made via *arbitrary* linear operators with values in  $\mathbb{R}^n$  as opposed to *finite-dimensional* linear maps, has only been rigorously demonstrated for Gaussian *measures* on separable Hilbert spaces, not for the Gaussian *process* perspective, to the best of our knowledge. The two perspectives are closely related, but there are thorny technical difficulties to consider. We intentionally frame the problem from the Gaussian process perspective to make use of the expressive modeling capabilities provided by the kernel. Our framework at its core relies on this result, which we explain in detail in Section 4 and prove in Appendix B.3.

### 3.1.1 ENCODING PRIOR KNOWLEDGE ABOUT THE SOLUTION

We can infer the solution of a linear PDE more quickly by specifying inductive biases in the prior, which can encode both provable and approximately known properties of the solution.<sup>5</sup>

**Function Space of the Solution** The most basic known property derived from the PDE is an appropriate choice of function space for the solution. For strong solutions, this can be done by inspecting the differential operator  $\mathcal{D}$  and keeping track of the partial derivatives. In fact, in implementation this can be automatically derived solely from the problem definition, e.g. by compositionally defining differential operators and storing information on the necessary differentiability. Let  $\beta_i \in \mathbb{N}_0$  be the number of times any partial derivative in the differential operator  $\mathcal{D}$  differentiates w.r.t. the variable  $x_i$ .<sup>6</sup> Then a sensible choice of solution space is the space  $\mathbb{U} = C^\beta(\overline{\mathbb{D}})$  (see Appendix B.2). To define a prior with paths in this solution space, a common choice of prior covariance function is the tensor product of one-dimensional half-integer Matérn kernels  $k_{\nu_i}$  with  $\nu_i = \beta_i + \frac{1}{2}$  (see Appendix B.4). For weak solutions, the Sobolev spaces  $\mathbb{U} = H^m(\mathbb{D})$  are prototypical choices of solution spaces. In this case, a (multivariate) Matérn kernel with smoothness parameter  $\nu = m + \frac{1}{2}$  is a useful default prior covariance function. In both cases, a parametric kernel  $k(\mathbf{x}_0, \mathbf{x}_1) = \phi(\mathbf{x}_0)^\top \Sigma \phi(\mathbf{x}_1)$  is also a valid choice if  $\phi_i \in \mathbb{U}$ . See Appendix B.4.2 for a detailed account on how to choose priors for physics-informed GP regression.

**Symmetries, In- and Equivariances** Many solutions of PDEs exhibit a-priori known symmetries. For example, to calculate the strength of a magnet rotated by  $\mathbf{R} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ , one can equivalently compute the field  $\mathbf{B}$  of the magnet in its original position and rotate the field, i.e.  $\mathbf{B}(\mathbf{R}\mathbf{x}) = \mathbf{R}\mathbf{B}(\mathbf{x})$ . Inductive biases reflecting symmetries can be encoded via kernels that are *invariant*  $\mathbf{k}(\rho_g \mathbf{x}_0, \rho_g \mathbf{x}_1) = \mathbf{k}(\mathbf{x}_0, \mathbf{x}_1)$ , or *equivariant*  $\mathbf{k}(\rho_g \mathbf{x}_0, \rho_g \mathbf{x}_1) = \rho_g \mathbf{k}(\mathbf{x}_0, \mathbf{x}_1) \rho_g^*$ , where  $\rho_g$  is a unitary group representation. The most commonly used kernels are stationary, i.e. translation invariant, but one can also construct invariant (Haasdonk and Burkhardt, 2007; Azangulov et al., 2022), as well as equivariant kernels (Reisert and Burkhardt, 2007; Holderrieth et al., 2021) for many other group actions.

**Related Problems** Given a set of solutions from related problems, the prior mean function can be set to a combination thereof and the prior kernel can then be chosen to reflect how related the problems are. For example, if we have an approximate solution of a PDE

5. In the special case of GP regression, if the prior smoothness matches the smoothness of the target function  $\mathbf{u}$ , the convergence rate is optimal in the number of observations (Kanagawa et al., 2018, Thm. 5.1).

6. Formally,  $\beta \in \mathbb{N}_0^d$  is the “smallest” multi-index such that  $\alpha \leq \beta$  for every multi-index of a partial derivative occurring in  $\mathcal{D}$  (see Definitions C.1 and C.2).

computed on a coarser mesh, we can condition our function space prior on the coarse solution with a noise level reflecting the fidelity of the discretization. Similarly, if we solved the same PDE with different parameters, we can condition on the available solutions with a noise level chosen according to how similar the parameters are to the one of interest.

**Domain Expertise** Domain experts often have approximate knowledge of what solutions can be expected, either from experience, previous experiments, or familiarity with the physical interpretation of the solution  $\mathbf{u}$ . For example, an engineer who designs electrical components is likely able to give realistic temperature ranges for a component for which we aim to simulate the temperature distribution. This can be included by choosing the (initial) kernel hyperparameters, such as the output- and lengthscales based on this expertise.

### 3.1.2 (INDIRECTLY) OBSERVING THE SOLUTION

From a computational perspective, the most important reason for choosing Gaussian processes is that when conditioning on linear observations, the resulting posterior is again a Gaussian process with closed form mean and covariance function (Bishop, 2006). We extend this classic result from observations via a finite-dimensional linear map to general  $\mathbb{R}^n$ -valued linear operators in Theorem 1. This is crucial to condition on the different types of observations, most importantly the PDE itself, made via the information operator in (3.2). Given such an affine observation defined via a linear operator  $\mathcal{L}: \mathbb{U} \rightarrow \mathbb{R}^n$  and an independent Gaussian random variable  $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ , we can condition our prior belief using Theorem 1 on the observations to obtain a posterior of the form  $\mathbf{u} \mid (\mathcal{L}[\mathbf{u}] + \boldsymbol{\epsilon} = \mathbf{y}) \sim \mathcal{GP}(\mathbf{m}^{\mathbf{u}|\mathbf{y}}, \mathbf{k}^{\mathbf{u}|\mathbf{y}})$  with mean and covariance function given by

$$m_i^{\mathbf{u}|\mathbf{y}}(\mathbf{x}) = m_i(\mathbf{x}) + \mathcal{L}[k_{:,i}(\cdot, \mathbf{x})]^\top (\mathcal{L}\mathbf{k}\mathcal{L}' + \boldsymbol{\Sigma})^\dagger (\mathbf{y} - (\mathcal{L}[\mathbf{m}] + \boldsymbol{\mu})), \quad (3.3)$$

$$k_{i,j}^{\mathbf{u}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) = k_{i,j}(\mathbf{x}_1, \mathbf{x}_2) + \mathcal{L}[k_{:,i}(\cdot, \mathbf{x}_1)]^\top (\mathcal{L}\mathbf{k}\mathcal{L}' + \boldsymbol{\Sigma})^\dagger \mathcal{L}[k_{:,j}(\cdot, \mathbf{x}_2)]. \quad (3.4)$$

We will now look more closely at how we can condition on the boundary conditions, the PDE itself and direct measurements of the solution.

**Observing the Solution via the PDE** The differential operator  $\mathcal{D}$  in Equation (2.1) is linear and therefore it is tempting to define the information operator  $\mathcal{I}[\mathbf{u}] = \mathcal{D}[\mathbf{u}] - f$  and attempt to condition on  $\mathcal{I}[\mathbf{u}] = 0$ . Under some assumptions on  $\mathbb{U}$ ,  $\mathcal{D}$ , and  $\mathbf{u}$ , one can even show that this is well-defined. Unfortunately, it turns out that computing the posterior moments is then at least as hard as solving the PDE directly and thus typically intractable in practice. Loosely speaking, this is because  $f$  is a function and hence  $\mathcal{D}[\mathbf{u}] = f$  corresponds to an infinite number of observations. However, by only enforcing the PDE at a finite number of points in the domain, we can immediately give a canonical example of an approximation to this intractable information operator. Concretely, we can condition  $\mathbf{u}$  on the fact that the PDE holds at a finite sequence of well-chosen domain points  $\mathbf{X}_{\text{PDE}} = (\mathbf{x}_i)_{i=1}^n \in \mathbb{D}^n$ , i.e. we compute  $\mathbf{u} \mid (\mathcal{D}[\mathbf{u}](\mathbf{X}_{\text{PDE}}) - f(\mathbf{X}_{\text{PDE}}) = 0)$  by choosing  $\mathcal{L} = \delta_{\mathbf{X}_{\text{PDE}}} \circ \mathcal{D}$  and  $\mathbf{y} = f(\mathbf{X}_{\text{PDE}})$ . If the set  $\mathbf{X}_{\text{PDE}}$  of domain points is dense enough, we obtain a good approximation to the exact conditional process. This approach, known as the *probabilistic meshless method* (Cockayne et al., 2017), is analogous to existing non-probabilistic approaches to solving PDEs, commonly referred to as *collocation methods*, wherein the points  $\mathbf{X}$  are called *collocation points*.

Satisfying the PDE at a set of collocation points is far from the only choice within our general framework. For example, we can choose a set of test functions  $l^{(i)} \in \mathbb{V}'$ , which we use to observe the PDE with, such that  $\mathcal{L}_i[\mathbf{u}] = l^{(i)}[\mathcal{D}[\mathbf{u}]]$  and  $\mathbf{y}_i = l^{(i)}[f]$ . For efficient evaluation of the differential operator we can further represent the solution in a basis of trial functions from a subspace  $\hat{\mathbb{U}}$ , resulting in  $\mathcal{L}_i[\mathbf{u}] = l^{(i)}[\mathcal{D}[\mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{u}]]]$ . This turns out to be very powerful and is analogous to some of the most successful classical PDE solvers. In fact, for certain priors and choices of subspaces, our framework recovers several important classic solvers in the posterior mean (see Section 3.3.4). The above can be applied to both time-dependent and time-independent PDEs and regardless of the type of linear PDE (e.g. elliptic, parabolic, hyperbolic). Moreover, an extension to systems of linear PDEs is straightforward.

**Observing the Solution at the Boundary** As for the PDE, we could attempt to directly condition on the boundary conditions by choosing  $\mathcal{L} = \mathcal{B}$  and  $\mathbf{y} = g$ . However, we are faced with the same intractability issues that we discussed above. Instead, we observe that the boundary conditions hold at a finite set of points  $\mathbf{X}_{\text{BC}} \subset \partial\mathbb{D}$ , i.e.  $\mathcal{L} = \delta_{\mathbf{X}_{\text{BC}}} \circ \mathcal{B}$  and  $\mathbf{y} = g(\mathbf{X}_{\text{BC}})$ . In practice, sometimes the boundary conditions are only known at a finite set of points making this a natural choice.

**Observing the Solution Directly** Finally, as in standard GP regression, we can directly condition on (noisy) measurements of the solution, for example from a real world experiment, by choosing  $\mathcal{L} = \delta_{\mathbf{X}_{\text{MEAS}}}$  and  $\mathbf{y} = \mathbf{u}^*(\mathbf{X}_{\text{MEAS}})$ .

In summary, the probabilistic viewpoint allows us to

- encode prior information about the solution,
- condition on various kinds of (partial) information, such as the boundary condition, the PDE itself, or direct measurements, and
- output a structured error estimate, reflecting all obtained information and performed computation.

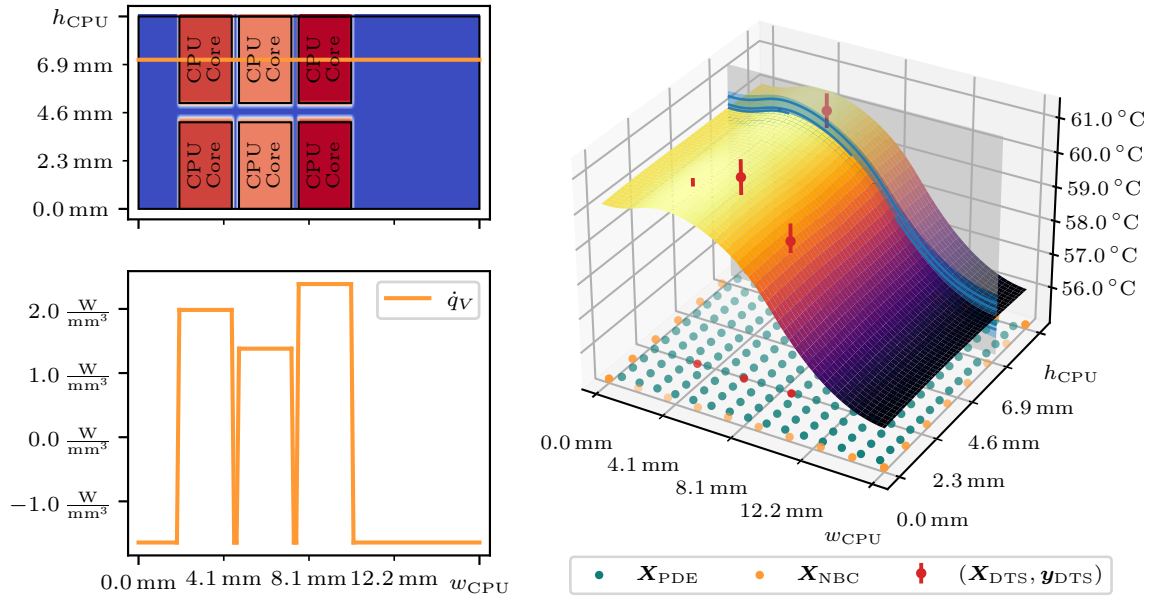
We will now give concrete examples for some of the possible modeling choices described above in a case study.

### 3.2 Case Study: Modeling the Temperature Distribution in a CPU

Central processing units (CPUs) are pieces of computing hardware that are constrained by the vast amounts of heat they dissipate under computational load. Surpassing the maximum temperature threshold of a CPU for a prolonged period of time can result in reduced longevity or even permanent hardware damage (Michaud, 2019). To counteract overheating, cooling systems are attached to the CPU, which are controlled by digital thermal sensors (DTS). For simplicity, assume that the CPU is under sustained computational load and that the cooling device is controlled in a way such that the die reaches thermal equilibrium.

**Example 3.2** (Stationary Heat Equation). *The temperature distribution of a solid at thermal equilibrium, i.e.  $\frac{\partial u}{\partial t} = 0$  in Example 3.1, is described by the linear PDE*

$$-\kappa\Delta u - \dot{q}_V = 0, \tag{3.5}$$



(a) Top: CPU die with CPU cores as heat sources and uniform cooling over the whole surface. Bottom: Magnitude of heat sources and sinks  $\dot{q}_V$  in the 1D slice in the upper subplot (—).


(b) Gaussian process integrating *prior information* about the temperature distribution, a *mechanistic model* of heat conduction in the form of a linear PDE, and *empirical measurements* ( $\mathbf{X}_{\text{DTS}}, \mathbf{y}_{\text{DTS}}$ ) taken by limited-precision sensors (DTS). The plot shows the GP mean and a 1D slice illustrating the posterior uncertainty along with a few samples.

Figure 2: Physics-informed Gaussian process model of the stationary temperature distribution in an idealized hexa-core CPU die under sustained computational load.

known as the stationary heat equation (Lienhard and Lienhard, 2020). For our choice of material parameters Equation (3.5) is equivalent to the Poisson equation with  $f = \frac{\dot{q}_V}{\kappa}$ .

While the sensors control cooling, they only provide local, limited-precision measurements of the CPU temperature. This is problematic, since the chip may reach critical temperature thresholds in unmonitored regions. Therefore, our goal will be to infer the temperature in the entire CPU. We will use our framework to integrate the physics of heat flow, the controlled cooling at the boundary, and the noisy temperature measurements from the sensors. See Figure 2(b) for an illustration of the result. During manufacturing, the resulting belief over the temperature distribution could then help decide whether the CPU design needs to be changed to avoid premature failure. From here on out, we focus on a 1D slice across the CPU surface, as shown in Figure 2(a) (top), to easily visualize uncertainty.

**Encoding Prior Knowledge** By inspecting the PDE’s differential operator  $\mathcal{D} = -\kappa\Delta = -\kappa \sum_{i=1}^d \frac{\partial^2}{\partial x_i^2}$ , we can deduce that the paths of our Gaussian process need to be twice-differentiable in every input variable  $x_i$ . The construction in Appendix B.4.1 tells us that that a GP prior whose covariance function is a tensor product  $k_{\nu}$  of one-dimensional Matérn




(a) Gaussian process prior with a Matérn- $\frac{5}{2}$  kernel over the temperature distribution of the CPU. (b) Prior under the differential operator  $\mathcal{D} = -\kappa\Delta$  along with heat sources and sinks  $\dot{q}_V$ .

Figure 3: Prior model for the stationary temperature distribution of a CPU die under load.

kernels  $k_{\nu_i}$  with  $\nu_i = 2 + \frac{1}{2} = \frac{5}{2}$  fulfills the desired path properties. Assume we also know what temperature ranges are plausible from similar CPU architectures, meaning we set the kernel output scale to  $\sigma_{\text{out}}^2 = 9$ . Figure 3 shows the prior process  $u$  on along with its image  $\mathcal{D}[u] \sim \mathcal{GP}(\mathcal{D}[m], \sigma_{\text{out}}^2 \mathcal{D}k\mathcal{D}')$  under the differential operator. A draw from  $\mathcal{D}[u]$  can be interpreted as the heat sources and sinks that generated the corresponding temperature distribution draw from  $u$ .

**Conditioning on the PDE** We can now inform our belief about the physics of heat conduction using the mechanistic model defined by the stationary heat equation. We choose a set of collocation points  $\mathbf{X}_{\text{PDE}} \in \mathbb{D}^n$  and then condition on the observation that the PDE holds (exactly) at these points. In other words, we compute the physically-informed Gaussian process  $u \mid \text{PDE} := u \mid (\mathcal{I}^{\text{PDE}}[u] = \mathbf{0})$  with  $\mathcal{I}^{\text{PDE}}[u] := -\kappa\Delta u(\mathbf{X}_{\text{PDE}}) - \dot{q}_V(\mathbf{X}_{\text{PDE}})$  visualized in Figure 4. We can see that the resulting conditional process indeed satisfies the PDE exactly at the collocation points (see Figure 4(b)). The remaining uncertainty in Figure 4(b) is due to the approximation error introduced by only conditioning on a finite number of collocation points. However, while the samples from our belief about the solution in Figure 4(a) exhibit much more similarity to the mean function and less spatial variation, the marginal uncertainty hardly decreases. The latter is explained by the PDE not identifying a unique solution, since adding any affine function to  $u^*$  does not alter its image under the differential operator, i.e.  $\Delta(\mathbf{a}^\top \mathbf{x} + b) = 0$ . There is an at least two-dimensional subspace of functions which can not be observed. This ambiguity can be resolved by introducing boundary conditions.

**Conditioning on the Boundary Conditions** We assume that the CPU cooler extracts heat uniformly from all exposed parts of the CPU, in particular also from the sides, rather than just from the top. Instead of directly specifying the value of the temperature distribution at the edge points of the CPU slice, we only know the density  $\dot{q}_A$  of heat flowing out of each point on the CPU's boundary based on the cooler specification. We can use another thermodynamical law to turn this assumption into information about the temperature distribution  $u$ .



(a) Belief about the solution after conditioning on the PDE at a set of collocation points.


(b) Belief about heat sources and sinks after conditioning on the PDE at collocation points.

Figure 4: We integrate mechanistic knowledge about the system by conditioning on PDE observations  $-\kappa\Delta u(\mathbf{X}_{\text{PDE}}) - \dot{q}_V(\mathbf{X}_{\text{PDE}}) = \mathbf{0}$  at the collocation points  $\mathbf{X}_{\text{PDE}}$ , resulting in the conditional process  $u \mid \text{PDE}$ . The large remaining uncertainty in Figure 4(a) illustrates that the PDE by itself does not identify a unique solution.

**Example 3.1** (continuing from p. 8). *Fourier’s law states that the local density of heat  $\dot{q}_A$  flowing through a surface with normal vector  $\boldsymbol{\eta}$  is proportional to the inner product of the negative temperature gradient and the surface normal  $\boldsymbol{\eta}$ , i.e.  $\dot{q}_A = -\kappa\langle\boldsymbol{\eta}, \nabla u\rangle$ , where  $\kappa$  is the material’s thermal conductivity in  $\text{W m}^{-1} \text{K}$  (Lienhard and Lienhard, 2020).*

Assuming sufficient differentiability of  $u$ , the inner product above is equal to the directional derivative  $\partial_{\boldsymbol{\eta}}u$  of  $u$  in direction  $\boldsymbol{\eta}$ . We can assign an outward-pointing vector  $\boldsymbol{\eta}(x)$  (almost) everywhere on the boundary of the domain. Since the boundary of the CPU domain is its surface, we can summarize the above in a Neumann boundary condition  $-\kappa\partial_{\boldsymbol{\eta}(x)}u(x) = \dot{q}_A(x)$  for  $x \in \partial\mathbb{D}$ . Applying Corollary 2 once more, we can inform our estimate of the solution about the boundary conditions by computing  $u \mid \text{PDE, NBC} := (u \mid \text{PDE}) \mid \mathcal{I}^{\text{NBC}}[(u, \dot{q}_A)] = \mathbf{0}$ , where  $\mathcal{I}^{\text{NBC}}[u] := -\kappa\partial_{\boldsymbol{\eta}(\mathbf{X}_{\text{NBC}})}u(\mathbf{X}_{\text{NBC}}) - \dot{q}_A(\mathbf{X}_{\text{NBC}})$  with  $\mathbf{X}_{\text{NBC}} = \{0, w_{\text{CPU}}\}$  is the information operator induced by the boundary conditions. The result is visualized in Figure 5(a). The structure of the samples illustrates that most of the remaining uncertainty about the solution lies in a one-dimensional subspace of  $\mathbb{U}$  corresponding to constant functions. This is due to the fact that two Neumann boundary conditions on both sides of the domain only determine the solution of the PDE up to an additive constant. We need an additional source of information to address the remaining degree of freedom.

