

Fluids You Can Trust: Property-Preserving Operator Learning for Incompressible Flows

Ramansh Sharma¹, Matthew Lowery¹, Houman Owhadi², Varun Shankar^{1*}

¹Kahlert School of Computing, University of Utah, UT, USA.

²Department of Computing and Mathematical Sciences, California Institute of Technology, CA, USA.

*Corresponding author(s). E-mail(s): shankar@cs.utah.edu;

Contributing authors: ramansh@cs.utah.edu; mlowery@cs.utah.edu; owhadi@caltech.edu;

Abstract

We present a novel property-preserving kernel-based operator learning method for incompressible flows governed by the incompressible Navier–Stokes equations. Traditional numerical solvers incur significant computational costs to respect incompressibility. Operator learning offers efficient surrogate models, but current neural operators fail to exactly enforce physical properties such as incompressibility, periodicity, and turbulence. Our method maps input functions to expansion coefficients of output functions in a property-preserving kernel basis, ensuring that predicted velocity fields *analytically* and *simultaneously* preserve the aforementioned physical properties. We evaluate the method on challenging 2D and 3D, laminar and turbulent, incompressible flow problems. Our method achieves up to six orders of magnitude lower relative ℓ_2 errors upon generalization and trains up to five orders of magnitude faster compared to neural operators. Moreover, while our method enforces incompressibility analytically, neural operators exhibit very large deviations. Our results show that our method provides an accurate and efficient surrogate for incompressible flows.

Keywords: Operator Learning, Surrogate Modeling, Incompressible Flow, Kernel Methods

Incompressible fluid flows arise in an enormous range of engineering and scientific applications, such as the study of flow past airfoils & wings [1], aerodynamics [2–4], weather prediction [5], chemical mixing [6], and hemodynamics [7–12]. These flows are typically modeled with the incompressible Navier–Stokes (INS) equations given by:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = \nu \nabla^2 \mathbf{u} - \frac{1}{\rho} \nabla p + \frac{1}{\rho} \mathbf{f}, \quad \text{on } \Omega \times (0, T], \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad \text{on } \Omega \times (0, T], \quad (2)$$

$$\mathcal{B}\mathbf{u} = \mathbf{g}, \quad \text{on } \partial\Omega \times (0, T], \quad (3)$$

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x}), \quad \text{on } \Omega, \quad (4)$$

where \mathbf{u} is the divergence-free velocity field, p is the pressure, ν is the kinematic viscosity, ρ is the fluid density, \mathbf{f} is the external body forces, \mathcal{B} denotes the boundary operator, and \mathbf{u}_0 and \mathbf{g} are the initial and boundary conditions respectively. These equations constitute a system of nonlinear partial differential equations (PDEs) and encompass a variety of physical phenomena such as boundary layers, vortex dynamics, flow separation, and turbulence. The challenging nature of these PDEs necessitates specialized numerical discretizations of which there are four dominant classes: (1)

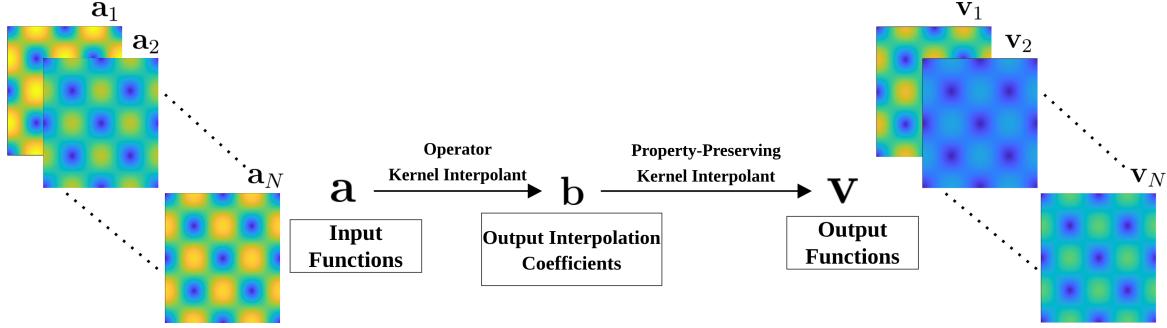


Fig. 1: Schematic diagram of the proposed property-preserving kernel method.

saddle-point methods, which solve the coupled PDE system directly and reach arbitrary orders of accuracy [13–16]; (2) high-order methods that treat the pressure explicitly and recover it via the pressure Poisson equation (“PPE” methods) [15, 17, 18]; (3) artificial compressibility methods [19, 20]; and (4) temporally low-order projection-based methods [21–27]. All of these methods come with their own computational bottlenecks, typically requiring problem-specific preconditioners [15, 16, 28], expensive Poisson solves [17, 29, 30] or vector Helmholtz solves [20, 31]. Further, explicit treatment of the nonlinear advection term leads to the Courant–Friedrichs–Lewy (CFL) stability constraint, severely limiting admissible time-step size [32]. These issues have led to a growing interest in developing *surrogate models* that can reduce the cost of obtaining solutions for new problem configurations [33–36]. Surrogate models for PDEs include reduced order models [37, 38], machine learning (ML) based methods [39–45]. More recently, operator learning models [46] that act directly on function spaces have emerged as a competitive class of surrogates. Operator learning seeks to approximate the *solution operator* of the PDE as a mapping from input functions (geometry, PDE parameters, initial conditions, and/or boundary conditions) to output functions (solutions). Prominent neural operator architectures include deep operator networks (DeepONets) [47], which train two separate neural network in conjunction; Fourier neural operators (FNOs) [48] and their geometrically-flexible variants such as the Geo-FNO [48], which involve kernel-based integral operators computed implicitly via the fast Fourier transform(FFT); the kernel neural operator (KNO) [49] which uses explicit kernels and quadrature instead of the FFT; deep Green networks (DGN) [50, 51] and their graph neural operator (GNO) counterparts [52], which aim to learn the Green’s kernel either globally or locally; and more recently, transformer-based neural operators that leverage the attention mechanism [53–55]. Complementary to these methods, recent work [56, 57] has shown that kernel/Gaussian process (GP) regression can be competitive for operator learning as well. Kernel methods for operator learning are meshless and typically use only a single trainable parameter; note that operator-valued kernels were first introduced in Kadri et al. [58, 59]. Despite these advances, to the best of our knowledge, current operator learning techniques are unable to *analytically* and *simultaneously* satisfy multiple fluid properties such as incompressibility, periodicity, and turbulence-related power laws. Recent attempts in this direction with either soft constraints [60–63] or hard constraints [64, 65] have only succeeded in approximate enforcement of incompressibility. Unfortunately, a surrogate that fails to analytically satisfy these properties can produce physically-inconsistent predictions, accumulate spurious divergence, or be unable to replicate key flow features even if its pointwise errors are small.

In this work, we address this gap by introducing a novel property-preserving kernel-based operator learning method for incompressible flows. Unlike the kernel method in Batlle et al. [56], we learn a map from input function samples to *interpolation coefficients* associated with output functions, where these coefficients are associated with a kernel basis that is analytically property-preserving. Our framework uses two kernel interpolants: (1) one that interpolates the output functions using a property-preserving kernel basis, and (2) one that consumes input functions and fits the interpolation coefficients from (1). The resulting method only uses two trainable parameters. This approach decouples pointwise generalization errors in our surrogate models from their ability to respect physical constraints; while the former still depends on the training data, the latter is always guaranteed to machine precision; see Figure 1 for an illustration of our method.

We illustrate the capabilities of our method with extensive numerical experiments on an extensive suite of 2D and 3D problems involving incompressible flows. We construct surrogates that are all analytically incompressible on generalization, but that also analytically satisfy periodicity or

turbulence-related power laws when applicable. Across these problems, we compare our method against the vanilla kernel method in Batlle et al. [56], and two state-of-the-art neural operator baselines, Geo-FNO [66] and Transolver [67]. The resulting method is meshless; interpretable; typically orders of magnitude more accurate than existing neural operator approaches even in pointwise generalization errors; analytically property-preserving (if the properties are known in advance); more efficient than neural operators in terms of trainable parameter counts (two for our method vs millions for neural operators); and orders of magnitude faster in terms of training time.

The remainder of the paper is organized as follows. Section 1 presents results on numerical experiments on incompressible flow problems, comparing our method against competitors in terms of accuracy, divergence, and runtime. In Section 2, we discuss the implications of the results in detail and discuss future directions. In Section 3, we present the mathematical formulation of the property-preserving kernel method, with a focus on incompressibility, periodicity, and turbulence-related power laws. We also present approximation theorems regarding convergence rates of the method.

1 Results

We now present experimental results comparing our method with operator learning baselines on a diverse set of challenging 2D and 3D operator learning problems involving incompressible flows. We demonstrate that the property-preserving kernel method (PPKM) consistently outperforms state-of-the-art neural operators by several orders of magnitude in accuracy while *analytically* preserving incompressibility and other properties across the board. We compared our method against three baselines: (i) the “vanilla” kernel method (VKM) [56], (ii) the Geo-FNO (geometry-aware FNO) [66], and (iii) the Transolver [67]; implementation details in Appendix E. Several variations of the PPKM were used in the experiments; the π -PPKM is spatially periodic, the η -PPKM incorporates turbulence power laws, and the $\pi\eta$ -PPKM includes both properties. The PPKM (and its variations) and VKM were trained on the NVIDIA RTX 4080 GPU (with the selection of the shape parameter ϵ done on an AMD Ryzen 16-core CPU) using double precision while the neural operators were trained on the stronger NVIDIA A100 and A40 GPUs in single precision; these neural operators typically had too many parameters to allow for double precision training. Despite our use of double precision, our method consistently trains orders of magnitude faster than the neural operators, though it exhibits an order of magnitude slower inference times.

Experimental details: All datasets (except for the Taylor–Green vortices problems, which have analytical solutions) were produced using the SU2 numerical solver [68] for incompressible flows [69], which uses a second-order accurate finite volume method on unstructured meshes. Moreover, SU2 supports a variety of boundary conditions (BCs), including but not limited to periodic BCs, no-slip BCs, and wall functions for turbulent flow. For the turbulent problems explored, we solved the incompressible Reynolds averaged Navier–Stokes (RANS) equations in SU2 with the Shear Stress Transport (SST) turbulence model. We produced domain geometry meshes using gmsh [70] and used them within SU2; these scripts will be provided in the codebase upon publication of the manuscript. Table 2 lists the details pertaining to the flow in each problem; the relevant flow regime, forcing term f , initial (IC) and boundary conditions, time domain, and operator learning map. For each problem, we generated 10,200 time dependent simulations of which $N = 10,000$ were used for training and $N_{\text{test}} = 200$ for testing. Following the mathematical notation introduced in Section 3, the input functions $a \in \mathcal{A}$ may be scalar- or vector-valued; in contrast, the output functions are always vector-valued incompressible velocity fields. We used the truncated normal distribution for randomly sampling governing parameters such as initial velocity, viscosity, leading coefficients, etc. We define it as $\mathcal{N}_{[a,b]}(\mu, \sigma)$, where μ is the mean, σ^2 is the variance, and the random variables are sampled in the range $[a, b]$. This section reports the results for four important 2D and 3D experiments (two each); results for additional classical benchmark problems are reported in Appendix B. In Appendix C, we also present results on the effect of the ridge parameter θ on the accuracy of our method. Section 1.3.1 discusses results from a GP experiment for uncertainty quantification using the 3D species transport example as a case study (see Section 3.3.1 for details).

Metrics reported: Let $\tilde{\mathbf{v}} \in \mathbb{R}^m$ and $\mathbf{v} \in \mathbb{R}^m$ be the vectors containing the pointwise velocity magnitudes of the predicted and true velocity fields respectively. We report the spatial relative ℓ_2 error computed as $(\|\tilde{\mathbf{v}} - \mathbf{v}\|_2)/\|\mathbf{v}\|_2$, *i.e.*, the relative error in the magnitudes. Further, we report the maximum pointwise interior divergence (to remove possible edge effects). Both metrics are averaged over the test functions. We additionally performed ablation studies over the number of training functions N (uniformly sampled at random without replacement) and the number of spatial samples

m and report the results with the combination that yielded the lowest errors for our method. For the Geo-FNO, we reported errors across all N using the largest available m for each problem. Additionally, at the value of N yielding the lowest error for the property-preserving kernel method, we reported Geo-FNO and Transolver errors using the same spatial resolution m . Table 1 reports the relative ℓ_2 errors, maximum pointwise divergence, training and inference (on training set) times, and the θ used by the proposed method. Section 3.4.3 describes how training times are computed for the property-preserving and vanilla kernel methods. We now present each of our benchmark problems.

Table 1: Accuracy, divergence, and runtime results. For each problem, we reported the spatial rel. (relative) ℓ_2 error and maximum pointwise divergence on the test set, total training time, and inference time (on the training set) corresponding to m points and N training functions they were trained with, and the θ used by the PPKM and the VKM. The runtimes are reported in seconds. We report the (statistically significant) mean results for the Geo-FNO and the Transolver models over three random seeds. The choice of the PPKM is shown next to the relative ℓ_2 error (see Section 3.2 for notation). The PPKM and the VKM were trained in double precision while the neural operators were trained in single precision. \sim indicates a blowup in the divergence magnitude ($> 10^3$).

Problem	m	N	θ	Metric	PPKM	VKM	Geo-FNO	Transolver
2D Flow Past a Cylinder (no vortex shedding)	1000	100	10^{-6}	rel. ℓ_2 error	$3.95 \times 10^{-5}, \Phi$	5.32×10^{-5}	9.88×10^{-4}	8.02×10^{-3}
				Divergence	0	0.788	0.725	0.652
				Training time	0.80	0.81	176	55
				Inference time	0.25	0.25	0.19	0.09
2D Flow Past a Cylinder (vortex shedding)	1000	10,000	10^{-6}	rel. ℓ_2 error	$3.14 \times 10^{-7}, \Phi$	2.34×10^{-6}	4.62×10^{-5}	1.19×10^{-4}
				Divergence	0	2.89	2.498	2.497
				Training time	77	77	11,462	4139
				Inference time	4.19	4.19	0.14	0.069
2D Lid-Driven Cavity Flow	1000	10,000	10^{-6}	rel. ℓ_2 error	$1.39 \times 10^{-6}, \Phi$	1.45×10^{-6}	2.19×10^{-4}	1.89×10^{-4}
				Divergence	0	7.96	29.596	29.537
				Training time	77	77	6314	3934
				Inference time	4.19	4.19	0.085	0.065
2D Backward-Facing Step	1000	500	10^{-6}	rel. ℓ_2 error	$1.51 \times 10^{-6}, \Phi$	2.9×10^{-6}	1.055×10^{-3}	1.065×10^{-3}
				Divergence	0	0.907	0.765	0.782
				Training time	0.91	0.92	191	239
				Inference time	0.26	0.26	0.056	0.079
2D Buoyancy-Driven Cavity Flow	5000	10,000	10^{-4}	rel. ℓ_2 error	$6.27 \times 10^{-6}, \Phi$	6.14×10^{-6}	1.44×10^{-4}	1.69×10^{-4}
				Divergence	0	4.28	3.213	3.213
				Training time	92	92	21,373	16,277
				Inference time	12	12	0.30	0.29
2D Taylor-Green Vortices	500	5000	10^{-8}	rel. ℓ_2 error	$8.76 \times 10^{-10}, \Phi^\pi$	1.68×10^{-9}	5.31×10^{-4}	1.22×10^{-4}
				Divergence	0	1.42	12.313	1.426
				Training time	25	26	2419	1988
				Inference time	0.78	0.78	0.066	0.050
2D Taylor-Green Vortices, Spacetime	500	5000	10^{-8}	rel. ℓ_2 error	$8.65 \times 10^{-10}, \Phi^\pi$	8.67×10^{-10}	8.75×10^{-4}	3.07×10^{-4}
				Divergence	0	14.58	\sim	\sim
				Training time	26	26	5966	4580
				Inference time	1.25	1.17	0.16	0.17
2D Merging Vortices	500	500	10^{-4}	rel. ℓ_2 error	$1.09 \times 10^{-4}, \Phi^\pi$	1.09×10^{-4}	5.8×10^{-3}	3.44×10^{-3}
				Divergence	0	0.56	1.048	1.034
				Training time	0.87	0.89	308	200
				Inference time	0.25	0.24	0.080	0.063
3D Species Transport	7000	10,000	10^{-4}	rel. ℓ_2 error	$2.38 \times 10^{-4}, \Phi^\eta$	2.35×10^{-4}	5.11×10^{-4}	2.26×10^{-3}
				Divergence	0	\sim	\sim	\sim
				Training time	90	83	80,738	23,901
				Inference time	21	20	1.13	0.43
3D Flow Past an Airfoil	500	7000	10^{-4}	rel. ℓ_2 error	$0.506, \Phi^{\pi, \eta}$	0.726	0.477	0.600
				Divergence	0	1.44	\sim	\sim
				Training time	53	53	31,535	5721
				Inference time	0.25	0.28	0.55	0.19

1.1 2D flow past a cylinder

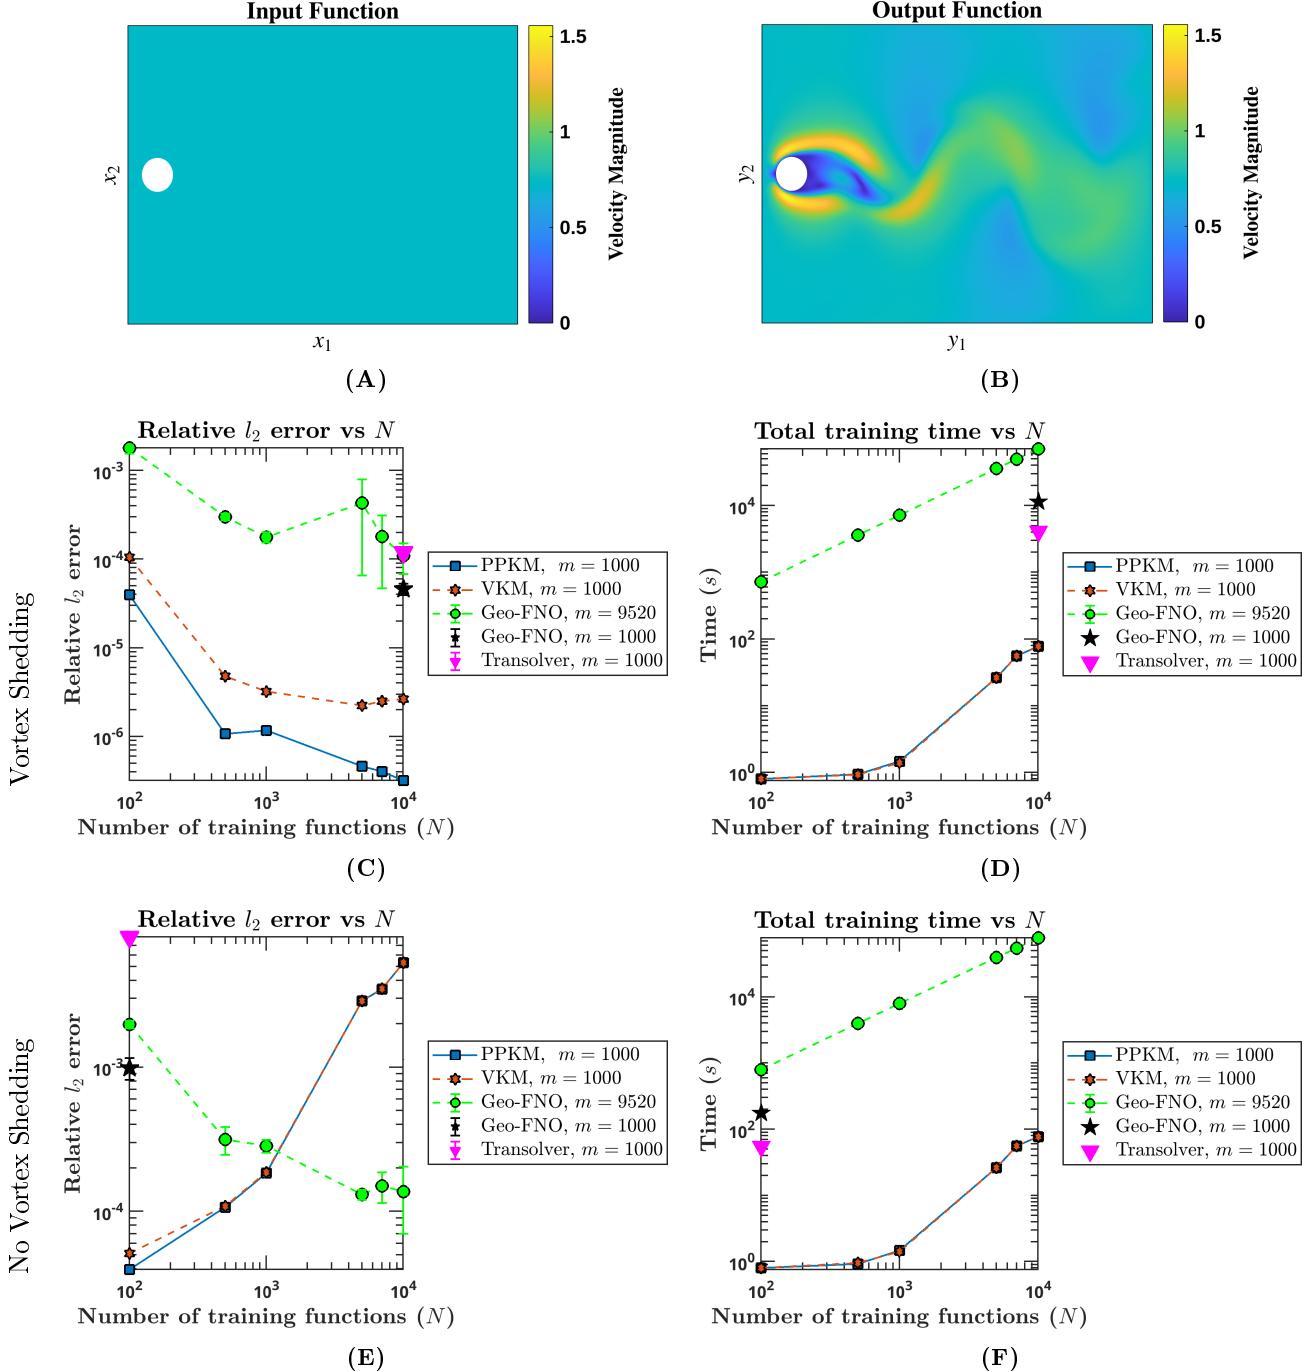


Fig. 2: The 2D laminar flow past a cylinder problem. (A) and (B) show examples of an input function (the initial velocity) and an output function (the final velocity), respectively, from the vortex shedding regime. (C) and (D) show the test relative ℓ_2 errors and training runtimes as functions of N for the vortex shedding regime. (E) and (F) show the same results for the regime without vortex shedding.

This classical problem [71–74] describes laminar fluid flow past a 2D cylinder of radius 0.8 centered at $(1.5, 10)$ within a rectangular domain $\Omega_a = \Omega_v = [0, 20] \times [0, 14]$ representing a channel. We generated a 2D mesh with 9520 points such that the refinement near the cylinder was at least twice as fine as it was elsewhere. We prescribed an inlet BC on the left boundary, free stream BC on the

top and bottom boundaries (set to be the same as the inlet velocity), a no slip BC on the cylinder, and a zero pressure BC on the right boundary. We generated simulations in two flow regimes, with and without vortex shedding [75] in the wake of the cylinder. The IC was varied across simulations, but always set to a constant velocity everywhere sampled from $\mathcal{N}_{[0.1,0.26]}(0.18, 0.026)$ for the case without vortex shedding; and from $\mathcal{N}_{[0.45,0.8]}(0.625, 0.1)$ with vortex shedding. The simulations were run to $T = 10$ with a time-step $\Delta t = 10^{-3}$. We then learned the operator map $\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 10)$. The Reynolds numbers for these flow regimes and the relevant configuration details are provided in row 1 in Table 2. Figures 2A and 2B show examples of input and output functions respectively.

The errors reported in Figure 2C show that in the vortex shedding case, our method consistently performed the best, with lower errors than the neural operators by up to 2 orders of magnitude for the same m and by up to 2.5 orders of magnitude when the neural operators use a much larger m . In the flow without any vortex shedding, our method showed errors (Figure 2E) increasing with N . We suspect that importance sampling of the input functions would restore the proper convergence behavior. However, the lowest error was still achieved by our method at $N = 100$. Further, Figures 2D and 2F show that our method trained orders of magnitude faster than all the baseline methods.

1.2 2D spacetime Taylor–Green vortices

Taylor–Green vortex flow [76, 77] describes the 2D laminar flow of a decaying periodic vortex in the torus $\Omega_a = \Omega_v = [0, 2\pi]^2$ (due to periodic boundary conditions). The flow admits an analytic solution $\mathbf{u} = [u_1, u_2]$ where $u_1(x) = A \sin(x_1) \cos(x_2) e^{-2\nu t}$, $u_2(x) = A \cos(x_1) \sin(x_2) e^{-2\nu t}$. We randomly sampled $A \sim \mathcal{N}_{[0.1,80]}(40, 30)$ and $\nu \sim \mathcal{N}_{[0.0001,1]}(0.006, 0.1)$; for the domain, we used a 2D triangular mesh with 7477 points. Here, the PPKM analytically enforces both incompressibility and spatial periodicity; see Section 3.2 for details. We then sought to recover the operator map given by $\mathcal{G}_1 : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, t)$, $t \in T$, which maps an initial condition to $\Omega_v \times T$ where $T = [0.7, 0.8, 0.9, 1]$. We also alternatively learned the parametric operator map $\mathcal{G}_2 : \tau \rightarrow \mathbf{u}(y, t)$, $t \in T$ where $\Omega_a = \tau \subset \mathbb{R}^2$ is the parameter space of A and ν . We report the relative ℓ_2 error and the maximum pointwise divergence averaged over the four timesteps. Example input and output (at $T = 1$) functions are shown in Figures 3A and 3B. Results for the problem focusing on the spatial operator map are shown in Appendix B.1.

