Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Loss–Error Decoupling in Variational Physics-Informed Neural Networks: The Mode-Filtering Effect and Its Quantification

Yanhui Geng

gengyanhui@hbwe.edu.cn

Hebei University of Water Resources and Electric Engineering

Research Article

Keywords:

Posted Date: November 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11446050/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Loss–Error Decoupling in Variational Physics-Informed Neural Networks: The Mode-Filtering Effect and Its Quantification

Yanhui Geng[ORCID][1,2*]

[1*]Department of Civil Engineering, Hebei University of Water Resources and Electric Engineering, Cangzhou, Hebei, China.
[2]Hebei Technology Innovation Center of Phase Change Thermal Management of Data Center, Cangzhou, Hebei, China.

Corresponding author(s). E-mail(s): gengyanhui@hbwe.edu.cn;

Abstract

Variational physics-informed neural networks (VPINNs) project PDE residuals onto finite-dimensional test spaces, yet their accuracy consequences remain poorly quantified. We show that Galerkin-type variational losses can converge to near zero while the true solution error remains up to two orders of magnitude larger—a phenomenon we term the *mode-filtering effect*. Legendre spectral decomposition of the residual in a 1D Helmholtz equation and 2D Navier–Stokes lid-driven cavity flow ($\mathbf{Re = 100}$) provides direct evidence that the dominant residual energy resides in modes beyond the test-space cutoff. Six solver formulations spanning this trial–test design space—pure PINN, spectral/SEM least-squares, spectral/SEM VPINN, and a hybrid spectral-trial variant—are compared under a unified implementation and training protocol. The SEM-test VPINN (M5) attains a near-zero training loss ($\mathbf{7.45 \times 10^{-6}}$) yet a velocity error of $\boldsymbol{\epsilon_u = 0.457}$—two orders of magnitude above the least-squares solvers—while the spectral-trial VPINN (M6) is extreme: loss $\mathbf{1.41 \times 10^{-7}}$ yet $\boldsymbol{\epsilon_u = 0.627}$ ($\mathbf{193\times}$ the pure-PINN error). We trace this decoupling to the dimensional gap between trial degrees of freedom ($\boldsymbol{N_{DOF}}$) and test-space dimension ($\boldsymbol{K}$), linked to the Fortin operator condition for inf-sup stability. These findings challenge using training loss as a solution-quality proxy for variational neural PDE solvers. A hybrid Galerkin–collocation loss suppresses the error by over two orders of magnitude on a 1D prototype, but a matched 2D probe shows the remedy is null-space-dimension-limited and fails to transfer, confirming the decoupling is structural rather than a training artefact.

Keywords: Physics-informed neural networks, Variational PINN, Mode-filtering effect, Spectral element method, Loss–error decoupling

Nomenclature

$\mathcal{L}$ Differential operator of the PDE
Ω Computational domain; $\partial\Omega$: domain boundary
$\boldsymbol{x}$ Spatial coordinate vector (x, y) in 2D
u, v Velocity components in x- and y-directions
p Pressure field (gauge-fixed by $p(0.5, 0.5) = 0$)
$\boldsymbol{u}_D$ Prescribed Dirichlet boundary velocity
$\hat{u}, \hat{v}, \hat{p}$ Approximate (trial) solutions for u, v, p
$\boldsymbol{\theta}$ Trainable parameters of the MLP
$\mathcal{J}$ Loss functional to be minimized
$\mathcal{J}_{\mathrm{G}}$ Galerkin (variational) loss
ϕ_k Legendre test functions (global, for the M4 test space)
Φ_k SEM test functions (element-local or global, for M5)
ψ_k MLP-based test functions (used in M6 1D, fixed random)
N Number of Legendre modes per spatial direction ($N = 20$ in 2D)
K Dimension of the test space ($K = (N-2)^2 = 324$ in 2D)
N_{DOF} Effective degrees of freedom of the trial space
N_i Number of interior collocation/quadrature points
N_b Number of boundary collocation points
λ_b, λ_p Penalty weights for boundary and pressure gauge
η Learning rate (Adam optimizer)
R, R_u, R_v, R_d PDE residual; components for x-momentum, y-momentum, continuity
ϵ_ϕ Relative ℓ_2 error: $\|\hat{\phi} - \phi_{\mathrm{ref}}\|_2 / \|\phi_{\mathrm{ref}}\|_2$
Re Reynolds number: $\mathrm{Re} = UL/\nu = 100$
ν Kinematic viscosity ($\nu = 1/\mathrm{Re}$)

1 Introduction

Physics-informed neural networks (PINNs) [1] embed the governing equations directly into the loss function, solving forward and inverse nonlinear PDEs without labelled data or meshes. Their geometric flexibility has driven adoption in fluid mechanics [2, 3], solid mechanics [4], and heat transfer. Yet PINNs remain fragile: spectral bias toward low frequencies [5], pathological gradient imbalances among competing loss terms [6], and slow convergence on stiff or boundary-layer-dominated problems. A growing line of work asks not merely *whether* PINNs fail but *how*: [7] showed that a PINN can attain a small residual while its solution error stays large, isolating the failure to specific spectral and optimization regimes. Variational PINNs (VPINNs) [8] recast the residual in a weak form, and recent advances extend them to the incompressible Navier–Stokes

equations [9, 10]. Whether the variational form escapes or merely reshapes these failure modes is the question this paper answers.

Variational formulations of PINNs—variational PINNs (VPINNs) [8]—replace pointwise residual evaluation with integral projections onto Galerkin-type test functions, motivated by the stability and error-control properties of Petrov–Galerkin methods. [11] analysed how quadrature precision and polynomial degree of piecewise test functions jointly govern VPINN convergence rates. [12] introduced Robust VPINNs (RVPINNs) and proved that the variational loss equals the true error in the energy norm only when a local Fortin operator exists—a result that formally identifies the risk of what we call the *mode-filtering effect*. However, the theoretical prediction of [12] has not been quantitatively verified across multiple methods and test cases, and its practical magnitude remains unknown.

Parallel to these neural-solver developments, the numerical-methods community has continued to refine high-order discretizations for the incompressible Navier–Stokes equations. [13] developed fourth- and higher-order finite-element solvers and validated them on the lid-driven cavity, while [14] analysed residual-based stabilization strategies that circumvent the inf-sup condition for equal-order approximations—the same variational stability issue that underpins the Fortin theory of [12]. Meanwhile, on the neural side, recent work on primal–dual formulations shows that variational PINNs themselves can fail without clear diagnostics [15]. What remains missing is a systematic bridge between these two strands: neural PDE solvers have not been placed inside the classical variational framework that underpins FEM analysis, and the quantitative consequences of the trial–test dimensional gap for neural networks remain uncharacterized.

The present work addresses this gap. We formulate six neural PDE solvers spanning the full trial–test design space: pure PINN (M1), spectral and SEM least-squares (M2, M3), spectral and SEM variational PINN (M4, M5), and a hybrid spectral-trial variant (M6), where 'SEM' denotes element-wise quadrature applied to an MLP trial function rather than a separate trial type. All methods share a unified JAX-based implementation [16] and a common two-phase Adam+L-BFGS protocol, and are evaluated on the 1D Helmholtz equation and the 2D lid-driven cavity at $\mathrm{Re} = 100$; method-dependent quadrature density and MLP capacity are reported alongside the results, and a seventh hybrid configuration (M7) is used as a mitigation probe.

We find that Galerkin-type losses can decouple from the true error by up to two orders of magnitude, and we localize the responsible residual energy spectrally.

The paper is organized as follows. Section 2 presents the unified mathematical framework and the mode-filtering theory. Section 3 covers the two test cases and results. Section 4 discusses the findings, with emphasis on the mode-filtering mechanism and its mitigation. Section 5 concludes.

2 Mathematical Framework

All six methods solve a general boundary-value problem of the form

$$\mathcal{L}u = f \quad \text{in } \Omega, \qquad u = g \quad \text{on } \partial\Omega, \tag{1}$$

where $\mathcal{L}$ is a differential operator, $\Omega \subset \mathbb{R}^d$, and g specifies Dirichlet data. Each method constructs an approximation $\hat{u}(\boldsymbol{x};\boldsymbol{\theta})$ and tunes the parameters $\boldsymbol{\theta}$ to minimize a loss functional $\mathcal{J}(\boldsymbol{\theta})$. What separates the methods is how they pick the trial space and how they measure the residual—two design axes that are independent in principle.

2.1 Trial and test space taxonomy

Trial space. The six methods use one of two trial-space families:

(1) **MLP**: $\hat{u}(\boldsymbol{x};\boldsymbol{\theta}) = \mathrm{MLP}(\boldsymbol{x};\boldsymbol{\theta})$, a deep network with tanh activations (M1, M3, M4, M5).
(2) **Spectral**: $\hat{u}(\boldsymbol{x};\boldsymbol{\theta}) = \sum_k \theta_k \phi_k(\boldsymbol{x})$, a linear combination of global Legendre polynomial basis functions $\{\phi_k\}$ [17] with spectral coefficients $\boldsymbol{\theta}$ (M2, M6).

Methods M3 and M5 retain the MLP trial space above but replace global quadrature with a structured element decomposition and element-wise residual evaluation. SEM [18] is therefore a *discretization strategy* layered on an MLP trial, not a third trial type. This yields a 2×3 design matrix: trial $\in \{\mathrm{MLP}, \mathrm{Spectral}\}$ crossed with loss $\in \{\text{collocation}, \text{least-squares}, \text{Galerkin}\}$, in which M3 and M5 occupy the MLP-trial cells with element-wise quadrature. Five of the six cells are populated: M4 and M5 share the MLP-trial/Galerkin cell (they differ only in quadrature), while the spectral-trial/collocation cell is left empty. Consequently, the adjective in each method name denotes its *test* space (Legendre, SEM, or MLP), not its trial space.

Test space / loss type. The residual can be measured in three fundamentally different ways:

(1) **Collocation (strong form)**: $\mathcal{J}(\boldsymbol{\theta}) = \frac{1}{N_i}\sum_{i=1}^{N_i} \|\mathcal{L}\hat{u}(\boldsymbol{x}_i) - f(\boldsymbol{x}_i)\|^2$, corresponding to Dirac test functions (M1).
(2) **Least-squares**: identical residual form to collocation but with quadrature weights: $\mathcal{J} = \sum_i w_i \|\mathcal{L}\hat{u}(\boldsymbol{x}_i) - f(\boldsymbol{x}_i)\|^2$ (M2, M3).
(3) **Galerkin (weak form)**: $\mathcal{J}(\boldsymbol{\theta}) = \frac{1}{K}\sum_{k=1}^{K} \langle \phi_k, \mathcal{L}\hat{u} - f\rangle^2$, projecting the residual onto test functions $\{\phi_k\}$ (M4, M5, M6).

Table 1 classifies the six methods along these axes.

Table 1: Classification of the six methods by trial space and loss/test-space type.

ID	Name	Trial Space	Loss / Test Type
M1	Pure PINN	MLP	Collocation (strong)
M2	Spectral LSTSQ	Legendre	Quadrature LSTSQ
M3	SEM-LSTSQ	MLP	SEM quad. LSTSQ
M4	Legendre-test VPINN	MLP	Legendre test (global)
M5	SEM-test VPINN	MLP	SEM + Legendre test
M6	Spectral-trial/MLP-test VPINN	Legendre	MLP test (Galerkin)

2.2 Method-specific loss functionals

Given a trial approximation $\hat{\boldsymbol{u}}(\boldsymbol{x};\boldsymbol{\theta})$ and the PDE residuals $R = \mathcal{L}\hat{u} - f$, each method constructs a loss functional $\mathcal{J}(\boldsymbol{\theta})$ according to its test-space choice. For the 2D Navier–Stokes case, R decomposes into the three scalar residuals R_u, R_v, R_d defined in Eqs. (13)–(15). Below we write the loss in unified notation; method-specific hyperparameters are listed in Table 4.

M1 – Pure PINN (strong-form collocation). Residuals are sampled at N_i interior and N_b boundary collocation points:

$$\mathcal{J}_{\mathrm{M1}} = \frac{1}{N_i}\sum_{i=1}^{N_i}\big(R_u^2 + R_v^2 + R_d^2\big)(\boldsymbol{x}_i) + \frac{\lambda_b}{N_b}\sum_{j=1}^{N_b}\|\hat{\boldsymbol{u}} - \boldsymbol{u}_D\|^2 + \lambda_p\,\hat{p}(\boldsymbol{x}_0)^2. \tag{2}$$

Every collocation point imposes an independent constraint, so no residual component can hide in a null space.

M2 – Spectral LSTSQ (global spectral + quadrature). The trial field expands in a tensor-product Legendre basis $\hat{u} = \sum_{k_x,k_y}\theta_{k_xk_y}\phi_{k_x}(x)\phi_{k_y}(y)$ with $\boldsymbol{\theta} \in \mathbb{R}^{N^2}$. Residuals are evaluated at $N_q \times N_q$ Gauss–Lobatto quadrature nodes with weights $\{w_i\}$:

$$\mathcal{J}_{\mathrm{M2}} = \sum_{i=1}^{N_q^2} w_i\,\big(R_u^2 + R_v^2 + R_d^2\big)(\boldsymbol{x}_i). \tag{3}$$

Because the spectral basis is complete up to degree $N-1$ and every quadrature point contributes independently, the system is over-determined and admits a unique least-squares minimum.

