



ARTICLE

Adaptive Physics-Informed Neural Networks with Overlapping Domain Decomposition for Vascular Flow Velocity Modeling

Muzhou Hou^{1,2,*}, Dezhou Lin², Kaituo Xu² and Wen Hou³

¹School of Computer Science and Engineering, Hunan University of Information Technology, Changsha, China

²School of Mathematics and Statistics, Central South University, Changsha, China

³School of Traffic Transportation of Engineering, Changsha University of Science & Technology, Changsha, China

*Corresponding Author: Muzhou Hou. Email: houlmuzhou@sina.com

Received: 07 August 2026; Accepted: 10 September 2026; Published: 21 October 2026

ABSTRACT: Hemodynamic information is crucial for studying cardiovascular pathogenesis and risk assessment. When measurement data are sparse, boundary data are incomplete, and branching geometries are involved, classical computational fluid dynamics (CFD) remains computationally intensive, while standard physics-informed neural networks (PINNs) often suffer from training instability due to poor gradient balance among competing loss terms and inadequate resolution of high-frequency flow features near bifurcations. This paper proposes three innovations to address these challenges. First, a residual-gradient-based adaptive weighting scheme is developed to mitigate gradient competition between boundary terms and data-driven terms, embedded into a meshless PINN for three-dimensional incompressible Navier-Stokes flow. Second, an overlapping subdomain decomposition with zero- to second-order interface conditions is constructed to enhance physical consistency between subdomains and suppress non-physical oscillations near bifurcations. Third, the same strategy is extended to one-dimensional pulse wave propagation, bridging local three-dimensional reconstruction and network-scale analysis. Numerical experiments on porcine abdominal aorta geometries demonstrate that the adaptive weighting yields more uniform gradient contributions and stable optimization paths; the overlapping decomposition improves reconstructed flow and accelerates convergence; and the one-dimensional formulation confirms that partial differential equation (PDE), boundary, and interface terms can be jointly minimized, providing empirical support for multiscale hemodynamic modeling.

KEYWORDS: Physics-informed neural networks; hemodynamics; Navier-Stokes equations; adaptive weighting; overlapping domain decomposition

1 Introduction

Cardiovascular disease remains one of the leading causes of mortality and morbidity worldwide [1], and its prevalence continues to grow with population aging, lifestyle changes, and the rising incidence of chronic metabolic diseases [2]. Clinical research has consistently concluded that vascular pathology cannot be explained solely by abnormal anatomy: how blood flows through diseased segments is critically important. At bifurcations, curves, or stenoses, blood flow deviates from simple patterns, velocity and pressure gradients surge over short arcs, and secondary motions recur. Hemodynamic indices such as wall shear stress (WSS), oscillatory shear index (OSI), and relative residence time (RRT) deviate from their physiological ranges [3], and clinicians increasingly rely on these signals when assessing atherosclerotic risk, aneurysm evolution, and endothelial damage.

Methodologically, hemodynamics research has developed along two main routes: image-based measurement and CFD-based simulation [4]. Image measurement is attractive at the bedside because it is minimally invasive, but its products are limited by voxel size, noise, cost, and reconstruction pipeline robustness. When digital anatomy, meshes, and boundary data are all well-conditioned, CFD can reproduce detailed velocity and pressure fields, but this dependency, coupled with high computational cost, makes classical workflows unsuitable for rapid screening of heavily bifurcated trees, for scenarios with only a few probes, or for analyses that must be tailored and refreshed per patient.

Scientific machine learning has recently introduced physics-informed neural networks (PINNs) as a method that folds mechanistic laws and observational data into the same loss landscape [5,6]. During training, PDE residuals are explicitly penalized alongside initial, boundary, and data terms, so the network seeks a single continuous field that respects every ingredient [7]. Compared with black-box regressors, PINNs retain more physical meaning when labels are scarce; compared with mesh-based discretizations, they avoid explicit meshing and enable more direct inversion [8]. Applying this idea to bifurcating vascular flow can therefore strengthen reconstruction under sparse measurements [9] and provide a computationally lighter scaffold for estimating hemodynamic biomarkers [10].

Despite the promise, several challenges persist. Early efforts by Psychogios and Ungar [10] and Lagaris et al. [11] combined neural networks with governing equations through automatic differentiation. Raissi et al. [5–7] formalized the PINN framework and demonstrated its potential in hidden fluid mechanics applications [12]. Jagtap et al. [13] introduced locally adaptive activation functions to improve convergence. However, as applications grew more demanding, it became clear that standard PINNs suffer from gradient competition among multiple loss terms, slow convergence, and insufficient resolution of local high-frequency features [14], and eigenvector bias in Fourier feature networks [15]. Lu et al. [16] employed Monte Carlo residual estimation for adaptive point placement, and Mao et al. [17] combined adaptive sampling with operator approximation. Kharazmi et al. [18] proposed variational PINNs with domain decomposition, while Yu et al. [19] developed gradient-enhanced PINNs. Multi-fidelity approaches [20,21] further advanced the field. Shin et al. [22] analyzed convergence for elliptic and parabolic problems, and Wang et al. [23] automated PINN topology through architecture search. Reynolds stress modeling with physics-informed machine learning [24] and cardiovascular flow prediction from 4D flow MRI [25] have also been explored. Recent studies have further extended PINNs to inverse problems in cardiovascular fluid mechanics, gradient-enhanced cardiovascular flow modeling with sparse data, and scalable hemodynamic simulations on GPU clusters.