**Conditioning on Direct Measurements** Fortunately, CPUs are equipped with digital thermal sensors (DTS) located close to each of the cores, which provide (noisy) local measurements of the core temperatures (Michaud, 2019). These measurements can be straightforwardly accounted for in our model by performing standard GP regression using  $u \mid \text{PDE, NBC}$  from Figure 5(a) as a prior. The resulting belief about the temperature distribution is visualized in Figure 5(b). We can see that integrating the interior measurements effectively reduces the uncertainty due to the remaining degree of freedom, albeit not completely. The remaining uncertainty is due to the model’s consistent accounting for noise in



(a) Belief about the solution after conditioning on the PDE and boundary conditions (BCs). (b) Belief about the solution after conditioning on the PDE, BCs and noisy sensor data.

Figure 5: Neumann boundary conditions encoding mechanistic knowledge about the heat flux across the boundary of the CPU and a sparse set of limited-precision measurements of the temperature distribution made by digital thermal sensors (DTS) located at the points  $\mathbf{X}_{\text{DTS}}$  further constrain the solution of the PDE. The remaining uncertainty is due to measurement noise and discretization error.

the thermal sensor readings, the uncertainty about the cooling, i.e. the boundary conditions, and the discretization error incurred by only choosing a small set of collocation points.

**Uncertainty in the Right-hand Side and the Boundary Function** Above, we assumed the true heat source term  $\dot{q}_V$ , i.e. the right-hand side of the PDE, and the boundary heat flux  $\dot{q}_A$  to be known exactly. However, in practice, this is rarely the case. Fortunately, our probabilistic viewpoint admits a straightforward relaxation of this assumption. Namely, we can replace  $\dot{q}_V$  and  $\dot{q}_A$  by a joint Gaussian process prior  $(\dot{q}_V, \dot{q}_A)$ , whose means are given by estimates of  $\dot{q}_V$  and  $\dot{q}_A$ .<sup>7</sup> Above, we assumed that the cooler is controlled in such a way, that the temperature distribution in the CPU does not change over time. However, a naive prior  $(\dot{q}_V, \dot{q}_A)$  may break this assumption. We need to encode that the amount of heat entering the CPU is equal to the amount of heat leaving the CPU via its boundary, i.e.

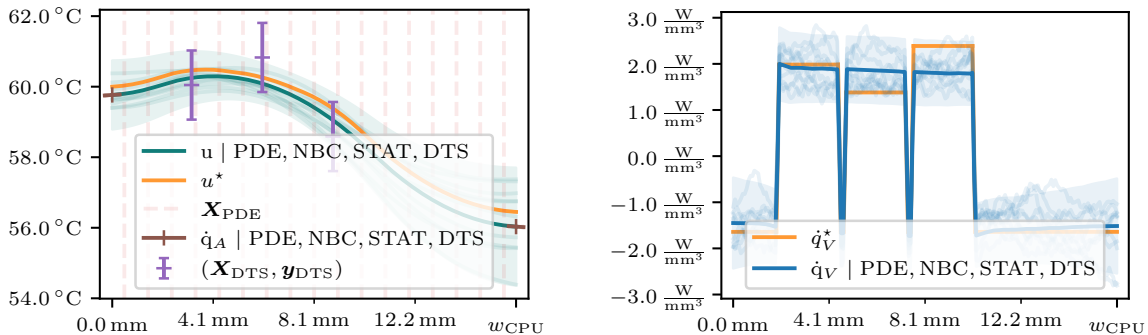
$$\mathcal{I}^{\text{STAT}}[(\dot{q}_V, \dot{q}_A)] := \int_{\mathbb{D}} \dot{q}_V(\mathbf{x}) \, d\mathbf{x} - \int_{\partial\mathbb{D}} \dot{q}_A(\mathbf{x}) \, dA = 0, \quad (3.6)$$

The (jointly) linear information operator  $\mathcal{I}^{\text{STAT}}$  computes the net amount of thermal energy that the CPU gains per unit time. Using Theorem 1 we can construct a multi-output GP prior  $(u, \dot{q}_V, \dot{q}_A)$ , which is consistent with the assumption of thermal stationarity by conditioning on  $\mathcal{I}^{\text{STAT}}[(\dot{q}_V, \dot{q}_A)] = 0$ . Here, we assume a-priori that  $u$ ,  $\dot{q}_V$ , and  $\dot{q}_A$  are pairwise independent. In the one-dimensional model, we can simplify Equation (3.6) by assuming that heat is drawn uniformly from the sides of the CPU. By encoding this information in the prior  $\dot{q}_A$ , the information operator corresponding to thermal stationarity resolves to

$$\mathcal{I}^{\text{STAT}}[(\dot{q}_V, \dot{q}_A)] = h_{\text{CPU}} \int_0^{w_{\text{CPU}}} \dot{q}_V(x) \, dx - h_{\text{CPU}} (\dot{q}_A(0) + \dot{q}_A(w_{\text{CPU}})). \quad (3.7)$$

7. Technically speaking, if the right-hand-side of the PDE is given as a Gaussian process, the PDE turns into a stochastic partial differential equation (SPDE).





(a) Posterior belief about the temperature distribution physically consistent with the assumption of stationarity.

(b) Posterior belief about the heat sources and sinks after conditioning on the corresponding uncertain right-hand-side  $\dot{q}_V$  of the PDE.

Figure 6: We integrate information from the joint prior  $(u, \dot{q}_V, \dot{q}_A) | \text{STAT}$  over the solution, the right-hand side of the PDE, and the values of the Neumann boundary conditions into our belief about the temperature distribution by conditioning on said PDE and boundary conditions.




Figure 7: Representation of the CPU model as a directed graphical model. The inference procedure described in Section 3.2 is equivalent to the *junction tree algorithm* (Bishop, 2006, Section 8.4.6) applied to the graphical model above. This example shows that the language of information operators is a powerful tool for aggregating heterogeneous sources of partial information in a joint probabilistic model.

As above, we can now use Corollary 2 to condition our physically-consistent GP prior  $(u, \dot{q}_V, \dot{q}_A) | \text{STAT} := (u, \dot{q}_V, \dot{q}_A) | (\mathcal{I}^{\text{STAT}}[(\dot{q}_V, \dot{q}_A)] = 0)$  on  $\mathcal{I}^{\text{PDE}}[(u, \dot{q}_V)] = \mathbf{0}$ , as well as  $\mathcal{I}^{\text{NBC}}[(u, \dot{q}_A)] = \mathbf{0}$  and the noisy measurements of the temperature distribution. Here, it is important to keep track of the cross-covariances in  $(u, \dot{q}_V, \dot{q}_A)$ , since the outputs in  $(\dot{q}_V, \dot{q}_A) | \text{STAT}$  become correlated. The resulting process  $u | \text{PDE, NBC, STAT, DTS}$  (or rather its marginals) is shown in Figure 6. Comparing Figures 5(b) and 6(a), we can see that, due to the uncertainty in the right-hand side  $\dot{q}_V$  of the PDE, the samples of  $u | \text{PDE, NBC, STAT, DTS}$  exhibit much more spatial variation. Moreover, the samples of the GP posterior over  $\dot{q}_V$  fulfill the stationarity constraint we imposed.

**Summary** Stepping back, we can view the problem of modeling the CPU under computational load as a scientific inference problem, where we need to aggregate heterogeneous

sources of information in a joint probabilistic model. This inference task is illustrated as a directed graphical model in Figure 7. Our physics-informed regression framework is a local computation in the global inference procedure on the graph. Importantly, its implementation does not change based on what happens to the solution estimate and the input data in either upstream or downstream computations. All this information is already handily encoded in the structured uncertainties of the Gaussian processes.

### 3.3 A General Class of Tractable Information Operators for Linear PDEs

Recall that conditioning on the linear PDE directly via the information operator  $\mathcal{I}[\mathbf{u}] = \mathcal{D}[\mathbf{u}] - f$  is usually intractable. Instead, in Section 3.1.2 we approximated this information operator by  $\mathcal{I}: \mathbb{U} \rightarrow \mathbb{R}^n$  with  $\mathcal{I}_i[\mathbf{u}] := \mathcal{D}[\mathbf{u}](\mathbf{x}_i) - f(\mathbf{x}_i)$  where  $\mathbf{x}_i \in \mathbb{D}$ . This implicitly assumes that point evaluation on both  $\mathcal{D}[\mathbf{u}]$  and  $f$  is well-defined, which crucially means that this approach applies *only* to strong solutions of PDEs. In this section, we extend this approximation scheme to a general class of tractable information operators aimed at approximating *both* weak and strong solutions to linear PDEs. Our framework is inspired by the method of weighted residuals (MWR) (see Section 2.1.2). In fact, in Section 3.3.4 we will show that GP inference with information operators in this class reproduces any weighted residual method in the posterior mean while additionally providing an estimate of the inherent approximation error.

In the following, we will consider both weak and strong formulations of linear PDEs, which is why we introduce the unifying notation  $\mathcal{D}^{(w)}[\mathbf{u}] = f^{(w)}$ . For a strong formulation,  $\mathcal{D}^{(w)} := \mathcal{D}$ , where  $\mathcal{D}: \mathbb{U} \mapsto \mathbb{V}$  is a linear differential operator (see Definition C.2), and  $f^{(w)} := f \in \mathbb{V}$  is the right-hand side function. In the context of a weak formulation,  $\mathcal{D}^{(w)} := \mathcal{D}^w$ , where  $\mathcal{D}^w: \mathbb{U} \mapsto \mathbb{V}'$  is the weak differential operator, and  $f^{(w)} := f^w \in \mathbb{V}'$  is the right-hand side functional (see Section 2.1.1). Following Section 2.1.2, we will apply linear functionals to the PDE residual. To facilitate notation, we define the shorthand  $\mathbb{L}^{(w)}$  for the space of continuous linear functionals on the image space of  $\mathcal{D}^{(w)}$ , i.e.  $\mathbb{L}^{(w)} := \mathbb{V}'$  in the context of a strong formulation and  $\mathbb{L}^{(w)} := \mathbb{V}''$  in the context of a weak formulation. We additionally require that  $l \circ \mathcal{D}^{(w)}$  is continuous for every  $l \in \mathbb{L}^{(w)}$ .

Let  $\mathbf{u} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$  be a Gaussian process prior over the solution  $\mathbf{u}$  of the PDE, whose path space can be continuously embedded into the solution space  $\mathbb{U}$  (see Appendix B.4 for more details on the latter assumption). It is intractable to condition the GP prior on the full information provided by the PDE via the family  $\{\mathcal{I}_l\}_{l \in \mathbb{L}^{(w)}}$  of affine information operators  $\mathcal{I}_l[\mathbf{u}] := (l \circ \mathcal{D}^{(w)})[\mathbf{u}] - l[f^{(w)}]$ , since  $\mathbb{L}^{(w)}$  is typically infinite-dimensional. To identify tractable families of information operators, we take inspiration from the method of weighted residuals.

#### 3.3.1 INFINITE-DIMENSIONAL TRIAL FUNCTION SPACES

Using Theorem 1 we can tractably condition on a finite subfamily  $\{\mathcal{I}_{l^{(i)}}\}_{i=1}^n \subset \{\mathcal{I}_l\}_{l \in \mathbb{L}^{(w)}}$  of information operators, where  $\{l^{(i)}\}_{i=1}^n \subset \mathbb{L}^{(w)}$  is a finite subset of test functionals, as long as we can compute  $\mathcal{I}_{l^{(i)}}[\mathbf{m}]$ ,  $\mathcal{L}_{l^{(i)}}[\mathbf{k}_{:,j}(\cdot, \mathbf{x})]$ , and  $\mathcal{L}_{l^{(i)}} \mathbf{k} \mathcal{L}'_{l^{(i)}}$ , where  $\mathcal{L}_{l^{(i)}} = l^{(i)} \circ \mathcal{D}^{(w)}$ . This might not always be possible in closed-form, since  $\mathcal{L}_{l^{(i)}}$  often involves computing integrals. However, in these cases one could fall back to an efficient numeric quadrature method, since

the integrals are often low-dimensional (typically at most four-dimensional). A prominent example of this approach is the *probabilistic meshless method* used in Section 3.

**Example 3.3** (Symmetric Collocation). *If the PDE is in strong formulation, then  $l^{(i)} = \delta_{\mathbf{x}_i} \in \mathbb{V}'$  with  $\mathbf{x}_i \in \mathbb{D}$  is a valid test functional, which induces the information operator*

$$\mathcal{I}_{l^{(i)}}[\mathbf{u}] = \mathcal{D}[\mathbf{u}](\mathbf{x}_i) - f(\mathbf{x}_i),$$

*i.e. we recover the probabilistic meshless method by Cockayne et al. (2017). They show that the conditional mean of this approach reproduces symmetric collocation (Fasshauer, 1997, 1999), a well-known method to approximate strong solutions of PDEs.*

The probabilistic meshless method can only be used to approximate strong solutions of linear PDEs, since point evaluation functionals are not well-defined on the image space  $\mathbb{V}'$  of  $\mathcal{D}^w$ . However, other choices of the  $l^{(i)}$  lead to approximation schemes for weak solutions.

**Example 3.4** (Weak Formulations). *Consider a linear PDE in weak formulation. As mentioned in Section 2.1.2, it is customary to use test functionals  $l^{(i)}$ , which are induced by test functions  $\psi^{(i)} \in \mathbb{V}$ , i.e.*

$$\mathcal{I}_{l^{(i)}}[\mathbf{u}] = \mathcal{D}^w[\mathbf{u}][\psi^{(i)}] - f^w[\psi^{(i)}]. \quad (3.8)$$

*For instance, if  $\mathbb{D} = ]l, r[ \subset \mathbb{R}$ , then a valid set of test functions for the weak formulation from Section 2.1.1 is given by*

$$\psi^{(i)} = \begin{cases} \frac{x-x_{i-1}}{x_i-x_{i-1}} & \text{if } x_{i-1} \leq x \leq x_i, \\ \frac{x_{i+1}-x}{x_{i+1}-x_i} & \text{if } x_i \leq x \leq x_{i+1}, \\ 0 & \text{otherwise.} \end{cases} \in H_0^1(]l, r[), \quad (3.9)$$

*where  $l = x_0 < \dots < x_{n+1} = r$ . The test functions are visualized in Figure 8(a). For the weak formulation in Section 2.1.1, the information operator from Equation (3.8) is equivalent to  $\mathcal{I}_{l^{(i)}}[\mathbf{u}] = B[\mathbf{u}, \psi^{(i)}] - \langle f, \psi^{(i)} \rangle_{L_2}$ .*

### 3.3.2 FINITE-DIMENSIONAL TRIAL FUNCTION SPACES

As opposed to the methods outlined in Section 2.1.2, we did not need to choose a finite-dimensional subspace of trial functions to arrive at tractable information operators in Section 3.3.1. Nevertheless, in practice, it might still be desirable to specify a finite-dimensional trial function basis  $\phi^{(1)}, \dots, \phi^{(m)}$ , e.g. because

- we want to reproduce the output of a classical method in the posterior mean to use the GP solver as an uncertainty-aware drop-in replacement (see Corollary 3.3);
- the trial basis encompasses problem-specific knowledge, which is difficult to encode in the prior; or
- we want to solve the problem in a coarse-to-fine scheme, allowing for mesh refinement strategies, which are informed by the GP's uncertainty estimate.

Naively, one might achieve this goal by defining the prior over  $\mathbf{u}$  as a parametric Gaussian process with features  $\phi^{(i)}$ . However, this means the posterior can not quantify the inherent approximation error, since the GP has no support outside of the finite subspace of  $\mathbb{U}$  spanned by the trial functions. Consequently, we need to take a different approach. Starting from a general, potentially nonparametric prior over  $\mathbf{u}$ , we consider a bounded (potentially oblique) projection  $\mathcal{P}_{\hat{\mathbb{U}}}: \mathbb{U} \rightarrow \hat{\mathbb{U}}$  onto a subspace  $\hat{\mathbb{U}} \subset \mathbb{U}$ , i.e.  $\mathcal{P}_{\hat{\mathbb{U}}}^2 = \mathcal{P}_{\hat{\mathbb{U}}}$ ,  $\|\mathcal{P}_{\hat{\mathbb{U}}}\| < \infty$ , and  $\text{ran}(\mathcal{P}_{\hat{\mathbb{U}}}) = \hat{\mathbb{U}}$ . In general, this subspace need not be finite-dimensional. We apply  $\mathcal{P}_{\hat{\mathbb{U}}}$  to our GP prior over  $\mathbf{u}$ , which, by Corollary 2, results in another GP

$$\hat{\mathbf{u}} := \mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{u}] \sim \mathcal{GP} \left( \mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{m}], \mathcal{P}_{\hat{\mathbb{U}}}\mathbf{k}\mathcal{P}'_{\hat{\mathbb{U}}} \right),$$

with sample paths in  $\hat{\mathbb{U}}$ . This discards prior information about  $\ker(\mathcal{P}_{\hat{\mathbb{U}}})$ . Hence, especially in case  $\dim \hat{\mathbb{U}} < \infty$ , applying the information operators  $\mathcal{I}_{l^{(i)}}$  from Section 3.3.1 directly to  $\hat{\mathbf{u}}$  would suffer from similar problems as choosing a parametric prior. However,

$$\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}} := \mathcal{I}_{l^{(i)}} \circ \mathcal{P}_{\hat{\mathbb{U}}} = (l^{(i)} \circ \mathcal{D}^{(w)} \circ \mathcal{P}_{\hat{\mathbb{U}}})[\cdot] - l^{(i)}[f^{(w)}]$$

is a valid information operator for  $\mathbf{u}$ , which leads to a probabilistic generalization of the method of weighted residuals. This is why we refer to  $\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}$  as an *MWR information operator*.

The similarity to the method of weighted residuals is particularly prominent if we choose a finite-dimensional subspace  $\hat{\mathbb{U}} = \text{span}(\phi^{(1)}, \dots, \phi^{(m)})$  as in Section 2.1.2. In this case, there is a bounded linear operator  $\mathcal{P}_{\mathbb{R}^m}: \mathbb{U} \rightarrow \mathbb{R}^m$  such that

$$\mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{u}] = \sum_{i=1}^m c_i \phi^{(i)} =: \mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbb{U}}}[\mathbf{c}],$$

where the  $\mathbf{c} := \mathcal{P}_{\mathbb{R}^m}[\mathbf{u}] \in \mathbb{R}^m$  are the coordinates of  $\mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{u}]$  in  $\hat{\mathbb{U}}$  and  $\mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbb{U}}}: \mathbb{R}^m \rightarrow \hat{\mathbb{U}}$  is the canonical isomorphism between  $\mathbb{R}^m$  and  $\hat{\mathbb{U}}$ . Hence, we get the factorization

$$\mathcal{P}_{\hat{\mathbb{U}}} = \mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbb{U}}} \mathcal{P}_{\mathbb{R}^m}, \quad (3.10)$$

which implies that  $\hat{\mathbf{u}}$  is a parametric Gaussian process. Moreover,  $l^{(i)}[f^{(w)}] = \hat{f}_i$  and

$$(l^{(i)} \circ \mathcal{D}^{(w)} \circ \mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbb{U}}})[\mathbf{c}] = \sum_{i=1}^m c_i (l^{(i)} \circ \mathcal{D}^{(w)})[\phi^{(i)}] = (\hat{\mathbf{D}}\mathbf{c})_i,$$

where  $\hat{\mathbf{D}}$  and  $\hat{\mathbf{f}}$  are defined as in Section 2.1.2. Consequently, the MWR information operator is given by  $\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}[\mathbf{u}] = (\mathcal{I}_{\mathbb{R}^m} \circ \mathcal{P}_{\hat{\mathbb{U}}})[\mathbf{u}]_i$ , where  $\mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] := \hat{\mathbf{D}}\mathbf{c} - \hat{\mathbf{f}}$ . This illustrates that we are dealing with the hierarchical model

$$\begin{aligned} \mathbf{u} &\sim \mathcal{GP}(\mathbf{m}, \mathbf{k}) \\ \mathbf{c} \mid \mathbf{u} &\sim \delta_{\mathcal{P}_{\mathbb{R}^m}[\mathbf{u}]} \end{aligned}$$

with observations  $\mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] = \mathbf{0}$ , where  $\mathbf{c} \sim \mathcal{N}(\mathcal{P}_{\mathbb{R}^m}[\mathbf{m}], \mathcal{P}_{\mathbb{R}^m}\mathbf{k}\mathcal{P}'_{\mathbb{R}^m})$ . Inference in this model can be broken down into two steps. First, we update our belief about the solution's coordinates in  $\hat{\mathbb{U}}$  by computing the conditional random variable  $\mathbf{c} \mid \mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] = \mathbf{0}$ , which is

also Gaussian. If  $\hat{\mathbf{D}}$  is invertible and  $\mathbf{c}$  has full support on  $\mathbb{R}^m$ , then the law of  $\mathbf{c} \mid \mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] = \mathbf{0}$ , is a Dirac measure whose mean is given by the coordinates of the MWR approximation  $\mathbf{c}^{\text{MWR}} = \hat{\mathbf{D}}^{-1} \hat{\mathbf{f}}$  from Equation (2.8). Next, we can reuse precomputed quantities from the conditional moments of  $\mathbf{c} \mid \mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] = \mathbf{0}$ , such as the representer weights  $\mathbf{w} = (\hat{\mathbf{D}} \mathcal{P}_{\mathbb{R}^m} \mathbf{k} \mathcal{P}'_{\mathbb{R}^m} \hat{\mathbf{D}}^\top)^\dagger (\hat{\mathbf{f}} - \hat{\mathbf{D}} \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}])$  to efficiently compute the conditional random process

$$(\mathbf{u} \mid (\mathcal{I}_{\mathbb{R}^m} \circ \mathcal{P}_{\mathbb{R}^m})[\mathbf{c}] = \mathbf{0}) = (\mathbf{u} \mid \{\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}[\mathbf{u}] = 0\}_{i=1}^n),$$

i.e. the main quantity of interest. Assuming once more that  $\hat{\mathbf{D}}$  is invertible and  $\mathbf{c}$  has full support on  $\mathbb{R}^m$ , the remaining uncertainty of the conditional process lies in the kernel of  $\mathcal{P}_{\hat{\mathbb{U}}}$ , since the law of  $\mathbf{c} \mid \mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] = \mathbf{0}$ , is a Dirac measure and

$$(\mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{u}] \mid \{\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}[\mathbf{u}] = 0\}_{i=1}^n) = (\mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbb{U}}}[\mathbf{c}] \mid \mathcal{I}_{\mathbb{R}^m}[\mathbf{c}] = \mathbf{0}).$$

Thus, all remaining uncertainty must be due to  $(\text{id}_{\mathbb{U}} - \mathcal{P}_{\hat{\mathbb{U}}})[\mathbf{u}] \mid \{\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}[\mathbf{u}] = 0\}_{i=1}^n$ . Note the striking similarity of this property to the notion of *Galerkin orthogonality* (Logg et al., 2012, Equation 2.63).