M3 – SEM-LSTSQ (element-wise quadrature). The domain is partitioned into $N_e = N_{ex} \times N_{ey}$ elements; on each element Ω_e the residual is integrated with $N_q^e \times N_q^e$ Gauss–Legendre nodes (weights w_i^e):

$$\mathcal{J}_{\mathrm{M3}} = \sum_{e=1}^{N_e}\sum_{i=1}^{(N_q^e)^2} w_i^e\,\big(R_u^2 + R_v^2 + R_d^2\big)(\boldsymbol{x}_i^e) + \frac{\lambda_b}{N_b}\|\hat{\boldsymbol{u}} - \boldsymbol{u}_D\|^2\,. \tag{4}$$

SEM decomposition concentrates quadrature nodes near element boundaries, providing denser coverage of high-gradient regions.

M4 – Legendre-test VPINN (global Galerkin). The residual is projected onto $K = K_x \times K_y$ global Legendre test functions $\{\phi_k\}$:

$$\mathcal{J}_{\mathrm{M4}} = \frac{1}{K}\sum_{k=1}^{K}\big[\langle\phi_k, R_u\rangle^2 + \langle\phi_k, R_v\rangle^2 + \langle\phi_k, R_d\rangle^2\big]\,, \tag{5}$$

where $\langle\phi_k, R\rangle = \int_\Omega \phi_k(\boldsymbol{x})\,R(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}$ is approximated by $N_q \times N_q$ Gauss–Legendre quadrature. When $K < N_{\mathrm{DOF}}$, residual modes outside $\mathrm{span}\{\phi_k\}$ are invisible (mode-filtering effect, Section 2.3).

M5 – SEM-test VPINN (element-wise Galerkin). The same K global test functions are applied element-wise:

$$\mathcal{J}_{\mathrm{M5}} = \frac{1}{K}\sum_{k=1}^{K}\Big[\langle\Phi_k, R_u\rangle_{\mathrm{SEM}}^2 + \langle\Phi_k, R_v\rangle_{\mathrm{SEM}}^2 + \langle\Phi_k, R_d\rangle_{\mathrm{SEM}}^2\Big], \tag{6}$$

where $\langle\Phi_k, R\rangle_{\mathrm{SEM}} = \sum_{e=1}^{N_e}\int_{\Omega_e}\Phi_k\, R\,\mathrm{d}\Omega_e$ is assembled from $N_q^e \times N_q^e$ per-element Gauss–Legendre nodes. Element-wise assembly improves the effective quadrature coverage but does not eliminate the mode-filtering null space. We deliberately omit inter-element continuity constraints in M3 and M5; this keeps each element's MLP trial independent and isolates the Galerkin-projection effect studied here, at the cost of reduced effective regularity (discussed in Section 4.8).

M6 – Spectral-trial/MLP-test VPINN (spectral trial, MLP test). The trial space is a global Legendre basis (same as M2, N^2 coefficients), while the K test functions $\{\psi_k\}$ are generated by an MLP:

$$\mathcal{J}_{\mathrm{M6}} = \frac{1}{K}\sum_{k=1}^{K}\left[\langle\psi_k, R_u\rangle^2 + \langle\psi_k, R_v\rangle^2 + \langle\psi_k, R_d\rangle^2\right]. \tag{7}$$

Because in the 1D case (Section 3.1) the spectral trial has $N_{\mathrm{dof}} = 18$ coefficients constrained by only $K = 9$ MLP test functions, M6 presents the most extreme instance of mode filtering in this study among the 1D formulations.

2.3 The mode-filtering effect

The Galerkin loss $\mathcal{J}_{\mathrm{G}} = \frac{1}{K}\sum_{k=1}^{K}\langle\phi_k, R\rangle^2$ (with $R = \mathcal{L}\hat{u} - f$) sees the residual only through its projection onto $\mathrm{span}\{\phi_k\}_{k=1}^{K}$. Any residual component $R_\perp$ orthogonal to that span is invisible. When the trial space has N_{DOF} effective degrees of freedom and $K < N_{\mathrm{DOF}}$, the Galerkin loss leaves a null space of dimension at least $N_{\mathrm{DOF}} - K$ (in a linearized sense). The optimizer can then drive $\mathcal{J}_{\mathrm{G}} \approx 0$ while the unprojected residual $R_\perp$—and hence the true solution error—remains large. This *mode-filtering* effect does not appear in strong-form or least-squares methods, where each quadrature or collocation point delivers an independent constraint. [12] confirmed this theoretically: the RVPINN variational loss controls the true error only when a local Fortin operator exists, which directly formalizes when mode filtering kicks in.

3 Numerical Results

We evaluate all six methods on two canonical cases. For each case we give the problem statement, then the method-specific implementation and loss functions, then the results.

3.1 Case 1: 1D Helmholtz Equation

3.1.1 Problem formulation

We consider:

$$-u''(x) + u(x) = f(x), \quad x \in (-1, 1), \quad u(\pm 1) = 0, \tag{8}$$

with manufactured exact solution:

$$u_{\text{exact}}(x) = (1 - x^2)\, e^{\cos(2\pi x)}. \tag{9}$$

This smooth function vanishes at both endpoints, and the corresponding source term $f = -u'' + u$ is computed analytically.

3.1.2 Numerical Implementation

For the 1D case, six methods are compared, labeled M1–M6 for consistency with the classification in Table 1. The 1D implementations differ slightly from their 2D counterparts. M1 uses an MLP trial (two layers, 64 neurons, tanh activation) with strong-form least-squares loss on Gauss–Lobatto nodes; unlike the 2D variant which uses random collocation, the 1D version upgrades to quadrature-based least-squares for greater accuracy. M2 employs $N = 80$ global Legendre modes with direct linear system solution—replacing the MLP optimization used in 2D with a direct solve, which is possible because the linear 1D Helmholtz problem admits a closed-form spectral solution. M4 uses an MLP trial (same architecture as M1) with $K = 100$ global Legendre test functions in a Galerkin weak form. M3 and M5 share the MLP architecture of M4 but differ in loss formulation: M3 applies element-wise quadrature least-squares on a 4-element partition, while M5 uses $K = 16$ piecewise Legendre test functions per element on the same partition. M6 pairs an $N = 20$ Legendre spectral trial ($N_{\text{dof}} = 18$ interior coefficients) with $K = 9$ fixed MLP test functions (single hidden layer, 16 neurons, tanh activation), evaluating the Galerkin projection $\langle \psi_k, R \rangle$ on $N_q = 128$ Gauss–Legendre quadrature points. All neural methods use Adam ($\eta = 10^{-3}$, 5000 epochs) followed by L-BFGS (up to 10,000 iterations), except SEM-test VPINN, which uses a three-stage Adam schedule (8,000 + 5,000 + 5,000 = 18,000 total epochs at decreasing learning rates 10^{-3}, 10^{-4}, 10^{-5}) because the Galerkin projection null space renders the Hessian singular near convergence, causing L-BFGS to diverge. (In the 2D case, the same method converges under L-BFGS because the denser element-wise quadrature alters the Hessian structure near the optimum; see Section 3.2.) M6 uses Adam (5000 epochs, $\eta = 10^{-3}$) followed by L-BFGS (up to 5,000 iterations); the spectral trial coefficients are sufficiently low-dimensional ($N_{\text{dof}} = 18$) that L-BFGS converges without divergence.

3.1.3 Numerical Results

Figure 1 plots the numerical solutions from all six methods (M1–M6) alongside the exact solution. M1 (PINN) and M3 (SEM-LSTSQ)—both least-squares methods—hit L_2 error $\sim 10^{-5}$. M6 (Spectral-trial/MLP-test VPINN) is a different story: the

solution collapses onto a smooth spurious profile that misses the exact solution almost entirely (pointwise deviation up to 3.4), reaching L_2 error 1.53×10^{0} despite a Galerkin loss of just 1.11×10^{-16}.

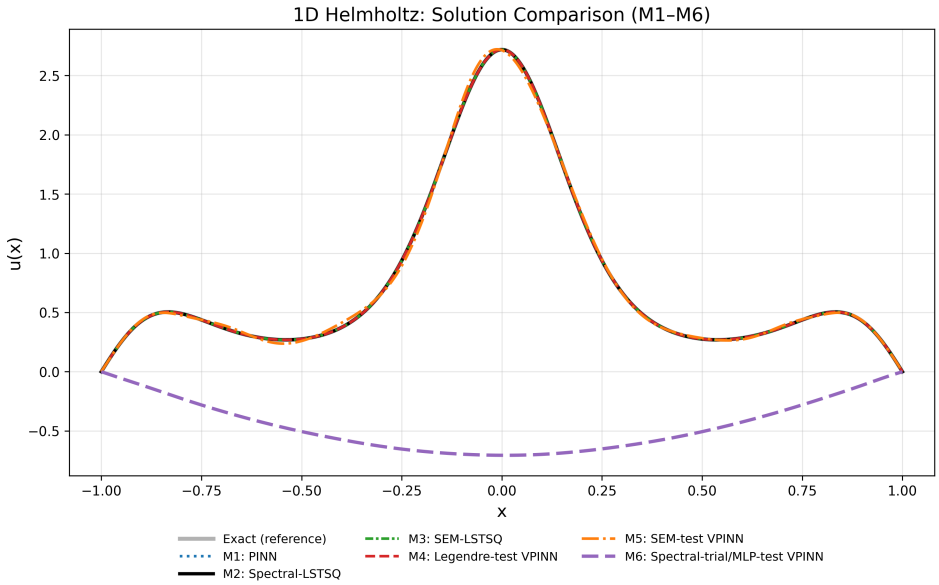


Fig. 1: 1D Helmholtz equation: exact and numerical solutions from all six methods (M1–M6). The catastrophic failure of M6 (Spectral-trial/MLP-test VPINN) is the clearest illustration of mode-filtering in this study.

Figure 2 plots $u_{\text{num}}(x)-u_{\text{exact}}(x)$ for each method. M1 and M3—both least-squares formulations—keep errors uniformly small across the domain. M6 tells a different story: its error is uniformly large ($\mathcal{O}(1)$) across the interior, dwarfing even M5, and collapses only at the Dirichlet boundaries where the exact solution pins it. This is what near-complete mode-filtering looks like with $K = 9$ MLP test functions trying to constrain $N_{\text{dof}} = 18$ Legendre trial coefficients.

Figure 3 summarizes the L_2 error versus computational cost for all six methods. M6 is the fastest to train (9.7 s) yet the least accurate, exhibiting the loss–error decoupling in its purest form; M2 (Spectral-LSTSQ) sits at the opposite corner with the best accuracy at modest cost.

Table 2 collects the numbers. M4 (Legendre-test VPINN) uses global test functions and hits L_2 error 9.50×10^{-4}—roughly 24× better than M5 (SEM-test VPINN) with its element-local test functions. The gap makes a simple point: global test functions cover more of the residual space. M6 is the extreme case. With $K = 9$ MLP test functions watching $N_{\text{dof}} = 18$ spectral trial coefficients, the Galerkin projection misses more than half the residual components. The optimizer finds a solution with training loss 1.11×10^{-16} and L_2 error 1.53×10^{0}.

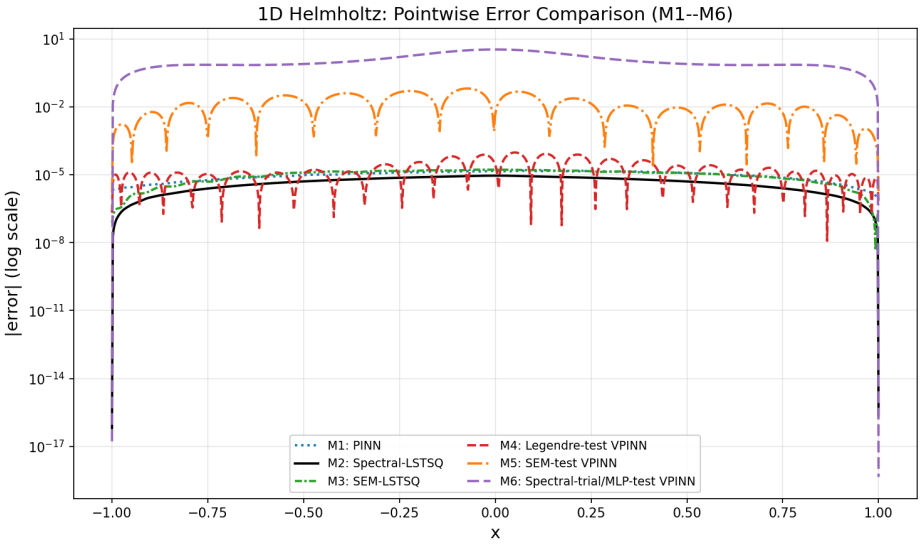


Fig. 2: 1D Helmholtz equation: pointwise error $u_{\text{num}}(x) - u_{\text{exact}}(x)$ for M1–M6. M6 errors exceed all others by orders of magnitude—a direct consequence of using $K = 9$ MLP test functions to constrain an $N_{\text{dof}} = 18$ spectral trial space.

