
Probabilistic Numerical Method of Lines for Time-Dependent Partial Differential Equations

Nicholas Krämer
University of Tübingen
Tübingen, Germany

Jonathan Schmidt
University of Tübingen
Tübingen, Germany

Philipp Hennig
University of Tübingen and
Max Planck Institute for Intelligent Systems
Tübingen, Germany

Abstract

This work develops a class of probabilistic algorithms for the numerical solution of nonlinear, time-dependent partial differential equations (PDEs). Current state-of-the-art PDE solvers treat the space- and time-dimensions separately, serially, and with black-box algorithms, which obscures the interactions between spatial and temporal approximation errors and misguides the quantification of the *overall* error. To fix this issue, we introduce a probabilistic version of a technique called *method of lines*. The proposed algorithm begins with a Gaussian process interpretation of finite difference methods, which then interacts naturally with filtering-based probabilistic ordinary differential equation (ODE) solvers because they share a common language: Bayesian inference. Joint quantification of space- and time-uncertainty becomes possible without losing the performance benefits of well-tuned ODE solvers. Thereby, we extend the toolbox of probabilistic programs for differential equation simulation to PDEs.

1 INTRODUCTION

This work develops a class of probabilistic numerical (PN) algorithms for the solution of initial value problems based on partial differential equations (PDEs). PDEs are a widely-used way of modelling physical interdependencies between temporal and spatial variables. With the recent advent of physics-informed neural networks (Raissi et al., 2019), neural operators (Li et al.,

2021; Lu et al., 2021), and neural ordinary/partial differential equations (Chen et al., 2018; Gelbrecht et al., 2021), PDEs have rapidly gained popularity in the machine learning community, too. Let F , h , and g be given nonlinear functions. Let the domain $\Omega \subset \mathbb{R}^d$ with boundary $\partial\Omega$ be sufficiently well-behaved so that the PDE (below) has a unique solution.¹ \mathcal{D} shall be a differential operator. The reader may think of the Laplacian, $\mathcal{D} = \sum_{i=1}^d \frac{\partial^2}{\partial x_i^2}$, but the algorithm is not restricted to this case. The goal is to approximate an unknown function $u : \Omega \rightarrow \mathbb{R}^L$ that solves

$$\frac{\partial}{\partial t} u(t, x) = F(t, x, u(t, x), \mathcal{D}u(t, x)), \quad (1)$$

for $t \in [t_0, t_{\max}]$ and $x \in \Omega$, subject to initial condition

$$u(t_0, x) = h(x), \quad x \in \Omega, \quad (2)$$

and boundary conditions $\mathcal{B}u(t, x) = g(x)$, $x \in \partial\Omega$. The differential operator \mathcal{B} is usually the identity (Dirichlet conditions) or the derivative along normal coordinates (Neumann conditions). Except for only a few problems, PDEs do not admit closed-form solutions, and numerical approximations become necessary. This also affects machine learning strategies such as physics-informed neural networks or neural operators, for example, because they rely on fast generation of training data.

One common strategy for solving PDEs, called the *method of lines* (MOL; Schiesser, 2012), first discretises the spatial domain Ω with a grid $\mathbb{X} := (x_0, \dots, x_N)$, and then uses this grid to approximate the differential operator \mathcal{D} with a matrix-vector product

$$(\mathcal{D}u)(t, \mathbb{X}) \approx Du(t, \mathbb{X}), \quad D \in \mathbb{R}^{(N+1) \times (N+1)}, \quad (3)$$

where we use the notation $u(t, \mathbb{X}) = (u(t, x_n))_{n=0}^N$. Replacing the differential operator \mathcal{D} with the matrix D turns the PDE into a system of ordinary differential

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

¹One common assumption is that Ω must be open and bounded, and that $\partial\Omega$ must be differentiable. But requirements vary across differential equations (Evans, 2010).




Figure 1: *PNMOL fixes bad calibration of the MOL/ODE filter combination*: Posterior means and error/uncertainty ratios of PNMOL (left) and MOL with a conventional PN solver (right) on a fine time grid (y -axis) and a coarse space grid (x -axis) for the heat equation. The means are indistinguishable (*/*2). PN w/ MOL is poorly calibrated (error/uncertainty ratios $\sim 10^5$; */1), but PNMOL acknowledges all inaccuracies.

equations (ODEs; more information about D follows). Standard ODE solvers can then numerically solve the resulting initial value problem.

This approach has one central problem. Discretising the spatial domain, and only then applying an ODE solver, turns the PDE solver into a *pipeline of two numerical algorithms* instead of a single algorithm. This is bad because as a result of this serialisation, the error estimates returned by the ODE solver are unreliable. The solver lacks crucial information about whether the spatial grid consists of, say, $N = 4$ or $N = 10^7$ points. It is intuitive that a coarse spatial grid puts a lower bound on the overall precision, even if the ODE solver uses small time steps. But since this is not “known” to the ODE solver, not even to a probabilistic one, it may waste computational resources by needlessly decreasing its step-size and may deliver (severely) overconfident uncertainty estimates (for example, Fig. 1). This phenomenon will be confirmed by the experiments in Section 7. Such overconfident uncertainty estimates are essentially useless to a practitioner, which complicates the usage of method of lines approaches within probabilistic simulation of differential equations; for example, in the context of latent force inference or inverse problems.

The present work provides a solution to this problem by revealing a way of tracking spatiotemporal correlations in the PDE solutions while preserving the computational advantages of traditional MOL through a probabilistic numerical solver. The main idea of the present paper is the following. The error that is introduced by approximating the differential operator \mathcal{D} with a matrix D can be quantified probabilistically, and its time-evolution acts as a latent force on the resulting ODE. The combined posterior over the PDE solution and the latent force can be computed efficiently with a PN ODE solver. More specifically, we contribute:

Probabilistic discretisation Sections 2 and 3 present a probabilistic technique for constructing finite-

dimensional approximations of \mathcal{D} that quantify the discretisation error. It translates an unsymmetric collocation method (Kansa, 1990) to the language of Gaussian processes and admits efficient, sparse approximations akin to radial-basis-function-generated finite-difference formulas (Tolstykh and Shirobokov, 2003).

Calibrated PDE filter Building on the probabilistic discretisation, Section 4 explains how the discretisation and the ODE solver need not be treated as two separate entities. By applying the idea behind the method of lines to a *global* Gaussian process prior, it becomes evident how the discretisation uncertainty naturally evolves over time. With a filtering-based probabilistic ODE solver (“ODE filter”), posteriors over the latent error process and the PDE solution can be inferred in linear time, and without ignoring spatial uncertainties like traditional MOL algorithms do.

Section 5 explains hyperparameter choices and calibration. Section 6 discusses connections to existing work. Section 7 demonstrates the advantages of the approach over conventional MOL. Altogether, PNMOL enriches the toolbox of probabilistic simulations of differential equations by a calibrated and efficient PDE solver.

2 PN DISCRETISATION

This section derives how to approximate a differential operator \mathcal{D} by a matrix D . Let $u_x(z) \sim \mathcal{GP}(0, \gamma^2 k_x)$ be a Gaussian process on Ω with covariance kernel $k_x(z, z')$ and output scale $\gamma > 0$. The presentation in Section 4 will assume a Gaussian process prior $u(t, x) \sim \mathcal{GP}(0, \gamma^2 k_t \otimes k_x)$, thus the notation “ u_x ” motivates u_x as “the x -part of $u(t, x)$ ”.² The boundary operator \mathcal{B} is ignored in the remainder because it can be treated in the same way as \mathcal{D} . Complete formulas including \mathcal{B} , are in Supplement A.

²Here, u_x does not refer to the partial derivative $\frac{\partial}{\partial x} u$, which it sometimes does in the PDE literature.

The objective is to approximate the PDE dynamics in a way that circumvents the differential operator \mathcal{D} ,

$$f(t, \mathbb{X}, u_x(\mathbb{X})) \approx F(t, \mathbb{X}, u_x(\mathbb{X}), (\mathcal{D}u_x)(\mathbb{X})), \quad (4)$$

that is, $\mathcal{D}u_x$ disappears from Eq. (1) because f replaces F . For linear differential operators \mathcal{D} , this reduces to replacing $(\mathcal{D}u_x)(\mathbb{X})$ with a matrix-vector product $\mathcal{D}u_x(\mathbb{X})$, $D \in \mathbb{R}^{(N+1) \times (N+1)}$ (recall Eq. (3)), for example, by using the method proposed below. \mathcal{D} must not be present in the discretised model, because otherwise, the PDE does not translate into a system of ODEs, and we cannot proceed with the method of lines. But with an approximate derivative based on only matrix-vector arithmetic, the computational efficiency of ODE solvers can be exploited to infer an approximate PDE solution.