Due to limitations in computational resources, we were unable to report neural operator errors for different N for \mathcal{G}_2 in this problem and in Section B.1. Our method significantly outperformed the neural operator baselines in this problem, achieving up to 4.5 orders of magnitude lower errors (Figures 3C and 3E) and up to 2.5 orders of magnitude faster training times (Figures 3D and 3F). Since this problem had no numerical noise in the dataset, all variations of the PPKM and the VKM achieve similar errors as they are all interpolatory; however, only the PPKM analytically satisfies periodicity and incompressibility.

1.3 3D species transport

Next, we investigated operator learning on a 3D turbulent flow problem ¹ inspired by the work of Ubal et al. [78]. This problem models the mixing of air and methane in a static Kenics mixer [78] with three perpendicular blades twisted along the z axis; see Figure 4C. At the $z = 0$ plane, we specified the air’s velocity at $y < 0$ and methane’s velocity at $y > 0$, sampled from $\mathcal{N}_{[1,40]}(20, 20)$; the velocity everywhere else was initialized to zero. The three blades, inner wall, and outer wall used no-slip BCs; the outlet boundary at the plane $z = 0.26$ used a zero pressure BC. We used a maximum of 7000 approximate kernel Fekete points due to memory constraints. We ran the simulation until $T = 0.5$ with $\Delta t = 0.005$. We learned the operator $\mathcal{G} : \mathbf{u}^{\text{inlet}} \rightarrow \mathbf{u}(y, 0.5)$. We show examples of input and output functions in Figures 4A and 4B respectively.

We report the errors in Figures 4D. The Geo-FNO achieved lower errors at smaller values of N , but was overtaken by the PPKM for larger N . Despite the added computational cost due to the 3D problem domain, our method was orders of magnitude faster to train than the neural operators (as shown in Figure 4E).

1.3.1 Uncertainty quantification

We also demonstrate the GP-based uncertainty quantification capabilities of our framework here, described in Section 3.3.1. We solved (24) and (25) to compute the posterior mean $\bar{\mathbf{b}}$ and covariance

¹https://su2code.github.io/tutorials/Inc_Species_Transport_Composition_Dependent_Model

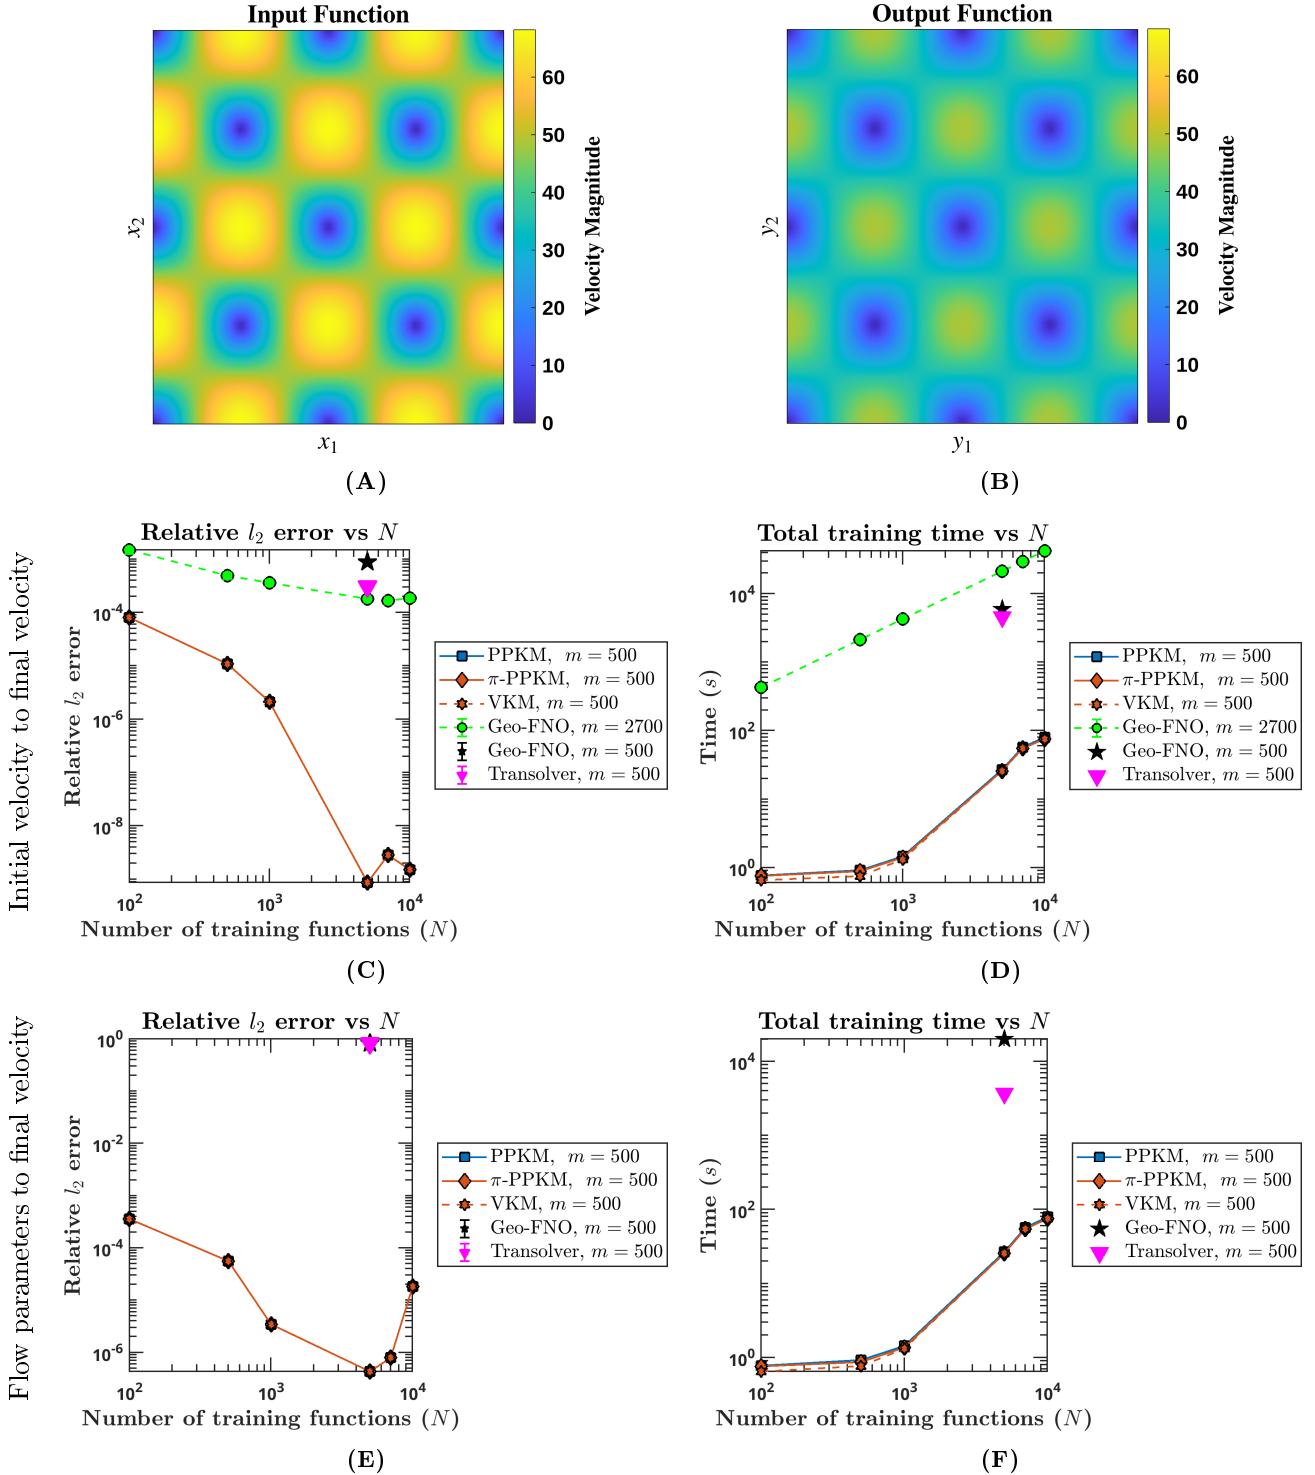


Fig. 3: The 2D laminar Taylor–Green vortices problem for the **spacetime** operator map. (A) and (B) show examples of an input function (the initial velocity) and an output function (the final velocity) at time $T = 1$, respectively. Here, the output functions are snapshots of the velocity field at four timesteps (see Section 1.2 for details). (C) and (D) show the test relative ℓ_2 errors and training runtimes as functions of N for the operator map from the initial velocity to the final velocity at the four timesteps. (E) and (F) show the same results for the operator map from the flow parameters to the final velocity at the four timesteps.

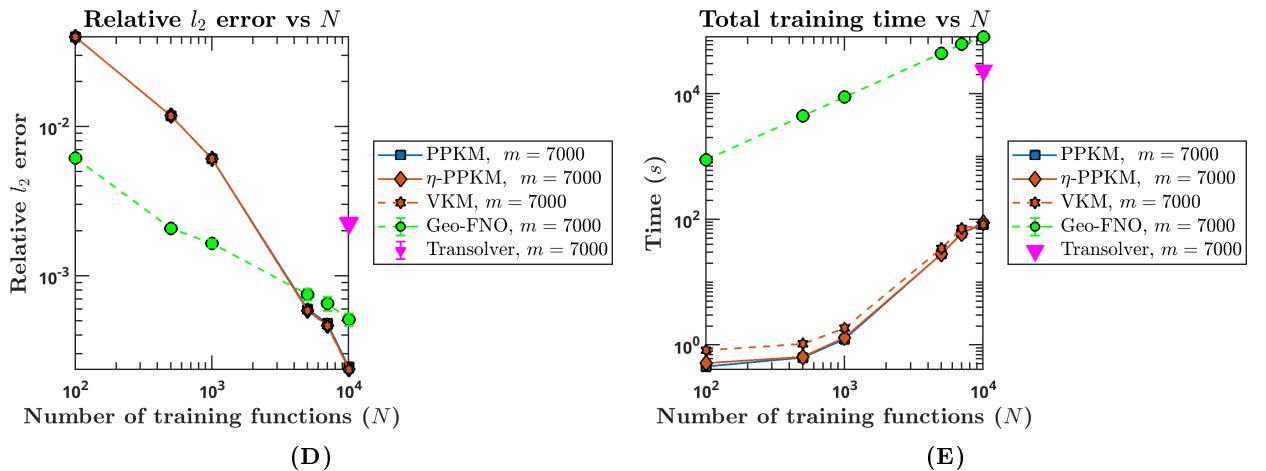
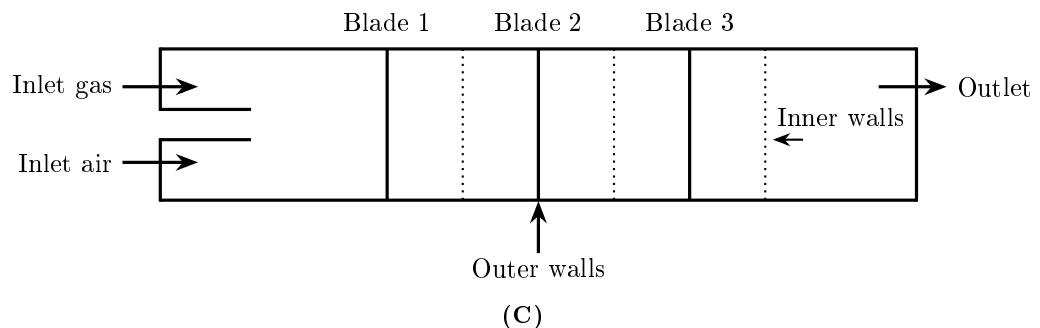
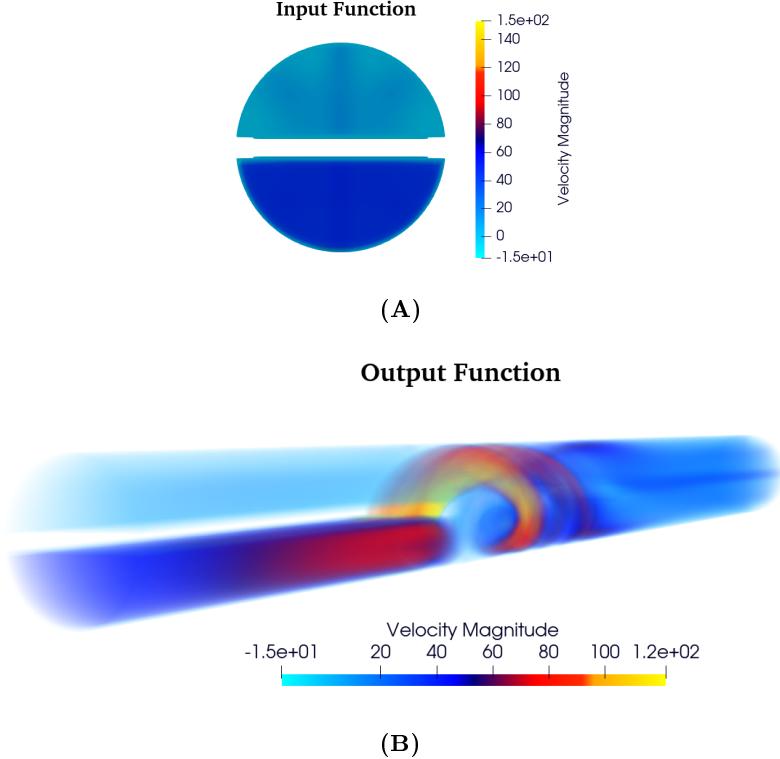


Fig. 4: The 3D turbulent species transport example. (A) and (B) show examples of an input function (the inlet velocity of the two gaseous species at the plane $z = 0$) and an example output function (the final velocity), respectively. (C) shows a top-down yz view of the domain and its placement of the three blades and the inner walls. (D) and (E) show the test relative ℓ_2 errors and training runtimes as functions of N .

Σ^* , respectively, of ξ . Then, we applied a Cholesky factorization on the covariance; $\Sigma^* = LL^T$. Finally, we used the Φ^η kernel to obtain velocity fields corresponding to ± 1 standard distribution with the perturbed coefficients $\bar{\mathbf{b}} + Lz$ and $\bar{\mathbf{b}} - Lz$. For one of the 200 test functions, we show the posterior mean velocity field in Figure 11A and the corresponding velocity fields within ± 1 standard deviation in Figures 11B and 11C. The maximum spatial variance for this test function within one standard deviation was 10^{-5} . All three velocity fields are divergence-free and encoded with turbulence power laws by construction.

1.4 3D flow past an airfoil

Arguably the most difficult (and possibly ill-posed) problem in this work, this problem attempts to recover the turbulent flow field in the wake of a 3D airfoil as a function of the shape of the airfoil on the domain $\Omega_v = [-7, 10] \times [-7, 7] \times [0, 3]$. Let $\Omega_a = \tau$ denote the parametrization of the airfoil geometry, parametrized using the popular NACA 4-digit series [79]. This format had three parameters; the position of the camber (measure of airfoil curvature), maximum camber, and airfoil thickness. We sampled these three parameters from the random distributions $\mathcal{N}_{[1,8]}(3, 1.4)$, $\mathcal{N}_{[3,7]}(6, 2)$, and $\mathcal{N}_{[10,25]}(22, 4)$, respectively. Then, we sampled an airfoil at 2000 two-dimensional points (used as input functions) and then extruded it in the z -direction with 20 uniform slices to obtain a 3D airfoil. The left and right walls had inlet BCs and zero pressure BCs, respectively; the front and back walls had periodic BCs; the top and bottom walls and the airfoil surface had no-slip BCs. We ran the simulation until $T = 1$ with $\Delta t = 10^{-2}$ and measured the velocity on the plane $x = 3$ at $m = 1000$ points. Each simulation used a different airfoil shape causing the resulting mesh and the number of input points to slightly vary. We obtained the velocity on the plane using local divergence-free interpolation; a paper on this technique is forthcoming.

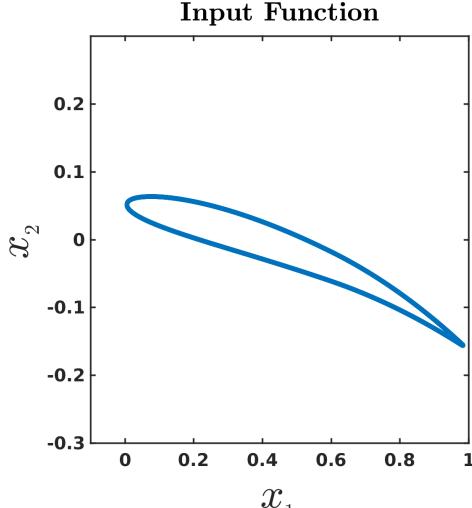
We learned the operator map $\mathcal{G} : \tau \rightarrow \mathbf{u}(y, 1)$. Example input and output functions are shown in Figures 5A and 5B respectively. The errors and runtimes are reported in Figures 5C and 5D respectively. While all methods generally struggled to be accurate here, the $\pi\eta$ -PPKM and π -PPKM outperformed the PPKM and VKM as N increased. These results show the accuracy benefits of simultaneously enforcing multiple fluid properties in the property-preserving kernel. Using the efficiency techniques described in Section 3.4, our method trains orders of magnitude faster than the neural operators.

2 Discussion

The underlying motivation for this work is the observation that a trustworthy surrogate must actually satisfy incompressibility and other properties analytically, much like the PDE solvers that these surrogates aim to replace; while pointwise generalization errors are important, lower errors alone are insufficient.

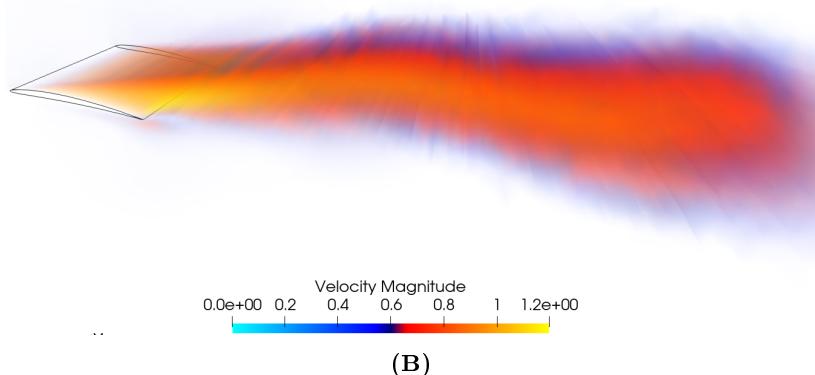
Our property-preserving operator learning framework achieves this through a key idea: learning maps from input function samples to expansion coefficients of output functions in a property-preserving basis. The core of the method involves (1) a property-preserving kernel interpolant for each training output function that analytically preserves incompressibility, periodicity, and turbulence (when applicable); and (2) an operator kernel interpolant that maps from input function samples to the interpolation/expansion coefficients obtained from (1). We chose kernels for (1) because they allowed for the *simultaneous* and *analytical* enforcement of multiple properties without any feature engineering; we chose kernels for (2) because kernel methods allow for rapid training at scale. Our method also allows for natural uncertainty quantification via Gaussian processes (GPs), though this is not the primary focus of this work. Our method also admits pessimistic and optimistic worst-case *a priori* error estimates, which we discuss in Section 3.5.

Remarkably, our results from Section 1 show that the PPKM almost always outperforms state-of-the-art neural operators (and the vanilla kernel method) even on the metric of pointwise generalization error across a suite of benchmark problems (see Table 2 for experimental configuration details). Our method was anywhere from one to six orders of magnitude more accurate on generalization. In addition, our method produces *analytically* incompressible flow fields, in contrast with neural operators; Table 1 shows the latter typically exhibited absolute divergences of $O(1) - O(10^4)$ across this benchmark suite. Surprisingly, our method achieves these superior results with **only two trainable parameters**—one for the property-preserving interpolant, and one for the

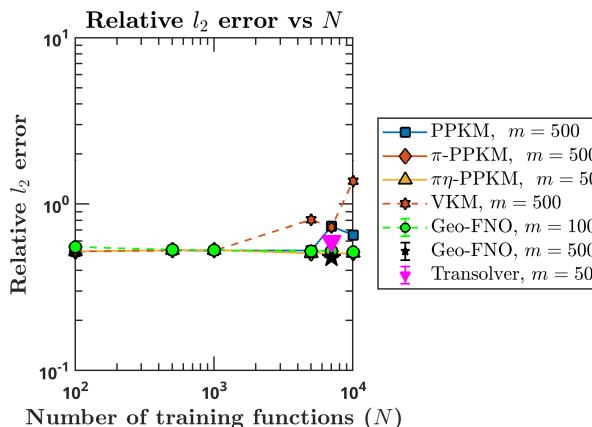


(A)

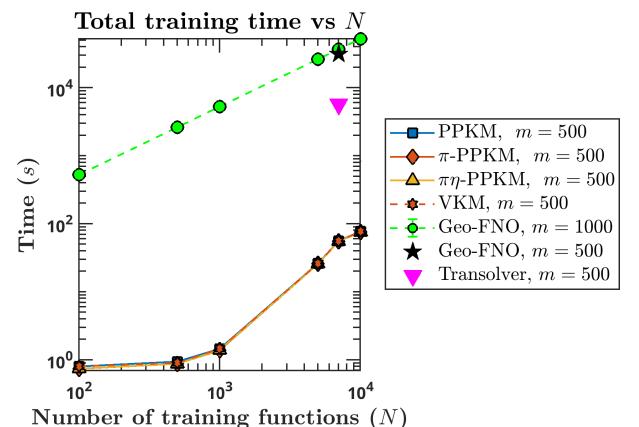
Output Function



(B)



(C)



(D)

Fig. 5: The 3D turbulent flow past an airfoil. (A) and (B) show examples of an input function (the 2D set of points constituting an airfoil shape) and an example output function (the final velocity), respectively. (C) and (D) show the test relative ℓ_2 errors and training runtimes as functions of N .

operator interpolant; in contrast, neural operators typically required $O(10^5)$ to $O(10^6)$ trainable parameters.

The benefits of our method extend beyond superior accuracy. The PPKM consistently trains up to five orders of magnitude faster than the neural operators, despite the use of double precision floating point (fp64) for our method and single precision floating point (fp32) for neural operators; note that it is in general inadvisable to use fp64 with neural operators due to the difficulties in storing trainable parameters. This training speedup is even more remarkable when one takes into account the fact that our method was timed on desktop GPUs while the neural operators were trained on high-end GPU servers. However, this leads to a **limitation** of our method: our inference times were typically an order of magnitude slower than neural operators; see Table 1. We are currently exploring if this was merely due to the difference in hardware, but it is also plausible that this was due to the difference in speeds of fp32 and fp64 on modern GPUs. Nevertheless, our method remains competitive for inference.

Neural operators are able to tackle very large numbers of training functions (N) primarily through batching techniques, but kernel methods do not directly allow for batching in N . However, our simple and powerful streaming technique allowed us to apply the kernel method on the GPU efficiently even for $N = 10,000$ training functions. In addition, our recursive Schur complement techniques in conjunction with careful interpolation node selection allow us to reduce the preprocessing times to $O(dm^3)$ (over a naive $O(d^3m^3)$); storage costs to $O(dm^2)$ (over a naive $O(d^2m^2)$); and inference times to $O(N^2 + Ndm^2)$ (over a naive $O(N^2 + Nd^2m^2)$). These critical details allowed us to actually tackle large, real-world problems in 3D, using only desktop GPUs. Another contribution of this work, though not major, is the meticulous and careful documentation of domain geometries, solver details from the SU2 solver, initial and boundary conditions, and flow regimes for all of our problems in Table 2.

We conclude with a discussion on the limitations of our method and avenues for future work. The primary limitation of kernel methods— which we partially worked around by streaming and efficient linear algebra— is the need to store and compute with large dense interpolation and evaluation matrices. Further, kernel methods typically perform best when at the edge of ill-conditioning (at least when using kernel translates as the basis), which necessitates the use of fp64 to ensure high accuracy; however, this in turn leads to slower GPU computation. Further speedups may be made possible by the use of multipole expansions [80], treecodes [81], sparse kernels [82], and low-rank approximation [83, 84]. Finally, our approach leverages *global* property-preserving kernel interpolation for the output functions, which inflates memory requirements; local interpolation techniques offer a way around this while potentially allowing for greater accuracy. We will also explore generalizations to compressible flow, electromagnetism, and magnetohydrodynamics (MHD), each of which come with their own unique properties and constraints.

3 Methods

3.1 Overview

We now briefly describe kernel-based operator learning and present the high-level algorithm used in this work.