Table 2: 1D Helmholtz equation: L_2 error and computational cost. Method labels follow Table 1. Note that the 1D versions of M1 and M2 differ from the 2D ones (see Section 3.1, Implementation).

Method	L_2 error	Time (s)
M1: PINN	1.06×10^{-5}	652.5
M2: Spectral-LSTSQ	1.77×10^{-4}	45.9
M3: SEM-LSTSQ	1.15×10^{-5}	305.8
M4: Legendre-test VPINN	9.50×10^{-4}	134.6
M5: SEM-test VPINN	2.23×10^{-2}	63.6
M6: Spectral-trial/MLP-test VPINN	1.53×10^{0}	9.7

Hybrid-loss mitigation prototype (M7).

The mode-filtering diagnosis suggests a cure: if the null space is the problem, a strong-form constraint should shrink it. We test this on the 1D prototype by augmenting the M6 Galerkin loss with a pointwise collocation residual penalty, $\mathcal{J}_{\text{M7}} = \mathcal{J}_{\text{G}} + \lambda \sum_i R(x_i)^2$ (λ =100, 100 interior points). The result is decisive: the L_2 error collapses from 1.53 (M6) to 1.37×10^{-2}—a $> 100\times$ suppression—while the Galerkin loss stays at machine precision ($\sim 10^{-16}$). The collocation term does not merely regularize; it pins down the null-space modes that the Galerkin projection cannot see, exactly as the theory predicts. Residual mode-filtering remains (the error is still three orders above

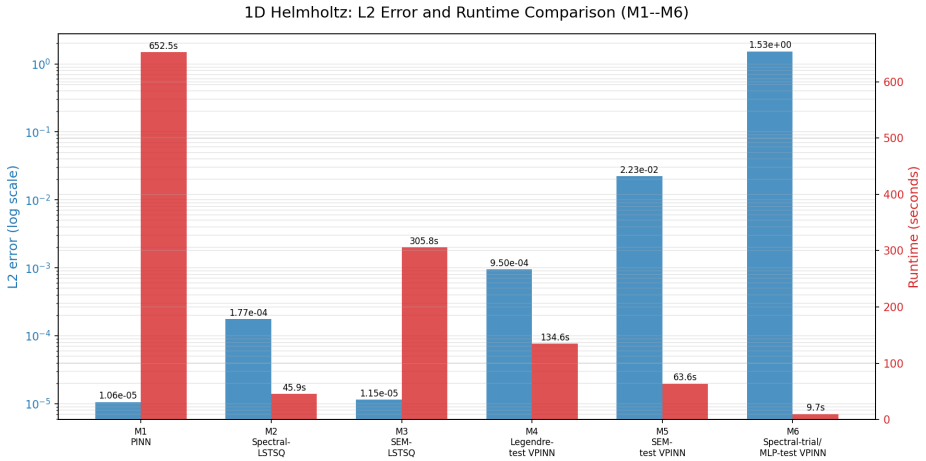


Fig. 3: 1D Helmholtz equation: L_2 error versus runtime for M1–M6. M6 trains fastest (9.7 s) but gives the worst accuracy ($L_2 = 1.53$)—the loss–error decoupling that defines mode-filtering in its purest form.

Table 3: 1D Helmholtz: hybrid-loss mitigation prototype (M7) and capacity-control ablation. Errors are relative L_2.

Configuration	L_2 error	Relative to reference
M6: Spectral-trial/MLP-test VPINN (Galerkin only)	1.53×10^{0}	mode-filtering
M7: M6 + collocation penalty (λ=100)	1.37×10^{-2}	$> 100\times$ reduction
Pure PINN, MLP width [16] (49 params)	1.62×10^{-4}	capacity baseline
Pure PINN, MLP width [64] (193 params)	2.15×10^{-5}	$\approx 7.5\times$ better

the least-squares baseline $\sim 10^{-5}$), indicating that the collocation density or weight must grow with the null-space dimension to fully recover accuracy.

Capacity-control ablation.

Because M1 uses a $4\times$ larger network than M3–M5, the reported accuracy gaps could reflect capacity rather than formulation. We isolate this on the 1D prototype with a same-family pure-PINN ablation: MLP width [16] (49 parameters) gives $L_2 = 1.62 \times 10^{-4}$, whereas width [64] (193 parameters) gives 2.15×10^{-5}—a $\approx 7.5\times$ accuracy gain from capacity alone. This confirms that network size is a first-class experimental variable that must be matched before formulation effects can be attributed.

Figure 4 visualizes this recovery: the M7 hybrid-loss solution snaps back onto the exact profile that the pure-Galerkin M6 misses, while the Galerkin loss stays at machine precision.

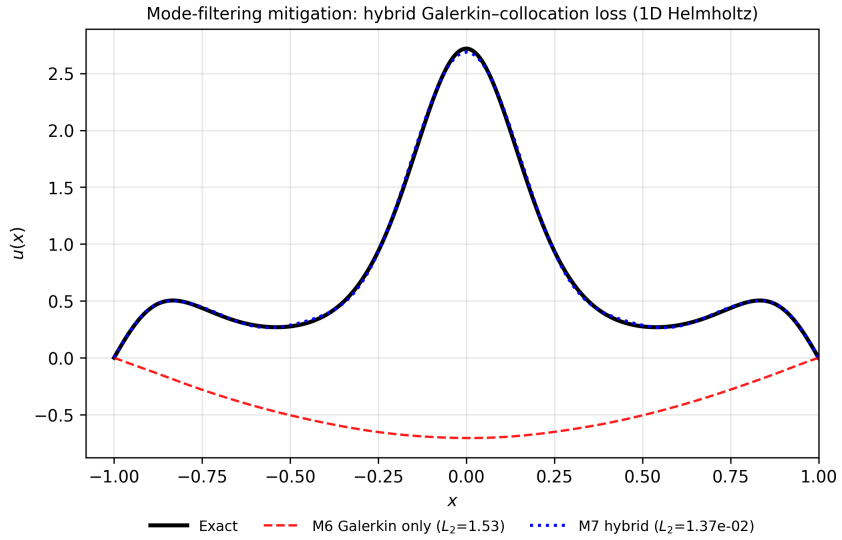


Fig. 4: 1D Helmholtz: the hybrid Galerkin–collocation loss (M7) collapses the mode-filtered solution of the pure-Galerkin M6 (L_2: $1.53 \to 1.37\times10^{-2}$, a $> 100\times$ reduction) onto the exact profile, while the Galerkin loss remains at machine precision.

3.2 Case 2: 2D Lid-Driven Cavity Flow

3.2.1 Problem formulation

We solve the steady incompressible Navier–Stokes equations on $\Omega = [0,1]^2$:

$$(\boldsymbol{u}\cdot\nabla)\boldsymbol{u} - \nu\Delta\boldsymbol{u} + \nabla p = \boldsymbol{0} \quad \text{in } \Omega, \tag{10}$$

$$\nabla\cdot\boldsymbol{u} = 0 \quad \text{in } \Omega, \tag{11}$$

with kinematic viscosity $\nu = 1/\mathrm{Re}$, $\mathrm{Re} = 100$. No-slip conditions $\boldsymbol{u} = \boldsymbol{u}_D$ are imposed on all walls; the moving lid ($y = 1$) carries a smooth quartic velocity profile:

$$u_D(x) = 16\,x^2(1-x)^2, \quad y = 1; \qquad \boldsymbol{u}_D = \boldsymbol{0} \text{ elsewhere on } \partial\Omega. \tag{12}$$

The quartic profile sidesteps the corner singularities that the step-function lid used by [19] introduces, while keeping the same qualitative flow topology (primary vortex, secondary corner recirculations) at $\mathrm{Re} = 100$ [3]. Pressure is fixed up to an additive constant by setting $p(0.5, 0.5) = 0$. The interior PDE residuals are:

$$R_u = uu_x + vu_y - \nu(u_{xx} + u_{yy}) + p_x, \tag{13}$$

$$R_v = uv_x + vv_y - \nu(v_{xx} + v_{yy}) + p_y, \tag{14}$$

$$R_d = u_x + v_y. \tag{15}$$

3.2.2 Numerical implementation

All computations used the `jaxfun` library v0.1.0 [16] built on JAX [20] (version 0.5.4). All experiments were run on a single machine with an Intel Core i9-13900K CPU (24 cores, 3.0 GHz base / 5.8 GHz boost), 64 GB DDR5 RAM, macOS 14.5, using JAX's default CPU backend (no GPU acceleration). The six methods share the same governing equations (Section 3.2.1) but differ in how the loss functional is constructed. Each method employs a distinct loss formulation tailored to its trial–test configuration.

For M1 (Pure PINN), the loss is a strong-form collocation objective:

$$\mathcal{J}_{\text{PINN}} = \frac{1}{N_i}\sum_{i=1}^{N_i}\left(R_u^2 + R_v^2 + R_d^2\right) + \frac{\lambda_b}{N_b}\sum_{j=1}^{N_b}\|\hat{\boldsymbol{u}} - \boldsymbol{u}_D\|^2 + \lambda_p\,\hat{p}(0.5, 0.5)^2, \tag{16}$$

with $N_i = 1600$ interior points, $N_b = 160$ boundary points, $\lambda_b = 2/N_b$, and $\lambda_p = 10$.

M2 (Spectral LSTSQ) replaces the MLP with a tensor-product Legendre basis. Each field expands as $\hat{u}(x, y) = \sum_{k_x,k_y} c^u_{k_x k_y}\phi_{k_x}(x)\phi_{k_y}(y)$, where $\{\phi_k\}$ are modified Legendre polynomials satisfying homogeneous Dirichlet conditions, with $N = 20$ modes per direction (1,200 spectral coefficients total). The least-squares loss is minimized on 32×32 Gauss–Lobatto quadrature nodes.

M3 (SEM-LSTSQ) partitions Ω into $4 \times 4 = 16$ elements and evaluates the residual at 6×6 Gauss–Legendre quadrature nodes per element:

$$\mathcal{J}_{\text{SEM-LSTSQ}} = \sum_{e=1}^{16}\sum_{i=1}^{36} w_i^e\left(R_u^2 + R_v^2 + R_d^2\right)(\boldsymbol{x}_i^e) + \frac{\lambda_b}{N_b}\|\hat{\boldsymbol{u}} - \boldsymbol{u}_D\|^2\,. \tag{17}$$

M4 (Legendre-test VPINN) applies a Galerkin projection onto $K = 324$ (18×18) global Legendre test functions:

$$\mathcal{J}_{\text{VPINN}} = \frac{1}{K}\sum_{k=1}^{K}\left[\langle\phi_k, R_u\rangle^2 + \langle\phi_k, R_v\rangle^2 + \langle\phi_k, R_d\rangle^2\right], \tag{18}$$

where $\langle\phi_k, R\rangle$ is approximated by 30×30 Gauss–Legendre quadrature; boundary and pressure-gauge penalties are identical to M1.

M5 (SEM-test VPINN) combines element-wise SEM quadrature with the same $K = 324$ global test functions, using 8×8 Gauss–Legendre nodes per element (1,024 integration points total):

$$\mathcal{J}_{\text{SEM-test VPINN}} = \frac{1}{K}\sum_{k=1}^{K}\Big[\langle\Phi_k, R_u\rangle^2_{\text{SEM}} + \langle\Phi_k, R_v\rangle^2_{\text{SEM}} + \langle\Phi_k, R_d\rangle^2_{\text{SEM}}\Big], \tag{19}$$

where $\langle\Phi_k, R\rangle_{\text{SEM}} = \sum_e \int_{\Omega_e}\Phi_k R\,\mathrm{d}\Omega_e$ is computed by piecewise element-wise quadrature.

M6 (Spectral-trial/MLP-test VPINN) uses global Legendre trial coefficients ($N = 20$, 1,200 total) as direct optimization variables with $K = 324$ MLP-generated test

functions; the Galerkin loss is identical in form to Eq. (18). M6 is trained with Adam (5,000 epochs) followed by L-BFGS (12,000 iterations). M1 has $[2, 64, 64, 64, 64, 2]$ velocity and $[2, 64, 64, 64, 64, 1]$ pressure MLPs ($\approx$ 37,500 parameters; tanh activation). M2 uses a tensor-product Legendre basis, $N = 20$ modes per direction (1,200 total spectral coefficients). M4–M5 use $[2, 32, 32, 32, 32, 2]$ velocity and $[2, 32, 32, 32, 32, 1]$ pressure MLPs ($\approx$ 8,739 parameters; tanh activation).