A canonical choice for  $\mathcal{P}_{\hat{\mathbb{U}}}$  would arguably be an orthogonal projection w.r.t. the RKHS inner product of the sample space of  $\mathbf{u}$  (see e.g. Kanagawa et al. 2018). However, this inner product is generally difficult to compute. Fortunately, we can use the  $L_2$  inner products or Sobolev inner products on the samples to induce a (usually non-orthogonal) projection  $\mathcal{P}_{\hat{\mathbb{U}}}$ .

**Example 3.5.** *If the elements of  $\mathbb{U}$  are square-integrable, then the linear operator*

$$\mathcal{P}_{\mathbb{R}^m}[\mathbf{u}] := \mathbf{P}^{-1} \left( \int_{\mathbb{D}} \langle \phi^{(i)}(\mathbf{x}), \mathbf{u}(\mathbf{x}) \rangle_{\mathbb{R}^d} d\mathbf{x} \right)_{i=1}^m,$$

where

$$P_{ij} := \int_D \langle \phi^{(i)}(\mathbf{x}), \phi^{(j)}(\mathbf{x}) \rangle_{\mathbb{R}^d} d\mathbf{x},$$

induces a projection  $\mathcal{P}_{\hat{\mathbb{U}}} = \mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbb{U}}} \mathcal{P}_{\mathbb{R}^m}$  onto  $\hat{\mathbb{U}} \subset \mathbb{U}$ , even if  $\mathbb{U}$  is not a Hilbert space with inner product  $\langle \cdot, \cdot \rangle_{L_2(\mathbb{D})}$ .

At first glance, information operators restricting  $\hat{\mathbb{U}}$  to be finite-dimensional might seem fundamentally inferior to the information operators from Section 3.3.1. However, the conditional mean of a Gaussian process prior conditioned on  $\{\mathcal{I}_{l^{(i)}}[\mathbf{u}] = 0\}_{i=1}^n$  is updated by a linear combination of  $n$  functions, while the covariance function receives an at most rank  $n$  downgrade. This means that, implicitly, conditioning a Gaussian process on an information operator with  $\mathcal{P}_{\hat{\mathbb{U}}} = \text{id}_{\mathbb{U}}$  also constructs a finite-dimensional trial function space, which depends on the test function basis, the bilinear form  $B$  and the prior covariance function  $\mathbf{k}$ .

MWR information operators with finite-dimensional trial function bases can be used to realize a GP-based analogue of the finite element method.

**Example 3.6** (A 1D Finite Element Method). *Finite element methods are (generalized) Galerkin methods, where the functions in the test and trial bases have compact support, i.e. they are nonzero only in a highly localized region of the domain. The archetype of a*

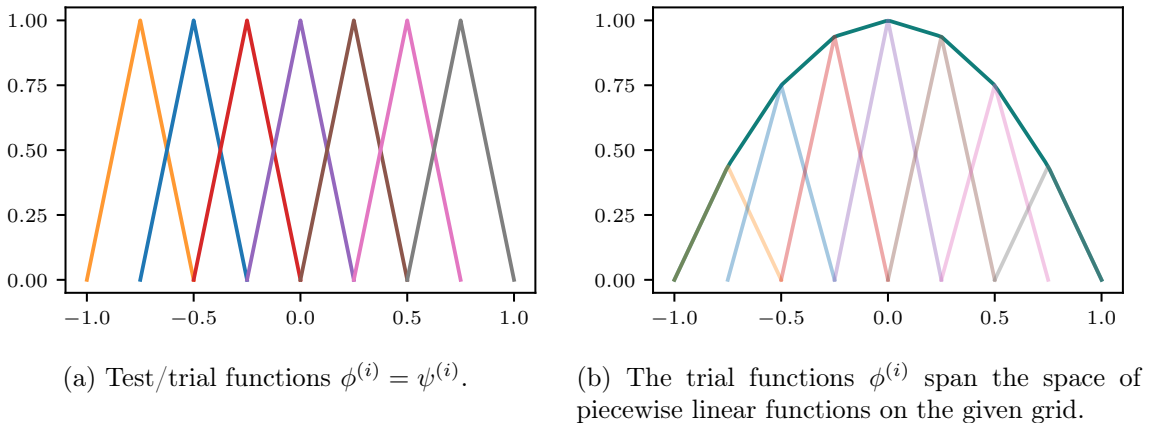


Figure 8: Linear Lagrange test and trial functions as used by the finite element method.

finite element method for the weak formulation from Section 2.1.1 uses linear Lagrange elements (Logg et al., 2012, Section 3.3.1) as test and trial functions, i.e.  $\phi^{(i)}(x) = \psi^{(i)}(x)$  and  $m = n$ . Linear Lagrange elements are piecewise linear on a triangulation of the domain. For instance, on a one-dimensional domain  $\mathbb{D} = ]-1, 1[$ , the linear Lagrange elements are given by Equation (3.9) from Example 3.4. Multiplying a coordinate vector  $\mathbf{c} \in \mathbb{R}^m$  with these basis functions leads to a piecewise linear interpolation between the points

$$(x_0, 0), (x_1, c_1), \dots, (x_n, c_n), (x_{n+1}, 0),$$

since, for  $x \in [x_i, x_{i+1}]$ ,


$$\sum_{i=1}^m c_i \phi^{(i)}(x) = c_i \frac{x_{i+1} - x}{x_{i+1} - x_i} + c_{i+1} \frac{x - x_i}{x_{i+1} - x_i} = \left(1 - \frac{x - x_i}{x_{i+1} - x_i}\right) c_i + \left(\frac{x - x_i}{x_{i+1} - x_i}\right) c_{i+1}.$$

The basis functions and an element in their span are visualized in Figure 8. The Lagrange elements at the boundary of the domain can also be easily modified such that arbitrary piecewise linear boundary conditions are fulfilled by construction. The effect of MWR information operators based on this set of test and trial functions is visualized in Figure 9(a).

### 3.3.3 MWR INFORMATION OPERATORS

Even though the class of information operators introduced above is constructed for linear PDEs, it can naturally be applied to the weak form of an arbitrary operator equation. In particular, we can use MWR information operators for the boundary conditions in an (I)BVP. Moreover, it is straightforward to extend  $\mathcal{I}_{l, \mathcal{P}_{\hat{\mathfrak{c}}}}$  to a joint GP prior over  $(\mathbf{u}, \mathbf{f}^{(w)})$  if the right-hand side  $\mathbf{f}^{(w)}$  of the operator equation is unknown as in Section 2.1. In this case,  $\mathcal{I}_{l, \mathcal{P}_{\hat{\mathfrak{c}}}}$  is jointly linear in  $(\mathbf{u}, \mathbf{f}^{(w)})$ . Summarizing Sections 3.3.1 and 3.3.2 and incorporating the extensions discussed here, we define an MWR information operator as follows:

**Definition 3.1** (MWR Information Operator). *Let  $\mathcal{D}^{(w)}[\mathbf{u}] = \mathbf{f}^{(w)}$  be an operator equation in strong or weak formulation. An MWR information operator for said operator equation is*



(a) Posterior process corresponding to a Matérn- $3/2$  prior. The sample paths of the process embed continuously into the Sobolev space  $H^1(\mathbb{D})$  (see Appendix B.4).

(b) Posterior process corresponding to an MWR Recovery Prior constructed from a Matérn- $3/2$  prior via Proposition 3.4. The posterior mean corresponds to the point estimate produced by the classical MWR.

Figure 9: Conditioning two different Gaussian process priors on the MWR information operators  $\{\mathcal{I}_{\psi^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}\}_{i=1}^n$  corresponding to the weak formulation of the Poisson equation, i.e. Equation (2.3), and  $m = 3$  linear Lagrange elements as test functions  $\psi^{(i)}$  and trial functions  $\phi^{(i)}$  (see Example 3.6). The trial functions  $\phi^{(1)}$  and  $\phi^{(m)}$  were modified to fulfill the non-zero boundary conditions exactly.

a continuous affine functional

$$\mathcal{I}_{l, \mathcal{P}_{\hat{\mathbb{U}}}} := (l \circ \mathcal{D}^{(w)} \circ \mathcal{P}_{\hat{\mathbb{U}}})[\cdot] - l[f^{(w)}]$$

parameterized by a test functional  $l \in \mathbb{L}^{(w)}$  and a bounded (potentially oblique) projection  $\mathcal{P}_{\hat{\mathbb{U}}}$  onto a subspace  $\hat{\mathbb{U}} \subset \mathbb{U}$ . We also write  $\mathcal{I}_l := \mathcal{I}_{l, \text{id}_{\mathbb{U}}}$ . The input of  $\mathcal{I}_{l, \mathcal{P}_{\hat{\mathbb{U}}}}$  can be extended to the right-hand side  $f^{(w)}$  of the operator equation, i.e.

$$\mathcal{I}_{l, \mathcal{P}_{\hat{\mathbb{U}}}}[(\mathbf{u}, f^{(w)})] := (l \circ \mathcal{D}^{(w)} \circ \mathcal{P}_{\hat{\mathbb{U}}})[\mathbf{u}] - l[f^{(w)}],$$

which is jointly linear in  $(\mathbf{u}, f^{(w)})$ .

### 3.3.4 RECOVERY OF CLASSICAL METHODS

In this section we will show that, under certain assumptions, the posterior mean of a GP prior conditioned on a set of MWR information operators is identical to the approximation generated by the corresponding traditional method of weighted residuals, examples of which are given in Table 1. More precisely, we will show that there is a flexible family of GP priors  $\mathbf{u} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$  whose posterior means after conditioning on  $\{\mathcal{I}_{l^{(i)}, \mathcal{P}_{\hat{\mathbb{U}}}}\}_{i=1}^m$  are identical to the corresponding classical MWR approximation  $\mathbf{e}^{\text{MWR}}$  to the solution of the same weak form linear PDE, where we use the same trial functions  $\phi^{(1)}, \dots, \phi^{(m)}$  and test functionals  $l^{(1)}, \dots, l^{(m)}$  in both cases, i.e.  $\hat{\mathbb{U}} = \text{span}(\phi^{(1)}, \dots, \phi^{(m)})$ . As in Section 2.1.2, we assume that the trial functions are already constructed in such a way that the boundary conditions

	Method	Trial Functions $\phi^{(i)}$	Test Functionals $l^{(i)}$ / Functions $\psi^{(i)}$
Strong Solutions	Collocation	arbitrary	$l^{(i)} = \delta_{\mathbf{x}_i}$ for $\mathbf{x}_i \in \mathbb{D}$ $\Rightarrow (l^{(i)} \circ \mathcal{D})[\mathbf{u}] = \mathcal{D}[\mathbf{u}](\mathbf{x}_i)$
	Subdomain (Finite Volume)	arbitrary	$\psi^{(i)} = \chi_{\mathbb{D}_i}$ for $\mathbb{D}_i \subset \mathbb{D}$ $\Rightarrow (l^{(i)} \circ \mathcal{D})[\mathbf{u}] = \int_{\mathbb{D}_i} \mathcal{D}[\mathbf{u}](\mathbf{x}) \, d\mathbf{x}$
	Pseudospectral	orthogonal and globally supported (e.g. Fourier basis or Chebychev polynomials)	$l^{(i)} = \delta_{\mathbf{x}_i}$ for $\mathbf{x}_i \in \mathbb{D}$ $\Rightarrow (l^{(i)} \circ \mathcal{D})[\mathbf{u}] = \mathcal{D}[\mathbf{u}](\mathbf{x}_i)$
Weak & Strong Solutions	Generalized Galerkin	arbitrary	arbitrary, but in general $\psi^{(i)} \neq \phi^{(i)}$
	Finite Element	locally supported (e.g. piecewise polynomial)	same class as trial functions, but in general $\psi^{(i)} \neq \phi^{(i)}$
	Spectral (Galerkin)	orthogonal and globally supported (e.g. Fourier basis or Chebychev polynomials)	same class as trial functions, but in general $\psi^{(i)} \neq \phi^{(i)}$
	(Ritz-)Galerkin	arbitrary	$\psi^{(i)} = \phi^{(i)}$

Table 1: Trial and test function(al)s defining commonly used methods of weighted residuals. If used as part of an MWR information operator, the GP posterior mean recovers the corresponding classic method (see Corollary 3.3).

are fulfilled. However, it is possible to extend the results below to the general case by adding MWR information operators corresponding to the boundary conditions and using

$$\mathbf{c}^{\text{MWR}} = \begin{pmatrix} \hat{\mathbf{D}}_{\text{PDE}} \\ \hat{\mathbf{D}}_{\text{BC}} \end{pmatrix}^{-1} \begin{pmatrix} \hat{\mathbf{f}}_{\text{PDE}} \\ \hat{\mathbf{f}}_{\text{BC}} \end{pmatrix}$$

as coordinates for the reference solution generated by the traditional MWR.

**Proposition 3.2.** *If  $\hat{\mathbf{D}} \in \mathbb{R}^{n \times m}$  and  $\Sigma_{\mathbf{c}} := \mathcal{P}_{\mathbb{R}^m} \mathbf{k} \mathcal{P}'_{\mathbb{R}^m} \in \mathbb{R}^{m \times m}$  are invertible, then*

$$\mathbf{c} \mid \hat{\mathbf{D}}\mathbf{c} - \hat{\mathbf{f}} = \mathbf{0} \sim \delta_{\mathbf{c}^{\text{MWR}}}$$

and the conditional mean  $\mathbf{m}^{\mathbf{u} \mid \hat{\mathbf{D}}, \hat{\mathbf{f}}}$  of  $\mathbf{u} \mid \hat{\mathbf{D}}\mathcal{P}_{\mathbb{R}^m}[\mathbf{u}] - \hat{\mathbf{f}} = \mathbf{0}$  admits a unique additive decomposition

$$\mathbf{m}^{\mathbf{u} \mid \hat{\mathbf{D}}, \hat{\mathbf{f}}} = \mathbf{u}^{\text{MWR}} + \mathbf{u}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})} \quad (3.11)$$

with  $\mathbf{u}^{\text{MWR}} \in \hat{\mathbf{U}}$  and  $\mathbf{u}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})} \in \ker(\mathcal{P}_{\hat{\mathbf{U}}})$ .

**Corollary 3.3** (MWR Generalization). *If, additionally,  $\mathbf{m} \in \hat{\mathbf{U}}$  and  $\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})} \mathbf{k} \mathcal{P}'_{\mathbb{R}^m} = \mathbf{0}$ , then the conditional mean function  $\mathbf{m}^{\mathbf{u} \mid \hat{\mathbf{D}}, \hat{\mathbf{f}}}$  is equal to the MWR approximation, i.e.*

$$\mathbf{m}^{\mathbf{u} \mid \hat{\mathbf{D}}, \hat{\mathbf{f}}} = \mathbf{u}^{\text{MWR}}.$$



It turns out that it is possible to transform any admissible GP prior over the (weak) solution of the PDE into a prior that fulfills the assumptions of Corollary 3.3.

**Proposition 3.4** (MWR Recovery Prior). *Let  $\tilde{\mathbf{u}} \sim \mathcal{GP}(\tilde{\mathbf{m}}, \tilde{\mathbf{k}})$  with mean and sample paths in  $\mathbb{U}$ . Then  $\mathbf{u} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$  with  $\mathbf{m} := \mathcal{P}_{\hat{\mathbb{U}}}[\tilde{\mathbf{m}}]$  and*

$$\begin{aligned} \mathbf{k} &:= \mathcal{P}_{\hat{\mathbb{U}}}\tilde{\mathbf{k}}\mathcal{P}'_{\hat{\mathbb{U}}} + \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}\tilde{\mathbf{k}}\mathcal{P}'_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})} \\ &= \mathcal{P}_{\hat{\mathbb{U}}}\tilde{\mathbf{k}}\mathcal{P}'_{\hat{\mathbb{U}}} + (\text{id}_{\mathbb{U}} - \mathcal{P}_{\hat{\mathbb{U}}})\tilde{\mathbf{k}}(\text{id}_{\mathbb{U}} - \mathcal{P}_{\hat{\mathbb{U}}})' \\ &= \tilde{\mathbf{k}} - \mathcal{P}_{\hat{\mathbb{U}}}\tilde{\mathbf{k}} - \tilde{\mathbf{k}}\mathcal{P}'_{\hat{\mathbb{U}}} + 2\mathcal{P}_{\hat{\mathbb{U}}}\tilde{\mathbf{k}}\mathcal{P}'_{\hat{\mathbb{U}}} \end{aligned}$$

has sample paths in  $\mathbb{U}$ ,  $\mathbf{m} \in \hat{\mathbb{U}}$ , and  $\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}\mathbf{k}\mathcal{P}'_{\mathbb{R}^m} = \mathbf{0}$ .

Figure 9(b) visualizes how a prior of this form reproduces a 1D finite element method in the posterior mean and Figure 9 as a whole contrasts the difference between  $\tilde{\mathbf{u}}$  and  $\mathbf{u}$ . Intuitively speaking, the construction for the covariance from Proposition 3.4 enforces statistical independence between the subspaces  $\hat{\mathbb{U}}$  and  $\ker(\mathcal{P}_{\hat{\mathbb{U}}})$  of the GP's path space. This way, an observation of the GP prior in the subspace  $\hat{\mathbb{U}}$  gains no information about  $\ker(\mathcal{P}_{\hat{\mathbb{U}}})$ , which means that the posterior process will not be updated along  $\ker(\mathcal{P}_{\hat{\mathbb{U}}})$ . Since  $\mathbf{m} \in \hat{\mathbb{U}}$ , i.e.  $\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}[\mathbf{m}]$ , it follows that the posterior mean will also lie in  $\hat{\mathbb{U}}$ . Even though this choice of prior is somewhat restrictive, there are good reasons to use it in practice, arguably the most important of which is that the uncertainty quantification provided by the GP can be added on top of traditional MWR solvers in existing pipelines in a plug-and-play fashion. This is because given the MWR recovery prior, the mean estimate of the probabilistic numerical method agrees with the point estimate produced by the classical solver.

### 3.4 Algorithm

Algorithm 1 summarizes our framework from an algorithmic standpoint. It outlines how a GP prior can be conditioned on heterogeneous sources of information such as mechanistic knowledge given in the form of a linear boundary value problem, and noisy measurement data by leveraging the notion of a linear information operator. All GP posteriors in this article were computed by this algorithm with different choices of prior, PDE, boundary conditions and policy. Modeling uncertainty over the right-hand side  $f^{(w)}$  of the PDE and the boundary function(al)  $g^{(w)}$  is achieved by specifying a joint prior  $(\mathbf{u}, f^{(w)}, g^{(w)}, \epsilon)$ . Therefore, Algorithm 1 also returns a multi-output Gaussian process posterior over the same objects. This means that our method can be used to solve PDE-constrained Bayesian inverse problems for the right-hand side  $f^{(w)}$  and the boundary function  $g^{(w)}$ , while computing a consistent distributional estimate for the corresponding solution  $\mathbf{u}$  of the forward problem. This is a generalization of a linear *latent force model* (Alvarez et al., 2009). If  $f^{(w)}$  and  $g^{(w)}$  are not uncertain, the corresponding covariance functions in the joint prior can simply be set to 0, which (in the absence of measurements) reduces the joint prior to a simple prior over the solution  $\mathbf{u}$ . To condition the GP on the PDE and the boundary conditions, we make use of MWR information operators (see Definition 3.1), where the test functions and projections are chosen by an arbitrary policy in each iteration of the method. An example of such a policy which reproduces Figure 1(c) chooses  $\mathcal{P}_{\hat{\mathbb{U}}}$  as the  $L_2$  projection onto the basis

---

 Algorithm 1: Solving PDEs via Gaussian Process Inference
 

---

**Input:** Joint GP prior  $(\mathbf{u}, f^{(w)}, g^{(w)}, \epsilon) \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$ , linear PDE  $(\mathcal{D}^{(w)}, f^{(w)})$ , boundary conditions  $(\mathcal{B}^{(w)}, g^{(w)})$ , (noisy) measurements  $(\mathbf{X}_{\text{MEAS}}, \mathbf{y}_{\text{MEAS}})$ , ...