Motivated by these gaps, this paper makes three main contributions. (1) We develop a residual-gradient-based adaptive weighting scheme that dynamically balances the gradient contributions of different loss terms, mitigating the gradient pathology that plagues standard PINNs in complex vascular flow. (2) We construct an overlapping subdomain decomposition with zero- to second-order interface conditions that enforces physical consistency across subdomains, suppressing non-physical oscillations near bifurcations. (3) We extend the same strategy to one-dimensional pulse wave propagation, providing a coherent pathway from local three-dimensional reconstruction to global network-scale analysis.

2 Mathematical Formulation

2.1 Governing Equations: 3D Incompressible Navier-Stokes

Under the continuum, incompressibility, and Newtonian fluid assumptions, blood flow in large arteries is governed by the three-dimensional incompressible Navier-Stokes equations. The continuity equation (mass conservation) requires that the velocity field be divergence-free, as given by Eq. (1):

$$\nabla \cdot \mathbf{u} = 0 \quad (1)$$

where $\mathbf{u} = (u, v, w)$ is the velocity vector. The momentum conservation equation expresses Newton's second law for fluid elements, as shown in Eq. (2):

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u} \quad (2)$$

where ρ is the fluid density ρ , p is the pressure, and μ is the dynamic viscosity μ . The nonlinear convective term $(\mathbf{u} \cdot \nabla) \mathbf{u}$ is the primary source of flow complexity, including turbulence and vortex formation, and is a major cause of numerical difficulty.

2.2 One-Dimensional Hemodynamic Model

To analyze pulse wave propagation across vascular networks at reduced computational cost, a one-dimensional model is derived by integrating the 3D equations over the vessel cross-section under the long-wavelength approximation [25]. The resulting hyperbolic system couples cross-sectional area A , volumetric flow Q , and average pressure p , as expressed in Eq. (3):

$$\frac{\partial A}{\partial t} + \frac{\partial Q}{\partial x} = 0, \quad \frac{\partial Q}{\partial t} + \frac{\partial}{\partial x} \left(\frac{Q^2}{A} \right) + \frac{A}{\rho} \frac{\partial p}{\partial x} = -K_f \frac{Q}{A} \quad (3)$$

where x is the axial coordinate, t is time, and f is the viscous friction coefficient approximated via the Darcy-Weisbach relation. A constitutive equation relating A and p , derived from Laplace's law for an isotropic elastic tube, closes the system.

2.3 PINN Framework and Composite Loss

A fully connected neural network with L hidden layers serves as the function approximator. For each layer l , the affine transformation ℓ and activation are defined by Eq. (4):

$$\mathbf{z}^{(l)} = \mathbf{W}^{(l)} \mathbf{a}^{(l-1)} + \mathbf{b}^{(l)}, \quad \mathbf{a}^{(l)} = \sigma(\mathbf{z}^{(l)}) \quad (4)$$

where $\mathbf{W}^{(l)}$ is the weight matrix, $\mathbf{b}^{(l)}$ is the bias vector, and $\sigma(\cdot)$ is a nonlinear activation function σ (Tanh in this work). The network takes spatiotemporal coordinates (x, y, z, t) as input and outputs velocity components and pressure. PINN training minimizes a composite loss function comprising data, PDE residual, and boundary condition terms, as given by Eq. (5):

Specifically, the data loss L_{data} measures the discrepancy between network predictions and available observational data (e.g., sparse velocity measurements) at data points $\{x_i^{\text{data}}\}$, defined as

$$L_{\text{data}} = \frac{1}{N_d} \sum_{i=1}^{N_d} \left\| \mathbf{u}_{\theta}(x_i^{\text{data}}) - \mathbf{u}_i^{\text{obs}} \right\|^2,$$

where $\mathbf{u}_{\theta}$ and $\mathbf{u}_i^{\text{obs}}$ denote the network-predicted and observed velocity vectors, respectively. The boundary condition loss L_{bc} enforces Dirichlet (velocity) and Neumann (stress/pressure gradient) conditions at inlet, outlet, and wall boundaries, defined as:

$$L_{\text{bc}} = \frac{1}{N_b} \sum_{x \in \Gamma_D \cup \Gamma_N} \left\| B(\mathbf{u}_{\theta}) - g \right\|^2,$$

where Γ_D and Γ_N denote Dirichlet and Neumann boundaries, and N_b is the total number of boundary collocation points.

$$\mathcal{L} = \lambda_{\text{data}} \mathcal{L}_{\text{data}} + \lambda_{\text{pde}} \mathcal{L}_{\text{pde}} + \lambda_{\text{bc}} \mathcal{L}_{\text{bc}} \quad (5)$$

where the PDE residual loss enforces the governing equations at collocation points distributed throughout the domain, as defined in Eq. (6):

$$\mathcal{L}_{\text{pde}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \| \mathbf{f}_r(\mathbf{x}_i, \mathbf{t}_i) \|^2 \quad (6)$$

here $\mathbf{r}$ denotes the Navier-Stokes residual operator computed via automatic differentiation, and N_r is the number of residual collocation points. In standard PINNs, the weights λ_{data} , λ_{pde} , and λ_{bc} are fixed $\lambda_{\text{data}} \lambda_{\text{pde}} \lambda_{\text{bc}}$, which leads to gradient imbalance when the loss terms differ in magnitude by several orders.