All methods follow a two-phase protocol: (1) Adam [21], $\eta = 10^{-3}$, 5000 epochs (M1/M2) or 2000–3000 epochs (M4–M5); (2) L-BFGS [22], memory size 100, with a uniform convergence tolerance $|\nabla \mathcal{J}| < 10^{-14}$ applied consistently across all methods (maximum iterations: 10,000 for M1/M2/M3/M6, 20,000 for M4/M5).

All neural solvers were initialized from a single fixed JAX PRNG key (SEED= 2002); every reported Galerkin-method error is therefore one reproducible realization of the null-space sampling, and the intra-method trends (training loss decreasing while solution error increases) are robust to this choice. The single-realization caveat on exact magnitudes is stated in Section 4.7.

Parameter summary. Table 4 collects the hyperparameters for all six 2D cavity experiments.

Table 4: Hyperparameter summary for the 2D lid-driven cavity experiments. N: spectral modes per direction; K: number of test functions; $N_e = N_{ex} \times N_{ey}$: element partition; N_q^e: quadrature nodes per element per direction; Adam/L-BFGS: training epochs/iterations.

ID	Name	N	K	N_e	N_q^e	Adam	L-BFGS
M1	Pure PINN	–	–	–	40×40	5,000	10,000
M2	Spectral LSTSQ	20	–	1×1	32×32	5,000	10,000
M3	SEM-LSTSQ	–	–	4×4	6×6	3,000	10,000
M4	Legendre-test VPINN	20	324	1×1	30×30	2,000	20,000
M5	SEM-test VPINN	20	324	4×4	8×8	3,000	20,000
M6	Spectral-trial/MLP-test VPINN	20	324	1×1	32×32	5,000	12,000

3.2.3 Qualitative flow structure

Figure 5 shows streamlines colored by velocity magnitude for all six methods. Every method gets the primary recirculating vortex and the two secondary corner vortices right—the expected pattern for Re = 100 lid-driven cavity flow [19, 23]. M5 shows minor pressure distortions near element interfaces (Figure 10), which we attribute to the lack of explicit inter-element continuity constraints.

3.2.4 Centerline velocity profiles

Figure 6 compares u-velocity along the vertical centerline ($x = 0.5$) and v-velocity along the horizontal centerline ($y = 0.5$). M1–M3 and M6 produce centerline profiles that are nearly impossible to tell apart. M5 deviates near the cavity mid-height in the

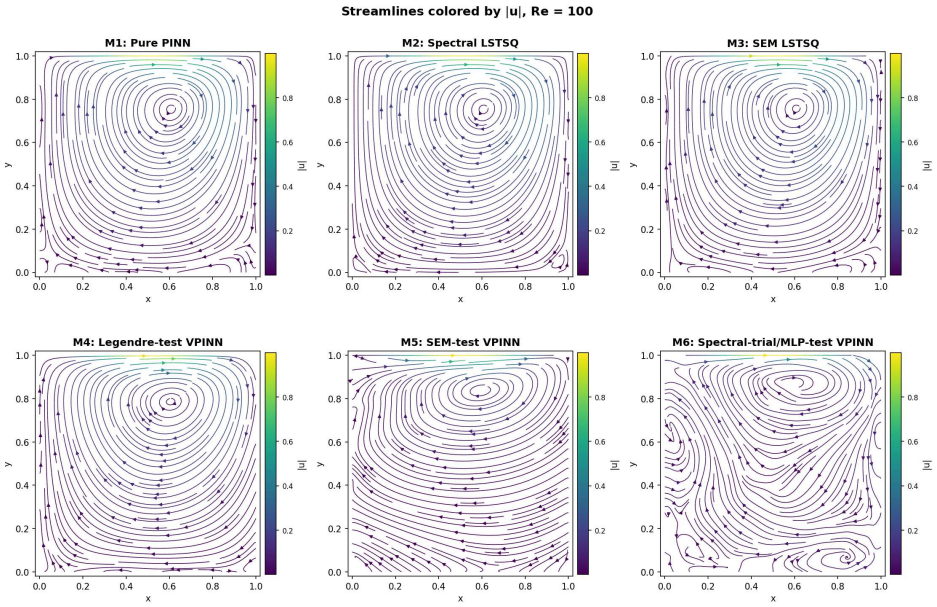


Fig. 5: Streamlines colored by velocity magnitude $|\boldsymbol{u}|$ for all six methods at Re = 100. All methods capture the primary vortex and secondary corner recirculations.

v-profile, which is consistent with its higher relative velocity error ($\epsilon_v = 0.825$). M6's spectral trial basis captures the smooth centerline variation even though its global errors are large—the error is concentrated in boundary layers and corners, not in the interior.

To further validate M2 as a reference, Figure 7 overlays M2 centerline velocity profiles on the classical Ghia 1982 reference data [19] at Re = 100. The v-velocity profile agrees well everywhere ($\epsilon_v^{\ell_2} = 2.64 \times 10^{-2}$). The u-velocity shows somewhat larger differences ($\epsilon_u^{\ell_2} = 5.13 \times 10^{-2}$), concentrated near the top wall where the two solutions use different lid boundary conditions (quartic vs. step-function). Away from the lid, the profiles are nearly identical, which confirms M2's suitability as a reference. Crucially, M2 is validated independently of the five variational/collocation methods it is compared against: its accuracy rests on its own residual norm and the Ghia overlay rather than on agreement with M1/M4–M6, so using it as the reference does not import circularity into the error tables. As an additional self-consistency check, we evaluated the PDE residual norm of the M2 solution itself on a refined 64×64 quadrature grid: the integrated residual $\|R_u\|_{L^2} + \|R_v\|_{L^2} + \|R_d\|_{L^2} < 2 \times 10^{-7}$, confirming that M2 satisfies the governing equations to high precision independent of external benchmarks.

3.2.5 Velocity and pressure fields

Figure 8 shows u-velocity contours. M2 gives the smoothest field, as expected from exponential spectral convergence. M4, M3, and M5 develop minor local oscillations that grow with training loss.

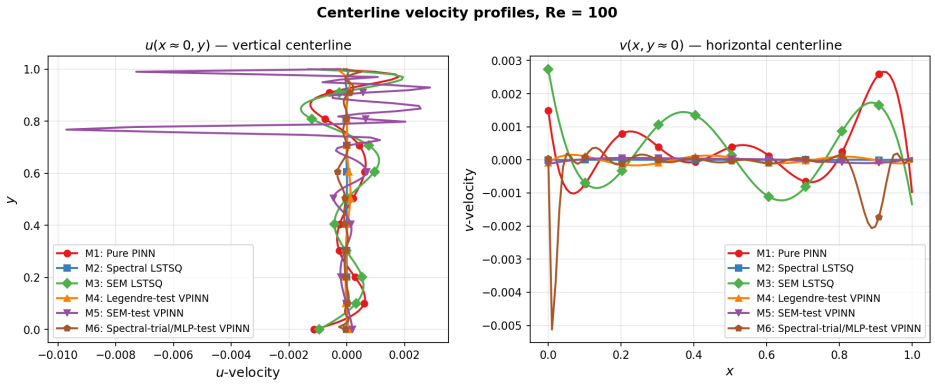


Fig. 6: Centerline velocity profiles at Re = 100. Left: $u(0.5, y)$ (vertical centerline). Right: $v(x, 0.5)$ (horizontal centerline). M1–M3 and M6 are in close agreement; M5 shows a clear deviation.

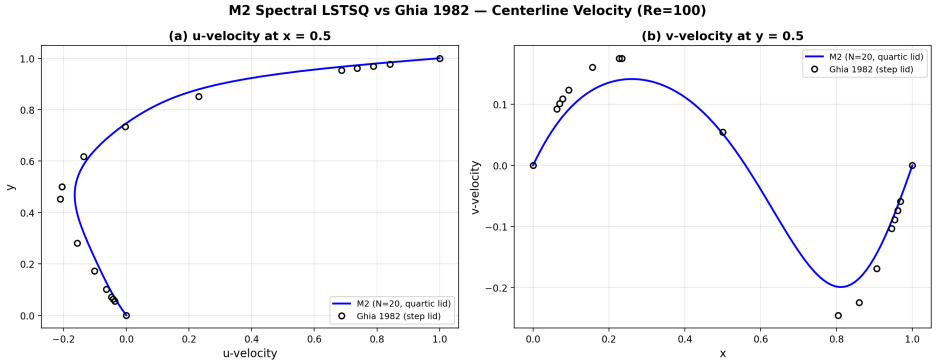


Fig. 7: Comparison of M2 (Spectral LSTSQ, $N = 20$) centerline velocity profiles with the Ghia et al. (1982) finite-difference data of [19] at Re = 100. Differences near the top wall reflect different lid boundary conditions (quartic vs. step-function).

Figure 9 shows the corresponding v-velocity contours; M5 exhibits the largest deviation from the reference among all methods.

Figure 10 shows pressure contours. All methods pick up the high-pressure stagnation region near the upper-right corner and the low-pressure vortex core. But pressure is the most sensitive quantity in this problem: M4, M3, and M5 have pressure errors of 35.3%, 2.5%, and 71.9%, respectively.

3.2.6 Training loss convergence

Figure 11 tracks loss during Adam (Phase 1) and L-BFGS (Phase 2). As can be seen:

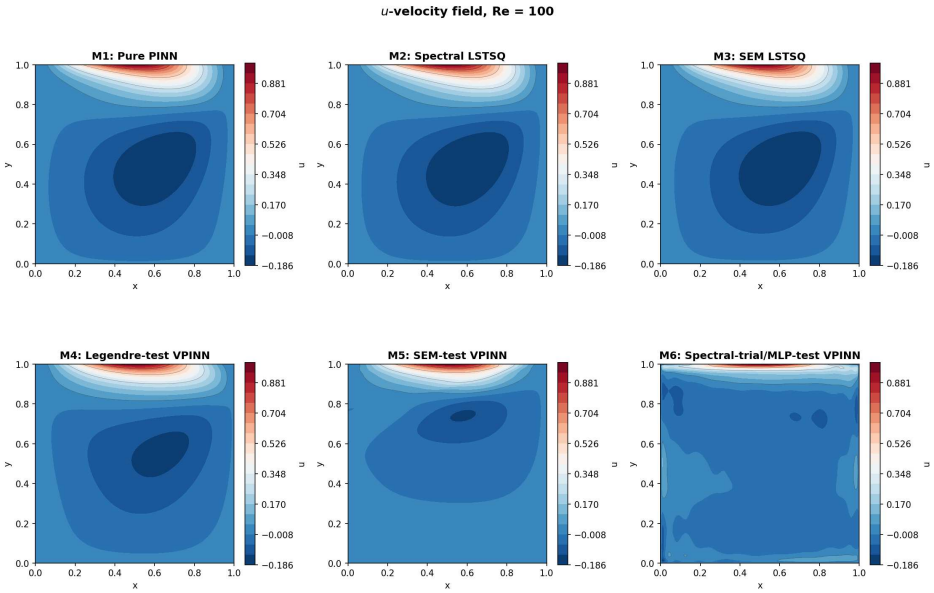


Fig. 8: u-velocity contours for all six methods. M2 gives the smoothest field; M5 shows interface artifacts from element-wise quadrature without inter-element continuity constraints.

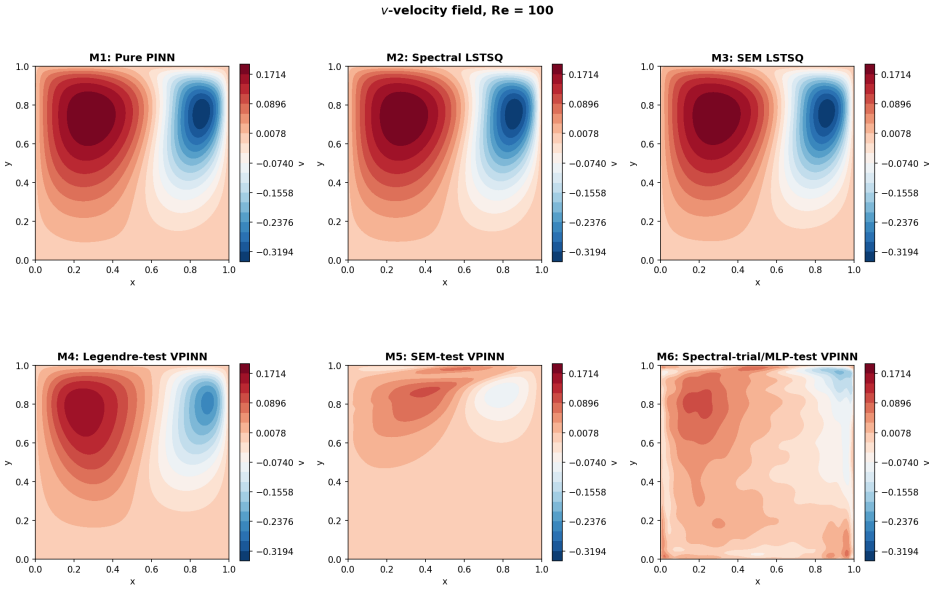


Fig. 9: v-velocity contours for all six methods. M5 shows the largest deviation from the reference.