3.1.1 An overview of operator learning

Let \mathcal{A} and \mathcal{V} be two separable Banach spaces of functions and let $\mathcal{G} : \mathcal{A} \rightarrow \mathcal{V}$ be an operator that maps functions $a \in \mathcal{A}$ (input functions) to functions $v \in \mathcal{V}$ (output functions), i.e. $\mathcal{G}(a) = v$. The goal of operator learning is to learn an approximation $\tilde{\mathcal{G}} \approx \mathcal{G}$ given N pairs of input and output functions, $\{(a_i, v_i)\}, i = 1, \dots, N$. Let Ω_a and Ω_v be the domains for the input and output functions respectively. The choice of these domains is arbitrary; they can be spatial, spatiotemporal, or parametric, for instance. In practice, the functions are provided as a finite number of samples on their respective domains; let n and m denote the number of samples provided for the input and output functions, respectively. Further, in the most general setting and the setting of this work, the functions can be vector-valued so that $a : \Omega_a \rightarrow \mathbb{R}^p$ and $v : \Omega_v \rightarrow \mathbb{R}^d$. In this work, we will focus on a kernel-based parametrization of $\tilde{\mathcal{G}}$, a generalization of the work by Batlle et al. [56]. The kernel method for operator learning (first proposed by Batlle et al. [56]) formulates $\tilde{\mathcal{G}}$ as

$$v \approx \tilde{\mathcal{G}}(a) = \chi \circ \tilde{f} \circ \varphi(a), \quad (5)$$

where $\chi: \mathbb{R}^m \rightarrow \mathcal{V}$ reconstructs v by interpolating its m samples, $\varphi: \mathcal{A} \rightarrow \mathbb{R}^n$ generates n samples of a , and $\tilde{f}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ maps from n samples of a to m samples of v . Following Battelle et al. [56], we endow both \mathcal{A} and \mathcal{V} with reproducing kernel Hilbert spaces (RKHSs) and choose kernel maps for χ and \tilde{f} . Specifically, we choose to work with positive-definite radial kernels (also called radial basis functions or RBFs) so that, assuming pairwise distinct interpolation points, the linear systems in (8), (18), and (21) admit unique interpolants [85–87].

3.1.2 Our contribution

Algorithm 1 Training and evaluation procedure for the property-preserving kernel method.

Require:

- 1: **Training dataset:** N pairs of input and output function evaluations, $\{(\mathbf{a}_i, \mathbf{v}_i)\}_{i=1}^N$.
- 2: **Inference example:** $(\mathbf{a}^*, \mathbf{v}^*)$ is an unseen pair of input and output function evaluations.
- 3: Property-preserving kernel Φ and operator kernel λ . Ridge parameter θ .

Let: $\mathbf{a} \in \mathbb{R}^{np}$ and $\mathbf{v} \in \mathbb{R}^{md}$ denote the column vectors of the training input and output function evaluations respectively.

Let: $\mathbf{v}^* \in \mathbb{R}^{m^*d}$ denote the column vector of the inference output function evaluations at m^* points.

Let: Φ denote the $md \times md$ Gramian matrix associated with Φ , see Section 3.2 for details.

Let: \mathbf{K} denote the $N \times N$ Gramian matrix associated with λ , see Section 3.3 for details.

Let: \mathbf{K}^* denote the $1 \times N$ operator kernel evaluation matrix and Φ^* denote the $m^*d \times md$ property-preserving kernel evaluation matrix.

Training:

- 1: Solve the block linear system (8) for \mathbf{b}_i , $i = 1, \dots, N$ using the techniques in Section 3.4.2.
- 2: Find ϵ as the root of the function $\kappa(\mathbf{K}_\epsilon) - 10^{15}$ using the approach outlined in Section 3.4.1.
- 3: Solve the linear system (22) for \mathbf{c}^j , $j = 1, \dots, md$, using ϵ and θ .

Generalization:

- 1: Obtain output coefficients as $\mathbf{b}^* \leftarrow \begin{bmatrix} \mathbf{K}^* \mathbf{c}^1 \\ \vdots \\ \mathbf{K}^* \mathbf{c}^{md} \end{bmatrix} \in \mathbb{R}^{md}$.
- 2: Obtain a property-preserving prediction $\tilde{\mathbf{v}}^* \leftarrow \Phi^* \mathbf{b}^*$.

Our method emerges from the observation that if χ is carefully constructed to be a matrix-valued kernel interpolant that *analytically* enforces desirable properties (such as incompressibility, periodicity, and turbulence) when interpolating v , then changing \tilde{f} to output expansion coefficients in that kernel interpolant ensures that all predictions made by $\tilde{\mathcal{G}}(a)$ for every a will automatically (and analytically) satisfy those properties as well. Our method therefore has two main components; (i) the vector-valued **operator kernel** interpolant \tilde{f} , and (ii) the vector-valued **property-preserving kernel** interpolant χ . We summarize the method in Algorithm 1; describe the two components in detail in Sections 3.2 and 3.3; and present relevant implementation details in Section 3.4. A schematic of our method is shown in Figure 1.

A note on block vector notation: We briefly describe the vector notation used throughout this section for input and output function evaluations. Let $v: \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a vector valued function that is evaluated at a set of points $\{y_j\}_{j=1}^m \subset \mathbb{R}^d$. We denote these evaluations by the block vector

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}^1 \\ \vdots \\ \mathbf{v}^d \end{bmatrix} \in \mathbb{R}^{md}, \quad (6)$$

where each block $\mathbf{v}^k = [v^k(y_1), \dots, v^k(y_m)]^T \in \mathbb{R}^m$ contains the evaluations of the k^{th} spatial component of v . This notation is helpful when discussing vector-valued kernel interpolants.

3.2 The property-preserving kernel interpolant

We now discuss how various properties can be analytically encoded into a kernel basis. For this work, the relevant function spaces \mathcal{V} always consist of incompressible velocity fields which are solutions to PDEs such as (1). Additionally, depending on the problem, these velocity fields can also

exhibit spatial periodicity (due to boundary conditions) and turbulence (due to modeling choices). In our method, we use a property-preserving kernel basis in χ such that the velocity fields recovered analytically preserve these spatial properties.

Let $Y = \{y_j\}_{j=1}^m \subset \Omega_v \subset \mathbb{R}^d$ be a set of spatial points where an output function $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is evaluated; further, let the block vector $\mathbf{v} = [\mathbf{v}^1, \dots, \mathbf{v}^d]^T \in \mathbb{R}^{md}$ denote the function evaluations of v at Y . Letting $\Phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ be a matrix-valued positive-definite property-preserving kernel (described in the following subsections), we explicitly write the resulting kernel interpolant as

$$\chi(y) = \sum_{j=1}^m \Phi(y, y_j) \underline{b}^j, \quad (7)$$

where $\underline{b}^j \in \mathbb{R}^d$. In order to interpolate v , we enforce $\chi(y) = v(y)$ for all $y \in Y$. These interpolation conditions give rise to the linear system

$$\Phi \mathbf{b} = \mathbf{v}, \quad (8)$$

where $\Phi_{ij} = \Phi(y_i, y_j)$ is the $md \times md$ block Gramian matrix arising from evaluations of Φ , and $\mathbf{b} = [\mathbf{b}^1, \dots, \mathbf{b}^d]^T \in \mathbb{R}^{md}$ contains the interpolation coefficient vectors for each spatial component. In the operator learning setting, there are multiple such target functions v_i ; we denote by $\mathbf{b}_i \in \mathbb{R}^{md}$ the coefficient vector for the i^{th} output function, and by b_i^j its j^{th} component. The system (8) has a unique solution if the points in Y are distinct [82, 86]. This $md \times md$ system is computationally expensive to solve ($O(d^3 m^3)$ operations for a Cholesky factorization, and $O(d^2 m^2)$ operations for subsequent solves for each right-hand side). However, we never form and store Φ , but instead compute only with the individual blocks in Φ , reducing the solution cost significantly. These details are described in Section 3.4. Once the interpolation coefficients in (7) are computed, the interpolant can be evaluated at any location (ideally within the hull of Y). Let Y^* be a set of m^* evaluation locations. Then, the evaluation of $\chi(y)$ at Y^* can be written as a matrix-vector product

$$\chi(y)|_{Y^*} = \Phi^* \mathbf{b}, \quad (9)$$

where $\Phi^*(y_k^*, y_j)$, $k = 1, \dots, m^*$, $j = 1, \dots, m$ is the $m^* d \times md$ (rectangular) evaluation matrix.

Thus far, our description of these property-preserving matrix-valued kernels has been abstract. In the following subsections, we describe in detail how to simultaneously encode incompressibility, periodicity, and turbulence in Φ .

3.2.1 Incompressibility

The primary goal of this work is the high-fidelity surrogate modeling of incompressible fluid flows. To that end, our matrix-valued kernel Φ and the corresponding property-preserving interpolant (7) always enforce incompressibility analytically.

The approximation of divergence-free (incompressible) vector fields has a long history in the numerical methods and approximation literature. Classical divergence-conforming finite element spaces, such as the Raviart–Thomas [88], Brezzi–Douglas–Marini [89], and Nédélec elements [90], enforce the divergence-free constraint by construction through their basis functions, yielding local and exactly divergence-free numerical solutions. This paradigm established the foundational principle of enforcing physical constraints by restricting the admissible solution space. However, these constructions rely on polynomial interpolation over simplices, thereby necessitating spatial meshes. Subsequent work in the meshfree approximation literature led to the construction of matrix-valued divergence-free kernels from scalar-valued positive-definite kernels [91–95], even leading to numerical methods for meshfree Helmholtz–Hodge decompositions and incompressible flow solvers [96–99]. While Algorithm 1 can easily leverage mesh-based divergence-free approximations, we choose to use kernel-based methods as they are (1) meshfree and therefore more general; and more importantly (2) can *simultaneously encode multiple properties in addition to incompressibility*.

We now describe the kernel construction used in this work. Let $\phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ be a scalar-valued positive-definite kernel that is at least $C^2(\mathbb{R}^d)$. The associated divergence-free matrix-valued kernel

is defined by

$$\Phi(x, y) = \nabla_y \times \nabla_x \times \phi(x, y), \quad (10)$$

where the differential operators act componentwise, and $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ are points in \mathbb{R}^d . For $d = 2$, we obtain the following matrix-valued kernel

$$\Phi(x, y) = \begin{bmatrix} \partial_{x_2} \partial_{y_2} & -\partial_{x_1} \partial_{y_2} \\ -\partial_{x_2} \partial_{y_1} & \partial_{x_1} \partial_{y_1} \end{bmatrix} \phi(x, y). \quad (11)$$

Whereas for $d = 3$ we obtain

$$\Phi(x, y) = \begin{bmatrix} \partial_{x_3} \partial_{y_3} + \partial_{x_2} \partial_{y_2} & -\partial_{x_1} \partial_{y_2} & -\partial_{x_1} \partial_{y_3} \\ -\partial_{x_2} \partial_{y_1} & \partial_{x_3} \partial_{y_3} + \partial_{x_1} \partial_{y_1} & -\partial_{x_2} \partial_{y_3} \\ -\partial_{x_3} \partial_{y_1} & -\partial_{x_3} \partial_{y_2} & \partial_{x_2} \partial_{y_2} + \partial_{x_1} \partial_{y_1} \end{bmatrix} \phi(x, y). \quad (12)$$

By construction, each column of Φ is divergence free with respect to the evaluation coordinate y , and the resulting kernel is positive-definite; a simplified form is also readily available for radial kernels [97]. An important consequence of this construction is that any kernel interpolant of the form (7) lies in a divergence-free space. Consequently, the interpolant χ satisfies $\nabla \cdot \chi = 0$ pointwise throughout the domain, independently of the sampling locations or interpolation coefficients. This provides an analytically exact enforcement of incompressibility.

3.2.2 Periodicity

In addition to incompressibility, many applications require velocity fields to satisfy periodic boundary conditions, typically for enabling the modeling of larger physical domains at lower computational cost but also for modeling truly periodic domain geometries. Given that we began with the divergence-free kernel Φ , this now opens up the issue of additionally incorporating periodicity into Φ .

In the kernel and RBF literature, periodic structure is imposed in settings ranging from interpolation on periodic domains [86, 100] to the numerical treatment of periodic boundary conditions [101] and closed curves [102]. Existing approaches fall into three broad categories: Fourier-based constructions, direct kernel and RBF modifications on periodic domains, and embedding-based techniques that enforce periodicity by construction. Fourier spectral and pseudospectral methods are classical tools for approximating periodic functions [103–105], while kernel and RBF methods impose periodicity via kernel periodization and lattice-point methods [106–108]. Embedding-based methods encode periodicity through an explicit embedding of Euclidean coordinates into a compact manifold, most commonly the circle or torus. This idea has been employed in numerical methods for PDEs on the sphere [109, 110]; for elastic surfaces in fluid-structure interaction problems in Shankar et al. [111], Shankar and Olson [112], Kassen et al. [113]; and for problems posed on tori in Owhadi [99], Fuselier and Wright [114].

Adopting the embedding-based approach, we enforce periodicity by composing kernels with an explicitly periodic embedding into a higher-dimensional Euclidean space. With this approach, periodicity is enforced by construction, allowing standard kernel constructions to be used without requiring kernel periodization. Specifically, we define an embedding $h: \mathbb{T}^d \rightarrow \mathbb{R}^{2d}$ as

$$h(y) = (\cos(y_1), \sin(y_1), \dots, \cos(y_d), \sin(y_d)). \quad (13)$$

Composition with h induces 2π -periodicity in each spatial coordinate. Let $\phi^\pi: \mathbb{R}^{2d} \times \mathbb{R}^{2d} \rightarrow \mathbb{R}$ be a scalar-valued positive-definite kernel defined on the embedding space, *i.e.*, $\phi^\pi(x, y) = \phi(h(x), h(y))$, where ϕ is some positive-definite kernel. Applying the standard curl-curl construction from (10) on ϕ^π gives the *periodic, divergence-free, matrix-valued kernel*

$$\Phi^\pi(x, y) = \nabla_y \times \nabla_x \times \phi^\pi(h(x), h(y)), \quad (14)$$

which inherits periodicity directly from the embedding and is therefore also 2π -periodic in each spatial coordinate; note that this kernel is no longer radial. If periodicity must be imposed in only one coordinate, we modify $h(y)$ to only incorporate the torus embedding into that coordinate.

3.2.3 Turbulence through power laws

Finally, we turn our attention to creating reliable surrogates for turbulent incompressible flows. Turbulence in incompressible flow is a high-Reynolds-number regime of the Navier-Stokes equations characterized by nonlinear vortex stretching, broadband energy cascades, and anomalous dissipation [115, 116]. It is ubiquitous in geophysical and engineering flows, including atmospheric and oceanic boundary layers [117, 118], riverine and open-channel flows [119], and internal and external aerodynamic flows [120]. Turbulent flows arise when the Reynolds number $Re = UL/\nu \gg 1$ and exhibit multiscale structure [116, 121]. Modeling therefore requires statistical closures (RANS) [122], large-eddy simulation (LES) [123], or direct numerical simulation (DNS) [123] depending on the resolved range of scales [116]; a range of other approaches also exist [123–126]. In the high-Reynolds-number regime, kinetic energy injected at large scales is transferred to progressively smaller eddies through nonlinear interactions—Richardson’s cascade—until viscous dissipation dominates at the Kolmogorov scale $\eta \sim (\nu^3/\varepsilon)^{1/4}$, yielding the inertial-range spectrum $E(k) \sim \varepsilon^{2/3}k^{-5/3}$ [115, 121, 127].

Due to the ubiquity of turbulence in practical applications, we also incorporate turbulence preservation into our property-preserving kernels by leveraging the fact that Kolmogorov’s inertial-range prediction $E(k) \sim \varepsilon^{2/3}k^{-5/3}$ implies that second-order statistics scale as power laws in separation r (e.g., $S_2(r) = \mathbb{E}|u(x+r) - u(x)|^2 \sim (\varepsilon r)^{2/3}$), so that velocity fields exhibit self-similar, scale-invariant correlations [115, 121]. More specifically, we follow Owhadi [99, Section 5.1] and define an additive multiscale kernel Φ^η of the form

$$\Phi^\eta = \sum_{s=1}^q \alpha_s \Phi_s, \quad (15)$$

where $\sigma_s = \frac{\sigma_0}{2^s}$ is the shape parameter for Φ_s (see Section 3.4 for details on how σ_0 is picked), $\alpha_s = \sigma_s^\gamma$, and

$$\gamma = \begin{cases} 4, & \text{if } d = 2, \\ \frac{2}{3} + 2, & \text{if } d = 3. \end{cases} \quad (16)$$

We set $q = 5$ for all problems involving turbulence. Here, q denotes the number of modes in the additive kernel, with each mode associated with a characteristic length scale corresponding to eddies at that scale [128]. In general, increasing q improves the approximation accuracy of the multiscale kernel [99, Section 7]. The key idea here is that Φ^η can analytically mimic the Richardson cascade; in the continuum limit, this corresponds to representing Φ^η as a scale mixture whose spectral density matches the Kolmogorov target.

Simultaneous property-preservation

Though we described Φ , Φ^π , and Φ^η separately, our approach allows for the simultaneous property-preserving kernel $\Phi^{\pi,\eta}$, which is matrix-valued, positive-definite, and more importantly, analytically divergence-free, analytically periodic, and analytically approximating the Richardson Cascade through power laws. We use this kernel within the property-preserving interpolant for turbulent, periodic, incompressible flows. In general, our work involves utilizing known problem features and the analytic enforcement of these features in our surrogates. We defer *detection* of these features to future work.

3.2.4 Spacetime kernel interpolation

In problems where Ω_v is a spacetime domain, the interpolant χ must now interpolate spacetime data. As an example, let $\Omega_v = \Omega \times \Gamma$, where Ω denotes the spatial domain and Γ denotes the temporal domain. Consider a function $v(x, t)$, where $x \in \Omega \subset \mathbb{R}^d$ and $t \in \Gamma \subset \mathbb{R}$. Assume that we are provided with a set of spatial snapshots of v on Ω at discrete time instances $\{t_1, \dots, t_T\}$ as part of our training data. We now describe how to incorporate the ability to produce such spacetime predictions (upon generalization) into our surrogates.

We draw from the literature on spacetime kernel design, which employs products of spatial and temporal kernels [129–139]. Let $\psi : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a scalar-valued positive-definite kernel; we select $\psi = \phi$, but other choices are possible. We then define a spacetime interpolant using the product

kernel $\Phi\psi$; this interpolant takes the form

$$\chi(y, t) = \sum_{j=1}^{mT} (\Phi(y, y_j) \psi(t, t_j)) \underline{b}^j, \quad (17)$$

where (y_j, t_j) ranges over all space-time evaluations and the coefficients. As in the spatial case, we enforce the interpolation conditions $\chi(y, t) = v(y, t)$ at all $y \in Y$ and $t \in \{t_1, \dots, t_T\}$, thereby obtaining a block linear system of the form

$$\begin{bmatrix} \Phi \psi_{1,1} & \Phi \psi_{1,2} & \dots & \Phi \psi_{1,T} \\ \Phi \psi_{2,1} & \Phi \psi_{2,2} & \dots & \Phi \psi_{2,T} \\ \vdots & \vdots & \ddots & \vdots \\ \Phi \psi_{T,1} & \Phi \psi_{T,2} & \dots & \Phi \psi_{T,T} \end{bmatrix} \begin{bmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \\ \vdots \\ \mathbf{b}^T \end{bmatrix} = \begin{bmatrix} \mathbf{v}^1 \\ \mathbf{v}^2 \\ \vdots \\ \mathbf{v}^T \end{bmatrix}, \quad (18)$$

where $\psi_{ij} = \psi(t_i, t_j)$, Φ is the Gramian matrix obtained by evaluating the matrix-valued spatial kernel Φ on Y , and the superscripts on \mathbf{b} and \mathbf{v} denote the timestep. Since Φ and ψ are both positive-definite kernels and the spatial and temporal sampling locations are distinct, the kernel interpolant χ is uniquely determined [82, 86]. It is important to note that Φ here can be replaced by Φ^π or $\Phi^{\pi,\eta}$, as needed. Much as in the case of (8), (18) requires the use of efficient linear algebra to contend with the $d + 1$ dimensionality of the problem. We describe these techniques in Section 3.4.2.

3.2.5 Approximate kernel Fekete points

The global kernel interpolation performed throughout this work in (7) and (17) is susceptible to the Runge phenomenon, a well known instability affecting global interpolation of smooth target functions [140, 141]. This issue has been observed both for global polynomial interpolation and for RBF methods [142, 143]. Since the velocity fields considered here are smooth and defined on bounded domains, naive global kernel interpolation on equally-spaced or quasi-uniform points may exhibit large oscillations near the domain boundary.

Most work in the literature addresses this difficulty through careful choice of interpolation points or bases. Approaches include boundary-clustered points [144], spatially varying shape parameters in RBF interpolation [143], and extension-based techniques such as RBF extension [145], which are closely related to Fourier extension methods [146, 147]. Within the RBF literature, greedy point selection algorithms [148, 149] and inducing-point strategies [150] are commonly employed to identify subsets of points that give improved stability and approximation accuracy.

A complementary and well-established strategy in polynomial interpolation is the use of Fekete points [151, 152]. These points are defined as the maximizers of the absolute value of the determinant of the polynomial Vandermonde matrix [153], a property that is closely connected to favorable conditioning and reduced Lebesgue constants, and hence improved interpolation stability. However, exact Fekete points are expensive to compute and are known analytically only in very restricted settings. In practice, *approximate* Fekete points are typically computed via rank-revealing, column-pivoted QR factorizations of Vandermonde matrices [151, 154]. Recent work has begun extending these ideas to kernel Gramian matrices in the univariate setting [155]. Motivated by these developments, we adopt an analogous approach for kernel-based interpolation by computing approximate Fekete points directly from our kernel Gramian matrices. We describe this approach below.

Let Y be the set of candidate points from which we wish to pick approximate Fekete points; typically, this is the set of all data locations available to us. Let ϕ be a scalar-valued positive-definite kernel and let $\mathbf{A}_{ij} = \phi(y_i, y_j)$ be its associated Gramian matrix on Y . Since \mathbf{A} is positive-definite, it admits the Cholesky factorization $\mathbf{A} = LL^T$ where L is lower triangular. The columns of L can be interpreted as a set of feature vectors induced by the kernel ϕ at the points Y . We seek to find a subset of points whose feature vectors are maximally linearly independent. To this end, we apply a rank revealing QR factorization with column pivoting to L

$$LP = QR, \quad (19)$$

where Q has orthogonal columns, R is upper triangular, and P is a permutation matrix containing column pivot indices. The column pivoting orders the columns of L according to their contribution to

the numerical rank. We take the points corresponding to the first m entries of P as the *approximate kernel Fekete points*. These points are used by the property-preserving kernel in all experiments. Example sets of approximate Fekete points are shown in the appendix in Figure 14. In cases where $\Omega_a = \Omega_v$, the same set of points is used for both input and output fields; $n = m$.

3.3 The operator kernel interpolant

Once the interpolation coefficients in (7) or (17) are obtained, the next step in our framework is to treat those coefficients as a function of the input functions and fit them in turn (via the object \tilde{f} mentioned previously); this enables the prediction of new coefficients on generalization (when an unseen input function is supplied). We now describe the construction of \tilde{f} , which is an operator-valued kernel interpolant. Let $a : \mathbb{R}^d \rightarrow \mathbb{R}^p$ be an input function and $X = \{x_1, \dots, x_n\} \subset \Omega_a \subset \mathbb{R}^d$ the set of points on its domain where it is sampled; we denote the block vector of evaluations of a at X by $\mathbf{a} = [\mathbf{a}^1, \dots, \mathbf{a}^p]^T \in \mathbb{R}^{np}$. The operator-valued kernel interpolant (or simply operator kernel interpolant) requires a matrix-valued positive-definite kernel $\Lambda : \mathbb{R}^{np} \times \mathbb{R}^{np} \rightarrow \mathbb{R}^{md \times md}$ that is capable of consuming these \mathbb{R}^{np} block vectors as arguments and outputting a matrix in $\mathbb{R}^{md \times md}$. Equipped with this matrix-valued kernel, we may now write the operator kernel interpolant \tilde{f} as

$$\tilde{f}(\mathbf{a}) = \sum_{i=1}^N \Lambda(\mathbf{a}, \mathbf{a}_i) \underline{c}_i, \quad (20)$$

In order to reduce the dimensionality of the associated interpolation problem, we first select Λ to be a diagonal kernel of the form $\Lambda = \lambda I$, where I is the $md \times md$ identity matrix and $\lambda : \mathbb{R}^{np} \times \mathbb{R}^{np} \rightarrow \mathbb{R}$ is some scalar-valued positive definite kernel. This choice decouples all solves for the components of the coefficients \underline{c}_i . After making this choice, we enforce the interpolation conditions $\tilde{f}(\mathbf{a}_i) = \mathbf{b}_i$, where $i = 1, \dots, N$ indexes training functions and \mathbf{b}_i is the coefficient vector associated with the property-preserving interpolant χ for the i -th output function. This results in the following linear system with N right-hand sides and the shared Gramian of λ

$$\underbrace{\begin{bmatrix} \lambda_{1,1} & \lambda_{1,2} & \dots & \lambda_{1,N} \\ \lambda_{2,1} & \lambda_{2,2} & \dots & \lambda_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{N,1} & \lambda_{N,2} & \dots & \lambda_{N,N} \end{bmatrix}}_{\mathbf{K}} \underbrace{\begin{bmatrix} c_1^1 & c_1^2 & \dots & c_1^{md} \\ c_2^1 & c_2^2 & \dots & c_2^{md} \\ \vdots & \vdots & \ddots & \vdots \\ c_N^1 & c_N^2 & \dots & c_N^{md} \end{bmatrix}}_C = \underbrace{\begin{bmatrix} b_1^1 & b_1^2 & \dots & b_1^{md} \\ b_2^1 & b_2^2 & \dots & b_2^{md} \\ \vdots & \vdots & \ddots & \vdots \\ b_N^1 & b_N^2 & \dots & b_N^{md} \end{bmatrix}}_B \quad (21)$$

The system (21) has a unique solution if the input function sample vectors are all distinct [56, 82, 86]. To account for possible irregularities in the input function samples, we include a ridge parameter θ on the Gramian's diagonal for regularization [56, 156]. This ensures positive-definiteness and serves to control the magnitude of the columns of C . This results in the following linear system:

$$(\mathbf{K}_\epsilon + \theta I) \mathbf{c}^j = \mathbf{b}^j, \quad j = 1, \dots, md, \quad (22)$$

where \mathbf{c}^j are the columns of C , \mathbf{b}^j are the columns of B , ϵ is the shape parameter of the operator kernel, and I is the $N \times N$ identity matrix.

