GeoPI-LNODE: Geometry-conditioned Physics-Informed Latent Neural ODEs for Long-Horizon Prediction

Anonymous authors
Paper under double-blind review

Abstract

We propose a Geometry-conditioned Physics-Informed Latent Neural Ordinary Differential Equations (GeoPI-LNODE), a continuous-time surrogate modeling framework for long-term prediction of complex unsteady systems governed by known partial differential equations (PDEs). High-fidelity numerical simulations are essential for engineering analysis and design, but they often incur prohibitive computational costs, motivating the development of learning-based surrogate models to enable efficient prediction. However, accurate long-horizon prediction using surrogate models remains challenging due to accumulated rollout errors, physical inconsistencies, and difficulties in handling complex geometries. GeoPI-LNODE combines parameter- and boundary-conditioned latent dynamics, geometry-conditioned field reconstruction, and physics-informed regularization to model complex spatiotemporal dynamics while preserving physical consistency over long-horizon rollouts. A latent neural ODE evolves a low-dimensional state, while a coordinate-based decoder reconstructs physical fields and enables physics-informed regularization by using residuals from the governing equations to regularize the latent dynamics in physical space. We evaluate GeoPI-LNODE on coupled fluid–scalar transport, nonlinear hyperelastic dynamics, and three-dimensional aneurysm flow—challenging regimes involving tightly coupled processes, strong nonlinearities, and complex three-dimensional geometries, where surrogate models often struggle to maintain accurate and stable long-term predictions. Within each fixed computational domain, GeoPI-LNODE improves long-term prediction accuracy over representative surrogate baselines and approximately halves the flow-rate imbalance in the three-dimensional flow problem relative to the physics-informed variant without geometry conditioning. These results demonstrate the potential of physics-informed latent dynamics for accurate long-horizon surrogates.

1 Introduction

High-fidelity numerical PDE solvers, such as finite element method (FEM) and finite difference method (FDM), accurately resolve complex spatiotemporal dynamics, but repeated simulations at high spatial resolutions, over long time horizons, and across many parameter configurations create a major computational bottleneck in design optimization, uncertainty quantification, and inverse problems (Peherstorfer et al., 2018). Data-driven surrogate models offer a promising alternative by reducing computational cost and enabling rapid predictions under previously unseen operating conditions (Lu et al., 2021). However, constructing surrogate models that remain both accurate and physically consistent over long-horizon rollouts is still challenging, particularly for nonlinear, coupled, and geometrically complex systems (Sholokhov et al., 2023).

To address these challenges, physics-informed neural networks (PINNs) incorporate governing-equation residuals into the training objective to enforce physical consistency (Raissi et al., 2019), but often require retraining for new parameters or boundary conditions and remain difficult to optimize for strongly nonlinear and long-time dynamics (Krishnapriyan et al., 2021). Neural operators, including DeepONet (Lu et al., 2021) and the Fourier Neural Operator (FNO) (Li et al., 2021), learn mappings from input to solution function spaces, enabling prediction across parameterized inputs

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

with a single trained model. However, learning high-dimensional spatiotemporal solutions directly remains challenging for complex three-dimensional systems and long horizons. In contrast, latent dynamics models instead compress physical fields into low-dimensional states that capture dominant spatiotemporal structures, reducing the complexity of learning and integrating the dynamics without evolving every spatial degree of freedom. In particular, latent neural ODEs (Regazzoni et al., 2024; Nair et al., 2025) model these dynamics in continuous time, providing a flexible framework for long-horizon prediction. However, reconstruction-based training can accurately fit observed trajectories without learning faithful latent dynamics, allowing small dynamical errors to accumulate during extrapolative rollouts.

To overcome these limitations, we propose a geometry-conditioned physics-informed latent neural ordinary differential equation (GeoPI-LNODE). The proposed framework consists of two neural networks: a latent dynamics network and a coordinate-based decoder. The latent dynamics network evolves a low-dimensional latent state conditioned on static physical parameters and time-dependent boundary inputs. The coordinate-based decoder reconstructs the physical fields from the latent state and spatial coordinates, while incorporating geometric information such as a distance function. To impose physical constraints on the learned dynamics, temporal derivatives of the reconstructed fields are obtained by propagating the latent dynamics through the decoder Jacobian via the chain rule, whereas spatial derivatives are evaluated by automatic differentiation with respect to the coordinate inputs. The resulting governing-equation and boundary-condition residuals are then used to regularize the latent trajectories. This formulation constrains the latent evolution through the reconstructed physical fields without explicitly evolving the full high-dimensional discretized state. Furthermore, because the decoder operates on spatial coordinates, the same framework can be applied to both structured grids and unstructured spatial discretizations.

We evaluate the proposed framework on three problems: coupled fluid and scalar transport, nonlinear hyperelastic structural dynamics, and unsteady flow in a complex three-dimensional aneurysm geometry. These problems respectively test the framework on multiphysics coupling, strongly nonlinear deformation, and spatiotemporal field reconstruction in a complex three-dimensional domain. For each problem, we assess generalization to input parameters and time-dependent input waveforms that are not included in the training data, as well as long-term rollout performance over temporal extrapolation intervals extending from twice to as much as ten times the training time range.

Contributions. This paper elucidates the different roles of *geometry* and *physics*-informed surrogate regularization for the latent neural ODE, based on the following contributions.

- **Latent modeling with geometry-conditioned reconstruction.** We formulate a surrogate that evolves a compact latent state conditioned on static physical parameters and prescribed time-dependent boundary inputs. Geometry-related spatial features condition coordinate-based field reconstruction, while a physics-informed surrogate loss regularizes the predicted fields without requiring full-state evolution.
- **Improved prediction over the tested extrapolation horizons.** Across separately trained models for flow–transport, hyperelastic dynamics, and aneurysm flow, GeoPI-LNODE achieves the lowest mean velocity-magnitude errors among the evaluated methods, together with lower concentration and displacement errors in the corresponding problems.
- **Flow-balance improvements beyond pointwise accuracy.** In the aneurysm experiment, adding geometry conditioning to the physics-informed model reduces mean normalized flow-rate imbalance from 22.9% to 12.1%, while mean velocity-magnitude error (root mean squared error) changes only from 0.4033 to 0.4026 (Table 3). This empirical distinction motivates evaluating integral flow balance alongside pointwise prediction error.

2 Related Work

Physics-informed learning. Recent studies have increasingly extended PINNs to complex three-dimensional geometries and more realistic physical systems. In fluid-flow problems, particular attention has been paid to enforcing conservation laws over complex computational domains. MD-PINNs (Hwang, 2025) partition the domain and impose flow consistency across subdomains, while control-volume-based methods such as ICPINN (Liu et al., 2026) and MUSA-PINN (Zhang et al., 2026) enforce integral conservation of mass and momentum to improve physical consistency in com-

plex flow geometries. These studies highlight the importance of maintaining conservation properties beyond local pointwise PDE residuals. Nevertheless, PINNs remain primarily solution-oriented methods, limiting their efficiency for repeated forward predictions. This limitation has motivated increasing interest in incorporating physics-informed learning into surrogate modeling approaches, including neural operators and latent neural ODEs, which aim to provide rapid predictions across varying physical parameters and boundary conditions using a single trained model.

Neural operators. Representative physics-informed neural operators include Physics-Informed DeepONet (PI-DeepONet) (Wang et al., 2021) and Physics-Informed Neural Operators (PINO) (Li et al., 2022), which incorporate governing-equation residuals into operator learning for parameterized PDEs. In parallel, geometry-aware extensions have been developed for complex domains: Geo-FNO learns mappings between irregular physical and regular latent domains (Li et al., 2023a), while GINO incorporates geometric representations, such as signed distance functions and point clouds, to enable large-scale prediction on complex three-dimensional geometries (Li et al., 2023b). More recently, PI-GANO combines physics-informed training with geometry-aware operator learning to generalize across both PDE parameters and domain geometries (Zhong & Meidani, 2025). These methods motivate incorporating physical and geometric information into surrogate learning. Our study examines a complementary formulation: explicitly evolving compact latent dynamics under prescribed time-dependent forcing and assessing how geometry-conditioned reconstruction and physics-informed regularization affect temporal extrapolation and integral flow balance.

Latent dynamical models. Latent dynamical models represent high-dimensional spatiotemporal fields through low-dimensional latent states and learn their temporal evolution from data. Neural ODEs provide a natural continuous-time formulation for such dynamics (Chen et al., 2018), while Latent Dynamics Networks (LDNets) combines parameter- and input-conditioned latent dynamics with coordinate-based field reconstruction (Regazzoni et al., 2024). More recently, CNF-ROM, DLDMF (Liang & Gao, 2026), and PIDO have incorporated governing physical laws into latent neural ODE frameworks to improve generalization and temporal extrapolation (Kim et al., 2024; Wang et al., 2025). Despite these advances, existing evaluations have largely focused on relatively low-dimensional or geometrically regular PDE settings, and the combination of physics-informed latent dynamics with explicit geometry conditioning, unstructured spatial discretizations, multi-physics coupling, and strongly nonlinear dynamics remains less explored. A detailed positioning of the present work relative to these methods is provided in Table A.1 of Appendix A. Relative to the input-conditioned LDNets formulation, GeoPI-LNODE incorporates geometry-related features into field reconstruction and adds a physics-informed surrogate loss. Within the broader family of physics-informed latent models, including PIDO, CNF-ROM, and DLDMF, our focus is the joint treatment of time-dependent forcing and geometry-conditioned reconstruction, and the distinct effects of these choices on extrapolation error and inlet–outlet flow balance.

3 METHODS

3.1 PROBLEM FORMULATION

Our goal is to construct a surrogate model for numerical simulators that takes static physical parameters $\boldsymbol{\mu} \in \mathcal{P}$ of the governing equations and/or time-dependent boundary conditions $\mathbf{b}(t) \in \mathbb{R}^{d_b}$ as inputs and outputs unsteady physical fields

$$\mathbf{u}_{\boldsymbol{\mu},\mathbf{b}} : \Omega \times [t_0, T] \to \mathbb{R}^C, \tag{1}$$

where $\Omega \subset \mathbb{R}^d$ is the spatial domain, $d \in \{2, 3\}$ is the spatial dimension, and $[t_0, T]$ is the time interval. We set $t_0 = 0$ in all experiments. Let $\mathbf{x} \in \Omega$ and $t \in [t_0, T]$ denote the spatial coordinate and time, respectively. The vector $\mathbf{u}_{\mu,\mathbf{b}}(\mathbf{x}, t) \in \mathbb{R}^C$ contains the C physical-field components at $(\mathbf{x}, t)$.