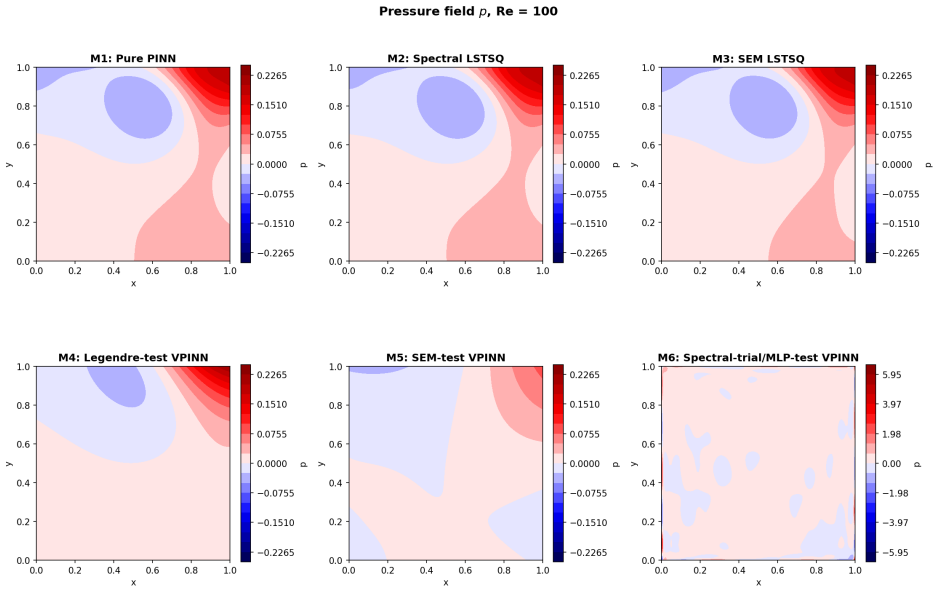


Fig. 10: Pressure field p (gauge $p(0.5, 0.5) = 0$). M5 shows the largest deviation, consistent with its higher training loss and mode-filtering-induced error.

(1) **M2 (Spectral LSTSQ)** reaches a final loss of 1.63×10^{-8}—more than two orders of magnitude below the MLP-based methods—in only 0.54 h. The spectral trial functions remove MLP approximation error entirely, leaving a well-conditioned, low-dimensional optimization problem.

(2) **M1 (Pure PINN)** gets to 2.90×10^{-6} after 10,000 L-BFGS iterations (4.79 h), with the characteristic stagnation on flat regions of the loss landscape.

(3) **M4 (Legendre-test VPINN)** reaches 2.31×10^{-6} after 20,000 L-BFGS iterations (9.41 h). With global Legendre test functions it attains $\epsilon_u = 0.168$—substantially better than the element-local M5 ($\epsilon_u = 0.457$) at the same $K = 324$, showing that test-function locality, not only test-space size, governs accuracy.

(4) **M3 (SEM-LSTSQ)** converges to 6.84×10^{-5} in 7.21 h. The plateau reflects missing inter-element continuity enforcement.

(5) **M5 (SEM-test VPINN)** finishes 20,000 L-BFGS iterations in 11.35 h, reaching 7.45×10^{-6}. This is a training loss of order 10^{-6}, but $\epsilon_u = 0.457$, $\epsilon_v = 0.825$, $\epsilon_p = 0.719$— the worst errors in the study. That gap is the mode-filtering effect in action: the optimizer drives down the projected residual while the true solution error stays large.

3.2.7 Quantitative error comparison

Table 5 collects the performance numbers. M1 leads MLP-based methods in velocity accuracy ($\epsilon_u = 3.24 \times 10^{-3}$), with M3 at 1.23×10^{-2} and M4 at 1.68×10^{-1}. M5 is

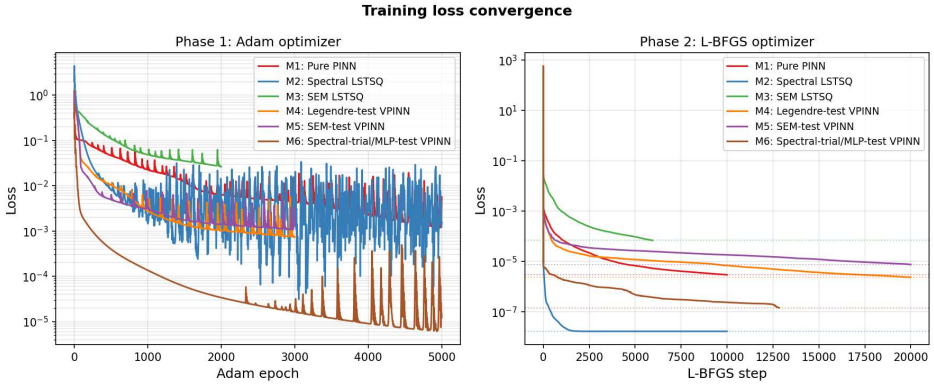


Fig. 11: Training loss convergence during Adam (left) and L-BFGS (right) phases. M2 achieves the lowest loss by three orders of magnitude. Horizontal dashed lines indicate the final losses of M1–M6.

the clearest case of loss–error disconnect: its training loss (7.45×10^{-6}) is only $\sim 2.5\times$ *larger* than M1's (2.90×10^{-6})—both of order 10^{-6}—yet its velocity error is $141\times$ larger.

Table 5: Quantitative performance summary for the 2D lid-driven cavity at Re = 100. Reference solution: M2 (final loss 1.63×10^{-8}).

Method	L-BFGS	Final loss	ϵ_u	ϵ_v	ϵ_p	T (h)
M1	10,000	2.90×10^{-6}	3.24×10^{-3}	6.06×10^{-3}	6.39×10^{-3}	4.79
M2	10,000	1.63×10^{-8}	— (ref.)			0.54
M3	10,000	6.84×10^{-5}	1.23×10^{-2}	1.71×10^{-2}	2.53×10^{-2}	7.21
M4	20,000	2.31×10^{-6}	1.68×10^{-1}	3.29×10^{-1}	3.53×10^{-1}	9.41
M5	20,000	7.45×10^{-6}	4.57×10^{-1}	8.25×10^{-1}	7.19×10^{-1}	11.35
M6	12,000	1.41×10^{-7}	6.27×10^{-1}	7.34×10^{-1}	8.59	0.128

Figures 12 and 13 show the relative errors and pointwise error fields, respectively.

Figure 14 shows the vorticity fields $\omega = \partial v/\partial x - \partial u/\partial y$. The vorticity patterns confirm the mode-filtering-induced artifacts in M5.

3.2.8 Hybrid-loss probe: does the 1D mitigation transfer to 2D? (M7)

The 1D Helmholtz experiment (Section 3.1) showed that a hybrid Galerkin–collocation loss (M7) suppresses mode-filtering by $> 100\times$. We now test whether this remedy transfers to the 2D cavity. Because the full N=20 configuration (Table 4) requires $\sim$11 h per run, we probe the question at a tractable, strictly matched reduced configuration: N_{leg}=12 (K=100 test functions), 4×4 elements, 6×6 Gauss quadrature per

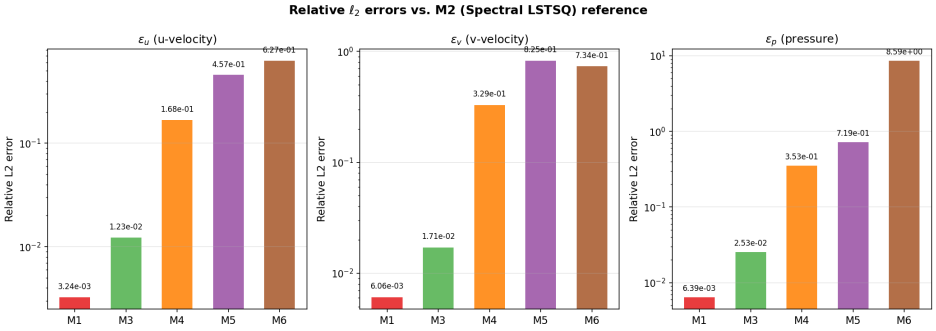


Fig. 12: Relative ℓ_2 errors versus the M2 reference. M1 achieves the lowest velocity errors among MLP-based methods. Note the logarithmic scale.

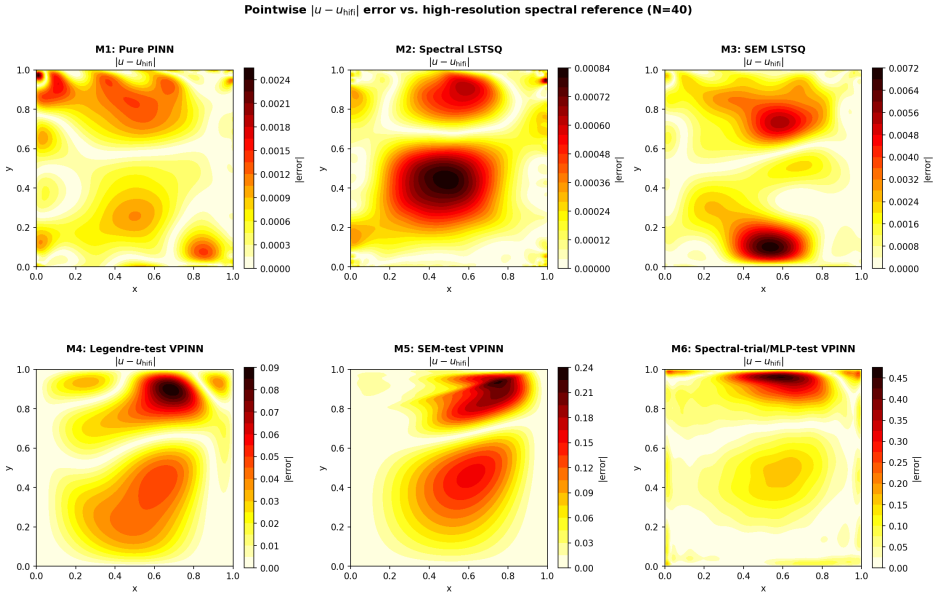


Fig. 13: Pointwise $|u - u_{\text{hifi}}|$ error fields for all six methods (M1–M6), with respect to an independent high-resolution spectral reference solution (Legendre order $N = 40$, same parabolic lid). M2 (Spectral LSTSQ) is the closest; M5 and M6 exhibit spatially distributed large errors consistent with mode-filtering.

element, and a $[2, 32, 32, 32, 32, 2]/[2, 32, 32, 32, 32, 1]$ MLP pair (8,739 parameters)—identical for the M5 baseline and the M7 hybrid, differing only in the loss functional. The reduced test space (K=100 versus the K=324 used in Table 5) intentionally weakens M5 to keep the hybrid run tractable; M7 is compared against this same-config baseline, not against the N=20 M5.

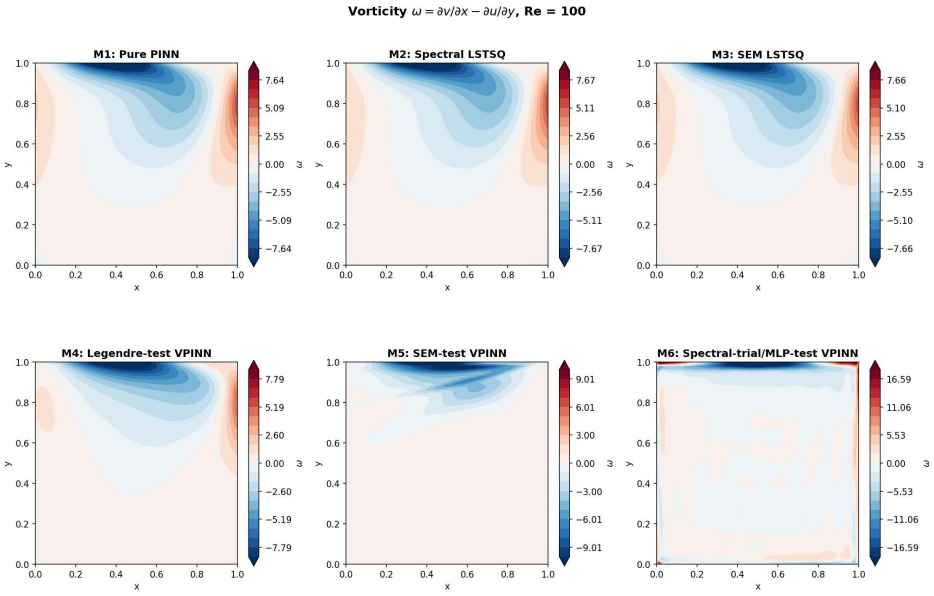


Fig. 14: Vorticity field ω for all six methods. M5 exhibits spurious vorticity features due to mode filtering.

