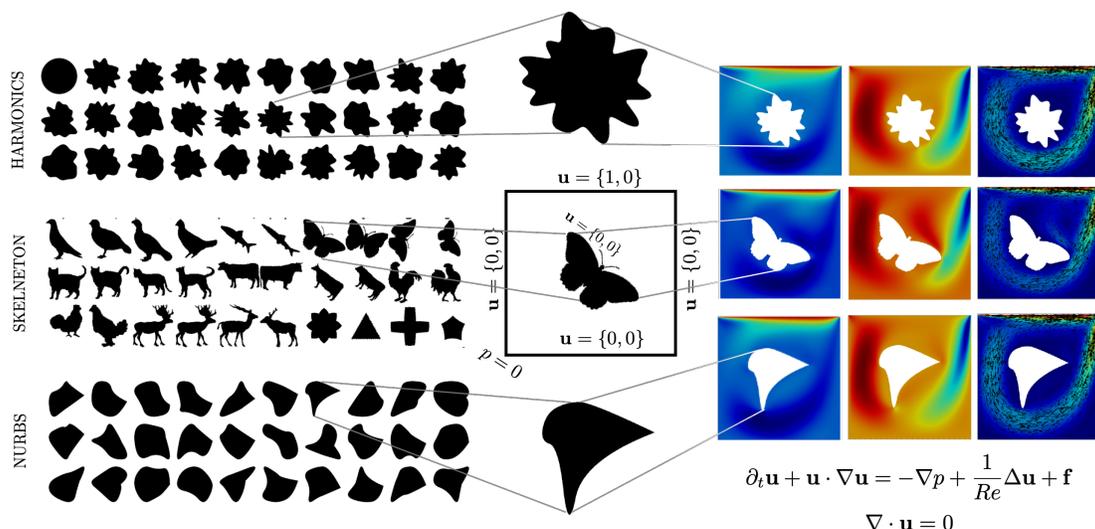


# FlowBench: A Large Scale Benchmark for Flow Simulation over Complex Geometries

Ronak Tali<sup>1†</sup>, Ali Rabeh<sup>1†</sup>, Cheng-Hau Yang<sup>1†</sup>, Mehdi Shadkhah<sup>1</sup>, Samundra Karki<sup>1</sup>, Abhisek Upadhyaya<sup>2</sup>, Suriya Dhakshinamoorthy<sup>1</sup>, Marjan Saadati<sup>1</sup>, Soumik Sarkar<sup>1</sup>, Adarsh Krishnamurthy<sup>1</sup>, Chinmay Hegde<sup>2</sup>, Aditya Balu<sup>1</sup>, Baskar Ganapathysubramanian<sup>1\*</sup>

{rtali — arabeh — chenghau — mehdish — samundra}@iastate.edu, au2216@nyu.edu  
 {snarayan — marjansd — soumiks — adarsh}@iastate.edu, chinmay.h@nyu.edu  
 {baditya — baskarg}@iastate.edu

<sup>1</sup>Iowa State University, Ames, IA 50011, USA  
<sup>2</sup>New York University, New York, NY 10012, USA



**Figure 1:** *FlowBench* offers comprehensive datasets and metrics for assessing neural PDE solvers designed to model flow phenomena around complex objects. It includes three sets of application-relevant geometries with varying complexities and high-fidelity flow simulation data under different forcing conditions. The left panel in the figure above showcases 30 randomly selected shapes from each geometry group. The middle panel provides a close-up of one geometry within the computational domain, highlighting the boundary conditions. The right panel displays the simulation outputs, including velocity results for three samples.

## Abstract

Simulating fluid flow around arbitrary shapes is key to solving various engineering problems. However, simulating flow physics across complex geometries remains numerically challenging and computationally resource-intensive, particularly when using conventional PDE solvers. Machine learning methods offer attractive opportunities to create fast and adaptable PDE solvers. However, benchmark datasets to measure the performance of such methods are scarce, especially for flow physics across complex geometries. We introduce **FlowBench**, a dataset for neural simulators with over 10K samples, which is currently larger than any publicly available flow physics dataset. **FlowBench** contains flow simulation data across complex geometries (*parametric vs. non-parametric*), spanning a range of flow conditions (*Reynolds number and Grashoff number*), capturing a diverse array of flow phenomena (*steady vs. transient; forced vs. free convection*), and for both 2D

† These authors contributed equally to this work.  
 \* Corresponding author

and 3D. **FlowBench** contains over 10K data samples, with each sample the outcome of a fully resolved, direct numerical simulation using a well-validated simulator framework designed for modeling transport phenomena in complex geometries. For each sample, we include velocity, pressure, and temperature field data at 3 different resolutions and several summary statistics features of engineering relevance (such as coefficients of lift and drag, and Nusselt numbers). We envision that **FlowBench** will enable evaluating the interplay between complex geometry, coupled flow phenomena, and data sufficiency on the performance of current, and future, neural PDE solvers. We enumerate several evaluation metrics to help rank order the performance of current (and future) neural PDE solvers. We benchmark the performance of several methods, including Fourier Neural Operators (FNO), Convolutional Neural Operators (CNO), DeepONets, and recent foundational models. This dataset ([here](#)) will be a valuable resource for developing and evaluating AI-for-science approaches, specifically neural PDE solvers, that model complex fluid dynamics around 2D and 3D objects.

## 1 Introduction

Accurate modeling of fluid flow around complicated objects is central to a plethora of applications. In aerospace and automotive applications, flow around wings, and car bodies can significantly impact design and performance [Greenblatt and Wygnanski \(2000\)](#); [Spohn and Gilliéron \(2002\)](#); [You and Moin \(2008\)](#). In civil and environmental engineering applications, understanding fluid flow patterns around structures like buildings [\(Ramponi and Blocken, 2012\)](#) and bridges [\(Helgedagsrud et al., 2019\)](#) is critical from a safety perspective, while flow patterns inside buildings are important from a safety and comfort perspective [\(Tan et al., 2023\)](#). Bio-flow applications, such as flow around fish [\(Seo and Mittal, 2022; Zhu et al., 2021\)](#), birds [\(Song et al., 2014\)](#), insects [\(Engels et al., 2016\)](#), and heart valves [\(Xu et al., 2021\)](#), are all important fields of study. Sports engineering scientists use fluid flow simulations to optimize equipment design, such as bike helmets and golf ball shapes [\(Ting, 2003\)](#). *All these applications involve flow through and across **complex geometries**.*

Fluid flow is also often connected with thermal effects. A thermal mismatch between the fluid medium and the object geometry can produce buoyancy-driven phenomena, cause thermal plumes, and impact mixing and thermal transport with significant engineering implications, necessitating a multiphysics approach, where fluid flow and heat transfer are modeled concurrently. For instance, in the semiconductor industry, the design of heat exchangers must account for how the geometry influences both fluid flow and thermal transport [\(Bhutta et al., 2012; Abeykoon, 2020\)](#). Similarly, addressing urban heat island effects requires a deep understanding of the interaction between thermal and flow dynamics in complex urban settings [\(Priyadarsini et al., 2008; Allegrini and Carmeliet, 2018\)](#). Additionally, ensuring indoor comfort and safety, particularly in mitigating the transmission of infectious diseases, depends on accurately modeling the interplay of airflow and temperature within built environments [\(Bhattacharyya et al., 2020; Tan et al., 2023\)](#). *These examples further underscore the necessity of considering multiphysics flow simulations through and across **complex geometries**.*

A central challenge in accurately simulating flow patterns in complex geometries is the high resource cost of traditional simulation approaches [\(Rabeh et al., 2024b\)](#). High-fidelity flow (and thermal) simulations often require hours to days (or sometimes even months [\(Saurabh et al., 2023\)](#)) on high-performance computing (HPC) systems. Scientific machine learning (SciML) has emerged as a promising path towards resolving this challenge. By combining training data with domain-specific information (e.g., physical constraints and smoothness assumptions), SciML approaches offer fast simulation, better extrapolation capabilities, and lower data requirements. This includes impactful applications like weather prediction [\(Rasp et al., 2020\)](#) and canonical flows in simple geometries [\(Khara et al., 2024; Bonnet et al., 2022a; Luo et al., 2024; Xu et al., 2023; Janny et al., 2023; Bonnet et al., 2022b\)](#).

Despite these successes, however, there is a notable lack of datasets for flow (and flow-thermal) interactions with complicated geometries. While databases exist for flow past simple shapes such as cylinders [\(Luo et al., 2024; Xu et al., 2023\)](#) and airfoils [\(Bonnet et al., 2022a,b\)](#), there remains a significant gap in datasets involving more diverse and complex shapes. We note the availability of a few datasets that include other geometries, but these geometries are limited to drone shapes [\(Janny et al., 2023\)](#), and canonical car shapes [\(Ashton et al., 2024b\)](#).

In addition to the need for more complex geometries in existing flow datasets, there is also a shortage of multiphysics datasets. By multiphysics, we mean applications that are modeled as a set of (tightly) coupled partial differential equations (PDE), with each PDE modeling a specific physical phenomenon – for instance, Navier-Stokes that models flow phenomena and the advective heat equation that models thermal phenomena. This gap exists because coupling different types of PDEs is inherently challenging and computationally expensive. Solving multiphysics problems requires sophisticated numerical methods and substantial computational resources to accurately simulate each subproblem and capture the interactions between various physical phenomena.

**FlowBench** seeks to enable the ML community to build the next generation of SciML neural PDE solvers by filling these gaps – complex geometries and multiphysics phenomena. In particular, **FlowBench** offers:

- **Flow across complex geometries:** We simulate flow and thermal-flow phenomena across a wide range of complex shapes – both parametric and non-parametric – in 2D and 3D. These include simple shapes like ellipses, more complex blobs, and geometries like insects, animals, and birds. Flow across this spectrum of complex objects exhibits a rich array of vortex formation, flow separation, and a range of lift and drag profiles. The diverse shapes and the flow interactions with them provide a rich and intricate dataset for training geometry-aware SciML solvers that can be used for various applications.
- **Multiphysics simulations:** For each shape, we perform a variety of simulations representing flow (i.e., incompressible Navier-Stokes, NS) as well as thermal flow (i.e., coupled Navier-Stokes and Heat Transfer using the Boussinesque coupling, NS-HT) scenarios. Here, we span steady-state and transient behaviors, offering a comprehensive dataset to test and evaluate solvers under various scenarios.
  - For the steady-state case, we consider a variant of the canonical fluid dynamics problem of lid-driven cavity flow (LDC), which is an example of internal flow. The cavity has an object (complex geometry) placed inside (see [Figure 8](#)). Additionally, we consider a temperature difference between the cavity walls and the object, producing thermal and flow coupling. For each object, we simulate across a range of Reynolds number  $Re \in [10^1, 10^3]$ , and Grashof number  $Gr \in [10^1, 10^7]$ . This range offers a variety of forced ( $Gr/Re^2 \approx 0.1$ ), mixed, and free convection ( $Gr/Re^2 \approx 10$ ) scenarios. In this case, there are nearly 9000 unique samples for 2D and 500 unique samples for 3D.
  - For the transient case, we consider a variant of another canonical fluid dynamics problem – flow past a bluff object (FPO) – which is an example of external flow. Here, we consider flow moving past a stationary (complex geometry) object that is placed in a large domain. The fluid exhibits various intricate time-dependent patterns as it moves past the object. For each object, we simulate across a range of Reynolds number  $Re \in [10^2, 10^3]$ , offering an array of vortex-shedding frequencies and other time-dependent patterns. In this case, there are over 1000 unique samples in 2D.

We provide field data of velocity, pressure, and temperature for each of the over  $10K$  simulations. Additionally, we provide summary engineering features, including the coefficient of lift ( $C_L$ ), drag ( $C_D$ ), and the average heat transfer (Nusselt number,  $Nu$ ) from the surface of the complex object. We provide our dataset on huggingface at <https://huggingface.co/datasets/BGLab/FlowBench/tree/main> as a benchmark for others interested in the development and evaluation of SciML models.

- **Benchmark metrics and comparisons:** Besides the dataset, we also include workflows to train select group of neural operators – Fourier Neural Operators (FNO), Convolutional Neural Operators (CNO), and Deep Operator Networks (DeepONets). Neural operators can be trained using our code for the steady-state case and study the comprehensive evaluation metrics that we recommend in [Section 4.6](#). We also suggest a hierarchy of in-distribution and out-of-distribution tests to evaluate the generalizability of these models.

The dataset (consisting of over  $10K$  samples), evaluation metrics, workflows, and trained models together make **FlowBench** a valuable tool for the ML community to create SciML solvers of coupled phenomena involving complex geometries. All data, visualizations, models, and model evaluations are available on the [FlowBench website](#).

**Table 1:** Comparison of incompressible flow simulation data in Curated Dataset, Graph-Mesh, CFDBench, AirfRANS, PDEBench, Eagle, MegaFlow, Fluid Flow Dataset, ScalarFlow, WindsorML, AhmedML, DrivAerML, BubbleML, and our FlowBench. An orange check mark indicates a restricted family of shapes (Airfoil: AirfRANS and Graph-Mesh; Cylinder: CFDBench, Cylinders and Ellipses: MegaFlow; Drone geometry: Eagle; Windsor bodies: WindsorML; Ahmed bodies: AhmedML; DrivAer bodies: DrivAerML).

Name	Dim.	Size	Geometry	# Geometries	Multiphysics
Curated Dataset (McConkey et al., 2021)	2	116	✗	0	✗
Graph-Mesh (Bonnet et al., 2022b)	2	230	✓	Few airfoil shapes	✗
CFDBench (Luo et al., 2024)	2	739	✓	1	✗
AirfRANS (Bonnet et al., 2022a)	2	1000	✓	375 airfoil shapes	✗
PDEBench (Takamoto et al., 2023)	2	1000	✗	0	✗
Eagle (Janny et al., 2023)	2	1200	✓	600 drone shapes	✗
MegaFlow (Xu et al., 2023)	2	3000	✓	3000 ellipses and circles	✗
Fluid Flow (Jakob et al., 2020)	2	8000	✗	0	✗
ScalarFlow (Eckert et al., 2019)	3	100	✗	0	✓
WindsorML (Ashton et al., 2024a)	3	355	✓	355 Windsor bodies	✗
AhmedML (Ashton et al., 2024b)	3	500	✓	500 Ahmed car bodies	✗
DrivAerML (Ashton et al., 2024c)	3	500	✓	500 DrivAer car bodies	✗
BubbleML (Hassan et al., 2023)	2,3	79	✗	0	✓
FlowBench (ours)	2,3	<b>10,000</b>	✓	300 different shapes	✓

## 2 Related Work

Most publicly available benchmark datasets are summarized and compared in Table 1. PDEBench (Takamoto et al., 2023) presents a dataset consisting of simulations run for a wide range of PDEs, not just flow physics. It includes simulations for both compressible and incompressible Navier-Stokes problems, in both two and three dimensions. CFDBench (Luo et al., 2024) is a fluid flow-focused dataset containing a total of 739 cases. These cases are distributed across various setups: lid-driven cavity (with varying density and viscosity over 25 different length and width combinations), tube flow (varying density, viscosity, and geometry, with 50 different inlet velocity conditions), dam flow, and cylinder flow specialties. MegafLOW2D (Xu et al., 2023) is a comprehensive collection of 3000 cases for 2D Navier-Stokes problems. Each case features different geometrical configurations, including circles, ellipses, and nozzles. All simulations are performed at a fixed Reynolds number of 300.

McConkey et al. (2021) focuses exclusively on 2D turbulence simulations, covering 841 different cases across 29 flow scenarios such as a periodic hill, square duct, parametric bump, converging-diverging channel, and curved backward-facing step. Each scenario includes 29 simulations with varying parameters, such as Reynolds numbers, producing 841 data samples. The Graph-Mesh (Bonnet et al., 2022b) dataset simulates incompressible 2D Navier-Stokes problems at high Reynolds numbers (greater than  $10^6$ ) exclusively for Aerofoils. It varies the angle of attack, inlet velocity, Reynolds number, and Mach number for a total of 230 simulations. A distinguishing feature is the calculation of drag and lift after fitting the SciML models, not just the solutions. WeatherBench (Rasp et al., 2020) condenses raw weather data from the ERA5 dataset by reducing resolution levels to fit within a GPU. Along with discrete weather-based variables, WeatherBench runs simulations for a range of  $u, v$  velocities, temperature, and vorticity. It uses the entire world as its grid and provides both 2D and 3D data post-processed from the ERA5 dataset.

ScalarFlow (Eckert et al., 2019) is the multiphysics simulation dataset coupling Navier-Stokes with density distribution solving, focusing on buoyancy-driven smoke plume reconstructions. It provides volumetric 3D flow reconstructions for complex buoyancy-driven flows transitioning to turbulence. The Fluid Flow Dataset (Jakob et al., 2020) contains 8000 unsteady 2D fluid flow simulations, each with 1001 time steps, parameterized by Reynolds number. BubbleML (Hassan et al., 2023) is another multiphysics dataset to couple two-phase

Navier-Stokes and energy equations. It includes 79 simulations covering nucleate pool boiling, flow boiling, and sub-cooled boiling under various gravity conditions, flow rates, sub-cooling levels, and wall superheat. EAGLE (Janny et al., 2023) is a large (over 1.1M 2D simulations) dataset that models the airflow generated by 2D UAV moving in a 2D environment with different boundary geometries. This dataset captures the highly turbulent flow consisting of non-periodic eddies induced by the varying geometries.

