RAA-PINN: RESIDUAL ANTI-ALIASING FOR PHYSICS-INFORMED NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Physics-informed neural networks (PINNs) enforce differential equations at finitely many collocation points, but small residuals at those points can conceal unresolved errors between them. We study this discrepancy as residual aliasing, focusing on high-frequency structure in the PDE residual that pointwise training may miss. We propose RAA-PINN, a frequency-aware framework that couples residual spectral regularization with spectrum-aware adaptive sampling. The regularizer penalizes high-frequency residual energy rather than the solution spectrum, while the sampler combines residual magnitude with local spectral energy to allocate new collocation points. Across five PDE benchmarks covering Allen–Cahn, convection–diffusion, and Kovasznay flow, RAA-PINN achieves the lowest raw high-frequency residual energy among the evaluated methods under both full-core and equal-time settings. A strict spectral guard admits RAA-PINN candidates to auxiliary-metric comparisons only when their raw high-frequency residual energy is lower than that of the strongest applicable non-RAA baseline. Under this protocol, the selected candidates also achieve the lowest relative L_2 error, worst-region error, and dense residual on all five tasks. Ablation studies examine the contributions of spectral regularization, adaptive sampling, and guarded refinement. Independent two- and three-dimensional FEM/COMSOL reference comparisons further show lower aggregate solution errors. These results support residual-spectrum control as an effective way to improve physics-informed learning beyond pointwise residual minimization.

1 INTRODUCTION

Partial differential equations (PDEs) underpin models in fluid mechanics, heat transfer, reaction–diffusion systems, electromagnetics, and materials science. Classical finite-element and spectral discretizations remain standard numerical tools Trefethen (2000); Fornberg (1998); Hughes (2000), while deep solvers represent solutions or operators with trainable models Lagaris et al. (1998); Sirignano & Spiliopoulos (2018); E & Yu (2018); Han et al. (2018); Lu et al. (2021a); Li et al. (2021); Kovachki et al. (2023). Physics-informed neural networks (PINNs) directly encode PDE residuals and boundary or initial conditions in the training objective, enabling mesh-free forward and inverse modeling with limited supervision Raissi et al. (2019); Karniadakis et al. (2021); libraries such as DeepXDE have made this formulation broadly reusable Lu et al. (2021b).

A small empirical residual does not by itself determine the residual over the domain or the error relative to a reference solution. PINN training depends on the interaction between the differential operator, the loss terms, optimization, and the distribution of collocation points. Previous studies identify loss imbalance, gradient and neural-tangent-kernel effects, temporal ordering, and ill-conditioned optimization as sources of training difficulty Wang et al. (2021; 2022; 2024); Rathore et al. (2024). The distribution of residual evaluations raises a related question: which features of the residual field are represented by the points used during training?

We focus on residual structure that is poorly represented by finite collocation. Oscillations or localized variations may occur between sampled points even when the residual is small at those points. Figure 1 illustrates this sampling concern, motivated by frequency-resolution considerations in spectral and collocation methods. Here the frequency content belongs to the PDE residual, rather than to

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

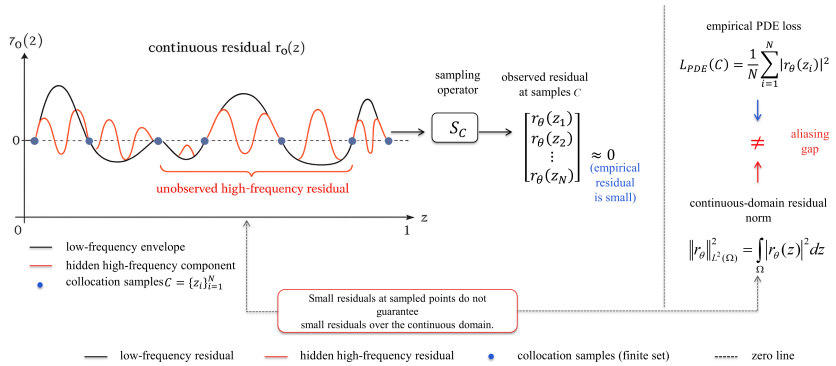


Figure 1: Conceptual illustration of residual structure between collocation points. Small residuals at sampled locations can coexist with larger variations between samples, motivating residual spectral probing.

the predicted solution itself. Spectral probes provide an additional view of this residual field over a frequency band represented by the probe grid.

Existing approaches improve collocation through residual-based sampling, adjust the relative weights of loss terms, or impose additional residual and weak-form constraints. Frequency-aware representations address how the network expresses a solution, while region-based training broadens the locations used to enforce the equation. These approaches motivate the use of information beyond an isolated residual value. Our interest is in using residual-frequency information both to define a training penalty and to allocate new collocation points.

RAA-PINN combines these two uses of the residual spectrum. A spectral penalty acts on a selected high-frequency band, and a local spectral score supplements residual magnitude when ranking candidate points. The two components can be used separately or together, which we examine through the A0–A5 ablation series. The experiments compare residual-band energy, solution error, and spatial error distributions on five benchmark cases, followed by comparisons with two- and three-dimensional FEM reference fields. The central contribution is the joint use of residual spectral information in regularization and adaptive collocation.