**Output:** GP posterior  $\mathcal{GP}(\mathbf{m}^{(i)}, \mathbf{k}^{(i)})$

```

1  procedure LINPDE-GP( $\mathbf{m}, \mathbf{k}, \mathcal{I}_{\cdot}^{\text{PDE}}, \mathcal{I}_{\cdot}^{\text{BC}}, \mathbf{X}_{\text{MEAS}}, \mathbf{y}_{\text{MEAS}}$ )
2       $i \leftarrow 0$ 
3       $(\mathbf{m}^{(0)}, \mathbf{k}^{(0)}) \leftarrow (\mathbf{m}, \mathbf{k})$ 
4       $\mathbf{w}^{(0)} \leftarrow ()$ 
5       $\mathbf{G}^{(0)} \leftarrow ()$ 
6      while not STOPPINGCRITERION() do
7           $i \leftarrow i + 1$ 
8           $(l_{\text{PDE}}^{(i)}, l_{\text{BC}}^{(i)}, \mathcal{P}_{\hat{\mathbf{U}}}^{(i)}, \dots, \mathbf{v}_{\text{MEAS}}^{(i)}) \leftarrow \text{POLICY}(\mathbf{m}^{(i)}, \mathbf{k}^{(i)})$  ▷ Action
9           $\mathcal{I}^{(i)} \leftarrow (\mathbf{u}, f^{(w)}, g^{(w)}, \epsilon) \mapsto \begin{pmatrix} \mathcal{I}_{l_{\text{PDE}}^{(i)}, \mathcal{P}_{\hat{\mathbf{U}}}^{(i)}}^{\text{PDE}}[(\mathbf{u}, f^{(w)})] \\ \mathcal{I}_{l_{\text{BC}}^{(i)}, \mathcal{P}_{\hat{\mathbf{U}}}^{(i)}}^{\text{BC}}[(\mathbf{u}, g^{(w)})] \\ \vdots \\ \langle \mathbf{v}_{\text{MEAS}}^{(i)}, \mathbf{u}(\mathbf{X}_{\text{MEAS}}) + \epsilon \rangle \end{pmatrix}$  ▷ Information operator
10          $\mathbf{y}^{(i)} \leftarrow \begin{pmatrix} 0 & 0 & \dots & \langle \mathbf{v}_{\text{MEAS}}^{(i)}, \mathbf{y}_{\text{MEAS}} \rangle \end{pmatrix}^{\top}$  ▷ Observations
11          $\mathbf{G}^{(i)} \leftarrow \begin{pmatrix} \mathbf{G}^{(i-1)} & \mathcal{I}^{(1:i-1)} \mathbf{k}(\mathcal{I}^{(i)})' \\ \mathcal{I}^{(i)} \mathbf{k}(\mathcal{I}^{(1:i-1)})' & \mathcal{I}^{(i)} \mathbf{k}(\mathcal{I}^{(i)})' \end{pmatrix}$  ▷ Update Gram matrix
12          $\mathbf{w}^{(i)} \leftarrow (\mathbf{G}^{(i)})^{\dagger} (\mathbf{y}^{(1:i)} - \mathcal{I}^{(1:i)}[\mathbf{m}])$  ▷ Update representer weights
13          $m_j^{(i)} \leftarrow \mathbf{x} \mapsto m_j(\mathbf{x}) + \mathcal{I}^{(1:i)}[\mathbf{k}_{:,j}(\cdot, \mathbf{x})]^{\top} \mathbf{w}^{(i)}$  ▷ Belief Update
14          $k_{j_1, j_2}^{(i)} \leftarrow (\mathbf{x}_1, \mathbf{x}_2) \mapsto k_{j_1, j_2}(\mathbf{x}_1, \mathbf{x}_2) - \mathcal{I}^{(1:i)}[\mathbf{k}_{:,j_1}(\cdot, \mathbf{x}_1)]^{\top} (\mathbf{G}^{(i)})^{\dagger} \mathcal{I}^{(1:i)}[\mathbf{k}_{:,j_2}(\cdot, \mathbf{x}_2)]$ 
15     return  $\mathcal{GP}(\mathbf{m}^{(i)}, \mathbf{k}^{(i)})$ 

```

---

from Example 3.6 in every iteration, the test functions  $l_{\text{BC}} \in \{\delta_{-1}, \delta_1\}$ , and  $l_{\text{PDE}} = 0$  in the first two iterations; and  $l_{\text{PDE}}$  is induced by  $\psi^{(i-2)} = \phi^{(i-2)}$  (and  $l_{\text{BC}} = 0$ ) from iteration 3 onward. The ellipses in the information operator  $\mathcal{I}^{(i)}$  and the observations  $\mathbf{y}^{(i)}$  indicate that adding additional information operators is possible in the same fashion. For instance, adding additional PDE information operators enables the solution of systems of linear PDEs.

**Performance Considerations** Instead of naively conditioning the previous conditional process on the new observation in each iteration, Algorithm 1 always conditions the prior on the accumulated observations. This is because the naive expressions for the conditional moments become more and more complex over time. While, in principle, it is possible to use automatic differentiation (AD) to compute  $\mathcal{I}^{(i)}[\mathbf{m}^{(i)}]$ ,  $\mathcal{I}^{(i)}[\mathbf{k}_{:,j}^{(i-1)}(\cdot, \mathbf{x})]$ , and  $\mathcal{I}^{(i)} \mathbf{k}^{(i-1)}(\mathcal{I}^{(i)})'$  in each iteration and then evaluate Equations (4.10) and (4.11) naively, we found that this is detrimental to the performance of the algorithm. In Algorithm 1, we only need to compute  $\mathcal{I}^{(i)}[\mathbf{m}]$ ,  $\mathcal{I}^{(i)}[\mathbf{k}_{:,j}(\cdot, \mathbf{x})]$ , and  $\mathcal{I}^{(i)} \mathbf{k}(\mathcal{I}^{(i)})'$  on the prior moments, which are much less complex and cheaper to evaluate. For maximum efficiency, for many information operator

/ kernel combinations one can compute optimized closed-form expressions for these terms, alleviating the need for automatic differentiation or quadrature. We can avoid unnecessary recomputation of the representer weights at every iteration of the method by means of block-matrix inversion. For instance, if a Cholesky decomposition is used to invert the Gramian  $\mathbf{G}^{(i)}$ , we can use a variant of the block Cholesky decomposition (Golub and Van Loan, 2013) to update the Cholesky factor of  $\mathbf{G}^{(i-1)}$ .

**Code** A Python implementation of Algorithm 1 based on ProbNum (Wenger et al., 2021) and JAX (Bradbury et al., 2018) is available at:

<https://github.com/marvinpfoertner/linpde-gp>

### 3.5 Related Work

The area of physics-informed machine learning (Karniadakis et al., 2021) aims at augmenting machine learning models with mechanistic knowledge about physical phenomena, mostly in the form of ordinary and partial differential equations. Recently, there has been growing interest in deep learning-based approaches (Raissi et al., 2019; Li et al., 2020, 2021). However, this model choice makes it inherently difficult to quantify the uncertainty about the solution induced by noise-corrupted input data and inevitable approximation error. Instead, we approach the problem through the lens of *probabilistic numerics* (Hennig et al., 2015; Cockayne et al., 2019b; Oates and Sullivan, 2019; Owhadi et al., 2019; Hennig et al., 2022), which frames numerical problems as statistical estimation tasks. Probabilistic numerical methods for the solution of PDEs are predominantly based on Gaussian process priors. Our work builds upon and extends these works. Many existing methods aim to find a strong solution to a linear PDE using a collocation scheme (e.g. Graepel 2003; Cockayne et al. 2017; Raissi et al. 2017). Unfortunately, many practically relevant (linear) PDEs only admit weak solutions. Our framework extends existing collocation approaches to weak formulations. Probabilistic numerical methods approximating weak formulations are primarily based on discretization. For example, Cockayne et al. (2019a); Wenger and Hennig (2020) apply a probabilistic linear solver to the linear system arising from discretization. Girolami et al. (2021) propose a statistical version of the finite element method (statFEM), which uses a specific parametric GP prior. However, these approaches do not quantify the inherent discretization error – often the largest source of uncertainty about the solution. In contrast, our framework models this error and additionally admits a broader class of discretizations. Wang et al. (2021); Krämer et al. (2022) propose GP-based solvers for strong formulations of time-dependent nonlinear PDEs by leveraging finite-difference approximations to the differential operator and linearization-based approximate inference. While it is possible to apply such methods to linear PDEs, the finite difference approximation of the differential operator introduces additional estimation error. In contrast, the evaluation of the differential operator in our method is exact. Cockayne et al. (2017); Raissi et al. (2017); Girolami et al. (2021) also apply their methods to solve PDE-constrained (Bayesian) inverse problems. Särkkä (2011) directly infers the right-hand side of a linear PDE in strong formulation by observing measurements of the solution through the associated Green’s function. Our approach also builds a belief over an unknown right-hand side without requiring access to a Green’s function. The aforementioned methods use the closure of Gaussian processes under conditioning on observations of the sample paths through a linear operator without proof. Owhadi and

Scovel (2018) show how to condition Gaussian *measures* on an orthogonal direct sum of separable Hilbert spaces on observations of one of the summands. However, this result does not apply to separable *Banach* spaces such as Hölder spaces, which are ubiquitous in the study of strong solutions of linear PDEs. Furthermore, when it can be applied, it does not translate to Gaussian *processes* without significant effort.<sup>8</sup> Our work therefore provides the theoretical basis for conditioning Gaussian *processes* on observations of their sample paths made through an arbitrary bounded linear operator with values in  $\mathbb{R}^n$ . Recent results about the sample spaces of GPs (Steinwart, 2019; Kanagawa et al., 2018) ensure the applicability of our work to practical GP regression problems. From a practitioner’s perspective this allows the modeling flexibility of Gaussian processes via the kernel, while ensuring that conditioning on observations of the sample paths through a linear operator is possible. To our knowledge this is the first complete proof of this widely used property of GPs. Thus, Theorem 1 provides the theoretical basis for physics-informed GP regression, including the aforementioned methods for the solution of PDEs. In our work, it enables conditioning on information operators constructed from e.g. PDEs, boundary conditions and general integral equations.

#### 4. Gaussian Process Inference with Linear Operator Observations

Our framework fundamentally relies on the fact that when a Gaussian process prior is conditioned on linear observations of its paths, one obtains a closed-form posterior. This section provides the theoretical foundation for this result. While this property is used widely in the literature (see e.g. Graepel (2003); Rasmussen and Williams (2006); Särkkä (2011); Särkkä et al. (2013); Cockayne et al. (2017); Raissi et al. (2017); Agrell (2019); Albert (2019); Krämer et al. (2022)), no proof of its general form where observations are made via *bounded linear operators* mapping a separable Banach function spaces into  $\mathbb{R}^n$ , instead of *finite-dimensional linear maps* on a finite number of point evaluations exists, to the best of our knowledge. Owhadi and Scovel (2018) give a proof of a related property for Gaussian *measures* on separable Hilbert spaces. Here, we extend their results to the case of Gaussian *processes*. While these perspectives are closely related, significant technical attention needs to be paid for this result to transfer to the GP case. For our framework this is essential such that we can leverage the modeling capabilities provided by specifying a kernel as described in Section 3.1.1.

To state the result, let  $\mathbf{f} \sim \mathcal{GP}(\mathbf{m}, \mathbf{k})$  be a (multi-output) GP prior with index set  $\mathbb{X}$ ,  $\mathcal{L}: \text{paths}(\mathbf{f}) \rightarrow \mathbb{R}^n$  a linear operator acting on the paths of  $\mathbf{f}$ , and  $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  a Gaussian random vector in  $\mathbb{R}^n$  with  $\boldsymbol{\epsilon} \perp \mathbf{f}$ . We need to compute the conditional random process

$$\mathbf{f} \mid \mathcal{L}[\mathbf{f}] + \boldsymbol{\epsilon} = \mathbf{y}$$

for some  $\mathbf{y} \in \mathbb{R}^n$ . This object is defined as the family  $(\mathbf{f} \mid \mathcal{L}[\mathbf{f}] + \boldsymbol{\epsilon} = \mathbf{y}) := \{\mathbf{f}(x, \cdot) \mid E\}_{x \in \mathbb{X}}$ , of conditional random variables<sup>9</sup>, where  $(\Omega, \mathcal{B}(\Omega), \mathbb{P})$  is the probability space on which both  $\mathbf{f}$  and  $\boldsymbol{\epsilon}$  are defined,  $E$  is the event  $E := \mathbf{h}^{-1}(\{\mathbf{y}\}) \in \mathcal{B}(\Omega)$ , and  $\mathbf{h}$  is the random variable

$$\mathbf{h}: \Omega \rightarrow \mathbb{R}^n, \omega \mapsto \mathcal{L}[\mathbf{f}(\cdot, \omega)] + \boldsymbol{\epsilon}(\omega).$$

---

8. The theoretical results of an earlier version of this work were based on the result by Owhadi and Scovel (2018). In order to generalize our framework to Banach spaces, we’ve adopted a different proof strategy.

We refer to Appendix B for definitions of the objects mentioned above. For instance, in Section 3, we use  $\mathcal{L} := (\mathcal{D}[\cdot](\mathbf{x}_i))_{i=1}^n$ , where  $\mathcal{D}$  is a linear differential operator, as well as  $\mathcal{L}[\mathbf{f}] := (\mathbf{f}(\mathbf{x}_i))_{i=1}^n$ , and, in Section 3.2, we additionally use

$$\mathcal{L}[\mathbf{f}] = \int_{\mathbb{X}} \mathbf{f}(\mathbf{x}) \, d\mathbf{x}.$$

It is well-known that  $\mathbf{h}$  is a Gaussian random vector  $\mathbf{h} \sim \mathcal{GP}(\mathcal{L}[\mathbf{m}] + \boldsymbol{\mu}, \mathcal{L}\mathbf{k}\mathcal{L}' + \boldsymbol{\Sigma})$ , where  $\mathcal{L}\mathbf{k}\mathcal{L}' \in \mathbb{R}^{n \times n}$  with  $(\mathcal{L}\mathbf{k}\mathcal{L}')_{i_1, i_2} = \mathcal{L}_{i_1}[\mathbf{x} \mapsto (\mathcal{L}_{i_2}[\mathbf{k}_{j_1, \cdot}(\mathbf{x}, \cdot)])_{j_1=1}^n]$ , and that the conditional random process is a (multi-output) Gaussian process

$$\mathbf{f} \mid \mathcal{L}[\mathbf{f}] + \boldsymbol{\epsilon} = \mathbf{y} \sim \mathcal{GP}(m^{\mathbf{f}|\mathbf{y}}, k^{\mathbf{f}|\mathbf{y}})$$

with conditional moments given by

$$\begin{aligned} m_i^{\mathbf{f}|\mathbf{y}}(\mathbf{x}) &= m_i(\mathbf{x}) + \mathcal{L}[\mathbf{k}_{:,i}(\cdot, \mathbf{x})]^\top (\mathcal{L}\mathbf{k}\mathcal{L}' + \boldsymbol{\Sigma})^{-1} (\mathbf{y} - (\mathcal{L}[\mathbf{m}] + \boldsymbol{\mu})), \quad \text{and} \\ k_{i_1, i_2}^{\mathbf{f}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) &= k_{i_1, i_2}(\mathbf{x}_1, \mathbf{x}_2) + \mathcal{L}[\mathbf{k}_{:,i_1}(\cdot, \mathbf{x}_1)]^\top (\mathcal{L}\mathbf{k}\mathcal{L}' + \boldsymbol{\Sigma})^{-1} \mathcal{L}[\mathbf{k}_{:,i_2}(\cdot, \mathbf{x}_2)]. \end{aligned}$$

Since the above are nontrivial claims about potentially ill-behaved infinite-dimensional objects, a proof is important, be it just to identify a precise set of assumptions about the objects at play, ensuring the result holds. For instance, the statement that  $\mathbf{h}$  is a random vector, i.e. a measurable function, is highly nontrivial. To remedy this situation, a major contribution of this work are Theorem 1 and Corollary 2 and their proof in Appendix B, which prove the claims above under realistic assumptions. Hence, besides being the theoretical basis for this work, Theorem 1 and Corollary 2 also provide theoretical backing for many of the publications cited above. Our results identify a set of mild assumptions, which are easy to verify and widely-applicable in practical applications. Assumption 1 constitutes the common set of assumptions shared by Theorem 1 and Corollary 2.

**Assumption 1.** *Let  $f \sim \mathcal{GP}(m, k)$  be a Gaussian process prior with index set  $\mathbb{X}$  on the probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , whose paths lie in a real separable reproducing kernel Banach space (RKBS)  $\mathbb{B} \subset \mathbb{R}^{\mathbb{X}}$  such that  $\omega \mapsto f(\cdot, \omega)$  is a  $\mathbb{B}$ -valued Gaussian random variable.*

For instance, for a 1D domain  $\mathbb{D} \subset \mathbb{R}$ , a GP prior with half-integer Matérn kernel with smoothness parameter  $\nu = p + \frac{1}{2}$  fulfills Assumption 1 with  $\mathbb{B} = C^p(\mathbb{D})$ , i.e. the space of  $p$ -times differentiable functions with bounded and uniformly continuous derivatives. Similar results hold in multiple dimensions and for other kernels. See Appendix B.4 for more information on prior selection.

Theorem 1 enables affine observations, in which the GP sample paths enter through one or multiple continuous linear functionals. For example, we used Theorem 1 in Section 3.2 to condition on observations of an integral of a GP's paths and in Section 3.3 to condition on projections of the paths. To state the result conveniently, we introduce some notation.

**Notation 1.** *Let Assumption 1 hold and let  $\mathcal{L}: \mathbb{B} \rightarrow \mathbb{R}^n$  and  $\tilde{\mathcal{L}}: \mathbb{B} \rightarrow \mathbb{R}^{\tilde{n}}$  be bounded linear operators. By  $\mathcal{L}\mathbf{k}\tilde{\mathcal{L}}' \in \mathbb{R}^{n_1 \times n_2}$  we denote the matrix with entries*

$$(\mathcal{L}\mathbf{k}\tilde{\mathcal{L}}')_{ij} := \mathcal{L}[\mathbf{x} \mapsto \tilde{\mathcal{L}}[k(\mathbf{x}, \cdot)]]_j.$$

---

9. Here, we need to work with *regular conditional probability measures* (Klenke, 2014), since the event  $E$  typically has probability 0.

Table 2: Theorem 1 provides the theoretical basis to condition on (affine) observations of a Gaussian process. While results like conditioning on derivative evaluations are used ubiquitously (e.g. for monotonic GPs, Bayesian optimization, probabilistic numerical PDE solvers, ...) a complete proof does not exist in the literature, to the best of our knowledge.

Observation	Information operator	Proof known?	Reference
Point evaluation	$\mathbf{f}(\mathbf{x})$	✓	Bishop (2006)
Finite-dim. affine map	$\mathbf{A}\mathbf{f}(\mathbf{X}) + \mathbf{b}$	✓	Bishop (2006)
Point evaluation of derivative	$\left. \frac{\partial^{ \alpha }}{\partial \mathbf{x}^\alpha} f_i \right _{\mathbf{x}=\mathbf{x}'}$	✗	Corollary 2
Integral	$\int_{\mathbb{X}} \langle \psi(\mathbf{x}), \mathbf{f}(\mathbf{x}) \rangle d\mu(\mathbf{x})$	✗	Theorem 1
General affine functionals	$\mathcal{L}[\mathbf{f}] + \mathbf{b}$	✗	Theorem 1

The order in which the operators  $\mathcal{L}$ ,  $\tilde{\mathcal{L}}$  are applied to the arguments of  $k$  does not matter, i.e.  $(\mathcal{L}k\tilde{\mathcal{L}}')_{ij} = \mathcal{L}[\mathbf{x} \mapsto \tilde{\mathcal{L}}[k(\mathbf{x}, \cdot)]_j]_i = \tilde{\mathcal{L}}[\mathbf{x} \mapsto \mathcal{L}[k(\cdot, \mathbf{x})]_i]_j$ . This motivates the parenthesis-free notation  $\mathcal{L}k\tilde{\mathcal{L}}'$  introduced above.