Generalization

Analogous to the evaluation of the property-preserving kernel interpolant in (9), \tilde{f} can be evaluated at an unseen test function \mathbf{a}^* to obtain the property-preserving kernel interpolation coefficients for the corresponding output function \mathbf{v}^* . Denote by \mathbf{K}^* the $1 \times N$ evaluation matrix for the kernel λ so that $(\mathbf{K}^*)_{1j} = \lambda(\mathbf{a}^*, \mathbf{a}_j)$, $j = 1, \dots, N$. Then,

$$\tilde{f}(\mathbf{a}^*) = \mathbf{K}^* C. \quad (23)$$

For N^* test functions, \mathbf{K}^* is simply an $N^* \times N$ matrix. However, note that $\tilde{f}(\mathbf{a}^*)$ merely returns an approximation to the output coefficient vector \mathbf{b}^* (in a property-preserving kernel expansion) to the output function \mathbf{v}^* . Samples of the output function \mathbf{v}^* are in turn obtained by (9).

3.3.1 Operator-valued Gaussian processes for uncertainty quantification

The operator kernel interpolant map also provides us with a natural method for estimating function-level uncertainty. Uncertainty quantification involves calculating the statistical uncertainty associated with a surrogate model's prediction [157]. The applications of this field vary widely, such as aerodynamic shape design [158], materials science [159], inverse problems [160, 161], and seismic wave propagation [160]. In particular, kernel interpolants (regularized or otherwise) automatically form the mean function for a Gaussian process (GP); alternatively, GPs may be thought of as a straightforward probabilistic generalization of kernel interpolation. GP methods have been used to approximate uncertainty in machine learning model predictions [162, 163]. GPs have also been used in kernel based methods [57, 164–166].

While not a major focus of this work, we briefly describe how to leverage the operator kernel interpolant \tilde{f} to quantify uncertainty. The GP naturally induced by the operator kernel is given by $\xi \sim \mathcal{N}(\tilde{\mathbf{b}}, \Lambda)$, where $\tilde{\mathbf{b}}$ is the prior mean and Λ the prior covariance kernel of the GP (and also the kernel of the interpolant (20)). Conditioning the GP on the training observations yields a posterior mean and posterior covariance kernel Σ^* given by

$$\bar{\mathbf{b}} = \Lambda^* (\Lambda + \theta I)^{-1} \mathbf{b}_{\text{train}}, \quad (24)$$

$$\Sigma^* = \Lambda^{**} - \Lambda^* (\Lambda + \theta I)^{-1} (\Lambda^*)^T, \quad (25)$$

where $\mathbf{b}_{\text{train}} = \begin{bmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_N \end{bmatrix} \in \mathbb{R}^{Nmd}$, Λ is the $Nmd \times Nmd$ matrix denoting the prior covariance on the training dataset, Λ^* is the $md \times Nmd$ cross-covariance matrix, and $\Lambda^{**} = \Lambda(\mathbf{a}^*, \mathbf{a}^*)$ is the $md \times md$ matrix denoting the prior covariance on the test dataset. The uncertainty in the predicted output field can be obtained by perturbing the posterior mean using the Cholesky factorization of the posterior covariance. Specifically, if $\Sigma^* = LL^T$, then samples from the predictive distribution within one standard deviation are given by

$$\tilde{\mathbf{b}}^* = \bar{\mathbf{b}} \pm Lz, \quad z \sim \mathcal{N}(0, I), \quad (26)$$

where $z \in \mathbb{R}^{md}$. The perturbed coefficient vector $\tilde{\mathbf{b}}^*$ is then used in the property-preserving kernel interpolant in (8) to obtain a sample velocity field from the posterior distribution (via (9)). Hence, our method provides a framework for *property-preserving uncertainty quantification*. While we leave a detailed exploration of the GP associated with our operator kernel interpolant for future work, we present preliminary results in Section 1.3.1 for a turbulent flow problem in 3D.

3.4 Implementation and training details

We now describe important implementation details for the two kernel interpolants used in our framework.

3.4.1 Kernels, shape parameters, and nuggets

Kernels

We use the Euclidean distance metric in all kernel arguments; $\phi(x, y) = \phi(\|x - y\|_2)$, thereby making all non-periodic kernels radial. In this special case, the divergence-free kernel can be written as $\Phi(x, y) = (-\Delta I + \nabla \nabla^T) \phi(\|x - y\|_2)$ where Δ and ∇ are the \mathbb{R}^d Laplacian and gradient operators respectively acting on y , and I is a $d \times d$ identity matrix. This does not work in the periodic case because the transformation h (see (13)) breaks the radial characteristic of ϕ .

We tested three different choices of radial kernels (RBFs): (i) the Gaussian kernel, (ii) the compactly supported C^4 Wendland kernel [167], and (iii) the C^4 Matérn kernel [168]. The C^4 Matérn kernel outperformed the other two kernels consistently across all experiments (the Gaussian was more ill-conditioned and the Wendland kernel was less accurate when supports were small). We therefore report results using only this kernel. It is given by $\phi_\epsilon(r) = e^{-(r\epsilon)}(3 + 3r\epsilon + (r\epsilon)^2)$, where $r = \|x - y\|_2$. Specifically, the kernels ϕ , ψ , and λ are all chosen to be the C^4 Matérn kernel. In Section 3.5, we present universal approximation results and theoretical convergence rates for the Matérn kernel.

Shape parameters

Like most RBFs, the Matérn kernel comes equipped with a *shape parameter* ϵ which enormously affects accuracy and conditioning of the corresponding kernel interpolants; in our case, we have three *different* shape parameters that must be selected (one corresponding to ϕ , ψ , and λ each). We empirically found that selecting these shape parameters to be as small as possible produced the most accurate generalization results; this corresponds to the so-called flat limit of the kernel. We hypothesize that this was because our training function fields output by the SU2 solver (a finite-volume solver) are naturally consistent with a finite-order Sobolev regularity class [169]. For fixed smoothness ν , a Matérn kernel induces a Sobolev (Bessel-potential) native space, and as the shape parameter goes to zero, its reproducing kernel Hilbert space (RKHS) norm approaches the homogeneous \dot{H}^m seminorm ($m = \nu + \frac{d}{2}$), so the interpolant converges to the polyharmonic (Duchon) spline that minimizes that Sobolev semi-norm among all interpolants [170, 171]; thus, using a near-flat Matérn kernel potentially aligns the operator kernel interpolant with the natural energy class of the finite-volume data.

Algorithms for selecting shape parameters have been studied extensively in the literature. One class of methods depends on the target functions, such as the leave-one-out cross-validation (LOOCV) method [172, 173], which is closely related to Rippa’s algorithm [172, 174, 175]. Another class of methods involves optimizing the shape parameter for a target condition number [176–180]. For convenience, we leverage the latter approach. Specifically, given a kernel Gramian \mathbf{A}_ϵ and a target condition number δ , we find the root of the function $s(\epsilon) = \kappa(\mathbf{A}_\epsilon) - \delta$, where κ is 1-norm condition number of \mathbf{A}_ϵ ; the (reciprocal) 1-norm condition number is much easier to estimate and therefore more computationally inexpensive than the 2-norm condition number. More specifically, we find shape parameters by rootfinding using the popular Brent–Dekker method [181, 182] (as implemented in Matlab’s `fzero` function). We consistently selected a target condition number of 10^{12} for the property-preserving kernel Gramian (or the product kernel in the spacetime case) and 10^{15} for the operator kernel Gramian; we were able to use a larger target condition number for the latter as it also included a ridge regularization term (aka nugget).

Nuggets

Finally, following standard practice in RBF interpolation of noisy data [162, 183], we also regularized the operator kernel by adding a small “nugget”, commonly known as the ridge parameter, θ to the diagonal of the Gramian in (22). Interestingly, we found that adding a ridge parameter to the property-preserving kernel degraded the overall generalization accuracy, whereas it was crucial for the operator kernel. We show results on the effect of varying θ in Appendix C and discuss these results in Section 2.

3.4.2 Efficient Linear Algebra

Each interpolant carries its own implementation challenges: the property-preserving interpolant brought with it computational cost issues due to the $O(d^3m^3)$ cost of the naive linear solve, especially for the large m explored in this work; the operator interpolant brings challenges in terms of GPU memory, and the inability to store all N training functions on GPU memory for large N . We now describe our specialized implementation techniques.

Property-preserving kernel interpolant

We now describe our approach to overcome the $O(d^3m^3)$ cost for finding the property-preserving interpolation coefficients, a necessary preprocessing step to obtaining the training data for the operator kernel interpolant. Our approach is centered around careful and recursive use of the Schur complement of the block matrix Φ (the property-preserving kernel Gramian). In the discussion below, we will focus on Φ ; the same techniques were used for the other variants.

We first consider the case when $d = 2$. Here, Φ is a 2×2 matrix-valued kernel. When evaluated at m points, its Gramian is a matrix of size $2m \times 2m$. The Gramian naturally separates into four blocks, but each block acts on both vector components of the coefficients; recall that each component of the coefficients corresponds to a different spatial dimension (in our case coordinate axes). Following Fuselier et al. [97], we first reorder the individual blocks *spatially*; evaluations for

each spatial component are grouped together and the resulting Gramian has the form

$$\Phi = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \quad A_{ij} \in \mathbb{R}^{m \times m}, \quad (27)$$

where, for notational convenience, A_{ij} is the (i, j) block in (11) evaluated at all m points. Then, the block linear system for interpolation is given by

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \end{bmatrix} = \begin{bmatrix} \mathbf{v}^1 \\ \mathbf{v}^2 \end{bmatrix}. \quad (28)$$

For the symmetric kernels ϕ used in this work, the mixed partial derivatives commute w.r.t. x and y . As a result, Φ exhibits symmetry in its block structure and $A_{12} = A_{21}$. This symmetry allows us to store only the upper (or lower) triangular blocks. Rather than forming and factorizing the full $2m \times 2m$ matrix, we apply a Schur complement reduction with respect to A_{22} . Defining

$$S = A_{11} - A_{12}A_{22}^{-1}A_{12}, \quad (29)$$

the coefficients are obtained with the following two steps:

$$\mathbf{b}^1 = S^{-1}(\mathbf{v}^1 - A_{12}A_{22}^{-1}\mathbf{v}^2), \quad (30)$$

$$\mathbf{b}^2 = A_{22}^{-1}(\mathbf{v}^2 - A_{12}\mathbf{b}^1). \quad (31)$$

This avoids explicitly forming and storing the full Gramian and reduces the memory usage from $\mathcal{O}(m^2d^2)$ to $\mathcal{O}(m^2)$. Additionally, since A_{22} is a positive-definite matrix, we apply a Cholesky factorization $A_{22} = LL^T$, and then perform forward and backward substitution solves using the lower triangular matrix L in (29), (30), and (31). This reduces the overall runtime complexity from $\mathcal{O}(m^3d^3)$ to $\mathcal{O}(m^3)$.

In $d = 3$, after a similar reorder by spatial coordinates, the symmetric block linear system takes the form

$$\begin{bmatrix} A_{11} & A_{12} & | & A_{13} \\ A_{21} & A_{22} & | & A_{23} \\ \hline A_{31} & A_{32} & | & A_{33} \end{bmatrix} \begin{bmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \\ \mathbf{b}^3 \end{bmatrix} = \begin{bmatrix} \mathbf{v}^1 \\ \mathbf{v}^2 \\ \mathbf{v}^3 \end{bmatrix}, \quad (32)$$

where A_{ij} refers to the (i, j) block in (12), $A_{ij} = A_{ji}$, and the horizontal and vertical bars indicate a partition of the Gramian into a 2×2 system. We apply the same Schur complement strategy recursively. First, we eliminate \mathbf{b}^3 via the Schur complement of Φ_{33} , giving the following 2×2 block system

$$\begin{bmatrix} \tilde{A}_{11} & \tilde{A}_{12} \\ \tilde{A}_{21} & \tilde{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{v}}^1 \\ \tilde{\mathbf{v}}^2 \end{bmatrix}, \quad (33)$$

where

$$\begin{aligned} \tilde{A}_{11} &= A_{11} - A_{13}A_{33}^{-1}A_{13}, & \tilde{A}_{12} &= A_{12} - A_{13}A_{33}^{-1}A_{23}, \\ \tilde{A}_{22} &= A_{22} - A_{23}A_{33}^{-1}A_{23}, & \tilde{\mathbf{v}}^1 &= \mathbf{v}^1 - A_{13}A_{33}^{-1}\mathbf{v}^3, \\ & & \tilde{\mathbf{v}}^2 &= \mathbf{v}^2 - A_{23}A_{33}^{-1}\mathbf{v}^3. \end{aligned}$$

(33) is then solved using the same Schur complement procedure described previously, after which \mathbf{b}^3 is recovered as

$$\mathbf{b}^3 = A_{33}^{-1}(\mathbf{v}^3 - A_{13}\mathbf{b}^1 - A_{23}\mathbf{b}^2). \quad (34)$$

We once again obtain an $\mathcal{O}(m^3)$ cost for this recursive solution and require $\mathcal{O}(m^2)$ storage. It is also likely that these systems are better conditioned in practice than the original full Gramian. We also

leverage a version of these algorithms at inference/generalization time: we only store blocks of the evaluation matrix Φ^* and apply them component-wise to the predicted coefficient (block) vectors.

Extension to spacetime kernels: In the spacetime experiment in Section 1.2, the output consists of velocity fields at four discrete timesteps, resulting in (18) having a 4×4 block Gramian. We recursively apply the same Schur complement strategy to reduce the system to first a 3×3 system, then a 2×2 system in turn.

Operator kernel interpolant

The main challenge with the operator kernel interpolant was in forming the $N \times N$ Gramian matrix $\mathbf{K}_{ij} = \lambda(\mathbf{a}_i, \mathbf{a}_j)$ where $\mathbf{a} \in \mathbb{R}^{np}$. Much of our computation is done with the JAX library [184] in Python which provides a useful vectorizable map, `vmap`, that can be customized to execute instructions in parallel on the GPU. Fortunately, forming the entries of \mathbf{K} is an embarrassingly parallelizable task. However, it is not possible to use `vmap` to form the entire matrix in one call due to the cost of storing all the input functions in memory simultaneously, even on a moderately powerful desktop GPU (NVIDIA RTX 4080); `vmap` also appears to incur a significant memory overhead of its own, as do JAX’s just-in-time (JIT) compilation features. In our experiments, n and N can each be as large as 10,000. Consequently, we required an efficient streaming method to compute the entries of \mathbf{K} . We used Algorithm 2 to form $B \times B$ blocks of \mathbf{K} in batches.

Algorithm 2 Streaming construction of the operator kernel Gramian

Require: Evaluations of the input functions $\{\mathbf{a}_i\}_{i=1}^N$, operator kernel λ , and batch size B .

```

1: Initialize  $\mathbf{K} = \mathbf{0}_{N \times N}$ 
2:  $P = \lceil N/B \rceil$ 
3: for  $i = 1$  to  $P$  do
4:    $I_{\text{start}} \leftarrow iB$ 
5:    $I_{\text{end}} \leftarrow \min((i+1)B, N)$ 
6:   for  $j = 1$  to  $P$  do
7:      $J_{\text{start}} \leftarrow jB$ 
8:      $J_{\text{end}} \leftarrow \min((j+1)B, N)$ 
9:      $\mathbf{K}_{I_{\text{start}}:I_{\text{end}}, J_{\text{start}}:J_{\text{end}}} = \lambda(\mathbf{a}_{I_{\text{start}}:I_{\text{end}}}, \mathbf{a}_{J_{\text{start}}:J_{\text{end}}})$  ▷ vmap call here.
10:    end for
11:   end for

```

3.4.3 Training details

We now describe our training method. Though the computation of our interpolation coefficients only requires linear solves, we also compute the shape parameters as described in Section 3.4.1.

Preprocessing

Before we begin training, we first preprocess our inputs to enable more intuitive and comparable selection of the shape parameter and the ridge parameter across all experiments; a discussion of the magnitude of the ridge parameter θ is provided in Appendix C. Specifically, we normalized all spatial functions to each have zero mean and unit standard deviation. Recalling that $\{a_i(x)\}_{i=1}^{N+N_{\text{test}}}$ are the input functions, let $\mu(x)$ be the estimated spatial mean and estimated $\zeta(x)$ be the spatial standard deviation computed from the training set at the set of points X . Both training and test input functions are then normalized as follows: $\frac{a_i(x) - \mu(x)}{\zeta(x)}$, $i = 1, \dots, N + N_{\text{test}}$.

We also conduct another critical preprocessing step: the property-preserving kernel interpolants to the N training output functions must be precomputed, and their coefficients must be stored as they form the right-hand side to the operator kernel interpolant in (22). We conducted this preprocessing step on the CPU for all experiments simply due to the memory constraints of storing the resulting Nmd coefficients; while a streaming approach is possible here also, we did not explore such an approach here as this was simply viewed as a preprocessing step.

Two-step training

We use a simple two-step training procedure. First, we optimize for the operator kernel’s shape parameter as discussed in Section 3.4.1: we used MATLAB’s `fzero` function on the CPU to find the

shape parameter for the target condition number of 10^{15} as it proved to be significantly faster than Python’s rootfinding routines; note that `fzero` (the Brent-Dekker method) is iterative and gradient-free. We then use this shape parameter to reform the Gramian \mathbf{K} and solve the linear system (22) on the GPU via the JAX Python library. We measure the total training time as the time taken for each of these steps. We measure total inference time for our method as the sum of the times taken to form the evaluation matrix \mathbf{K}^* , predict the output coefficients, and reconstruct the velocity fields from those coefficients.

3.5 Approximation Theorems and Convergence Rates

We now present a series of approximation theorems governing the existence and convergence rates of our interpolants. The corollaries establish the same results for cases where the operator kernel interpolant includes ridge terms during training (which is our setting). In all cases, we specialize for the C^4 Matérn kernel, which has a smoothness parameter $\nu = 9/2$.

The first theorem establishes universal approximation of output coefficients via Matérn kernels; we show that they are dense in the space of continuous functions defined over the vector space whose dimension is given by the length of the coefficient vector.

Theorem 1 (Universal Approximation) *Let $X \subset \mathbb{R}^n$ be compact, fix $\nu > 0$, let $k_\nu(x, x') := \kappa_\nu(x - x')$ be the translation-invariant Matérn kernel on \mathbb{R}^n , let $B \in \mathbb{R}^{M \times M}$ be symmetric positive definite, and define the separable matrix-valued kernel $K_X(x, x') := k_\nu(x, x')B$ with associated vector-valued RKHS $\mathcal{H}_{K_X} \subset C(X; \mathbb{R}^M)$. Then \mathcal{H}_{K_X} is dense in $C(X; \mathbb{R}^M)$ with respect to the uniform norm, i.e.*

$$\overline{\mathcal{H}_{K_X}}^{\|\cdot\|_{L^\infty(X; \mathbb{R}^M)}} = C(X; \mathbb{R}^M).$$

Proof Since k_ν is translation-invariant with spectral density strictly positive on \mathbb{R}^n , it is c_0 -universal, hence its RKHS H_{k_ν} is dense in $C(X)$ for every compact $X \subset \mathbb{R}^n$ [185]. For $B > 0$ and $K_X(x, x') = k_\nu(x, x')B$, the associated RKHS satisfies (as a set) $\mathcal{H}_{K_X} = \{B^{1/2}g : g \in H_{k_\nu}(X)^M\}$ with norm equivalence depending only on B [186]. Let $f \in C(X; \mathbb{R}^M)$ and set $g := B^{-1/2}f \in C(X; \mathbb{R}^M)$; by density of $H_{k_\nu}(X)$ in $C(X)$ applied componentwise, there exists $g_m \in H_{k_\nu}(X)^M$ with $\|g_m - g\|_{L^\infty(X; \mathbb{R}^M)} \rightarrow 0$. Then $f_m := B^{1/2}g_m \in \mathcal{H}_{K_X}$ and $\|f_m - f\|_{L^\infty(X; \mathbb{R}^M)} \rightarrow 0$, proving density. \square

The next theorem describes *worst-case* convergence rates for recovering output coefficients. This theorem is highly pessimistic (by no means offering tight upper bounds) and uses the sample dimension of the input function to determine the convergence rate; in practice the true rate is determined by low-rank structure associated with both the input functions and the operator.

Theorem 2 (Coefficient interpolation rate for vector-valued C^4 Matérn kernels in $L^2(X)$) *Let $X \subset \mathbb{R}^n$ be a bounded Lipschitz domain, let k_ν denote the scalar Matérn kernel of smoothness ν , let $B \in \mathbb{R}^{M \times M}$ be symmetric positive definite, define the separable matrix-valued kernel $K_X(x, x') := k_\nu(x, x')B$ with associated RKHS \mathcal{H}_{K_X} , and let $X_N \subset X$ be quasi-uniform with fill distance $h_X := \sup_{x \in X} \min_{x_i \in X_N} \|x - x_i\|$. If $c^\dagger \in H^{\nu+n/2}(X; \mathbb{R}^M)$ and $\widehat{c}_N \in \mathcal{H}_{K_X}$ denotes the K_X -interpolant of c^\dagger on X_N , then*

$$\|\widehat{c}_N - c^\dagger\|_{L^2(X)} \leq C_X h_X^{\nu+n/2} \|c^\dagger\|_{H^{\nu+n/2}(X)}.$$

If X_N is quasi-uniform, then $h_X \asymp N^{-1/n}$ and

$$\|\widehat{c}_N - c^\dagger\|_{L^2(X)} \leq \widetilde{C}_X N^{-(\nu+n/2)/n} \|c^\dagger\|_{H^{\nu+n/2}(X)}.$$

Proof For separable kernels $K_X(x, x') = k_\nu(x, x')B$ with $B > 0$, the associated vector-valued RKHS \mathcal{H}_{K_X} is isomorphic (with norm equivalence constants depending only on B) to the product space $H_{k_\nu}(X)^M$, and under this identification the K_X -interpolant corresponds to componentwise scalar k_ν -interpolation [186]. Since $H_{k_\nu}(X) \simeq H^{\nu+n/2}(X)$ [82, Thm. 10.47], the scalar Sobolev-type estimate in $L^2(X)$ implies

$$\|I_{X_N} g - g\|_{L^2(X)} \leq C h_X^{\nu+n/2} \|g\|_{H^{\nu+n/2}(X)}$$

for $g \in H^{\nu+n/2}(X)$, and applying this componentwise yields

$$\|\widehat{c}_N - c^\dagger\|_{L^2(X)} \leq C_X h_X^{\nu+n/2} \|c^\dagger\|_{H^{\nu+n/2}(X)}.$$

The quasi-uniform rate follows from $h_X \asymp N^{-1/n}$. \square

The following corollary extends this pessimistic worst-case convergence rate to the case where ridge regularization (a nugget) is used in the operator kernel interpolant.

Corollary 3 (Ridge perturbation of coefficient interpolation in $L^2(X)$) Under the assumptions of Theorem 2, let $\widehat{c}_{N,\lambda} \in \mathcal{H}_{K_X}$ denote the kernel ridge regression estimator defined by

$$\widehat{c}_{N,\lambda} \in \arg \min_{f \in \mathcal{H}_{K_X}} \frac{1}{N} \sum_{i=1}^N \|f(x_i) - c^\dagger(x_i)\|_2^2 + \lambda \|f\|_{\mathcal{H}_{K_X}}^2,$$

and let \widehat{c}_N denote the K_X -interpolant of c^\dagger on X_N . Then

$$\|\widehat{c}_{N,\lambda} - c^\dagger\|_{L^2(X)} \leq C_X h_X^{\nu+n/2} \|c^\dagger\|_{H^{\nu+n/2}(X)} + \|\widehat{c}_{N,\lambda} - \widehat{c}_N\|_{L^2(X)}.$$

If X_N is quasi-uniform, then the first term is $\mathcal{O}(N^{-(\nu+n/2)/n})$.

Proof Add and subtract \widehat{c}_N and apply the triangle inequality, then use Theorem 2. \square

Next, we prove a theorem regarding convergence of the operator learning procedure to target functions on generalization. The proof takes into account the error in predicting the coefficients and also error estimates for divergence-free approximation when using exact coefficients; we omit the proofs for the periodic and turbulent cases.