We identify limitations of these existing datasets that motivate the curation of FlowBench.

- **Restricted Geometries:** Many datasets are confined to simple or canonical geometries such as airfoils, cylinders, or standard shapes. For instance, AirFRANS focuses exclusively on 375 airfoil shapes, while CFDBench simulates only one geometry type (cylinders). Similarly, datasets like MegaFlow and AhmedML emphasize restricted geometry families like ellipses and Ahmed bodies, respectively. These datasets lack the diversity needed to explore flow dynamics over highly complex or non-parametric geometries, which are common in real-world engineering problems.
- **Limited Multiphysics Coverage:** While datasets like ScalarFlow and BubbleML include multiphysics simulations, their scope is narrow, focusing only on specific phenomena like buoyancy-driven flows or two-phase boiling. Most other datasets, including PDEBench and Eagle, completely omit the coupling of flow with thermal or other physical phenomena, which are critical in areas such as thermal management and environmental modeling.
- **Small Dataset Size:** The scale of existing datasets often limits their utility for training data-intensive machine learning models. For example, CFDBench includes only 739 cases, and Graph-Mesh contains 230 cases, both of which are inadequate for benchmarking modern neural PDE solvers that rely on large-scale data to achieve generalizability and robustness.
- **Lack of Diverse Flow Regimes:** Many datasets focus on specific flow regimes (e.g., laminar or turbulent) and neglect transitional flows or mixed convection scenarios. For instance, MegaFlow provides simulations at a fixed Reynolds number, missing the opportunity to examine the impact of varying flow conditions on solver performance.
- **Resolution and Usability:** While some datasets include high-resolution simulations, they often lack multi-resolution data that can facilitate super-resolution studies or efficiency testing. Additionally, the absence of consistent and reusable formats for dataset access hampers usability in machine learning workflows.

### 3 Mathematical Formulation of Simulation Approach for Creating FlowBench

Simulating incompressible flows over complex geometries is challenging, especially with traditional body-fitted mesh methods. These approaches often struggle to handle intricate shapes, as creating a mesh that accurately conforms to irregular geometries requires significant manual adjustment to ensure stability and convergence, making large-scale simulations labor-intensive and costly (Verzicco, 2023; Mittal and Seo, 2023; George et al., 1991; Owen, 1998; Persson, 2005; Boelens et al., 2009; Kurtulus, 2015). Although advances in automated mesh generation (George et al., 1991; Simonovski and Cizelj, 2011; Beaufort et al., 2020) have eased some difficulties, achieving high-quality meshes for complex structures without human intervention remains elusive.

A promising alternative is the Immersed Boundary Method (IBM), first developed by Peskin (Peskin, 1972), which enables simulations across complex geometries by embedding the shape within a fixed mesh. The Shifted Boundary Method (SBM), a robust variant of IBM, further improves this approach by imposing boundary conditions on a surrogate rather than the actual boundary. In SBM, the surrogate boundary aligns with a Cartesian mesh, and boundary conditions are adjusted using Taylor expansions, effectively transforming the problem to resemble a body-fitted domain while streamlining mesh generation (Main and Scovazzi, 2018b,a; Atallah et al., 2020, 2021; Yang et al., 2024a). This surrogate is carefully chosen to minimize numerical errors, ensuring accuracy and stability. Particularly effective when applied with adaptive octree meshes, SBM allows local refinement near complex geometries, reduces computational demands, and enhances scalability, making it ideal for complex fluid flow and multiphysics simulations across varied Reynolds numbers and configurations.

For completeness, we briefly detail the Finite Element approach – specifically, the additional terms needed in Shifted Boundary Method over the standard stabilized finite element formulation – used to create the dataset. In the formulations that follow,  $w_i$  is defined as the test function for the  $i$ -th component of the momentum equation, where  $i$  runs from 0 to  $\text{DIM} - 1$  (with  $\text{DIM}$  representing the spatial dimension), corresponding to the  $x$ -,  $y$ -, and  $z$ -directions. Similarly,  $u_i$  denotes the corresponding velocity components. These conventions are standard in the FEM literature. We refer the interested reader to (Yang et al., 2024b) for a detailed discussion of the approach, implementation details, and exhaustive validation of the accuracy of the SBM approach.

### 3.1 Formulation of SBM for Navier Stokes

In SBM, additional terms are incorporated into the standard stabilized finite element formulations for solving the Navier-Stokes equations (see Section 2 in (Yang et al., 2024b)). These terms include a *consistency* term (which appears due to integration by parts operation), an *adjoint consistency* term (which is included to ensure optimal convergence rates), and the *penalty* term (which ensures that as the mesh size is reduced, the boundary conditions asymptote to the true boundary conditions) :

$$\begin{aligned}
\widetilde{B}_{NS}^{VMS} = & B_{NS}^{VMS} - \underbrace{\left\langle w_i^{c,h}, \frac{1}{Re} \left( \frac{\partial u_i^{c,h}}{\partial x_j} + \frac{\partial u_j^{c,h}}{\partial x_i} \right) \tilde{n}_j - p^{c,h} \tilde{n}_i \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Consistency term}} \\
& - \underbrace{\left\langle \frac{1}{Re} \left( \frac{\partial w_i^{c,h}}{\partial x_j} + \frac{\partial w_j^{c,h}}{\partial x_i} \right) \tilde{n}_j + q^{c,h} \tilde{n}_i, u_i^{c,h} + \frac{\partial u_i^{c,h}}{\partial x_j} d_j \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Adjoint consistency term}} \\
& + \underbrace{\frac{\beta}{h \cdot Re} \left\langle w_i^{c,h} + \frac{\partial w_i^{c,h}}{\partial x_j} d_j, u_i^{c,h} + \frac{\partial u_i^{c,h}}{\partial x_j} d_j \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Penalty term}}, \tag{1}
\end{aligned}$$

and

$$\begin{aligned}
\widetilde{F}_{NS}^{VMS} = & F_{NS}^{VMS} - \underbrace{\left\langle \frac{1}{Re} \left( \frac{\partial w_i^{c,h}}{\partial x_j} + \frac{\partial w_j^{c,h}}{\partial x_i} \right) \tilde{n}_j + q^{c,h} \tilde{n}_i, g_i \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Adjoint Consistency Term}} \\
& + \underbrace{\frac{\beta}{h \cdot Re} \left\langle w_i^{c,h} + \frac{\partial w_i^{c,h}}{\partial x_j} d_j, g_i \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Penalty Term}}, \tag{2}
\end{aligned}$$

where  $g_i$  is the boundary condition we want to apply,  $\tilde{n}_i$  is the unite outward-pointing normal,  $h$  is the element size,  $\beta$  is the penalty parameter for the Navier-Stokes equation,  $B_{NS}^{VMS}$  is the bilinear weak form for Navier-Stokes without SBM, and  $F_{NS}^{VMS}$  is the linear weak for Navier-Stokes form without SBM. The formulation for Navier-Stokes with SBM can be expressed as:

$$\widetilde{B}_{NS}^{VMS} - \widetilde{F}_{NS}^{VMS} = 0. \tag{3}$$

Readers interested in the derivation of the consistency term, adjoint-consistency term, and penalty term for weakly enforcing shifted boundary conditions are referred to [Section A.1](#).

### 3.2 Formulation of SBM for Heat Transfer

Similar to the Navier-Stokes equations, we use SBM to apply Dirichlet boundary conditions in the energy equation. Essentially, we use SBM to set the temperature to a desired value ( $T_D$ ) on the geometries. The

formulation for energy equation with SBM is:

$$\widetilde{B_{HT}^{VMS}} - \widetilde{F_{HT}^{VMS}} = 0, \quad (4)$$

with

$$\begin{aligned} \widetilde{B_{HT}^{VMS}} = & B_{HT}^{VMS} - \underbrace{\frac{1}{Pe} \left\langle l^{c,h}, \frac{\partial T^{c,h}}{\partial x_j} \tilde{n}_j \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Consistency term}} - \underbrace{\frac{1}{Pe} \left\langle \frac{\partial l^{c,h}}{\partial x_j} \tilde{n}_j, T^{c,h} + \frac{\partial T^{c,h}}{\partial x_j} d_j \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Adjoint consistency term}} \\ & + \underbrace{\frac{\alpha}{h \cdot Pe} \left\langle l^{c,h} + \frac{\partial l^{c,h}}{\partial x_j} d_j, T^{c,h} + \frac{\partial T^{c,h}}{\partial x_j} d_j \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Penalty term}}, \end{aligned} \quad (5)$$

and

$$\begin{aligned} \widetilde{F_{HT}^{VMS}} = & F_{HT}^{VMS} - \underbrace{\frac{1}{Pe} \left\langle \frac{\partial l^{c,h}}{\partial x_j} \tilde{n}_j, T_D \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Adjoint consistency term}} + \underbrace{\frac{\alpha}{h \cdot Pe} \left\langle l^{c,h} + \frac{\partial l^{c,h}}{\partial x_j} d_j, T_D \right\rangle_{\tilde{\Gamma}_{D,h}}}_{\text{Penalty term}}, \end{aligned} \quad (6)$$

where  $T_D$  is the boundary condition we want to apply,  $\alpha$  is the penalty parameter for energy equation,  $Pe$  is Peclet number, which is  $0.7Re$  inside our simulations,  $B_{HT}^{VMS}$  is the bilinear weak form for Heat Transfer without SBM, and  $F_{HT}^{VMS}$  is the linear weak form for Heat Transfer without SBM.

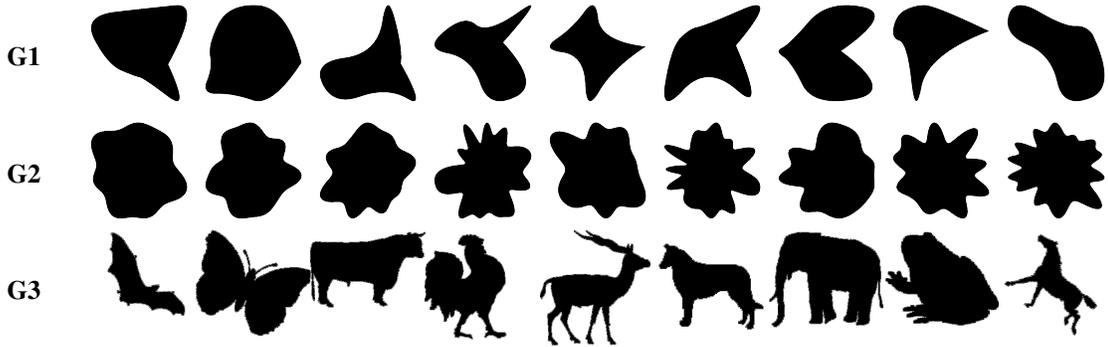
## 4 FlowBench

### 4.1 Geometries

Our dataset includes three distinct categories of geometries, namely **G1**, **G2**, and **G3** as illustrated in [Figure 2](#). The first set of geometries, **G1**, consists of parametric shapes generated using Non-Uniform Rational B-Splines (NURBS) curves. NURBS are mathematical representations used in computer graphics and CAD systems to generate and represent curves and surfaces. They offer great flexibility and precision in modeling complex shapes. Each NURBS curve is defined by a set of control points, the degree of the basis function, and knot vectors ([Piegl and Tiller, 2012](#)). We use a uniform knot vector with a second-order (quadratic) basis function, which remains fixed. However, the positions of eight control points are randomly varied to produce a variety of curves. We ensure that the shapes are smooth and do not have any discontinuities or self-intersections. All shapes are normalized to fit within the unit hypercube,  $[0, 1]^2$ . We provide the shapes as well as code for recreating these geometries.

The next set of geometries, **G2**, consists of parametric shapes generated using spherical harmonics ([Wei et al., 2018](#)). We randomly select  $N = 8, \dots, 15$  harmonics with amplitudes  $(a_n, b_n)$  ranging from 0 to 0.2. The radial function  $r(t) = 0.5 + \sum_{n=1}^N (a_n \cos(nt) + b_n \sin(nt))$  then defines the shape; and is computed at 500 evenly spaced points in  $t \in [0, 2\pi]$ . We normalize  $r(t)$  so that any surface point is within a distance of 0.5 from the center of the shape,  $r(t) = 0.5 \left( \frac{r(t)}{r_{\max}} \right)$ . We provide shapes as well as code for recreating these geometries.

The last set of geometries, **G3**, consist of non-parametric shapes sampled from the grayscale dataset in SkelNetOn ([Demir et al., 2019](#); [Atienza, 2019](#)). We apply a Gaussian blur filter with a scale of 2 to smoothen out some of the thin features of the object. This choice of  $\sigma = 2$  was empirically determined to balance the smoothing of small-scale, potentially noisy features—which, if left unfiltered, could lead to jagged or irregular contours during contour extraction and meshing—against the need to preserve the primary geometric characteristics of the shapes. This ensures the shape remains consistent across the three resolutions we provide data for. We then scale the shape to ensure it is contained in the unit hypercube  $[0, 1]^2$ .



**Figure 2:** Examples of the diverse and complex geometries in *FlowBench* using 9 samples from each of the three groups. The first row corresponds to geometries from the nurbs group G1, the second row to the spherical harmonics group G2, and the third row to the skelneton group G3.

## 4.2 Simulation Framework and Compute effort

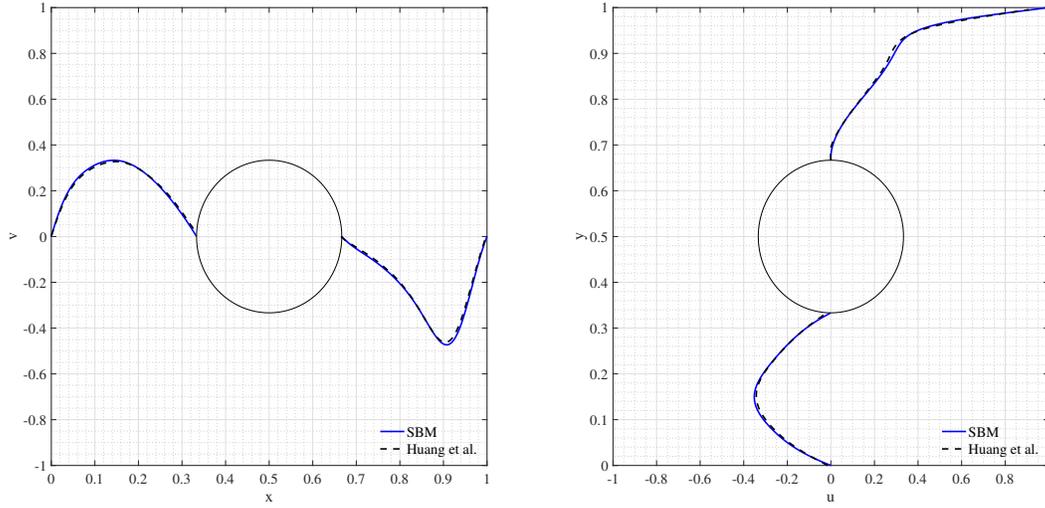
Our simulation framework is a highly parallel octree-based CFD and multiphysics code. We account for complex geometries using a robust variant of the immersed boundary method, called the Shifted Boundary Method (SBM) (Main and Scovazzi, 2018b,a). For readers interested in our framework’s SBM formulation and implementation, please refer to (Yang et al., 2024a,c). Our octree-based framework has been extensively validated and used in various applications like industrial-scale CFD simulations (Saurabh et al., 2021a), two-phase flow (Khanwale et al., 2023), buoyancy-driven flows (Xu et al., 2019), and coupled multiphysics application (Kim et al., 2024). For additional validation of the incompressible flow and thermal incompressible flow solving using SBM, please refer to Section 4.3 in the appendix. All the 2D LDC cases are simulated using a uniform mesh of  $512 \times 512$ . All the 2D FPO cases are simulated using an adaptive mesh that resolved to the Kolmogorov length scale close to the object. All the 3D LDC cases are simulated using an adaptive mesh that resolved to  $2 \times$  the Kolmogorov scale. Illustrations of the mesh are provided in Section A.2. We deployed this framework on one of the largest academic supercomputing clusters in the US, called TACC Frontera (Stanzione et al., 2020). These simulations required about 65K node hours of runtime.

## 4.3 Validation of the numerical framework used to create the dataset

We rigorously validate our SBM framework using results from existing literature. See (Yang et al., 2024c) where we have an exhaustive set of validations. In particular, for cases lacking validation within the literature, such as complex geometry simulations, we compare our SBM results with those obtained from the Boundary-Fitted Method (BFM). This comparison is especially valuable as the BFM requires a tedious and time-consuming process of manually creating body-fitted meshes to capture the geometry accurately. Crafting these meshes is often challenging and labor-intensive, demanding extensive domain knowledge and iterative refinement to ensure mesh quality and alignment with complex boundaries. For completeness, we showcase a representative set of validation examples of the SBM approach.