3 Methodology

3.1 Residual-Gradient-Based Adaptive Weighting

In vascular flow, the velocity gradient near the wall can differ from that in the core flow by several orders of magnitude. When standard PINNs are applied to such problems, the PDE residual gradient often dominates the boundary condition gradient during early training, causing the optimizer to satisfy boundary requirements quickly while physical residuals stagnate. This phenomenon, known as gradient pathology, results in physically inconsistent predictions in the lumen despite accurate near-wall fitting.

To address this, we redefine the total loss with dynamically updated weights, as formulated in Eq. (7):

$$\mathcal{L}_{\text{total}} = \sum_{n=1}^N \hat{\lambda}_n \mathcal{L}_n \quad (7)$$

where $\hat{\lambda}_n$ is the adaptive weight for the n -th loss term $\mathcal{L}_n$, updated at each iteration based on gradient statistics. Specifically, we compute the gradient of each loss term with respect to the last hidden layer's weight parameters and derive a scaling factor that aligns gradient magnitudes, as shown in Eq. (8):

$$\hat{\lambda}_n = \frac{\max |\nabla_{\theta} \mathcal{L}_{\text{pde}}|}{\text{mean}(|\nabla_{\theta} \mathcal{L}_n|)} \quad (8)$$

The numerator uses the maximum absolute gradient of the PDE residual loss as a reference, while the denominator uses the mean gradient magnitude of the target loss term. This ensures that terms with small gradients receive larger weights, forcing the optimizer to allocate more capacity to under-fitted physical constraints.

Because different collocation point samples can cause the raw weights to oscillate, an exponential moving average (EMA) smoothing mechanism is introduced, as described by Eq. (9):

$$\hat{\lambda}_n^{(k)} = \alpha \hat{\lambda}_n^{(k-1)} + (1 - \alpha) \bar{\lambda}_n^{(k)} \quad (9)$$

where α is the momentum smoothing factor (set to 0.9 in this work), $\bar{\lambda}_n^{(k)}$ is the previous smoothed weight, and $\hat{\lambda}_n^{(k)}$ is the newly computed raw weight α . This mechanism suppresses gradient noise and prevents numerical instability when transitioning to the second-order L-BFGS optimizer. The overall improved PINN framework is illustrated in Fig. 1. The choice of $\alpha = 0.9$ was determined through a sensitivity study: values

of $\alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9, 0.95, 0.99\}$ were tested, and $\alpha = 0.9$ yielded the best trade-off α between responsiveness and stability. When $\alpha = 0.1$, the smoothed weights α track the raw gradient ratios too closely, resulting in high-frequency oscillations that destabilize the Adam–L-BFGS transition and increase the final L_{total} by approximately 40%. When $\alpha = 0.5$, the smoothing α is insufficient to suppress sampling-induced noise, and the convergence curve exhibits periodic spikes. Conversely, $\alpha \geq 0.95$ makes the weights $\alpha \geq$ too sluggish to adapt to rapidly changing gradient statistics during early training, delaying convergence by roughly 1500 iterations. Thus $\alpha = 0.9$ is selected α as the optimal default.




Figure 1: Improved PINN method: adaptive gradient-based weighting scheme with momentum smoothing.

3.2 Overlapping Domain Decomposition

Porcine abdominal aorta exhibits multi-bifurcation topology with secondary motions and steep gradients near the carina. A single global network often cannot simultaneously represent the relatively smooth trunk flow and the high-frequency non-linear features near bifurcations, leading to spectral bias and stress discontinuities at informal interfaces. To overcome this limitation, we propose an overlapping subdomain decomposition with strong physical coupling.

The total computational domain Ω is decomposed into a trunk subdomain Ω_1 and branch subdomains Ω_2, Ω_3 . Between adjacent subdomains, an overlapping buffer region of thickness δ is constructed, with $0.5D < \delta < 1.0D$ where D is the characteristic diameter. Within this overlap, three levels of interface constraints are enforced. The zero-order constraint ensures variable continuity, the first-order constraint enforces gradient (flux/stress) continuity, and the second-order constraint imposes residual consistency. The second-order residual consistency loss is defined as Eq. (10):

The complete composite loss for the overlapping domain decomposition is formulated as

$$L_{\text{total}} = \sum_k \left[\lambda_{\text{pde}}^{(k)} L_{\text{pde}}^{(k)} + \lambda_{\text{bc}}^{(k)} L_{\text{bc}}^{(k)} + \lambda_{\text{data}}^{(k)} L_{\text{data}}^{(k)} \right] + \lambda_0 L_0 + \lambda_1 L_1 + \lambda_2 L_2,$$

where the first sum runs over all subdomains k , and the last three terms are the interface coupling losses. The zero-order constraint loss enforces continuity of the primary variables (velocity and pressure) across the overlapping interface:

$$L_0 = \frac{1}{N_{\text{ov}}} \sum_{i \in \Omega_{\text{ov}}} [\| \mathbf{u}^{(k)}(\mathbf{x}_i, t_i) - \mathbf{u}^{(k+1)}(\mathbf{x}_i, t_i) \|^2 + |p^{(k)}(\mathbf{x}_i, t_i) - p^{(k+1)}(\mathbf{x}_i, t_i)|^2],$$