Since u is a Gaussian process, applying the linear operator \mathcal{D} yields another Gaussian process $\mathcal{D}u_x \sim \mathcal{GP}(0, \mathcal{D}^2 k_x)$, where we abbreviate $\mathcal{D}^2 k_x(z, z') = \mathcal{D}\mathcal{D}^* k_x(z, z')$ and \mathcal{D}^* is the adjoint of \mathcal{D} (in the present context, this means applying \mathcal{D} to z' instead of z). Conditioning u_x on realisations of $u_x(\mathbb{X})$, and then applying \mathcal{D} , yields another Gaussian process with moments

$$\hat{m}(z) = W_{\mathbb{X}}(z)u_x(\mathbb{X}), \quad (5a)$$

$$\hat{k}(z, z') = (\mathcal{D}^2 k_x)(z, z') - W_{\mathbb{X}}(z)k_x(\mathbb{X}, \mathbb{X})W_{\mathbb{X}}(z')^\top, \quad (5b)$$

$$W_{\mathbb{X}}(z) := (\mathcal{D}k_x)(z, \mathbb{X})k_x(\mathbb{X}, \mathbb{X})^{-1}. \quad (5c)$$

Let $\xi_x \sim \mathcal{GP}(0, \gamma^2 \hat{k})$. ξ_x is independent of u_x , since \hat{k} only depends on \mathcal{D} , k_x , and \mathbb{X} , not on $u_x(\mathbb{X})$. Then,

$$W_{\mathbb{X}}(z)u_x(\mathbb{X}) + \xi_x(z) \sim \mathcal{GP}(0, \mathcal{D}\mathcal{D}^* k_x) \quad (6)$$

follows, which implies

$$p(\mathcal{D}u_x(\cdot)) = p(W_{\mathbb{X}}(\cdot)u_x(\mathbb{X}) + \xi(\cdot)). \quad (7)$$

In particular, when evaluated on the grid, we obtain the matrix-vector formulation of the differential operator

$$(\mathcal{D}u_x)(\mathbb{X}) = W_{\mathbb{X}}(\mathbb{X})u_x(\mathbb{X}) + \xi_x(\mathbb{X}). \quad (8)$$

Eq. (8) yields $\mathcal{D}u_x$ solely as a transformation of values of u_x with known, additive Gaussian noise ξ_x . Altogether, the above derivation identifies the *differentiation matrix* D and the *error covariance* E

$$D := W_{\mathbb{X}}(\mathbb{X}), \quad E := \hat{k}(\mathbb{X}, \mathbb{X}). \quad (9)$$

If we know $u_x(\mathbb{X})$, we have an approximation to $(\mathcal{D}u_x)(\mathbb{X})$ that is wrong with covariance $\gamma^2 E$. The matrix E makes the approach probabilistic (and new).

There are two possible interpretations for D (and E). D appears in a method for solving PDEs $\mathcal{D}u_x = f$ with radial basis functions, called *unsymmetric collocation*,

or *Kansa's method* (Hon and Schaback, 2001; Kansa, 1990; Schaback, 2007). A related method, called symmetric collocation (Fasshauer, 1997, 1999), has been built on and translated into a Bayesian probabilistic numerical method by Cockayne et al. (2017). The following derivation shows that a translation is possible for the probabilistic discretisation, respectively unsymmetric collocation, as well: Eq. (7) explains

$$(\mathcal{D}u_x)(\mathbb{X}) \mid u_x(\mathbb{X}), \xi_x \sim \delta[\mathcal{D}u_x(\mathbb{X}) + \xi_x] \quad (10a)$$

$$= \mathcal{N}(\mathcal{D}u_x(\mathbb{X}), \gamma^2 E). \quad (10b)$$

δ is the Dirac delta. In other words, D and E (respectively D and ξ) estimate $\mathcal{D}u_x$ from values of u_x at a set of grid-points – just like finite difference formulas do. By construction, it is a Bayesian probabilistic numerical algorithm in the technical sense of Cockayne et al. (2019). Supplement B contains a formal proof of this statement.

The limitations of the above probabilistic discretisation are twofold. First, computation of the differentiation matrix D involves inverting the kernel Gram matrix $k_x(\mathbb{X}, \mathbb{X})$. For large point sets, this matrix is notoriously ill-conditioned (Schaback, 1995). Second, computation – and application – of D is expensive because D is a dense matrix with as many rows and columns as there are points in \mathbb{X} . Since loosely speaking, a derivative is a “local” phenomenon (other than e.g. computing integrals), intuition suggests that $\mathcal{D}u_x$ can be approximated more cheaply by exploiting this locality. This thought leads to localised differentiation matrices and probabilistic numerical finite differences.

3 PN FINITE DIFFERENCES

The central idea of an efficient approximation of the probabilistic discretisation is to approximate the derivative of u_x at each point x_n in \mathbb{X} individually. This leads to a row-by-row assembly of D , which can be sparsified naturally and without losing the ability to quantify numerical discretisation/differentiation uncertainty.

Write the matrix D as a vertical stack of $N + 1$ row-vectors, $D = (d_0, \dots, d_N)^\top$, and denote by e_n the n th diagonal element of E , i.e. the variance of the approximation of $\mathcal{D}u_x$ at the n th grid-point in \mathbb{X} . Thus,

$$\mathcal{D}u_x(x_n) \mid u_x(\mathbb{X}), \xi_x \sim \delta[d_n^\top u_x(\mathbb{X}) + \xi_{x,n}] \quad (11)$$

with the scalar random variable $\xi_{x,n} \sim \mathcal{N}(0, \gamma^2 e_n)$. If, in Eq. (11), we replace $u_x(\mathbb{X})$ by the values of u_x at only a local neighbourhood of x_n (a “stencil”),

$$x_{\text{loc}(n)} = \{x_{n-k}, \dots, x_n, \dots, x_{n+k}\}, \quad (12)$$

the general form of Eq. (11) is preserved, but d_n consists of only $2k + 1$ elements, with $k \ll N$,

$$d_n = (\mathcal{D}k_x)(x_n, x_{\text{loc}(n)})k_x(x_{\text{loc}(n)}, x_{\text{loc}(n)})^{-1} \quad (13)$$




Figure 2: *Discretise the Laplacian with a local and global approximation:* The target is the Laplacian $\mathcal{D} = \Delta$ of $u(x) = \sin(\|x\|^2)$ (f). *Left:* Sparsity pattern of the differentiation matrix and error covariance matrix for the localised approximation (a, b) and the global approximation (c, d) on a mesh with $N = 25$ points. The approximation is least certain at the boundaries. *Centre:* The root-mean-square error between Δu and its approximation decreases with an increased stencil size but the approximation breaks down for larger stencils (e), likely due to ill-conditioned kernel Gram matrices. A maximum likelihood estimate of the input scale $r \in \mathbb{R}$ of the square exponential kernel $k(x, y) = e^{-r^2\|x-y\|^2}$ based on data $u_x(x)$ alone does not necessarily lead to well-conditioned system matrices, nor does it generally imply a low RMSE (e). *Right:* Samples from the prior GP u_x for both length scales are shown next to the solution and the target function (f; the colours match the colours in the RMSE plot). *Increasing stencil sizes improves the accuracy until stability concerns arise.*

instead of $N + 1$ elements. The error e_n becomes

$$e_n := (\mathcal{D}^2 k_x)(x_n, x_n) - d_n k_x(x_{\text{loc}(n)}, x_{\text{loc}(n)}) d_n^\top. \quad (14)$$

If the coefficients in the new d_n are all embedded into D according to the indices of $x_{\text{loc}(n)}$ in \mathbb{X} , D becomes sparse and E becomes diagonal. More precisely, $D_{ij} \neq 0$ if and only if $j \in \text{loc}(i)$. On a one-dimensional domain Ω , and with $k = 1$, D is a banded matrix with bandwidth 3. Furthermore, due to the point-by-point construction, choosing $k = N/2$ implies that the original D from Section 2 is recovered; however, E remains diagonal. The sparse approximation resolves many of the performance issues that persist with the global approximation (Fig. 2).