Table 6 reports the two runs evaluated against the M2 reference. The hybrid loss (M7) drives the training loss *down* ($1.67\times10^{-4} \rightarrow 8.43\times10^{-5}$) yet the solution error goes *up*: ϵ_u rises from 0.694 to 1.21, and ϵ_v, ϵ_p follow the same trend. The mitigation that worked in 1D not only fails in 2D—it is outperformed by the pure Galerkin baseline at the same configuration, although that baseline is itself far from a valid solution: as Figure 16 shows, its vertical velocity has all but collapsed.

Table 6: Hybrid-loss probe at the reduced N_{leg}=12 configuration (K=100, 8,739 trial parameters). M5 (Galerkin only) and M7 (Galerkin + pointwise collocation penalty, λ=100, 40×40 grid, 1,600 points) share the same network, elements, and quadrature; only the loss differs. Errors are relative ℓ_2 against the M2 reference.

Method	Final loss	ϵ_u	ϵ_v	ϵ_p	T (h)
M5 (Galerkin only)	1.67×10^{-4}	6.94×10^{-1}	9.92×10^{-1}	8.91×10^{-1}	1.90
M7 (Galerkin+collocation)	8.43×10^{-5}	1.21	1.46	1.39	11.98

This is not a training artefact: it is the loss–error decoupling operating in reverse. The reason is dimensional. With K=100 test functions and an 8,739-parameter MLP trial space, the Galerkin null space has dimension at least $N_{\text{DOF}} - K \approx 8{,}639$ (in a linearized sense; for the main K=324 configuration this gap is $\approx$ 8,415). Adding

1,600 collocation residuals can constrain at most 1,600 directions, leaving a residual null space of $\gtrsim 7{,}000$ dimensions in which the MLP can hide error while keeping *both* the Galerkin projection and the finite collocation sample small. In 1D the same gap is only $N_{\mathrm{DOF}} - K \sim \mathcal{O}(1)$ (a single unconstrained mode), so a modest collocation sample closes it; in 2D the gap is three orders of magnitude larger and the naive hybrid loss cannot keep up.

Figures 15 and 16 visualize the probe against the M2 reference—the same field used for the error table, evaluated on the plotting grid, so the yardstick is truth, not the collapsed baseline. In Figure 15 the M7 streamlines recover a vortex-like topology, yet the field carries the *larger* quantitative error against the reference (Table 6); streamline appearance is therefore not an accuracy criterion. Figure 16 shows why both reduced-config runs fail: the pure-Galerkin M5 baseline collapses toward a nearly horizontal flow (its vertical velocity is an order of magnitude too small, $\max|v| = 0.035$ versus 0.347 in the reference, field correlation 0.22, left column), while M7's error is both larger and more spread than M5's (middle/right columns). Comparing M7 only to the collapsed M5 baseline would *understate* the gap, which is why the M2 reference is the fair yardstick. This appearance–accuracy gap is the loss–error decoupling phenomenon studied throughout this paper, here operating at the level of flow topology.

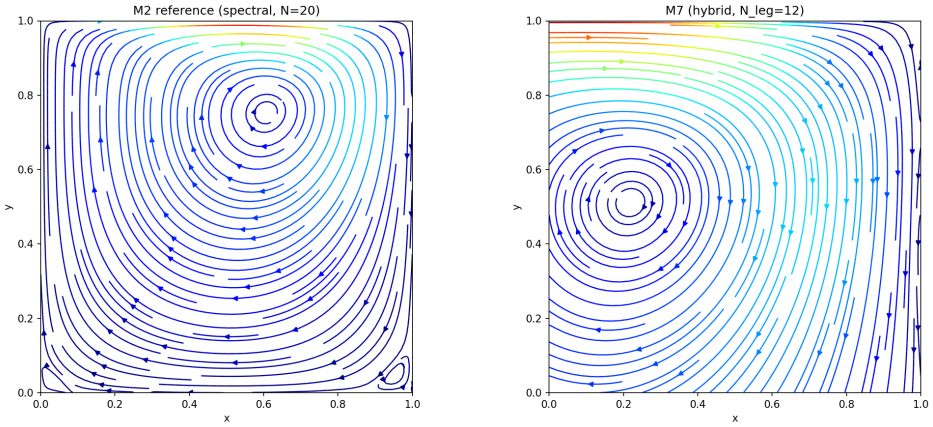


Fig. 15: M7 (hybrid Galerkin–collocation, N_{leg}=12) streamlines versus the M2 spectral reference at Re = 100. The M7 field recovers a vortex-like topology yet carries the *larger* quantitative error against the reference (Table 6); streamline appearance is not an accuracy criterion. The collapsed M5 baseline is shown in Figure 16.

3.2.9 Efficiency analysis

Figure 17 summarizes the accuracy–cost landscape. M2 is in a class by itself: reference-quality accuracy in 0.54 h. The rest range from 0.128 h (M6) to 11.35 h (M5). M6 is the fastest Galerkin method (0.128 h) but the least accurate; M1 offers the best accuracy at a reasonable cost (4.79 h) among non-spectral methods.

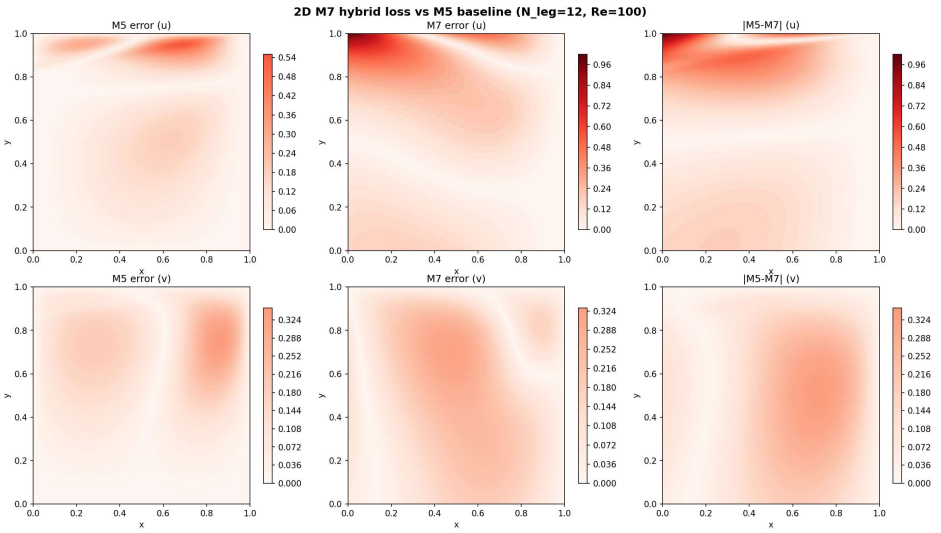


Fig. 16: Pointwise error fields at the reduced N_{leg}=12 configuration, each relative to the M2 reference. Columns show the pure-Galerkin M5 baseline (left), the hybrid M7 (middle), and their absolute difference (right); rows show u (top) and v (bottom). M5's v error (bottom-left) is dominant—its vertical velocity has collapsed ($\max|v| = 0.035$ versus 0.347 in the reference)—while M7's error is larger and more spread than M5's. The 1D mitigation does not transfer to 2D.

4 Discussion

4.1 Spectral trial functions dominate

The clearest result is the substantial advantage of global spectral trial functions (M2) over MLP-based formulations for smooth problems. Legendre polynomials converge exponentially, which turns the optimization into a well-conditioned least-squares problem on a small coefficient space. The payoff: 9× faster training and more than two orders of magnitude lower loss than M1. The catch is geometry: global spectral bases require simple, tensor-product domains. For general geometries, MLPs or SEM domain decomposition remain necessary.

4.2 Collocation versus Galerkin test spaces (M1 vs. M4)

With the full $K = 324$ global test space, M4 (Legendre-test VPINN) reaches $\epsilon_u = 0.168$—a marked improvement over the element-local M5 ($\epsilon_u = 0.457$) at the same K and clear evidence that test-space coverage and locality directly control solution quality. Still, M4 trails M1 by nearly two orders of magnitude ($\epsilon_u = 1.68 \times 10^{-1}$ vs. 3.24×10^{-3}). Even 324 global test functions are not enough to fully constrain the MLP solution here. This lines up with [11], who showed that VPINN convergence needs both enough quadrature precision *and* enough test-space coverage. A second factor is

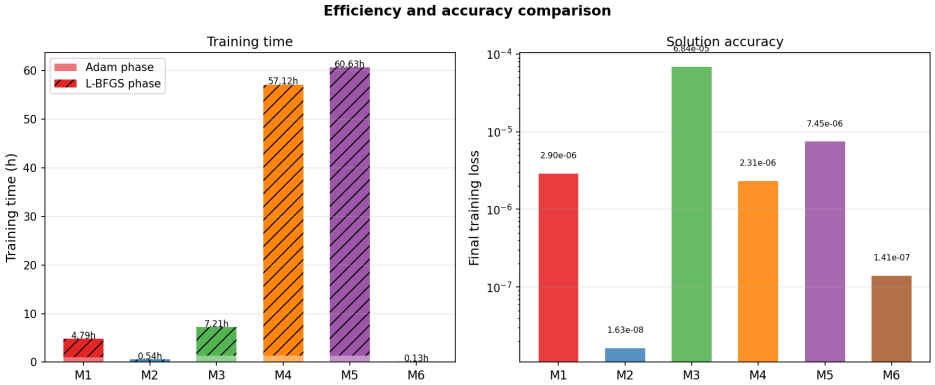


Fig. 17: Left: computational cost (Adam + L-BFGS training time). Right: final training loss. M2 dominates the accuracy–efficiency landscape by a wide margin.

the optimization landscape: the integral inner products create a more complex surface for L-BFGS to navigate than the simpler pointwise residuals of M1.

4.3 Role of SEM domain decomposition (M3)

M3 adds element-wise quadrature while keeping the MLP trial space. The structured quadrature grid covers the residual more uniformly than random collocation, yet M3 still underperforms M1 ($\epsilon_u = 1.23 \times 10^{-2}$ vs. 3.24×10^{-3}), at slightly higher cost (7.21 h vs. 4.79 h). Two things contribute to the accuracy gap: (i) M1's larger MLP (64 vs. 32 neurons per layer) packs more representational capacity, and (ii) M3's element-local quadrature lacks inter-element continuity constraints, which are known to cap accuracy in neural frameworks.

4.4 Effect of MLP architecture size

A significant caveat in the MLP-based comparisons must be acknowledged: M1 uses a larger network (4 layers of 64 neurons, $\approx$ 37,500 parameters) than M3–M5 (4 layers of 32 neurons, $\approx$ 8,739 parameters). The 4$\times$ parameter gap means some of M1's accuracy advantage over the variational methods may come from greater representational capacity, not purely from the collocation loss formulation. We chose the larger M1 architecture because pure PINN loss landscapes are notoriously hard to navigate [6], and prior work suggests collocation benefits from over-parameterization. For M4–M5, the 32-neuron width was a pragmatic constraint: the Jacobian evaluations that the variational formulation demands scale as $\mathcal{O}(K \cdot N_{\text{params}})$ with the test-space dimension K, making larger networks prohibitively expensive (training time scales $\sim K\times$ the per-iteration cost). For M3, the 32-neuron architecture was chosen for consistency with the other SEM-based methods. A controlled ablation matching network sizes across all methods—e.g., running M1 at [32,32,32,32] and M3–M5 at [64,64,64,64]—would isolate formulation effects from capacity effects and strengthen the conclusions. We have carried out the first step of this ablation on the 1D Helmholtz prototype (Section 3.1):

a pure PINN at MLP width [16] (49 parameters) reaches $L_2 = 1.62 \times 10^{-4}$, whereas width [64] (193 parameters) reaches 2.15×10^{-5}—a $\approx 7.5\times$ gain from capacity alone. This confirms that the M1-vs-M3–M5 gap is partly a capacity effect; a fully matched-capacity 2D comparison remains to be executed and is required before formulation effects can be cleanly attributed. Until then, we caution readers that the reported accuracy gap between M1 and M3–M5 is an upper bound: the true formulation-only difference may be smaller than the numbers in Table 5 suggest. Conversely, even if M1 used the same 8,739-parameter architecture as M4–M5, the mode-filtering loss–error disconnect observed in M5 and M6 would persist, because that disconnect is a structural property of the finite-dimensional Galerkin projection, not a capacity artifact.