For each physical system, we assume that the governing differential operators $\{\mathcal{F}_m\}_{m=1}^M$, the associated initial and boundary conditions, and the parameterization of the prescribed inputs are known. Here, M denotes the number of governing equations for the physical system. Numerical reference solutions are available for a finite collection of parameter and boundary-condition configurations. We use these solutions together with the known governing equations to train a surrogate for each physical system. For these configurations, the following ground-truth samples are available for

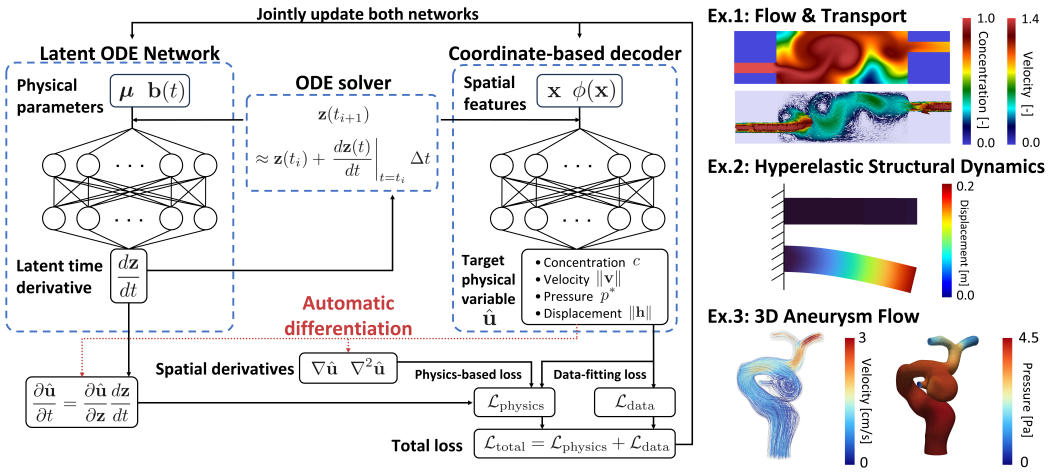

Figure 1: Overview of the proposed physics-informed conditional latent neural ODE. Static physical parameters $\boldsymbol{\mu}$ and time-dependent inputs $\mathbf{b}(t)$ condition the latent dynamics. A coordinate-based decoder reconstructs the physical fields from the latent state and spatial features. Governing-equation residuals evaluated on the reconstructed fields are incorporated into the training objective.

training:

$$\mathcal{U} = \left\{ \left(\boldsymbol{\mu}_n, \mathbf{b}_n(\cdot), \mathbf{x}_{ni}, t_{ni}, \mathbf{u}_{\boldsymbol{\mu}_n,\mathbf{b}_n}(\mathbf{x}_{ni}, t_{ni})\right) \;\middle|\; n \in \{1,\dots,N\},\ i \in \{1,\dots,I_n\} \right\}. \tag{2}$$

Here, N denotes the number of parameter and boundary-condition configurations, and I_n denotes the number of spatiotemporal data points sampled for the n-th configuration. Thus, the total number of training samples is $|\mathcal{U}| = \sum_{n=1}^{N} I_n$. The notation $\mathbf{b}_n(\cdot)$ represents the complete time-dependent boundary condition for the n-th configuration. To simplify notation, we omit the dependence of $\mathbf{u}$ on $\boldsymbol{\mu}$, $\mathbf{b}$, $\mathbf{x}$, and t whenever clear from the context.

3.2 Conditional Latent Neural ODE

The proposed model is composed of two neural networks: a latent ODE network f_θ and a coordinate-based decoder D_ψ, where θ and ψ denote the trainable parameters of these networks. The latent ODE network f_θ takes the physical parameters $\boldsymbol{\mu}$, current boundary conditions $\mathbf{b}(t)$, and current latent state $\mathbf{z}(t)$ as inputs and defines the latent dynamics as

$$\frac{d\mathbf{z}(t)}{dt} = f_\theta\big(\mathbf{z}(t), \boldsymbol{\mu}, \mathbf{b}(t)\big), \tag{3}$$

where $\mathbf{z}(t)$ is the low-dimensional representation of the system state.

The continuous latent trajectory $\mathbf{z}(t)$ is obtained by solving the initial-value problem with $\mathbf{z}(t_0) = \mathbf{0}$, since all simulations considered in this study start from the same quiescent physical state. To obtain a discrete sequence of latent states, which is the input of the coordinate-based decoder $\mathcal{D}_\psi$, we numerically integrate the ODE. For example, using the Euler method with time step Δt, the sequence $\{\mathbf{z}(t_i)\}_{i=0}^{I}$, where $t_I = T$, is computed as

$$\mathbf{z}(t_{i+1}) \approx \mathbf{z}(t_i) + \frac{d\mathbf{z}(t)}{dt}\bigg|_{t=t_i} \Delta t = \mathbf{z}(t_i) + f_\theta\big(\mathbf{z}(t_i), \boldsymbol{\mu}, \mathbf{b}(t_i)\big)\,\Delta t, \quad i \in \{0,\dots,I-1\}. \tag{4}$$

The coordinate-based decoder D_ψ reconstructs the physical fields $\hat{\mathbf{u}}(\mathbf{x}, t)$ from the latent state $\mathbf{z}(t)$, the spatial coordinate $\mathbf{x}$, and auxiliary spatial features $\boldsymbol{\phi}(\mathbf{x})$:

$$\hat{\mathbf{u}}(\mathbf{x}, t) = D_\psi\Big(\mathbf{z}(t), \mathbf{x}, \operatorname{stopgrad}\big(\boldsymbol{\phi}(\mathbf{x})\big)\Big). \tag{5}$$

Here $\boldsymbol{\phi}(\mathbf{x})$ denotes the time-independent spatial features. For the flow problems, we use geometry-derived signed-distance features, whereas for the hyperelastic beam, we use a normalized reference-coordinate feature. These features provide explicit geometry-related or boundary-relative information to the decoder in addition to the spatial coordinates. Their definitions are provided in Section 4 and Appendix C.

For each benchmark, the spatial features are fixed across all parameter configurations; thus, we evaluate parameter and temporal generalization within fixed domains rather than across geometries. The $\mathrm{stopgrad}$ operator leaves $\boldsymbol{\phi}(\mathbf{x})$ unchanged in the forward pass but blocks its derivative path with respect to $\mathbf{x}$. Consequently, geometry information remains available for field reconstruction, while spatial derivatives are taken only through the explicit coordinate input; gradients with respect to ψ and through $\mathbf{z}(t)$ remain intact. Although $\boldsymbol{\phi}(\mathbf{x})$ varies across query locations, its spatial derivative contributions are omitted. The resulting residuals are therefore frozen-feature surrogates that generally differ from the physical strong-form residuals of the composite field $\mathbf{x} \mapsto D_\psi(\mathbf{z}(t), \mathbf{x}, \boldsymbol{\phi}(\mathbf{x}))$, and the associated loss is interpreted as a physics-informed surrogate regularizer. All GeoPI-LNODE results in the main text use this convention, motivated by its observed numerical stability and predictive performance. The benchmark-specific spatial features are shown in Figs. C.1, C.2, and C.3 of Appendix C.

3.3 Training Objective and Physics-Informed Regularization

We use two types of loss terms for training: a data-fitting loss $\mathcal{L}_{\mathrm{data}}$ and a physics-based loss $\mathcal{L}_{\mathrm{physics}}$. The data-fitting loss $\mathcal{L}_{\mathrm{data}}$ is introduced to reduce the difference between the reconstructed physical fields $\hat{\mathbf{u}}(\mathbf{x}, t)$ and the ground-truth physical fields $\mathbf{u}(\mathbf{x}, t)$, and is defined as follows:

$$\mathcal{L}_{\mathrm{data}} = \frac{1}{C|\mathcal{U}|} \sum_{n=1}^{N} \sum_{i=1}^{I_n} \|\hat{\mathbf{u}}_n(\mathbf{x}_{ni}, t_{ni}) - \mathbf{u}_n(\mathbf{x}_{ni}, t_{ni})\|_2^2 \,. \tag{6}$$

The physics-based loss $\mathcal{L}_{\mathrm{physics}}$ regularizes the reconstructed fields using the frozen-feature surrogates of the governing-equation residuals. The governing equations are written as

$$\mathcal{F}_m[\mathbf{u}; \boldsymbol{\mu}, \mathbf{b}(t)](\mathbf{x}, t) = 0, \qquad (m = 1, 2, \ldots, M). \tag{7}$$

To make the frozen-feature differentiation explicit, we introduce the auxiliary reconstructed field

$$\hat{\mathbf{u}}_{\boldsymbol{\eta},n}(\mathbf{x}, t) := D_\psi(\mathbf{z}_n(t), \mathbf{x}, \boldsymbol{\eta}), \tag{8}$$

where $\boldsymbol{\eta}$ is treated as independent of $\mathbf{x}$ during spatial differentiation. The frozen-feature surrogate residual of the m-th governing equation for the n-th configuration is defined as

$$\mathbf{r}_{mn}(\mathbf{x}, t) := \mathcal{F}_m[\hat{\mathbf{u}}_{\boldsymbol{\eta},n}; \boldsymbol{\mu}_n, \mathbf{b}_n(t)](\mathbf{x}, t)|_{\boldsymbol{\eta}=\boldsymbol{\phi}(\mathbf{x})} \,, \tag{9}$$

where $\mathbf{r}_{mn}(\mathbf{x}, t)$ denotes the frozen-feature surrogate residual of the m-th governing equation for the n-th configuration. For the n-th configuration, let J_{mn} denote the number of collocation points used to evaluate the m-th governing equation, and let $(\mathbf{x}_{mnj}, t_{mnj})$ denote the j-th such point. Then, we define the physics-based loss $\mathcal{L}_{\mathrm{physics}}$ as

$$\mathcal{L}_{\mathrm{physics}} = \sum_{m=1}^{M} \lambda_m \mathcal{L}_m, \qquad \mathcal{L}_m = \frac{1}{\sum_{n=1}^{N} J_{mn}} \sum_{n=1}^{N} \sum_{j=1}^{J_{mn}} \|\mathbf{r}_{mn}(\mathbf{x}_{mnj}, t_{mnj})\|_2^2, \tag{10}$$

where $\mathcal{L}_m$ is the mean-squared residual of the m-th governing equation over all corresponding collocation points, and λ_m controls its contribution to physics.