**Theorem 1.** *Let Assumption 1 hold and let  $\mathcal{L}: \mathbb{B} \rightarrow \mathbb{R}^n$  be a bounded linear operator. Then*

$$\mathcal{L}[\mathbf{f}] \sim \mathcal{N}(\mathcal{L}[m], \mathcal{L}k\mathcal{L}'). \quad (4.1)$$

Let  $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  be an  $\mathbb{R}^n$ -valued Gaussian random vector with  $\boldsymbol{\epsilon} \perp \mathbf{f}$ . Then, for any  $\mathbf{y} \in \mathbb{R}^n$ ,

$$\mathbf{f} \mid \mathcal{L}[\mathbf{f}] + \boldsymbol{\epsilon} = \mathbf{y} \sim \mathcal{GP}(m^{\mathbf{f}|\mathbf{y}}, k^{\mathbf{f}|\mathbf{y}}), \quad (4.2)$$

with conditional mean and covariance function given by

$$m^{\mathbf{f}|\mathbf{y}}(\mathbf{x}) = m(\mathbf{x}) + \mathcal{L}[k(\mathbf{x}, \cdot)]^\top (\mathcal{L}k\mathcal{L}' + \boldsymbol{\Sigma})^\dagger (\mathbf{y} - (\mathcal{L}[m] + \boldsymbol{\mu})), \quad (4.3)$$

and

$$k^{\mathbf{f}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) = k(\mathbf{x}_1, \mathbf{x}_2) - \mathcal{L}[k(\mathbf{x}_1, \cdot)]^\top (\mathcal{L}k\mathcal{L}' + \boldsymbol{\Sigma})^\dagger \mathcal{L}[k(\cdot, \mathbf{x}_2)]. \quad (4.4)$$

Finally, we turn to Corollary 2, which is the result that is most widely-used throughout the literature (Graepel, 2003; Särkkä, 2011; Särkkä et al., 2013; Cockayne et al., 2017; Raissi et al., 2017; Agrell, 2019; Albert, 2019; Krämer et al., 2022). It shows how Gaussian processes can be conditioned on point evaluations of the image of their paths under a linear operator, provided that the linear operator is bounded and maps into a separable Banach function space, on which point evaluation is continuous. Moreover, it shows that, under these conditions, the image of the GP under the linear operator is itself a Gaussian process. Again, we introduce some notation to facilitate stating the result.

**Notation 2.** *Let Assumption 1 hold and let  $\mathcal{L}_i: \mathbb{B} \rightarrow \mathbb{B}_i$  for  $i = 1, 2$  be bounded linear operators mapping into real separable RKBSs  $\mathbb{B}_i \subset \mathbb{R}^{\mathbb{X}_i}$ , respectively. In analogy to Notation 1, we define the bivariate functions*

$$k\mathcal{L}'_2: \mathbb{X} \times \mathbb{X}_2 \rightarrow \mathbb{R}, (\mathbf{x}, \mathbf{x}_2) \mapsto \mathcal{L}_2[k(\mathbf{x}, \cdot)](\mathbf{x}_2), \quad (4.5)$$

$$\mathcal{L}_1 k: \mathbb{X}_1 \times \mathbb{X} \rightarrow \mathbb{R}, (\mathbf{x}_1, \mathbf{x}) \mapsto \mathcal{L}_1[k(\cdot, \mathbf{x})](\mathbf{x}_1), \quad \text{and} \quad (4.6)$$

$$\mathcal{L}_1 k \mathcal{L}'_2: \mathbb{X}_1 \times \mathbb{X}_2 \rightarrow \mathbb{R}, (\mathbf{x}_1, \mathbf{x}_2) \mapsto \mathcal{L}_1[(k \mathcal{L}'_2)(\cdot, \mathbf{x}_2)](\mathbf{x}_1). \quad (4.7)$$

**Corollary 2.** *Let Assumption 1 hold and let  $\mathcal{L}: \mathbb{B} \rightarrow \tilde{\mathbb{B}}$  be a bounded linear operator mapping into another real separable RKBS  $\tilde{\mathbb{B}} \subset \mathbb{R}^{\tilde{\mathbb{X}}}$ . Then*

$$\mathcal{L}[f] \sim \mathcal{GP}(\mathcal{L}[m], \mathcal{L}k \mathcal{L}'). \quad (4.8)$$

Let  $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  with values in  $\mathbb{R}^n$  and  $\boldsymbol{\epsilon} \perp f$ . Then, for  $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_i)_{i=1}^n \in \tilde{\mathbb{X}}^n$  and  $\mathbf{y} \in \mathbb{R}^n$ ,

$$f | \mathcal{L}[f](\tilde{\mathbf{X}}) + \boldsymbol{\epsilon} = \mathbf{y} \sim \mathcal{GP}(m^{\mathbf{f}|\mathbf{y}}, k^{\mathbf{f}|\mathbf{y}}) \quad (4.9)$$

with

$$m^{\mathbf{f}|\mathbf{y}}(\mathbf{x}) := m(\mathbf{x}) + (k \mathcal{L}')(\mathbf{x}, \tilde{\mathbf{X}})^\top \left( (\mathcal{L}k \mathcal{L}')(\tilde{\mathbf{X}}, \tilde{\mathbf{X}}) + \boldsymbol{\Sigma} \right)^\dagger (\mathbf{y} - (\mathcal{L}[m](\tilde{\mathbf{X}}) + \boldsymbol{\mu})) \quad (4.10)$$

and

$$k^{\mathbf{f}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) := k(\mathbf{x}_1, \mathbf{x}_2) - (k \mathcal{L}')(\mathbf{x}_1, \tilde{\mathbf{X}})^\top \left( (\mathcal{L}k \mathcal{L}')(\tilde{\mathbf{X}}, \tilde{\mathbf{X}}) + \boldsymbol{\Sigma} \right)^\dagger (\mathcal{L}k)(\tilde{\mathbf{X}}, \mathbf{x}_2) \quad (4.11)$$

If additionally  $\tilde{\mathbb{X}} = \mathbb{X}$ , then

$$\begin{pmatrix} f \\ \mathcal{L}[f] \end{pmatrix} \sim \mathcal{GP} \left( \begin{pmatrix} m \\ \mathcal{L}[m] \end{pmatrix}, \begin{pmatrix} k & k \mathcal{L}' \\ \mathcal{L}k & \mathcal{L}k \mathcal{L}' \end{pmatrix} \right). \quad (4.12)$$

This corollary is the theoretical basis for most of Section 3.2. For  $\mathcal{L} = \text{id}_{\mathbb{B}}$ , we recover standard GP regression as a special case in Corollary 2. Finally, both Theorem 1 and Corollary 2 apply also to vector-valued Gaussian processes.

**Remark 4.1** (Multi-Output Gaussian Processes). *Theorem 1 and Corollary 2 also apply to multi-output GPs  $\mathbf{f}$ . In this case, we interpret the sample paths  $\mathbf{f}(\cdot, \omega): \mathbb{X} \rightarrow \mathbb{R}^{d'}$  of the multi-output GP as sample paths  $\tilde{\mathbf{f}}(\cdot, \omega): I \times \mathbb{X} \rightarrow \mathbb{R}$ ,  $\tilde{\mathbf{f}}((i, \mathbf{x}), \omega) := f_i(\mathbf{x}, \omega)$  of a regular GP with index set  $I \times \mathbb{X} \rightarrow \mathbb{R}$ , where  $I = \{1, \dots, d'\}$  (see Section 2.2). We also generalize notation like  $\mathcal{L}k \mathcal{L}'$  accordingly.*

## 5. Conclusion

In this work, we developed a probabilistic framework for the solution of (systems of) linear partial differential equations, which can be interpreted as physics-informed Gaussian process regression. It enables the seamless fusion of (1) a-priori known, provable properties of the system of interest, (2) exact and partial mechanistic information, (3) subjective domain expertise, as well as, (4) noisy empirical measurements into a unified scientific model. This model outputs a consistent uncertainty estimate, which quantifies the inherent approximation error in addition to the uncertainty arising from partially-known physics, as well as limited-precision measurements. Our framework fundamentally relies on the closure of Gaussian processes under conditioning on observations of their sample paths through an

arbitrary bounded linear operator. While this result has been used ubiquitously in the literature, a rigorous proof for linear operator observations, as needed in the PDE setting, did not exist prior to this work to the best of our knowledge. Our work generalizes and unifies several related formulations of GP-PDE inference. Importantly, our formulation extends these ideas to virtually all popular methods for PDE simulation, revealing them to be a form of Gaussian process inference and in turn clarifying the underlying (probabilistic) assumptions. More specifically, by choosing a specific prior and information operator in our framework, it recovers methods of weighted residuals, a popular family of numerical methods for the solution of (linear) PDEs, which includes generalized Galerkin methods such as finite element and spectral methods. This demonstrates that classical linear PDE solvers can be generalized in their functionality to include approximate input data and equipped with a structured uncertainty estimate. Our work outlines a general framework for the integration of mechanistic building blocks in the form of information operators derived from e.g. linear PDEs into probabilistic models. Our case study shows that the language of information operators is a powerful toolkit for aggregating heterogeneous sources of partial information in a joint probabilistic model, especially in the context of physics-informed machine learning. This opens up several interesting lines of research. For example, the choice of prior and information operator are not fixed and can be specifically chosen for the problem at hand. The design of adaptive information operators, which actively collect information based on the current belief about the solution could prove to be a promising research direction. Further, the uncertainty estimate about the solution could be used to inform experimental design choices. For example, in the case study from Section 3.2, the posterior belief can be used to optimize the locations of the digital thermal sensors in future CPU designs. Finally, it remains an open question whether this framework can be adapted to nonlinear partial differential equations in a similar manner to how many classic methods solve a sequence of linearized problems to approximate the solution of a nonlinear PDE.

## Acknowledgments

MP, PH and JW gratefully acknowledge financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence “Machine Learning - New Perspectives for Science”, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting MP and JW. IS thanks the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for supporting this work by funding EXC 2075-390740016 under Germany’s Excellence Strategy. IS also acknowledges support by the Stuttgart Center for Simulation Science (SimTech).

Finally, the authors are grateful to Filip Tronarp for many helpful discussions concerning the theoretical part of this work.



## Appendix A. Proofs for Section 3.3

**Proof of Proposition 3.2** By Theorem 1, we have

$$\begin{aligned}
 m_i^{\mathbf{u}|\hat{D},\hat{\mathbf{f}}}(\mathbf{x}) &= m_i(\mathbf{x}) + (\hat{D}\mathcal{P}_{\mathbb{R}^m})[\mathbf{k}_{:,i}(\cdot, \mathbf{x})]^\top \left( (\hat{D}\mathcal{P}_{\mathbb{R}^m})\mathbf{k}(\hat{D}\mathcal{P}_{\mathbb{R}^m})' \right)^{-1} \left( \hat{\mathbf{f}} - \hat{D}\mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \right) \\
 &= m_i(\mathbf{x}) + \mathcal{P}_{\mathbb{R}^m}[\mathbf{k}_{:,i}(\cdot, \mathbf{x})]^\top \hat{D}^\top \left( \hat{D}\Sigma_{\mathbf{c}}\hat{D}^\top \right)^{-1} \hat{D} \left( \hat{D}^{-1}\hat{\mathbf{f}} - \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \right) \\
 &= m_i(\mathbf{x}) + \mathcal{P}_{\mathbb{R}^m}[\mathbf{k}_{:,i}(\cdot, \mathbf{x})]^\top \Sigma_{\mathbf{c}}^{-1} \left( \hat{D}^{-1}\hat{\mathbf{f}} - \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \right).
 \end{aligned}$$

Since  $\mathcal{P}_{\hat{\mathbb{U}}}$  is a bounded projection, we have  $\mathbb{U} = \text{ran}(\mathcal{P}_{\hat{\mathbb{U}}}) \oplus \ker(\mathcal{P}_{\hat{\mathbb{U}}}) = \hat{\mathbb{U}} \oplus \ker(\mathcal{P}_{\hat{\mathbb{U}}})$  (see Rudin, 1991, Section 5.16), where each  $\mathbf{u} \in \mathbb{U}$  decomposes uniquely into  $\mathbf{u} = \mathbf{u}_{\hat{\mathbb{U}}} + \mathbf{u}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}$  with  $\mathbf{u}_{\hat{\mathbb{U}}} \in \hat{\mathbb{U}}$  and  $\mathbf{u}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})} \in \ker(\mathcal{P}_{\hat{\mathbb{U}}})$ . It is clear that  $\mathbf{u}_{\hat{\mathbb{U}}} = \mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{u}]$ , and  $\mathbf{u}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})} = (\text{id} - \mathcal{P}_{\hat{\mathbb{U}}})[\mathbf{u}] = \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}[\mathbf{u}]$ . This implies

$$\begin{aligned}
 \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}^{\mathbf{u}|\hat{D},\hat{\mathbf{f}}}] &= \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] + \underbrace{\mathcal{P}_{\mathbb{R}^m}\mathbf{k}\mathcal{P}'_{\mathbb{R}^m}}_{=\Sigma_{\mathbf{c}}} \Sigma_{\mathbf{c}}^{-1} \left( \hat{D}^{-1}\hat{\mathbf{f}} - \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \right) \\
 &= \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] + \hat{D}^{-1}\hat{\mathbf{f}} - \mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \\
 &= \hat{D}^{-1}\hat{\mathbf{f}} = \mathbf{c}^{\text{MWR}}.
 \end{aligned}$$

Hence, we have

$$\mathcal{P}_{\hat{\mathbb{U}}}[\mathbf{m}^{\mathbf{u}|\hat{D},\hat{\mathbf{f}}}] = \sum_{i=1}^m \left( \mathcal{P}_{\mathbb{R}^n}[\mathbf{m}^{\mathbf{u}|\hat{D},\hat{\mathbf{f}}}]_i \right) \phi^{(i)} = \sum_{i=1}^m c_i^{\text{MWR}} \phi^{(i)} = \mathbf{u}^{\text{MWR}} \in \hat{\mathbb{U}}$$

and since  $\mathbb{U} = \hat{\mathbb{U}} \oplus \ker(\mathcal{P}_{\hat{\mathbb{U}}})$ , the statement follows. Moreover,  $\mathcal{P}_{\mathbb{R}^n}[\mathbf{m}^{\mathbf{u}|\hat{D},\hat{\mathbf{f}}}]$  is the mean of  $\mathbf{c} \mid \hat{D}\mathbf{c} - \hat{\mathbf{f}} = \mathbf{0}$  and its covariance matrix is given by

$$\begin{aligned}
 \Sigma^{\mathbf{c}|\hat{D},\hat{\mathbf{f}}} &= \Sigma_{\mathbf{c}} - \Sigma_{\mathbf{c}}\hat{D}^\top \left( \hat{D}\Sigma_{\mathbf{c}}\hat{D}^\top \right)^{-1} \hat{D}\Sigma_{\mathbf{c}} \\
 &= \Sigma_{\mathbf{c}} - \Sigma_{\mathbf{c}}\hat{D}^\top (\hat{D}^\top)^{-1} \Sigma_{\mathbf{c}}^{-1} \hat{D}^{-1} \hat{D}\Sigma_{\mathbf{c}} \\
 &= \Sigma_{\mathbf{c}} - \Sigma_{\mathbf{c}}\Sigma_{\mathbf{c}}^{-1}\Sigma_{\mathbf{c}} = \mathbf{0}.
 \end{aligned}$$

Consequently,  $\mathbf{c} \mid \hat{D}\mathbf{c} - \hat{\mathbf{f}} = \mathbf{0} \sim \delta_{\mathbf{c}^{\text{MWR}}}$ . ■

### Proof of Corollary 3.3

$$\begin{aligned}
 &\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}[\mathbf{m}^{\mathbf{u}|\hat{D},\hat{\mathbf{f}}}] (\mathbf{x}) \\
 &= \underbrace{\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}[\mathbf{m}]}_{=0}(\mathbf{x}) + (\delta_{\mathbf{x}} \circ \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})})\mathbf{k}(\hat{D}\mathcal{P}_{\mathbb{R}^m})' \left( (\hat{D}\mathcal{P}_{\mathbb{R}^m})\mathbf{k}(\hat{D}\mathcal{P}_{\mathbb{R}^m})' \right)^{-1} \left( \hat{\mathbf{f}} - \hat{D}\mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \right) \\
 &= \delta_{\mathbf{x}} \underbrace{[\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbb{U}}})}\mathbf{k}\mathcal{P}'_{\mathbb{R}^m}]}_{=0} \hat{D}^\top \left( (\hat{D}\mathcal{P}_{\mathbb{R}^m})\mathbf{k}(\hat{D}\mathcal{P}_{\mathbb{R}^m})' \right)^{-1} \left( \hat{\mathbf{f}} - \hat{D}\mathcal{P}_{\mathbb{R}^m}[\mathbf{m}] \right) = \mathbf{0}
 \end{aligned}$$

■

**Proof of Proposition 3.4** The process  $\mathbf{u}$  can be constructed as the sum of independent samples from the processes  $\mathcal{P}_{\hat{\mathbf{U}}}[\tilde{\mathbf{u}}]$  and  $\mathcal{P}_{\hat{\mathbf{U}}}[\tilde{\mathbf{u}} - \tilde{\mathbf{m}}]$ . This proves that the sample paths lie in  $\mathbb{U}$ . Since  $\mathcal{P}_{\hat{\mathbf{U}}}$  is idempotent, we have  $\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\mathcal{P}_{\hat{\mathbf{U}}} = \mathcal{P}_{\hat{\mathbf{U}}} - \mathcal{P}_{\hat{\mathbf{U}}}^2 = \mathcal{P}_{\hat{\mathbf{U}}} - \mathcal{P}_{\hat{\mathbf{U}}} = 0$  and  $\mathcal{P}_{\mathbb{R}^n}\mathcal{P}_{\hat{\mathbf{U}}} = (\mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbf{U}}})^{-1}\mathcal{P}_{\hat{\mathbf{U}}}^2 = (\mathcal{I}_{\mathbb{R}^m}^{\hat{\mathbf{U}}})^{-1}\mathcal{P}_{\hat{\mathbf{U}}} = \mathcal{P}_{\mathbb{R}^n}$ . It follows that

$$\begin{aligned} \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\tilde{\mathbf{k}}\mathcal{P}_{\mathbb{R}^n}^* &= \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\tilde{\mathbf{k}}\mathcal{P}_{\mathbb{R}^n}^* - \underbrace{\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\mathcal{P}_{\hat{\mathbf{U}}}\tilde{\mathbf{k}}}_{=0} \\ &\quad - \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\tilde{\mathbf{k}}\mathcal{P}_{\hat{\mathbf{U}}}\mathcal{P}_{\mathbb{R}^n}^* + 2\underbrace{\mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\mathcal{P}_{\hat{\mathbf{U}}}\tilde{\mathbf{k}}\mathcal{P}_{\hat{\mathbf{U}}}\mathcal{P}_{\mathbb{R}^n}^*}_{=0} \\ &= \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\tilde{\mathbf{k}}\mathcal{P}_{\mathbb{R}^n}^* - \mathcal{P}_{\ker(\mathcal{P}_{\hat{\mathbf{U}}})}\tilde{\mathbf{k}}\underbrace{(\mathcal{P}_{\mathbb{R}^n}\mathcal{P}_{\hat{\mathbf{U}}})^*}_{=\mathcal{P}_{\mathbb{R}^n}} = 0. \end{aligned}$$

■

## Appendix B. Proofs for Section 4

Using the rules of linear-Gaussian inference (Bishop, 2006), we can easily see that

$$\begin{aligned} \mathbf{f} &\sim \mathcal{GP}(m, k) \\ \mathbf{A}\mathbf{f}(\mathbf{X}) &\sim \mathcal{N}\left(\mathbf{A}m(\mathbf{X}), \mathbf{A}k(\mathbf{X}, \mathbf{X})\mathbf{A}^\top\right) \\ \mathbf{f} \mid \mathbf{A}\mathbf{f}(\mathbf{X}) + \mathbf{b} = \mathbf{y} &\sim \mathcal{GP}\left(m^{\mathbf{f}|\mathbf{y}}, k^{\mathbf{f}|\mathbf{y}}\right), \end{aligned}$$

where  $\mathbf{A} \in \mathbb{R}^{m \times n}$ ,  $\mathbf{X} = (\mathbf{x}_i)_{i=1}^n \in \mathbb{X}$ ,  $\mathbf{b} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  with  $\mathbf{b} \perp \mathbf{f}$  and

$$\begin{aligned} m^{\mathbf{f}|\mathbf{y}}(\mathbf{x}) &:= m(\mathbf{x}) + k(\mathbf{x}, \mathbf{X})\mathbf{A}^\top(\mathbf{A}k(\mathbf{X}, \mathbf{X})\mathbf{A}^\top + \boldsymbol{\Sigma})^\dagger(\mathbf{y} - (\mathbf{A}m + \boldsymbol{\mu})) \\ k^{\mathbf{f}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) &:= k(\mathbf{x}_1, \mathbf{x}_2) - k(\mathbf{x}_1, \mathbf{X})\mathbf{A}^\top(\mathbf{A}k(\mathbf{X}, \mathbf{X})\mathbf{A}^\top + \boldsymbol{\Sigma})^\dagger\mathbf{A}k(\mathbf{X}, \mathbf{x}_2). \end{aligned}$$

It is tempting to think that the above also extends to more general linear transformations of  $\mathbf{f}$  such as differentiation at a point  $\mathbf{x} \in \mathbb{X}$  and integration. Unfortunately, this is not the case, since the result from (Bishop, 2006) heavily uses the fact that, by definition, evaluations of the Gaussian process at a *finite* set of points follow a joint Gaussian distribution. However, differentiation at a point and integration are examples of linear functionals, i.e. linear maps from a vector space of functions to the real numbers, which operate on an (uncountably) infinite subset of random variables.

To generalize the result above to general linear operators  $\mathcal{L}$  mapping the paths of  $\mathbf{f}$  into  $\mathbb{R}^n$ , we will take the following route:

1. In Appendix B.2, we will show that, under certain conditions on  $\mathbf{f}$  and  $\mathbb{X}$ , the function  $\omega \mapsto \mathbf{f}(\cdot, \omega)$  is a Gaussian random variable with values in a separable Banach space  $\mathbb{B}$  of real-valued functions on  $\mathbb{X}$ . We introduce Gaussian random variables on separable Banach spaces and their essential properties in Appendix B.1.
2. Under the assumption that  $\mathcal{L}$  is continuous, we can use the transformation properties of Gaussian random variables on separable Banach spaces (see Lemma B.4) to show

that  $\mathcal{L}[f]$  and for  $\mathbf{X} \in \mathbb{X}^m$  also

$$\begin{pmatrix} f(\mathbf{X}) \\ \mathcal{L}[f] \end{pmatrix}$$

are Gaussian random variables on  $\mathbb{R}^n$  and  $\mathbb{R}^{m+n}$ , respectively.