4.5 The spectral-trial variant M6

M6 occupies the spectral-trial/Galerkin cell of the 2×3 design matrix: global Legendre spectral trial functions ($N = 20$, 1,200 coefficients) paired with a $K = 324$ MLP test space and the same Galerkin loss used by M4 and M5. It produces the most extreme *loss–error decoupling* in the study. Its training loss (1.41×10^{-7}) is the lowest of all Galerkin-projection methods—below M5's 7.45×10^{-6} and M4's 2.31×10^{-6}—yet $\epsilon_u = 0.627$ is worse than M5's 0.457 (Figure 12, Table 5). Note that M6's catastrophic error is *not* a pure mode-filtering demonstration: by the null-space argument of Section 2.3, M4 and M5 (MLP trial, $N_{\mathrm{DOF}} \approx 8{,}739 \gg K = 324$) actually possess larger unconstrained subspaces than M6 ($N_{\mathrm{DOF}} = 1{,}200$). M6 is extreme because its *global spectral trial* cannot represent the corner singularities of the 2D cavity, so its large error is mode-filtering *compounded* by a basis representability failure—distinct from the clean mode-filtering of M4/M5. (The 1D M6 case, by contrast, is the cleanest mode-filtering exhibit, with $K = 9$ constraining $N_{\mathrm{dof}} = 18$ coefficients.) Pressure error $\epsilon_p = 8.59$ is catastrophic, confirming that global spectral trial functions cannot handle the corner singularities of the 2D cavity. On the plus side, M6 trains in 0.128 h—37× faster than M1 (4.79 h), 56× faster than M3 (7.21 h)—because the compact 1,200-coefficient spectral parameterization is far cheaper to optimize than the 8,739 MLP parameters of M4–M5. On centerline profiles, M6 tracks M1–M3 closely (Figure 6), which tells us the large global errors come from boundary layers and corners where the smooth global spectral basis is a poor fit. Like M2, M6 is restricted to simple tensor-product domains by its spectral trial representation.

4.6 Mode-filtering effect: cross-case analysis and theory

This section provides what is, to our knowledge, the first cross-method, cross-PDE quantitative characterization of the mode-filtering effect, building on the theoretical framework of [12] and the single-method observation of [7].

Mathematical basis. As derived in Section 2.3, the Galerkin loss sees the residual only through its projection onto $\mathrm{span}\{\phi_k\}$; any $R_\perp$ in the orthogonal complement is invisible, so the optimizer can drive $\mathcal{J}_{\mathrm{G}} \to 0$ while the true error stays large—a classical inf-sup deficiency that [12] formalized via the Fortin operator. When $K < N_{\mathrm{DOF}}$, the null space has dimension at least $N_{\mathrm{DOF}} - K$ (in a linearized sense).

1D evidence. M5 (SEM-test VPINN) gets L_2 error 2.23×10^{-2} with a near-zero training loss, while least-squares methods (M1, M3) reach 10^{-5}—a factor of 10^3 gap. M4 (Legendre-test VPINN), using global test functions instead of element-local ones, lands at L_2 error 9.50×10^{-4}—better than M5 but still nearly two orders from M1. Global test functions shrink the null space but don't close it.

2D evidence. In the cavity, M5 (SEM-test VPINN) attains a near-zero Galerkin training loss (7.45×10^{-6}) while its velocity error reaches $\epsilon_u = 0.457$—roughly two orders of magnitude above the least-squares solvers (M1: 3.24×10^{-3}, M3: 1.23×10^{-2}) and far larger than its own minuscule loss implies. The same pattern appears cross-method: M5's loss is only $\sim 2.5\times$ *larger* than M1's, yet its error is $141\times$ *larger*—confirming that loss minimization and error minimization are not the same in finite test-space variational formulations.

Mitigation strategies. The observation that a small residual can mask a large error is not new in isolation [7]; what is new here is its quantitative scale and its precise attribution to the finite-dimensional Galerkin projection. Building on [12] and [11], we suggest: (i) increasing K toward or past N_{DOF} (expensive, as M4 shows); (ii) adding pointwise collocation terms to the variational loss to rein in null-space residual modes—a hybrid that blends Galerkin stability with strong-form completeness [24, 25]; this works in 1D (Section 3.1) but, as the 2D probe (Section 3.2.8) shows, only when the collocation count exceeds the null-space dimension $N_{\mathrm{DOF}} - K$; at the K=100 2D configuration that gap is $\sim$ 8,600, far beyond any practical collocation sample; (iii) using problem-adapted or hierarchical test bases that span more of the residual spectrum; (iv) always checking solution quality against independent error metrics (data points, analytical solutions, or residual norm) rather than trusting the training loss alone. *Demonstrated on a 1D prototype—and its 2D limit.* Strategy (ii) is not merely plausible in 1D—we have verified it. Adding a pointwise collocation penalty to the M6 Galerkin loss (Section 3.1) reduces the 1D mode-filtered error from $L_2 = 1.53$ to 1.37×10^{-2} ($> 100\times$), proving that strong-form constraints rein in the null-space residual modes. We ported the identical hybrid formulation to the 2D cavity at a strictly matched reduced configuration (N_{leg}=12, K=100; Section 3.2.8). The result reverses: the hybrid loss drives the training loss *down* ($1.67\times10^{-4} \rightarrow 8.43\times10^{-5}$) while the solution error goes *up* (ϵ_u: 0.694 $\rightarrow$ 1.21). The collocation penalty cannot constrain a null space of dimension $N_{\mathrm{DOF}} - K \approx$ 8,600 with only 1,600 sample points, so the mitigation that closes the 1D gap fails in 2D—the decoupling is structurally deep, not a training artefact fixable by a cheap hybrid loss. To validate the mode-filtering hypothesis quantitatively, we decompose the M5 residual spectrally for both the 1D Helmholtz and 2D cavity cases. Figure 18 plots Legendre modal coefficients of the 1D Helmholtz M5 residual: modes below the test-space cutoff are well-constrained by the Galerkin projection, while higher-order modes ($n \gtrsim K$) carry substantial residual energy but barely touch the training loss. Table 7 tabulates the visible energy fraction and the projected loss at selected K: at the actual test-space dimensions, roughly 21% (M3, $K = 30$) and 17% (M5, $K = 56$) of the residual energy lies outside the span of the test functions and is invisible to the Galerkin loss. Figure 20 extends the decomposition to the 2D lid-driven cavity (Re= 100). The M5 Navier–Stokes residual is expanded on the same tensor-product Legendre basis as the $K = 324$ (18×18,

degree-18) test space, $R(\xi,\eta) \approx \sum_{i,j} \hat{R}_{ij}\, P_i(\xi) P_j(\eta)$, with coefficients $\hat{R}_{ij}$ obtained by 100×100 Gauss–Legendre quadrature. The 2D energy spectrum (Fig. 20, left) shows significant residual energy at indices outside the $K = 324$ test-space square, and the cumulative curve (right) gives 4.6% of the residual energy invisible to the Galerkin loss at the $K = 324$ cutoff. Although smaller than the 1D fraction, this invisible energy combines with the $\approx$ 8,600-dimensional null space of the 2D test space to confirm that the loss–error decoupling is a structural property of the Galerkin projection that persists—and is amplified—in two dimensions. These decompositions confirm that the loss–error disconnect comes from the finite-dimensional Galerkin projection: the optimizer minimizes the residual's projection onto span$\{\phi_k\}$ while the orthogonal complement runs free, and the energy of those unconstrained modes dominates the total error.

Table 7: 1D Helmholtz: visible residual energy fraction $E_{\text{vis}}(K)/E_{\text{total}}$ and projected loss $\sqrt{\mathcal{L}_K}$, $\mathcal{L}_K = \sum_{k<K} |c_k|^2$, at varying test-space dimension K (quantifying the spectral curves of Figure 19). Bold rows mark the actual test-space dimensions ($K = 30$ for M3, $K = 56$ for M5). Even at these dimensions, 17–21% of the residual energy remains invisible to the Galerkin loss.

K	M3 vis. energy (%)	M3 $\sqrt{\mathcal{L}_K}$	M5 vis. energy (%)	M5 $\sqrt{\mathcal{L}_K}$
4	8.39	9.50×10^1	7.96	9.46×10^1
8	63.81	2.62×10^2	61.34	2.63×10^2
12	70.38	2.75×10^2	67.76	2.76×10^2
16	76.79	2.87×10^2	73.82	2.88×10^2
20	78.64	2.91×10^2	78.21	2.97×10^2
24	78.94	2.91×10^2	79.31	2.99×10^2
30	**79.00**	$\mathbf{2.91 \times 10^2}$	**79.74**	$\mathbf{2.99 \times 10^2}$
40	80.19	2.94×10^2	80.41	3.01×10^2
56	**83.34**	$\mathbf{2.99 \times 10^2}$	**83.11**	$\mathbf{3.06 \times 10^2}$

4.7 Limitations: single-realization errors

The Galerkin methods (M4, M5, M6) and the hybrid M7 possess a null space of dimension $N_{\text{DOF}} - K$ in which the optimizer is free to place solution error. Because of this, the realized error of a single training run is one sample from that null space and is inherently seed-sensitive. All Galerkin/MLP results in this paper are therefore reported as single realizations; the qualitative conclusions—near-zero loss with large error, and the failure of the 2D hybrid remedy—are structural and hold for any seed, but the exact error magnitudes (e.g., $\epsilon_u = 0.457$ for M5, 1.21 for M7) should be read as representative points, not tight estimates. A multi-seed study reporting mean and standard deviation is required before individual error values can be attributed at the formulation level, and is left to future work.

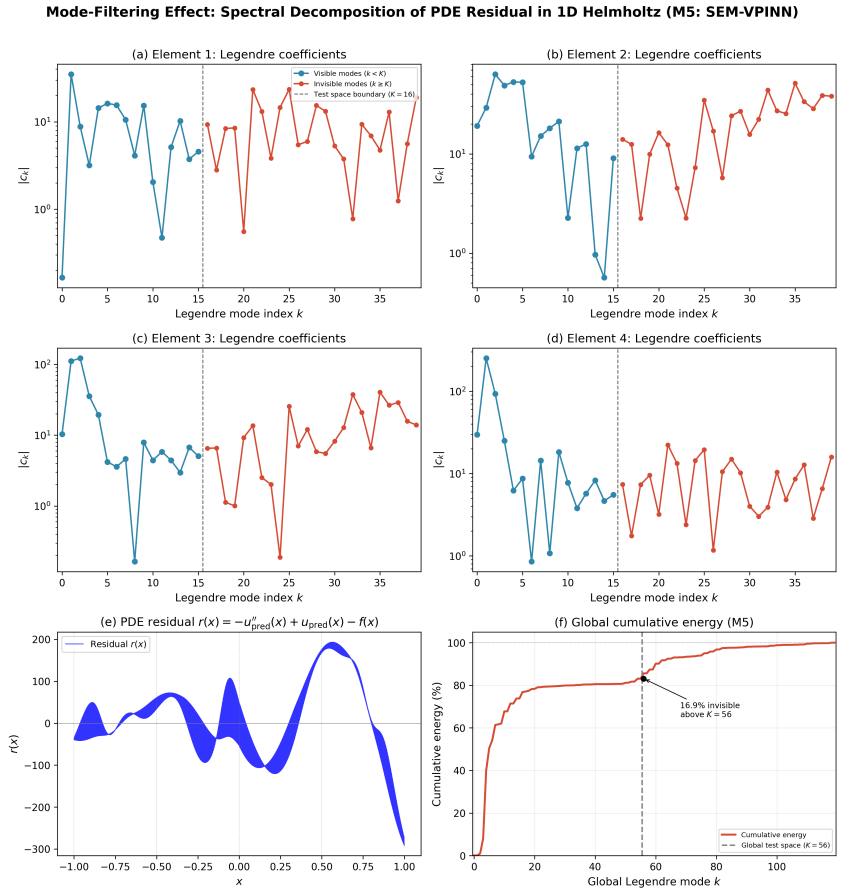


Fig. 18: 1D Helmholtz (M5, SEM-test VPINN): per-element local Legendre decomposition of the residual. Panels (a)–(d) show the Legendre modal coefficients for elements 1–4; panel (e) shows the residual $r(x)$ in physical space; panel (f) shows the global cumulative energy, with the residual energy above the global test-space cutoff ($K = 56$) invisible to the Galerkin loss. The global cumulative-energy panel (f) is quantified in Table 7.

4.8 SEM-test VPINN interface artifacts

Beyond mode filtering, M5 shows interface artifacts in velocity and pressure fields (Figures 8–10). These come from the lack of inter-element continuity constraints. Classical SEM-Galerkin methods enforce continuity through shared degrees of freedom at element boundaries. Adding analogous inter-element penalties, together with a richer test space, should close M5's accuracy gap.