where N_{overlap} is the number of collocation points in the overlapping region N_{ov} , and $\mathbf{u}^{(k)}, p^{(k)}$ denote the velocity and pressure predicted by subnetwork k . The first-order constraint loss enforces continuity of gradients (fluxes and stresses) across the interface:

$$L_1 = \frac{1}{N_{\text{ov}}} \sum_{i \in \Omega_{\text{ov}}} [\| \nabla \mathbf{u}^{(k)} - \nabla \mathbf{u}^{(k+1)} \|^2 + \| \nabla p^{(k)} - \nabla p^{(k+1)} \|^2],$$

which ensures that the viscous stress tensor and pressure gradient remain continuous at the subdomain boundaries, suppressing spurious flux jumps. The second-order residual consistency loss L_2 is defined in Eq. (10). All interface weights $\lambda_0, \lambda_1, \lambda_2$ are dynamically updated via the adaptive weighting scheme described in Section 3.1.

$$\mathcal{L}_2 = \frac{1}{N_{\text{ov}}} \sum_{i=1}^{N_{\text{ov}}} \| f_{r,1}(\mathbf{x}_i) - f_{r,2}(\mathbf{x}_i) \|^2 \quad (10)$$

where $f_{r,1}$ and $f_{r,2}$ are the Navier-Stokes residuals computed by neighboring subnetworks at overlapping collocation points, and N_{ov} is the number of overlap points. This constraint forces the two subnetworks to share the same physical logic during convergence, fundamentally eliminating the field discontinuity caused by independent parameterization. The adaptive weighting mechanism from Section 3.1 is extended to multi-domain training by dynamically balancing interface and interior physics losses. The subdomain design is shown in Fig. 2.




Figure 2: Overlapping subdomain decomposition design for bifurcating vascular geometry.

3.3 Extension to 1D Pulse Wave Propagation

To connect local three-dimensional reconstruction with network-scale analysis, the overlapping decomposition and adaptive weighting are extended to one-dimensional pulse wave propagation in elastic vessels. The 1D governing equations (Eq. (3)) are recast as PINN residuals suitable for automatic differentiation. Spatial coordinate masks define 1D overlapping subdomains along the vessel tree, and interface losses analogous to those in Section 3.2 (variable continuity, flux smoothness, and residual consistency) are

applied at segment boundaries. The adaptive gradient weighting coordinates the optimization of PDE, boundary, and interface terms, ensuring that the wave front's high-frequency discontinuities receive adequate computational attention. This formulation provides a coherent bridge between fine-scale 3D bifurcation modeling and network-scale 1D wave propagation.

4 Numerical Experiments

4.1 Experimental Setup

A three-dimensional geometric model simulating the porcine abdominal aorta and its iliac bifurcation was constructed. The vessel centerline advances along the x -direction with a sinusoidal bend in y , and branch radii decrease linearly with flow distance. Key geometric parameters: radius $R = 0.5$, length $L = 5.0$, bending amplitude $A = 0.6$. Collocation points include interior points, wall points, inlet points, and outlet points. The 3D sampling strategy is shown in Fig. 3. The computational domain is fixed in a Lagrangian frame attached to the vessel centerline; although the vessel wall undergoes small radial displacements under pulsatile pressure, the wall motion is treated as a prescribed boundary condition rather than a free-surface deformation, and the domain geometry remains stationary throughout the simulation. The inlet is located at $x = 0$, the two outlets at $x = L$ with bifurcated branches, and the wall boundary enforces the no-slip condition. The fluid parameters are set to blood-like values: density $\rho = 1060 \text{ kg/m}^3$, dynamic viscosity $\mu = 3.5 \times 10^{-3} \text{ Pa}\cdot\text{s}$, and the flow is assumed incompressible and isothermal. The inlet boundary prescribes a pulsatile velocity waveform $u_{\text{in}}(t) = U_0 \left[1 + \cos\left(\frac{2\pi t}{T}\right) \right]$ with peak velocity $U_0 = 0.45 \text{ m/s}$ and period $T = 0.8 \text{ s}$, corresponding to a heart rate of 75 bpm. The two outlets impose a zero-traction (stress-free) condition $\frac{\partial u}{\partial n} = 0$ and $p = 0$ (reference pressure). The no-slip condition $u = 0$ is enforced at the vessel wall. Initial conditions at $t = 0$ are set to zero velocity and zero pressure throughout the domain, with the network trained from $t = 0$ to $t = 2T$ to allow transient effects to settle.




Figure 3: Three-dimensional sampling strategy for the curved vascular geometry.

A convergence analysis with respect to the number of collocation points was performed for the 3D model. The interior collocation points were varied across $N_r \in \{500, 1000, 2000, 4000, 8000\}$, with boundary points scaled proportionally ($N_{bc} \approx 0.3 N_r$). The final PDE residual L_{pde} decreased monotonically from $2.34e-04$ ($N_r = 500$) to $8.76e-07$ ($N_r = 8000$), following an approximate convergence rate of $O(N_r^{-0.82})$. The velocity MSE against the reference solution decreased from $1.52e-03$ to $4.31e-05$ over the same range. The standard configuration uses $N_r = 2000$, which provides a favorable balance between accuracy and computational cost; doubling to $N_r = 4000$ yields only a $1.7\times$ residual reduction at $2.1\times$ the training time. For the 1D model, each vessel segment uses a uniform allocation of 200 collocation points per unit length, with additional refinement (up to 400 points/unit length) near bifurcation junctions where gradients are steep.