The sparsified differentiation matrix is known as the radial-basis-function-generated finite difference matrix (Driscoll and Fornberg, 2002; Fornberg and Flyer, 2015; Shu et al., 2003; Tolstykh and Shirobokov, 2003). The correspondence to finite-difference formulas stems from the fact that if k_x were a polynomial kernel and the mesh \mathbb{X} were equidistant and one-dimensional, the coefficients in D would equal the standard finite difference coefficients (Fornberg, 1988). The advantage of the more general, kernel-based finite difference approximation over polynomial coefficients is a more robust approximation, especially for non-uniform grid points

and in higher dimensions (Fornberg and Flyer, 2015; Tolstykh and Shirobokov, 2003). The Bayesian point of view does not only add uncertainty quantification in the form of E but also reveals that a practitioner may choose suitable kernels k_x to include prior information into the PDE simulation.

4 PN METHOD OF LINES

The previous two sections have been concerned with approximating the derivative of a function, purely from observations of this function on a grid. Next, we use these strategies to solve time-dependent PDEs. To this end, we combine the probabilistic discretisation with an ODE filter. As opposed to non-probabilistic MOL, this combination quantifies the leak of information between the space discretisation and the ODE solution. Simply put, what is commonly treated as a pipeline of disjoint solvers, becomes more of a single algorithm.

4.1 Spatiotemporal Prior Process

For any function $\varphi = \varphi(t)$, let $\vec{\varphi}$ be the stack of φ and its first ν derivatives, $\vec{\varphi}(t) := (\varphi(t), \dot{\varphi}(t), \dots, \varphi^{(\nu)}(t))$. This abbreviation is convenient because some stochastic processes (like the integrated Wiener or Matérn process) do not have the Markov property, but the stack of

state and derivatives does. The following assumption is integral for the probabilistic method of lines.

Assumption 1. *Assume a Gaussian process prior with separable covariance structure,*

$$u = u(t, x) \sim \mathcal{GP}(0, \gamma^2 k_t \otimes k_x) \quad (15)$$

for some output-scale $\gamma > 0$. $k_t \otimes k_x$ is the product kernel $(k_t \otimes k_x)(t, t', x, x') = k_t(t, t')k_x(x, x')$.

Compared to the traditional method of lines, where the temporal and spatial dimensions are treated independently and with black-box methods, Assumption 1 is mild: even though the covariance is separable, the algorithm still starts with a single, global Gaussian process. Assumption 1 allows choosing temporal kernels that eventually lead to a fast algorithm:

Assumption 2. *Assume k_t is a covariance kernel that allows conversion to a linear, time-invariant stochastic differential equation (SDE) in the following sense: For any $v \sim \mathcal{GP}(0, \gamma^2 k_t)$, \vec{v} solves the SDE*

$$d\vec{v}(t) = A\vec{v}(t) dt + B dw(t) \quad (16)$$

subject to Gaussian initial conditions

$$\vec{v}(t_0) \sim \mathcal{N}(m_0, \gamma^2 C_0), \quad (17)$$

for A , B , m_0 , and C_0 that derive from k_t , and for a one-dimensional Wiener process w with diffusion γ^2 .

Assumption 2 is satisfied, for instance, for the integrated Wiener process or the Matérn process; many more examples are given in Chapter 12 of the book by Särkkä and Solin (2019). Assumptions 1 and 2 unlock the machinery of probabilistic ODE solvers.

Next, we add spatial correlations into the prior SDE model. The following Lemma 3 will simplify the derivations further below (Soln, 2016).

Lemma 3. *Let k_t be a covariance function that satisfies Assumption 2. Let $M \in \mathbb{R}^{q \times q}$ be a matrix. For a process $\varphi \sim \mathcal{GP}(0, \gamma^2 k_t \otimes M)$, $\vec{\varphi}$ solves*

$$d\vec{\varphi}(t) = (A \otimes I)\vec{\varphi}(t) dt + (B \otimes I) d\hat{w}(t), \quad (18)$$

subject to Gaussian initial conditions

$$\vec{\varphi}(t_0) \sim \mathcal{N}(m_0 \otimes \mathbf{1}, \gamma^2 C_0 \otimes M), \quad (19)$$

where \hat{w} is a q -dimensional Wiener process with constant diffusion $\gamma^2 M$, and $I \in \mathbb{R}^{q \times q}$ is the identity. $m_0 \otimes \mathbf{1}$ is a vertical stack of d copies of m_0 .

Lemma 3 suggests that a spatiotemporal prior (Assumption 1) may be restricted to a spatial grid without losing the computational benefits that SDE priors provide. Abbreviate $U(t) := u(t, \mathbb{X})$. Lemma 3 gives

rise to an SDE representation for $U(t)$ by choosing $M = k_x(\mathbb{X}, \mathbb{X})$. The same holds for the error model: Recall the definition of the error covariance E from Eq. (9) (or Eq. (14) respectively, if the localised version is used). $\xi \sim \mathcal{GP}(0, \gamma^2 k_t \otimes E)$ admits a state-space formulation due to Lemma 3. Through

$$Du(t, \mathbb{X}) + \xi(t) \quad (20a)$$

$$\sim \mathcal{GP}(0, \gamma^2 k_t \otimes [Dk_x(\mathbb{X}, \mathbb{X})D^\top + E]) \quad (20b)$$

$$= \mathcal{GP}(0, \gamma^2 k_t \otimes (DD^*k_x(\mathbb{X}, \mathbb{X}))) \quad (20c)$$

it is evident that ξ is an appropriate prior model for the time-evolution of the spatial discretisation error. This error being part of the probabilistic model is the advantage of PNMOL over non-probabilistic versions.

4.2 Information Model

The priors over U and ξ become a probabilistic numerical PDE solution by conditioning U and ξ on “solving the PDE” at a number of grid points as follows. Recall from Eq. (8) how $Du(t, \mathbb{X})$ is approximated by $Du(t, \mathbb{X}) + \xi$. The residual process

$$r(t) := \dot{U}(t) - F(t, \mathbb{X}, U(t), DU(t) + \xi) \quad (21)$$

measures how well realisations of \vec{U} and $\vec{\xi}$ solve the PDE. \dot{U} and U are components in \vec{U} , therefore all operations are tractable given the extended state vectors \vec{U} and $\vec{\xi}$. Conditioning \vec{U} and $\vec{\xi}$ on $r(t) = 0$ for all t yields a probabilistic PDE solution. However, in practice, we need to discretise the time variable first in order to be able to compute the (approximate) posterior.

4.3 Time-Discretisation

Let $\mathbb{T} := (t_0, \dots, t_K)$ be a grid on $[t_0, t_{\max}]$. Define the time-steps $h_k := t_{k+1} - t_k$. Restricted to the time grid,

$$\vec{U}(t_{k+1}) | \vec{U}(t_k) \sim \mathcal{N}(\Phi(h_k)\vec{U}(t_k), \gamma^2 \Sigma_U(h_k)) \quad (22a)$$

$$\vec{\xi}(t_{k+1}) | \vec{\xi}(t_k) \sim \mathcal{N}(\Phi(h_k)\vec{\xi}(t_k), \gamma^2 \Sigma_\xi(h_k)) \quad (22b)$$

with transition matrix

$$\Phi(h_k) := \check{\Phi}(h_k) \otimes I := \exp(Ah_k) \otimes I \quad (23a)$$

and process noise covariance matrices

$$\Sigma_U(h_k) := \check{\Sigma}(h_k) \otimes k_x(\mathbb{X}, \mathbb{X}) \quad (24a)$$

$$\Sigma_\xi(h_k) := \check{\Sigma}(h_k) \otimes E \quad (24b)$$

$$\check{\Sigma}(h_k) := \int_0^{h_k} \check{\Phi}(h_k - \tau) B B^\top \check{\Phi}(h_k - \tau)^\top d\tau. \quad (24c)$$

$\check{\Phi}$ and $\check{\Sigma}$ can be computed efficiently with matrix fractions (Särkkä and Solin, 2019). Integrated Wiener




Figure 3: *Graphical visualisation of PNMOL*: The states $\vec{\xi}$ and \vec{U} are conditionally independent given $\theta = (k_x, k_t, \mathbb{X}, \mathcal{D}, \gamma, \mathbb{T})$. The existence of the bottom row distinguishes PNMOL from other PDE solvers. The dependencies between the $\xi(t_k)$, indicated by the dashed lines, are optional; removing them improves the efficiency of the algorithm; details are in Section 4.5.

process priors, their time-discretisations, and practical considerations for implementation of high-orders are discussed in the probabilistic ODE solver literature (e.g. Krämer and Hennig, 2020; Tronarp et al., 2021). On the discrete-time grid, the information model reads

$$r(t_k) \mid \vec{U}(t_k), \vec{\xi}(t_k) \sim \delta[\dot{U}(t_k) - F(t, \mathbb{X}, U(t_k), DU(t_k) + \xi(t_k))]. \quad (25)$$

The complete setup is depicted in Fig. 3.