The spatial derivatives $\nabla\hat{\mathbf{u}}$ and $\nabla^2\hat{\mathbf{u}}$ are computed using automatic differentiation. All spatial derivatives in the training residuals are evaluated using the frozen-feature convention defined in Section 3.2. Because the decoder has no explicit time dependence and the auxiliary spatial features are time independent, the temporal derivative $\partial\hat{\mathbf{u}}/\partial t$ is obtained through the latent state using the chain rule:

$$\frac{\partial \hat{\mathbf{u}}}{\partial t} = \frac{\partial \hat{\mathbf{u}}}{\partial \mathbf{z}} \frac{d\mathbf{z}}{dt} = \frac{\partial D_\psi}{\partial \mathbf{z}} \, f_\theta(\mathbf{z}(t), \boldsymbol{\mu}, \mathbf{b}(t)) \,. \tag{11}$$

We train the proposed model by minimizing the total loss $\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{data}} + \mathcal{L}_{\mathrm{physics}}$. The collocation-point sampling strategy used for residual evaluation and the loss weights $\{\lambda_m\}_{m=1}^{M}$ are provided in Appendix B.

Table 1: Summary of the three evaluation problems. Static parameters are supplied to the latent dynamics throughout the rollout, whereas time-dependent inputs represent prescribed boundary conditions or external forcing.

Problem	Dimension	Predicted fields	Static parameters	Time-dependent input	Physics residuals
Flow–transport	2D	$\mathbf{v}^*, p^*, c$	Re, Sc	$U^*_{\mathrm{mean}}(t)$	continuity, momentum, transport
Hyperelastic beam	2D	$\mathbf{h}, \mathbf{v}$	E	$q(t)$	kinematics, momentum
Aneurysm flow	3D	$\mathbf{v}^*, p^*$	Re, Wo	$U^*_{\mathrm{mean}}(t)$	continuity, momentum

Notation: $\mathbf{v}^*$ denotes nondimensional velocity; p^*, nondimensional pressure; c, scalar concentration; $\mathbf{u}$, displacement; $\mathbf{v}$, velocity; Re, Reynolds number; Sc, Schmidt number; Wo, Womersley number; E, Young's modulus; $U^*_{\mathrm{mean}}(t)$, nondimensional mean inlet velocity; and $q(t)$, time-dependent external traction.

4 Experiments

4.1 Experimental Setup

We evaluated the proposed framework on three problems: (i) *flow and transport*, (ii) *nonlinear hyperelastic structural dynamics*, and (iii) *three-dimensional aneurysm flow*, because they probe complementary challenges in long-horizon surrogate modeling: multiphysics coupling, strongly nonlinear deformation, and field reconstruction on a complex unstructured domain, respectively. A separate model was trained for each system. The predicted variables, conditioning inputs, and physics residuals used for each problem are summarized in Table 1. Complete governing equations, boundary conditions, parameter configurations, and numerical simulation settings are provided in Appendix C.

Evaluation protocol. Normalization statistics for the input coordinates and physical parameters were computed using the training set only and subsequently applied to the validation and test sets. Each model was trained over $t_{\mathrm{train}} \in [t_0, T_{\mathrm{train}}]$ and evaluated for long-term prediction over the temporal extrapolation interval $t_{\mathrm{test}} \in (T_{\mathrm{train}}, T_{\mathrm{test}}]$.

For a scalar or vector field $\mathbf{u}$, the root mean squared error (RMSE) over t_{test} was computed over all test cases, time steps, and spatial points as

$$\mathrm{RMSE}_{\mathrm{ext}}(a) = \frac{1}{K}\sum_{k=1}^{K}\sqrt{\frac{1}{N_t N_s}\sum_{n=1}^{N_t}\sum_{i=1}^{N_s}\left(\hat{a}_{k,i}(t_n) - a_{k,i}(t_n)\right)^2}, \tag{12}$$

where K is the number of test cases, N_t is the number of evaluation time points in t_{test}, and N_s is the number of spatial points. The reported scalar field a is p^*, c, the velocity magnitude $\|\mathbf{v}\|_2$, or the displacement magnitude $\|\mathbf{u}\|_2$, as applicable; $\hat{a}$ denotes the corresponding prediction. To further assess global mass conservation over the entire computational domain, we evaluated the flow-rate imbalance between the inlet and outlet boundaries in the three-dimensional aneurysm problem (Fig. 5). The normalized imbalance was defined as $\epsilon_Q = \frac{\sum_{n=1}^{N_t}|Q_{\mathrm{out}}(t_n) - Q_{\mathrm{in}}(t_n)|}{\sum_{n=1}^{N_t}|Q_{\mathrm{in}}(t_n)|}$, where Q_{in} and Q_{out} denote the volumetric flow rates at the inlet and outlet boundaries, respectively.

Comparison methods. We compared GeoPI-LNODE with Baseline LNODE, DeepONet, FNO, and CNF-ROM (Kim et al., 2024) on the two-dimensional problems. For three-dimensional aneurysm flow, FNO was replaced by MeshGraphNet (Pfaff et al., 2021), which supports unstructured meshes. Baseline LNODE uses the same architecture and optimization settings as GeoPI-LNODE but omits the physics-based losses. Further details are provided in Appendix D.

Ablation study. To assess the contribution of three components of GeoPI-LNODE, namely the latent neural ODE, physics-informed regularization, and geometry conditioning, we conducted ablation experiments in which each component is removed independently. All ablated variants were trained and evaluated using the same data splits and hyperparameter settings as the full model.

Table 2: Comparison with existing surrogate models and ablation study over t_{test}. RMSE values are reported as the mean ± standard deviation across test cases.

Problem	Method	p^*	c	$\|\mathbf{v}\|$	$\|\mathbf{h}\|$
Flow–transport	Baseline LNODE	0.252 ± 0.191	0.104 ± 0.00768	0.0719 ± 0.0110	–
	DeepONet (Lu et al., 2021)	0.801 ± 0.910	0.194 ± 0.0745	0.156 ± 0.0344	–
	FNO (Li et al., 2021)		Diverged		–
	CNF-ROM (Kim et al., 2024)	0.214 ± 0.226	0.101 ± 0.0486	0.0860 ± 0.0390	–
	w/o spatial feature	0.612 ± 0.473	0.167 ± 0.00823	0.125 ± 0.0180	–
	w/o latent dynamics	1.76 ± 1.06	0.692 ± 0.0322	0.431 ± 0.0212	–
	w/o physics-informed	0.259 ± 0.291	0.0961 ± 0.0199	0.0908 ± 0.0252	–
	GeoPI-LNODE	$\mathbf{0.173 \pm 0.220}$	$\mathbf{0.0834 \pm 0.0217}$	$\mathbf{0.0521 \pm 0.0126}$	–
Hyperelastic beam	Baseline LNODE	–	–	0.130 ± 0.107	0.00729 ± 0.00602
	DeepONet (Lu et al., 2021)	–	–	1.79 ± 0.0945	0.133 ± 0.00728
	FNO (Li et al., 2021)	–	–	81.9 ± 13.4	2.52 ± 0.592
	CNF-ROM (Kim et al., 2024)	–	–	Diverged	
	w/o spatial feature	–	–	0.0897 ± 0.0877	0.00429 ± 0.00203
	w/o latent dynamics	–	–	0.185 ± 0.129	0.0105 ± 0.00288
	w/o physics-informed	–	–	0.124 ± 0.110	0.00897 ± 0.00638
	GeoPI-LNODE	–	–	$\mathbf{0.0794 \pm 0.0516}$	$\mathbf{0.00392 \pm 0.00243}$
Aneurysm flow	Baseline LNODE	10.75 ± 2.057	–	4057 ± 4022	–
	DeepONet (Lu et al., 2021)		Diverged		–
	MeshGraphNet (Pfaff et al., 2021)	9.350 ± 2.283	–	0.4800 ± 0.08098	–
	CNF-ROM (Kim et al., 2024)	10.73 ± 2.041	–	0.4931 ± 0.06224	–
	w/o spatial feature	$\mathbf{9.107 \pm 2.203}$	–	0.4033 ± 0.03450	–
	w/o latent dynamics	192.7 ± 17.43	–	24.78 ± 0.4330	–
	w/o physics-informed	9.223 ± 2.338	–	0.4141 ± 0.03786	–
	GeoPI-LNODE	9.216 ± 2.348	–	$\mathbf{0.4026 \pm 0.03377}$	–

Notation: p^* denotes the nondimensional pressure; c denotes the concentration; $\|\mathbf{v}\|$ denotes the velocity magnitude; and $\|\mathbf{h}\|$ denotes the displacement magnitude. "–" indicates that the corresponding variable is not applicable. "Diverged" indicates that $\mathrm{RMSE}_{\text{ext}}$ exceeded 10^4.

4.2 Three practical scenarios

Flow and transport. The first problem considers two-dimensional incompressible flow coupled with passive scalar transport in a fixed rectangular computational domain containing an implicitly defined chamber geometry represented by a signed distance function (SDF) shown in Fig. C.1. GeoPI-LNODE predicts the nondimensional velocity $\mathbf{v}^*(\mathbf{x}, t)$, pressure $p^*(\mathbf{x}, t)$, and concentration $c(\mathbf{x}, t)$, conditioned on the Reynolds and Schmidt numbers, the time-dependent mean inlet velocity, and a signed distance function describing the geometry. The physics-informed regularization includes continuity, momentum, and advection–diffusion residuals. T_{train} and T_{test} are set to 20.0 s and 40.0 s. Governing equations and simulation details are provided in Appendix C.1 and dataset details in Appendix E.1.

Nonlinear hyperelastic structural dynamics. The second problem considers the transient deformation of a two-dimensional hyperelastic cantilever subjected to a time-dependent external traction $q(t)$. The model predicts displacement $\mathbf{h}(\mathbf{X}, t)$ and velocity $\mathbf{v}(\mathbf{X}, t)$, conditioned on Young's modulus E, the applied traction, and a boundary-related spatial feature. Here, $\mathbf{X}$ denotes the material coordinate in the reference configuration. The material is modeled using a compressible Neo-Hookean constitutive law, and the physics-informed regularization enforces kinematic consistency and momentum balance. T_{train} and T_{test} were set to 10.0 s and 40.0 s. Governing equations and simulation details are provided in Appendix C.2 and dataset details in Appendix E.2.