The network architecture comprises 8 hidden layers with 256 neurons per layer, using the Tanh activation function. A two-stage optimization strategy is employed: Adam for global exploration with adaptive weights (5000 iterations), followed by Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) for high-precision refinement (5000 iterations). All experiments were conducted on a machine with Intel Core i7 CPU and NVIDIA GeForce GTX 1060 Max-Q GPU, using PyTorch 2.5.1 (CUDA 12.1). The Tanh activation function is selected for its smoothness and bounded output range $[-1, 1]$, which facilitates gradient-based optimization and avoids the unbounded activations that can cause exploding gradients in deep networks. While Tanh can suffer from vanishing gradients when inputs fall into the saturated regions ($|x| > 3$), this issue is mitigated in the present framework by: (i) input normalization to zero mean and unit variance, which keeps most activations in the linear regime; (ii) the adaptive weighting scheme, which prevents any single loss term from dominating and pushing weights into saturation; and (iii) the two-stage optimization (Adam followed by L-BFGS), where L-BFGS performs quasi-Newton steps that are less sensitive to mild gradient vanishing. Alternative activations (ReLU, SiLU, GELU) were tested in preliminary experiments; ReLU produced non-smooth pressure fields due to its kink, while SiLU and GELU offered comparable accuracy but slightly higher computational cost. Tanh is therefore retained as the default. All hyperparameters were selected through a systematic grid search over the following ranges: number of hidden layers $L \in \{4, 6, 8, 10\}$, neurons per layer $\in \{64, 128, 256, 512\}$, Adam learning rate $\in \{1e-4, 5e-4, 1e-3, 5e-3\}$, L-BFGS max iterations $\in \{3000, 5000, 8000\}$, and collocation point density $\in \{500, 1000, 2000, 4000\}$ interior points. The selected configuration (8 layers, 256 neurons, Adam lr = $1e-3$, 5000 Adam + 5000 L-BFGS iterations, 2000 interior points) achieved the lowest validation L_{total} on a held-out set of 500 points. The batch size is set to the full collocation set (full-batch training), and the weight initialization follows the Xavier uniform scheme. The exponential moving average momentum $\alpha = 0.9$ is selected as described in [Section 3.1](#).

The evaluation metrics used throughout [Section 4](#) are defined as follows. (i) The PDE residual loss L_{pde} measures the mean-squared violation of the Navier-Stokes equations at interior collocation points. (ii) The boundary condition loss L_{bc} measures the mean-squared error at inlet, outlet, and wall boundaries. (iii) The total loss $L_{total} = \lambda_{data}L_{data} + \lambda_{pde}L_{pde} + \lambda_{bc}L_{bc}$ aggregates all terms with adaptive weights. (iv) In addition to residual-based metrics, we report the mean squared error (MSE) between the network-predicted solution and a high-resolution reference solution obtained from a stabilized finite-element solver on a fine mesh, computed separately for velocity and pressure:

$$MSE_u = \frac{1}{N_{test}} \sum \| \mathbf{u}_{net} - \mathbf{u}_{ref} \|^2 \text{ and } MSE_p = \frac{1}{N_{test}} \sum |p_{net} - p_{ref}|^2,$$

where $N_{test} = 5000$ independently sampled test points not used during training. The gradient ratio metric, defined as the median ratio of the L_2 norm of the PDE-residual gradient to the boundary-condition gradient with respect to the last hidden layer weights, quantifies the degree of gradient balance.

4.2 Single-Domain Adaptive Weighting Results

Under the standard collocation setting, the gradient ratio median improved from 0.0671 (before weighting) to 0.9576 (after weighting), and the 10%–90% percentile range shifted from [0.0247, 0.1973] to [0.9198, 0.9882]. This demonstrates that the adaptive weighting effectively aligns gradient magnitudes across loss terms. The convergence curves comparing the improved and baseline PINNs are shown in Fig. 4. Specifically, Fig. 4 plots the total loss L_{total} (left y-axis, log scale) and the gradient ratio (right y-axis) against training iteration for both methods. The baseline PINN exhibits a steep initial descent followed by a plateau with oscillatory behavior after approximately 3000 iterations, while the adaptive PINN maintains a smooth, monotonic descent throughout the Adam phase and transitions seamlessly to the L-BFGS phase without loss spikes. The gradient ratio for the baseline fluctuates between 0.01 and 0.6, indicating persistent imbalance, whereas the adaptive method stabilizes the ratio near 1.0 within the first 500 iterations.




Figure 4: Convergence curve comparison between adaptive-weighted PINN and baseline PINN.

Table 1 compares the key metrics between the improved and baseline PINNs under the standard collocation setting. While the final error levels are similar, the improved method exhibits significantly smoother convergence, demonstrating enhanced gradient balancing capability.

Table 1: Comparison of key metrics between improved and baseline PINN under standard collocation.