**LDC (Lid-driven cavity flow), NS:** We simulate a disk with a diameter  $D = \frac{L}{3}$  placed at the center of the lid-driven cavity, where  $L$  is the chamber’s box length. Our CFD results match against detailed simulations available in literature (Huang and Lim, 2020) for this canonical problem, as shown in Figure 3.

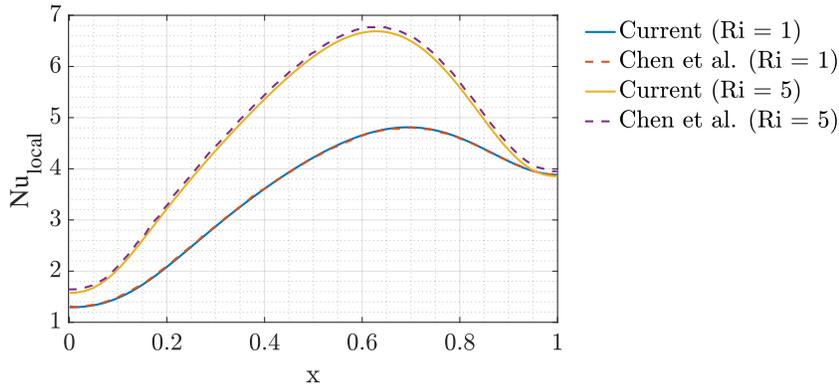
**LDC (Lid-driven cavity flow), NSHT:** To validate our multiphysics simulation framework, we select a case from (Chen et al., 2020) to compare with. Here, a heated circle is placed at the center of the chamber, with a radius of  $0.2L$ , where  $L$  represents the length of the chamber. We evaluate the local Nusselt number on the bottom wall and report an excellent match with (Chen et al., 2020) in Figure 4 at two distinct operating conditions.



(a) Horizontal velocity profile.

(b) Vertical velocity profile.

**Figure 3:** Comparison of velocity profiles for lid-driven cavity flow with a circular disk at  $Re = 1000$  against results from the literature.

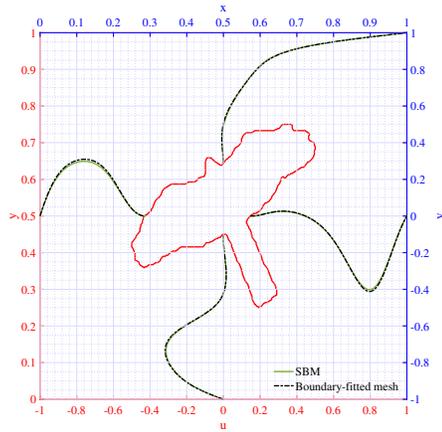


**Figure 4:** Comparison of the Nusselt number profile along the bottom wall of the domain. The simulation results are validated against data from (Chen et al., 2020) to demonstrate the accuracy of the current multiphysics (NSHT) model.

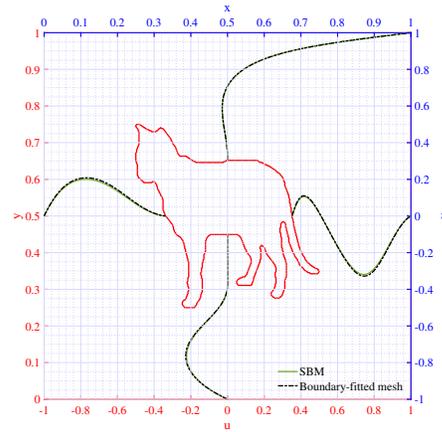
**FPO (Flow past object):** For the flow past (bluff) geometries, we tested several shapes. We first compared against the canonical case of flow past a circular disk and evaluated at two different Reynolds numbers: 100 and 1000. Our drag coefficients and Strouhal numbers matched well with the literature, as shown in Table 2 and Table 3.

**Complex Geometries:** To evaluate our framework’s performance with complex geometries, we first present further validation of the Shifted Boundary Method (SBM) by comparing it with the Boundary-Fitted Method (BFM) in the context of lid-driven cavity simulations shown in Figure 5. For the boundary-fitted mesh, we used a mesh size of approximately  $\frac{1}{29}$ , matching the element size employed in the SBM simulations.

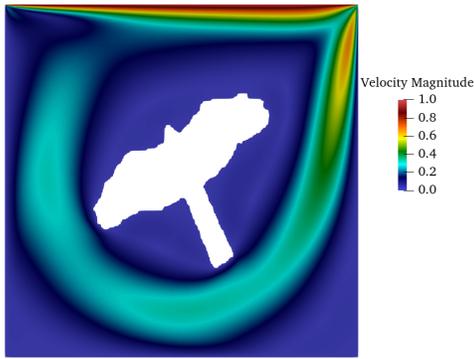
Additionally, we validated the solver’s performance for complex geometries by comparing our results to previous research. Specifically, we simulated flow past a D-shaped cylinder at various Reynolds numbers and compared the results with those from (Shao et al., 2020), as shown in Table 4. The flow visualizations are



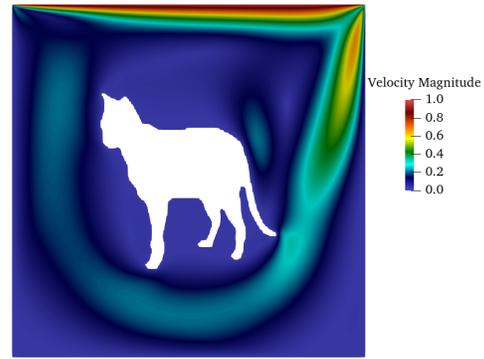
(a) SBM and BFM for the bird at  $Re = 912$ .



(b) SBM and BFM for the cat at  $Re = 729$ .



(c) Flow visualization for the bird at  $Re = 912$ .



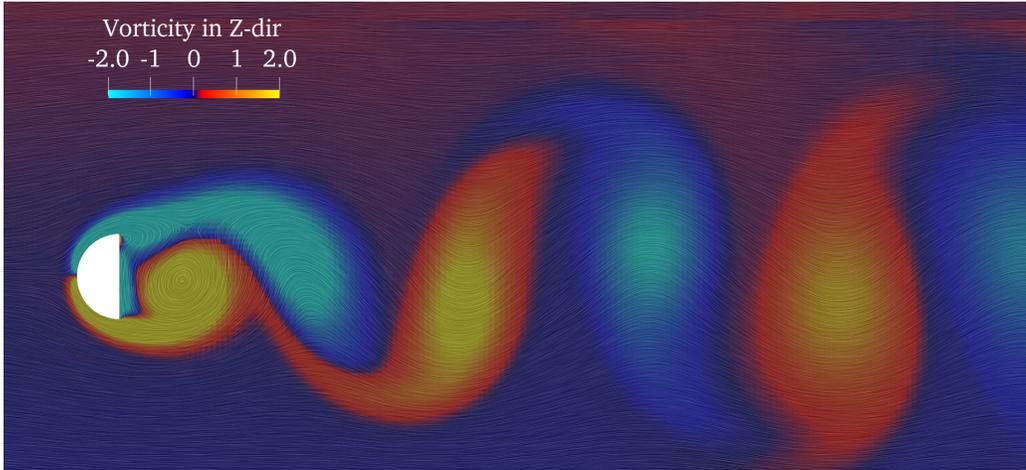
(d) Flow visualization for the cat at  $Re = 729$ .

**Figure 5:** Comparison between results from BFM and SBM in lid-driven cavity simulations for the bird and cat cases.

**Table 2:** Comparison of drag coefficients and strouhal numbers for flow past a 2D cylinder at  $Re = 100$ .

	<b>Cd</b>	<b>St</b>
Current	1.35	0.167
Mittal <i>et al.</i> (Mittal <i>et al.</i> , 2008)	1.35	0.165
Henderson <i>et al.</i> (Henderson, 1995)	1.35	-
Luo <i>et al.</i> (Luo <i>et al.</i> , 2009)	1.35	0.159
Kamensky <i>et al.</i> (Kamensky <i>et al.</i> , 2015)	1.386	0.170
Main <i>et al.</i> (Main and Scovazzi, 2018a)	1.36	0.169
Kang <i>et al.</i> (Kang and Masud, 2021)	1.374	0.168

illustrated in Figure 6. Our findings closely align with the literature, demonstrating that an increase in Reynolds number corresponds to a higher Strouhal number. Readers interested in further validating our simulation framework for complex geometries are encouraged to read (Yang *et al.*, 2024c) and visit our [website](#).



**Figure 6:** Flow visualization for a D-shaped cylinder at  $Re = 100$ .

**Table 3:** Comparison of drag coefficients and strouhal numbers for flow past a 2D cylinder at  $Re = 1000$ .

	<b>Cd</b>	<b>St</b>
Current	1.52	0.239
Mittal <i>et al.</i> (Mittal <i>et al.</i> , 2008)	1.45	0.230
Henderson <i>et al.</i> (Henderson, 1995)	1.51	-
Luo <i>et al.</i> (Luo <i>et al.</i> , 2009)	1.56	0.235
Cheny <i>et al.</i> (Cheny and Botella, 2010)	1.61	0.251
Jester <i>et al.</i> (Jester and Kallinderis, 2003)	1.51	0.25

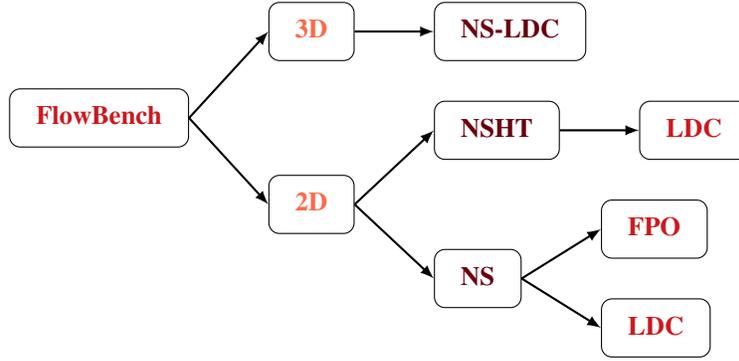
**Table 4:** Strouhal numbers for flow past a D-shaped cylinder

<b>Re</b>	<b>Current</b>	<b>(Shao <i>et al.</i>, 2020)</b>
46	0.1363	0.1316
100	0.2006	0.1910
160	0.2112	0.2042

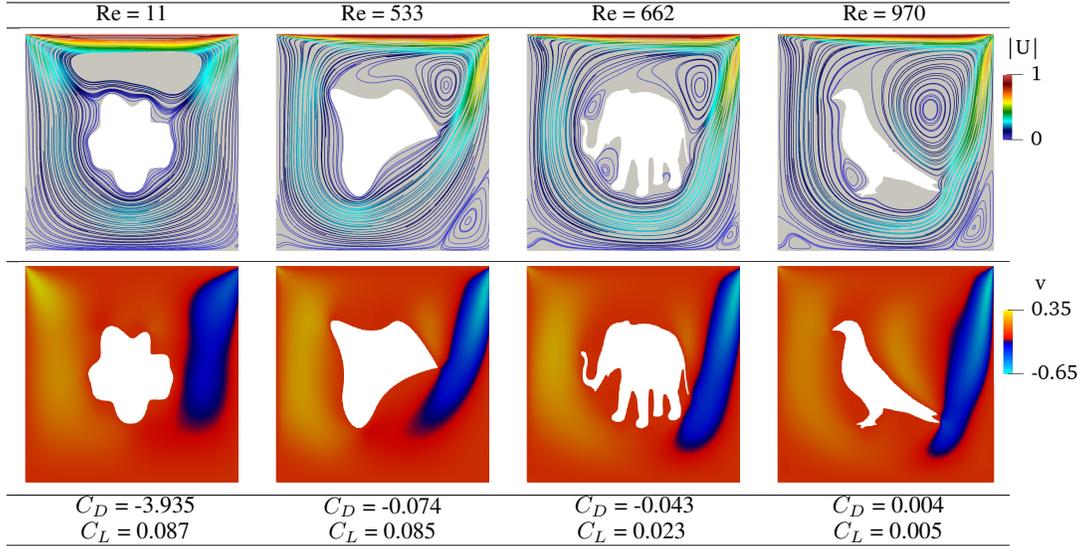
#### 4.4 Flow Physics Problems: Domain, Boundary Conditions, and Outputs

Our **FlowBench** dataset presents a comprehensive suite of simulations across 300 complex geometries, capturing a diverse range of physical phenomena, including Navier-Stokes and thermally coupled flows. This dataset addresses a spectrum of canonical problems—such as flow past obstacles and lid-driven cavity flow with internal objects—spanning both two-dimensional and three-dimensional domains. The dataset accommodates various flow regimes, encompassing forced and natural convection as well as transitions from laminar to turbulent flows. A structured, hierarchical overview of the dataset is provided in [Figure 7](#), illustrating its extensive applicability in benchmarking complex fluid dynamics and heat transfer scenarios.

In the 2D LDC setup, a square features three stationary walls and one moving lid, with a domain size of  $[0, 2] \times [0, 2]$ . An object is placed in the middle of the flow within the chamber. Examples of LDC simulations showing streamlines and y-direction velocities, along with the drag coefficient ( $C_D$ ) and lift coefficient ( $C_L$ ) values, are shown in [Figure 8](#). Increasing the Reynolds number brings the vortices closer to the right wall, with additional vortices forming at the bottom-left and bottom-right corners. With increasing Reynolds numbers, the smaller viscous forces acting on the geometries decrease both  $C_D$  and  $C_L$ .

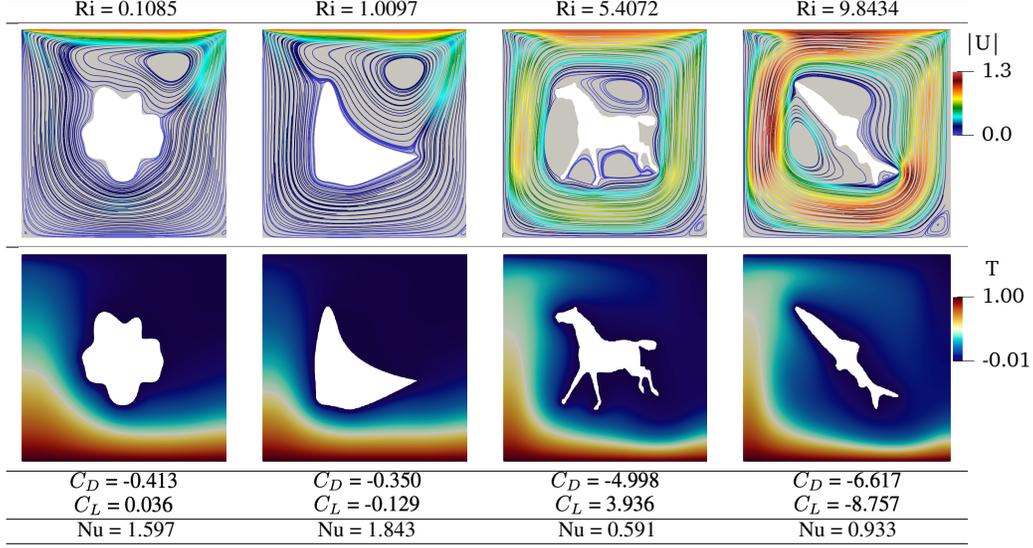


**Figure 7:** Hierarchical structure of the FlowBench framework. The diagram illustrates the categorization into 2D and 3D simulations, with further subdivisions based on Navier-Stokes (NS) and Navier-Stokes coupled with Heat Transfer (NSHT) under different test cases, such as the Lid-Driven Cavity (LDC) and Flow Past Object (FPO). Each category represents a specific configuration for benchmarking complex fluid dynamics and heat transfer simulations.



**Figure 8:** Drag coefficients ( $C_D$ ), and lift coefficients ( $C_L$ ) for different shapes and different Reynolds numbers in pure-NS LDC simulations.

For the 2D Navier-Stokes and heat transfer (NSHT) lid-driven cavity (LDC) problem, the bottom wall is set to a temperature of 1, while the top wall is set to 0. The left and right walls have zero-flux temperature boundary conditions. The object surface is set to a temperature of 0. This setup produces a combination of forced convection (flow due to the moving lid) and Rayleigh-Benard instabilities (flow due to buoyancy-driven natural convection). Examples of NSHT-LDC simulations with fixed Reynolds number, showing streamlines, temperatures, and the values of  $C_D$ ,  $C_L$ , and  $Nu$ , are presented in Figure 9. We use a non-dimensional number, Richardson number, defined as  $Ri = Gr/Re^2$ , representing the ratio between buoyancy and inertial force. As  $Ri$  increases, we observe higher values of  $C_D$  and  $C_L$ . A higher Richardson number indicates that buoyancy effects are more significant than the forced flow. The heated fluid rises more strongly, creating greater circulation within the chamber. The increased circulation results in stronger forces acting on the object, leading to higher  $C_D$  and  $C_L$ .