3. Finally, in Appendix B.3, we can then use the well-known linear-Gaussian inference theorem (Bishop, 2006) to show that  $f \mid \mathcal{L}[f] = \mathbf{y}$  is again a Gaussian process.

In the following,  $\mathcal{B}(\mathbb{B})$  denotes the Borel  $\sigma$ -algebra generated by the norm topology on a Banach space  $\mathbb{B}$ .

### B.1 Gaussian Measures on Separable Banach Spaces

As stated before, in many cases, the function  $\omega \mapsto f(\cdot, \omega)$  will often turn out to be a Gaussian random variable with values in an infinite-dimensional separable Banach space  $\mathbb{B} \supseteq$  paths ( $f$ ) of real-valued functions on  $\mathbb{X}$  (see Proposition B.11).

**Definition B.1.** *Let  $\mathbb{B}$  be a real separable Banach space. A Borel probability measure  $\mu$  on  $(\mathbb{B}, \mathcal{B}(\mathbb{B}))$  is called Gaussian if every continuous linear functional  $l \in \mathbb{B}'$  is a univariate Gaussian random variable. A  $\mathbb{B}$ -valued random variable is called Gaussian if its law is Gaussian.*

Just as for Gaussian random variables on Euclidean vector space  $\mathbb{R}^n$ , we can define a mean and covariance (operator) for their counterparts on general separable Banach spaces.

**Proposition B.2.** *Let  $b$  be a Gaussian random variable on  $(\Omega, \mathcal{F}, P)$  with values in a real separable Banach space  $\mathbb{B}$ . Then there is a unique  $m_b \in \mathbb{B}$  such that  $l[m_b] = \mathbb{E}_b[l[b]]$  for any  $l \in \mathbb{B}'$ . We refer to  $m_b$  as the mean (vector) of  $b$ . Moreover, there is a unique bounded linear operator  $\mathcal{C}_b: \mathbb{B}' \rightarrow \mathbb{B}$  such that  $l_2[\mathcal{C}_b[l_1]] = \text{Cov}_b[l_1[b], l_2[b]]$  for any  $l_1, l_2 \in \mathbb{B}'$ , the so-called covariance operator of  $b$ .*

**Proof** Fernique's theorem (Da Prato and Zabczyk, 1992, Theorem 2.7) implies that  $\|b\|_{\mathbb{B}} \in L_p(\Omega, P)$  for all  $p \in \mathbb{N}_{\geq 1}$ . By assumption,  $b$  is measurable and  $\mathbb{B} \supset \text{ran}(b)$  is separable, which means that  $b$  is strongly measurable (Yosida, 1995, Section V.4, Pettis' Theorem). Since  $\|b\|_{\mathbb{B}} \in L_1(\Omega, P)$ , it follows that  $b$  is Bochner integrable (Yosida, 1995, Section V.5, Theorem 1). Let  $l \in \mathbb{B}'$ . By Corollary 2 in Yosida (1995, Section V.5) we then have

$$\mathbb{E}_b[l[b]] = \int_{\Omega} l[b(\omega)] dP(\omega) = l \left[ \underbrace{\int_{\Omega} b(\omega) dP(\omega)}_{=: m_b} \right].$$

Now assume that there is another  $\tilde{m}_b \in \mathbb{B}$  with  $l[\tilde{m}_b] = \mathbb{E}_b[l[b]]$ . Then

$$0 = \mathbb{E}_b[l[b]] - \mathbb{E}_b[l[b]] = l[\tilde{m}_b] - l[m_b] = l[\tilde{m}_b - m_b]$$

and since this holds for all  $l \in \mathbb{B}'$ , it follows that  $\tilde{m}_b = m_b$ .

Let  $l_1 \in \mathbb{B}'$ . Then  $\omega \mapsto l_1[b(\omega) - m_b](b(\omega) - m_b)$  is clearly weakly measurable and, since  $\mathbb{B}$  is separable, also strongly measurable (Yosida, 1995, Section V.4, Pettis' Theorem). As

above, Fernique's theorem shows that  $\|b\|_{\mathbb{B}} \in L_2(\Omega, \mathbb{P})$ . By the triangle inequality in  $\mathbb{B}$  and the fact that  $\mathbb{P}$  is a probability measure, we also have  $\|b(\cdot) - m_b\|_{\mathbb{B}} \in L_2(\Omega, \mathbb{P})$ . It follows that

$$\begin{aligned} \int_{\Omega} \|l_1[b(\omega) - m_b](b(\omega) - m_b)\|_{\mathbb{B}} \, d\mathbb{P}(\omega) &= \int_{\Omega} |l_1[b(\omega) - m_b]| \|b(\omega) - m_b\|_{\mathbb{B}} \, d\mathbb{P}(\omega) \\ &\leq \|l_1\|_{\mathbb{B}'} \int_{\Omega} \|b(\omega) - m_b\|_{\mathbb{B}} \|b(\omega) - m_b\|_{\mathbb{B}} \, d\mathbb{P}(\omega) \\ &= \|l_1\|_{\mathbb{B}'} \|\omega \mapsto \|b(\omega) - m_b\|_{\mathbb{B}}\|_{L_2(\Omega, \mathbb{P})}^2 < \infty, \end{aligned}$$

where  $\|l_1\|_{\mathbb{B}'} < \infty$ , since  $l_1$  is continuous. Let  $l_2 \in \mathbb{B}'$ . Again by Corollary 2 in Yosida (1995, Section V.5), we find that

$$\begin{aligned} \text{Cov}[l_1[b], l_2[b]] &= \int_{\Omega} l_1[b(\omega) - m_b] l_2[b(\omega) - m_b] \, d\mathbb{P}(\omega) \\ &= l_2 \left[ \int_{\Omega} \underbrace{l_1[b(\omega) - m_b](b(\omega) - m_b)}_{=: \mathcal{C}[l_1]} \, d\mathbb{P}(\omega) \right]. \end{aligned}$$

$\mathcal{C}_b$  is bounded, since

$$\begin{aligned} \|\mathcal{C}_b[l_1]\|_{\mathbb{B}} &\leq \int_{\Omega} \|l_1[b(\omega) - m_b](b(\omega) - m_b)\|_{\mathbb{B}} \, d\mathbb{P}(\omega) \\ &\leq \|l_1\|_{\mathbb{B}'} \underbrace{\|\omega \mapsto \|b(\omega) - m_b\|_{\mathbb{B}}\|_{L_2(\Omega, \mathbb{P})}^2}_{=: \|\mathcal{C}_b\|}, \end{aligned}$$

where  $\|\mathcal{C}_b\| < \infty$  because  $\|b(\cdot) - m_b\|_{\mathbb{B}} \in L_2(\Omega, \mathbb{P})$ . Uniqueness follows from an argument analogous to the one used to prove uniqueness of the mean.  $\blacksquare$

**Remark B.3.** *One can show that the mean and the covariance operator of a Gaussian random variable with values in a separable Banach space identify its law uniquely. Hence, we often write  $\mathcal{N}(m, \mathcal{C})$  to denote Gaussian measures on separable Banach spaces.*

### B.1.1 CONTINUOUS AFFINE TRANSFORMATIONS

Just as their finite-dimensional counterparts, Gaussian random variables with values in separable Banach spaces are closed under linear transformations as long as they are continuous. Moreover, the expressions for the transformed mean and covariance operator are analogous to the finite-dimensional case. For instance, we will use this result to show that  $\mathcal{L}[f]$  is an  $\mathbb{R}^n$ -valued random variable.

**Lemma B.4.** *Let  $b \sim \mathcal{N}(m, \mathcal{C})$  be a Gaussian random variable on  $(\Omega, \mathcal{F}, \mathbb{P})$  with values in a real separable Banach space  $\mathbb{B}$ . Let  $\mathcal{L}: \mathbb{B} \rightarrow \tilde{\mathbb{B}}$  be a bounded linear operator mapping into another real separable Banach space  $\tilde{\mathbb{B}}$ . Then  $\mathcal{L}[b] \sim \mathcal{N}(\mathcal{L}[m], \mathcal{L}\mathcal{C}\mathcal{L}')$ .*

**Proof**  $\mathcal{L}$  is continuous and hence Borel measurable, which means that  $\mathcal{L}[b]$  is a  $\tilde{\mathbb{B}}$ -valued random variable. Moreover, for  $\tilde{l} \in \tilde{\mathbb{B}}'$ , we have  $l := \tilde{l} \circ \mathcal{L} \in \mathbb{B}'$  and hence  $l[b] = \tilde{l}[\mathcal{L}[b]]$  is Gaussian. This implies that  $\mathcal{L}[b]$  is a  $\tilde{\mathbb{B}}$ -valued Gaussian random variable. Moreover, we

have  $\mathbb{E}_b[l[b]] = l[m] = \tilde{l}[\mathcal{L}[m]]$ , i.e.  $\mathcal{L}[m]$  is the mean of  $\mathcal{L}[b]$ . Let  $\tilde{l}_1, \tilde{l}_2 \in \mathbb{B}'$  and define  $l_i := \tilde{l}_i \circ \mathcal{L} \in \mathbb{B}'$  for  $i = 1, 2$ . Then

$$\text{Cov}_b[l_1[b], l_2[b]] = l_2[\mathcal{C}[l_1]] = (\tilde{l}_2 \circ \mathcal{L})[\mathcal{C}[\tilde{l}_1 \circ \mathcal{L}]] = \tilde{l}_2[(\mathcal{L}\mathcal{C}\mathcal{L}')[\tilde{l}_1]],$$

i.e.  $\mathcal{L}\mathcal{C}\mathcal{L}'$  is the covariance operator of  $\mathcal{L}[b]$ . ■

## B.2 Gaussian Processes as Gaussian Random Functions

We now introduced all necessary preliminaries to show that, under certain assumptions on  $f$  and  $\mathbb{X}$ , the function  $\omega \mapsto f(\cdot, \omega)$  is a Gaussian random variable with values in a special kind of separable Banach space, namely a *reproducing kernel Banach space (RKBS)*:

**Definition B.5** (Lin et al. 2022, Definition 2.1). *A reproducing kernel Banach space (RKBS)  $(\mathbb{B}, \|\cdot\|_{\mathbb{B}})$  is a Banach space of real-valued functions on a nonempty set  $\mathbb{X}$ , on which the point evaluation functionals  $\delta_{\mathbf{x}}$  for  $\mathbf{x} \in \mathbb{X}$  are continuous.*

We formulate Theorem 1 and Corollary 2, under the following set of assumptions:

**Assumption 1.** *Let  $f \sim \mathcal{GP}(m, k)$  be a Gaussian process prior with index set  $\mathbb{X}$  on the probability space  $(\Omega, \mathcal{F}, \mathbb{P})$ , whose paths lie in a real separable reproducing kernel Banach space (RKBS)  $\mathbb{B} \subset \mathbb{R}^{\mathbb{X}}$  such that  $\omega \mapsto f(\cdot, \omega)$  is a  $\mathbb{B}$ -valued Gaussian random variable.*

In the following, we prove that Assumption 1 is fulfilled for three important classes of Banach spaces, which often arise in the study of Gaussian Processes and PDEs.

**Proposition B.6.** *Assumption 1 holds for any GP with paths in a separable RKHS.*

**Proof** By definition,  $f(\mathbf{x}, \cdot)$  is a Gaussian random variable for every  $\mathbf{x} \in \mathbb{X}$ . Hence, Corollary 12 in (Berlinet and Thomas-Agnan, 2004, Chapter 4, Section 2, p.195) ensures that  $\omega \mapsto f(\cdot, \omega)$  is Borel measurable and thus a random variable, which is Gaussian by Theorem 91 in (Berlinet and Thomas-Agnan, 2004, Chapter 4, Section 3.1, p.196). ■

**Proposition B.7.** *Assumption 1 holds for any GP with paths in the space of continuous functions  $\mathbb{B} = C(\mathbb{X})$  with compact domain  $\mathbb{X} \subset \mathbb{R}^d$ .*

**Proof** In the following, we generalize the proof of Theorem 1(a) in Rajput and Cambanis (1972). By definition,  $f(\mathbf{x}, \cdot)$  is measurable for every  $\mathbf{x} \in \mathbb{X}$ . Moreover, since the paths of  $f$  lie in  $\mathbb{B}$ ,  $f(\cdot, \omega)$  is continuous for all  $\omega \in \Omega$ . Hence,  $f(\cdot, \cdot)$  is a Carathéodory function, which implies that  $\omega \mapsto f(\cdot, \omega)$  is Borel measurable (Aliprantis and Border, 2006, Definition 4.50 and Theorem 4.55). It remains to show that  $\omega \mapsto l[f]$  is Gaussian for every  $l \in \mathbb{B}'$ . To do so, we modify a construction from Alt (2012, Paragraph 4.22). For  $S \subset \mathbb{R}^d$ , define  $\text{diam}(S) := \sup_{\mathbf{x}, \mathbf{x}_0 \in S} (\|\mathbf{x} - \mathbf{x}_0\|)$ . Since  $\mathbb{X}$  is compact, there is a finite open cover  $\{\tilde{S}_i^{(k)}\}_{i=1}^{n^{(k)}} \subset \mathbb{R}^d$  with  $\text{diam}(\tilde{S}_i^{(k)}) < \frac{1}{k}$  for every  $k \in \mathbb{N}_{>0}$ . Then  $\{S_i^{(k)}\}_{i=1}^{n^{(k)}} \subset \mathbb{X}$  with

$$S_i^{(k)} := \mathbb{X} \cap \tilde{S}_i^{(k)} \setminus \bigcup_{j < i} \tilde{S}_j^{(k)}$$

is also a cover of  $\mathbb{X}$  and  $S_i^{(k)} \in \mathcal{B}(\mathbb{X})$ . Now choose  $\mathbf{x}_i^{(k)} \in S_i^{(k)}$  for all  $i = 1, \dots, n^{(k)}$  (w.l.o.g.  $S_i^{(k)} \neq \emptyset$ ). For  $f \in C(\mathbb{X})$ , we define

$$f_k = \sum_{i=1}^{n^{(k)}} f(\mathbf{x}_i^{(k)}) \chi_{S_i^{(k)}}.$$

We will now show that the sequence  $\{f_k\}_{k \in \mathbb{N}_{>0}}$  converges uniformly to  $f$ . By the Heine-Cantor theorem,  $f$  is uniformly continuous. Thus, for every  $\epsilon > 0$ , there is  $\delta(\epsilon) > 0$  such that  $|f(\mathbf{x}) - f(\mathbf{x}_0)| < \epsilon$  for  $\mathbf{x}, \mathbf{x}_0 \in \mathbb{X}$  with  $\|\mathbf{x} - \mathbf{x}_0\| < \delta(\epsilon)$ . Not let  $\epsilon > 0$ . Then for  $k > \frac{1}{\delta(\epsilon)}$  and any  $\mathbf{x} \in S_i^{(k)}$ , we have

$$\|\mathbf{x} - \mathbf{x}_i^{(k)}\| < \text{diam}(S_i^{(k)}) < \frac{1}{k} < \delta(\epsilon)$$

and thus  $|f(\mathbf{x}) - f(\mathbf{x}_i^{(k)})| < \epsilon$ . Consequently,

$$\|f - f_k\|_\infty = \max_{i=1, \dots, n^{(k)}} \sup_{\mathbf{x} \in S_i^{(k)}} |f(\mathbf{x}) - f(\mathbf{x}_i^{(k)})| < \max_{i=1, \dots, n^{(k)}} \epsilon = \epsilon.$$

Hence,  $f_k \rightarrow f$  uniformly.

Let  $l \in \mathbb{B}'$ . By the Riesz-Markov theorem, there is a signed Borel measure  $\lambda$  on  $\mathbb{X}$  such that  $l[f] = \int_{\mathbb{X}} f(\mathbf{x}) d\lambda(\mathbf{x})$ . Moreover, by the bounded convergence theorem and the fact that the  $f_k$  are simple with uniform limit  $f$ , we have

$$l[f] = \int_{\mathbb{X}} f(\mathbf{x}) d\lambda(\mathbf{x}) = \lim_{k \rightarrow \infty} \int_{\mathbb{X}} f_k(\mathbf{x}) d\lambda(\mathbf{x}) = \lim_{k \rightarrow \infty} \underbrace{\sum_{i=1}^{n^{(k)}} f(\mathbf{x}_i^{(k)}) \lambda(S_i^{(k)})}_{=l[f_k]}.$$

Now let  $f = f(\cdot, \omega)$  for some  $\omega \in \Omega$ . Since  $f(\mathbf{x}_1^{(k)}, \cdot), \dots, f(\mathbf{x}_{n^{(k)}}^{(k)}, \cdot)$  is jointly Gaussian, it follows that  $\omega \mapsto l[f_k(\cdot, \omega)]$  is a univariate Gaussian random variable. As the pointwise limit of Gaussian random variables is again a Gaussian random variable, we conclude that  $\omega \mapsto l[f(\cdot, \omega)]$  is Gaussian.  $\blacksquare$

The Banach spaces  $C^k(\overline{\mathbb{X}})$  of  $k$ -times differentiable functions on an open and bounded domain  $\mathbb{X} \subset \mathbb{R}^d$  with bounded and uniformly continuous partial derivatives and their subspaces (in particular the Hölder spaces) appear naturally in the study of strong solutions to linear PDEs. However, to allow for a flexible prior construction, we define a generalization of these spaces.

**Definition B.8.** Let  $\mathbb{X} \subset \mathbb{R}^d$  be open and bounded and let  $A \subset \mathbb{N}_0^d$  be a non-empty set of multi-indices such that  $\boldsymbol{\beta} \in A$  implies  $\boldsymbol{\alpha} \in A$  for every  $\boldsymbol{\alpha} \in \mathbb{N}_0^d$  with  $\boldsymbol{\alpha} \leq \boldsymbol{\beta}$ . We define  $C^A(\overline{\mathbb{X}})$  as the space of real-valued functions  $f$  on  $\mathbb{X}$ , for which all partial derivatives  $D^\alpha f$  with  $\boldsymbol{\alpha} \in A$  are bounded and uniformly continuous.

**Remark B.9.** One can show that  $C^A(\overline{\mathbb{X}})$  is a separable Banach space with norm

$$\|f\|_{C^A(\overline{\mathbb{X}})} := \max_{\boldsymbol{\alpha} \in A} \sup_{\mathbf{x} \in \mathbb{X}} |D^\alpha f(\mathbf{x})|.$$

Since every partial derivative  $D^\alpha f$  with  $\alpha \in A$  is bounded and uniformly continuous, it has a unique, bounded, continuous extension to the closure  $\overline{\mathbb{X}}$  of  $\mathbb{X}$  (Adams and Fournier, 2003).

**Proposition B.10.** *Assumption 1 holds for any GP with paths in the space  $\mathbb{B} = C^A(\overline{\mathbb{X}})$  of real-valued functions on an open and bounded domain  $\mathbb{X} \subset \mathbb{R}^d$  with bounded and uniformly continuous partial derivatives  $D^\alpha f$  for all  $f \in C^A(\overline{\mathbb{X}})$  and  $\alpha \in A$ .*

**Proof** In the following, we adapt the proof of Theorem 3.9 in (Adams and Fournier, 2003). We choose an arbitrary ordering  $\alpha_1, \dots, \alpha_n$  of the multi-indices in  $A$ , i.e.  $A = \{\alpha_1, \dots, \alpha_n\}$ , where  $n = |A|$ . Let

$$\mathbb{W} := \{(D^{\alpha_1} f, \dots, D^{\alpha_n} f) : f \in C^A(\overline{\mathbb{X}})\} \subset C(\overline{\mathbb{X}})^n,$$

where we interpret  $D^{\alpha_i} f$  as a function defined on the closure  $\overline{\mathbb{X}}$  of  $\mathbb{X}$  by the unique continuous extension mentioned in Remark B.9. Then  $(\mathbb{W}, \|\cdot\|_{\mathbb{W}})$  with

$$\|\mathbf{w}\|_{\mathbb{W}} := \max_{i=1, \dots, n} \|w_i\|_{C(\overline{\mathbb{X}})}$$

is a separable Banach space (Adams and Fournier, 2003, Theorem 1.23). Let  $\mathcal{I}: C^A(\overline{\mathbb{X}}) \rightarrow \mathbb{W}$  be the linear operator defined by  $\mathcal{I}[f]_i = D^{\alpha_i} f$ . The operator  $\mathcal{I}$  is surjective and an isometry by construction, and therefore it is also an isomorphism. It follows that  $\mathbb{W}$  is closed. Let  $l \in C^A(\overline{\mathbb{X}})'$ . Then  $l \circ \mathcal{I}^{-1} \in \mathbb{W}'$  and  $\|l\|_{C^A(\overline{\mathbb{X}})'} = \|l \circ \mathcal{I}^{-1}\|_{\mathbb{W}'}$ . By the Hahn-Banach theorem, there is a norm-preserving extension  $\tilde{l}$  of  $l \circ \mathcal{I}^{-1}$  to  $(C(\overline{\mathbb{X}})^n)'$ . This means that there are  $\tilde{l}_1, \dots, \tilde{l}_n \in C(\overline{\mathbb{X}})'$  such that

$$l[f] = (l \circ \mathcal{I}^{-1})[\mathcal{I}[f]] = \tilde{l}[\mathcal{I}[f]] = \sum_{i=1}^n \tilde{l}_i[D^{\alpha_i} f].$$

By the Riesz-Markov theorem, there are signed Borel measures  $\lambda_1, \dots, \lambda_n$  such that  $\tilde{l}_i[f] = \int_D f(\mathbf{x}) d\lambda_i(\mathbf{x})$ . If we apply the construction from the proof of Proposition B.7 to each summand, we arrive at

$$l[f] = \lim_{k \rightarrow \infty} \sum_{i=1}^n \sum_{j=1}^{n^{(k)}} D^{\alpha_i} f(\mathbf{x}_j^{(k)}) \lambda_i(S_j^{(k)}). \quad (\text{B.1})$$

Now let  $f = f(\cdot, \omega)$  for some  $\omega \in \Omega$ . Applying Theorems 1.91 and 1.92 from (Klenke, 2014) to Equation (B.1) inductively shows that  $\omega \mapsto l[f(\cdot, \omega)]$  is measurable for every  $l \in C^A(\overline{\mathbb{X}})'$ , i.e.  $\omega \mapsto f(\cdot, \omega)$  is weakly measurable. By the separability of  $C^A(\overline{\mathbb{X}})$ , this implies that  $\omega \mapsto f(\cdot, \omega)$  is also Borel measurable (Bogachev, 1998, Theorem A.3.7). Since pointwise limits of Gaussian random variables are Gaussian random variables, the family

$$\{\omega \mapsto D^{\alpha_i} f(\mathbf{x}_j^{(k)}, \omega)\}_{(i,j) \in \{1, \dots, n\} \times \{1, \dots, n^{(k)}\}}$$

of random variables is jointly Gaussian and by an argument analogous to the one made in the proof of Proposition B.7,  $l[f]$  is Gaussian.  $\blacksquare$

Having investigated conditions under which  $\omega \mapsto f(\cdot, \omega)$  is a Gaussian random variable, it remains to show what its mean and covariance operator are. Perhaps unsurprisingly, it turns out that they are strongly related to the mean and covariance function of the GP.