Theorem 4 (Divergence-free Matérn reconstruction error in $L^2(\Omega)$) Let $\Omega \subset \mathbb{R}^d$ be a bounded Lipschitz domain, let ϕ_ν denote the scalar Matérn kernel on \mathbb{R}^d , define the divergence-free matrix-valued kernel $\Phi_{\text{div}}(x) := (-\Delta I + \nabla \nabla^\top) \phi_\nu(x)$, fix quasi-uniform sites $Y_M \subset \Omega$ with fill distance $h_Y := \sup_{y \in \Omega} \min_{y_j \in Y_M} \|y - y_j\|$, fix vectors $b_j \in \mathbb{R}^d$, define basis functions $\psi_j(x) := \Phi_{\text{div}}(x - y_j) b_j$ and reconstruction operator $(\mathcal{R}c)(x) := \sum_{j=1}^M c_j \psi_j(x)$, and define the $L^2(\Omega)$ mass matrix $\mathbf{T}_{ij} := \langle \psi_i, \psi_j \rangle_{L^2(\Omega)}$. If $v \in H^{\nu+d/2}(\Omega; \mathbb{R}^d)$ is divergence-free and $v_M = \mathcal{R}c^*$ denotes its divergence-free Matérn interpolant on Y_M , then for any $\widehat{c} \in \mathbb{R}^M$,

$$\|v - \mathcal{R}\widehat{c}\|_{L^2(\Omega)} \leq C_Y h_Y^{\nu+d/2} \|v\|_{H^{\nu+d/2}(\Omega)} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c} - c^*\|_2.$$

If Y_M is quasi-uniform, then $h_Y \asymp M^{-1/d}$ and

$$\|v - \mathcal{R}\widehat{c}\|_{L^2(\Omega)} \leq \widetilde{C}_Y M^{-(\nu+d/2)/d} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c} - c^*\|_2.$$

Proof The divergence-free construction admits Sobolev-type approximation rates [93], yielding

$$\|v - v_M\|_{L^2(\Omega)} \leq C_Y h_Y^{\nu+d/2} \|v\|_{H^{\nu+d/2}(\Omega)}.$$

Moreover,

$$\|\mathcal{R}(c^* - \widehat{c})\|_{L^2(\Omega)}^2 = (c^* - \widehat{c})^\top \mathbf{T} (c^* - \widehat{c}) \leq \lambda_{\max}(\mathbf{T}) \|\widehat{c} - c^*\|_2^2.$$

By the triangle inequality,

$$\|v - \mathcal{R}\widehat{c}\|_{L^2(\Omega)} = \|(v - v_M) + \mathcal{R}(c^* - \widehat{c})\|_{L^2(\Omega)} \leq \|v - v_M\|_{L^2(\Omega)} + \|\mathcal{R}(c^* - \widehat{c})\|_{L^2(\Omega)},$$

which yields the result. \square

The following corollary modifies the error estimate from Theorem 6 to account for the error introduced by ridge regression.

Corollary 5 (Ridge operator bound in $L^2(X; L^2(\Omega))$) Under the assumptions of Theorem 4, replace \widehat{c}_N by the kernel ridge regression estimator $\widehat{c}_{N,\lambda}$ defined in the previous corollary and define $\widehat{\mathcal{G}}_{N,M,\lambda}(x) := \mathcal{R}(\widehat{c}_{N,\lambda}(x))$. Then

$$\|\widehat{\mathcal{G}}_{N,M,\lambda} - \mathcal{G}\|_{L^2(X; L^2(\Omega))} \leq \widetilde{C}_Y M^{-(\nu+d/2)/d} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c}_{N,\lambda} - c^\dagger\|_{L^2(X)}.$$

In particular,

$$\|\widehat{\mathcal{G}}_{N,M,\lambda} - \mathcal{G}\|_{L^2(X; L^2(\Omega))} \leq \widetilde{C}_Y M^{-(\nu+d/2)/d} + \widetilde{C}_X \sqrt{\lambda_{\max}(\mathbf{T})} N^{-(\nu+n/2)/n} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c}_{N,\lambda} - \widehat{c}_N\|_{L^2(X)}.$$

Proof Repeat the proof of Theorem 6 with \widehat{c}_N replaced by $\widehat{c}_{N,\lambda}$, then substitute the previous corollary. \square

Theorem 6 (below) now combines the results of the previous theorems in order to bound the operator error. Once again, this is a pessimistic (not sharp) worst-case error estimate, working in the sampling dimension of the input functions rather than the unknown intrinsic dimension. In general, operator learning techniques (kernel or otherwise) appear to converge faster. Nevertheless, Theorem 6 gives us a baseline result that we hope to improve upon.

Theorem 6 (Operator error in $L^2(X; L^2(\Omega))$) Under the assumptions of the previous two theorems, if $v(\cdot; x) \in H^{\nu+d/2}(\Omega; \mathbb{R}^d)$ satisfies $\|v\|_{L^\infty(X; H^{\nu+d/2}(\Omega))} := \sup_{x \in X} \|v(\cdot; x)\|_{H^{\nu+d/2}(\Omega)} < \infty$, if $c^\dagger(x)$ is defined by $v_M(\cdot; x) = \mathcal{R}(c^\dagger(x))$ where $v_M(\cdot; x)$ denotes the divergence-free Matérn interpolant of $v(\cdot; x)$ on Y_M , and if \widehat{c}_N denotes the K_X -interpolant of c^\dagger on X_N , then defining $\widehat{\mathcal{G}}_{N,M}(x) := \mathcal{R}(\widehat{c}_N(x))$ and $\mathcal{G}(x) := v(\cdot; x)$ yields

$$\|\widehat{\mathcal{G}}_{N,M} - \mathcal{G}\|_{L^2(X; L^2(\Omega))} \leq \widetilde{C}_Y M^{-(\nu+d/2)/d} + \widetilde{C}_X \sqrt{\lambda_{\max}(\mathbf{T})} N^{-(\nu+n/2)/n}.$$

Proof For each $x \in X$, apply the divergence-free reconstruction estimate with $v = v(\cdot; x)$, $c^* = c^\dagger(x)$, and $\widehat{c} = \widehat{c}_N(x)$ to obtain

$$\|v(\cdot; x) - \mathcal{R}(\widehat{c}_N(x))\|_{L^2(\Omega)} \leq C_Y h_Y^{\nu+d/2} \|v(\cdot; x)\|_{H^{\nu+d/2}(\Omega)} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c}_N(x) - c^\dagger(x)\|_2.$$

Taking the $L^2(X)$ norm in x and using $\|v(\cdot; \cdot)\|_{L^\infty(X; H^{\nu+d/2}(\Omega))} < \infty$ gives

$$\|\widehat{\mathcal{G}}_{N,M} - \mathcal{G}\|_{L^2(X; L^2(\Omega))} \leq C_Y h_Y^{\nu+d/2} |X|^{1/2} \|v\|_{L^\infty(X; H^{\nu+d/2}(\Omega))} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c}_N - c^\dagger\|_{L^2(X)}.$$

Substituting the coefficient estimate from Theorem 2 and $h_Y \asymp M^{-1/d}$ yields the claim (absorbing $|X|^{1/2} \|v\|_{L^\infty(X; H^{\nu+d/2})}$ into \widetilde{C}_Y). \square

We now present two theorems that clearly show error estimates in terms of the unknown intrinsic dimension. The result shows the same algebraic decay in the fill distance, albeit when that fill distance is measured within that intrinsic dimension. If the intrinsic dimension $r \ll n$ (in the notation of this section), then this implies a much faster rate of approximation in practice. The first theorem is for recovering coefficients.

Theorem 7 (Intrinsic-dimension variant of coefficient interpolation in $L^2(X_{\text{eff}})$) *Let $X \subset \mathbb{R}^n$ be a bounded Lipschitz domain, let k_ν be the scalar Matérn kernel of smoothness ν , let $B \in \mathbb{R}^{M \times M}$ be symmetric positive definite, define $K_X(x, x') := k_\nu(x, x')B$ with RKHS \mathcal{H}_{K_X} , and assume there exist an integer $r \ll n$, a bounded Lipschitz domain $\widetilde{X} \subset \mathbb{R}^r$, and a bi-Lipschitz bijection $\Phi : \widetilde{X} \rightarrow X_{\text{eff}} \subset X$ such that $X_N = \Phi(\widetilde{X}_N)$ for a quasi-uniform set $\widetilde{X}_N \subset \widetilde{X}$ with fill distance $h_{\widetilde{X}} := \sup_{\xi \in \widetilde{X}} \min_{\xi_i \in \widetilde{X}_N} \|\xi - \xi_i\|$. If c^\dagger factors through Φ as $c^\dagger(x) = \widetilde{c}(\Phi^{-1}(x))$ for $x \in X_{\text{eff}}$ with $\widetilde{c} \in H^{\nu+r/2}(\widetilde{X}; \mathbb{R}^M)$ and \widehat{c}_N denotes the K_X -interpolant of c^\dagger on X_N , then*

$$\|\widehat{c}_N - c^\dagger\|_{L^2(X_{\text{eff}})} \leq C_\Phi h_{\widetilde{X}}^{\nu+r/2} \|\widetilde{c}\|_{H^{\nu+r/2}(\widetilde{X})}.$$

If \widetilde{X}_N is quasi-uniform, then $h_{\widetilde{X}} \asymp N^{-1/r}$ and

$$\|\widehat{c}_N - c^\dagger\|_{L^2(X_{\text{eff}})} \leq \widetilde{C}_\Phi N^{-(\nu+r/2)/r} \|\widetilde{c}\|_{H^{\nu+r/2}(\widetilde{X})}.$$

Proof Define $\widetilde{c} : \widetilde{X} \rightarrow \mathbb{R}^M$ by $\widetilde{c}(\xi) := c^\dagger(\Phi(\xi))$ and define $\widetilde{X}_N := \Phi^{-1}(X_N)$. Since Φ is bi-Lipschitz, the pullback identifies $L^2(X_{\text{eff}})$ with $L^2(\widetilde{X})$ up to a constant depending only on Φ , and likewise $H^{\nu+r/2}(X_{\text{eff}})$ with $H^{\nu+r/2}(\widetilde{X})$ up to a constant depending only on Φ . Moreover, because $X_N = \Phi(\widetilde{X}_N)$, the K_X -interpolant of c^\dagger restricted to X_{eff} corresponds, under pullback by Φ , to Matérn interpolation of \widetilde{c} on the r -dimensional set \widetilde{X}_N . Applying the L^2 Matérn/Sobolev interpolation estimate on \widetilde{X} yields

$$\|\widehat{c}_N - c^\dagger\|_{L^2(X_{\text{eff}})} \leq C_\Phi h_{\widetilde{X}}^{\nu+r/2} \|\widetilde{c}\|_{H^{\nu+r/2}(\widetilde{X})}.$$

If \widetilde{X}_N is quasi-uniform then $h_{\widetilde{X}} \asymp N^{-1/r}$, giving the stated algebraic rate. \square

Finally, our last theorem establishes operator convergence rates in terms of the unknown intrinsic dimension $r \ll n$.

Theorem 8 (Intrinsic-dimension operator bound in $L^2(X_{\text{eff}}; L^2(\Omega))$) *Under the assumptions of the divergence-free reconstruction theorem and the intrinsic-dimension coefficient theorem, if $c^\dagger(x)$ is defined by $v_M(\cdot; x) = \mathcal{R}(c^\dagger(x))$ where $v_M(\cdot; x)$ denotes the divergence-free Matérn interpolant of $v(\cdot; x)$ on Y_M , then defining $\widehat{\mathcal{G}}_{N,M}(x) := \mathcal{R}(\widehat{c}_N(x))$ and $\mathcal{G}(x) := v(\cdot; x)$ yields*

$$\|\widehat{\mathcal{G}}_{N,M} - \mathcal{G}\|_{L^2(X_{\text{eff}}; L^2(\Omega))} \leq \widetilde{C}_Y M^{-(\nu+d/2)/d} + \widetilde{C}_\Phi \sqrt{\lambda_{\max}(\mathbf{T})} N^{-(\nu+r/2)/r}.$$

Proof For each $x \in X_{\text{eff}}$, apply the divergence-free reconstruction theorem with $v = v(\cdot; x)$, $c^* = c^\dagger(x)$ (defined by $v_M(\cdot; x) = \mathcal{R}(c^\dagger(x))$), and $\widehat{c} = \widehat{c}_N(x)$ to obtain

$$\|v(\cdot; x) - \mathcal{R}(\widehat{c}_N(x))\|_{L^2(\Omega)} \leq C_Y h_Y^{\nu+d/2} \|v(\cdot; x)\|_{H^{\nu+d/2}(\Omega)} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c}_N(x) - c^\dagger(x)\|_2.$$

Take the $L^2(X_{\text{eff}})$ norm in x and use $\|v\|_{L^\infty(X_{\text{eff}}; H^{\nu+d/2}(\Omega))} < \infty$ to get

$$\|\widehat{\mathcal{G}}_{N,M} - \mathcal{G}\|_{L^2(X_{\text{eff}}; L^2(\Omega))} \leq C_Y h_Y^{\nu+d/2} |X_{\text{eff}}|^{1/2} \|v\|_{L^\infty(X_{\text{eff}}; H^{\nu+d/2}(\Omega))} + \sqrt{\lambda_{\max}(\mathbf{T})} \|\widehat{c}_N - c^\dagger\|_{L^2(X_{\text{eff}})}.$$

Substitute the intrinsic-dimension coefficient estimate from Theorem 7 and $h_Y \asymp M^{-1/d}$, absorbing $|X_{\text{eff}}|^{1/2} \|v\|_{L^\infty(X_{\text{eff}}; H^{\nu+d/2})}$ into \widetilde{C}_Y . \square

References

- [1] Thwaites, B., Street, R.: Incompressible aerodynamics: an account of the theory and observation of the steady flow of incompressible fluid past aerofoils, wings, and other bodies. *Physics Today* **13**(12), 60–61 (1960)
- [2] Kwak, D., Chang, J.L., Shanks, S.P., Chakravarthy, S.R.: A three-dimensional incompressible navier-stokes flow solver using primitive variables. *AIAA journal* **24**(3), 390–396 (1986)
- [3] Chang, J., Kwak, D., Rogers, S., Yang, R.-J.: Numerical simulation methods of incompressible flows and an application to the space shuttle main engine. *International Journal for numerical methods in fluids* **8**(10), 1241–1268 (1988)
- [4] Kwak, D., Kiris, C.: Cfd for incompressible flows at nasa ames. *Computers & fluids* **38**(3), 504–510 (2009)
- [5] Castorrini, A., Gentile, S., Geraldi, E., Bonfiglioli, A.: Investigations on offshore wind turbine inflow modelling using numerical weather prediction coupled with local-scale computational fluid dynamics. *Renewable and Sustainable Energy Reviews* **171**, 113008 (2023)
- [6] Najm, H.N., Knio, O.M.: Modeling low mach number reacting flow with detailed chemistry and transport. *Journal of Scientific Computing* **25**(1), 263–287 (2005)
- [7] Janela, J., Moura, A., Sequeira, A.: A 3d non-newtonian fluid–structure interaction model for blood flow in arteries. *Journal of Computational and Applied Mathematics* **234**(9), 2783–2791 (2010)
- [8] Canic, S., Mikelic, A.: Effective equations modeling the flow of a viscous incompressible fluid through a long elastic tube arising in the study of blood flow through small arteries. *SIAM Journal on Applied Dynamical Systems* **2**(3), 431–463 (2003)
- [9] Secomb, T.W.: Hemodynamics. *Comprehensive physiology* **6**(2), 975–1003 (2016)
- [10] Womersley, J.R.: Method for the calculation of velocity, rate of flow and viscous drag in arteries when the pressure gradient is known. *The Journal of physiology* **127**(3), 553 (1955)
- [11] McDonald, D.: The relation of pulsatile pressure to flow in arteries. *The Journal of physiology* **127**(3), 533 (1955)
- [12] Perktold, K., Rappitsch, G.: Computer simulation of local blood flow and vessel mechanics in a compliant carotid artery bifurcation model. *Journal of biomechanics* **28**(7), 845–856 (1995)
- [13] Orszag, S.A., Israeli, M., Deville, M.O.: Boundary conditions for incompressible flows. *Journal of Scientific Computing* **1**(1), 75–111 (1986)
- [14] Karniadakis, G.E., Israeli, M., Orszag, S.A.: High-order splitting methods for the incompressible navier-stokes equations. *Journal of computational physics* **97**(2), 414–443 (1991)
- [15] John, V.: Higher order finite element methods and multigrid solvers in a benchmark problem for the 3d navier–stokes equations. *International Journal for Numerical Methods in Fluids* **40**(6), 775–798 (2002)
- [16] Ahmed, N., Bartsch, C., John, V., Wilbrandt, U.: An assessment of solvers for saddle point problems emerging from the incompressible navier–stokes equations. *Computational Methods in Applied Mechanics and Engineering* **331**, 492–513 (2018). Based on WIAS Preprint 2408 (2017)
- [17] Johnston, H., Liu, J.-G.: Accurate, stable and efficient navier–stokes solvers based on explicit treatment of the pressure term. *Journal of Computational Physics* **199**(1), 221–259 (2004)

[18] Rosales, R.R., Seibold, B., Shirokoff, D., Zhou, D.: High-order finite element methods for a pressure poisson equation reformulation of the navier–stokes equations with electric boundary conditions. *Computer Methods in Applied Mechanics and Engineering* **373**, 113451 (2021)

[19] Chorin, A.J.: A numerical method for solving incompressible viscous flow problems. *Journal of computational physics* **2**(1), 12–26 (1967)

[20] Guermond, J.-L., Minev, P.: High-order time stepping for the incompressible navier–stokes equations. *SIAM Journal on Scientific Computing* **37**(6), 2656–2681 (2015)

[21] Chorin, A.J., Marsden, J.E.: A Mathematical Introduction to Fluid Mechanics. Texts in Applied Mathematics, vol. 4. Springer, New York, NY (1993)

[22] Brown, D.L., Cortez, R., Minion, M.L.: Accurate projection methods for the incompressible navier–stokes equations. *Journal of computational physics* **168**(2), 464–499 (2001)

[23] Goda, K.: A multistep technique with implicit difference schemes for calculating two-or three-dimensional cavity flows. *Journal of computational physics* **30**(1), 76–95 (1979)

[24] Kan, J.: A second-order accurate pressure-correction scheme for viscous incompressible flow. *SIAM Journal on Scientific and Statistical Computing* **7**(3), 870–891 (1986)

[25] Bell, J.B., Colella, P., Glaz, H.M.: A second-order projection method for the incompressible navier–stokes equations. *Journal of Computational Physics* **85**(2), 257–283 (1989)

[26] Guy, R.D., Fogelson, A.L.: Stability of approximate projection methods on cell-centered grids. *Journal of Computational Physics* **203**(2), 517–538 (2005)

[27] Guermond, J.-L., Minev, P., Shen, J.: An overview of projection methods for incompressible flows. *Computer methods in applied mechanics and engineering* **195**(44-47), 6011–6045 (2006)

[28] Elman, H.C.: Preconditioners for saddle point problems arising in computational fluid dynamics. *Applied Numerical Mathematics* **43**(1-2), 75–89 (2002)

[29] Jung, H.-R., Kim, S.-T., Noh, J., Hong, J.-M.: A heterogeneous cpu–gpu parallel approach to a multigrid poisson solver for incompressible fluid simulation. *Computer Animation and Virtual Worlds* **24**(3-4), 185–193 (2013)

[30] Dick, C., Rogowsky, M., Westermann, R.: Solving the fluid pressure poisson equation using multigrid—evaluation and improvements. *IEEE transactions on visualization and computer graphics* **22**(11), 2480–2492 (2015)

[31] Guermond, J.-L., Minev, P.D.: High-order time stepping for the navier–stokes equations with minimal computational complexity. *Journal of Computational and Applied Mathematics* **310**, 92–103 (2017)

[32] Kress, W., Lötstedt, P.: Time step restrictions using semi-explicit methods for the incompressible navier–stokes equations. *Computer methods in applied mechanics and engineering* **195**(33-36), 4433–4447 (2006)

[33] Ionita, A.C., Antoulas, A.C.: Data-driven parametrized model reduction in the loewner framework. *SIAM Journal on Scientific Computing* **36**(3), 984–1007 (2014)

[34] Ohlberger, M., Rave, S.: Reduced basis methods: Success, limitations and future challenges. *arXiv preprint arXiv:1511.02021* (2015)

[35] Tezzele, M., Demo, N., Stabile, G., Mola, A., Rozza, G.: Enhancing cfd predictions in shape design problems by model and parameter space reduction. *Advanced Modeling and Simulation in Engineering Sciences* **7**(1), 40 (2020)

[36] Bhattacharya, K., Hosseini, B., Kovachki, N.B., Stuart, A.M.: Model reduction and neural

networks for parametric pdes. *The SMAI journal of computational mathematics* **7**, 121–157 (2021)

- [37] Holmes, P., Lumley, J.L., Berkooz, G., Rowley, C.W. (eds.): *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, Cambridge, UK; New York, NY, USA (2012)
- [38] Quarteroni, A., Manzoni, A., Negri, F.: *Reduced Basis Methods for Partial Differential Equations: An Introduction* vol. 92. Springer, Cham, Switzerland (2016)
- [39] Lee, K., Carlberg, K.T.: Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders. *Journal of Computational Physics* **404**, 108973 (2020)
- [40] Fresca, S., Manzoni, A., Dedè, L., Quarteroni, A.: Deep learning-based reduced order models in cardiac electrophysiology. *PLoS one* **15**(10), 0239416 (2020)
- [41] Sirignano, J., Spiliopoulos, K.: Dgm: A deep learning algorithm for solving partial differential equations. *Journal of computational physics* **375**, 1339–1364 (2018)
- [42] Thuerey, N., Weissenow, K., Prantl, L., Hu, X.: Deep learning methods for reynolds-averaged navier–stokes simulations of airfoil flows. *AIAA journal* **58**(1), 25–36 (2020)
- [43] Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., Battaglia, P.: Learning to simulate complex physics with graph networks. In: *International Conference on Machine Learning*, pp. 8459–8468 (2020). PMLR
- [44] Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., Battaglia, P.: Learning mesh-based simulation with graph networks. In: *International Conference on Learning Representations* (2020)
- [45] Raissi, M., Perdikaris, P., Karniadakis, G.E.: 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)
- [46] Boullé, N., Townsend, A.: A mathematical guide to operator learning. In: *Handbook of Numerical Analysis* vol. 25, pp. 83–125. Elsevier, Amsterdam, The Netherlands (2024)
- [47] Lu, L., Jin, P., Pang, G., Zhang, Z., Karniadakis, G.E.: Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence* **3**(3), 218–229 (2021)
- [48] Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895* (2020)
- [49] Lowery, M., Turnage, J., Morrow, Z., Jakeman, J.D., Narayan, A., Zhe, S., Shankar, V.: Kernel neural operators (knos) for scalable, memory-efficient, geometrically-flexible operator learning. *arXiv preprint arXiv:2407.00809* (2024)
- [50] Gin, C.R., Shea, D.E., Brunton, S.L., Kutz, J.N.: Deepgreen: Deep learning of green’s functions for nonlinear boundary value problems (2020)
- [51] Boullé, N., Earls, C.J., Townsend, A.: Data-driven discovery of green’s functions with human-understandable deep learning. *Scientific reports* **12**(1), 4824 (2022)
- [52] Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Neural operator: Graph kernel network for partial differential equations. *arXiv preprint arXiv:2003.03485* (2020)
- [53] Hao, Z., Wang, Z., Su, H., Ying, C., Dong, Y., Liu, S., Cheng, Z., Song, J., Zhu, J.: GNOT: A General Neural Operator Transformer for Operator Learning (2023)

[54] Liu, Q., Zhong, W., Meidani, H., Abueidda, D., Koric, S., Geubelle, P.: Geometry-informed neural operator transformer. arXiv preprint arXiv:2504.19452 (2025)

[55] Liu, X., Xu, B., Cao, S., Zhang, L.: Mitigating spectral bias for the multiscale operator learning. *Journal of Computational Physics* **506**, 112944 (2024)

[56] Batlle, P., Darcy, M., Hosseini, B., Owhadi, H.: Kernel methods are competitive for operator learning. *Journal of Computational Physics* **496**, 112549 (2024)

[57] Mora, C., Yousefpour, A., Hosseini, S., Owhadi, H., Bostanabad, R.: Operator learning with gaussian processes. *Computer Methods in Applied Mechanics and Engineering* **434**, 117581 (2025)

[58] Kadri, H., Preux, P., Duflos, E., Canu, S.: Operator-valued kernels for nonparametric operator estimation. PhD thesis, INRIA (2011)

[59] Kadri, H., Duflos, E., Preux, P., Canu, S., Rakotomamonjy, A., Audiffren, J.: Operator-valued kernels for learning from functional response data. *Journal of Machine Learning Research* **17**(20), 1–54 (2016)

[60] Goswami, S., Bora, A., Yu, Y., Karniadakis, G.: Physics-informed deep neural operator networks (2022)

[61] Li, W., Bazant, M.Z., Zhu, J.: Phase-field deeponet: Physics-informed deep operator neural network for fast simulations of pattern formation governed by gradient flows of free-energy functionals. *Computer Methods in Applied Mechanics and Engineering* **416**, 116299 (2023)

[62] Li, Z., Zheng, H., Kovachki, N., Jin, D., Chen, H., Liu, B., Azizzadenesheli, K., Anandkumar, A.: Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science* **1**(3), 1–27 (2024)