Model	Records	Best L_{total}	Best Iter	End L_{total}	End L_{pde}	End L_{bc}
Baseline PINN	10,000	9.39e−11	4803	2.73e−9	6.40e−10	2.53e−9
Adaptive PINN	10,000	1.32e−10	5001	3.32e−9	7.65e−10	3.19e−9

Under extreme under-constrained conditions (boundary and interior points reduced to 120 and 900, respectively), the adaptive PINN maintained smoother descent compared to the baseline, whose convergence curve exhibited significant oscillations. This demonstrates the robustness of the adaptive weighting scheme when observational data are scarce.

4.3 Domain Decomposition Results

For the bifurcating geometry, three subnetworks were deployed with the overlapping decomposition. The proposed method achieved a minimum L_{total} of $1.82\text{e-}11$, compared to $1.41\text{e-}10$ for the baseline single-network approach, representing nearly an order-of-magnitude improvement. The gradient ratio (PDE/BC) stabilized at $0.88\text{--}1.03$ for the proposed method, vs. $0.0003\text{--}0.58$ for the baseline. Table 2 summarizes the loss and convergence stability comparison.

Table 2: Loss and convergence stability comparison for domain decomposition experiments.

Metric	Baseline	Proposed
Min L_{total}	$1.41\text{e-}10$	$1.82\text{e-}11$
Min L_{pde}	$4.64\text{e-}09$	$3.46\text{e-}09$
Convergence stability	Gradient imbalance causes oscillation	Rapid stabilization in Adam phase

In the carina bifurcation region, the proposed method reduced the PDE residual mean from $4.7386\text{e-}06$ (baseline) to $1.6729\text{e-}06$, and the maximum residual from $1.2319\text{e-}05$ to $7.7756\text{e-}06$. The overlapping decomposition successfully captured secondary flow features at the bifurcation ridge, while the baseline exhibited non-physical oscillations and discontinuities. Table 3 presents the regional metrics comparison.

Table 3: Regional metrics comparison between baseline and proposed methods at the bifurcation and global scale.

Region	Metric	Baseline	Proposed
Carina	PDE residual mean	$4.74\text{e-}06$	$1.67\text{e-}06$
Carina	PDE residual max	$1.23\text{e-}05$	$7.78\text{e-}06$
Carina	log10 (residual) std	0.18	0.23
Global	PDE residual mean	$8.32\text{e-}06$	$2.04\text{e-}06$
Global	PDE residual 99%ile	$5.73\text{e-}05$	$1.41\text{e-}05$
Global	log10 (residual) std	0.39	0.35

Regarding computational efficiency, the three-subnetwork decomposition increases the total trainable parameters by approximately a factor of 2.8 compared with a single global network of equivalent per-subnetwork capacity (8 layers $\times$ 256 neurons). On the NVIDIA GeForce GTX 1060 Max-Q GPU, a single training run of 10,000 iterations required approximately 42 min for the proposed three-network method vs. 18 min for the baseline single network, representing a $2.3\times$ increase in wall-clock time. However, the proposed method achieves nearly an order-of-magnitude lower residual ($L_{\text{total}} = 1.82\text{e-}11$ vs. $1.41\text{e-}10$), and the per-iteration overhead is largely offset by the faster convergence stabilization in the Adam phase. For the adaptive weighting scheme alone (Section 4.2), the computational overhead is negligible—less than 3% additional wall-clock time—because the gradient statistics are computed using already-available backward-pass gradients and the exponential moving average (EMA) update involves only scalar operations.

4.4 One-Dimensional Network Results

The 1D model was tested on an 11-segment tree-like bifurcating network. Under nondimensionalization and proper initialization, the PDE residual loss reached convergence levels comparable to the 3D case. The convergence curve, adaptive weight evolution, gradient balance, and residual spatial distribution formed a complete analysis chain. Fig. 5 shows the 1D convergence behavior, confirming that the improved framework

maintains tractable optimization behavior after dimensional reduction. The 11-segment tree-like network comprises one trunk vessel (segment 1, length $L_1 = 5.0$ cm, radius $R_1 = 0.5$ cm) that bifurcates into two daughter vessels (segments 2–3, $L = 3.5$ cm, $R = 0.35$ cm), each of which bifurcates again (segments 4–7, $SL = 2.5$ cm, $R = 0.25$ cm), followed by a third generation (segments 8–11, $L = 1.8$ cm, $R = 0.18$ cm). The vessel wall is modeled as an isotropic elastic tube with Young's modulus $E = 5.0 \times 10^5$ Pa, wall thickness $h = 0.1R$, and Poisson's ratio $\nu = 0.5$, yielding a pressure–area relation $p(A) = \frac{Eh}{R_0^2} (\sqrt{A} - \sqrt{A_0})$ via Laplace's law. The blood density is $\rho = 1060$ kg/m³ and dynamic viscosity $\mu = 3.5 \times 10^{-3}$ Pa·s. At the inlet (segment 1, $x = 0$), a pulsatile flow waveform $Q_{in}(t) = Q_0 [1 + 0.3 \sin(\frac{2\pi t}{T})]$ is prescribed with $Q_0 = 5.0$ cm³/s and $T = 0.8$ s. At all terminal outlets (segments 8–11), a three-element Windkessel boundary condition is applied with proximal resistance $R_1 = 1.2 \times 10^8$ Pa·s/m³, distal resistance $R_2 = 4.8 \times 10^8$ Pa·s/m³, and compliance $C = 2.5 \times 10^{-10}$ m³/Pa. At each bifurcation, continuity of pressure ($p_{parent} = p_{daughter1} = p_{daughter2}$) and conservation of flow ($Q_{parent} = Q_{daughter1} + Q_{daughter2}$) are enforced. Initial conditions set $Q(x, 0) = 0$ and $A(x, 0) = A_0$ (equilibrium cross-sectional area) throughout the network. Each segment uses 200 collocation points per unit length, with 400 points/unit length near bifurcation interfaces.