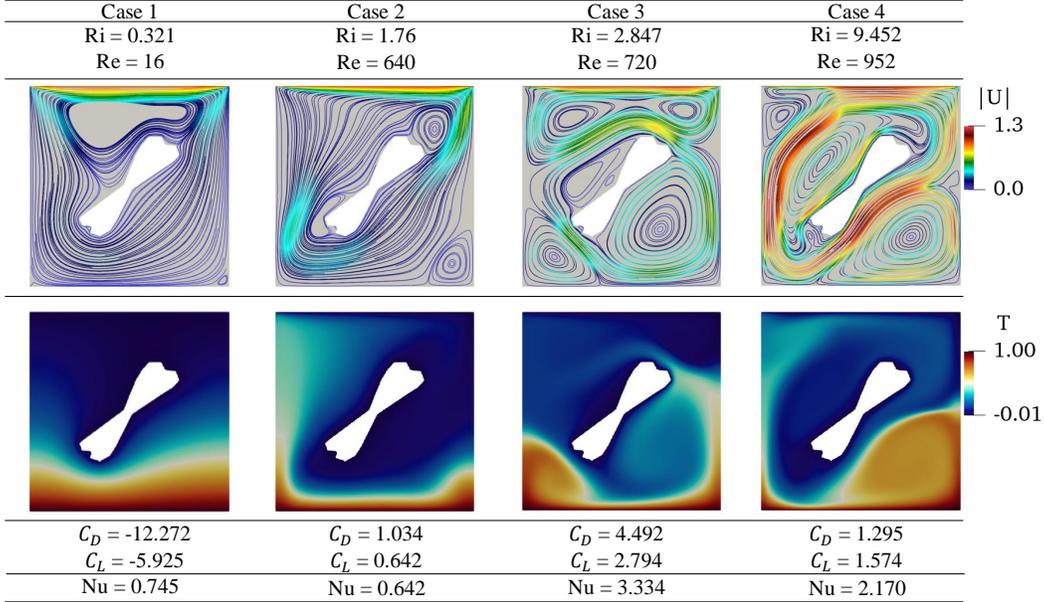


**Figure 9:** Drag coefficients ( $C_D$ ), lift coefficients ( $C_L$ ), and Nusselt numbers ( $Nu$ ) for different shapes and different Richardson numbers in NSHT LDC simulations (fixed  $Re = 100$ ).

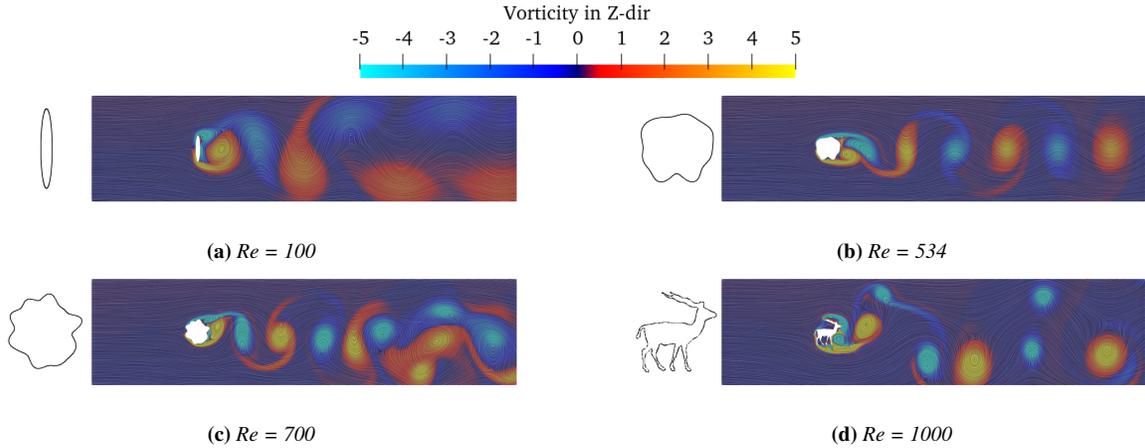
Apart from the NSHT-LDC simulations with a fixed Reynolds number, we also conduct simulations with randomly varying Reynolds and Richardson numbers. We pick up one geometry to demonstrate in Figure 10. This showcase examines four cases with progressively increasing Reynolds and Richardson numbers. In Case 1, the high-speed flow predominantly occurs at the top of the geometry, resulting in a negative  $C_D$  due to the flow applying force on the upper part of the geometry. As we increase the Reynolds number in Case 2, the flow shifts towards the bottom part of the geometry, applying force in the positive y-direction, leading to a change in the sign of  $C_D$  from Case 1 to Case 2. The same as pure-NS flow, we observe a high-speed jet in the y-direction near the left-hand side wall due to the higher Reynolds number in Case 2. The magnitudes of  $C_D$  and  $C_L$  decrease from Case 1 to Case 2 because the viscous forces are reduced as the Reynolds number increases. Given that the Reynolds numbers in Case 2 and Case 3 are close, the reduction in force due to the increasing Reynolds number is mitigated, and the Richardson number plays a crucial role in enhancing flow circulation, leading to an increase in  $C_D$  and  $C_L$ . However, from Case 3 to Case 4, the drag force decreases due to the higher Reynolds number, as observed in Figure 8. Additionally, as the Richardson number increases, the temperature convect to higher positions, progressing from Case 1 to Case 4.

For the FPO setup, we consider a domain  $[0, 64] \times [0, 16]$  with the object placed at (6, 8). A parabolic velocity inlet boundary condition is applied on the left, no-slip boundary conditions are applied on the top and bottom walls, and a zero pressure boundary condition is applied on the right. The large domain size for the FPO problem is chosen to capture as much physical detail as possible in our dataset. We report a smaller cropped-out region of size  $[0, 16] \times [0, 4]$  from this dataset, representing a tradeoff between dataset size and the amount of physics captured. Figure 11 shows time snapshots of representative shapes showing vortex shedding, as the flow rotates and stretches around the object. Movies of the time-dependent flow can be seen on the FlowBench website.

For the simulations we performed in 2D, we ensure the fluid mesh's resolution is fine enough for the Direct Numerical Simulation (DNS) simulations. DNS is an expensive high-fidelity computational approach that solves the full Navier-Stokes equations, resolving all scales of motion in a fluid without additional turbulence modeling. The Kolmogorov scale is crucial in turbulence theory as it represents the smallest length scale that must be resolved in DNS to capture all aspects of turbulence accurately. This scale can be expressed in terms of the Reynolds number as  $\eta \sim L/Re^{3/4}$ , where  $L$  is the characteristic large-eddy length scale. In our



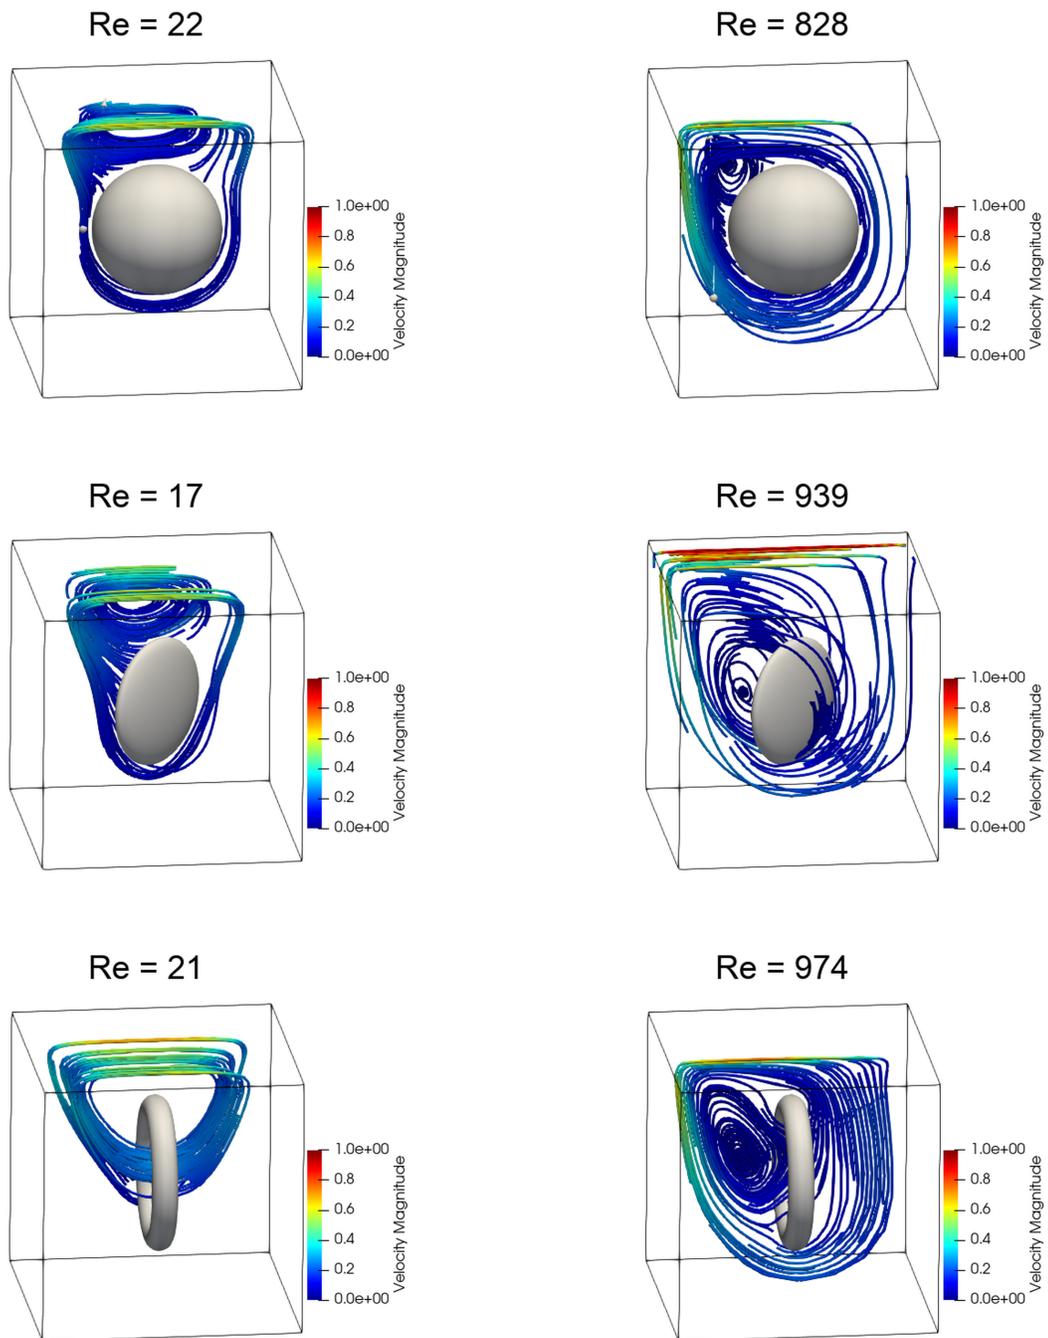
**Figure 10:** Drag coefficients ( $C_D$ ), lift coefficients ( $C_L$ ), and Nusselt numbers ( $Nu$ ) for different shapes and different Richardson numbers in NSHT LDC simulations (random Reynolds number).



**Figure 11:** Time snapshots of vorticity for the FPO problem, illustrating vortex shedding profiles across four samples of geometries and Reynolds numbers. Each sample showcases distinct flow characteristics and vortex dynamics.

2D simulation, we ensure that the mesh size is smaller than the Kolmogorov scale, calculated based on the Reynolds number, to capture the full range of turbulent motions.

The 3D LDC cases are similarly defined and simulated. In the 3D LDC setup, a cube features five stationary walls and one moving lid, with a domain size of  $[0, 2] \times [0, 2] \times [0, 2]$ . An object is placed in the middle of the flow within the chamber. Due to the computational expense of performing DNS on 3D geometries, we report 500 simulations of complex objects. Here, we focused on parametric shapes, specifically ellipsoids and tori, exhibiting various aspect ratios and orientations. Examples of LDC simulations showing streamlines and around the object is shown in Figure 12. As Reynolds number increase, more circulation regions around the object appear.



**Figure 12:** Streamlines around a torus/ellipsoid for pure Navier-Stokes LDC simulations in 3D. Each row represents one object, with a low Reynolds number result on the left and a high Reynolds number result on the right.

## 4.5 MetaData

**Input Fields:** We have provisioned the following input fields:

1. **Reynolds Number and Grashof Number:** We feed the Reynolds and Grashof numbers as a concordant (matching the dimensions of the flow domain, e.g., for 2D lid-driven cavity problem,  $512 \times 512$ ) array comprising of a single integer value everywhere.
2. **Geometry Mask:** A geometry mask ( $g$ ) is a binary representation of a shape or object within a spatial domain. Each element in the mask can have one of two values, typically 0 or 1, where: 0 indicates that the point is outside the object and 1 indicates that the point is inside the object. The geometry mask helps identify and isolate the region of interest (the object) from the background or surrounding space. This information is again packaged as a concordant array with each entry marked as 0 or 1.
3. **Signed Distance Field (SDF):** A Signed Distance Field ( $s$ ) is a scalar field that represents the shortest distance from any point in space to the surface of a given shape. The "signed" aspect of the signed distance field indicates whether the point is inside or outside the shape. A negative value indicates the point is inside the shape, while a positive value indicates the point is outside the shape. A value of zero indicates the point is exactly on the surface of the shape. The SDF provides additional information through the distance values, which is important for understanding the spatial relationship between locations and the geometric boundary. This method offers richer geometric information compared to the binary geometry mask, as it includes both positional and distance information. Yet again, this information is packaged as a concordant array, with each entry representing the nearest distance to the geometrical shape.

**Output Fields:** We are interested in obtaining field solutions (i.e., solutions at every point in the interior of the domain) for certain cardinal fields. For a 2D solution domain, these are:  $u$  - velocity in  $x$  direction,  $v$  - velocity in  $y$  direction,  $p$  - pressure and  $\theta$  - temperature. Additionally, depending on whether we are solving a steady-state or a time-dependent problem, we would have either have one final snapshot of these cardinal fields or a sequence of these fields distributed uniformly over time.

**Resolution:** We perform simulations at DNS resolution. **FlowBench** contains postprocessed simulation results at three different resolutions. This serves multiple purposes: First, it allows the community to systematically explore tradeoffs in the amount of data vs resolution vs accuracy. Second, since we perform non-interpolatory sub-sampling, this data allows the community to systematically test PDE super-resolution approaches. Finally, the lower-resolution datasets offer easier opportunities to train the data. For instance, at the time of submission, we could not train the 3D SciML models at the highest resolution available. For the steady state cases, we provide data at resolutions  $512 \times 512$ ,  $256 \times 256$ ,  $128 \times 128$  ( $\times 128$  - for 3D LDC NS). For the time dependant cases, we provide raw data at resolutions,  $1024 \times 256$ , and  $512 \times 128$ . We provide 240 snapshots for each time-dependant case. With only a light amount of postprocessing, the end user can choose from these 240 timesteps to create more managable datasets for machine learning, e.g., taking snapshots every sixth time step.

**Dataset Format:** All of our datasets are provided as numpy compressed (`.npz`) files. For each steady-state problem, we provide two `.npz` files - one for input and one for output. The input files are suffixed with the marker "`_X.npz`", and similarly, the corresponding output files are suffixed with the marker "`_Y.npz`". In 2D, each of these `.npz` files contains a 4D numpy tensor of the following form:

`[samples][number_of_channels][resolution_x][resolution_y]`

Similarly, in 3D, each of these `.npz` files contains a 5D numpy tensor of the following form:

`[samples][number_of_channels][resolution_x][resolution_y][resolution_z]`

For time-dependent problems, we provide a single file for each problem. This decision was taken to allow maximum flexibility to the end user in deciding what and how many time steps they want to use to train their models, as these time-dependent problems often take the shape of sequence-to-sequence formulations. In 2D, the resulting `.npz` files take the following form:

`[samples][number_of_time_steps][number_of_channels][resolution_x][resolution_y]`

We provide a total of over 10K samples spread across four families of datasets. [Table 9](#) in the Appendix provides a detailed formulaic description of the packaging of the input and output `numpy` tensors for each of these four families.

#### 4.6 Evaluation Metrics and Test Datasets

**Errors:** We compute the following errors to evaluate the accuracy of all the models trained on [FlowBench](#):

- *Mean Squared Error:*

$$\sum_{i=1}^n (y_i - y'_i)^2$$

- *L<sub>2</sub> Relative Error:*

$$L_2 = \sqrt{\frac{\sum_{i=1}^n (y_i - y'_i)^2}{\sum_{i=1}^n y_i'^2}}$$

- *L<sub>∞</sub> Relative Error:*

$$L_\infty = \frac{\max |y_i - y'_i|}{\max |y'_i|}$$

*Note:* In all the above definitions,  $y, y'$  and  $n$  indicate the predicted value from a model, the ground truth and the test dataset size respectively.