4.4 Inference

Altogether, the probabilistic method of lines targets

$$p\left(\vec{U}, \vec{\xi} \mid [r(t_k) = 0]_{k=0}^K\right) \quad (26)$$

which we will refer to as the *probabilistic numerical PDE solution* of Eq. (1). The precise form of the PN PDE solution is intractable because F (and therefore the likelihood in Eq. (25)) are nonlinear. It can be approximated efficiently with techniques from extended Kalman filtering (Särkkä and Solin, 2019), which is based on approximating the nonlinear PDE vector field F with a first-order Taylor series. Inference in the linearised model is feasible with Kalman filtering and smoothing. Särkkä and Solin (2019) summarise the details, and Tronarp et al. (2019) explain specific considerations for ODE solvers.

The role of ξ and its impact on the posterior distribution (i.e. the PDE solution) is comparable to that of a latent force in the ODE (Alvarez et al., 2009; Hartikainen et al., 2012; Schmidt et al., 2021). The coarser \mathbb{X} is, the larger is E , respectively ξ . A large ξ indicates how increasingly severe the misspecification of the ODE becomes. Unlike traditional PDE solvers, inference according to the information model in Eq. (25)

embraces latent discretisation errors. It does so by an algorithm that shares similarities with the latent force inference algorithm by Schmidt et al. (2021) (compare Fig. 3 to Schmidt et al. (2021, Fig. 3)), but the sources of misspecification are different in both works.

Inference of ξ is expensive because the complexity of a Gaussian filter scales as $O(Kd^3)$ where d is the dimension of the state-space model. In the present case, $d = 2(N+1)(\nu+1)$ because for $N+1$ spatial grid points, PNMOL tracks ν time-derivatives of U and ξ . Both U and ξ are $(N+1)$ -dimensional. If the only purpose of ξ is to incorporate a measure of spatial discretisation uncertainty into the information model of the PDE solver, tracking ξ in the state space is not required, as long as we introduce another approximation.

4.5 White-Noise Approximation

Recall $\xi \sim \mathcal{GP}(0, \gamma^2 k_t \otimes E)$, i.e. ξ is an integrated Wiener process in time and a Gaussian $\mathcal{N}(0, \gamma^2 E)$ in space. One may relax the temporal integrated Wiener process prior to a white noise process prior; that is, $\xi(t)$ is independent of $\xi(s)$ for $s \neq t$, and $\xi(t) \sim \mathcal{N}(0, \gamma^2 E)$ for all t . The state-space realisation of a white noise process is trivial because there is no temporal evolution (recall the dashed lines in Fig. 3). To understand the impact of E in the white-noise formulation on the information model, consider the semi-linear PDE

$$F_{\text{semi}}(t, x, u, Du) := Du(t, x) + f(u(t, x)) \quad (27)$$

for some nonlinear f . Eq. (25) becomes

$$r_{\text{white}}(t_k) \mid \vec{U}(t_k) \sim \mathcal{N}(\rho(\vec{U}(t_k)), \gamma^2 E) \quad (28a)$$

$$\rho(\vec{U}(t)) := \dot{U}(t) - DU(t) - f(U(t)) \quad (28b)$$

The Dirac likelihood in Eq. (25) becomes a Gaussian with the measurement noise E . Versions of Eq. (28) for other nonlinearities are in Supplement D.

The advantage of the white-noise approximation over the latent-force version is that the state-space model is precisely half the size because ξ is not a state variable anymore. Since the complexity of PNMOL depends cubically on the dimension of the state-space, half as many state variables improve the complexity of the algorithm by a factor $2^3 = 8$. Arguably, $\gamma^2 E$ entering the information model as a measurement covariance may also provide a more intuitive explanation of the impact that the statistical quantification of the discretisation error has on the PDE solution: the larger E , the less strictly is a zero PDE residual enforced during inference. The error covariance E regularises the impact of the information model, and, in the limit of $E \rightarrow \infty$, yields the prior as a PDE solution. Intuitively put, for $E > 0$, the algorithm puts less trust in the residual information $r(\cdot)$ than for $E = 0$.

5 HYPERPARAMETERS

Kernels As common in the literature on probabilistic ODE solvers, we use temporal integrated Wiener process priors (e.g. Bosch et al., 2021; Tronarp et al., 2019). In the experiments, the order of integration is $\nu \in \{1, 2\}$. Using low order ODE solvers is not unusual for MOL implementations (Cash and Psihoyios, 1996). Spatial kernels need to be sufficiently differentiable to admit the formula in Eq. (9). In this work, we use squared exponential kernels. Choosing their input scale is not straightforward (recall Fig. 2; we found $r = 0.25$ to work well across experiments). Other sufficiently regular kernels, e.g. rational quadratic or Matérn kernels, would work as well. Polynomial kernels recover traditional finite difference weights (Fornberg, 1988), but like for any other feature-based kernel, $k(x, y) = \Phi(x)^\top \Phi(y)$ holds, thus both summands in Eq. (9) (and in Eq. (14)) cancel out. The discretisation uncertainty E would be zero. In this case, PNMOL gains similarity to the traditional method of lines combined with a probabilistic ODE solver. (The boundary conditions would be treated slightly differently; we refer to Supplement A.)

Spatial Grid The spatial grid can be any set of freely scattered points. Spatial neighbourhoods can, for example, be queried from a KD tree (Bentley and Ottmann, 1979). For the simulations in the present paper, we use equispaced grids. PNMOL’s requirements on the grid differ from, for instance, finite element methods, in that there needs to be no notion of connectivity or triangulation between the grid points. PN finite differences only require stencils; in light of the stability results in Fig. 2, we choose them maximally small.

Output-Scale The output-scale γ calibrates the width of the posterior and can be tuned with quasi-maximum likelihood estimation. Omitting the boundary conditions, this means (recall r from Eq. (25))

$$(\hat{\gamma})^2 := \frac{1}{(N+1)(K+1)} \sum_{k=0}^K \|\mathbb{E}[r(t_k)]\|_{\mathbb{C}[r(t_k)]^{-1}}^2, \quad (29)$$

where the mean $\mathbb{E}[r(t_k)]$ and the covariance between the output-dimensions of $r(\cdot)$ at time t_k , $\mathbb{C}[r(t_k)]$, emerge from the same Gaussian approximation that computes the approximate posterior (Supplement C). Eq. (29) uses the Mahalanobis norm $\|x\|_A^2 = x^\top Ax$.

6 RELATED WORK

Connections to non-probabilistic numerical approximation have been discussed in Sections 2 and 3. Recall from there that the strongest connections are to unsymmetric collocation (Hon and Schaback, 2001; Kansa,

1990; Schaback, 2007), radial-basis-function-generated finite differences (Driscoll and Fornberg, 2002; Fornberg and Flyer, 2015; Shu et al., 2003; Tolstykh and Shirobokov, 2003), and collocation methods in general. We refer to Fasshauer (2007); Fornberg and Flyer (2015) for a more comprehensive overview. The literature on the method of lines is covered by, e.g., Schiesser (2012). Dereli and Schaback (2013); Hon et al. (2014) combine collocation with MOL. None of the above exploits the correlations between spatial and temporal errors. The significance of estimating the interplay of both error sources for MOL has been recognised by Berzins (1988); Berzins et al. (1991); Lawson et al. (1991).