Three-dimensional aneurysm flow. The final problem considers unsteady incompressible flow in a complex three-dimensional aneurysm geometry. GeoPI-LNODE predicts the nondimensional velocity $\mathbf{v}^*(\mathbf{x}, t)$ and pressure $p^*(\mathbf{x}, t)$, conditioned on the Reynolds and Womersley numbers, the time-dependent mean inlet velocity, and a signed distance feature describing the vessel geometry. The physics-informed regularization consists of continuity and momentum residuals. T_{train} and T_{test} were set to 10.0 s and 100.0 s. Governing equations and simulation details are provided in Appendix C.3 and dataset details in Appendix E.3.

5 Results

As summarized in Table 2, GeoPI-LNODE consistently achieved the lowest $\mathrm{RMSE}_{\text{ext}}$ for the primary predicted fields across all three benchmark problems among the comparison methods and

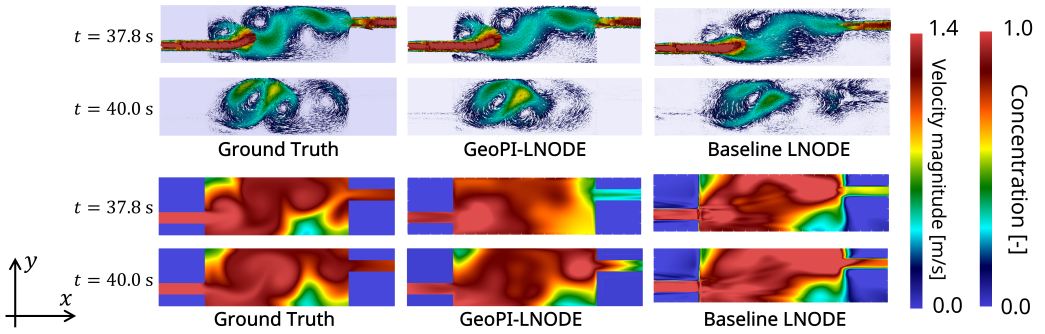

Figure 2: Snapshots of the velocity magnitude and scalar concentration at $t = 37.8$ s and $t = 40.0$ s. Results are shown for the reference solution, GeoPI-LNODE, and Baseline LNODE.

ablation variants. Each entry represents the mean $\pm$ standard deviation of the case-wise $\mathrm{RMSEext}$ over the test cases. For the flow–transport problem, Fig. 2 compares the ground-truth and predicted distributions of velocity magnitude and concentration at $t = 37.8$ s and $t = 40.0$ s, GeoPI-LNODE more accurately preserved the dominant flow structures, including the vortex location and swirling-flow pattern, and reproduced the spatial concentration distribution, whereas the Baseline LNODE exhibited larger deviations from the ground truth in both the velocity and concentration fields. The corresponding time-resolved RMSE over the temporal extrapolation interval is provided in Fig. F.1. Additionally, snapshots of the spatial distributions of the predicted velocity and concentration fields are provided in Fig. F.2.

For the hyperelastic beam, GeoPI-LNODE exhibited less error accumulation over the temporal extraporation interval t_{test} than the Baseline LNODE (Fig. 3). The spatial distribution of deformation for a representative case is provided in Fig. F.3.

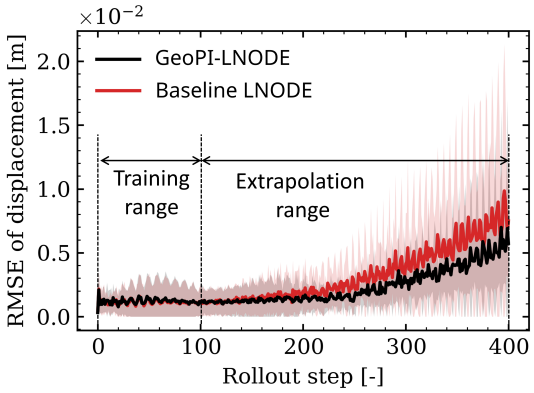

Figure 3: Time-resolved displacement RMSE during the rollout for GeoPI-LNODE and the Baseline LNODE on the hyperelastic beam problem. Shaded regions indicate one standard deviation across test configurations.

For the three-dimensional aneurysm flow, the rollout RMSE of velocity magnitude and streamline snapshots for the ground truth, GeoPI-LNODE, and Baseline LNODE are shown in Fig. 4. GeoPI-LNODE exhibited lower RMSE than the Baseline LNODE over the t_{test}. In addition, the streamline snapshots show that GeoPI-LNODE maintained coherent and physically plausible flow structures over time, whereas the streamlines predicted by the Baseline LNODE became increasingly distorted over time and failed to preserve the reference flow organization at both $t = 60.0$ s and $t = 99.4$ s. These results indicate that GeoPI-LNODE better preserves the large-scale flow structure during long-horizon extrapolation, even for complex three-dimensional unsteady flow.

Regarding the ablation study, GeoPI-LNODE achieved the lowest RMSEs across all evaluated quantities except pressure in the three-dimensional aneurysm flow problem, as shown in Table 2. In addition, the evaluation of flowrate imbalance showed that GeoPI-LNODE yielded the lowest flow-rate mismatch among all ablation variants, reducing the flow-rate imbalance to approximately half that of the physics-informed variant without SDF conditioning (Table 3). These results indicate that the proposed framework improves not only local physical consistency enforced through the physics-informed regularization but also global mass conservation across the computational domain.

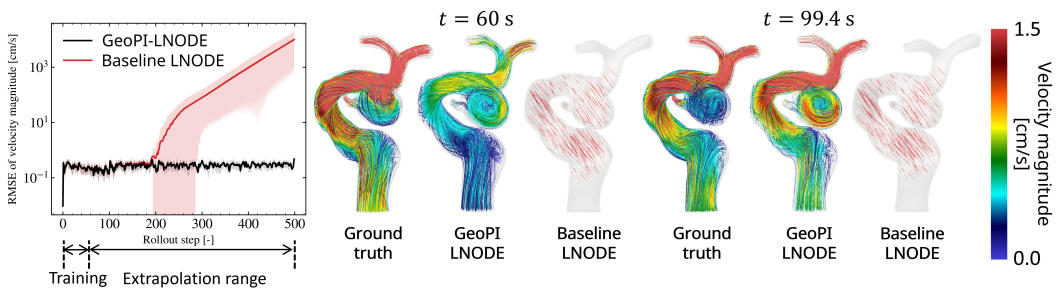

Figure 4: Rollout RMSE for velocity magnitude in the three-dimensional aneurysm problem (left), and streamline visualizations at $t = 60.0\,\mathrm{s}$ and $t = 99.4\,\mathrm{s}$ (right). Results are shown for the reference solution, GeoPI-LNODE, and the Baseline LNODE.

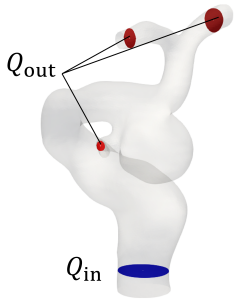

Figure 5: The inlet and outlet boundaries used to compute the flow-rate imbalance.

Table 3: Ablation of physics-informed regularization (P) and SDF-based geometry conditioning (G) for the three-dimensional aneurysm problem. All variants retain the latent ODE. Values are reported as the mean $\pm$ one standard deviation across test configurations, with the best results shown in bold.

Physics reg. P	SDF cond. G	Velocity-magnitude $\mathrm{RMSE}_{\mathrm{ext}} \downarrow$	Flow-rate imbalance ϵ_Q [%] $\downarrow$
–	–	$(4.06 \pm 4.02) \times 10^3$	117.9 ± 44.9
–	✓	0.4141 ± 0.0379	102.9 ± 38.7
✓	–	0.4033 ± 0.0345	22.9 ± 8.3
✓	✓	$\mathbf{0.4026 \pm 0.0338}$	$\mathbf{12.1 \pm 2.5}$

6 Discussion and Conclusion

As discussed in Section 2, previous latent reduced-order models have primarily been evaluated on relatively simple two-dimensional domains, while their applicability to geometrically complex fixed domains remains less well studied. Our results show that continuous latent dynamics can be combined with coordinate-based, geometry-aware decoding for long-horizon prediction within three fixed benchmark domains. This provides evidence that low-dimensional latent representations can remain effective beyond simple domains.

Geometry representations have also been incorporated into previous reduced-order model (ROM) frameworks such as CNF-ROM (Kim et al., 2024). However, whereas such approaches primarily use geometric information to enforce boundary constraints, GeoPI-LNODE directly conditions the reconstruction of the internal physical field on the geometry representation. Our ablation results show that this conditioning improves physical consistency, particularly the inlet–outlet flow-rate balance, even when differences in pointwise velocity error remain relatively modest. This suggests that geometric information contributes not only to boundary conformity but also to the physically consistent reconstruction of the interior solution.

Overall, these findings highlight the complementary roles of geometry-aware conditioning and physics-informed regularization for long-horizon surrogate modeling within fixed domains. Notably, pointwise prediction accuracy and integral flow balance respond differently to geometry conditioning, indicating that both local and global metrics should be considered when evaluating physics-based surrogates.

Scope and future directions. This study focuses on parameter and temporal generalization within fixed domains, enabling a controlled evaluation of geometry-aware decoding and physics-informed regularization across complementary physical systems. Extending the framework to varying geometries may require transferable geometry representations and cross-domain latent alignment, providing a distinct direction for future work.

REPRODUCIBILITY STATEMENT

The source code, configuration files, and scripts for data generation and evaluation will be made publicly available upon acceptance.

ETHICS STATEMENT

This study involves only computational modeling and does not include new experiments on human participants or animals. No personally identifiable information was used.

AI USE STATEMENT

Generative AI tools were used to improve the English writing of this manuscript, including polishing sentence structure, grammar, clarity, and overall presentation. All AI-assisted revisions were carefully reviewed and edited by the authors to ensure that they accurately reflected the intended scientific content. The authors take full responsibility for the final content of this work, including all text, claims, and conclusions.

REFERENCES

Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/69386f6bb1dfed68692a24c8686939b9-Paper.pdf.

Christophe Geuzaine and Jean-François Remacle. Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *International Journal for Numerical Methods in Engineering*, 79(11):1309–1331, 2009. doi: 10.1002/nme.2579. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/nme.2579.

Leslie K Hwang. Physics-informed machine learning with domain decomposition and global dynamics for three-dimensional intersecting flows. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=rOuRpOA6pm.

Minji Kim, Tianshu Wen, Kookjin Lee, and Youngsoo Choi. Physics-informed reduced order model with conditional neural fields, 2024. URL https://arxiv.org/abs/2412.05233v1. Comparison based on arXiv version 1.

Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael Mahoney. Characterizing possible failure modes in physics-informed neural networks. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 26548–26560. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/df438e5206f31600e6ae4af72f2725f1-Paper.pdf.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

Zongyi Li, Hongkai Zheng, Nikola Borislavov Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Andrew Stuart, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations, 2022. URL https://openreview.net/forum?id=dtYnHcmQKeM.

Zongyi Li, Daniel Zhengyu Huang, Burigede Liu, and Anima Anandkumar. Fourier neural operator with learned deformations for PDEs on general geometries. *Journal of Machine Learning Research*, 24(388):1–26, 2023a. URL http://jmlr.org/papers/v24/23-0064.html.

Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, and Animashree Anandkumar. Geometry-informed neural operator for large-scale 3d pdes. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 35836–35854. Curran Associates, Inc., 2023b. doi: 10.52202/075280-1556. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/70518ea42831f02afc3a2828993935ad-Paper-Conference.pdf.

Zhangyong Liang and Huanhuan Gao. Parameterized temporal extrapolation of PDEs via disentangled latent dynamics manifold fusion, 2026. URL https://arxiv.org/html/2603.12676v5. Title and comparison follow the version-5 full text.

Youqiong Liu, Li Cai, Yaping Chen, Jing Xue, Wangwei He, Wenxian Xie, and Jie Wei. Integral conservation physics-informed neural networks with different network architectures for patient-specific aortic flow simulations. *International Journal of Heat and Fluid Flow*, 117:110011, 2026. ISSN 0142-727X. doi: https://doi.org/10.1016/j.ijheatfluidflow.2025.110011. URL https://www.sciencedirect.com/science/article/pii/S0142727X25002693.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021. URL https://doi.org/10.1038/s42256-021-00302-5.

Ashish S. Nair, Shivam Barwey, Pinaki Pal, Jonathan F. MacArt, Troy Arcomano, and Romit Maulik. Understanding latent timescales in neural ordinary differential equation models of advection-dominated dynamical systems. *Physica D: Nonlinear Phenomena*, 476:134650, 2025. ISSN 0167-2789. doi: https://doi.org/10.1016/j.physd.2025.134650. URL https://www.sciencedirect.com/science/article/pii/S0167278925001290.

Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Survey of multifidelity methods in uncertainty propagation, inference, and optimization. *SIAM Review*, 60(3):550–591, 2018. doi: 10.1137/16M1082469. URL https://doi.org/10.1137/16M1082469.

Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=roNqYL0_XP.

M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2018.10.045. URL https://www.sciencedirect.com/science/article/pii/S0021999118307125.

Francesco Regazzoni, Stefano Pagani, Mattia Salvador, et al. Learning the intrinsic dynamics of spatio-temporal processes through latent dynamics networks. *Nature Communications*, 15:1834, 2024. doi: 10.1038/s41467-024-45323-x. URL https://doi.org/10.1038/s41467-024-45323-x.

Aleksei Sholokhov, Yuying Liu, Hassan Mansour, and Saleh Nabi. Physics-informed neural ODE (PINODE): Embedding physics into models using collocation points. *Scientific Reports*, 13:10166, 2023. doi: 10.1038/s41598-023-36799-6. URL https://doi.org/10.1038/s41598-023-36799-6.

Honghui Wang, Yifan Pu, Shiji Song, and Gao Huang. Advancing generalization in PINNs through latent-space representations, 2024. URL https://arxiv.org/abs/2411.19125v1. Comparison based on arXiv version 1.

Honghui Wang, Yifan Pu, Shiji Song, and Gao Huang. Advancing generalization in pinns through latent-space representations. *IEEE Transactions on Neural Networks and Learning Systems*, 36(12):20243–20257, 2025. doi: 10.1109/TNNLS.2025.3598617.

Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric partial differential equations with physics-informed deeponets. *Science Advances*, 7(40):eabi8605, 2021. doi: 10.1126/sciadv.abi8605.

Nathan M. Wilson, Ana K. Ortiz, and Allison B. Johnson. The vascular model repository: A public resource of medical imaging data and blood flow simulation results. *Journal of Medical Devices*, 7(4):040923, 2013. doi: 10.1115/1.4025983. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC4023857/.

Weizheng Zhang, Xunjie Xie, Hao Pan, Xiaowei Duan, Bingteng Sun, Qiang Du, and Lin Lu. MUSA-PINN: Multi-scale weak-form physics-informed neural networks for fluid flow in complex geometries. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=1gYQX9L1cb.

Weiheng Zhong and Hadi Meidani. Physics-informed geometry-aware neural operator. *Computer Methods in Applied Mechanics and Engineering*, 434:117540, 2025. doi: 10.1016/j.cma.2024.117540.

A Position of this work

Table A.1: Positioning of GeoPI-LNODE relative to selected latent-dynamics and physics-informed surrogate formulations.

	Model formulation					Demonstrated evaluation	
Method	**Continuous latent dynamics**	**Time-varying input $b(t)$**	**Coordinate field decoder**	**Explicit geometry feature**	**PDE loss**	**Future-time rollout**	**Complex geometry**
LDNets (Regazzoni et al., 2024)	✓	✓	✓	—	—	✓	ND
PIDO (Wang et al., 2024)	✓	—	✓	—	✓	✓	ND
CNF-ROM (Kim et al., 2024)	✓	—	✓	✓[†]	✓	✓	ND
DLDMF (Liang & Gao, 2026)	✓	—	✓	—	✓	✓	ND
GeoPI-LNODE (ours)	✓	✓	✓	✓[‡]	✓	✓	✓

Notation and scope. ✓ denotes a component included in the reported formulation or an evaluation demonstrated in the cited work; — denotes a component not used in the compared formulation; and ND denotes that the corresponding evaluation was not demonstrated, rather than being inherently unsupported. [†]CNF-ROM incorporates geometry through an approximate distance function (ADF)-based trial representation. [‡]GeoPI-LNODE incorporates geometry through an SDF/spatial feature representation.

B Collocation-point sampling strategy and loss weights for each residual

The physics-informed loss was evaluated using a finite set of spatiotemporal collocation points sampled at each training epoch. Because evaluating the governing-equation residuals at all available spatiotemporal points would be computationally and memory intensive, temporal points were randomly sampled from the training interval, and a fixed number of spatial points were randomly selected at each sampled time. For each benchmark problem, the same sampling strategy and number of collocation points were used for all ablation studies. The numbers of sampled temporal and spatial points are summarized in Table B.1.

Table B.1: Collocation-point sampling strategy and loss weights used for the physics-informed residuals in each benchmark problem. N_t denotes the number of sampled temporal points per epoch, and N_s denotes the number of sampled spatial collocation points at each selected time.

Benchmark	N_t	N_s	Residual	Weight
Flow–transport	16	256	Continuity	$\lambda_{\mathrm{div}} = 10^{-2}$
			Navier–Stokes	$\lambda_{\mathrm{NS}} = 10^{-3}$
			Advection–diffusion	$\lambda_{\mathrm{ADE}} = 10^{-3}$
Hyperelastic beam	3	512	Kinematic	$\lambda_{\mathrm{kin}} = 10^{-3}$
			Momentum	$\lambda_{\mathrm{mom}} = 10^{-4}$
Aneurysm flow	15	512	Continuity	$\lambda_{\mathrm{div}} = 10^{-4}$
			Navier–Stokes	$\lambda_{\mathrm{NS}} = 10^{-5}$

C GOVERNING EQUATIONS AND NUMERICAL SIMULATIONS

C.1 FLOW AND TRANSPORT

The incompressible flow and passive scalar transport are governed by the nondimensional Navier–Stokes and advection–diffusion equations:

$$\begin{aligned} \nabla \cdot \mathbf{v}^* &= 0, \\ \frac{\partial \mathbf{v}^*}{\partial t} + (\mathbf{v}^* \cdot \nabla)\,\mathbf{v}^* + \nabla p^* - \frac{1}{\mathrm{Re}} \nabla^2 \mathbf{v}^* &= \mathbf{0}, \\ \frac{\partial c}{\partial t} + \mathbf{v}^* \cdot \nabla c - \frac{1}{\mathrm{Re}\,\mathrm{Sc}} \nabla^2 c &= 0. \end{aligned} \tag{C.1}$$

Here, $\mathbf{v}^*$ and p^* denote the nondimensional velocity and pressure, respectively, c denotes the scalar concentration, and Re and Sc denote the Reynolds and Schmidt numbers, respectively.

The Reynolds and Schmidt numbers are defined as

$$\mathrm{Re} = \frac{U_{\mathrm{ref}} L_{\mathrm{ref}}}{\nu}, \qquad \mathrm{Sc} = \frac{\nu}{D}, \tag{C.2}$$

where U_{ref} is the reference velocity, L_{ref} is the reference length, ν is the kinematic viscosity, and D is the scalar diffusivity. The reference velocity U_{ref} was defined as the maximum cross-sectional mean inlet velocity.

The prescribed inlet waveform is expressed as

$$U^*_{\mathrm{mean}}(t) = \frac{U_{\mathrm{mean}}(t)}{U_{\mathrm{ref}}} = \frac{1 - \cos(2\pi f t)}{2}. \tag{C.3}$$

A parabolic velocity profile with cross-sectional mean velocity $U_{\mathrm{mean}}(t)$ was prescribed at the inlet. A no-slip boundary condition was imposed on the walls, and the nondimensional pressure was set to $p^* = 0$ at the outlet.

For the scalar transport problem, a prescribed concentration $c = 1.0$ was imposed at the inlet, while a zero-flux condition was applied on the walls and a zero-gradient condition was imposed at the outlet. The governing equations were spatially discretized using the finite element method. For the incompressible Navier-Stokes equations, P2/P1 Taylor–Hood elements were used for velocity and pressure fields, respectively, and time integration was performed using a semi-implicit incremental pressure-correction scheme (IPCS). The advection term was treated explicitly, while the viscous term was discretized using a Crank–Nicolson scheme. The advection-diffusion equation was discretized using P1 elements and integrated in time using an implicit backward Euler scheme. Under these conditions, the governing equations were solved on the two-dimensional computational domain shown in Fig. C.1. The time step was set to 0.2 s, and the computational mesh was generated with a spatial resolution of $\Delta x = 0.025$ m.