**Evaluation metrics:** We recommend a hierarchy of metrics ( $M1, M2, M3, M4$ ) to comprehensively assess the performance of trained models using [FlowBench](#):

- *M1: Global metrics:* Computing all three errors over the entire domain is a good primary metric to measure the accuracy of predicted velocity, pressure, and temperature fields, reflecting how closely the model’s predictions match the true values.
- *M2: Boundary layer metrics:* The errors in velocity, pressure, and temperature fields over a narrow region around the object. We define the boundary layer by considering the solution conditioned on the Signed Distance Field ( $0 \leq SDF \leq 0.2$ ). This tougher metric provides insight into the model’s accuracy in predicting near-surface phenomena, which are crucial for applications including flow diagnostics, shape design, and dynamic control.
- *M3: Property metrics:* An application-driven metric is to evaluate the accuracy of the summary statistics –  $C_D, C_L, Nu$ . The coefficients of lift and drag are critical for assessing the forces acting on the object, while the average Nusselt number predicts heat transfer rates, all of which are key for flow-thermal management and engineering applications. This metric represents spatially averaged properties.
- *M4: Residual metrics:* One can also evaluate how well the predicted field satisfies the underlying (set of) PDE. Evaluating how well continuity ( $\nabla \cdot u = 0$ ) is satisfied all over the domain is a measure of respecting the conservation of mass. Similarly, the global average of the PDE residual evaluates how well the model satisfies the underlying PDE in the domain. For a detailed explanation of the PDE residuals calculation, see [Appendix B.3](#).

**Test Datasets:** We recommend evaluating trained models on a held out dataset using the standard 80-20 random split of the prepared dataset.

Together, these metrics and datasets provide a comprehensive evaluation framework, allowing practitioners to evaluate model accuracy, physical consistency, and practical reliability.

## 5 Experiments

We report baseline results for training a suite of the most common neural PDE solvers. We explored the following Neural Operators and Foundation Models, reporting results on the 2D LDC-NS: (a) Fourier Neural Operator (FNO) (Li et al., 2021), (b) DeepONet (Lu et al., 2021), (c) Poseidon (Herde et al., 2024). Plots of field solutions for  $u$ ,  $v$ , and  $p$ , as well as training and validation losses, are available on our [website](#). For training, we adhered closely to the published code examples. All the aforementioned models were trained on a single A100 80GB GPU using the Adam optimizer with a learning rate of  $10^{-3}$  and were run for 400 epochs. The validation loss for all models converged and stabilized by 400 epochs.

**Table 5:** The mean squared errors of various neural operators trained on the 2D LDC dataset. All errors are reported on the testing dataset.

Model	Geometry	M1	M2	M3
FNO	G1	$5.8 \times 10^{-3}$	$7.1 \times 10^{-4}$	$2.7 \times 10^{-3}$
	G2	$6.8 \times 10^{-3}$	$4.9 \times 10^{-4}$	$2.9 \times 10^{-3}$
	G3	$9.7 \times 10^{-3}$	$1.2 \times 10^{-3}$	$1.3 \times 10^{-2}$
DeepONet	G1	$4.9 \times 10^{-3}$	$1.8 \times 10^{-3}$	$1.8 \times 10^{-2}$
	G2	$2.3 \times 10^{-3}$	$7.0 \times 10^{-4}$	$7.0 \times 10^{-3}$
	G3	$6.0 \times 10^{-3}$	$1.4 \times 10^{-3}$	$3.7 \times 10^{-2}$
poseidon	G1	$2.2 \times 10^{-4}$	$7.7 \times 10^{-5}$	$5.2 \times 10^{-4}$
	G2	$1.4 \times 10^{-4}$	$4.9 \times 10^{-5}$	$4.1 \times 10^{-4}$
	G3	$4.0 \times 10^{-4}$	$1.4 \times 10^{-4}$	$5.4 \times 10^{-3}$

**Table 6:** The relative  $L_2$  error of various neural operators trained on the 2D LDC dataset. All metrics are reported on the testing dataset.

Model	Geometry	M1	M2	M3
FNO	G1	$3.3 \times 10^{-1}$	$3.6 \times 10^{-1}$	$7.8 \times 10^{-1}$
	G2	$3.8 \times 10^{-1}$	$4.3 \times 10^{-1}$	$6.8 \times 10^{-1}$
	G3	$4.9 \times 10^{-1}$	$6.0 \times 10^{-1}$	$8.7 \times 10^{-1}$
DeepONet	G1	$3.9 \times 10^{-1}$	$5.1 \times 10^{-1}$	$1.4 \times 10^0$
	G2	$4.9 \times 10^{-1}$	$5.7 \times 10^{-1}$	$8.4 \times 10^{-1}$
	G3	$6.2 \times 10^{-1}$	$7.1 \times 10^{-1}$	$1.4 \times 10^0$
poseidon	G1	$1.9 \times 10^{-1}$	$2.7 \times 10^{-1}$	$2.5 \times 10^{-1}$
	G2	$2.4 \times 10^{-1}$	$2.7 \times 10^{-1}$	$3.7 \times 10^{-1}$
	G3	$3.9 \times 10^{-1}$	$4.9 \times 10^{-1}$	$3.7 \times 10^{-1}$

**Table 7:** The relative  $L_\infty$  error of various neural operators trained on the 2D LDC dataset. All metrics are reported on the testing dataset.

Model	Geometry	M1	M2	M3
FNO	G1	$6.1 \times 10^{-1}$	$7.2 \times 10^{-1}$	$9.8 \times 10^{-1}$
	G2	$7.4 \times 10^{-1}$	$8.1 \times 10^{-1}$	$8.5 \times 10^{-1}$
	G3	$9.8 \times 10^{-1}$	$1.1 \times 10^0$	$1.2 \times 10^0$
DeepONet	G1	$8.3 \times 10^{-1}$	$9.1 \times 10^{-1}$	$3.3 \times 10^0$
	G2	$9.0 \times 10^{-1}$	$1.0 \times 10^0$	$2.6 \times 10^0$
	G3	$1.1 \times 10^0$	$1.3 \times 10^0$	$3.1 \times 10^0$
poseidon	G1	$3.3 \times 10^{-1}$	$4.3 \times 10^{-1}$	$2.9 \times 10^{-1}$
	G2	$4.1 \times 10^{-1}$	$5.7 \times 10^{-1}$	$4.5 \times 10^{-1}$
	G3	$6.4 \times 10^{-1}$	$7.7 \times 10^{-1}$	$5.1 \times 10^{-1}$

Table 5 through Table 7 present a comprehensive evaluation of the performance of all the trained models. The models were assessed using the four metrics outlined in Section 4.6, across each geometry class and at the highest resolution ( $512 \times 512$ ) for all seven models. Our findings indicate that in terms of both *mean squared error* and  $L_\infty$  error, the foundation models consistently outperformed the neural operators across all geometries and metrics. For the  $L_2$  error, the neural operators performed comparably to the foundation models. We also include examples of the field prediction for  $u$ ,  $v$  and  $p$  using the three scientific machine learning models in Appendix B.4. We encourage further exploration and validation of these findings by the research community.

**Table 8:** The residual errors ( $M4$ ) of various neural operators trained on the 2D LDC dataset. All metrics are reported on the testing dataset. Errors reported on the test dataset (per sample) normalized by the total number of mesh points.

Model	Geometry	momentum	continuity
FNO	G1	$1.03 \times 10^{-4}$	$9.11 \times 10^{-8}$
	G2	$5.80 \times 10^{-5}$	$7.05 \times 10^{-8}$
	G3	$7.62 \times 10^{-5}$	<b><math>6.76 \times 10^{-8}</math></b>
DeepONet	G1	$5.09 \times 10^{-5}$	$8.96 \times 10^{-8}$
	G2	<b><math>4.51 \times 10^{-5}</math></b>	$1.39 \times 10^{-7}$
	G3	$5.72 \times 10^{-5}$	$5.07 \times 10^{-7}$
poseidon	G1	$9.62 \times 10^{-5}$	$1.35 \times 10^{-7}$
	G2	$8.35 \times 10^{-5}$	$1.39 \times 10^{-7}$
	G3	$9.64 \times 10^{-5}$	$1.46 \times 10^{-7}$

## 6 Conclusions

We introduce a comprehensive benchmark dataset designed for evaluating neural solvers of flow simulations over complex geometries. **FlowBench** encompasses 2D and 3D simulations, covering many scenarios, from steady-state problems to time-dependent problems. **FlowBench** aims to help develop scientific machine learning models by offering a complex dataset designed to challenge and benchmark their performance. Neural PDE solvers that can account for the effect of complex geometries can have a major impact on various applications ranging from bioengineering and power production to automotive and aerospace engineering defined by the interaction of complex geometrical objects with a fluid medium.

As a step toward reproducibility and ease of use, we have released an end-to-end tutorial in our follow-up work (Rabeh et al., 2024a) and an accompanying GitHub repository available (here). This repository provides step-by-step instructions and code for training and evaluating Fourier Neural Operator, DeepONet, and the pretrained poseidon model on the 2D LDC subset of FlowBench.

**Limitations:** (1) Our evaluation of existing neural PDE solvers is limited to one of the four **FlowBench** datasets and on seven neural PDE models. We encourage the community to contribute by evaluating a wider range of approaches using this dataset and the proposed metrics. Additionally, we plan to expand **FlowBench** with more data—particularly for 3D simulations—and invite the CFD community to do the same, extending the dataset to cover higher  $Re$  and  $Gr$  operating conditions.

## Acknowledgements

We gratefully acknowledge support from the NAIRR pilot program for computational access to TACC Frontera. This work is supported by the AI Research Institutes program supported by NSF and USDA-NIFA under AI Institute: for Resilient Agriculture, Award No. 2021-67021-35329. We also acknowledge partial support through NSF 2053760.

## References

- Chamil Abeykoon. Compact heat exchangers—design and optimization with CFD. *International Journal of Heat and Mass Transfer*, 146:118766, 2020.
- Jonas Allegrini and Jan Carmeliet. Simulations of local heat islands in zürich with coupled CFD and building energy models. *Urban climate*, 24:340–359, 2018.
- Neil Ashton, Jordan B Angel, Aditya S Ghate, Gaetan KW Kenway, Man Long Wong, Cetin Kiris, Astrid Walle, Danielle C Maddix, and Gary Page. Windsorml—high-fidelity computational fluid dynamics dataset for automotive aerodynamics. *arXiv preprint arXiv:2407.19320*, 2024a.
- Neil Ashton, Danielle C Maddix, Samuel Gundry, and Parisa M Shabestari. Ahmedml: High-fidelity computational fluid dynamics dataset for incompressible, low-speed bluff body aerodynamics. *arXiv preprint arXiv:2407.20801*, 2024b.
- Neil Ashton, Charles Mockett, Marian Fuchs, Louis Fliessbach, Hendrik Hetmann, Thilo Knacke, Norbert Schonwald, Vangelis Skaperdas, Grigoris Fotiadis, Astrid Walle, et al. Drivaerml: High-fidelity computational fluid dynamics dataset for road-car external aerodynamics. *arXiv preprint arXiv:2408.11969*, 2024c.
- Nabil Atallah, Claudio Canuto, and Guglielmo Scovazzi. The shifted boundary method for solid mechanics. *International Journal for Numerical Methods in Engineering*, 122(20):5935–5970, 2021.
- Nabil M Atallah, Claudio Canuto, and Guglielmo Scovazzi. The second-generation shifted boundary method and its numerical analysis. *Computer Methods in Applied Mechanics and Engineering*, 372:113341, 2020.
- Rowel Atienza. Pyramid u-network for skeleton extraction from shape points. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2019.
- Utkarsh Ayachit. *The paraview guide: a parallel visualization application*. Kitware, Inc., 2015.
- Pierre-Alexandre Beaufort, Christophe Geuzaine, and Jean-François Remacle. Automatic surface mesh generation for discrete models—a complete and automatic pipeline based on reparametrization. *Journal of Computational Physics*, 417:109575, 2020.
- Suvanjan Bhattacharyya, Kunal Dey, Akshoy Ranjan Paul, and Ranjib Biswas. A novel CFD analysis to minimize the spread of COVID-19 virus in hospital isolation room. *Chaos, Solitons & Fractals*, 139:110294, 2020.
- Muhammad Mahmood Aslam Bhutta, Nasir Hayat, Muhammad Hassan Bashir, Ahmer Rais Khan, Kanwar Naveed Ahmad, and Sarfaraz Khan. CFD applications in various heat exchangers design: A review. *Applied Thermal Engineering*, 32:1–12, 2012.
- OJ Boelens, KJ Badcock, Stefan Görtz, S Morton, W Fritz, SL Karman Jr, T Michal, and John E Lamar. F16-xl geometry and computational grids used in cranked-arrow wing aerodynamics project international. *Journal of Aircraft*, 46(2):369–376, 2009.
- Florent Bonnet, Jocelyn Mazari, Paola Cinnella, and Patrick Gallinari. Airfrans: High fidelity computational fluid dynamics dataset for approximating reynolds-averaged navier–stokes solutions. *Advances in Neural Information Processing Systems*, 35:23463–23478, 2022a.
- Florent Bonnet, Jocelyn Ahmed Mazari, Thibaut Munzer, Pierre Yser, and Patrick Gallinari. An extensible benchmarking graph-mesh dataset for studying steady-state incompressible navier-stokes equations, 2022b.
- Zhen Chen, Chang Shu, L. Yang, Xiang Zhao, and Ningyu Liu. Immersed boundary—simplified thermal lattice Boltzmann method for incompressible thermal flows. *Physics of Fluids*, 32, 01 2020. doi: 10.1063/1.5138711.

- Yoann Cheny and Olivier Botella. The ls-stag method: A new immersed boundary/level-set method for the computation of incompressible viscous flows in complex moving geometries with good conservation properties. *Journal of Computational Physics*, 229(4):1043–1076, 2010.
- Ilke Demir, Camilla Hahn, Kathryn Leonard, Geraldine Morin, Dana Rahbani, Athina Panotopoulou, Amelie Fondevilla, Elena Balashova, Bastien Durix, and Adam Kortylewski. Skelneton 2019: Dataset and challenge on deep learning for geometric shape understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.
- Marie-Lena Eckert, Kiwon Um, and Nils Thuerey. ScalarFlow: a large-scale volumetric data set of real-world scalar transport flows for computer animation and machine learning. *ACM Transactions on Graphics (TOG)*, 38(6):1–16, 2019.
- T Engels, D Kolomenskiy, Kai Schneider, F-O Lehmann, and J Sesterhenn. Bumblebee flight in heavy turbulence. *Physical review letters*, 116(2):028103, 2016.
- Dhruv Gamdha, Kumar Saurabh, Baskar Ganapathysubramanian, and Adarsh Krishnamurthy. Geometric modeling and physics simulation framework for building a digital twin of extrusion-based additive manufacturing. *arXiv preprint arXiv:2305.07120*, 2023.
- Paul-Louis George, Frédéric Hecht, and Éric Saltel. Automatic mesh generator with specified boundary. *Computer methods in applied mechanics and engineering*, 92(3):269–288, 1991.
- David Greenblatt and Israel J Wygnanski. The control of flow separation by periodic excitation. *Progress in aerospace Sciences*, 36(7):487–545, 2000.
- Sheikh Md Shakeel Hassan, Arthur Feeney, Akash Dhruv, Jihoon Kim, Youngjoon Suh, Jaiyoung Ryu, Yoonjin Won, and Aparna Chandramowliswaran. BubbleML: A multiphase multiphysics dataset and benchmarks for machine learning. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=0Wmglu8zak>.
- Eric Heisler, Cheng-Hau Yang, Aadesh Deshmukh, Baskar Ganapathysubramanian, and Hari Sundar. Generating finite element codes combining adaptive octrees with complex geometries. *arXiv preprint arXiv:2305.19398*, 2023.
- Tore A Helgedagsrud, Yuri Bazilevs, Kjell M Mathisen, and Ole A Øiseth. ALE-VMS methods for wind-resistant design of long-span bridges. *Journal of Wind Engineering and Industrial Aerodynamics*, 191: 143–153, 2019.
- Ronald D Henderson. Details of the drag curve near the onset of vortex shedding. *Physics of Fluids*, 7(9): 2102–2104, 1995.
- Maximilian Herde, Bogdan Raonić, Tobias Rohner, Roger Käppeli, Roberto Molinaro, Emmanuel de Bézenac, and Siddhartha Mishra. Poseidon: Efficient foundation models for pdes, 2024.
- Tingting Huang and Hee-Chang Lim. Simulation of lid-driven cavity flow with internal circular obstacles. *Applied Sciences*, 10(13):4583, 2020.
- Jakob Jakob, Markus Gross, and Tobias Günther. A fluid flow data set for machine learning and its application to neural flow map interpolation. *IEEE Transactions on Visualization and Computer Graphics*, 27(2): 1279–1289, 2020.
- Steeven Janny, Aurélien Bénateau, Madiha Nadri, Julie Digne, Nicolas Thome, and Christian Wolf. Eagle: Large-scale learning of turbulent fluid dynamics with mesh transformers, 2023.