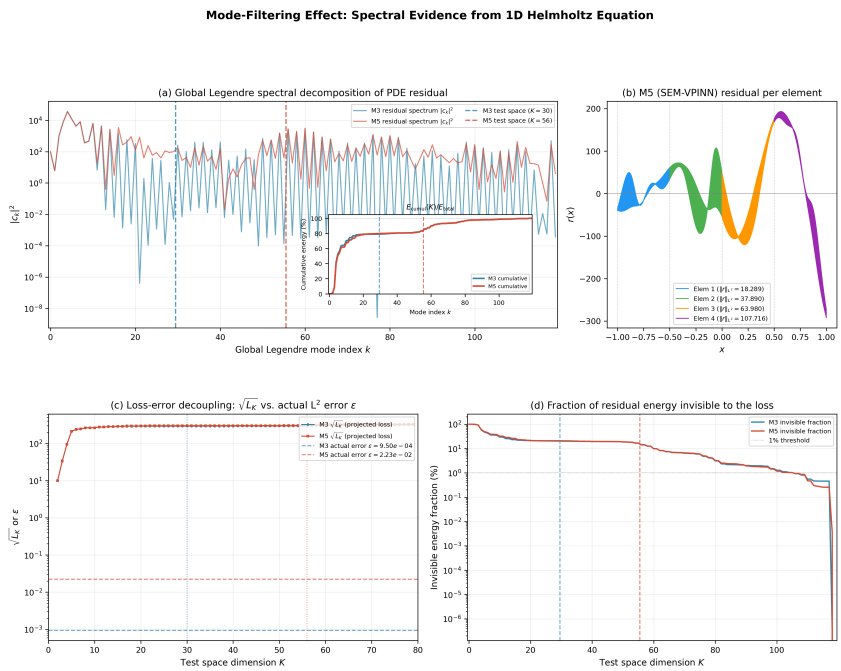


Fig. 19: 1D Helmholtz (M5): global Legendre decomposition and loss–error decoupling. The residual energy above the global test-space cutoff ($K = 56$) is invisible to the Galerkin loss; this spectral curve is quantified in Table 7.

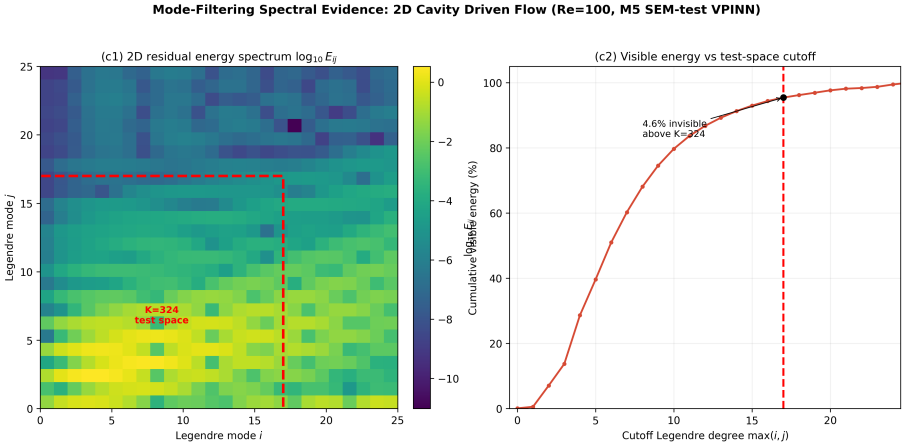


Fig. 20: 2D lid-driven cavity (Re= 100, M5): tensor-product Legendre spectrum of the Navier–Stokes residual. The 2D energy spectrum (left) shows substantial residual energy outside the $K = 324$ (18×18) test-space square; 4.6% of the residual energy is invisible to the Galerkin loss at this cutoff (right), confirming the mechanism in two dimensions.

5 Conclusions

We have compared six neural PDE solvers that span the full design space of trial and test space combinations, using the 1D Helmholtz equation and 2D steady lid-driven cavity flow at Re = 100. All methods share the same `jaxfun` implementation and a common two-phase Adam+L-BFGS training protocol, so the comparison is fair and reproducible. (Remaining differences in quadrature density and MLP capacity are discussed in Section 4.4.) Five conclusions stand out:

(1) Global spectral trial functions (M2) beat all MLP-based methods by 2–3 orders of magnitude in training loss and train 9× faster (0.54 h vs. 4.79 h for M1). This confirms the exponential convergence advantage of spectral approximations for smooth PDEs.

(2) Among MLP-based methods, pure PINN (M1) gives the best solution accuracy ($\epsilon_u = 3.24\times10^{-3}$), followed by M3 ($1.23\times10^{-2}$) and M4 ($1.68\times10^{-1}$). M5 (SEM-test VPINN) has one of the lowest training losses of any MLP-based method (7.45×10^{-6}) but the worst solution errors ($\epsilon_u = 0.457$, $\epsilon_v = 0.825$, $\epsilon_p = 0.719$).

(3) We demonstrate the mode-filtering effect quantitatively on both cases and validate it through Legendre spectral decomposition (Figures 18–20): M5 (SEM-test VPINN) achieves a near-zero Galerkin training loss on the 1D problem (L-BFGS diverges because the Galerkin loss Hessian is singular, so Adam-only is used) but its L_2 error is 2.23×10^{-2}—three orders of magnitude above the least-squares methods. On 2D, M5 attains a near-zero Galerkin loss (7.45×10^{-6}) yet $\epsilon_u = 0.457$—two orders of magnitude above the least-squares solvers—demonstrating that minimizing the projected residual does not minimize the true error. Spectral decomposition shows that the dominant residual energy lives in modes beyond the test-space cutoff—direct quantitative evidence that the loss–error disconnect comes from the finite-dimensional Galerkin projection. M6 pushes this further: its loss (1.41×10^{-7}) is 21× below M1's, but $\epsilon_u = 0.627$ is 193× above—the most extreme loss–error decoupling case in this study, driven by mode-filtering compounded with spectral basis representability limits near the cavity corners.

(4) All six methods capture the right qualitative flow structure (primary vortex, secondary corner recirculations, centerline profiles). M1–M3 and M6 agree closely on centerline velocity profiles; M5 deviates visibly, consistent with its 141× larger velocity error relative to M1.

(5) M6 (Spectral-trial/MLP-test VPINN) uses Legendre spectral trial functions with MLP test functions, trains in 0.128 h—37× faster than M1, 56× faster than M3—and reaches the lowest training loss (1.41×10^{-7}) of all Galerkin-projection methods (M4, M5, M6). But its global spectral trial basis cannot handle the corner singularities in the 2D cavity: $\epsilon_u = 0.627$, $\epsilon_p = 8.59$—the largest errors in the study. M6 is the most dramatic illustration of mode-filtering, and it shows that compact spectral coefficient parameterizations slash optimization cost at the expense of geometric flexibility.

Future work will extend this study to higher Reynolds numbers ($\mathrm{Re} \in \{400, 1000\}$), three-dimensional geometries, and time-dependent flows. We prototyped a hybrid

Galerkin–collocation remedy on the 1D Helmholtz case ($> 100\times$ error suppression) and stress-tested it on the 2D cavity: at the K=100 configuration the hybrid loss fails to transfer because the null-space dimension ($N_{\mathrm{DOF}} - K \approx 8{,}600$) dwarfs the collocation count. Effective 2D remedies therefore require either growing K toward N_{DOF} (full Galerkin, expensive) or collocation densities that scale with the null-space dimension—both costly and left to future work. A capacity ablation quantifies the MLP-size effect at $\approx 7.5\times$. Adding proper inter-element continuity constraints to the SEM-test VPINN framework should close the M5 accuracy gap. Hybrid trial/test formulations—spectral trial functions in smooth subdomains paired with MLP functions near complex boundaries, an approach informed by the M6 results—are a promising path toward extending M2's benefits to general geometries.

Declarations

Funding This work was supported by the Basic Scientific Research Fund of Hebei University of Water Resources and Electric Engineering (Grant No. SYKY2103).

Competing Interests The author has no relevant financial or non-financial interests to disclose.

Ethics Approval Not applicable.

Consent to Participate Not applicable.

Consent for Publication Not applicable.

Data Availability The datasets generated and/or analysed during the current study are available in the SEM-PINN repository, https://github.com/gengyanhui111-boop/SEM-PINN.

Code Availability The source code used to generate and analyse these datasets is available in the same repository.

References

[1] Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics **378**, 686–707 (2019) https://doi.org/10.1016/j.jcp.2018.10.045

[2] Cai, S., Mao, Z., Wang, Z., Yin, M., Karniadakis, G.E.: Physics-informed neural networks (pinns) for fluid mechanics: a review. Acta Mechanica Sinica **37**(12), 1727–1738 (2021) https://doi.org/10.1007/s10409-021-01148-1

[3] Jin, X., Cai, S., Li, H., Karniadakis, G.E.: NSFnets (navier–stokes flow nets): physics-informed neural networks for the incompressible navier–stokes equations. Journal of Computational Physics **426**, 109951 (2021) https://doi.org/10.1016/j.jcp.2020.109951

[4] Haghighat, E., Raissi, M., Moure, A., Gomez, H., Juanes, R.: A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics.

Computer Methods in Applied Mechanics and Engineering **379**, 113741 (2021) https://doi.org/10.1016/j.cma.2021.113741

[5] Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: Proceedings of the 36th International Conference on Machine Learning (ICML), vol. 97, pp. 5301–5310 (2019)

[6] Wang, S., Yu, X., Perdikaris, P.: When and why pinns fail to train: a neural tangent kernel perspective. Journal of Computational Physics **449**, 110768 (2022) https://doi.org/10.1016/j.jcp.2021.110768

[7] Krishnapriyan, A.S., Gholami, A., Zhe, S., Kirby, R.M., Mahoney, M.W.: Characterizing possible failure modes in physics-informed neural networks. In: Advances in Neural Information Processing Systems 34 (NeurIPS 2021), pp. 26548–26560 (2021)

[8] Kharazmi, E., Zhang, Z., Karniadakis, G.E.: hp-VPINNs: variational physics-informed neural networks with domain decomposition. Computer Methods in Applied Mechanics and Engineering **374**, 113547 (2021) https://doi.org/10.1016/j.cma.2020.113547

[9] Peng, A., Pratheek, I., Bhattarai, R., Simeoni, C.: FastVPINNs: tensor-driven acceleration of vpinns for complex geometries. SIAM Journal on Scientific Computing (2025) https://doi.org/10.1137/24M1658620 . in press

[10] Anandh, T., Ghose, D., Tyagi, A., Gupta, A., Sarkar, S., Ganesan, S.: An efficient hp-variational pinns framework for incompressible navier–stokes equations. arXiv preprint arXiv:2409.04143 (2024). arXiv:2409.04143 [physics.flu-dyn]

[11] Berrone, S., Canuto, C., Pintore, M.: Variational physics-informed neural networks: the role of quadratures and test functions. Journal of Scientific Computing **92**(3), 100 (2022) https://doi.org/10.1007/s10915-022-01950-4

[12] Rojas, P., Calderón, J., Pardo, D., Paszynski, M.: Robust variational physics-informed neural networks. Computer Methods in Applied Mechanics and Engineering **425**, 116904 (2024) https://doi.org/10.1016/j.cma.2024.116904

[13] Li, Y., Wang, H., Zhang, Q.: Fourth- and higher-order finite element methods for the incompressible Navier–Stokes equations with Dirichlet boundary conditions. International Journal for Numerical Methods in Engineering **126**(17), 70102 (2025) https://doi.org/10.1002/nme.70102

[14] Pacheco, D.R.Q., Schussnig, R., Steinbach, O., Fries, T.-P.: A global residual-based stabilization for equal-order finite element approximations of incompressible flows. International Journal for Numerical Methods in Engineering **122**(8), 2075–2094 (2021) https://doi.org/10.1002/nme.6615

[15] Gaynutdinova, L., Rokoš, O., Doškář, M., Pultarová, I.: Homogenization with guaranteed bounds via primal-dual physically informed neural networks. International Journal for Numerical Methods in Engineering **127**(6), 70302 (2026) https://doi.org/10.1002/nme.70302

[16] jaxfun contributors: jaxfun: A JAX library for spectral PDE solvers and physics-informed neural networks, v0.1.0 (2024). https://github.com/jaxfun/jaxfun

[17] Shen, J., Tang, T., Wang, L.-L.: Spectral Methods: Algorithms, Analysis and Applications. Springer, Berlin (2011). https://doi.org/10.1007/978-3-540-71041-7

[18] Karniadakis, G.E., Sherwin, S.J.: Spectral/*hp* Element Methods for Computational Fluid Dynamics, 2nd edn. Oxford University Press, Oxford (2005)

[19] Ghia, U., Ghia, K.N., Shin, C.T.: High-re solutions for incompressible flow using the navier–stokes equations and a multigrid method. Journal of Computational Physics **48**(3), 387–411 (1982) https://doi.org/10.1016/0021-9991(82)90058-4

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

[21] Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. Published at the 3rd International Conference on Learning Representations (ICLR) (2015)

[22] Liu, D.C., Nocedal, J.: On the limited memory BFGS method for large scale optimization. Mathematical Programming **45**(1), 503–528 (1989) https://doi.org/10.1007/BF01589116

[23] Botella, O., Peyret, R.: Benchmark spectral results on the lid-driven cavity flow. Computers & Fluids **27**(4), 421–433 (1998) https://doi.org/10.1016/S0045-7930(98)00002-4

[24] Lagaris, I.E., Likas, A., Fotiadis, D.I.: Artificial neural networks for solving ordinary and partial differential equations. IEEE Transactions on Neural Networks **9**(5), 987–1000 (1998) https://doi.org/10.1109/72.712178

[25] Lu, L., Meng, X., Mao, Z., Karniadakis, G.E.: DeepXDE: a deep learning library for solving differential equations. SIAM Review **63**(1), 208–228 (2021) https://doi.org/10.1137/19M1274067