C.2 NONLINEAR HYPERELASTIC DYNAMICS

The structural dynamics in the reference configuration are described by the kinematic relation and the balance of linear momentum:

$$\begin{aligned} \frac{\partial \mathbf{h}}{\partial t} - \mathbf{v} &= \mathbf{0}, \\ \rho \frac{\partial \mathbf{v}}{\partial t} - \nabla \cdot \mathbf{P} &= \mathbf{0}. \end{aligned} \tag{C.4}$$

Here, $\mathbf{h}$ and $\mathbf{v}$ denote the displacement and velocity vectors, respectively, ρ is the mass density, ∇ denotes the differential operator with respect to the reference coordinates $\mathbf{X}$, and $\mathbf{P}$ is the first Piola-Kirchhoff stress tensor.

To account for nonlinear material behavior, a compressible Neo-Hookean constitutive model is adopted. The deformation gradient, Jacobian, strain-energy density function, first invariant of the

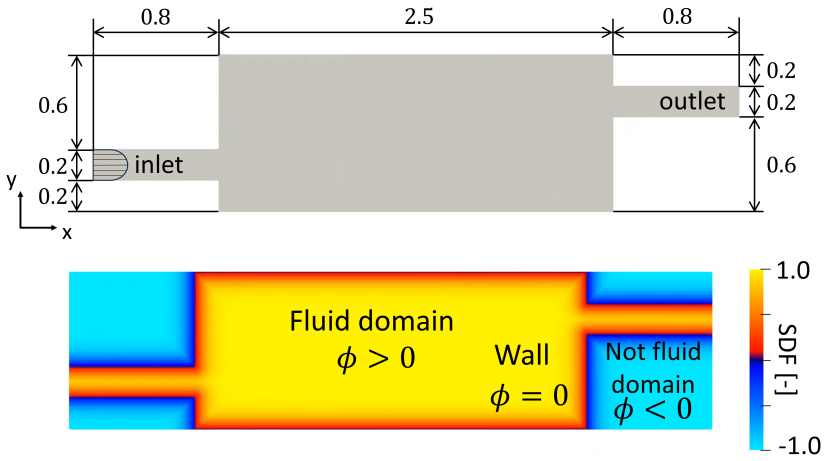

Figure C.1: Computational domain and boundary conditions for the coupled flow–transport problem (Upper). The signed distance feature supplied to the decoder (Lower).

right Cauchy-Green tensor, and first Piola-Kirchhoff stress tensor are given by

$$\begin{aligned} \mathbf{F} &= \mathbf{I} + \nabla \mathbf{h}, \\ J &= \det(\mathbf{F}), \\ \Psi &= \frac{\mu}{2}(I_C - d) - \mu \ln J + \frac{\lambda}{2}(\ln J)^2, \\ I_C &= \mathrm{tr}(\mathbf{F}^T \mathbf{F}), \\ \mathbf{P} &= \mu \left(\mathbf{F} - \mathbf{F}^{-T}\right) + \lambda \ln J \, \mathbf{F}^{-T}. \end{aligned} \tag{C.5}$$

Here, $\mathbf{F}$ denotes the deformation gradient tensor, $J = \det(\mathbf{F})$ represents the local volume ratio, Ψ is the strain-energy density function, I_C is the first invariant of the right Cauchy-Green deformation tensor, d denotes the spatial dimension, and μ and λ are the Lamé parameters.

The Lamé parameters are expressed in terms of Young's modulus E and Poisson's ratio ν as

$$\mu = \frac{E}{2(1+\nu)}, \qquad \lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}. \tag{C.6}$$

The left boundary is fixed, while a time-dependent traction is applied to the right boundary, and the remaining boundaries are traction-free. The loading frequency is defined relative to the first natural frequency f_1 as

$$f_{\text{load}} = \omega f_1. \tag{C.7}$$

where ω denotes the dimensionless loading-frequency ratio. Under these conditions, the governing equations were spatially discretized using the FEM with P2 elements for the displacement field. Time integration was performed using the implicit Newmark-β method with $\beta = 0.25$ and $\gamma = 0.5$. The resulting nonlinear system at each time step was solved using the Newton-Raphson method. A schematic illustration of the computational setup is shown in Fig C.2. The time step was set to 0.5×10^{-4} s, and the computational mesh was generated with a spatial resolution of $\Delta x = 0.01$ m.

C.3 Three-Dimensional Aneurysm Flow

The aneurysm flow is governed by the three-dimensional nondimensional incompressible Navier–Stokes equations:

$$\begin{aligned} \nabla \cdot \mathbf{v}^* &= 0, \\ \frac{\partial \mathbf{v}^*}{\partial t} + (\mathbf{v}^* \cdot \nabla)\, \mathbf{v}^* + \nabla p^* - \frac{1}{\mathrm{Re}} \nabla^2 \mathbf{v}^* &= \mathbf{0}. \end{aligned} \tag{C.8}$$

Here, $\mathbf{v}^*$ and p^* denote the nondimensional velocity and pressure, respectively, and Re is the Reynolds number. In addition to the Reynolds number, the Womersley number Wo was used as

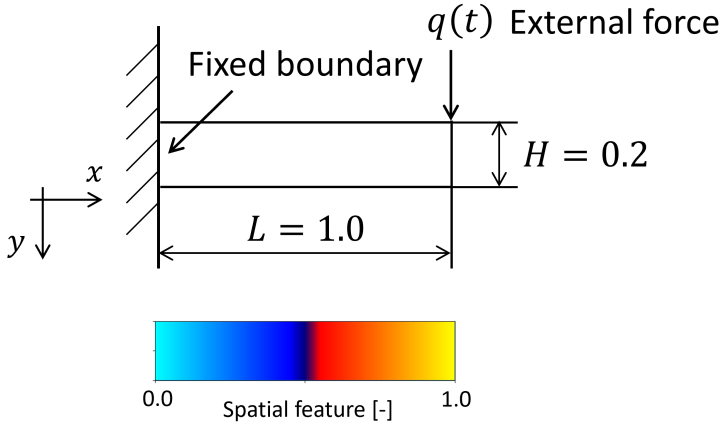

Figure C.2: Reference configuration and boundary conditions for the hyperelastic cantilever. The left boundary is fixed, and the time-dependent traction $q(t)$ is applied to the right boundary. The normalized coordinate x/L is used as a spatial input feature.

an input parameter. The Womersley number is a dimensionless quantity that characterizes the degree of unsteadiness in pulsatile flow and is defined as

$$\mathrm{Wo} = R_{\mathrm{ref}} \sqrt{\frac{\omega \rho}{\mu}}, \tag{C.9}$$

where R_{ref} denotes the characteristic vessel radius, ω is the angular frequency of the pulsatile flow, ρ is the fluid density, and μ is the dynamic viscosity. For the numerical simulations, a prescribed pulsatile velocity was imposed over the inlet cross-section as

$$U_{\mathrm{mean}} = A\left[\sin\left(2\pi f t\right) + 1\right], \tag{C.10}$$

where f denotes the pulsation frequency and A denotes the velocity amplitude. The values of A and f were adjusted according to the prescribed Re and Wo, respectively. A no-slip boundary condition was imposed on the vessel walls, while a constant pressure of $0\,\mathrm{Pa}$ was prescribed at all outlets. Under these boundary conditions, the governing equations were spatially discretized using stabilized FEM with P1/P1 element method for the velocity and pressure fields. The velocity and pressure were solved simultaneously using a monolithic formulation, and time integration was performed with a fully implicit backward-Euler method. The computational domain and inlet and outlet positions are shown in Fig. C.3. The aneurysm surface geometry was obtained from the Vascular Model Repository (Wilson et al., 2013), and the corresponding tetrahedral mesh was generated using Gmsh (Geuzaine & Remacle, 2009). The resulting mesh consisted of 50,781 nodes and 255,676 tetrahedral elements.

D Training details and Model comparison

For GeoPI-LNODE, the model architecture and training hyperparameters used for each benchmark are summarized in Table D.1.

For the two-dimensional problems, we compared GeoPI-LNODE with Baseline LNODE, DeepONet, FNO, and CNF-ROM. For three-dimensional aneurysm flow, FNO was replaced by MeshGraphNet, while Baseline LNODE, DeepONet, and CNF-ROM were retained. DeepONet, FNO, MeshGraphNet, and CNF-ROM were implemented following their original studies (Lu et al., 2021; Li et al., 2021; Pfaff et al., 2021; Kim et al., 2024). Their main hyperparameters are summarized in Table D.2.

All models used in this study were trained on a single NVIDIA GeForce RTX 4090 GPU.

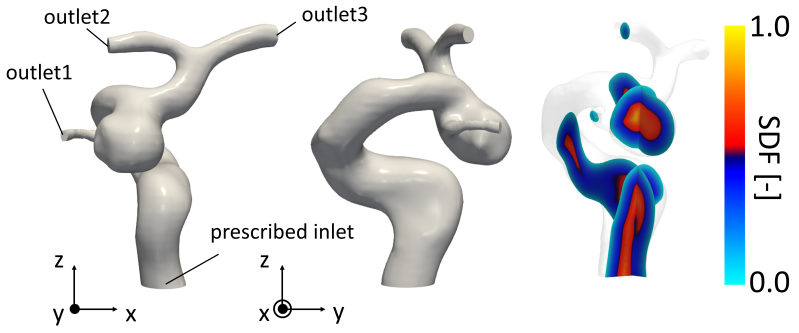

Figure C.3: Three-dimensional aneurysm geometry, computational mesh, and prescribed inlet, outlet, and wall boundaries (left). The spatial distribution of signed distance function (SDF) of the aneurysm geometry (right).

Table D.1: Model architecture and training hyperparameters used in the three benchmark problems.

Hyperparameter	Flow–transport	Hyperelastic dynamics	Aneurysm flow
Latent dimension	32	32	128
Latent ODE hidden dimension	128	128	128
Latent ODE layers	4	3	5
Reconstructor hidden dimension	128	128	128
Reconstructor layers	4	3	5
Latent ODE activation	Tanh	SiLU	SiLU
Reconstructor activation	SiLU	SiLU	SiLU
ODE solver	Euler	Euler	Euler
Training time steps	200	100	50
Rollout/test time steps	400	400	500
Batch size	1	8	4
Optimizer	Adam	Adam	Adam
Weight decay	0.0	0.0	0.0
Initial learning rate	0.001	0.001	0.001
Training epochs	20,000	10,000	20,000
Learning-rate scheduler	ReduceLROnPlateau	ReduceLROnPlateau	ReduceLROnPlateau

E Dataset Details

E.1 Flow and transport

A total of 18 datasets were generated, of which 80% were used for training and 20% for validation. The Reynolds number was varied from 100 to 500 in increments of 50, while the Schmidt number was varied from 0.1 to 1.0. For testing, 15 additional datasets were used, with Reynolds numbers ranging from 50 to 550 in increments of 50 and Schmidt numbers of 0.1, 0.5, and 1.0. The distribution of the training, validation, and test datasets is illustrated in Fig. E.1. Representative snapshots of the spatiotemporal distributions for the most diffusion-dominated and convection-dominated cases are shown in Fig. E.2.