- W. Jester and Y. Kallinderis. Numerical study of incompressible flow about fixed cylinder pairs. *Journal of Fluids and Structures*, 17(4):561–577, 2003. ISSN 0889-9746. doi: [https://doi.org/10.1016/S0889-9746\(02\)00149-4](https://doi.org/10.1016/S0889-9746(02)00149-4). URL <https://www.sciencedirect.com/science/article/pii/S0889974602001494>.
- D. Kamensky, M.-C. Hsu, D. Schillinger, J. A. Evans, A. Aggarwal, Y. Bazilevs, M. S. Sacks, and T. J. R. Hughes. An immersogeometric variational framework for fluid–structure interaction: Application to bioprosthetic heart valves. *Computer Methods in Applied Mechanics and Engineering*, 284:1005–1053, 2015.
- Soonpil Kang and Arif Masud. A variational multiscale method with immersed boundary conditions for incompressible flows. *Meccanica*, 56(6):1397–1422, 2021.
- Makrand A Khanwale, Kumar Saurabh, Masado Ishii, Hari Sundar, James A Rossmanith, and Baskar Ganapathysubramanian. A projection-based, semi-implicit time-stepping approach for the cahn-hilliard navier-stokes equations on adaptive octree meshes. *Journal of Computational Physics*, 475:111874, 2023.
- Biswajit Khara, Ethan Herron, Aditya Balu, Dhruv Gamdha, Chih-Hsuan Yang, Kumar Saurabh, Anushrut Jignasu, Zhanhong Jiang, Soumik Sarkar, Chinmay Hegde, et al. Neural pde solvers for irregular domains. *Computer-Aided Design*, 172:103709, 2024.
- Sungu Kim, Kumar Saurabh, Makrand A Khanwale, Ali Mani, Robbyn K Anand, and Baskar Ganapathysubramanian. Direct numerical simulation of electrokinetic transport phenomena in fluids: Variational multi-scale stabilization and octree-based mesh refinement. *Journal of Computational Physics*, 500:112747, 2024.
- Dilek Funda Kurtulus. On the unsteady behavior of the flow around NACA 0012 airfoil with steady external conditions at  $Re=1000$ . *International journal of micro air vehicles*, 7(3):301–326, 2015.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations, 2021.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3:218–229, 2021.
- Haoxiang Luo, Hu Dai, and Paulo Ferreira de Sousa. A hybrid formulation to suppress the numerical oscillations caused by immersed moving boundaries. In *APS Division of Fluid Dynamics Meeting Abstracts*, volume 62, pages EL–006, 2009.
- Yining Luo, Yingfa Chen, and Zhen Zhang. CFDBench: A large-scale benchmark for machine learning methods in fluid dynamics, 2024. arXiv preprint.
- Alex Main and Guglielmo Scovazzi. The shifted boundary method for embedded domain computations. Part II: Linear advection-diffusion and incompressible Navier-Stokes equations. *J. Comput. Phys.*, 372:996–1026, 2018a.
- Alex Main and Guglielmo Scovazzi. The shifted boundary method for embedded domain computations. Part I: Poisson and Stokes problems. *Journal of Computational Physics*, 372:972–995, 2018b.
- Ryley McConkey, Eugene Yee, and Fue-Sang Lien. A curated dataset for data-driven turbulence modelling. *Scientific Data*, 8(1), September 2021. ISSN 2052-4463. doi: 10.1038/s41597-021-01034-2. URL <http://dx.doi.org/10.1038/s41597-021-01034-2>.
- Rajat Mittal and Jung Hee Seo. Origin and evolution of immersed boundary methods in computational fluid dynamics. *Physical review fluids*, 8(10):100501, 2023.

- Rajat Mittal, Haibo Dong, Meliha Bozkurttas, FM Najjar, Abel Vargas, and Alfred Von Loebbecke. A versatile sharp interface immersed boundary method for incompressible flows with complex boundaries. *Journal of computational physics*, 227(10):4825–4852, 2008.
- Steven J Owen. A survey of unstructured mesh generation technology. *IMR*, 239(267):15, 1998.
- Per-Olof Persson. *Mesh generation for implicit geometries*. PhD thesis, Massachusetts Institute of Technology, 2005.
- Charles S Peskin. Flow patterns around heart valves: a numerical method. *Journal of Computational Physics*, 10(2):252–271, 1972.
- Les Piegl and Wayne Tiller. *The NURBS book*. Springer Science & Business Media, 2012.
- Rajagopalan Priyadarsini, Wong Nyuk Hien, and Cheong Kok Wai David. Microclimatic modeling of the urban thermal environment of singapore to mitigate urban heat island. *Solar energy*, 82(8):727–745, 2008.
- Ali Rabeh, Ethan Herron, Aditya Balu, Soumik Sarkar, Chinmay Hegde, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Geometry matters: Benchmarking scientific ml approaches for flow prediction around complex geometries. *arXiv preprint arXiv:2501.01453*, 2024a.
- Ali Rabeh, Makrand A Khanwale, Jonghyun Lee, and Baskar Ganapathysubramanian. Modeling and simulations of high-density two-phase flows using projection-based cahn-hilliard navier-stokes equations. *arXiv preprint arXiv:2406.17933*, 2024b.
- Rubina Ramponi and Bert Blocken. CFD simulation of cross-ventilation for a generic isolated building: Impact of computational parameters. *Building and environment*, 53:34–48, 2012.
- Stephan Rasp, Peter D. Dueben, Sebastian Scher, Jonathan A. Weyn, Soukayna Mouatadid, and Nils Thuerey. Weatherbench: A benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems*, 12(11), November 2020. ISSN 1942-2466. doi: 10.1029/2020ms002203. URL <http://dx.doi.org/10.1029/2020MS002203>.
- Kumar Saurabh, Boshun Gao, Milinda Fernando, Songzhe Xu, Makrand A Khanwale, Biswajit Khara, Ming-Chen Hsu, Adarsh Krishnamurthy, Hari Sundar, and Baskar Ganapathysubramanian. Industrial scale Large Eddy simulations with adaptive octree meshes using immersogeometric analysis. *Computers & Mathematics with Applications*, 97:28–44, 2021a.
- Kumar Saurabh, Masado Ishii, Milinda Fernando, Boshun Gao, Kendrick Tan, Ming-Chen Hsu, Adarsh Krishnamurthy, Hari Sundar, and Baskar Ganapathysubramanian. Scalable adaptive PDE solvers in arbitrary domains. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15, 2021b.
- Kumar Saurabh, Masado Ishii, Makrand A Khanwale, Hari Sundar, and Baskar Ganapathysubramanian. Scalable adaptive algorithms for next-generation multiphase flow simulations. In *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 590–601. IEEE, 2023.
- Jung-Hee Seo and Rajat Mittal. Improved swimming performance in schooling fish via leading-edge vortex enhancement. *Bioinspiration & Biomimetics*, 17(6):066020, 2022.
- Ze Shao, Hongjun Zhu, Wenhua Zhao, and Tongming Zhou. Numerical simulation of steady flow past a D-shaped cylinder at low Reynolds numbers. In *Australasian Fluid Mechanics Conference (AFMC)*. The University of Queensland, 2020.
- Igor Simonovski and Leon Cizelj. Automatic parallel generation of finite element meshes for complex spatial structures. *Computational Materials Science*, 50(5):1606–1618, 2011.

- Jialei Song, Haoxiang Luo, and Tyson L Hedrick. Three-dimensional flow and lift characteristics of a hovering ruby-throated hummingbird. *Journal of The Royal Society Interface*, 11(98):20140541, 2014.
- A Spohn and P Gilliéron. Flow separations generated by a simplified geometry of an automotive vehicle. In *IUTAM Symposium: unsteady separated flows*, volume 1. Kluwer Academic, 2002.
- Dan Stanzione, John West, R Todd Evans, Tommy Minyard, Omar Ghattas, and Dhabaleswar K Panda. Frontera: The evolution of leadership computing at the National Science Foundation. In *Practice and Experience in Advanced Research Computing*, pages 106–111, 2020.
- Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, Dan MacKinlay, Francesco Alesiani, Dirk Pflüger, and Mathias Niepert. PDEBENCH: An extensive benchmark for scientific machine learning, 2023.
- Kendrick Tan, Boshun Gao, Cheng-Hau Yang, Emily L Johnson, Ming-Chen Hsu, Alberto Passalacqua, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. A computational framework for transmission risk assessment of aerosolized particles in classrooms. *Engineering with Computers*, pages 1–22, 2023.
- LL Ting. Effects of dimple size and depth on golf ball aerodynamic performance. In *Fluids engineering division summer meeting*, volume 36975, pages 811–817, 2003.
- Roberto Verzicco. Immersed boundary methods: Historical perspective and future outlook. *Annual Review of Fluid Mechanics*, 55(1):129–155, 2023.
- Deheng Wei, Jianfeng Wang, and Budi Zhao. A simple method for particle shape generation with spherical harmonics. *Powder Technology*, 330:284–291, 2018. ISSN 0032-5910. doi: <https://doi.org/10.1016/j.powtec.2018.02.006>. URL <https://www.sciencedirect.com/science/article/pii/S0032591018301189>.
- Fei Xu, Emily L Johnson, Chenglong Wang, Arian Jafari, Cheng-Hau Yang, Michael S Sacks, Adarsh Krishnamurthy, and Ming-Chen Hsu. Computational investigation of left ventricular hemodynamics following bioprosthetic aortic and mitral valve replacement. *Mechanics Research Communications*, 112: 103604, 2021.
- Songzhe Xu, Boshun Gao, Ming-Chen Hsu, and Baskar Ganapathysubramanian. A residual-based variational multiscale method with weak imposition of boundary conditions for buoyancy-driven flows. *Computer Methods in Applied Mechanics and Engineering*, 352:345–368, 2019.
- Wenzhuo Xu, Noelia Grande Gutierrez, and Christopher McComb. MegaFlow2D: A parametric dataset for machine learning super-resolution in computational fluid dynamics simulations. In *Proceedings of Cyber-Physical Systems and Internet of Things Week 2023*. Association for Computing Machinery, 2023. URL <https://doi.org/10.1145/3576914.3587552>.
- Cheng-Hau Yang, Kumar Saurabh, Guglielmo Scovazzi, Claudio Canuto, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Optimal surrogate boundary selection and scalability studies for the shifted boundary method on octree meshes. *Computer Methods in Applied Mechanics and Engineering*, 419: 116686, 2024a.
- Cheng-Hau Yang, Guglielmo Scovazzi, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Simulating incompressible flows over complex geometries using the shifted boundary method with incomplete adaptive octree meshes. *arXiv preprint arXiv:2411.00272*, 2024b.
- Cheng-Hau Yang, Guglielmo Scovazzi, Adarsh Krishnamurthy, and Baskar Ganapathysubramanian. Simulating incompressible flows over complex geometries using the shifted boundary method with incomplete adaptive octree meshes, 2024c. URL <https://arxiv.org/abs/2411.00272>.

D You and Parviz Moin. Active control of flow separation over an airfoil using synthetic jets. *Journal of Fluids and structures*, 24(8):1349–1357, 2008.

Yi Zhu, Fang-Bao Tian, John Young, James C Liao, and Joseph CS Lai. A numerical study of fish adaption behaviors in complex environments with a deep reinforcement learning and immersed boundary–lattice boltzmann method. *Scientific Reports*, 11(1):1691, 2021.

## Appendix A. Details of the CFD simulation framework

Our CFD framework is a well-validated and massively scalable software suite that uses a variant of the immersed boundary method integrated with octree meshes to perform highly efficient and accurate Large-Eddy Simulations (LES) of flows around complex geometries. This framework demonstrates scalability of up to 32,000 processors, achieved through several key innovations, including (a) Rapid in-out tests: These tests quickly determine whether a point is inside or outside a given geometry, significantly speeding up the simulation process. (b) adaptive quadrature: This technique ensures accurate force evaluation by dynamically adjusting the numerical integration based on the local complexity of the geometry, and (c) Tensorized operations: These operations optimize performance by leveraging tensor algebra for computational efficiency. Additional details of the CFD framework, including implementation, are provided in [Saurabh et al. \(2021a, 2023\)](#).

Our datasets and code are licensed under CC-BY-NC-4.0. Next, we briefly detail the mathematics of the shifted boundary method.

### A.1 Heuristic derivation of the SBM formulation for Navier-Stokes

#### A.1.1 DERIVATION FOR NITSCHÉ'S METHOD FOR NAVIER-STOKES

Nitsché's method starts with the Lagrange multiplier method and adds a penalty term to obtain the augmented Lagrangian formulation. We begin with the Lagrange multiplier method, which is given by:

$$\underbrace{\int w \lambda \, d\Gamma}_{\text{consistency term}} + \underbrace{\int \delta \lambda (u - u_D) \, d\Gamma}_{\text{adjoint-consistency term}}. \quad (7)$$

After adding the penalty term  $\frac{\beta}{h \cdot Re} \int w (u - u_D) \, d\Gamma$ , we can obtain Nitsché's method:

$$\underbrace{\int w \lambda \, d\Gamma}_{\text{consistency term}} + \underbrace{\int \delta \lambda (u - u_D) \, d\Gamma}_{\text{adjoint-consistency term}} + \underbrace{\frac{\beta}{h \cdot Re} \int w (u - u_D) \, d\Gamma}_{\text{penalty term}}. \quad (8)$$

Our objective is to determine the Lagrange multiplier,  $\lambda$ , for the Navier-Stokes equation. This is achieved by ensuring that the consistency term matches the boundary term, which arises when integration by parts is applied to the weak form of the Navier-Stokes equation. To derive the weak form, we start from the strong form of the Navier-Stokes equation and utilize test functions  $w$  and  $q$ :

$$\begin{cases} \frac{\partial u}{\partial t} + u \cdot \nabla u - f - \nabla \cdot \sigma = 0 \\ \nabla \cdot u = 0 \end{cases} \Rightarrow \int w \left( \frac{\partial u}{\partial t} + u \cdot \nabla u - f \right) d\Omega - \underbrace{\int w \nabla \cdot \sigma \, d\Omega}_{\text{Cauchy stress term}} + \int q \nabla \cdot u \, d\Omega = 0. \quad (9)$$

Performing integration by parts ([Section A.1.2](#)) on the Cauchy stress term, we obtain the following:

$$- \int w \nabla \cdot \sigma \, d\Omega = \underbrace{\int \nabla w : \sigma \, d\Omega}_{\text{volume term}} - \underbrace{\int w \sigma \cdot \mathbf{n} \, d\Gamma}_{\text{boundary term}}. \quad (10)$$

The boundary term is the same as the consistency term in Nitsche's formulation so that we can obtain the following:

$$\int w \lambda \, d\Gamma = - \int w \sigma \cdot \mathbf{n} \, d\Gamma \Rightarrow \lambda = -\sigma \cdot \mathbf{n} = -\left(\frac{2}{Re} \nabla^s u - pI\right) \cdot \mathbf{n}. \quad (11)$$

And, we can obtain  $\delta\lambda = -\left(\frac{2}{Re} \nabla^s w + qI\right) \cdot \mathbf{n}$ .