Figure 5: Convergence curve for the one-dimensional pulse wave propagation model.

The 1D results confirm that the PDE, boundary, and interface terms can be jointly minimized within a single optimization objective, providing empirical support for linking local three-dimensional details with global one-dimensional analysis. The adaptive weighting forces the optimizer to focus on high-frequency discontinuous regions near the pulse wave front, ensuring coordinated convergence of the governing equations and topological connection conditions.

5 Key Findings and Limitations

The results demonstrate that gradient pathology is a fundamental bottleneck in applying PINNs to complex vascular flow. The residual-gradient-based adaptive weighting addresses this by dynamically aligning

gradient magnitudes, which is particularly effective when loss terms differ by several orders of magnitude. The momentum smoothing factor ($\alpha = 0.9$) provides a balance between responsiveness to current sampling and suppression of gradient noise, ensuring stable transitions between Adam and L-BFGS optimizers.

The overlapping domain decomposition with three-level interface constraints represents a significant advance over traditional non-overlapping approaches. By enforcing not only variable continuity but also gradient and residual consistency, the method ensures that neighboring subnetworks share the same physical logic. The second-order residual consistency constraint is particularly important: it forces both subnetworks to violate the governing equations to the same degree in the overlap, fundamentally eliminating the field discontinuities caused by independent parameterization.

The extension to 1D pulse wave propagation demonstrates the generality of the approach. While the 3D model captures fine-scale flow details, the 1D model enables network-scale analysis at minimal computational cost. The consistent use of overlapping decomposition and adaptive weighting across both scales provides a coherent multiscale modeling framework. However, several limitations remain: the geometric and physiological parameters are based on porcine abdominal aorta, the 3D-to-1D connection is methodological rather than quantitatively validated, and the adaptive weights still rely on empirical parameter choices.

6 Conclusion

This paper proposed three innovations for vascular flow velocity modeling using improved physics-informed neural networks. First, a residual-gradient-based adaptive weighting scheme was developed to mitigate gradient competition among multiple loss terms in 3D incompressible Navier-Stokes flow. Second, an overlapping subdomain decomposition with zero- to second-order interface conditions was constructed to enhance physical consistency and suppress non-physical oscillations near bifurcations. Third, the strategy was extended to 1D pulse wave propagation, bridging local 3D reconstruction and network-scale analysis.

Numerical experiments on porcine abdominal aorta geometries demonstrated that the adaptive weighting yields more uniform gradient contributions and stable optimization paths, the overlapping decomposition improves reconstructed flow and accelerates convergence by nearly an order of magnitude in residual reduction, and the 1D formulation confirms the feasibility of joint minimization across PDE, boundary, and interface terms. Future work will focus on integrating multi-subject imaging data, exploring residual-driven adaptive resampling, and connecting the 3D and 1D scales through quantitative transfer metrics. The proposed framework can serve as a general training enhancement module for multi-physics PINNs, particularly in scenarios with uneven constraint distributions or significant multi-task gradient scale differences. Several limitations should be acknowledged: the geometric and physiological parameters are based on porcine abdominal aorta and may not directly generalize to human cerebral or coronary vasculature; the 3D-to-1D coupling is methodological and has not been quantitatively validated against coupled multiscale simulations; and the adaptive weight hyperparameters (α , initial weights) remain empirically tuned rather than theoretically derived.

Acknowledgement: Not applicable.