E.2 Nonlinear Hyperelastic Dynamics

A total of 27 datasets were generated for model development, of which 80% were used for training and 20% for validation. The datasets were constructed by varying three physical parameters: the loading frequency ratio ω, the nondimensional loading amplitude A^*, and Young's modulus E. Specifically, the training and validation datasets were generated using $\omega \in \{0.05, 0.20, 0.40\}$, $A^* \in \{0.40, 0.85, 1.30\}$, and $E \in \{1.5 \times 10^7, 2.0 \times 10^7, 2.5 \times 10^7\}$ Pa. All combinations

Table D.2: Architecture and training hyperparameters used for the comparison models.

Model	Hyperparameter	Flow–transport	Hyperelastic beam	Aneurysm flow
DeepONet	Branch hidden dimension	128	256	256
	Branch hidden layers	4	3	4
	Branch activation	SiLU	GELU	SiLU
	Trunk hidden dimension	128	256	256
	Trunk hidden layers	4	3	4
	Trunk activation	SiLU	Tanh	SiLU
	Basis dimension	32	128	128
	Batch size	1	8	4
	Initial learning rate	0.001	0.001	0.001
	Training epochs	20,000	10,000	10,000
FNO	Fourier layers	4	4	–
	Fourier modes	12×20	12×16	–
	Hidden width	16	32	–
	Activation function	GELU	GELU	–
	Batch size	2	2	–
	Initial learning rate	0.001	0.001	–
	Training epochs	20,000	10,000	–
MeshGraphNet	Hidden dimension	–	–	96
	Message-passing layers	–	–	4
	MLP layers	–	–	3
	Activation function	–	–	SiLU
	Batch size	–	–	1
	Initial learning rate	–	–	0.001
	Training epochs	–	–	10,000
CNF-ROM	Latent dimension	32	32	128
	Latent ODE hidden dimension	128	128	128
	Latent ODE hidden layers	4	3	5
	Latent ODE activation	SiLU	SiLU	SiLU
	Neural-field hidden dimension	128	128	128
	Fourier modulation blocks	4	4	4
	Fourier frequency scale	64	64	64
	Trainable Fourier frequencies	No	NO	Trainable
	Batch size	1	8	4
	Initial learning rate	0.001	0.001	0.001
	Training epochs	20,000	10,000	20,000

of these parameter values yielded 27 cases. The distribution of the datasets in the parameter space is shown in Fig. E.3. For the interpolation test set, the parameter values were set to $\omega \in \{0.075, 0.15, 0.25, 0.35\}$, $A^* \in \{0.55, 0.85, 1.15\}$, and $E \in \{1.75 \times 10^7, 2.25 \times 10^7\}$ Pa. All combinations of these values yielded a total of 24 test cases. The spatial distribution of the deformation for a representative case is shown in Fig. E.4.

E.3 Three-Dimensional Aneurysm Flow

A total of 25 datasets were generated for model development, of which 80% were used for training and 20% for validation. The Reynolds number was varied from 100 to 200 in increments of 25, while the Womersley number was varied over $\{17.72, 19.82, 21.71, 23.45, 25.07\}$. All combinations of these parameter values were used, resulting in 25 cases. For testing, eight additional datasets were generated using parameter combinations not included in the training set. The corresponding $(\mathrm{Re}, \mathrm{Wo})$ pairs were $(112.5, 18.80)$, $(112.5, 24.26)$, $(137.5, 20.76)$, $(137.5, 22.58)$, $(162.5, 20.76)$, $(162.5, 22.58)$, $(187.5, 18.80)$, and $(187.5, 24.26)$. The distribution of the dataset is shown in Fig. E.5. The streamlines for a representative case are shown in Fig. E.6.

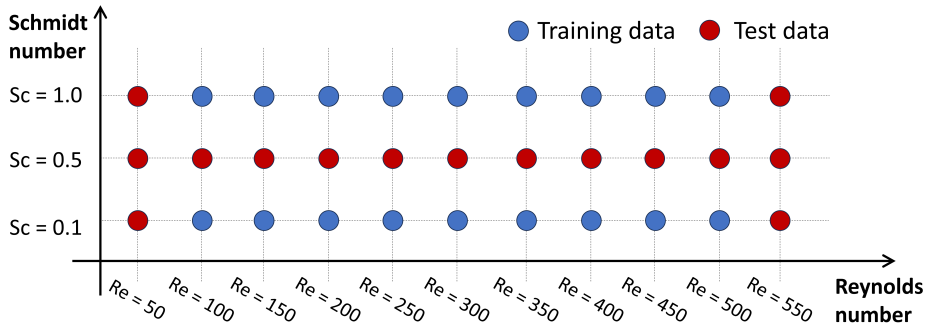

Figure E.1: The distribution of flow and transport dataset generated by numerical simulation.

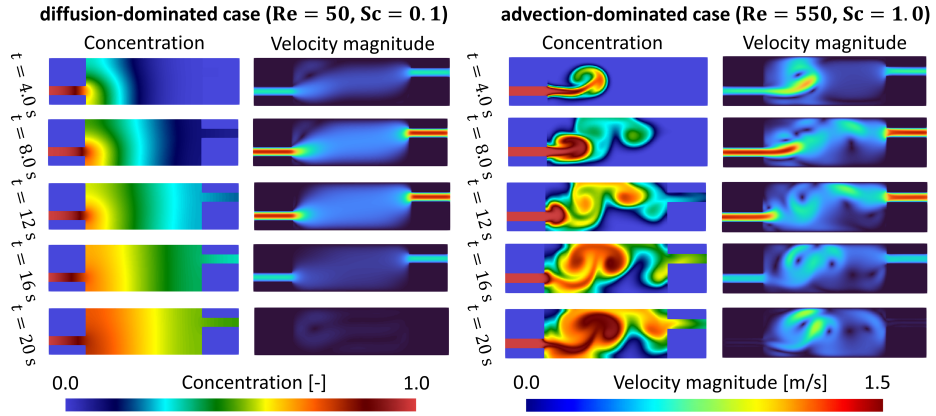

Figure E.2: Snapshots of the concentration and velocity-magnitude fields for the most diffusion-dominated case $\mathrm{Re} = 50$, $\mathrm{Sc} = 0.1$ and the most advection-dominated case $\mathrm{Re} = 550$, $\mathrm{Sc} = 1.0$.

F Additional Experimental Results

The component-wise $\mathrm{RMSE}_{\mathrm{ext}}$ values for the velocity field in the flow–transport problem, the displacement and velocity fields in the hyperelastic beam problem, and the velocity field in the aneurysm flow problem are reported in Tables F.1, F.2, and F.3, respectively.

The rollout RMSEs for velocity magnitude and concentration in the flow–transport problem are shown in Fig. F.1.

Snapshots of the spatial distributions of velocity magnitude and concentration for all comparison models and the ground truth are shown in Fig. F.2.

Snapshots of the hyperelastic beam deformation for the ground truth, GeoPI-LNODE, and Baseline LNODE are shown in Fig. F.3.

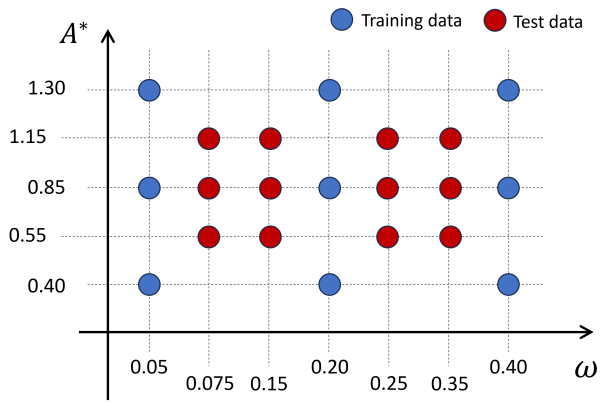

Figure E.3: Distribution of the nonlinear hyperelastic dataset generated by numerical simulations. A^* and ω denote the nondimensional amplitude and frequency ratio of the external force, respectively. The same parameter distribution was generated for two Young's moduli, $E = 1.75 \times 10^7$ and 2.25×10^7 Pa.

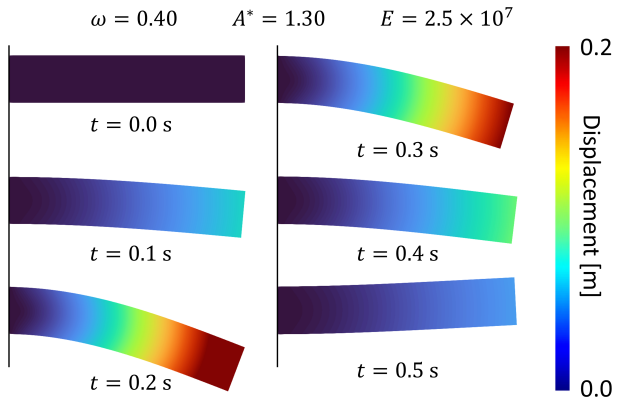

Figure E.4: Spatial distribution of the deformation of the hyperelastic beam over one loading cycle. The frequency ratio is $\omega = 0.40$, the nondimensional amplitude is $A^* = 1.30$, and Young's modulus is $E = 2.5 \times 10^7$ Pa.

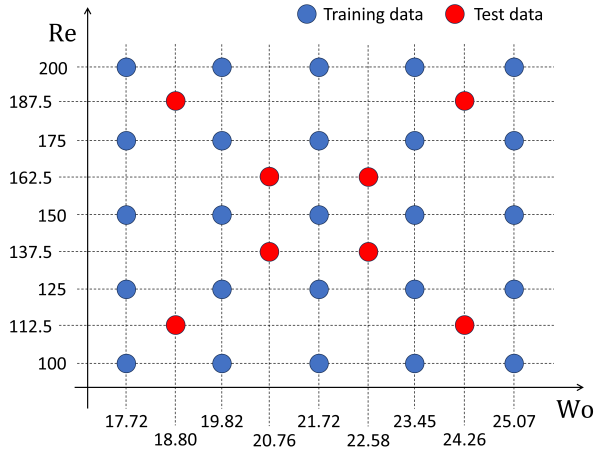

Figure E.5: Distribution of the aneurysm flow dataset generated by numerical simulations. Re and Wo denote the Reynolds number and Womersley number, respectively.