[63] Zhang, T., Zhai, W., Yann, R., Gao, J., Cao, H., Xing, X.: Floating-body hydrodynamic neural networks. arXiv preprint arXiv:2509.13783 (2025)

[64] Richter-Powell, J., Lipman, Y., Chen, R.T.: Neural conservation laws: A divergence-free perspective. *Advances in Neural Information Processing Systems* **35**, 38075–38088 (2022)

[65] Khorrami, M.S., Goyal, P., Mianroodi, J.R., Svendsen, B., Benner, P., Raabe, D.: A physics-encoded fourier neural operator approach for surrogate modeling of divergence-free stress fields in solids. arXiv preprint arXiv:2408.15408 (2024)

[66] Li, Z., Huang, D.Z., Liu, B., Anandkumar, A.: Fourier neural operator with learned deformations for pdes on general geometries. *Journal of Machine Learning Research* **24**(388), 1–26 (2023)

[67] Wu, H., Luo, H., Wang, H., Wang, J., Long, M.: Transolver: A fast transformer solver for pdes on general geometries. arXiv preprint arXiv:2402.02366 (2024)

[68] Economou, T.D., Palacios, F., Copeland, S.R., Lukaczyk, T.W., Alonso, J.J.: Su2: An open-source suite for multiphysics simulation and design. *Aiaa Journal* **54**(3), 828–846 (2016)

[69] Economou, T.D.: Simulation and adjoint-based design for variable density incompressible flows with heat transfer. *AIAA Journal* **58**(2), 757–769 (2020)

[70] Geuzaine, C., Remacle, J.-F.: Gmsh: A 3-d finite element mesh generator with built-in pre-and post-processing facilities. *International journal for numerical methods in engineering* **79**(11), 1309–1331 (2009)

[71] Tritton, D.J.: Experiments on the flow past a circular cylinder at low reynolds numbers. *Journal of Fluid Mechanics* **6**(4), 547–567 (1959)

[72] Jackson, C.: A finite-element study of the onset of vortex shedding in flow past variously shaped bodies. *Journal of fluid Mechanics* **182**, 23–45 (1987)

[73] Li, J., Chambarel, A., Donneaud, M., Martin, R.: Numerical study of laminar flow past one and two circular cylinders. *Computers & fluids* **19**(2), 155–170 (1991)

[74] Rajani, B., Kandasamy, A., Majumdar, S.: Numerical simulation of laminar flow past a circular cylinder. *Applied Mathematical Modelling* **33**(3), 1228–1247 (2009)

[75] Provansal, M., Mathis, C., Boyer, L.: Bénard-von kármán instability: transient and forced regimes. *Journal of Fluid Mechanics* **182**, 1–22 (1987)

[76] Taylor, G.I., Green, A.E.: Mechanism of the production of small eddies from large ones. *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences* **158**(895), 499–521 (1937)

[77] Kim, J., Moin, P.: Application of a fractional-step method to incompressible navier-stokes equations. *Journal of computational physics* **59**(2), 308–323 (1985)

[78] Ubal, C.M., Beishuizen, N., Kusch, L., Oijen, J.: Adjoint-based design optimization of a kenics static mixer. *Results in Engineering* **21**, 101856 (2024)

[79] Ladson, C.L., Cuyler W. Brooks, J.: Development of a computer program to obtain ordinates for naca 4-digit, 4-digit modified, 5-digit and 16 series airfoils. NASA Technical Memorandum NASA-TM-X-3284, National Aeronautics and Space Administration, Washington, DC (1975)

[80] Rokhlin, V.: Rapid solution of integral equations of classical potential theory. *Journal of computational physics* **60**(2), 187–207 (1985)

[81] Yokota, R., Barba, L.A.: Treecode and fast multipole method for n-body simulation with cuda. In: GPU Computing Gems Emerald Edition, pp. 113–132. Elsevier, Boston, MA (2011)

[82] Wendland, H.: Scattered Data Approximation. Cambridge University Press, Cambridge, UK (2005)

[83] Drineas, P., Mahoney, M.W.: Approximating a gram matrix for improved kernel-based learning. In: International Conference on Computational Learning Theory, pp. 323–337 (2005). Springer

[84] Gittens, A., Mahoney, M.: Revisiting the nyström method for improved large-scale machine learning. In: International Conference on Machine Learning, pp. 567–575 (2013). PMLR

[85] Chen, C., Hon, Y., Schaback, R.: Scientific computing with radial basis functions. Department of Mathematics, University of Southern Mississippi, Hattiesburg, MS **39406** (2005)

[86] Fasshauer, G.E.: Meshfree Approximation Methods with Matlab (With Cd-rom) vol. 6. World Scientific Publishing Company, Singapore; Hackensack, NJ, USA (2007)

[87] Baxter, B.: The interpolation theory of radial basis functions. arXiv preprint arXiv:1006.2443 (2010)

[88] Raviart, P.A., Thomas, J.M.: A mixed finite element method for 2-nd order elliptic problems. In: Galligani, I., Magenes, E. (eds.) Mathematical Aspects of Finite Element Methods, pp. 292–315. Springer, Berlin, Heidelberg (1977)

[89] Brezzi, F., Douglas Jr, J., Marini, L.D.: Two families of mixed finite elements for second order elliptic problems. *Numerische Mathematik* **47**(2), 217–235 (1985)

[90] Nédélec, J.-C.: Mixed finite elements in \mathbb{R}^3 . *Numerische Mathematik* **35**(3), 315–341 (1980)

[91] Narcowich, F.J., Ward, J.D.: Generalized hermite interpolation via matrix-valued conditionally positive definite functions. *Mathematics of Computation* **63**(208), 661–687 (1994)

[92] Narcowich, F.J., Ward, J.D., Wright, G.B.: Divergence-free rbfs on surfaces. *Journal of Fourier Analysis and Applications* **13**(6), 643–663 (2007)

[93] Fuselier, E.: Sobolev-type approximation rates for divergence-free and curl-free rbf interpolants. *Mathematics of Computation* **77**(263), 1407–1423 (2008)

[94] Fuselier, E., Narcowich, F., Ward, J., Wright, G.: Error and stability estimates for surface-divergence free rbf interpolants on the sphere. *Mathematics of Computation* **78**(268), 2157–2186 (2009)

[95] Wendland, H.: Divergence-free kernel methods for approximating the stokes problem. *SIAM Journal on Numerical Analysis* **47**(4), 3158–3179 (2009)

[96] Fuselier, E.J., Wright, G.B.: A radial basis function method for computing helmholtz–hodge decompositions. *IMA Journal of Numerical Analysis* **37**(2), 774–797 (2017)

[97] Fuselier, E.J., Shankar, V., Wright, G.B.: A high-order radial basis function (rbf) leray projection method for the solution of the incompressible unsteady stokes equations. *Computers & Fluids* **128**, 41–52 (2016)

[98] Drake, K.P., Fuselier, E.J., Wright, G.B.: A partition of unity method for divergence-free or curl-free radial basis function approximation. *SIAM Journal on Scientific Computing* **43**(3), 1950–1974 (2021)

[99] Owhadi, H.: Gaussian process hydrodynamics. *Applied Mathematics and Mechanics* **44**(7), 1175–1198 (2023)

[100] Jacob, M., Blu, T., Unser, M.: Sampling of periodic signals: A quantitative error analysis. *IEEE Transactions on Signal Processing* **50**(5), 1153–1159 (2002)

[101] Nguyen, V.-D., Béchet, E., Geuzaine, C., Noels, L.: Imposing periodic boundary condition on arbitrary meshes by polynomial interpolation. *Computational Materials Science* **55**, 390–406 (2012)

[102] Boissonnat, J., Teillaud, M. (eds.): *Effective Computational Geometry for Curves and Surfaces*, 1st edn. *Mathematics and Visualization*. Springer, Berlin, Heidelberg (2006)

[103] Canuto, C., Hussaini, M.Y., Quarteroni, A., Zang, T.A.: *Spectral Methods: Fundamentals in Single Domains*. *Scientific Computation*. Springer, Berlin, Heidelberg (2006)

[104] Boyd, J.P.: *Chebyshev and Fourier Spectral Methods*. Courier Corporation, Berlin, Heidelberg (2001)

[105] Fornberg, B.: *A Practical Guide to Pseudospectral Methods*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge, United Kingdom (1996)

[106] Xiao, J., Boyd, J.P.: Periodized radial basis functions, part i: theory. *Applied Numerical Mathematics* **86**, 43–73 (2014)

[107] Cools, R., Kuo, F.Y., Nuyens, D., Sloan, I.H.: Lattice algorithms for multivariate approximation in periodic spaces with general weight parameters. *arXiv preprint arXiv:1910.06604* (2019)

[108] Kaarnioja, V., Kazashi, Y., Kuo, F.Y., Nobile, F., Sloan, I.H.: Fast approximation by periodic kernel-based lattice-point interpolation with application in uncertainty quantification. *Numerische Mathematik* **150**(1), 33–77 (2022)

[109] Flyer, N., Wright, G.B.: Transport schemes on a sphere using radial basis functions. *Journal of Computational Physics* **226**(1), 1059–1084 (2007)

[110] Flyer, N., Wright, G.B.: A radial basis function method for the shallow water equations on a sphere. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **465**(2106), 1949–1976 (2009)

[111] Shankar, V., Wright, G.B., Kirby, R.M., Fogelson, A.L.: Augmenting the immersed boundary method with radial basis functions (rbfs) for the modeling of platelets in hemodynamic flows. *International Journal for Numerical Methods in Fluids* **79**(10), 536–557 (2015)

[112] Shankar, V., Olson, S.D.: Radial basis function (rbf)-based parametric models for closed and open curves within the method of regularized stokeslets. *International Journal for Numerical Methods in Fluids* **79**(6), 269–289 (2015)

[113] Kassen, A., Barrett, A., Shankar, V., Fogelson, A.L.: Immersed boundary simulations of cell-cell interactions in whole blood. *Journal of Computational Physics* **469**, 111499 (2022)

[114] Fuselier, E., Wright, G.B.: Order-preserving derivative approximation with periodic radial basis functions. *Advances in Computational Mathematics* **41**(1), 23–53 (2015)

[115] Frisch, U.: *Turbulence: The Legacy of A.N. Kolmogorov*. Cambridge University Press, Cambridge, UK (1995)

[116] Pope, S.B.: *Turbulent Flows*. Cambridge University Press, Cambridge, UK (2000)

[117] Stull, R.B.: *An Introduction to Boundary Layer Meteorology*. Kluwer Academic Publishers, Dordrecht, The Netherlands (1988)

[118] Vallis, G.K.: *Atmospheric and Oceanic Fluid Dynamics*, 2nd edn. Cambridge University Press, Cambridge, UK (2017)

[119] Nezu, I., Nakagawa, H.: Turbulence in open-channel flows. *IAHR Monograph Series* (2005)

[120] Schlichting, H., Gersten, K.: *Boundary-Layer Theory*, 8th edn. Springer, Berlin, Heidelberg (2000)

[121] Kolmogorov, A.N.: The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Doklady Akademii Nauk SSSR* **30**, 301–305 (1941)

[122] Alfonsi, G.: Reynolds-averaged navier–stokes equations for turbulence modeling. *Applied Mechanics Reviews* **62**(4), 040802 (2009)

[123] Spalart, P.R.: Strategies for turbulence modelling and simulations. *International journal of heat and fluid flow* **21**(3), 252–263 (2000)

[124] Alam, M., Thompson, D.S., Walters, D.K.: Hybrid reynolds-averaged navier–stokes/large-eddy simulation models for flow around an iced wing. *Journal of aircraft* **52**(1), 244–256 (2015)

[125] Cambon, C., Scott, J.F.: Linear and nonlinear models of anisotropic turbulence. *Annual review of fluid mechanics* **31**(1), 1–53 (1999)

[126] Durbin, P.A.: Some recent developments in turbulence closure modeling. *Annual Review of Fluid Mechanics* **50**(1), 77–103 (2018)

[127] Richardson, L.F.: *Weather Prediction by Numerical Process*. Cambridge University Press, Cambridge, UK (1922)

[128] Owhadi, H., Scovel, C., Yoo, G.R.: *Kernel Mode Decomposition and the Programming of Kernels*. Springer, Cham, Switzerland (2021)

[129] Posa, D.: A simple description of spatial-temporal processes. *Computational Statistics & Data Analysis* **15**(4), 425–437 (1993)

[130] Cressie, N., Huang, H.-C.: Classes of nonseparable, spatio-temporal stationary covariance functions. *Journal of the American Statistical Association* **94**(448), 1330–1339 (1999)

[131] Myers, D., De Iaco, S., Posa, D., De Cesare, L.: Space-time radial basis functions. *Computers & Mathematics with Applications* **43**(3-5), 539–549 (2002)

[132] Stein, M.L.: Space-time covariance functions. *Journal of the American Statistical Association* **100**(469), 310–321 (2005)

[133] Romero, D., Ioannidis, V.N., Giannakis, G.B.: Kernel-based reconstruction of space-time functions on dynamic graphs. *IEEE Journal of Selected Topics in Signal Processing* **11**(6), 856–869 (2017)

[134] Jing, X., Qiu, L., Wang, F., Gu, Y.: Modified space-time radial basis function collocation method for solving three-dimensional transient elastodynamic problems. *Engineering Analysis with Boundary Elements* **169**, 106027 (2024)

[135] Yue, X., Wang, F., Hua, Q., Qiu, X.-Y.: A novel space-time meshless method for nonhomogeneous convection-diffusion equations with variable coefficients. *Applied Mathematics Letters* **92**, 144–150 (2019)

[136] Ku, C.-Y., Xiao, J.-E., Liu, C.-Y.: Space-time radial basis function-based meshless approach for solving convection-diffusion equations. *Mathematics* **8**(10), 1735 (2020)

[137] Ku, C.-Y., Hong, L.-D., Liu, C.-Y., Xiao, J.-E.: Space-time polyharmonic radial polynomial basis functions for modeling saturated and unsaturated flows. *Engineering with Computers* **38**(6), 4947–4960 (2022)

[138] Li, Z., Mao, X.-Z., Li, T.S., Zhang, S.: Estimation of river pollution source using the space-time radial basis collocation method. *Advances in Water Resources* **88**, 68–79 (2016)

[139] Li, Z., Mao, X.-z.: Global multiquadric collocation method for groundwater contaminant source identification. *Environmental modelling & software* **26**(12), 1611–1621 (2011)

[140] Runge, C.: Über empirische funktionen und die interpolation zwischen äquidistanten ordinaten. *Zeitschrift für Mathematik und Physik* **46**, 224–243 (1901)

[141] Epperson, J.F.: On the runge example. *The American Mathematical Monthly* **94**(4), 329–341 (1987)

[142] Platte, R.B., Driscoll, T.A.: Polynomials and potential theory for gaussian radial basis function interpolation. *SIAM Journal on Numerical Analysis* **43**(2), 750–766 (2005)

[143] Fornberg, B., Zuev, J.: The runge phenomenon and spatially variable shape parameters in rbf interpolation. *Computers & Mathematics with Applications* **54**(3), 379–398 (2007)

[144] Berrut, J.-P., Trefethen, L.N.: Barycentric lagrange interpolation. *SIAM review* **46**(3), 501–517 (2004)

[145] Boyd, J.P.: Six strategies for defeating the runge phenomenon in gaussian radial basis functions on a finite interval. *Computers & Mathematics with Applications* **60**(12), 3108–3122 (2010)

[146] Boyd, J.P.: A comparison of numerical algorithms for fourier extension of the first, second, and third kinds. *Journal of Computational Physics* **178**(1), 118–160 (2002)

[147] Boyd, J.P., Ong, J.R.: Exponentially-convergent strategies for defeating the runge phenomenon for the approximation of non-periodic functions, part i: single-interval schemes. *Comput. Phys* **5**(2-4), 484–497 (2009)

[148] Schaback, R., Wendland, H.: Adaptive greedy techniques for approximate solution of large rbf systems. *Numerical Algorithms* **24**(3), 239–254 (2000)

[149] De Marchi, S., Schaback, R., Wendland, H.: Near-optimal data-independent point locations for radial basis function interpolation. *Advances in Computational Mathematics* **23**(3), 317–330 (2005)

[150] Galy-Fajou, T., Opper, M.: Adaptive inducing points selection for gaussian processes. arXiv preprint arXiv:2107.10066 (2021)

[151] Sommariva, A., Vianello, M.: Computing approximate fekete points by qr factorizations of vandermonde matrices. *Computers & Mathematics with Applications* **57**(8), 1324–1336 (2009)

[152] Bos, L.P., Levenberg, N., *et al.*: On the calculation of approximate fekete points: the univariate case. *Electron. Trans. Numer. Anal.* **30**, 377–397 (2008)

[153] De Marchi, S., Marchioro, M., Sommariva, A.: Polynomial approximation and cubature at approximate fekete and leja points of the cylinder. *Applied Mathematics and Computation* **218**(21), 10617–10629 (2012)

[154] Sommariva, A., Vianello, M., *et al.*: Approximate fekete points for weighted polynomial interpolation. *Electron. Trans. Numer. Anal.* **37**, 1–22 (2010)

[155] Karvonen, T., Särkkä, S., Tanaka, K.: Kernel-based interpolation at approximate fekete points. *Numerical Algorithms* **87**(1), 445–468 (2021)

[156] Hastie, T., Tibshirani, R., Friedman, J.H., Friedman, J.H.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* vol. 2. Springer, New York, NY (2009)

[157] Sudret, B., Marelli, S., Wiart, J.: Surrogate models for uncertainty quantification: An overview. In: 2017 11th European Conference on Antennas and Propagation (EUCAP), pp. 793–797 (2017). IEEE

[158] Asouti, V., Kontou, M., Giannakoglou, K.: Radial basis function surrogates for uncertainty quantification and aerodynamic shape optimization under uncertainties. *Fluids* **8**(11), 292 (2023)

[159] Kovachki, N., Liu, B., Sun, X., Zhou, H., Bhattacharya, K., Ortiz, M., Stuart, A.: Multiscale modeling of materials: Computing, data science, uncertainty and goal-oriented optimization. *Mechanics of Materials* **165**, 104156 (2022)

[160] Martin, J., Wilcox, L.C., Burstedde, C., Ghattas, O.: A stochastic newton mcmc method for large-scale statistical inverse problems with application to seismic inversion. *SIAM Journal on Scientific Computing* **34**(3), 1460–1487 (2012)

[161] Huang, D.Z., Schneider, T., Stuart, A.M.: Iterated kalman methodology for inverse problems. *Journal of Computational Physics* **463**, 111262 (2022)

[162] Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA (2005)

[163] Zhu, Y., Zabaras, N.: Bayesian deep convolutional encoder–decoder networks for surrogate modeling and uncertainty quantification. *Journal of Computational Physics* **366**, 415–447 (2018)

[164] Tharakan, S., March, W., Biros, G.: Scalable kernel methods for uncertainty quantification. *Recent Trends in Computational Engineering-CE2014: Optimization, Uncertainty, Parallel Algorithms, Coupled and Complex Problems*, 3–28 (2015)

[165] Csáji, B.C., Kis, K.B.: Distribution-free uncertainty quantification for kernel methods by gradient perturbations. *Machine Learning* **108**(8), 1677–1699 (2019)

[166] Narayan, A., Yan, L., Zhou, T.: Optimal design for kernel interpolation: Applications to uncertainty quantification. *Journal of Computational Physics* **430**, 110094 (2021)

[167] Wendland, H.: Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics* **4**(1), 389–396 (1995)

[168] De Marchi, S., Perracchione, E.: *Lectures on Radial Basis Functions*. Department of Mathematics “Tullio Levi-Civita”, University of Padua, Padua, Italy (2018). Department of Mathematics “Tullio Levi-Civita”, University of Padua

[169] LeVeque, R.J.: *Finite Volume Methods for Hyperbolic Problems* vol. 31. Cambridge university press, Cambridge, UK (2002)

[170] Stein, M.L.: *Interpolation of Spatial Data: Some Theory for Kriging*. Springer, New York, NY (1999)

[171] Duchon, J.: Splines minimizing rotation-invariant semi-norms in sobolev spaces. In: *Constructive Theory of Functions of Several Variables: Proceedings of a Conference Held at Oberwolfach April 25–May 1, 1976*, pp. 85–100 (2006). Springer

[172] Fasshauer, G.E., Zhang, J.G.: On choosing “optimal” shape parameters for rbf approximation. *Numerical Algorithms* **45**(1), 345–368 (2007)

[173] Cavoretto, R., De Rossi, A., Haider, A., Lancellotti, S.: Comparing deterministic and statistical optimization techniques for the shape parameter selection in rbf interpolation. *Dolomites Research Notes on Approximation* **17**(DRNA Volume 17.3), 48–55 (2024)

[174] Rippa, S.: An algorithm for selecting a good value for the parameter c in radial basis function interpolation. *Advances in Computational Mathematics* **11**(2), 193–210 (1999)

[175] Ling, L., Marchetti, F.: A stochastic extended rippa’s algorithm for lpocv. *Applied Mathematics Letters* **129**, 107955 (2022)

[176] Wang, J.G., Liu, G.R.: On the optimal shape parameters of radial basis functions used for 2-d meshless methods. *Computer Methods in Applied Mechanics and Engineering* **191**(23), 2611–2630 (2002)

[177] Chen, C.-S., Noorizadegan, A., Young, D.L., Chen, C.S.: On the selection of a better radial basis function and its shape parameter in interpolation problems. *Applied Mathematics and Computation* **442**, 127713 (2023)

[178] Krowiak, A., Podgóński, J.: On choosing a value of shape parameter in radial basis function collocation methods. *AIP Conference Proceedings* **2116**(1), 450020 (2019)

[179] Flyer, N., Fornberg, B.: Radial basis functions: Developments and applications to planetary scale flows. *Computers & Fluids* **46**(1), 23–32 (2011)

[180] Shankar, V.: *Radial Basis Function-based Numerical Methods for the Simulation of Platelet Aggregation*. The University of Utah, Utah, US (2014)

[181] Brent, R.P.: *Algorithms for Minimization Without Derivatives*. Prentice-Hall, Englewood Cliffs, NJ (1973)

[182] Dekker, T.J.: Finding a zero by means of successive linear interpolation. In: Dejon, B., Henrici, P. (eds.) *Constructive Aspects of the Fundamental Theorem of Algebra*. Wiley–Interscience, London (1969)

[183] Fasshauer, G., McCourt, M.: *Kernel-based Approximation Methods Using MATLAB*. World Scientific, Singapore; Hackensack, NJ, USA (2015)

[184] Bradbury, J., Frostig, R., Hawkins, P., Johnson, M.J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., Zhang, Q.: JAX: composable transformations of Python+NumPy programs (2018)

[185] Sriperumbudur, B.K., Fukumizu, K., Lanckriet, G.R.: Universality, characteristic kernels and rkhs embedding of measures. *Journal of Machine Learning Research* **12**(7) (2011)

[186] Micchelli, C.A., Pontil, M.: On learning vector-valued functions. *Neural computation* **17**(1), 177–204 (2005)

[187] Ramanan, N., Homsy, G.M.: Linear stability of lid-driven cavity flow. *Physics of Fluids* **6**(8), 2690–2701 (1994)

[188] Peng, Y.-F., Shiao, Y.-H., Hwang, R.R.: Transition in a 2-d lid-driven cavity flow. *Computers & Fluids* **32**(3), 337–352 (2003)

[189] Sahin, M., Owens, R.G.: A novel fully-implicit finite volume method applied to the lid-driven cavity problem—part ii: Linear stability analysis. *International journal for numerical methods in fluids* **42**(1), 79–88 (2003)

[190] Kuhlmann, H.C., Romanò, F.: In: Gelfgat, A. (ed.) *The Lid-Driven Cavity*, pp. 233–309. Springer, Cham (2019)

[191] Chiang, T., Sheu, T.W., Fang, C.: Numerical investigation of vortical evolution in a backward-facing step expansion flow. *Applied Mathematical Modelling* **23**(12), 915–932 (1999)

[192] Sockol, P.M.: Multigrid solution of the navier–stokes equations at low speeds with large temperature variations. *Journal of Computational Physics* **192**(2), 570–592 (2003)

[193] Dritschel, D.G.: A general theory for two-dimensional vortex interactions. *Journal of Fluid Mechanics* **293**, 269–303 (1995)

[194] Josserand, C., Rossi, M.: The merging of two co-rotating vortices: a numerical study. *European Journal of Mechanics-B/Fluids* **26**(6), 779–794 (2007)

[195] Trieling, R., Dam, C., Heijst, G.: Dynamics of two identical vortices in linear shear. *Physics of Fluids* **22**(11) (2010)

[196] Flyer, N., Fornberg, B., Bayona, V., Barnett, G.A.: On the role of polynomials in rbf-fd approximations: I. interpolation and accuracy. *Journal of Computational Physics* **321**, 21–38 (2016)

A Problem configurations

In Table 2, we present detailed configurations for all our experiments, including the domain, the flow regime, the forcing term, the initial and boundary conditions, the final time and time-step used in SU2, and the specific operator being learned. In cases where the flow was turbulent, we used the SST turbulence model in conjunction with RANS.

Table 2: Problem configurations; flow regime, domain, forcing term (f), initial and boundary conditions (IC and BC), time domain, and the operator of interest. Re stands for the Reynolds number and Ra for the Rayleigh number.