**Proposition B.11.** *Let Assumption 1 hold. Then  $m \in \mathbb{B}$ ,  $k(\mathbf{x}, \cdot) \in \mathbb{B}$  for all  $\mathbf{x} \in \mathbb{X}$ , and the mean and covariance operator of  $\omega \mapsto f(\cdot, \omega)$  are given by  $m$  and*

$$\mathcal{C}_k: \mathbb{B}' \rightarrow \mathbb{B}, \quad l \mapsto \mathcal{C}_k[l](\mathbf{x}) = l[k(\mathbf{x}, \cdot)], \quad (\text{B.2})$$

respectively.

**Proof** Since  $\omega \mapsto f(\cdot, \omega)$  is Gaussian, its mean and covariance operator  $m_f$  and  $\mathcal{C}_f$  exist by Proposition B.2 and we have  $m(\mathbf{x}) = \mathbb{E}_{\mathbb{P}}[f(\mathbf{x})] = \mathbb{E}_f[\delta_{\mathbf{x}}[f]] = \delta_{\mathbf{x}}[m_f]$  for all  $\mathbf{x} \in \mathbb{X}$  and

$$k(\mathbf{x}_1, \mathbf{x}_2) = \text{Cov}_{\mathbb{P}}[f(\mathbf{x}_1), f(\mathbf{x}_2)] = \text{Cov}_f[\delta_{\mathbf{x}_1}[f], \delta_{\mathbf{x}_2}[f]] = \mathcal{C}_f[\delta_{\mathbf{x}_1}](\mathbf{x}_2)$$

for all  $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{X}$ , since all point evaluation functionals are continuous on  $\mathbb{B}$ . Hence,  $m = m_f \in \mathbb{B}$  and  $k(\mathbf{x}, \cdot) = \mathcal{C}_f[\delta_{\mathbf{x}}] \in \mathbb{B}$  for all  $\mathbf{x} \in \mathbb{X}$ . Additionally, for all  $l \in \mathbb{B}'$  and  $\mathbf{x} \in \mathbb{X}$ ,

$$\mathcal{C}_f[l](\mathbf{x}) = \text{Cov}_f[l[f], \delta_{\mathbf{x}}[f]] = \text{Cov}_f[\delta_{\mathbf{x}}[f], l[f]] = l[\mathcal{C}_f[\delta_{\mathbf{x}}]] = l[k(\mathbf{x}, \cdot)] = \mathcal{C}_k[l](\mathbf{x}).$$

This shows that  $\mathcal{C}_f = \mathcal{C}_k$ . ■

### B.3 Proof of Theorem 1

Using the results from Appendices B.1 and B.2, particularly Proposition B.11 and Lemma B.4, we can now conduct the proof of Theorem 1 and Corollary 2 as outlined in the beginning of Appendix B.

The main theorem deals with the case in which we observe the GP through a finite number of linear functionals. This happens when conditioning on integral observations or on (Galerkin) projections as in Sections 3.2 and 3.3.

**Theorem 1.** *Let Assumption 1 hold and let  $\mathcal{L}: \mathbb{B} \rightarrow \mathbb{R}^n$  be a bounded linear operator. Then*

$$\mathcal{L}[f] \sim \mathcal{N}(\mathcal{L}[m], \mathcal{L}k\mathcal{L}'). \quad (4.1)$$

Let  $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  be an  $\mathbb{R}^n$ -valued Gaussian random vector with  $\boldsymbol{\epsilon} \perp f$ . Then, for any  $\mathbf{y} \in \mathbb{R}^n$ ,

$$f \mid \mathcal{L}[f] + \boldsymbol{\epsilon} = \mathbf{y} \sim \mathcal{GP}(m^{\mathbf{f}|\mathbf{y}}, k^{\mathbf{f}|\mathbf{y}}), \quad (4.2)$$

with conditional mean and covariance function given by

$$m^{\mathbf{f}|\mathbf{y}}(\mathbf{x}) = m(\mathbf{x}) + \mathcal{L}[k(\mathbf{x}, \cdot)]^{\top} (\mathcal{L}k\mathcal{L}' + \boldsymbol{\Sigma})^{\dagger} (\mathbf{y} - (\mathcal{L}[m] + \boldsymbol{\mu})), \quad (4.3)$$

and

$$k^{\mathbf{f}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) = k(\mathbf{x}_1, \mathbf{x}_2) - \mathcal{L}[k(\mathbf{x}_1, \cdot)]^{\top} (\mathcal{L}k\mathcal{L}' + \boldsymbol{\Sigma})^{\dagger} \mathcal{L}[k(\cdot, \mathbf{x}_2)]. \quad (4.4)$$



**Proof** By Lemma B.4,  $\mathcal{L}[f]$  is a Gaussian random variable with mean  $\mathcal{L}[m]$  and covariance matrix  $\Sigma$  with

$$\Sigma_{ij} = \mathcal{L}_i[\mathcal{C}[\mathcal{L}_j]] = \mathcal{L}_i[x \mapsto \mathcal{L}_j[k(x, \cdot)]] = (\mathcal{L}k\mathcal{L}')_{ij},$$

where we used Propositions B.2 and B.11. This proves Equation (4.1). Now let  $\mathbf{X} = (\mathbf{x}_i)_{i=1}^m \in \mathbb{X}^m$  and consider

$$\tilde{\mathcal{L}}: U \rightarrow \mathbb{R}^{m+n}, f \mapsto \begin{pmatrix} f(\mathbf{X}) \\ \mathcal{L}[f] \end{pmatrix}.$$

$\tilde{\mathcal{L}}$  is linear and bounded. Hence, by Lemma B.4 and the stability properties of independent Gaussian random variables on  $\mathbb{R}^{m+n}$ ,

$$\begin{pmatrix} f(\mathbf{X}) \\ \mathcal{L}[f] + \epsilon \end{pmatrix} = \tilde{\mathcal{L}}[f] + \begin{pmatrix} \mathbf{0}_{m \times n} \\ \mathbf{I}_n \end{pmatrix} \epsilon \sim \mathcal{N} \left( \begin{pmatrix} m(\mathbf{X}) \\ \mathcal{L}[m] + \mu \end{pmatrix}, \begin{pmatrix} k(\mathbf{X}, \mathbf{X}) & \Sigma^{\mathbf{X}, \mathcal{L}} \\ \Sigma^{\mathcal{L}, \mathbf{X}} & \mathcal{L}k\mathcal{L}' + \Sigma \end{pmatrix} \right),$$

where

$$\Sigma_{i,j}^{\mathbf{X}, \mathcal{L}} = \delta_{\mathbf{x}_i}[\mathcal{C}[\mathcal{L}_j]] = \mathcal{C}[\mathcal{L}_j](\mathbf{x}_i) = \mathcal{L}_j[k(\mathbf{x}_i, \cdot)] = (\delta_{\mathbf{X}}k\mathcal{L}')_{i,j}$$

and  $\Sigma^{\mathcal{L}, \mathbf{X}} = (\Sigma^{\mathbf{X}, \mathcal{L}})^\top$ . By the well-known conditioning theorem for Gaussian random variables in  $\mathbb{R}^{m+n}$ , we arrive at

$$f(\mathbf{X}) \mid \mathcal{L}[f] + \epsilon = \mathbf{y} \sim \mathcal{N} \left( \mu^{f(\mathbf{X})|\mathbf{y}}, \Sigma^{f(\mathbf{X})|\mathbf{y}} \right),$$

with

$$\mu^{f(\mathbf{X})|\mathbf{y}} = m(\mathbf{X}) + (\delta_{\mathbf{X}}k\mathcal{L}')(\mathcal{L}k\mathcal{L}' + \Sigma)^\dagger (\mathbf{y} - (\mathcal{L}[m] + \mu))$$

and

$$\Sigma^{f(\mathbf{X})|\mathbf{y}} = k(\mathbf{X}, \mathbf{X}) - (\delta_{\mathbf{X}}k\mathcal{L}')(\mathcal{L}k\mathcal{L}' + \Sigma)^\dagger (\mathcal{L}k\delta'_{\mathbf{X}}).$$

This shows that  $f = \{\omega \mapsto f(\mathbf{x}, \omega)\}_{\mathbf{x} \in \mathbb{X}}$  is a Gaussian process on the probability space

$$(\Omega, \mathcal{F}, \mathbb{P}(\cdot \mid \mathcal{L}[f] + \epsilon = \mathbf{y}))$$

where  $\mathbb{P}(\cdot \mid \mathcal{L}[f] + \epsilon = \mathbf{y})$  is a regular conditional probability whose existence is guaranteed, since  $\mathbb{R}^n$  is Polish. The mean and covariance function of the conditional process evaluated at  $\mathbf{x}_i$  are given by  $\mu_i^{f(\mathbf{X})|\mathbf{y}}$  and  $\Sigma_{i,i}^{f(\mathbf{X})|\mathbf{y}}$ . Since the points  $\mathbf{X}$  were chosen arbitrarily, this holds for any  $\mathbf{x} \in \mathbb{X}$ , which proves Equations (4.3) and (4.4).  $\blacksquare$

Finally, we address the archetypical case, in which both the prior  $f$  and the prior predictive  $\mathcal{L}[f] + \epsilon$  are Gaussian processes. This happens if the linear operator maps into a function space, in which point evaluation is continuous. In this article, this case occurred in Sections 3.1 and 3.2, where we inferred the strong solution of a PDE from observations of the PDE residual at a finite number of domain points.

**Corollary 2.** *Let Assumption 1 hold and let  $\mathcal{L}: \mathbb{B} \rightarrow \tilde{\mathbb{B}}$  be a bounded linear operator mapping into another real separable RKBS  $\tilde{\mathbb{B}} \subset \mathbb{R}^{\tilde{\mathbb{X}}}$ . Then*

$$\mathcal{L}[f] \sim \mathcal{GP}(\mathcal{L}[m], \mathcal{L}k\mathcal{L}'). \quad (4.8)$$

Let  $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  with values in  $\mathbb{R}^n$  and  $\boldsymbol{\epsilon} \perp f$ . Then, for  $\tilde{\mathbf{X}} = (\tilde{\mathbf{x}}_i)_{i=1}^n \in \tilde{\mathbb{X}}^n$  and  $\mathbf{y} \in \mathbb{R}^n$ ,

$$f | \mathcal{L}[f](\tilde{\mathbf{X}}) + \boldsymbol{\epsilon} = \mathbf{y} \sim \mathcal{GP}(m^{\mathbf{f}|\mathbf{y}}, k^{\mathbf{f}|\mathbf{y}}) \quad (4.9)$$

with

$$m^{\mathbf{f}|\mathbf{y}}(\mathbf{x}) := m(\mathbf{x}) + (k\mathcal{L}')(\mathbf{x}, \tilde{\mathbf{X}})^\top \left( (\mathcal{L}k\mathcal{L}')(\tilde{\mathbf{X}}, \tilde{\mathbf{X}}) + \boldsymbol{\Sigma} \right)^\dagger (\mathbf{y} - (\mathcal{L}[m](\tilde{\mathbf{X}}) + \boldsymbol{\mu})) \quad (4.10)$$

and

$$k^{\mathbf{f}|\mathbf{y}}(\mathbf{x}_1, \mathbf{x}_2) := k(\mathbf{x}_1, \mathbf{x}_2) - (k\mathcal{L}')(\mathbf{x}_1, \tilde{\mathbf{X}})^\top \left( (\mathcal{L}k\mathcal{L}')(\tilde{\mathbf{X}}, \tilde{\mathbf{X}}) + \boldsymbol{\Sigma} \right)^\dagger (\mathcal{L}k)(\tilde{\mathbf{X}}, \mathbf{x}_2) \quad (4.11)$$

If additionally  $\tilde{\mathbb{X}} = \mathbb{X}$ , then

$$\begin{pmatrix} f \\ \mathcal{L}[f] \end{pmatrix} \sim \mathcal{GP} \left( \begin{pmatrix} m \\ \mathcal{L}[m] \end{pmatrix}, \begin{pmatrix} k & k\mathcal{L}' \\ \mathcal{L}k & \mathcal{L}k\mathcal{L}' \end{pmatrix} \right). \quad (4.12)$$

**Proof** The linear operator  $\mathcal{L}[\cdot](\tilde{\mathbf{X}}): \mathbb{B} \rightarrow \mathbb{R}^n$  is bounded, since all point evaluation functionals on  $\tilde{\mathbb{B}}$  are continuous. Hence, Equations (4.8) to (4.11) follow directly from Theorem 1. Now let  $\mathbf{X} = (\mathbf{x}_i)_{i=1}^{m+m'} \in \mathbb{X}^{m+m'}$ . Then the linear operator

$$\mathcal{L}_{\mathbf{X}}: \mathbb{B} \rightarrow \mathbb{R}^{m+m'}, f \mapsto (f(\mathbf{x}_1) \ \dots \ f(\mathbf{x}_m) \ \mathcal{L}[f](\mathbf{x}_{m+1}) \ \dots \ \mathcal{L}[f](\mathbf{x}_{m+m'}))^\top$$

is bounded and  $\mathcal{L}_{\mathbf{X}}[f]$  is Gaussian by Theorem 1. This implies that  $\{\omega \mapsto \mathbf{f}^{\mathcal{L}}(\mathbf{x}, \omega)\}_{\mathbf{x} \in \mathbb{X}}$  with

$$\mathbf{f}^{\mathcal{L}}(\mathbf{x}, \omega) := \begin{pmatrix} f(\mathbf{x}, \omega) \\ \mathcal{L}[f(\cdot, \omega)](\mathbf{x}) \end{pmatrix}$$

is a 2-output Gaussian process. By Lemma B.4 and Propositions B.2 and B.11, its mean function is given by

$$\mathbf{m}^{\mathcal{L}}(\mathbf{x}) = \begin{pmatrix} \mathbb{E}_{\mathbb{P}}[\delta_{\mathbf{x}}[f]] \\ \mathbb{E}_{\mathbb{P}}[(\delta_{\mathbf{x}} \circ \mathcal{L})[f]] \end{pmatrix} = \begin{pmatrix} m(\mathbf{x}) \\ \mathcal{L}[m](\mathbf{x}) \end{pmatrix}$$

and its covariance function is given by

$$\begin{aligned} \mathbf{k}^{\mathcal{L}}(\mathbf{x}_1, \mathbf{x}_2) &= \begin{pmatrix} \text{COV}_{\mathbb{P}}[\delta_{\mathbf{x}_1}[f], \delta_{\mathbf{x}_2}[f]] & \text{COV}_{\mathbb{P}}[\delta_{\mathbf{x}_1}[f], (\delta_{\mathbf{x}_2} \circ \mathcal{L})[f]] \\ \text{COV}_{\mathbb{P}}[(\delta_{\mathbf{x}_1} \circ \mathcal{L})[f], \delta_{\mathbf{x}_2}[f]] & \text{COV}_{\mathbb{P}}[(\delta_{\mathbf{x}_1} \circ \mathcal{L})[f], (\delta_{\mathbf{x}_2} \circ \mathcal{L})[f]] \end{pmatrix} \\ &= \begin{pmatrix} \delta_{\mathbf{x}_1}[\mathcal{C}_k[\delta_{\mathbf{x}_2}]] & \delta_{\mathbf{x}_1}[\mathcal{C}_k[\delta_{\mathbf{x}_2} \circ \mathcal{L}]] \\ (\delta_{\mathbf{x}_1} \circ \mathcal{L})[\mathcal{C}_k[\delta_{\mathbf{x}_2}]] & (\delta_{\mathbf{x}_1} \circ \mathcal{L})[\mathcal{C}_k[\delta_{\mathbf{x}_2} \circ \mathcal{L}]] \end{pmatrix} \\ &= \begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_2) & (k\mathcal{L}')(\mathbf{x}_1, \mathbf{x}_2) \\ (\mathcal{L}k)(\mathbf{x}_1, \mathbf{x}_2) & (\mathcal{L}k\mathcal{L}')(\mathbf{x}_1, \mathbf{x}_2) \end{pmatrix}. \end{aligned}$$

This proves Equation (4.12). ■

## B.4 On Prior Selection

In order to apply Theorem 1 in practice, we need to construct our GP prior  $f$  such that

1.  $\omega \mapsto f(\cdot, \omega)$  is a Gaussian random variable on some suitably chosen RKBS  $\mathbb{B}$ , and
2.  $\mathcal{L}: \mathbb{B} \rightarrow \mathbb{R}^n$  is bounded.

Luckily, we can use existing results about the path spaces of Gaussian processes together with Propositions B.6 and B.10 to verify these assumptions. It is tempting to choose  $\mathbb{B} = \mathbb{H}_k$ , i.e. the RKHS of the GP's kernel  $k$ . However, this is only valid if  $\mathbb{H}_k$  is finite-dimensional.

**Remark B.12.** *Let  $f \sim \mathcal{GP}(m, k)$  be a Gaussian process with index set  $\mathbb{X}$  and let  $\mathbb{H}_k$  be the RKHS of the covariance function  $k$ . If  $\dim \mathbb{H}_k = \infty$ , then the sample paths of  $f$  almost surely do not lie in  $\mathbb{H}_k$ . We refer to (Kanagawa et al., 2018, Section 4) and Steinwart (2019) for more details on RKHS sample spaces.*

In the following, we will give example constructions of appropriate priors for GP regression tasks with linear operator observations.

### B.4.1 PRIORS FOR GP REGRESSION WITH LINEAR OPERATOR OBSERVATIONS

The spaces  $C^A(\overline{\mathbb{X}})$  from Definition B.8, particularly  $C^\beta(\overline{\mathbb{X}})$  with  $A := \{\alpha \in \mathbb{N}_0^d \mid \alpha \leq \beta\}$  (Wang et al., 2021) and  $C^k(\overline{\mathbb{X}})$  with  $A := \{\alpha \in \mathbb{N}_0^d \mid |\alpha| \leq k\}$ , are useful sample spaces for many GP regression tasks, since a large number of practically relevant observation functionals, including point evaluations of the paths and their partial derivatives, as well as integrals of the paths, are bounded on these spaces. Even though the functions in  $C^A(\overline{\mathbb{X}})$  are, technically speaking, only defined on the open and bounded set  $\mathbb{X} \subset \mathbb{R}^d$ , we can treat them as functions on the closure  $\overline{\mathbb{X}}$  of  $\mathbb{X}$  by continuous extension (see Remark B.9 for more details). In other words, we can evaluate functions in  $C^A(\overline{\mathbb{X}})$  on the boundary  $\partial\mathbb{X}$  of  $\mathbb{X}$ .

**Example B.1** (Tensor Products of Matérn Covariances). *Tensor products of 1D Matérn covariance functions*

$$k_\nu(\mathbf{x}_1, \mathbf{x}_2) = \prod_{i=1}^d k_{\nu_i}(x_{1,i}, x_{2,i})$$

are a particularly convenient choice of prior covariance function, since their hyperparameters directly control the differentiability of the sample paths. For an open and bounded domain  $\mathbb{X} \subset \mathbb{R}^d$ , Theorem 3.2 in Wang et al. (2021) implies that samples from a Gaussian process with mean function  $m$  and covariance function  $k_\nu$  lie in  $C^\beta(\overline{\mathbb{X}})$  (with probability 1) if  $\nu_i = \beta_i + \frac{1}{2}$  and  $m \in C^\beta(\overline{\mathbb{X}})$ . Any point-evaluated partial derivative  $f \mapsto D^\alpha f(\mathbf{x})$  with  $\alpha \leq \beta$  and  $\mathbf{x} \in \overline{\mathbb{X}}$  is continuous on  $C^\beta(\overline{\mathbb{X}})$ .