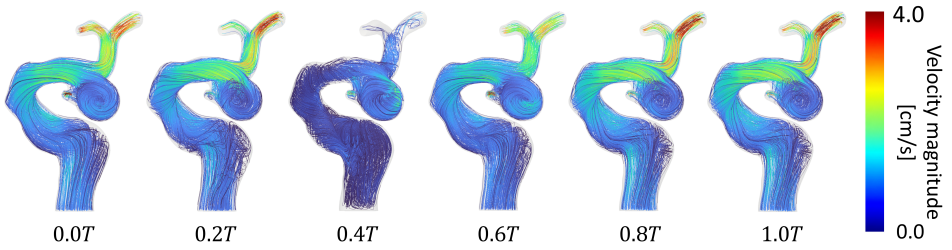

Figure E.6: Streamlines for the representative aneurysm flow case (Re = 200, Wo = 25.07) over one cardiac cycle. T denotes the duration of one cardiac cycle.

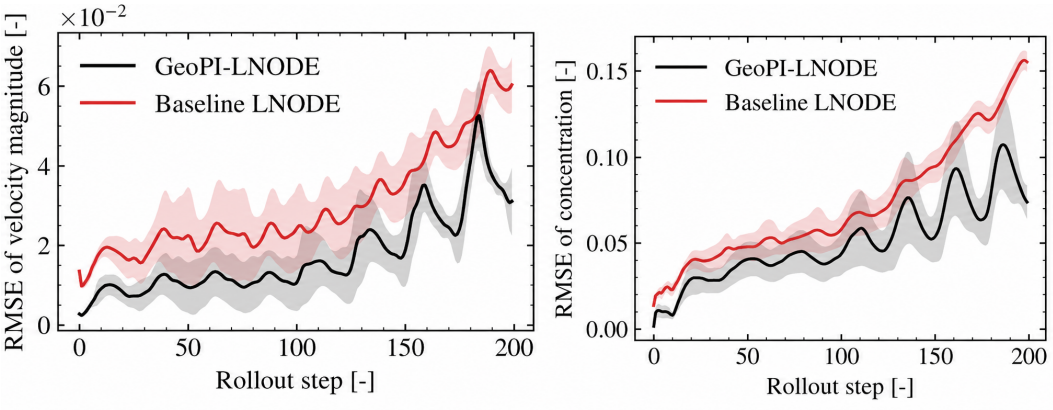

Figure F.1: Rollout RMSEs for the flow–transport problem: velocity magnitude (left) and concentration (right).

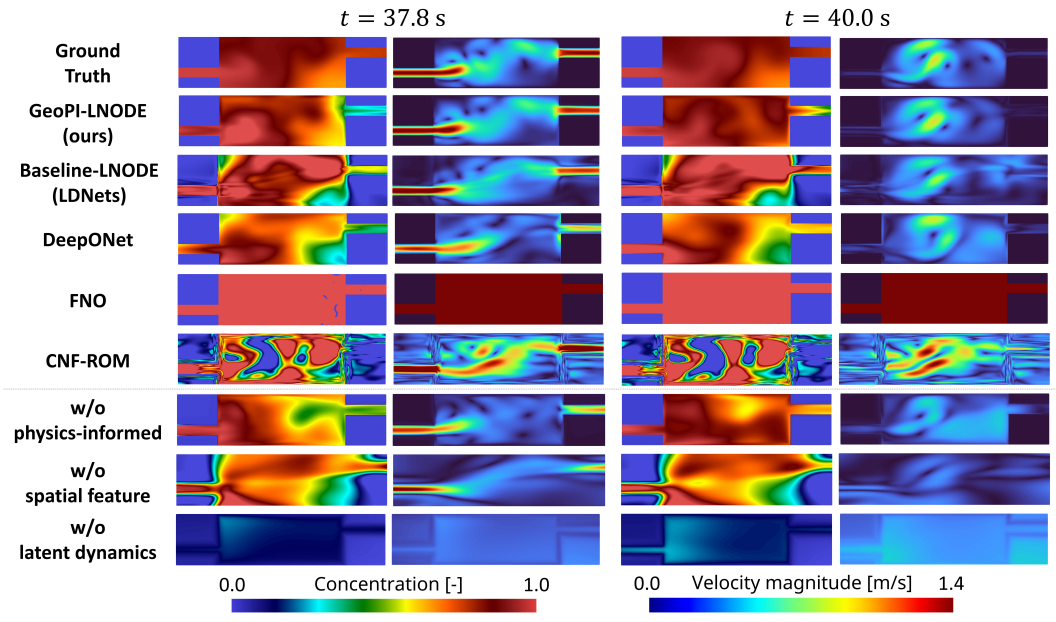

Figure F.2: Snapshots of the ground-truth and predicted velocity-magnitude and concentration fields for the flow–transport problem at $t = 37.8$ s and $t = 40.0$ s during the temporal extraporation interval. The rows correspond to the ground truth, GeoPI-LNODE, four comparison methods, and three ablation variants; the dashed horizontal line separates the comparison methods from the ablations. Identical color limits are used across all methods and times for each physical field.

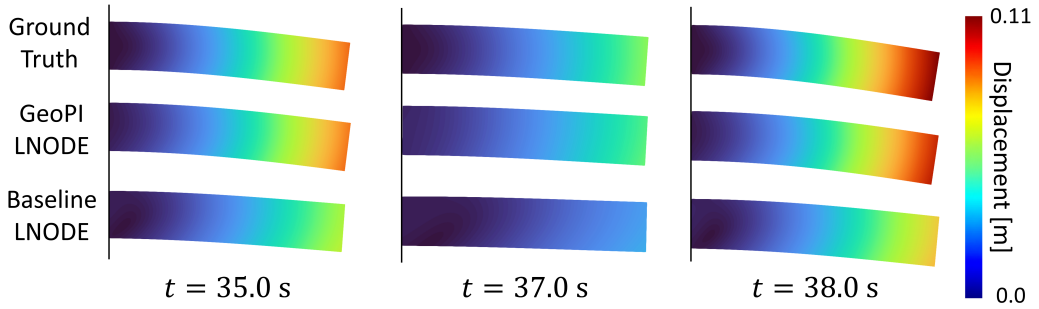

Figure F.3: Snapshots of the nonlinear hyperelastic deformation for the ground truth, GeoPI-LNODE, and Baseline LNODE at $t = 35.0$, 37.0, and 38.0 s.

Table F.1: Component-wise velocity RMSE for the flow–transport problem over t_{test}. Values are reported as the mean $\pm$ standard deviation across test cases.

Method	v_x	v_y
Baseline LNODE	0.0581 ± 0.00901	0.0422 ± 0.00759
DeepONet	0.109 ± 0.0232	0.0780 ± 0.0155
FNO	Diverged	
CNF-ROM (Kim et al., 2024)	0.0943 ± 0.0461	0.0692 ± 0.0321
w/o spatial feature	0.112 ± 0.0117	0.0550 ± 0.0167
w/o latent dynamics	0.322 ± 0.0138	0.179 ± 0.0105
w/o physics-informed	0.0809 ± 0.0157	0.0324 ± 0.0151
GeoPI-LNODE	$\mathbf{0.0385 \pm 0.00659}$	$\mathbf{0.0267 \pm 0.00903}$

Notation: v_x and v_y denote the velocity components. "Diverged" indicates that RMSE_{ext} exceeded 10^4.

Table F.2: Component-wise displacement and velocity RMSEs for the hyperelastic beam problem over t_{test}. Values are reported as the mean $\pm$ standard deviation across test cases.

	Displacement		Velocity	
Method	h_x	h_y	v_x	v_y
Baseline LNODE	0.00166 ± 0.00115	0.00625 ± 0.00523	0.0181 ± 0.0142	0.115 ± 0.0937
DeepONet	0.0589 ± 0.000944	0.166 ± 0.00841	1.31 ± 0.0113	1.22 ± 0.0210
FNO	0.536 ± 0.0777	2.33 ± 0.569	10.5 ± 1.35	75.8 ± 12.3
CNF-ROM (Kim et al., 2024)	Diverged			
w/o spatial feature	0.00158 ± 0.000833	0.00964 ± 0.00306	0.0347 ± 0.0199	0.131 ± 0.0897
w/o latent dynamics	0.00151 ± 0.000197	0.00581 ± 0.00112	0.00319 ± 0.0203	0.182 ± 0.122
w/o physics-informed	0.00894 ± 0.0316	0.0435 ± 0.147	0.118 ± 0.357	0.693 ± 1.97
GeoPI-LNODE	$\mathbf{0.00106 \pm 0.000586}$	$\mathbf{0.00454 \pm 0.00318}$	$\mathbf{0.00195 \pm 0.0173}$	$\mathbf{0.0991 \pm 0.0740}$

Notation: h_x and h_y denote the displacement components, while v_x and v_y denote the velocity components. "Diverged" indicates that RMSE_{ext} exceeded 10^4.

Table F.3: Component-wise velocity RMSEs for the aneurysm flow problem over t_{test}. Values are reported as the mean $\pm$ standard deviation across test cases.

Method	v_x	v_y	v_z
Baseline LNODE	$(1.155 \pm 3.476) \times 10^3$	$(1.138 \pm 3.409) \times 10^3$	$(2.956 \pm 8.903) \times 10^2$
DeepONet		Diverged	
MeshGraphNet	0.4781 ± 0.2666	0.5958 ± 0.7006	0.3584 ± 0.1432
CNF-ROM (Kim et al., 2024)	0.317 ± 0.1385	0.2928 ± 1.211	0.313 ± 0.1408
w/o spatial feature	0.346 ± 0.144	0.309 ± 0.1271	0.3199 ± 0.1454
w/o latent dynamics	$(1.573 \pm 1.231) \times 10^1$	8.204 ± 8.501	2.378 ± 1.291
w/o physics-informed	0.3245 ± 0.1320	0.2870 ± 0.1087	0.3224 ± 0.1421
GeoPI-LNODE	$\mathbf{0.2915 \pm 0.1251}$	$\mathbf{0.2661 \pm 0.1077}$	$\mathbf{0.2724 \pm 0.1069}$

Notation: v_x, v_y, and v_z denote the velocity components. "Diverged" indicates that RMSE_{ext} exceeded 10^4.