#	Problem	Domain	Flow regime	f	IC	BC	Time	Operator
1)	2D Flow Past a Cylinder	$[0, 20] \times [0, 14]$	Laminar	-	\mathbf{u} given	\mathbf{u} given at top/bottom/inlet	$T = 10$	$\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 10)$
			No shedding				$\Delta t = 10^{-3}$	
			$Re \in [25, 64]$			No-slip cylinder		
			Vortex shedding			outlet, $p = 0$		
2)	2D Lid-Driven Cavity Flow	$[0, 1]^2$	Laminar	-	\mathbf{u} given	\mathbf{u}^{top} given $\mathbf{u}_b = 0$, else	$T = 5$ $\Delta t = 10^{-3}$	$\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 5)$
3)	2D Backward Facing Step	$[0, 15] \times [-0.5, 0.5]$	Laminar	$-\mathbf{u}^{\text{inlet}} = \nu x_2 (0.5 - x_2)$	$\mathbf{u}^{\text{inlet}}$ given	$T = 5$	$\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 5)$	
			$Re \in [28, 900]$	$\mathbf{u} = 0$, else	$\mathbf{u} = 0$ at top, bottom, step outlet, $p = 0$	$\Delta t = 10^{-3}$		
4)	2D Buoyancy-Driven Cavity Flow	$[0, 1]^2$	Laminar	$-g$	$\mathbf{u} = 0$	No-slip walls	$T = 5$	$\mathcal{G} : \mathbf{T}(x, 0) \rightarrow \mathbf{u}(y, 5)$
			$Ra = 10^6$		$T_{\text{left}} = T_1$		$\Delta t = 10^{-3}$	
5)	2D Taylor-Green Vortices (spatial)	$[0, 2\pi]^2$	Laminar	-	$\mathbf{u} =$ $(A \sin x_1 \cos x_2, A \cos x_1 \sin x_2) e^{-2\nu t}$	$\mathbf{u}^{\text{left}} = \mathbf{u}^{\text{right}}$ $\mathbf{u}^{\text{top}} = \mathbf{u}^{\text{bottom}}$	$T = 1$	$\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 1)$
6)	2D Taylor-Green Vortices (spacetime)	$[0, 2\pi]^2$	Laminar	-	$\mathbf{u} =$ $(A \sin x_1 \cos x_2, A \cos x_1 \sin x_2) e^{-2\nu t}$	$\mathbf{u}^{\text{left}} = \mathbf{u}^{\text{right}}$ $\mathbf{u}^{\text{top}} = \mathbf{u}^{\text{bottom}}$	$T = [0.7, 0.8, 0.9, 1]$	$\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, T)$
7)	2D Merging Vortices	$[0, 2\pi]^2$	Laminar	-	\mathbf{u} given	$\mathbf{u}^{\text{left}} = \mathbf{u}^{\text{right}}$ $\mathbf{u}^{\text{top}} = \mathbf{u}^{\text{bottom}}$	$T = 0.4$	$\mathcal{G} : \omega(x, 0) \rightarrow \mathbf{u}(y, 0.4)$
					ω (vorticity) given		$\Delta t = 10^{-3}$	
8)	3D Species Transport	See Figure 4	Turbulent	-	$\mathbf{u}^{\text{inlet}}$, given	$\mathbf{u}^{\text{inlet}}$ given	$T = 0.5$	$\mathcal{G} : \mathbf{u}^{\text{inlet}} \rightarrow \mathbf{u}(y, 0.5)$
			$Re \in [10^5, 10^6]$		$\mathbf{u} = 0$, else	no-slip inner wall, outer wall, & blades outlet, $p = 0$	$\Delta t = 0.005$	
9)	3D Flow Past an Airfoil	$[-7, 10] \times [-7, 7] \times [0, 3]$	Turbulent	-	$\mathbf{u}^{\text{inlet}}$, given	$\mathbf{u}^{\text{inlet}}$ given	$T = 1$	$\mathcal{G} : \tau \rightarrow \mathbf{u}(y, 1)$
			$Re = 3 \times 10^6$		$\mathbf{u} = 0$, else	$\mathbf{u}_{\text{front}} = \mathbf{u}_{\text{back}}$ No-slip airfoil outlet, $p = 0$	$\Delta t = 10^{-2}$	
τ denotes the airfoil shape								

B Additional results

We present additional experimental results in this section on four classic incompressible fluid flow problems (including a temperature-driven flow) and a challenging one involving merging vortices.

B.1 2D Taylor–Green vortices

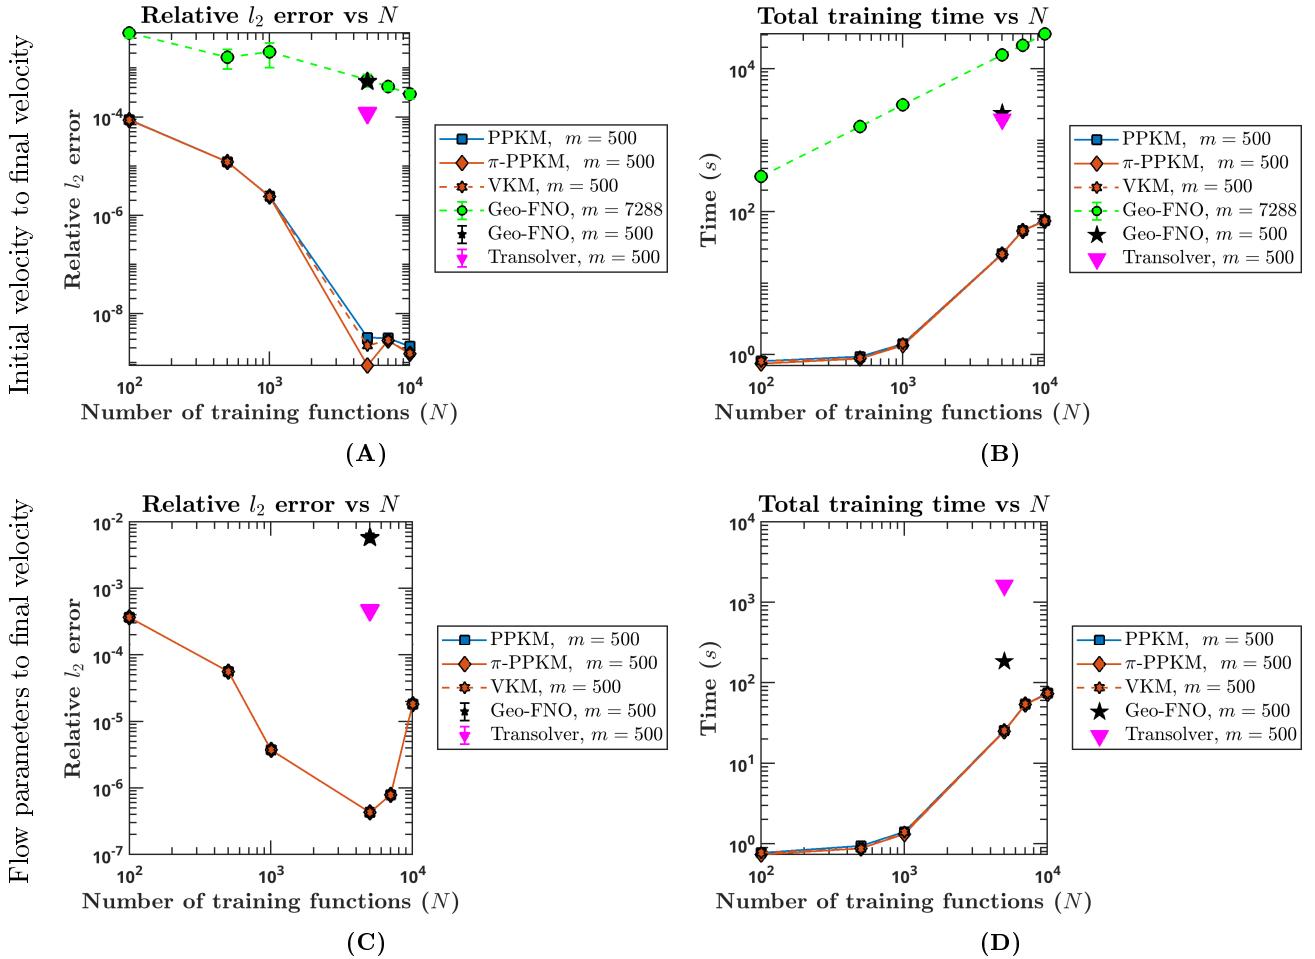


Fig. 6: The 2D laminar Taylor–Green vortices problem for the purely **spatial** operator map. The output functions are snapshots of the velocity at time $T = 1$. (A) and (B) show the test relative ℓ_2 errors and training runtimes as functions of N for the operator map from the initial velocity to the final velocity. (C) and (D) show the same results for the operator map from the flow parameters to the final velocity.

We considered a purely spatial variant of the classical Taylor-Green vortex problem described in Section 1.2, which we remind the reader has an exact solution. We are interested in two different operator maps being learned, (i) the spatial map $\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 1)$, and (ii) the parametric map $\mathcal{G} : \tau \rightarrow \mathbf{u}(y, 1)$ where $\Omega_a = \tau \subset \mathbb{R}^2$ is the parameter space of A and ν . All other flow details were identical to the setup in Section 1.2 and are provided in row 5 in Table 2. We show the results for both operator maps in Figure 6. We saw similar trends in this problem as the ones in the spacetime version of this problem in Section 1.2. However, here we saw the π-PPKM gain a mild accuracy advantage over the PPKM method for $N = 5000$, demonstrating the benefit of incorporating spatial periodicity in this periodic problem. Additionally, our method significantly outperforms the neural operator baselines in terms of both accuracy and training times.

B.2 2D lid-driven cavity flow

The classic 2D lid-driven cavity flow problem [187–190] describes laminar fluid flow in a square cavity $\Omega_a = \Omega_v = [0, 1]^2$ whose lid moves tangentially along the top boundary. We initialized the horizontal velocity u_1^{top} to be constant along the top boundary (hence an impulsive start), sampled from the random distribution $\mathcal{N}_{[0.8, 2.5]}(1.5, 1)$, and zero everywhere else. u_2 was initialized to zero everywhere. We used the moving wall option provided in SU2 to prescribe the BC on the top

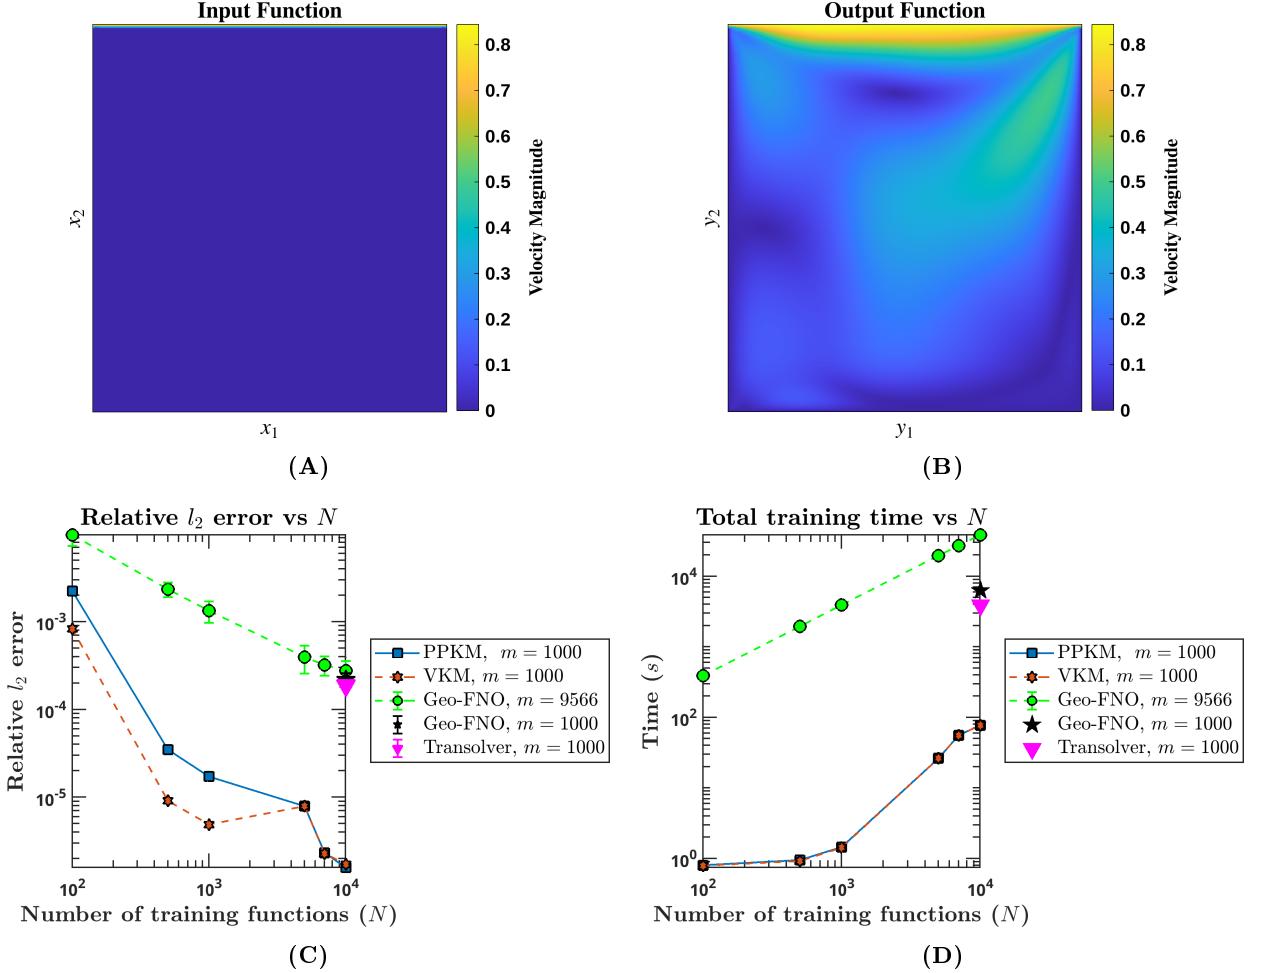


Fig. 7: The 2D laminar lid-driven cavity flow problem. (A) and (B) show examples of an input function (the initial velocity) and an output function (the final velocity), respectively. (C) and (D) show the test relative ℓ_2 errors and training runtimes as functions of N .

boundary. The other three boundaries had no-slip BCs. We used a triangular mesh with 9566 points. The simulation ran until $T = 5$ with $\Delta t = 10^{-3}$ and the operator map of interest was $\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 5)$. We also provide these details in row 2 in Table 2. We show example input and output functions in Figures 7A and 7B respectively.

The VKM method outperformed the PPKM method for small N as shown in Figure 7C, but both methods achieved the same errors as N increased. Our method achieved around 2.5 orders of magnitude lower errors and around 1.4 orders of magnitude faster training times (Figure 7D) than the neural operators.

B.3 2D backward-facing step

This problem focuses on 2D laminar fluid flow over an isothermal backward-facing step at $Re = 800$ [191]². The domain of the problem is $\Omega_a = \Omega_v = [0, 15] \times [-0.5, 0.5]$. Typically, flow simulations for this problem involve fluid flow in both the upstream and downstream channels. We simplified the geometry by only simulating the downstream channel of the step. Let the channel height and width be H and $15H$, respectively. The upper half of the left boundary was specified with the following initial condition,

$$u_1(x_2) = \nu x_2 * (0.5 - x_2), 0 \leq x_2 \leq 0.5, \quad (35)$$

²https://su2code.github.io/tutorials/Inc_Laminar_Step

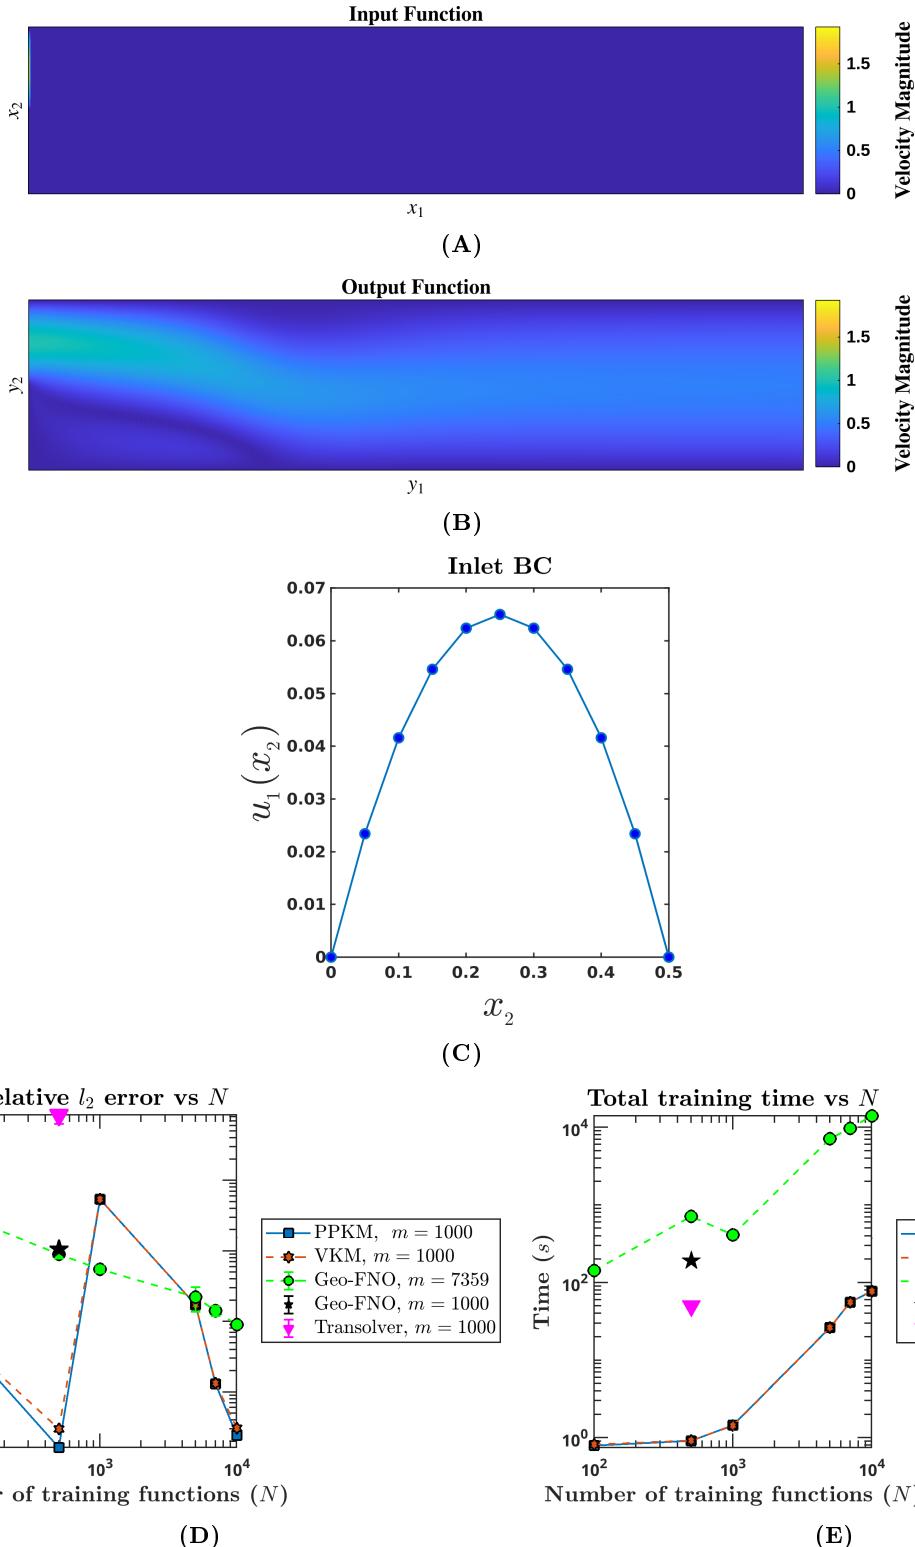


Fig. 8: The 2D laminar backward-facing step problem. (A) and (B) show examples of an input function (the initial velocity) and an output function (the final velocity), respectively. (C) shows the parabolic profile of the inlet velocity. (D) and (E) show the test relative ℓ_2 errors and training runtimes as functions of N .

where ν is the kinematic viscosity, u_1 is the horizontal velocity component, and x_2 is the vertical spatial coordinate. The velocity was initialized to be zero everywhere else. See Figure 8C for an example of the parabolic inlet BC. The top, bottom, and bottom half of the left (corresponding to the step face) boundaries were prescribed no-slip BCs. A zero pressure BC was imposed on the right boundary. We ran the SU2 simulation until $T = 5$ with $\Delta t = 10^{-3}$. ν was randomly sampled from the random distribution $\mathcal{N}_{[1,36]}(18, 18)$ which resulted in Reynolds number range [28, 900] across the simulations. A triangular mesh with 7359 points was used. The operator of interest was $\mathcal{G} : \mathbf{u}(x, 0) \rightarrow \mathbf{u}(y, 5)$. These configuration details are provided in row 3 in Table 2. We show examples of input and output functions in Figures 8A and 8B respectively.

Here, our method rapidly reached the lowest errors around $N = 1000$, possibly due to hidden low-rank structure in the input functions. A subsequent spike was then followed by another steady decrease in error for increasing N as shown in Figure 8D. Regardless, our method achieved close to three orders of magnitude lower errors and two orders of magnitude faster training times (Figure 8E) than the neural operators.

Out-of-distribution (OOD) test: We additionally performed an experiment to compare our methods against baseline operator learning methods in terms of OOD capability. We trained on the dataset which falls in the $\text{Re} \in [28, 900]$ regime and tested on 1000 test functions whose ν was sampled from $\mathcal{N}_{[38,70]}(54, 30)$. For $N = 500$ and $m = 1000$, both the kernel-based methods achieved a relative ℓ_2 error around 0.961 whereas the Geo-FNO and the Transolver models achieved errors of 0.221 and 0.0852, respectively. This is somewhat intuitive: the kernel methods are interpolatory and thus more sensitive to the range of the training data upon generalization, while neural operators are essentially trained with least-squares techniques, have more trainable parameters, and therefore better OOD behavior. We strongly suspect that building a viscosity-aware kernel will change this easily; it is also likely that making the kernels have more trainable parameters will help close the gap.

B.4 2D buoyancy-driven cavity flow

This (modern) classic problem describes the 2D laminar buoyancy-driven flow in a square cavity $\Omega_a = \Omega_v = [0, 1]^2$ with adiabatic top and bottom walls and constant temperature on the left and right walls [192]³. Additionally, we also prescribed the gravitational body force $f = -g$. All four walls had no-slip BCs prescribed. We randomly initialized the temperature (in Kelvin) at the left wall T^{left} by sampling from the distribution $\mathcal{N}_{[300,600]}(450, 100)$. T^{right} was set to $T^{\text{left}}/4$. The flow density was initialized to $5.97 \times 10^{-3} \text{ kg/m}^3$ which corresponds to a Rayleigh number of 10^6 . The temperature everywhere else in the domain was initialized to 288.15 K. The domain was discretized with a quadrilateral mesh containing 10,514 points. We ran time dependent simulations in SU2 until $T = 5$ with $\Delta t = 10^{-3}$. The operator map being learned was $\mathcal{G} : T(x, 0) \rightarrow \mathbf{u}(y, 5)$. These configuration details are provided in row 4 in Table 2. We show examples of input and output functions in Figures 9A and 9B. Our method achieved orders of magnitude smaller errors (Figure 9C) and training times (Figure 9D), and a faster convergence rate w.r.t N than the neural operator baselines.