Cockayne et al. (2017); Owhadi (2015, 2017); Raissi et al. (2017, 2018) describe a probabilistic solver for PDEs relating to symmetric collocation approaches from numerical analysis. Chen et al. (2021) extend the ideas to non-linear PDEs. Wang et al. (2021) continue the work of Chkrebti et al. (2016) in constructing an ODE/PDE initial value problem solver that uses (approximate) conjugate Gaussian updating at each time-step. Duffin et al. (2021) solve time-dependent PDEs by discretising the spatial domain with finite elements, and applying ensemble and extended Kalman filtering in time. They build on the paper by Girolami et al. (2021). Abdulle and Garegnani (2021); Conrad et al. (2017) compute probabilistic PDE solutions by randomly perturbing non-probabilistic solvers. All of the above discard the uncertainty associated with discretising $\mathcal{D} \approx D$. Some papers achieve ODE-solver-like complexity for time-dependent problems (Chkrebti et al., 2016; Duffin et al., 2021; Wang et al., 2021), while others compute a continuous-time posterior (Chen et al., 2021; Cockayne et al., 2017; Owhadi, 2015, 2017; Raissi et al., 2017, 2018). PNMOL does both.

The efficiency of the PDE filter builds on recent work on filtering-based probabilistic ODE solvers (Bosch et al., 2021; Kersting and Hennig, 2016; Kersting et al., 2020b; Krämer and Hennig, 2020; Schober et al., 2019; Tronarp et al., 2019, 2021) and their applications (Kersting et al., 2020a; Schmidt et al., 2021). Frank and Enßlin (2020) apply an ODE filter to solve discretised PDEs. Similar algorithms have been developed for other types of ODEs (Hennig and Hauberg, 2014; John et al., 2019; Krämer and Hennig, 2021). The papers by Abdulle and Garegnani (2020); Chkrebti et al. (2016); Conrad et al. (2017), prominently feature ODEs.

7 EXPERIMENTS

Code The code for the experiments³, and an implementation of the PN discretisation and PN finite dif-

³<https://github.com/schmidtjonathan/pnmol-experiments>

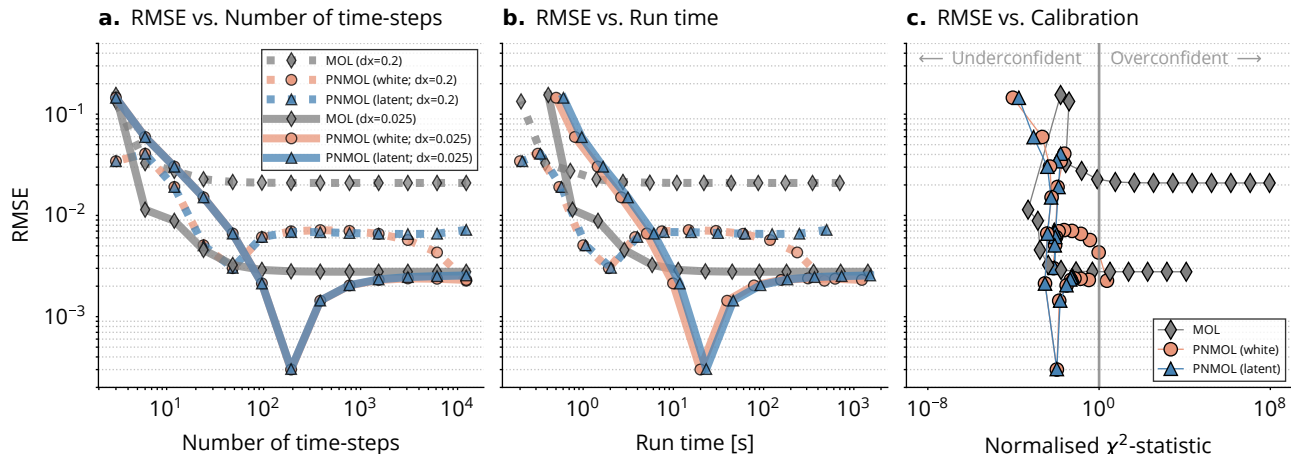


Figure 4: *Quantify the global error*: Work vs. precision vs. calibration of PNMOL in the latent-force version (blue) and the white-noise version (orange), compared to a traditional PN ODE solver combined with conventional MOL (grey), on the spatial Lotka-Volterra model. Two kinds of curves are shown: one for a coarse (dotted), and one for a fine spatial mesh (solid). A reference is computed by discretising the spatial domain with a ten times finer mesh and solving the ODE with backward differentiation formulas. The RMSE of both methods stagnates once a certain accuracy is reached, but PNMOL appears to reach a slightly lower RMSE for $dx = 0.2$ (left, middle; perhaps due to different treatment of boundary conditions; Supplement A). The run time of PNMOL-white is comparable to that of MOL, and the run time of PNMOL-latent is slightly longer (middle). The calibration of PNMOL, measured in the normalised χ^2 -statistic of the Gaussian posterior (so that the “optimum” is 1, not d), remains close to 1 but is slightly underconfident. With decreasing time-steps, MOL is poorly calibrated.

ferences described in Sections 2 and 3⁴ are on GitHub.

Quantify The Global Error We investigate how the PN method of lines impacts numerical uncertainty quantification. As a first experiment, we solve a spatial Lotka-Volterra model (Holmes et al., 1994), i.e. nonlinear predator-prey dynamics with spatial diffusion, on a range of temporal and spatial resolutions. From the results in Fig. 4, it is evident how the spatial accuracy limits the overall accuracy. But also how traditional ODE filters combined with MOL fail to quantify numerical uncertainty reliably. At any parameter configuration, it is *either* the spatial or the temporal discretisation that dominates the error. Decreasing the time-step alone lets the error not only stagnate but worsens the calibration because the ODE solver does not know how bad the spatial approximation is.

Which Error Dominates? To further examine which one of *either* Δx or Δt dominates the approximation, we consider a second example: a spatial SIR model (Gai et al., 2020). We investigate more formally how increasing either, the time-resolution vs. the space-resolution, leads to a low overall error. The results are in Fig. 5, and confirm the findings from Fig. 4 above. Traditional PN ODE solvers with conventional MOL

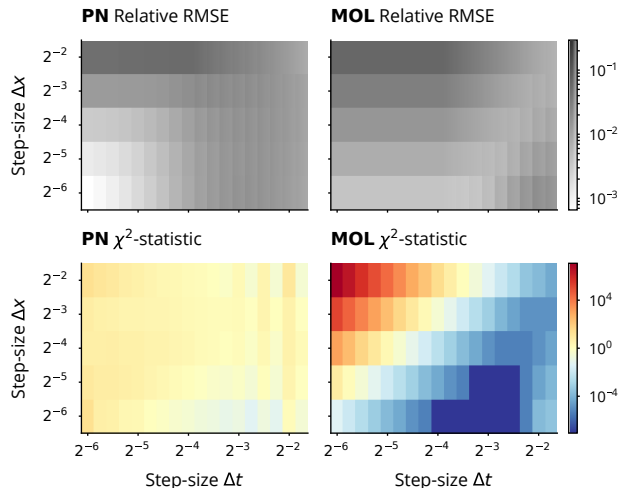


Figure 5: *Which error dominates?* The relative RMSE is only small if *both* Δt and Δx are small, which affects all solvers: PNMOL (top left) as well as traditional MOL combined with PN ODE solvers (top right). PN w/ MOL is severely overconfident for large Δx and small Δt (bottom right), while PNMOL delivers a calibrated posterior distribution (bottom left).

are unaware of the true, global approximation error. PNMOL is not, despite being equally accurate.

⁴<https://github.com/pnkraemer/probfindiff>

8 CONCLUSION

We presented probabilistic strategies for discretising PDEs, and for making use of the resulting quantification of spatial discretisation uncertainty in a probabilistic ODE solver. We discussed practical considerations, including sparsification of the differentiation matrices, keeping the dimensionality of the state-space low, and hyperparameter choices. Altogether, and unlike traditional PDE solvers, the probabilistic method of lines unlocks quantification of spatiotemporal correlations in an approximate PDE solution, all while preserving the efficiency of adaptive ODE solvers. This makes it a valuable algorithm in the toolboxes of probabilistic programs and differential equation solvers and may serve as a backbone for latent force models, inverse problems, and differential-equation-centric machine learning.

Acknowledgements

The authors gratefully acknowledge financial support by the German Federal Ministry of Education and Research (BMBF) through Project ADIMEM (FKZ 01IS18052B). They also 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. Moreover, the authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting Nicholas Krämer.

The authors thank Nathanael Bosch for helpful feedback on the manuscript.