In Section 3.2, we use tensor products of Matérn covariance functions to construct the GP priors. In particular, we choose  $\nu_i = \frac{5}{2} = 2 + \frac{1}{2}$ , which implies that the sample paths of the prior lie in  $C^{(2,2)}(\overline{\mathbb{X}})$  and that all point-evaluated differential operators of order  $\leq 2$  are continuous linear functionals on the sample space. Hence, the assumptions of Corollary 2 are fulfilled, which means that the inference procedure used in this section is supported by our theoretical results above.

We conjecture that a Gaussian process with a multivariate Matérn covariance function  $k_{\nu=p+\frac{1}{2}}$  induces a Gaussian measure on the space  $\mathbb{B} = C^p(\overline{\mathbb{X}})$ .

For Gaussian processes with *exponentiated quadratic* or *Gaussian* covariance function, Assumption 1 holds for  $\mathbb{B} = C^p(\overline{\mathbb{X}})$  for all  $p \in \mathbb{N}_0$  (Karvonen, 2022, Section 5.2). Informally speaking, this can be seen as a limit of the argument above, since a Matérn covariance function approaches the Gaussian covariance function for  $\nu \rightarrow \infty$ .

Gaussian processes with parametric covariance function  $k(\mathbf{x}_1, \mathbf{x}_2) = \boldsymbol{\phi}(\mathbf{x}_1)^\top \boldsymbol{\Sigma} \boldsymbol{\phi}(\mathbf{x}_2)$  with features  $\boldsymbol{\phi}: \mathbb{X} \rightarrow \mathbb{R}^m$  and  $\boldsymbol{\Sigma} \in \mathbb{R}^{m \times m}$  positive-(semi)definite have paths in  $\mathbb{B}$  if  $\phi_i \in \mathbb{B}$ . In this case, Assumption 1 is also satisfied, since the Gaussian measure can be explicitly constructed as the law of the random function  $\mathbf{w}^\top \boldsymbol{\phi}$ , where  $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ .

#### B.4.2 PRIORS FOR INFERRING WEAK SOLUTIONS OF LINEAR PDES

A typical choice for the solution spaces  $\mathbb{U}$  of linear PDEs in weak formulation (see Section 2.1.1), are *Sobolev spaces* (Adams and Fournier, 2003). Unfortunately, it is impossible to construct a Gaussian process prior  $\mathbf{u}$ , whose paths are elements of a Sobolev space  $\mathbb{U}$ . This is due to the fact that Sobolev spaces are, technically speaking, not function spaces, but rather spaces of equivalence classes  $[u]_\sim$  of functions, which are equal almost everywhere (Adams and Fournier, 2003). By contrast, the path spaces of Gaussian processes are proper function spaces, which means that, in this setting, paths  $(\mathbf{u}) \subseteq \mathbb{U}$  is impossible.

Fortunately, if the path space  $\mathbb{B} \supset \text{paths}(\mathbf{u})$  of  $\mathbf{u}$  can be continuously embedded in  $\mathbb{U}$ , i.e. there is a continuous and injective linear operator  $\iota: \mathbb{B} \rightarrow \mathbb{U}$ , commonly referred to as an *embedding*, then the inference procedure above can still be applied. If such an embedding exists, we can interpret the paths of the GP as elements of  $\mathbb{B}$  by applying  $\iota$  implicitly. For instance,  $B[\mathbf{u}, v]$  is then a shorthand notation for  $B[\iota[\mathbf{u}], v]$ . Fortunately, since the embedding is assumed to be continuous, the conditions for GP inference with linear operator observations are still met when applying  $\iota$  implicitly. The canonical choice for the embedding in the case of Sobolev spaces is  $\iota[u] = [u]_\sim$ .

**Example B.2** (Matérn covariances and Sobolev spaces). *Kanagawa et al. (2018) show that, under certain assumptions, RKHS sample spaces of GP priors with Matérn covariance functions are continuously embedded in Sobolev spaces whose smoothness depends on the parameter  $\nu$  of the covariance function. To be precise, let  $\mathbb{D} \subset \mathbb{R}^d$  be open and bounded with Lipschitz boundary such that the cone condition (Adams and Fournier, 2003, Definition 4.6) holds. Denote by  $k_{\nu,l}$  the Matérn kernel with smoothness parameter  $\nu > 0$  and lengthscale  $l > 0$ . Then, with probability 1, the sample paths of a Gaussian process  $\mathbf{u}$  with covariance function  $k_{\nu,l}$  are contained in any RKHS  $\mathbb{H}_{k_{\nu',l'}}$  with  $l' > 0$  and*

$$0 < \underbrace{\nu' + \frac{d}{2}}_{=: m'} < \nu \tag{B.3}$$

(Kanagawa et al., 2018, Corollary 4.15 and Remark 4.15), i.e.  $\text{paths}(\mathbf{u}) \subset \mathbb{H}_{k_{\nu',l'}}$ . Moreover, if  $m' \in \mathbb{N}$ , then the RKHS  $\mathbb{H}_{k_{\nu',l'}}$  is norm-equivalent to the Sobolev space  $H^{m'}(\mathbb{D})$  (Kanagawa et al., 2018, Example 2.6). This implies that the canonical embedding

$$\iota: \mathbb{H}_{k_{\nu',l'}} \rightarrow H^{m'}(\mathbb{D}), f(\cdot, \omega) \mapsto [f(\cdot, \omega)]_{\sim_{H^{m'}(\mathbb{D})}} \tag{B.4}$$

is continuous.

For  $\mathbb{U} = H^{m'}(\mathbb{D})$ , the example above shows that the Matérn covariance function  $k_{\nu,l}$  with  $\nu = m' + \epsilon$  for any  $\epsilon > 0$  leads to an admissible GP prior. The choice  $\epsilon = \frac{1}{2}$  makes evaluating the covariance function particularly efficient (Rasmussen and Williams, 2006). For instance, in Section 3.3, we used  $\nu = \frac{3}{2} = 1 + \frac{1}{2}$  for a weak form linear PDE with solution space  $\mathbb{U} = H^1(\mathbb{D})$ . However, the elements of the Sobolev space  $H^m(\mathbb{D})$  are only  $m$ -times weakly differentiable, which means that  $H^2(\mathbb{D})$  is not an admissible choice in Section 3.2.

## Appendix C. Linear Partial Differential Equations

**Definition C.1** (Multi-index). *Using a  $d$ -dimensional multi-index  $\alpha \in \mathbb{N}_0^d$ , we can represent (mixed) partial derivatives of arbitrary order as*

$$\frac{\partial^{|\alpha|}}{\partial \mathbf{x}^\alpha} := \frac{\partial^{|\alpha|}}{\partial x_1^{(\alpha_1)} \dots \partial x_d^{(\alpha_d)}},$$

where  $|\alpha| := \sum_{i=1}^d \alpha_i$ . If the variables w.r.t. which we differentiate are clear from the context, we also denote this (mixed) partial derivative by  $D^\alpha$ . For two multi-indices  $\alpha, \alpha' \in \mathbb{N}_0^d$ , we write  $\alpha \leq \alpha'$  iff  $\alpha_i \leq \alpha'_i$  for all  $i = 1, \dots, d$ .

**Definition C.2** (Linear differential operator). *A linear differential operator  $\mathcal{D}: \mathbb{U} \rightarrow \mathbb{V}$  of order  $k$  between a space  $\mathbb{U}$  of  $\mathbb{R}^d$ -valued functions and a space  $\mathbb{V}$  of real-valued functions defined on some common open domain  $\mathbb{D} \subset \mathbb{R}^d$  is a linear operator that linearly combines partial derivatives up to  $k$ -th order of its input function, i.e.*

$$\mathcal{D}[\mathbf{u}] := \sum_{i=1}^{d'} \sum_{\alpha \in \mathbb{N}_0^d, |\alpha| \leq k} A_{i,\alpha} D^\alpha \mathbf{u}_i,$$

where  $A_{i,\alpha} \in \mathbb{R}$  for every  $i \in \{1, \dots, d'\}$  and every multi-index  $\alpha \in \mathbb{N}_0^d$  with  $|\alpha| \leq k$ .

### C.1 Weak Derivatives and Sobolev Spaces

**Definition C.3** (Test Function). *Let  $\mathbb{D} \subset \mathbb{R}^d$  be open and let*

$$C_c^\infty(\mathbb{D}) := \{\phi \in C^\infty(\mathbb{D}, \mathbb{R}) \mid \text{supp}(\phi) \subset \mathbb{D} \text{ is compact}\}$$

be the space of smooth functions with compact support in  $\mathbb{D}$ . A function  $\phi \in C_c^\infty(\mathbb{D})$  is dubbed test function and we refer to  $C_c^\infty(\mathbb{D})$  as the space of test functions.

**Theorem C.4** (Sobolev Spaces<sup>10</sup>). *Let  $\mathbb{D} \subset \mathbb{R}^d$  be open,  $k \in \mathbb{N}_{>0}$ , and  $p \in [1, \infty) \cup \{\infty\}$ . The functional*

$$\|u\|_{k,p,\mathbb{D}} := \begin{cases} \left( \sum_{|\alpha| \leq k} \|D^\alpha u\|_{L_p(\mathbb{D})}^p \right)^{1/p} & \text{if } p < \infty, \\ \max_{|\alpha| \leq k} \|D^\alpha u\|_{L_\infty(\mathbb{D})} & \text{if } p = \infty, \end{cases} \quad (\text{C.1})$$

10. This theorem is a summary of (Adams and Fournier, 2003, Definitions 3.1 and 3.2 and Theorems 3.3 and 3.6)

where the  $D^\alpha$  are weak partial derivatives, is called a Sobolev norm. A Sobolev norm  $\|u\|_{k,p,\mathbb{D}}$  is a norm on subspaces of  $L_p(\mathbb{D})$ , on which the right-hand side is well-defined and finite. A Sobolev space of order  $k$  is defined as the subspace

$$W^{k,p}(\mathbb{D}) := \{u \in L_p(\mathbb{D}) \mid D^\alpha u \in L_p(\mathbb{D}) \text{ for } |\alpha| \leq k\}.$$

of  $L_p$ . Sobolev spaces  $W^{k,p}(\mathbb{D})$  are Banach spaces under the Sobolev norm  $\|\cdot\|_{k,p,\mathbb{D}}$ . The Sobolev space  $H^k(\mathbb{D}) := W^{2,k}(\mathbb{D})$  is a separable Hilbert space with inner product

$$\langle u_1, u_2 \rangle_{k,\mathbb{D}} := \sum_{|\alpha| \leq k} \langle D^\alpha u_1, D^\alpha u_2 \rangle_{L_2(\mathbb{D})} \quad (\text{C.2})$$

and norm  $\|\cdot\|_{k,\mathbb{D}} := \sqrt{\langle \cdot, \cdot \rangle_{k,\mathbb{D}}} = \|\cdot\|_{k,2,\mathbb{D}}$ .

## Bibliography

- Robert A. Adams and John J. F. Fournier. *Sobolev Spaces*, volume 140 of *Pure and Applied Mathematics*. Elsevier, second edition, 2003.
- Christian Agrell. Gaussian processes with linear operator inequality constraints. *Journal of Machine Learning Research*, 20(135):1–36, 2019.
- Christopher G. Albert. Gaussian processes for data fulfilling linear differential equations. *Proceedings of the 39th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, 33(1), 2019. doi:10.3390/proceedings2019033005.
- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer, Berlin, Heidelberg, third edition, 2006. doi:10.1007/3-540-29587-9.
- Hans Wilhelm Alt. *Lineare Funktionalanalysis: Eine anwendungsorientierte Einführung*. Springer, Berlin, Heidelberg, 2012. doi:10.1007/978-3-642-22261-0.
- Mauricio Alvarez, David Luengo, and Neil D. Lawrence. Latent force models. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 5, pages 9–16, Clearwater Beach, Florida, USA, 2009.
- Iskander Azangulov, Andrei Smolensky, Alexander Terenin, and Viacheslav Borovitskiy. Stationary kernels and Gaussian processes on Lie groups and their homogeneous spaces i: the compact case. *arXiv preprint arXiv:2208.14960*, 2022.
- Alain Berlinet and Christine Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer, first edition, 2004. doi:10.1007/978-1-4419-9096-9.
- Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. Springer, New York, first edition, 2006.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973. doi:10.1086/260062.

- Vladimir Igorevich Bogachev. *Gaussian Measures*, volume 62 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, Rhode Island, 1998.
- David Borthwick. *Introduction to Partial Differential Equations*. Universitext. Springer, first edition, 2018. doi:10.1007/978-3-319-48936-0.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/google/jax>.
- Jon Cockayne, Chris Oates, Tim Sullivan, and Mark Girolami. Probabilistic numerical methods for PDE-constrained Bayesian inverse problems. In Geert Verdoolaege, editor, *Proceedings of the 36th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 1853 of *AIP Conference Proceedings*, pages 060001–1 – 060001–8, 2017. doi:10.1063/1.4985359.
- Jon Cockayne, Chris J. Oates, Ilse C.F. Ipsen, and Mark Girolami. A Bayesian conjugate gradient method (with discussion). *Bayesian Analysis*, 14(3):937–1012, 2019a. doi:10.1214/19-BA1145.
- Jon Cockayne, Chris J. Oates, T. J. Sullivan, and Mark Girolami. Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789, 2019b. doi:10.1137/17M1139357.
- Guiseppe Da Prato and Jerzy Zabczyk. *Stochastic Equations in Infinite Dimensions*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1992. doi:10.1017/CBO9780511666223.
- Lawrence C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, Rhode Island, second edition, 2010.
- Gregory E. Fasshauer. Solving partial differential equations by collocation with radial basis functions. In Alain Le Méhauté, Christophe Rabut, and Larry L. Schumaker, editors, *Surface Fitting and Multiresolution Methods*, pages 131–138. Vanderbilt University Press, Nashville, TN, 1997.
- Gregory E. Fasshauer. Solving differential equations with radial basis functions: multilevel methods and smoothing. *Advances in Computational Mathematics*, 11:139–159, November 1999. doi:10.1023/A:1018919824891.
- C. A. J. Fletcher. *Computational Galerkin Methods*. Scientific Computation. Springer, Berlin, Heidelberg, first edition, 1984. doi:10.1007/978-3-642-85949-6.
- Jean Baptiste Joseph Fourier. *Théorie analytique de la chaleur*. Firmin Didot, 1822. doi:10.1017/CBO9780511693229.
- Mark Girolami, Eky Febrianto, Yin Ge, and Fehmi Cirak. The statistical finite element method (statFEM) for coherent synthesis of observation data and model predictions. *Computer Methods in Applied Mechanics and Engineering*, 275:113533, 2021. doi:10.1016/j.cma.2020.113533.

- Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. The Johns Hopkins University Press, Baltimore, fourth edition, 2013.
- Thore Graepel. Solving noisy linear operator equations by Gaussian processes: Application to ordinary and partial differential equations. In *Proceedings of the 20th International Conference on Machine Learning*, pages 234–241. AAAI Press, 2003.
- Bernard Haasdonk and Hans Burkhardt. Invariant kernel functions for pattern analysis and machine learning. *Machine learning*, 68(1):35–61, 2007.
- Philipp Hennig, Michael A. Osborne, and Mark Girolami. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society A*, 471(2179), 2015. doi:10.1098/rspa.2015.0142.
- Philipp Hennig, Michael A. Osborne, and Hans P. Kersting. *Probabilistic Numerics: Computation as Machine Learning*. Cambridge University Press, June 2022. doi:10.1017/9781316681411.
- David S. Holder, editor. *Electrical Impedance Tomography: Methods, History and Applications*. Institute of Physics Medical Physics Series. Institute of Physics Publishing, Bristol, 2005.
- Peter Holderrieth, Michael J Hutchinson, and Yee Whye Teh. Equivariant learning of stochastic fields: Gaussian processes and steerable conditional neural processes. In *International Conference on Machine Learning*, pages 4297–4307. PMLR, 2021.
- Motonobu Kanagawa, Philipp Hennig, Dino Sejdinovic, and Bharath K. Sriperumbudur. Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*, 2018.
- George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021. doi:https://doi.org/10.1038/s42254-021-00314-5.
- Toni Karvonen. Small Sample Spaces for Gaussian Processes, March 2022.
- Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the 4th Eurographics Symposium on Geometry Processing*, volume 7, 2006.
- Achim Klenke. *Probability Theory: A Comprehensive Course*. Universitext. Springer, London, second edition, 2014. doi:10.1007/978-1-4471-5361-0.
- Nicholas Krämer, Jonathan Schmidt, and Philipp Hennig. Probabilistic numerical method of lines for time-dependent partial differential equations. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 151, pages 625–639. PMLR, 2022.
- Benny Lautrup. The PDE’s of continuum physics. In *Proceedings of the Workshop on PDE methods in Computer Graphics*, 2005.



- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Graph kernel network for partial differential equations. In *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2020. doi:10.48550/arXiv.2003.03485.
- Zongyi Li, Nikola B. Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew M. Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. doi:10.48550/arXiv.2010.08895.
- John H. Lienhard, IV and John H. Lienhard, V. *A Heat Transfer Textbook*. Phlogiston Press, Cambridge, MA, fifth edition, 2020.
- Rong Rong Lin, Hai Zhang Zhang, and Jun Zhang. On reproducing kernel Banach spaces: Generic definitions and unified framework of constructions. *Acta Mathematica Sinica, English Series*, 38(8):1459–1483, August 2022. doi:10.1007/s10114-022-1397-7.
- Anders Logg, Kent-Andre Mardal, and Garth Wells, editors. *Automated Solution of Differential Equations by the Finite Element Method*, volume 84 of *Lecture Notes in Computational Science and Engineering*. Springer, Berlin, Heidelberg, 2012. doi:10.1007/978-3-642-23099-8.
- James Clerk Maxwell. A dynamical theory of the electromagnetic field. *Philosophical transactions of the Royal Society of London*, 155:459–512, 1865.
- Pierre Michaud. A simple model of processor temperature for deterministic turbo clock frequency. resreport RR-9308, Inria Rennes, 2019. URL <https://hal.inria.fr/hal-02391970>.
- Chris J. Oates and Tim J. Sullivan. A modern retrospective on probabilistic numerics. *Statistics and Computing*, 29:1335–1351, 2019. doi:10.1007/s11222-019-09902-z.
- Houman Owhadi and Clint Scovel. Conditioning Gaussian measure on Hilbert space. *Journal of Mathematical and Statistical Analysis*, 1(109), 2018.
- Houman Owhadi, Clint Scovel, and Florian Schäfer. Statistical numerical approximation. *Notices of the American Mathematical Society*, 66(10):1608–1617, 2019. doi:10.1090/noti1963.
- Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Machine learning of linear differential equations using Gaussian processes. *Journal of Computational Physics*, 348:683–693, 2017. doi:10.1016/j.jcp.2017.07.050.
- Maziar Raissi, Paris Perdikaris, and George Em 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. doi:<https://doi.org/10.1016/j.jcp.2018.10.045>.
- Balram S. Rajput and Stamatis Cambanis. Gaussian processes and Gaussian measures. *The Annals of Mathematical Statistics*, 43(6):1944–1952, 1972. doi:10.1214/aoms/1177690865.

- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, London, England, 2006.
- Marco Reisert and Hans Burkhardt. Learning equivariant functions with matrix valued kernels. *Journal of Machine Learning Research*, 8(3), 2007.
- Walter Rudin. *Functional Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, New York, second edition, 1991.
- Simo Särkkä. Linear operators and stochastic partial differential equations in Gaussian process regression. In *Artificial Neural Networks and Machine Learning – ICANN 2011*, pages 151–158, Berlin, Heidelberg, 2011. doi:10.1007/978-3-642-21738-8\_20.
- Simo Särkkä, Arno Solin, and Jouni Hartikainen. Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing: A look at Gaussian process regression through Kalman filtering. *IEEE Signal Processing Magazine*, 30(4):51–61, 2013. doi:10.1109/MSP.2013.2246292.
- Ingo Steinwart. Convergence types and rates in generic Karhunen-Loève expansions with applications to sample path properties. *Potential Analysis*, 51:361–395, 2019. doi:10.1007/s11118-018-9715-5.
- Bastian von Harrach. Numerik partieller differentialgleichungen. Lecture Notes, 2021. URL [https://www.math.uni-frankfurt.de/~harrach/lehre/Numerik\\_PDGL.pdf](https://www.math.uni-frankfurt.de/~harrach/lehre/Numerik_PDGL.pdf).
- Junyang Wang, Jon Cockayne, Oksana Chkrebtii, Tim J. Sullivan, and Chris J. Oates. Bayesian numerical methods for nonlinear partial differential equations. *Statistics and Computing*, 31(55), 2021. doi:10.1007/s11222-021-10030-w.
- Jonathan Wenger and Philipp Hennig. Probabilistic linear solvers for machine learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Jonathan Wenger, Nicholas Krämer, Marvin Pförtner, Jonathan Schmidt, Nathanael Bosch, Nina Effenberger, Johannes Zenn, Alexandra Gessner, Toni Karvonen, François-Xavier Briol, Maren Mahsereci, and Philipp Hennig. ProNum: Probabilistic numerics in python, 2021.
- Kôsaku Yosida. *Functional Analysis*, volume 123 of *Classics in Mathematics*. Springer, sixth edition, 1995. doi:10.1007/978-3-642-61859-8.