B.5 2D merging vortices

This challenging problem describes the 2D laminar flow resulting from the merging of two vortices [193–195] in a square domain $\Omega_a = \Omega_v = [0, 2\pi]^2$. We started by prescribing an incompressible velocity field in the polar coordinates centered at $x = (x_1^c, x_2^c)$. Let $r = \sqrt{(x_1 - x_1^c)^2 + (x_2 - x_2^c)^2}$ where $x = (x_1, x_2) \in \Omega$. Then we prescribed the following Cartesian velocity field components:

$$u_1(x) = -\frac{x_2 - x_2^c}{r} q(r), \quad (36)$$

$$u_2(x) = \frac{x_1 - x_1^c}{r} q(r), \quad (37)$$

with $q(r) = 4re^{-r^\alpha}$, ensuring $q(0) = 0$ for regularity at the vortex center. This construction yielded a divergence-free velocity field by construction. The associated scalar vorticity field ω is given by

$$\omega(r) = (8 - 4\alpha r^\alpha) e^{-r^\alpha}, \quad (38)$$

³https://su2code.github.io/tutorials/Inc_Laminar_Cavity

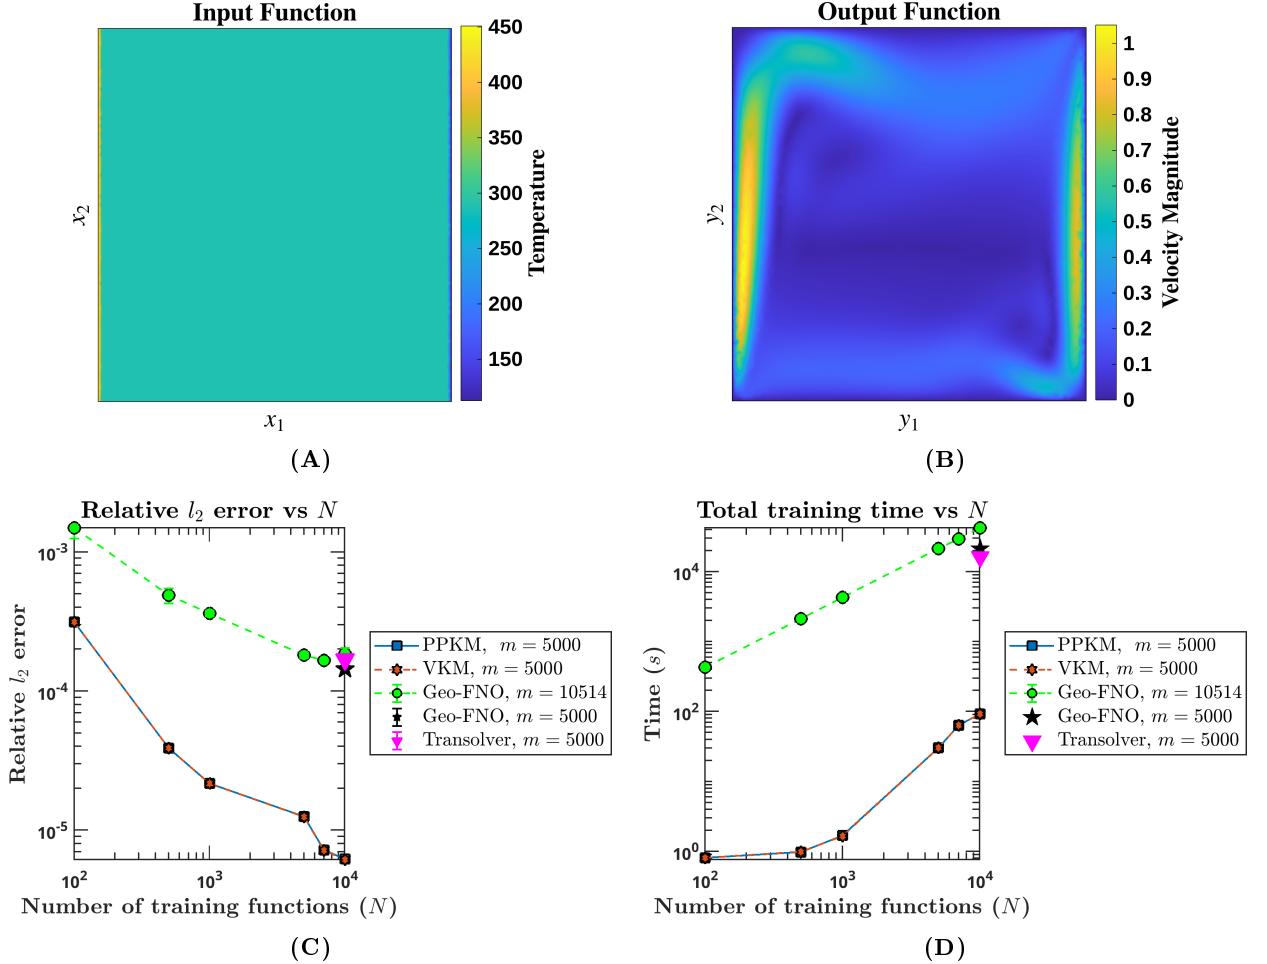


Fig. 9: The 2D laminar buoyancy-driven cavity flow problem. (A) and (B) show examples of an input function (the initial temperature) and an output function (the final velocity), respectively. (C) and (D) show the test relative ℓ_2 errors and training runtimes as functions of N .

where $\alpha > 0$ controls the sharpness of the vortex core. To the best of our knowledge as of the time of writing, SU2 does not allow the direct prescription of an initial vorticity field for time dependent simulations; we therefore initialized the flow through the velocity field in (37). We prescribed periodic BCs on all four walls. We sampled $\alpha \in [1, 14]$ from a uniform distribution. We used the same triangular mesh as the one in Section B.1. We ran the SU2 simulation until $T = 0.4$ with $\Delta t = 10^{-3}$. The operator of interest here was $\mathcal{G} : \omega(x, t = 0) \rightarrow \mathbf{u}(y, 0.4)$. We provide these details in row 7 in Table 2. We show examples of the input and output functions in Figures 10A and 10B.

The errors reported in Figure 10C show that our method rapidly achieves the lowest of all the errors just shy of $N = 1000$. The errors then began to climb as a function of N . The difficulty of this problem is further shown by the unusually large error bars on the neural operators and the fact that there are no discernible differences in the accuracy between π -PPKM and PPKM despite this being a periodic problem. It is likely that our method's accuracy deteriorates for large N due to the same reason as the no-vortex-shedding flow regime in Section 1.1; the output functions vary minimally with changes in the input functions. For a small training dataset, $N = 100$ or $N = 500$, our method outperforms the neural operator baselines. These results where errors climb with N also indicate the need for importance sampling and/or adaptive sampling of the input functions.

C Effect of the Ridge Parameter θ on Accuracy

As described in Section 3.4.1, we added a small regularization parameter (“nugget”) θ to the diagonal of the operator kernel. Across our experiments, we tested three different values of θ : 10^{-4} , 10^{-6} , and 10^{-8} , and reported the one with the best errors. We formed two insights from this

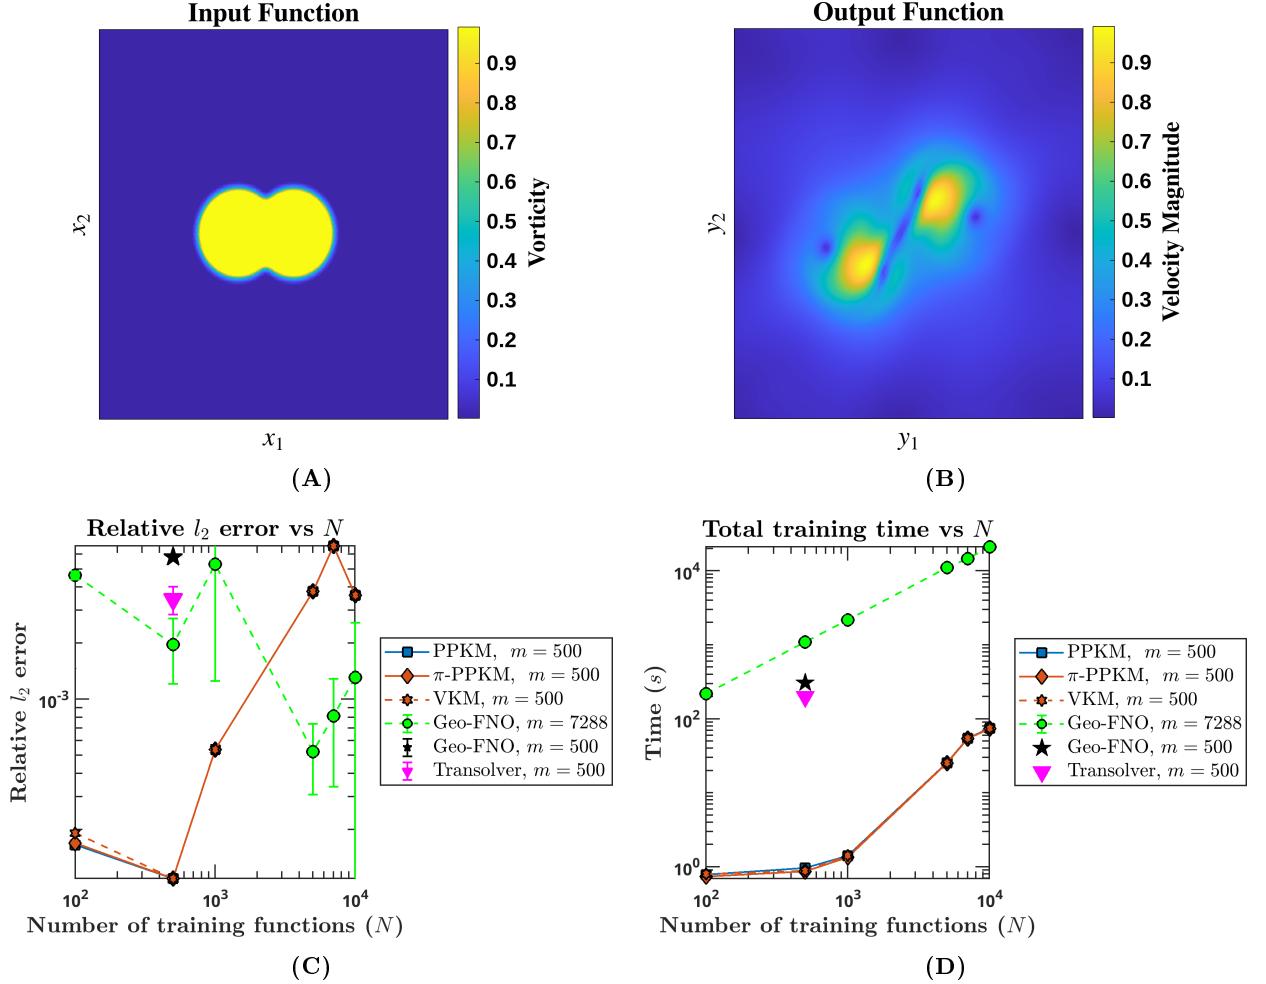


Fig. 10: The 2D laminar merging vortices problem. (A) and (B) show examples of an input function (the initial vorticity containing two adjacent vortices) and an output function (the final velocity), respectively. (C) and (D) show the test relative ℓ_2 errors and training runtimes as functions of N .

experiment. First, for problems where the dataset was generated using SU2, the operator kernel benefited greatly from a larger θ as opposed to problems where the dataset was generated from analytical solutions. This can be seen in Figures 12 and 13, which show the relative ℓ_2 errors at the coarsest and finest nodesets using all three magnitudes of θ for the 3D species transport and 2D Taylor–Green vortices problems, respectively. The effect of the magnitude of θ was more pronounced for the results on the coarser nodesets, shown in Figures 12A and 13A, than in the finer nodesets, shown in Figures 12B and 13B. This is likely due to the numerical truncation errors inherent in the solution fields, manifesting as both numerical dissipation and dispersion. These errors likely propagate to the spatial interpolation coefficients which are then interpolated by the operator kernel. In contrast, the Taylor–Green problems do not use SU2 as the solutions are known analytically. Consequently, adding a regularization parameter hurts the performance there. For most problems using SU2, $\theta = 10^{-4}$ or $\theta = 10^{-6}$ yielded the best errors. We also experimented with adding a similar θ to the property-preserving kernel to address the noise in the velocity fields, however, it exaggerated the regularization of the spatial interpolation coefficients which resulted in very poor relative ℓ_2 errors on generalization.

D Examples of approximate kernel Fekete points

Figure 14 shows example figures of $m = 100$, $m = 500$, and $m = 1000$ approximate kernel Fekete points in the domain $\Omega = [0, 1]^2$. These points were picked from the mesh of the 2D lid-driven cavity flow problem. The figures show very clearly the greater relative importance the algorithm puts on

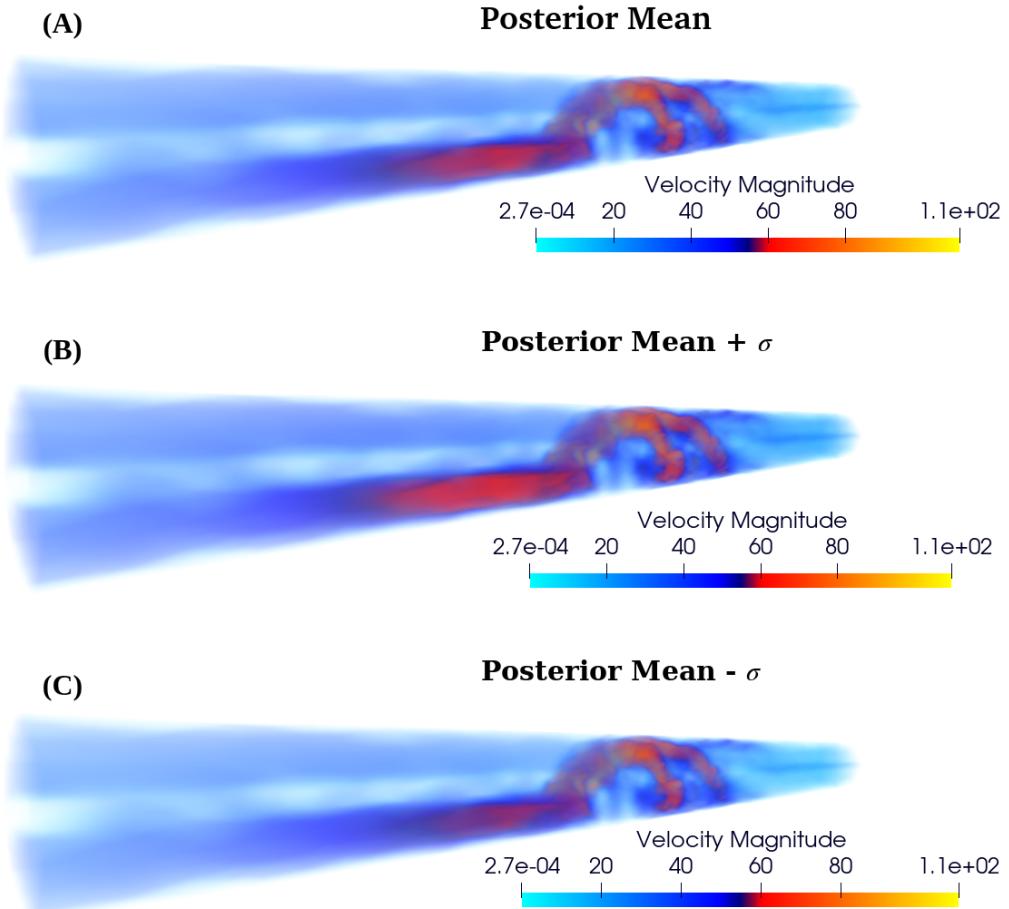


Fig. 11: Uncertainty quantification for the 3D turbulent species transport example. (A) shows the posterior mean velocity field. (B) and (C) show the velocity fields corresponding to the one standard deviation of the posterior predictive distribution ξ . Section 1.3.1 provides a discussion of these results.

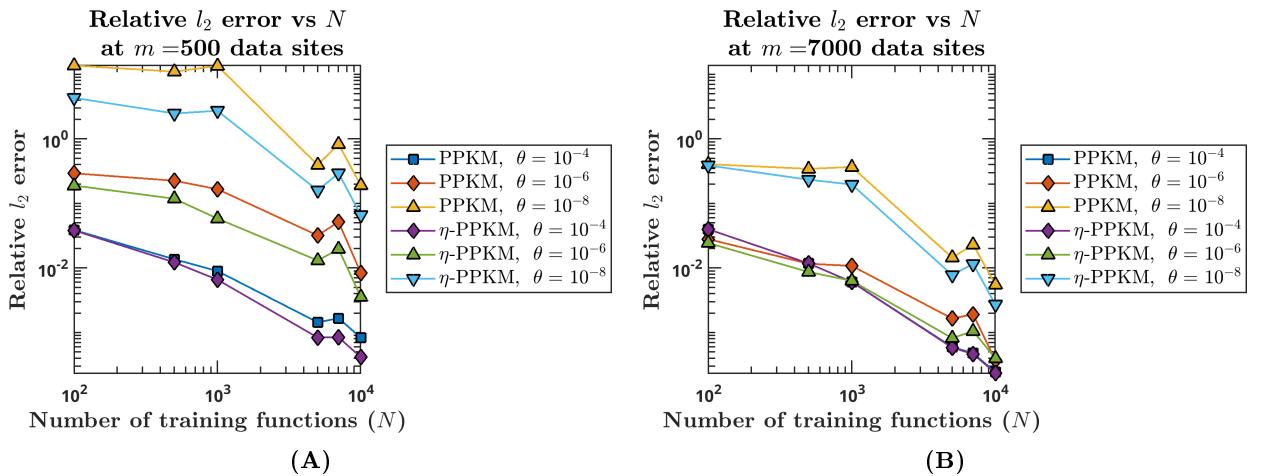


Fig. 12: The effect of the magnitude of θ in the 3D species transport problem. We tested using $\theta = 10^{-4}$, $\theta = 10^{-6}$, and $\theta = 10^{-8}$. (A) and (B) show the test relative ℓ_2 errors at 500 and 7000 points, respectively, using different θ .

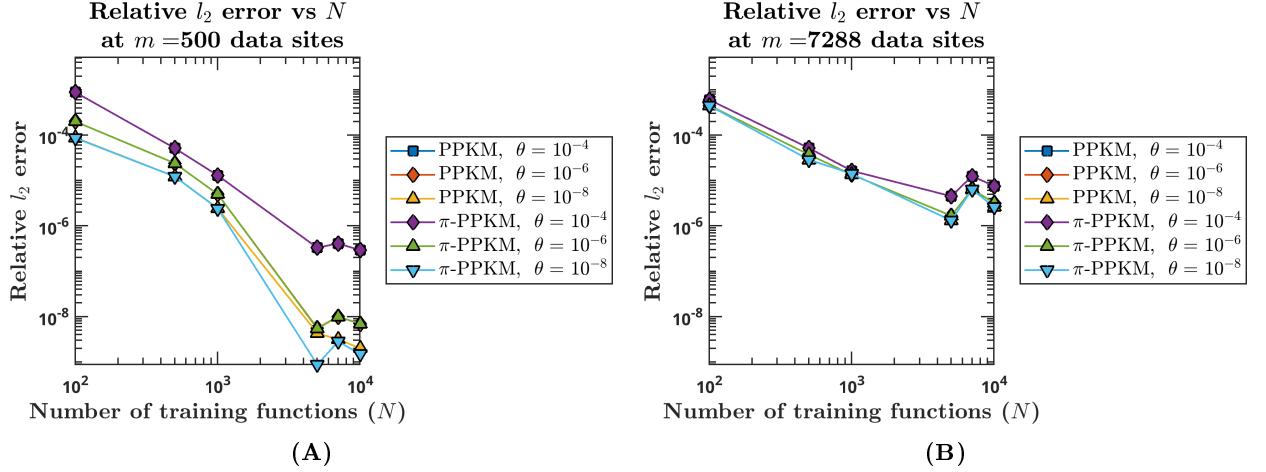


Fig. 13: The effect of the magnitude of θ in the 2D laminar Taylor–Green vortices problem for the purely spatial operator map. We tested using $\theta = 10^{-4}$, $\theta = 10^{-6}$, and $\theta = 10^{-8}$. (A) and (B) show the test relative ℓ_2 errors at 500 and 7288 points, respectively, using different θ .

the points closer to the boundary, which helps alleviate the Runge phenomenon and improves the regularity of the interpolation coefficients.

E Neural operator implementation details

We outline relevant implementation details for the Geo-FNO [66] and Transolver [67] models below.

Geo-FNO: The resolution is denoted by s in the publicly available code by the authors of [66]. We tested with $s = \{20, 30, 40, 50, 60\}$ for all the 2D problems and $s = \{10, 15, 20, 25\}$ for both 3D problems and picked the best performing s . We also tuned the number of modes used, searching over three distinct values for each problem including $\lfloor \frac{s}{2} \rfloor$. The channel dimension was varied over $\{32, 64, 128\}$. Following [66], we fixed the number of Fourier layers to four. Due to the Geo-FNO requiring the input and output functions to have the same d -dimensional domain, we made the following adjustments to the datasets:

- **2D Taylor–Green spacetime:** We treated this equivalently to a 3D spatial problem.
- **2D Taylor–Green spacetime parametric map:** This was also treated as a 3D spatial problem except since the input functions are the A and ν coefficients, we repeated them as vector-valued constant functions sampled on the same spacetime grid as the output functions.
- **3D Species transport and flow past an airfoil:** We padded the input functions in the z direction with zeros to make them functions of three variables.

Transolver: The hyperparameters in this model included the number of attention layers (varied over $\{3, 4, 5\}$), dimensions of the embeddings (varied over $\{32, 64, 128\}$), number of heads (varied over $\{4, 6, 8\}$), and number of slices (varied over $\{16, 32, 64\}$). In addition to matching the dimensionality of the input and output functions like the Geo-FNO, the Transolver model additionally requires that the input and output functions share the same underlying discretization. To achieve this for the 2D Taylor–Green vortices spacetime (both the spatiotemporal and parametric maps), 3D species transport, and 3D flow-past-an-airfoil problems, we constructed a new grid as the union of the input and output discretizations. Both the input and output functions were evaluated on this new grid with zeros where their respective data was missing. The loss was computed only on the original output function grid.

Training: For both models, we used the Adam optimizer and the default learning rate schedule from the publicly available code. A cosine annealing learning rate schedule was used for the Geo-FNO and a OneCycle⁴ learning rate schedule for the Transolver. For both, we set the maximum learning rate to 10^{-3} and the activation function to GELU. In all problems and for both models, we normalized the input spatial coordinates to $[0, 1]^d$; we observed that omitting this step degraded the

⁴https://docs.pytorch.org/docs/stable/generated/torch.optim.lr_scheduler.OneCycleLR.html

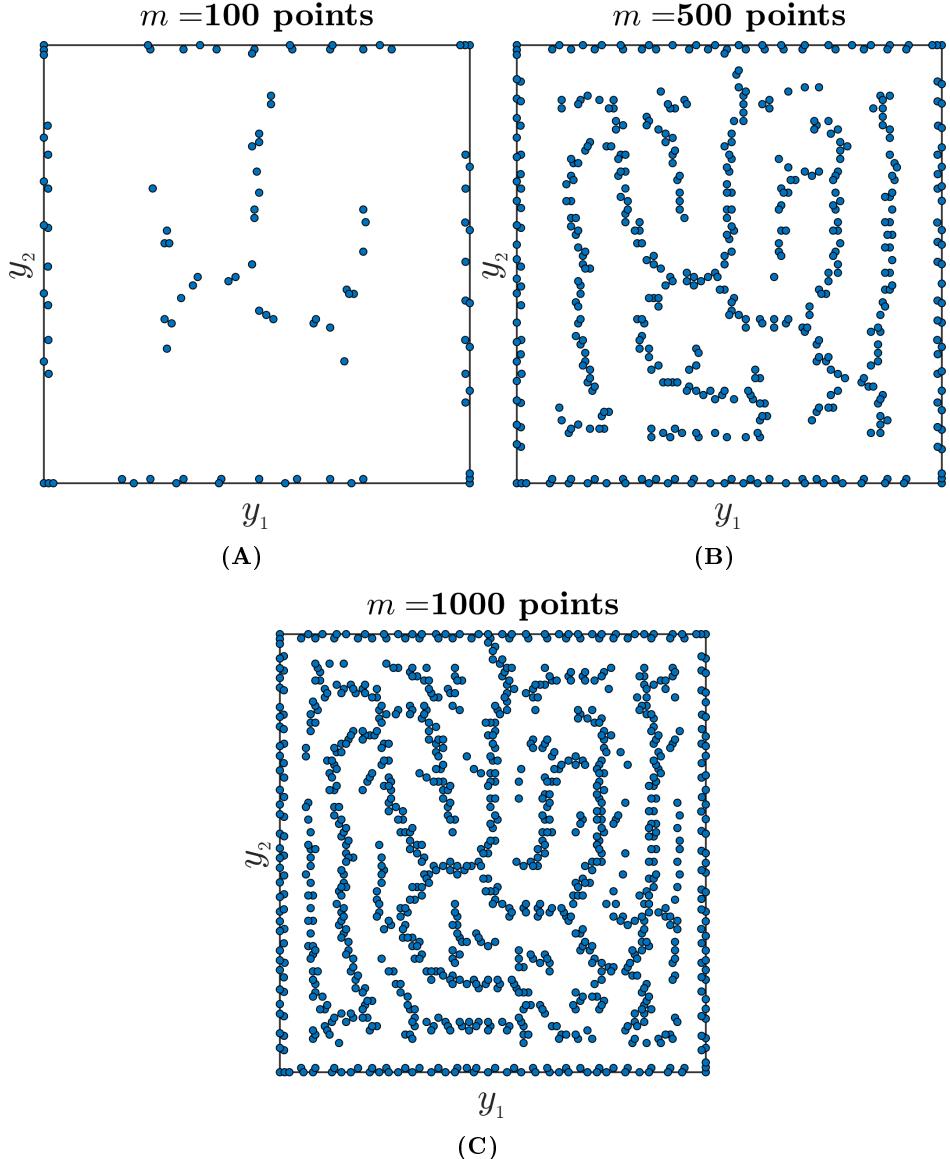


Fig. 14: (A) $m = 100$, (B) $m = 500$, and (C) $m = 1000$ approximate kernel Fekete points in the domain $\Omega = [0, 1]^2$.

performance. We used the default normalization scheme for the input function samples themselves; specifically, the Geo-FNO rescales the input functions pointwise across all training functions to zero mean and unit variance, and the Transolver rescales input functions with a global normalization across all points *and* training functions to zero mean and unit variance. Both were trained in single precision for 500 epochs with a batch size of 20 using the NCSA Delta GPU cluster ⁵ with NVIDIA A100 and A40 GPUs.

Divergence computation: We use high-order accurate local (stencil-wise) interpolation with polyharmonic spline (PHS) RBFs augmented with high-degree polynomials [196] to interpolate the neural operators' output functions. Then, the interpolant's bases were differentiated analytically to compute the divergence. The order of accuracy of the interpolation was varied over $\{5, 7, 9\}$ and the lowest resulting divergence (pointwise maximum divergence averaged across all test functions) was reported; we excluded the boundary divergences for safety as the stencils at the boundary were one-sided and hence not as reliable.

⁵<https://www.ncsa.illinois.edu/research/project-highlights/delta>