References

- Abdulle, A. and Garegnani, G. (2020). Random time step probabilistic methods for uncertainty quantification in chaotic and geometric numerical integration. *Statistics and Computing*, 30(4):907–932.
- Abdulle, A. and Garegnani, G. (2021). A probabilistic finite element method based on random meshes: A posteriori error estimators and Bayesian inverse problems. *Computer Methods in Applied Mechanics and Engineering*, 384:113961.
- Alvarez, M., Luengo, D., and Lawrence, N. D. (2009). Latent force models. In *AISTATS 2009*.
- Bentley, J. L. and Ottmann, T. A. (1979). Algorithms for reporting and counting geometric intersections. *IEEE Transactions on Computers*, 28(09):643–647.
- Berzins, M. (1988). Global error estimation in the method of lines for parabolic equations. *SIAM Journal on Scientific and Statistical Computing*, 9(4):687–703.
- Berzins, M., Baehmann, P., Flaherty, J., and Lawson, J. (1991). Towards an automated finite element solver for time-dependent fluid-flow problems. *The Mathematics of Finite Elements and Application*, 7:181–188.
- Bosch, N., Hennig, P., and Tronarp, F. (2021). Calibrated adaptive probabilistic ODE solvers. In *AISTATS 2021*.
- Cash, J. and Pshoyios, Y. (1996). The MOL solution of time dependent partial differential equations. *Computers & Mathematics with Applications*, 31(11):69–78.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. (2018). Neural ordinary differential equations. In *NeurIPS 2018*.
- Chen, Y., Hosseini, B., Owhadi, H., and Stuart, A. M. (2021). Solving and learning nonlinear PDEs with Gaussian processes. *Journal of Computational Physics*, 447.
- Chkrebtii, O. A., Campbell, D. A., Calderhead, B., and Girolami, M. A. (2016). Bayesian solution uncertainty quantification for differential equations. *Bayesian Analysis*, 11(4):1239–1267.
- Cockayne, J., Oates, C., Sullivan, T., and Girolami, M. (2017). Probabilistic numerical methods for PDE-constrained Bayesian inverse problems. In *AIP Conference Proceedings*, volume 1853.
- Cockayne, J., Oates, C. J., Sullivan, T. J., and Girolami, M. (2019). Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789.
- Conrad, P. R., Girolami, M., Särkkä, S., Stuart, A., and Zygalakis, K. (2017). Statistical analysis of differential equations: introducing probability measures on numerical solutions. *Statistics and Computing*, 27(4):1065–1082.
- Dereli, Y. and Schaback, R. (2013). The meshless kernel-based method of lines for solving the equal width equation. *Applied Mathematics and Computation*, 219(10):5224–5232.
- Driscoll, T. A. and Fornberg, B. (2002). Interpolation in the limit of increasingly flat radial basis functions. *Computers & Mathematics with Applications*, 43(3-5):413–422.
- Duffin, C., Cripps, E., Stemler, T., and Girolami, M. (2021). Statistical finite elements for misspecified models. *Proceedings of the National Academy of Sciences*, 118(2).

- Evans, L. C. (2010). *Partial Differential Equations*, volume 19. American Mathematical Society.
- Fasshauer, G. E. (1997). Solving partial differential equations by collocation with radial basis functions. In *Surface Fitting and Multiresolution Methods*, pages 131–138. University Press.
- Fasshauer, G. E. (1999). Solving differential equations with radial basis functions: multilevel methods and smoothing. *Advances in Computational Mathematics*, 11(2):139–159.
- Fasshauer, G. E. (2007). *Meshfree Approximation Methods with MATLAB*. World Scientific.
- Fornberg, B. (1988). Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of Computation*, 51(184):699–706.
- Fornberg, B. and Flyer, N. (2015). *A Primer on Radial Basis Functions with Applications to the Geosciences*. SIAM.
- Frank, P. and Enßlin, T. A. (2020). Probabilistic simulation of partial differential equations. *arXiv preprint arXiv:2010.06583*.
- Gai, C., Iron, D., and Kolokolnikov, T. (2020). Localized outbreaks in an S-I-R model with diffusion. *Journal of Mathematical Biology*, 80:1389–1411.
- Gelbrecht, M., Boers, N., and Kurths, J. (2021). Neural partial differential equations for chaotic systems. *New Journal of Physics*, 23(4):043005.
- Girolami, M., Febrianto, E., Yin, G., and Cirak, F. (2021). The statistical finite element method (stat-FEM) for coherent synthesis of observation data and model predictions. *Computer Methods in Applied Mechanics and Engineering*, 375.
- Hartikainen, J., Seppänen, M., and Särkkä, S. (2012). State-space inference for non-linear latent force models with application to satellite orbit prediction. *ICML 2012*.
- Hennig, P. and Hauberg, S. (2014). Probabilistic solutions to differential equations and their application to Riemannian statistics. In *AISTATS 2014*.
- Holmes, E., Lewis, M., Banks, J., and Veit, D. (1994). Partial differential equations in ecology: Spatial interactions and population dynamics. *Ecology*, 75:17–29.
- Hon, Y. and Schaback, R. (2001). On unsymmetric collocation by radial basis functions. *Applied Mathematics and Computation*, 119(2-3):177–186.
- Hon, Y., Schaback, R., and Zhong, M. (2014). The meshless kernel-based method of lines for parabolic equations. *Computers & Mathematics with Applications*, 68(12):2057–2067.
- John, D., Heuveline, V., and Schober, M. (2019). GOODE: A Gaussian off-the-shelf ordinary differential equation solver. In *ICML 2019*.
- Kansa, E. J. (1990). Multiquadrics—a scattered data approximation scheme with applications to computational fluid-dynamics—II solutions to parabolic, hyperbolic and elliptic partial differential equations. *Computers & Mathematics with Applications*, 19(8-9):147–161.
- Kersting, H. and Hennig, P. (2016). Active uncertainty calibration in Bayesian ODE solvers. *UAI 2016*.
- Kersting, H., Krämer, N., Schiegg, M., Daniel, C., Tiemann, M., and Hennig, P. (2020a). Differentiable likelihoods for fast inversion of ‘likelihood-free’ dynamical systems. In *ICML 2020*.
- Kersting, H., Sullivan, T. J., and Hennig, P. (2020b). Convergence rates of Gaussian ODE filters. *Statistics and Computing*, 30(6):1791–1816.
- Krämer, N. and Hennig, P. (2020). Stable implementation of probabilistic ODE solvers. *arXiv preprint arXiv:2012.10106*.
- Krämer, N. and Hennig, P. (2021). Linear-time probabilistic solutions of boundary value problems. In *NeurIPS 2021*.
- Lawson, J., Berzins, M., and Dew, P. M. (1991). Balancing space and time errors in the method of lines for parabolic equations. *SIAM Journal on Scientific and Statistical Computing*, 12(3):573–594.
- Li, Z., Kovachki, N. B., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. (2021). Fourier neural operator for parametric partial differential equations. In *ICLR 2021*.
- Lu, L., Jin, P., Pang, G., Zhang, Z., and Karniadakis, G. E. (2021). Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229.
- Owhadi, H. (2015). Bayesian numerical homogenization. *Multiscale Modeling & Simulation*, 13(3):812–828.
- Owhadi, H. (2017). Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games. *SIAM Review*, 59(1):99–149.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2017). Machine learning of linear differential equations using Gaussian processes. *Journal of Computational Physics*, 348:683–693.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2018). Numerical Gaussian processes for time-dependent and nonlinear partial differential equations. *SIAM Journal on Scientific Computing*, 40(1):A172–A198.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2019). 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.

Särkkä, S. and Solin, A. (2019). *Applied Stochastic Differential Equations*. Cambridge University Press.

Schaback, R. (1995). Error estimates and condition numbers for radial basis function interpolation. *Advances in Computational Mathematics*, 3(3):251–264.

Schaback, R. (2007). Convergence of unsymmetric kernel-based meshless collocation methods. *SIAM Journal on Numerical Analysis*, 45(1):333–351.

Schiesser, W. E. (2012). *The Numerical Method of Lines: Integration of Partial Differential Equations*. Elsevier.

Schmidt, J., Krämer, N., and Hennig, P. (2021). A probabilistic state space model for joint inference from differential equations and data.

Schober, M., Särkkä, S., and Hennig, P. (2019). A probabilistic model for the numerical solution of initial value problems. *Statistics and Computing*, 29(1):99–122.

Schweppe, F. (1965). Evaluation of likelihood functions for Gaussian signals. *IEEE Transactions on Information Theory*, 11(1):61–70.

Shu, C., Ding, H., and Yeo, K. (2003). Local radial basis function-based differential quadrature method and its application to solve two-dimensional incompressible Navier–Stokes equations. *Computer methods in applied mechanics and engineering*, 192(7-8):941–954.