Upon substitution of  $\lambda$  and  $\delta\lambda$  into [Equation 8](#), Nitsche's formulation for Navier-Stokes can be obtained, and is expressed as follows:

$$- \underbrace{\int w \left(\frac{2}{Re} \nabla^s u - pI\right) \, d\Gamma}_{\text{consistency term}} - \underbrace{\int \left(\frac{2}{Re} \nabla^s w + qI\right) (u - u_D) \, d\Gamma}_{\text{adjoint-consistency term}} + \underbrace{\frac{\beta}{h \cdot Re} \int w (u - u_D) \, d\Gamma}_{\text{penalty term}}. \quad (12)$$

### A.1.2 INTEGRATION BY PART FOR CAUCHY STRESS TERM

Doing integration by part for the Cauchy stress term in [Equation 9](#), we write the term in the form of Einstein notation so that it is easier to calculate:

$$\begin{aligned} & - \int w \nabla \cdot \sigma \, d\Omega = - \int w_i \frac{\partial \sigma_{ik}}{\partial x_k} \, d\Omega = - \int w_i \left( \frac{1}{Re} \frac{\partial \varepsilon_{ik}}{\partial x_k} - \delta_{ik} \frac{\partial p}{\partial x_k} \right) \, d\Omega = - \int w_i \left( \frac{1}{Re} \frac{\partial \varepsilon_{ik}}{\partial x_k} - \frac{\partial p}{\partial x_i} \right) \, d\Omega \\ & = \int \left( \underbrace{-\frac{1}{Re} \frac{\partial (w_i \varepsilon_{ik})}{\partial x_k}}_{\text{term 1}} + \underbrace{\frac{1}{Re} \varepsilon_{ik} \frac{\partial w_i}{\partial x_k}}_{\text{term 2}} + \underbrace{\frac{\partial (p w_i)}{\partial x_i}}_{\text{term 1}} - \underbrace{p \frac{\partial w_i}{\partial x_i}}_{\text{term 2}} \right) \, d\Omega \\ & = \int \left( \underbrace{-\frac{1}{Re} \frac{\partial (w_i \varepsilon_{ik})}{\partial x_k} + \frac{\partial (p w_i)}{\partial x_i}}_{\text{term 1}} \right) \, d\Omega + \int \left( \underbrace{\frac{1}{Re} \varepsilon_{ik} \frac{\partial w_i}{\partial x_k} - p \frac{\partial w_i}{\partial x_i}}_{\text{term 2}} \right) \, d\Omega \\ & = \int w_i \left( -\frac{1}{Re} \varepsilon_{ik} n_k + p n_i \right) \, d\Gamma + \int \frac{\partial w_i}{\partial x_k} \left( \frac{1}{Re} \varepsilon_{ik} - p \delta_{ik} \right) \, d\Omega = \int w_i (-\sigma_{ik} n_k) \, d\Gamma + \int \frac{\partial w_i}{\partial x_k} \sigma_{ik} \, d\Omega \\ & = - \int w \sigma \cdot \mathbf{n} \, d\Gamma + \int \nabla w : \sigma \, d\Omega. \end{aligned} \quad (14)$$

### A.1.3 SBM FOR NAVIER-STOKES

In the SBM, instead of enforcing boundary condition on the true boundary ( $\Gamma$ ), we enforce Nitsche's method on the surrogate boundary ( $\tilde{\Gamma}$ ) by changing the boundary condition from  $u_D$  to  $\tilde{u}_D$ :

$$- \underbrace{\int w \left(\frac{2}{Re} \nabla^s u - pI\right) \, d\tilde{\Gamma}}_{\text{consistency term}} - \underbrace{\int \left(\frac{2}{Re} \nabla^s w + qI\right) (u - \tilde{u}_D) \, d\tilde{\Gamma}}_{\text{adjoint-consistency term}} + \underbrace{\frac{\beta}{h \cdot Re} \int \tilde{w} (u - \tilde{u}_D) \, d\tilde{\Gamma}}_{\text{penalty term}}. \quad (15)$$

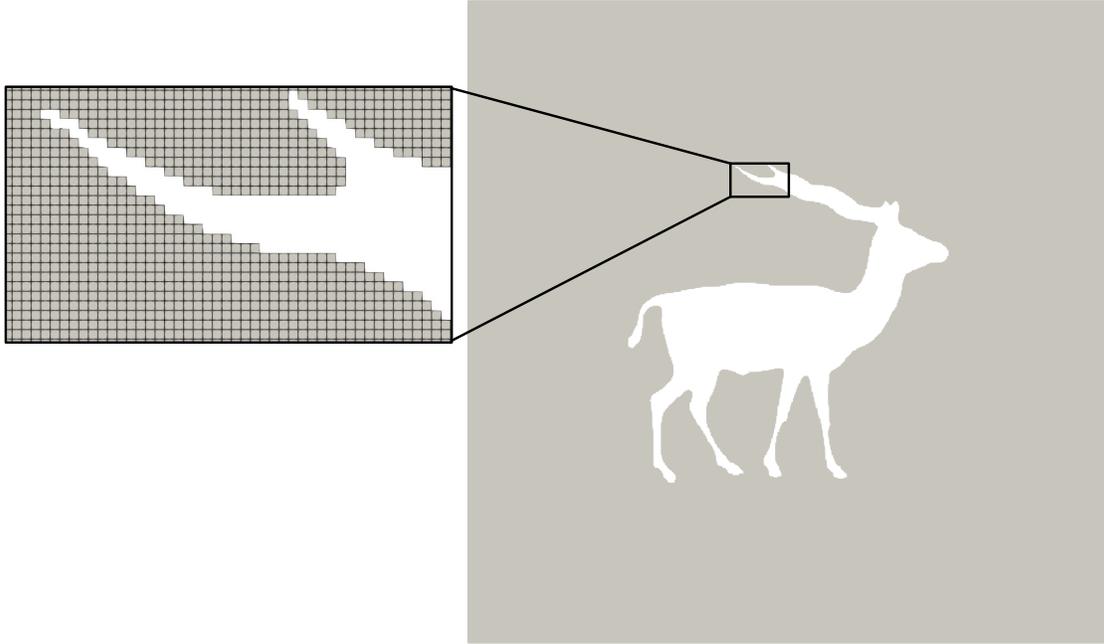
With Taylor expansion to account for the discrepancy between the surrogate and true boundary, the boundary condition becomes  $\tilde{u}_D = u_D - \nabla u \cdot d$ . Substituting this expression into [Equation 15](#), we obtain:

$$\begin{aligned} & - \underbrace{\int w \left(\frac{2}{Re} \nabla^s u - pI\right) \, d\tilde{\Gamma}}_{\text{consistency term}} - \underbrace{\int \left(\frac{2}{Re} \nabla^s w + qI\right) (u + \nabla u \cdot d - u_D) \, d\tilde{\Gamma}}_{\text{adjoint-consistency term}} \\ & + \underbrace{\frac{\beta}{h \cdot Re} \int (w + \nabla w \cdot d) (u + \nabla u \cdot d - u_D) \, d\tilde{\Gamma}}_{\text{penalty term}}. \end{aligned} \quad (16)$$

## A.2 Solving the CFD equations: Automated creation of meshes involving complex geometries

Tree-based mesh generation, using quadtrees in 2D and octrees in 3D, is common in computational sciences due to its simplicity and parallel scalability. These tree-based data structures enable efficient refinement of regions of interest, facilitating their deployment in large-scale multi-physics simulations. Our mesh generation tool, Dendro-kt (Saurabh et al., 2021b; Heisler et al., 2023; Gamdha et al., 2023), provides balanced, partitioned, and parallel tree structures, making it highly effective for large-scale numerical PDE discretizations.

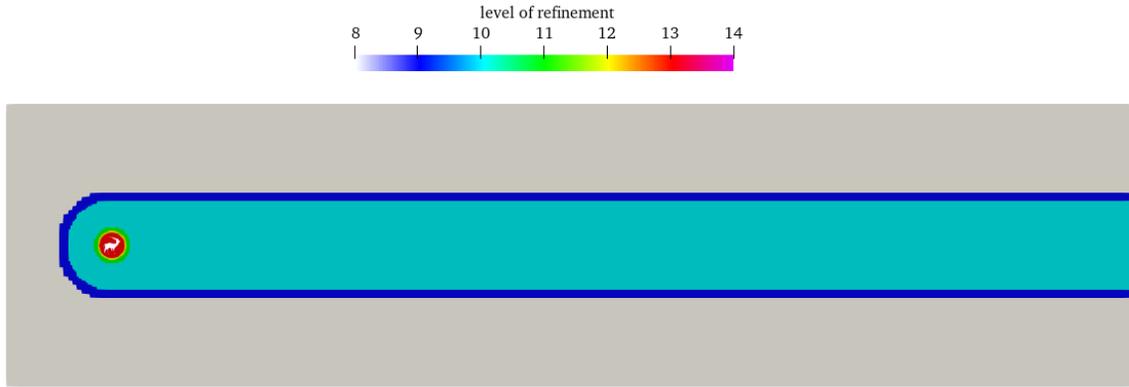
**Lid-driven cavity (LDC) flow:** For the LDC case, we use an octree-based mesh generation framework to create a uniform mesh with an element size of  $\frac{1}{2^9}$  over a  $[0, 2] \times [0, 2]$  domain. This produces a  $512 \times 512$  mesh resulting in total degrees-of-freedom of  $512 \times 512 \times 3 \sim 750K$ .



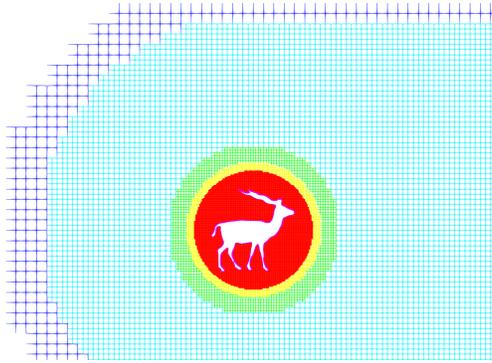
**Figure 13:** Uniform octree mesh for the lid-driven cavity flow. We illustrate the mesh with a representative shape from the Skelneton data. This mesh is carved out from a  $512 \times 512$  uniform tessellation of the domain.

**Flow past object (FPO):** Our computational domain is a rectangular region spanning  $[0, 64] \times [0, 16]$ , with the complex geometry centered at the coordinates  $(6, 8)$ . Our mesh refinement strategy involves several layers of progressively coarser mesh surrounding the geometry and along the wake. Specifically, we utilize five concentric circles centered at  $(6, 8)$ , with the innermost circle having a radius of 0.71 and a refinement level of 13 (yielding an element size of  $\frac{64}{2^{13}}$ ), the next circle with a radius of 0.8 and a refinement level of 12 ( $\frac{64}{2^{12}}$  element size), the third circle with a radius of 1 and a refinement level of 11 ( $\frac{64}{2^{11}}$  element size), the fourth circle with a radius of 2.5 and a refinement level of 10 ( $\frac{64}{2^{10}}$  element size), and the outermost circle with a radius of 3 and a refinement level of 9 ( $\frac{64}{2^9}$  element size). We use a non-dimensional time step of 0.01 for the simulation. Starting from a non-dimensional total time of 392, we begin outputting results every 0.05 non-dimensional time units until reaching a non-dimensional total time of 404. The period from non-dimensional time 392 to 404 is when we output results that are post-processed for use in FlowBench. This time step was selected to ensure that each vortex shedding cycle is captured with at least 100 snapshots, based on an evaluation of the shedding frequencies across the range of Reynolds numbers and shapes. Consequently, each sample contains at least two complete shedding cycles.

Additionally, we define two rectangular refinement regions aligned with the flow direction. The first rectangle has its bottom-left corner at  $(6, 5.5)$  and top-right corner at  $(64, 10.5)$ , with a refinement level of 10 ( $\frac{64}{2^{10}}$



(a) Overview of the mesh refinement levels.



(b) Zoomed-in view of the local mesh refinement. The five levels of progressively coarser meshes surrounding the geometry are clearly visible.



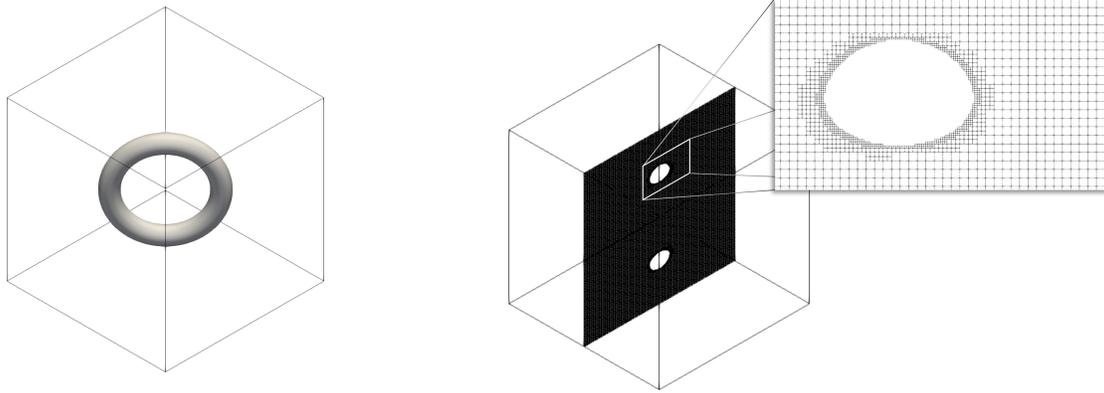
(c) Zoomed-in view of the local mesh refinement near the geometry.

**Figure 14:** Local refinement octree mesh for the flow past the geometry. The geometry is represented by a deer.

element size). The second rectangle has its bottom-left corner at (6, 5) and top-right corner at (64, 11), with a refinement level of 9 ( $\frac{64}{2^9}$  element size). Close to the geometry, to ensure detailed capture of the flow dynamics, we achieve a finer mesh with a refinement level of 14, resulting in an element size of  $\frac{64}{2^{14}}$ . The mesh with different refinement levels are shown in Figure 14.

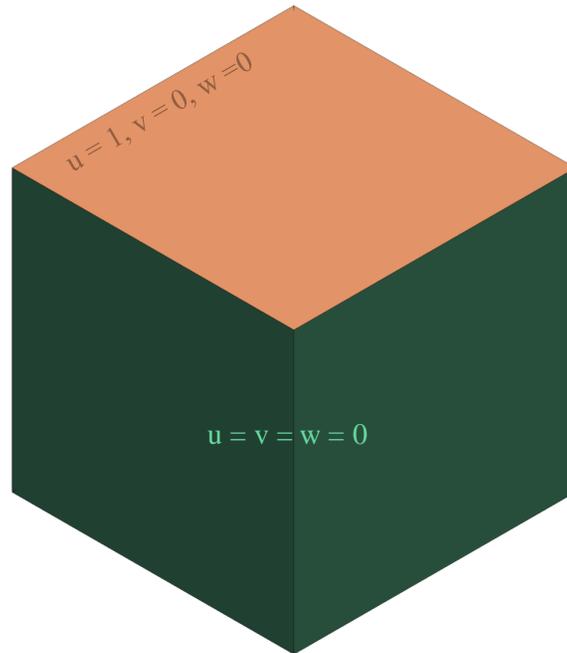
The refinement strategy leads to a total nodal point number of 117978. Consequently, the total degrees of freedom (DOFs) amount is  $117978 \times 3$  (353934). The combination of adaptive refinement and SBM formulation is critical to our ability to solve the PDEs in complex geometries. For comparison, if we were to use a uniform mesh (as in the case of the LDC) instead of an adaptively refined mesh, we would have had a problem with  $2^{14} \times 2^{12} \times 3 \sim 201M$  DOFs!

**3D LDC:** We use a similar adaptive refinement strategy for the 3D LDC case. See Figure 15 for an example of an object and the mesh used. Close to the object's surface, we use a refinement of level 9, producing elements of size  $2/2^9$ , and away from the object, we progressively coarsen the mesh to a refinement level of 7, for elements of size  $2/2^7$ . This produces a problem with 2.5M DOFs. For the boundary condition, a uniform velocity is applied along one direction to simulate the lid's movement, while the remaining two velocity components are set to zero. Specifically, the top boundary of the cavity moves with a velocity of  $u = 1$ , while the transverse components  $v = 0$  and  $w = 0$ , indicating no movement in those directions. All other boundaries are treated as stationary walls, where the no-slip condition applies, meaning  $u = v = w = 0$ . An illustration of the boundary conditions for the 3D lid-driven cavity (LDC) is shown in Figure 16.



(a) The object inside the lid driven cavity domain. (b) Octree mesh showing local refinement near the object boundary.

**Figure 15:** An example shape in the 3D LDC case, along with a slice of the computational mesh



**Figure 16:** Boundary conditions for the 3D lid-driven cavity flow. The top boundary has  $u = 1, v = 0, w = 0$ , while all other boundaries are no-slip.

**Solver configuration:** We use the Petsc linear algebra package for solving the system of equations. We utilize the BCGS solver, with an ASM pre-conditioner. The solver continues iterating until a relative tolerance,  $rtol$ , of  $10^{-8}$ , is reached.

### A.3 Postprocessing of the results to compute force coefficients and Nusselt number

After the CFD solve, we also compute three engineering summary variables. The drag coefficient ( $C_D$ ) represents the non-dimensional force exerted on an object in the direction of the flow. It is calculated using the

formula:

$$C_D = \frac{2F_x}{A_{py}}. \quad (17)$$

where  $F_x$  is the drag force and  $A_{py}$  is the projection area in the y-direction.

The lift coefficient ( $C_L$ ) indicates the non-dimensional lift force acting perpendicular to the flow direction. It is defined as:

$$C_L = \frac{2F_y}{A_{px}}. \quad (18)$$

where  $F_y$  is the lift force and  $A_{px}$  is the projection area in the x-direction.

Finally, we compute the Nusselt number to determine the amount of heat transfer. The local Nusselt number can be defined as:

$$Nu = \nabla T \cdot \mathbf{n}. \quad (19)$$

The averaged Nusselt number can be written as:

$$\overline{Nu} = \frac{\int Nu d\Gamma}{\int d\Gamma}. \quad (20)$$

We compute the average  $Nu$  across both the bottom wall as well as the surface of the object.

## Appendix B. Benchmark Setup

### B.1 Downsampling

We used our in-house FEM simulation software `DendroKT` to perform forward simulations. In order to capture the physics precisely, we use a very fine resolution – with mesh sizes at the Kolmogorov length scale in 2D, and mesh sizes at  $2\times$  the Kolmogorov length scale in 3D – to run the FEM simulations. While this approach allows us to model the physics very precisely, the resulting tensors generated at those resolutions are impractically large in size i.e., running into hundreds of gigabytes. We therefore use the ParaView tool (Ayachit, 2015) to downsample the original resolution to lower resolutions such as,  $512 \times 512$ ,  $256 \times 256$ ,  $128 \times 128 \times 128$  and  $1024 \times 256$ . It is important to note that this downsampling is on the fully resolved data, and thus still captures all the larger scale features (as well as the impact of the small scale features on the large scale features). Note that for the 3D LDC cases, although the forward simulations are performed on an adaptive mesh with refinements up to level-9 near the object and level-7 elsewhere, the final processed results (the .npz tensors) are uniformly downsampled to a uniform  $128 \times 128 \times 128$  resolution.