Funding Statement: This work was supported by The Key Research and Development Program of Hunan Provincial Science and Technology Department, China (Grant No. 2023DK2003), the Changsha Major Science and Technology Special Project (Grant No. kh2401008), the Key Scientific Research Project Plan of Institutions of Higher Education in Henan Province (No. 25A520062), the Doctoral Special Fund Project of Nanyang Normal University (Nos. 2025ZX010, 2025ZX014), and Hunan Provincial Natural Science Foundation of China (grant number 2026JJ90157). The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Author Contributions: The authors confirm contribution to the paper as follows: conceptualization, Muzhou Hou and Wen Hou; methodology, Muzhou Hou; software, Muzhou Hou; validation, Muzhou Hou, Dezhou Lin, Wen Hou and Kaituo Xu; formal analysis, Muzhou Hou; investigation, Muzhou Hou and Kaituo Xu; resources, Muzhou Hou; data curation, Muzhou Hou; writing—original draft preparation, Muzhou Hou; writing—review and editing, Dezhou Lin and Kaituo Xu; visualization, Muzhou Hou; supervision, Muzhou Hou; project administration, Muzhou Hou; funding acquisition, Muzhou Hou. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The data that support the findings of this study are available from the Corresponding Author, M. H., upon reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Wang Z. Report on cardiovascular health and diseases in China 2024: prevalence of cardiovascular diseases and their risk factors. *Chin J Cardiovasc Res.* 2025;23(9):771–93. (In Chinese). doi:10.37191/maps-ci-jccr-5(2)-094.
2. Liu F, Su S, Lu X. Epidemiological characteristics and trends of cardiovascular diseases in China. *Sci China Life Sci.* 2026;56:1–16. (In Chinese). doi:10.1360/SSV-2025-0221.
3. Acuna A, Berman AG, Damen FW, Meyers BA, Adelsperger AR, Bayer KC, et al. Computational fluid dynamics of vascular disease in animal models. *J Biomech Eng.* 2018;140(8):0808011–14. doi:10.1115/1.4039678.
4. Praharaj P, Sonawane CR, Bongale A. Advancement in CFD and responsive AI to examine cardiovascular pulsatile flow in arteries: a review. *Comput Model Eng Sci.* 2024;141(3):2021–64. doi:10.32604/cmesci.2024.056289.
5. Raissi M, Perdikaris P, Karniadakis GE. Physics informed deep learning (part I): data-driven solutions of nonlinear partial differential equations. *arXiv:1711.10561.* 2017.
6. Raissi M, Perdikaris P, Karniadakis GE. Physics informed deep learning (part II): data-driven discovery of nonlinear partial differential equations. *arXiv:1711.10566.* 2017.
7. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys.* 2019;378:686–707. doi:10.1016/j.jcp.2018.10.045.
8. Milano M, Koumoutsakos P. Neural network modeling for near wall turbulent flow. *J Comput Phys.* 2002;182(1):1–26. doi:10.1006/jcph.2002.7146.
9. Brunton SL, Noack BR, Koumoutsakos P. Machine learning for fluid mechanics. *Annu Rev Fluid Mech.* 2020;52:477–508. doi:10.1146/annurev-fluid-010719-060214.
10. Psychogios DC, Ungar LH. A hybrid neural network-first principles approach to process modeling. *AIChE J.* 1992;38(10):1499–511. doi:10.1002/aic.690381003.
11. Lagaris IE, Likas A, Fotiadis DI. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Trans Neural Netw.* 1998;9(5):987–1000. doi:10.1109/72.712178.
12. Raissi M, Yazdani A, Karniadakis GE. Hidden fluid mechanics: learning velocity and pressure fields from flow visualizations. *Science.* 2020;367(6481):1026–30. doi:10.1126/science.aaw4741.
13. Jagtap AD, Kawaguchi K, Karniadakis G. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. *Proc Math Phys Eng Sci.* 2020;476(2239):20200334. doi:10.1098/rspa.2020.0334.
14. Cai S, Mao Z, Wang Z, Yin M, Karniadakis GE. A review of physics-informed neural networks for fluid mechanics. *Comput Appl.* 2024;44(S1):133–41. (In Chinese). doi:10.11772/j.issn.1001-9081.2023060769.
15. Wang S, Wang H, Perdikaris P. On the eigenvector bias of Fourier feature networks: from regression to solving multi-scale PDEs with physics-informed neural networks. *Comput Meth Appl Mech Eng.* 2021;384:113938. doi:10.1016/j.cma.2021.113938.
16. Lu L, Meng X, Mao Z, Karniadakis GE. DeepXDE: a deep learning library for solving differential equations. *SIAM Rev.* 2021;63(1):208–28. doi:10.1137/19m1274067.

17. Mao Z, Lu L, Marxen O, Zaki TA, Karniadakis GE. DeepM&Mnet for hypersonics: predicting the coupled flow and finite-rate chemistry behind a normal shock using neural-network approximation of operators. *J Comput Phys.* 2021;447:110698. doi:10.1016/j.jcp.2021.110698.
18. Kharazmi E, Zhang Z, Karniadakis GEM. Hp-VPINNs: variational physics-informed neural networks with domain decomposition. *Comput Meth Appl Mech Eng.* 2021;374:113547. doi:10.1016/j.cma.2020.113547.
19. Yu J, Lu L, Meng X, Karniadakis GE. Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. *Comput Meth Appl Mech Eng.* 2022;393:114823. doi:10.1016/j.cma.2022.114823.
20. Meng X, Karniadakis GE. A composite neural network that learns from multi-fidelity data: application to function approximation and inverse PDE problems. *J Comput Phys.* 2020;401:109020. doi:10.1016/j.jcp.2019.109020.
21. Chakraborty S. Transfer learning based multi-fidelity physics informed deep neural network. *J Comput Phys.* 2021;426:109942. doi:10.1016/j.jcp.2020.109942.
22. Shin Y, Darbon J, Em Karniadakis G. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs. *CiCP.* 2025;28(5):2042–74. doi:10.4208/cicp.oa-2020-0193.
23. Wang Y, Han X, Chang CY, Zha D, Braga-Neto U, Hu X. Auto-PINN: understanding and optimizing physics-informed neural architecture. *arXiv:2205.13748.* 2022.
24. Xiao H, Wu JL, Wang JX, Sun R, Roy CJ. Quantifying and reducing model-form uncertainties in Reynolds-averaged Navier–Stokes simulations: a data-driven, physics-informed Bayesian approach. *J Comput Phys.* 2016;324:115–36. doi:10.1016/j.jcp.2016.07.038.
25. Kissas G, Yang Y, Hwuang E, Witschey WR, Detre JA, Perdikaris P. Machine learning in cardiovascular flows modeling: predicting arterial blood pressure from non-invasive 4D flow MRI data using physics-informed neural networks. *Comput Meth Appl Mech Eng.* 2020;358:112623. doi:10.1016/j.cma.2019.112623.