Solin, A. (2016). *Stochastic Differential Equation Methods for Spatio-Temporal Gaussian Process Regression*. Doctoral thesis, School of Science.

Tolstykh, A. and Shirobokov, D. (2003). On using radial basis functions in a “finite difference mode” with applications to elasticity problems. *Computational Mechanics*, 33(1):68–79.

Tronarp, F., Kersting, H., Särkkä, S., and Hennig, P. (2019). Probabilistic solutions to ordinary differential equations as non-linear Bayesian filtering: A new perspective. *Statistics and Computing*, 29(6):1297–1315.

Tronarp, F., Särkkä, S., and Hennig, P. (2021). Bayesian ODE solvers: The maximum a posteriori estimate. *Statistics and Computing*, 31(3):1–18.

Wang, J., Cockayne, J., Chkrebtti, O., Sullivan, T. J., and Oates, C. J. (2021). Bayesian numerical methods for nonlinear partial differential equations. *Statistics and Computing*, 31.

Supplementary Material: Probabilistic Numerical Method of Lines for Time-Dependent Partial Differential Equations

This supplement explains details regarding: (i) treatment of boundary conditions (Supplement A); (ii) the probabilistic numerical discretisation from Section 2 being a Bayesian probabilistic numerical method (Supplement B); (iii) quasi-maximum-likelihood-estimation of the output-scale γ (Supplement C); and (iv) inference schemes for fully non-linear systems of partial differential equations (Supplement D).

Recall the abbreviations from the main paper: partial differential equation (PDE), ordinary differential equation (ODE), probabilistic numerics (PN), method of lines (MOL), probabilistic numerical method of lines (PNMOL).

A BOUNDARY CONDITIONS

A.1 Setup

In this section, we explain how PNMOL treats boundary conditions. Recall the differential operator \mathcal{D} from Eq. (1), the differentiation matrix D and the error covariance matrix E from Eq. (9), respectively Eqs. (13) and (14). Also recall from Section 4.4 that in the ODE solve, each pair of states $(\vec{U}, \vec{\xi})$ is conditioned on a zero PDE residual

$$r(t) := \dot{U}(t) - F(t, \mathbb{X}, U(t), DU(t) + \xi) \stackrel{!}{=} 0 \quad (30)$$

where U , \dot{U} , and ξ are components of the extended state vectors \vec{U} and ξ . The remainder of this section extends this framework to include boundary conditions $\mathcal{B}u(t, x) = g(x)$ for all $x \in \partial\Omega$ and some function g (Eq. (1)).

A.2 Discretised Boundary Conditions

To augment the PDE residual information with boundary conditions, we begin by discretising \mathcal{B} probabilistically; either with global collocation as in Section 2 or with PN finite differences as in Section 3. Below, we show the former, because the latter becomes accessible with the same modifications from Section 3. Let $u_x \sim \mathcal{GP}(0, \gamma^2 k_x)$. Defining the differentiation matrix B and the error covariance R ,

$$B := (\mathcal{B}k_x)(\mathbb{X}, \mathbb{X})k_x(\mathbb{X}, \mathbb{X})^{-1}, \quad R := (\mathcal{B}^2 k_x)(\mathbb{X}, \mathbb{X}) - Bk_x(\mathbb{X}, \mathbb{X})B^\top, \quad (31)$$

we have access to boundary conditions: For any $y \in \mathbb{R}^{N+1}$,

$$p((\mathcal{B}u_x)(\mathbb{X}) \mid u_x(\mathbb{X}) = y) \sim \mathcal{N}(By, \gamma^2 R) \quad (32)$$

holds. For Dirichlet boundary conditions, \mathcal{B} is the identity, thus $B = I_{N+1}$ is the identity matrix and $R \equiv 0$ is the zero matrix. For Neumann conditions, B is the discretised derivative along normal coordinates, and R is generally nonzero (recall the explanation in Section 5 of the cases in which the error matrix is zero). The derivation surrounding Eq. (32) above suggests how the present framework would deal with boundary conditions that are subject to additive Gaussian noise.

A.3 Latent Force

Define the latent force $\vartheta = \vartheta(t) \sim \mathcal{GP}(0, \gamma^2 k_t \otimes R)$, which will play a role similar to ξ but for the boundary conditions. ϑ inherits a stochastic differential equation formulation from k_t just like ξ does (recall Lemma 3). Denote the stack of ϑ and its first $\nu \in \mathbb{N}$ time-derivatives by $\vec{\vartheta}$. Let $\mathbb{X}_B \subset \mathbb{X}$ be the subset of boundary points in

\mathbb{X} . The full information operator (i.e. an extended version of Eq. (30)) includes $\mathcal{B}u(t, x)|_{\partial\Omega} = g(x)$ as

$$r_{\text{full}}(t) := \begin{pmatrix} \dot{U}(t) - F(t, \mathbb{X}, U(t), DU(t) + \xi) \\ \mathcal{B}U(t) - g(\mathbb{X}_B) - \vartheta(t) \end{pmatrix} \stackrel{!}{=} 0. \quad (33)$$

r_{full} depends on \vec{U} , $\vec{\xi}$, and $\vec{\vartheta}$. Conditioning $(\vec{U}, \vec{\xi}, \vec{\vartheta})$ on $r_{\text{full}}(t) = 0$ yields a probabilistic PDE solution that knows boundary conditions. Notably, the bottom row in r_{full} is linear in $(\vec{U}, \vec{\xi}, \vec{\vartheta})$ so there is no linearisation required to enable (approximate) inference.

A.4 Comparison To MOL

Boundary conditions in PNMOL enter through the information operator, i.e. on the same level as the PDE vector field F , which is different to conventional MOL: In MOL, one only tracks the state variables in the interior of Ω , i.e. $u(t, \mathbb{X} \setminus \mathbb{X}_B)$ because boundary conditions can be inferred from the interior straightforwardly. For Dirichlet conditions, the boundary values are always dictated by $g(\mathbb{X}_B)$. Let $\frac{\partial}{\partial n}$ be the directional derivative taken in the direction normal to the boundary $\partial\Omega$. For Neumann conditions, for some small $\lambda > 0$, we can approximate

$$\frac{\partial}{\partial n} u(t, x) = \frac{u(t, x) - u(t, x - \lambda n)}{\lambda} \stackrel{!}{=} g(x). \quad (34)$$

This viewpoint suggests $u(t, x) = \lambda g(x) + u(t, x - \lambda n)$ and λ can be chosen such that $x - \lambda n$ is the nearest neighbor of $x \in \mathbb{X}$. While tracking only the state values in the interior of the domain has the advantage that the ODE system emerging from the method of lines is smaller than for PNMOL (which perhaps explains why in Fig. 4, PNMOL achieves a lower error than MOL), MOL has two disadvantages: (i) non-deterministic boundary conditions are not straightforward to include; (ii) the finite difference approximation of Neumann conditions introduces errors. PNMOL does not face the first issue and quantifies the error mentioned by the second issue.

B DISCRETISATION AS A BAYESIAN PROBABILISTIC NUMERICAL METHOD

Cockayne et al. (2019) introduce a formal definition of Bayesian and non-Bayesian probabilistic numerical methods. More specifically, they define a probabilistic numerical method to be a tuple of an *information operator* and a *belief update operator*. A PN method then becomes Bayesian if its output is the pushforward of a specific conditional distribution through the *quantity of interest*. All of those objects can be derived for the PN discretisation of \mathcal{D} . The same is true for the boundary conditions explained in Supplement A but is omitted in the following. The proof of the statement below mirrors the explanation why Bayesian quadrature is a Bayesian PN method in Section 2.2 of the paper by Cockayne et al. (2019).

Proposition 4. *The approximation of $(\mathcal{D}u_x)(\mathbb{X})$ with D and E as in Eqs. (7) and (9), using evaluations of u_x at a grid, is a Bayesian probabilistic numerical method.*

Proof. The prior measure is the GP prior $\mathcal{GP}(0, \gamma^2 k_x)$ and defined over some separable Banach space of sufficiently differentiable, real-valued functions. The information operator $\mathcal{I}[\varphi] := \varphi(\mathbb{X})$ evaluates a function φ at the grid \mathbb{X} . The belief update restricts the prior measure to the set of functions that interpolate $\varphi(\mathbb{X})$ (this is standard Gaussian process conditioning). The quantity of interest is the derivative $\mathcal{D}u_x$, which results in the posterior distribution in Eq. (5). \square

C QUASI-MAXIMUM-LIKELIHOOD ESTIMATION OF THE OUTPUT SCALE

Proving the validity of the estimator in Eq. (29) parallels similar statements for similar settings (Bosch et al., 2021; Krämer and Hennig, 2021; Tronarp et al., 2019) and consists of two phases: (i) showing that the posterior covariances are of the form $C_k = \gamma^2 \check{C}_k$ for some \check{C}_k , $k = 0, \dots, K$; (ii) deriving the maximum likelihood estimators.