### B.2 Machine Learning

Our Machine Learning pipeline consists of three modules. In the first module, we undertake the tensorization exercise. We group the available downsampled data to a four or five dimensional *numpy* tensor depending on the nature of the physics outlined in Table 9. Tensorization helps us to subset the data using the tensor notation, and it also makes the data ready to be fed into Neural Operators/Foundation models. Data at this point is fed into the second module, where we primarily train the Neural Operators/Foundation Models. We have two principal objectives in this module, a) Learn the best hyperparameters for a given dataset and a given Neural Operator/Foundation Model. b) Select the best hyperparameter and perform a final training exercise to record the performance of each Neural Operator/Foundation Model on every dataset. Finally, we move to the post-processing module, where we report a complete suite of goodness of fit statistics using the trained model on a held-out data sample. We make our code, for select Neural Operators, publicly available.

**Table 9:** Formulaic description of the input and output tensors. 3000/6000/1150/500 are sample sizes for the dataset. 240 is the number of equi-spaced time snapshots for the FPO case;  $x, y, z$  are the dimensions of a field. E.g., For LDC - NS  $Y[0, 1, :, :]$  indicates the pointwise  $v$  velocity over the entire grid for the first sample in the dataset. As stated in Section 4.5, the FPO data requires some light postprocessing to bring it to the desired `numpy` tensor format.

Dataset	Dim.	Input Tensor	Output Tensor
LDC - NS	2	$X[3000][Re, g, s][x][y]$	$Y[3000][u, v, p][x][y]$
LDC - NS+HT	2	$X[5990][Re, Gr, g, s][x][y]$	$Y[5990][u, v, p, \theta][x][y]$
FPO - NS	2	$X[1150][Re, g, s][x][y]$	$Y[1150][240][u, v, p][x][y]$
LDC - NS	3	$X[500][Re, g, s][x][y][z]$	$Y[500][u, v, p][x][y][z]$

### B.3 Residual Metrics (M4)

One can also evaluate how well the predicted field satisfies the underlying (set of) PDE. Evaluating how well continuity ( $\nabla \cdot \mathbf{u} = 0$ ) is satisfied over the domain provides a measure of mass conservation. Similarly, the global average of the PDE residual evaluates how well the model satisfies the underlying PDE in the domain. To compute these metrics, we adopt a finite element-based approach. For each element  $e$  in the finite element mesh  $\mathcal{K}_h$ , the local residuals for the Navier-Stokes momentum equations in the  $x$ - and  $y$ -directions are computed as:

$$r_x = \frac{\partial u_x}{\partial t} + u \cdot \nabla u_x - \eta \nabla^2 u_x + \frac{\partial p}{\partial x},$$

$$r_y = \frac{\partial u_y}{\partial t} + u \cdot \nabla u_y - \eta \nabla^2 u_y + \frac{\partial p}{\partial y}.$$

These residuals are evaluated at several quadrature points (using a  $2 \times 2$  Gauss quadrature in 2D) via linear basis function interpolation. The  $L_2$  norms of the local residuals for each element are computed as:

$$\|r_x\|_{L_2(e)} = \left( \int_e r_x^2 de \right)^{1/2}, \quad \|r_y\|_{L_2(e)} = \left( \int_e r_y^2 de \right)^{1/2}.$$

The total PDE residual is obtained by summing the squared contributions over all elements:

$$M4 \text{ (momentum)} = r_{\text{Total}} = \sum_{e \in \mathcal{K}_h} \left( \|r_x\|_{L_2(e)}^2 + \|r_y\|_{L_2(e)}^2 \right).$$

In addition, the continuity residual is computed to enforce the incompressibility condition. Specifically, the residual of the continuity equation,

$$\nabla \cdot \mathbf{u} = u_x + v_y + w_z,$$

is evaluated at the same quadrature points. The element-wise continuity residual is calculated as:

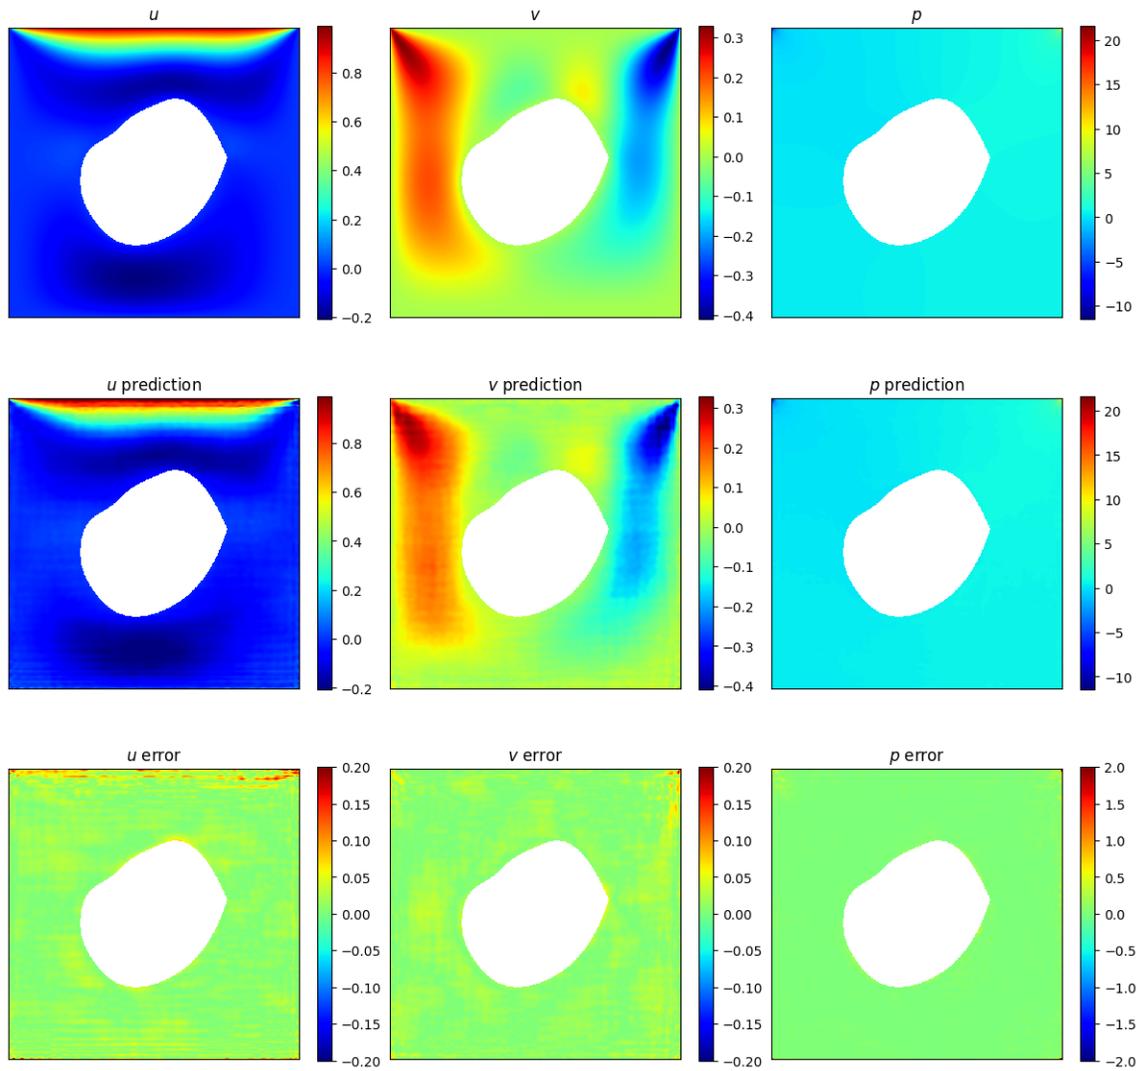
$$\|r_c\|_{L_2(e)} = \left( \int_e (\nabla \cdot \mathbf{u})^2 de \right)^{1/2},$$

and the global continuity residual is given by:

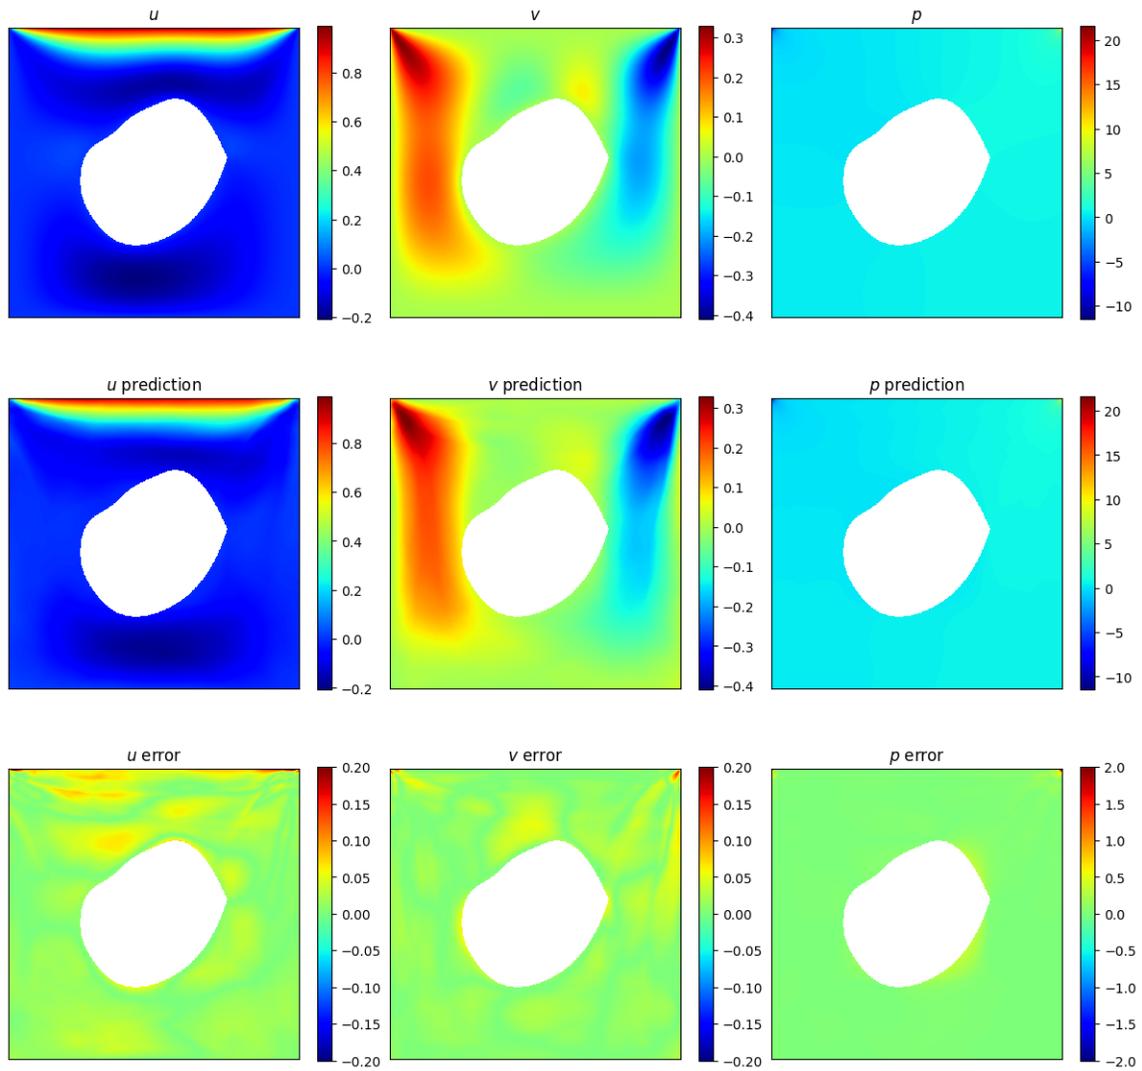
$$M4 \text{ (continuity)} = r_{c, \text{Total}} = \sum_{e \in \mathcal{K}_h} \|r_c\|_{L_2(e)}^2.$$

### B.4 Field Predictions

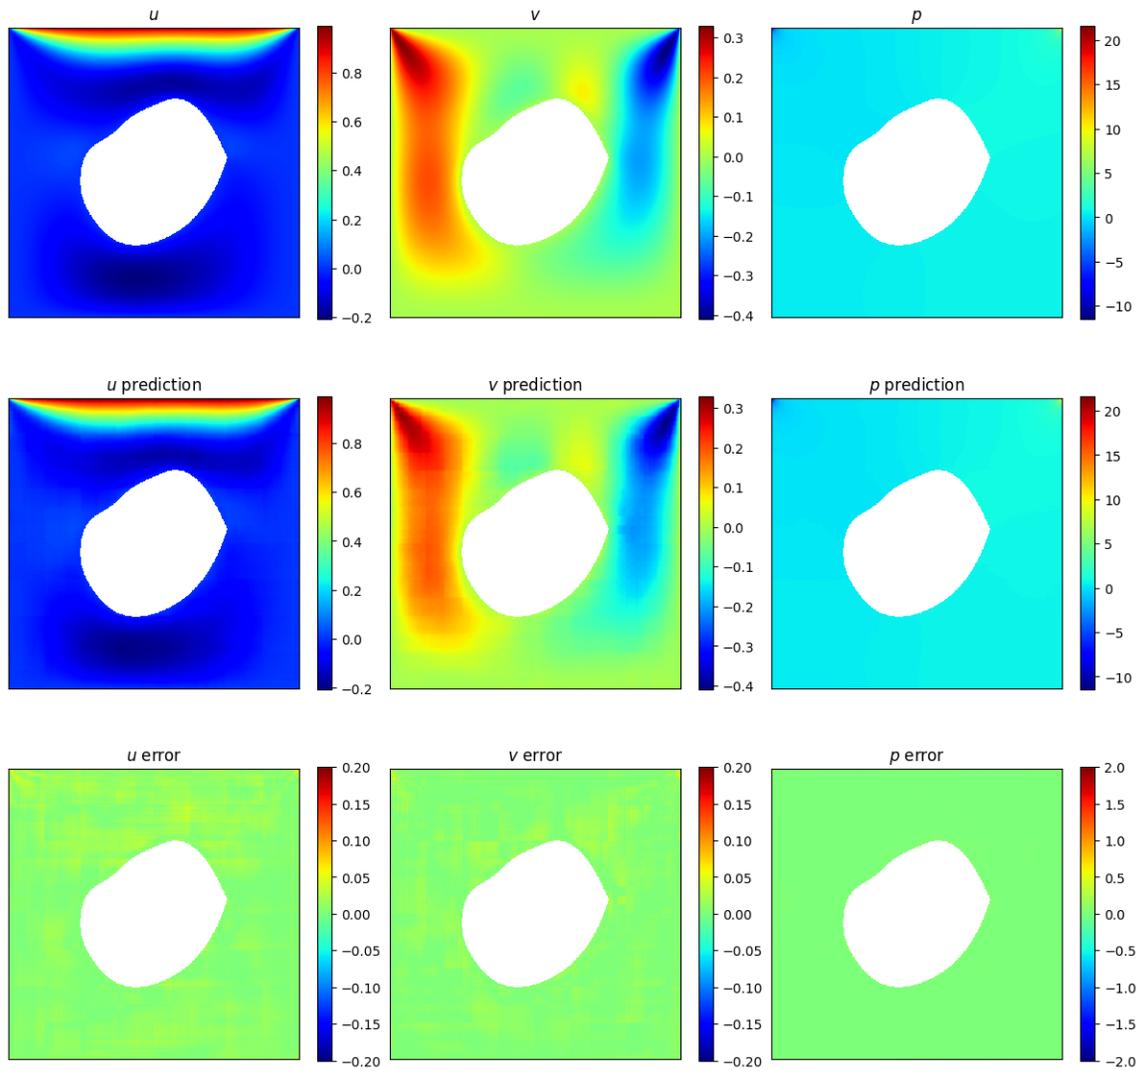
Below, we present the field predictions for a representative sample using scientific machine learning models (FNO, DeepONet, Poseidon-T). The top row illustrates the ground truth fields for  $u$ ,  $v$ , and  $p$ . The middle row displays the model predictions, while the bottom row depicts the corresponding error fields, defined as the difference between the predictions and the ground truth.



**Figure 17:** *Fourier Neural Operator: Prediction vs. Ground Truth and Error for  $u, v, p$ .*



**Figure 18:** DeepONet: Prediction vs. Ground Truth and Error for  $u$ ,  $v$ ,  $p$ .



**Figure 19:** *poseidon-T*: Prediction vs. Ground Truth and Error for  $u, v, p$ .