To show the first claim, recall the discrete-time transition of U and ξ from Eq. (22) and the information model from Eq. (25). Then, the fully discretised state-space model is

$$\vec{U}(t_0) \sim \mathcal{N}(m_0 \otimes \mathbf{1}, \gamma^2 C_0 \otimes k_x(\mathbb{X}, \mathbb{X})) \quad (35a)$$

$$\vec{\xi}(t_0) \sim \mathcal{N}(m_0 \otimes \mathbf{1}, \gamma^2 C_0 \otimes E) \quad (35b)$$

$$\vec{U}(t_{k+1}) | \vec{U}(t_k) \sim \mathcal{N}(\Phi(h_k)\vec{U}(t_k), \gamma^2 \Sigma_U(h_k)) \quad (35c)$$

$$\vec{\xi}(t_{k+1}) | \vec{\xi}(t_k) \sim \mathcal{N}(\Phi(h_k)\vec{\xi}(t_k), \gamma^2 \Sigma_\xi(h_k)) \quad (35d)$$

$$r(t_k) | \vec{U}(t_k), \vec{\xi}(t_k) \sim \delta \left(\dot{U}(t_k) - F(t, \mathbb{X}, U(t_k), DU(t_k) + \xi(t_k)) \right) \quad (35e)$$

All transitions have a process noise that depends multiplicatively on γ^2 . The Dirac likelihood is noise-free. Therefore, the posterior covariance depends multiplicatively on γ as well (which can be proved by an induction that is identical to the one in Appendix C of the paper by Tronarp et al. (2019)).

To show the second claim, let a linearised observation model be given by

$$r(t_k) | \vec{U}(t_k), \vec{\xi}(t_k) \sim \delta \left(H \begin{pmatrix} \vec{U}(t_k) \\ \vec{\xi}(t_k) \end{pmatrix} + b \right) \quad (36)$$

for appropriate H and b (which are explained in Supplement D below). Now, the distribution $p(r(t_k) | r(t_{k-1}))$ is Gaussian and thus fully defined by its mean $\mathbb{E}[r(t_k)]$ and its covariance $\mathbb{C}[r(t_k)]$. The covariance $\mathbb{C}[r(t_k)]$ is again of the form $\gamma^2 \check{S}_k$ for some \check{S}_k because every filtering covariance is, and the information operator is noise-free. Due to the prediction error decomposition (Schweppe, 1965),

$$p(r(t_0), \dots, r(t_K) | \gamma) = p(r(t_0) | \gamma) \prod_{k=1}^K p(r(t_k) | r(t_{k-1}), \gamma). \quad (37)$$

Recall that the dimension of $r(t_k)$ is $N + 1$ because \mathbb{X} consists of $N + 1$ grid points. Because everything is Gaussian, the negative log-likelihood (as a function of x) decomposes into

$$-\log p(x; r(t_0), \dots, r(t_K) | \gamma) = \frac{1}{2} \left(\sum_{k=0}^K \|x - \mathbb{E}[r(t_k)]\|_{\mathbb{C}[r(t_k)]^{-1}}^2 - (K + 1)(N + 1) \log \gamma^2 \right). \quad (38)$$

Setting the γ -derivative of the negative log-likelihood to zero, i.e. maximising it with respect to γ , yields the MLE from Eq. (29). Overall, the derivation is very similar to those provided by Bosch et al. (2021); Tronarp et al. (2019) for ODE initial value problems, and Krämer and Hennig (2021) for ODE boundary value problems.

D FULLY NONLINEAR (WHITE-NOISE) APPROXIMATE INFERENCE

The main paper explained white-noise approximations only using the example of semilinear PDEs.

This section recalls the basic concepts for linear PDEs and the more general case of fully nonlinear PDEs. The latter goes together with an explanation of Taylor-series linearisation of the nonlinear PDE vector field, so it also supports the statements in Section 4.4. The central idea is to regard the PDE vector field $F(t, x, u(t, x), \mathcal{D}u(t, x))$ as a function of U and ξ , instead of u and $\mathcal{D}u$, and linearise each non-linearity with a first order Taylor series.

D.1 Linear PDEs

At first, we explain the general idea for a linear PDE

$$F(t, x, u(t, x), \mathcal{D}u(t, x)) = Au(t, x) + B\mathcal{D}u(t, x) \quad (39)$$

for some coefficients A and B . If u is scalar-valued, A and B are scalars. If u is vector-valued, A and B are matrices. $\mathcal{D}u = Du + \xi$ allows for recasting the vector field F as a function of u and ξ , instead of u and $\mathcal{D}u$,

$$\tilde{F}(t, U(t), \xi(t)) := F(t, \mathbb{X}, U(t), \xi(t)) = AU(t) + BDU(t) + B\xi(t). \quad (40)$$

Since this equation is linear in u and $\mathcal{D}u$, thus also in U and ξ , there is no need for Taylor-series approximation. The linear measurement model emerges as

$$r_{\text{linear}}(t_k) \mid \vec{U}(t_k), \vec{\xi}(t_k) \sim \delta \left[\dot{U}(t_k) - AU(t_k) - BDU(t_k) - B\xi(t_k) \right]. \quad (41)$$

If the temporal evolution of ξ is ignored through replacing k_t with a white-noise approximation as in Section 4.5, Eq. (41) becomes

$$r_{\text{linear}}(t_k) \mid \vec{U}(t_k) \sim \mathcal{N} \left(\dot{U}(t_k) - AU(t_k) - BDU(t_k), BEB^\top \right). \quad (42)$$

The same principle can be generalised to fully nonlinear problems as follows.

D.2 Nonlinear PDEs

Next, consider a fully nonlinear PDE vector field $F(t, x, u(t, x), \mathcal{D}u(t, x))$. This general setting includes semilinear and quasilinear systems of equations. As before, we discretise the spatial domain using \mathbb{X} , use $\mathcal{D}u = DU + \xi$, rewrite

$$\tilde{F}(t, U(t), \xi(t)) := F(t, \mathbb{X}, u(t, \mathbb{X}), Du(t, \mathbb{X}) + \xi(t)), \quad (43)$$

and infer the solution via \tilde{F} . Since \tilde{F} is nonlinear, we have to linearise it before proceeding. Let $\nabla_U \tilde{F}$ be the derivative of \tilde{F} with respect to U , and $\nabla_\xi \tilde{F}$ be its ξ -counterpart. We linearise \tilde{F} at some $\eta = (\eta_U, \eta_\xi) \in \mathbb{R}^{2(N+1)}$,

$$\tilde{F}(t, U(t), \xi(t)) \approx \tilde{F}(t, \eta_U, \eta_\xi) + \nabla_U \tilde{F}(t, \eta_U, \eta_\xi)(U(t) - \eta_U) + \nabla_\xi \tilde{F}(t, \eta_U, \eta_\xi)(\xi(t) - \eta_\xi) \quad (44)$$

$$=: H_U U(t) + H_\xi \xi(t) + b(t). \quad (45)$$

η_U and η_ξ are commonly chosen as the predicted mean of ξ and U , which corresponds to extended Kalman filtering (Särkkä and Solin, 2019). As a result of the linearisation, the information model reads

$$r_{\text{nonlinear}}(t_k) \mid \vec{U}(t_k), \vec{\xi}(t_k) \sim \delta \left[\dot{U}(t_k) - H_U U(t_k) - H_\xi \xi(t_k) - b \right], \quad (46)$$

or, for the white-noise version,

$$r_{\text{nonlinear}}(t_k) \mid \vec{U}(t_k) \sim \mathcal{N} \left(\dot{U}(t_k) - H_U U(t_k) - b, H_\xi E H_\xi^\top \right). \quad (47)$$

Eqs. (46) and (47) are both linear in the states, and inference is possible with Kalman filtering and smoothing.