2 Related Work

2.1 PINN Training Pathologies and Residual Generalization

PINNs combine differential-equation residuals with boundary, initial, or observational constraints. Studies of their training behavior have examined gradient imbalance, neural-tangent-kernel effects, causality, and optimization landscapes (Wang et al., 2021; 2022; 2024; Rathore et al., 2024). These findings show that the empirical loss depends on both the physical problem and the training procedure. Related methods address where the constraints are evaluated: RoPINN extends pointwise optimization to local regions (Wu et al., 2024), while adversarial adaptive sampling uses optimal transport to construct a sampling distribution (Tang et al., 2024). RAA-PINN also changes the information used during training, but uses a residual-frequency band in the regularizer and local residual spectral energy in the collocation score.

2.2 Adaptive Sampling and Adaptive Weighting

Residual-based adaptive refinement (RAR) adds collocation points in regions with large PDE residuals. Its greedy form, RAR-G, ranks a candidate pool by residual magnitude and selects the highest-ranked points. Residual-based adaptive distribution (RAD) instead resamples points from a residual-dependent probability distribution (Wu et al., 2023). Self-adaptive weighting changes the contribution of individual training locations, whereas adaptive sampling changes the locations themselves.

RAA uses a candidate-ranking rule that combines residual magnitude with local high-frequency residual energy. Candidates with similar pointwise residuals can therefore receive different scores when the surrounding residual fields have different frequency content.

2.3 Spectral Representations, Residual Regularization, and Benchmarks

Fourier features modify the representation used to approximate high-frequency functions. Residual regularization instead acts on the governing equation evaluated at the network output. Gradient-enhanced PINNs add derivatives of the PDE residual to the loss and have also been combined with RAR (Yu et al., 2022). Weak-form methods constrain integrals of the residual, and regional methods extend pointwise training to local neighborhoods. These approaches differ in which properties of the residual enter the objective.

Residual-gradient and spectral penalties are closely related. For a sufficiently smooth periodic residual, Parseval's identity expresses its integrated squared gradient as frequency-squared-weighted spectral energy. This relation motivates a direct comparison of the frequency weighting, the selected band, and the sampling locations. RAA uses a band-restricted residual penalty together with a local spectral contribution to the collocation score. Residual-spectrum diagnostics have also been studied by Baheri (2026). The present work focuses on using such information in training and point selection. Since the penalty is applied to the differential residual, it does not directly penalize the frequency content of the predicted solution.

3 Residual Aliasing and RAA-PINN

3.1 Problem Formulation

Let z denote the input coordinate, $u_\theta(z)$ the neural approximation, $\mathcal{N}[\cdot]$ the differential operator, and $f(z)$ the forcing or source term. A standard PINN obtains the derivatives of u_θ by automatic differentiation and minimizes a weighted combination of the PDE, boundary, and initial-condition losses Raissi et al. (2019); Lu et al. (2021b); Lagaris et al. (1998); Baydin et al. (2018):

$$r_\theta(z) = \mathcal{N}[u_\theta](z) - f(z), \tag{1}$$

$$\mathcal{L}_{\mathrm{PINN}} = \mathcal{L}_{\mathrm{PDE}} + \lambda_{\mathrm{BC}}\mathcal{L}_{\mathrm{BC}} + \lambda_{\mathrm{IC}}\mathcal{L}_{\mathrm{IC}}. \tag{2}$$

We distinguish the collocation set $\mathcal{C}$, the spectral probe set $\mathcal{P}$, and the dense evaluation set $\mathcal{D}$. The collocation set defines the pointwise PDE loss, while probe evaluations supply the residual spectrum used during training. The training and dense residual estimates are

$$R_{\mathrm{train}} = \frac{1}{|\mathcal{C}|}\sum_{z_i\in\mathcal{C}} |r_\theta(z_i)|^2, \quad R_{\mathrm{dense}} = \frac{1}{|\mathcal{D}|}\sum_{z_j\in\mathcal{D}} |r_\theta(z_j)|^2. \tag{3}$$

The two estimates describe residual magnitude over different sets of coordinates. A small value on $\mathcal{C}$ can coexist with larger residuals elsewhere, for example when the sampling pattern misses a narrow peak or poorly represents an oscillation. Dense evaluation measures the residual on $\mathcal{D}$, while the spectral statistic below describes its energy in a selected frequency band. The probe set is part of the training procedure: its residual evaluations contribute to the spectral loss rather than serving solely as a diagnostic. Candidate selection and the scope of the reported evaluations are described in Section 3.4.

3.2 Raw HFRE and Residual Anti-Aliasing Loss

Raw high-frequency residual energy (raw HFRE) summarizes a selected band of the PDE-residual spectrum. Given the residual field evaluated on $\mathcal{P}$, Eq. (4) sums the spectral energy above the cutoff ω_c.

$$E_{\mathrm{HF}}^{\mathrm{raw}} = \sum_{\omega>\omega_c} |\widehat{r}_\theta(\omega)|^2. \tag{4}$$

The frequency coordinate is expressed as a normalized radial distance from zero, with the cutoff given relative to the Nyquist radius. The same task-level convention is used for the methods within a comparison. The term "raw" means that the tail energy is not divided by the total residual energy.

This distinction matters because a normalized tail fraction can decrease when low-frequency energy increases, even if the high-frequency energy stays unchanged. Raw HFRE describes the energy of the selected tail under the chosen transform convention. The frequencies represented by this calculation are determined by the probe resolution.

The spectral loss in Eq. (5) uses the same residual-frequency band, with weights $w(\omega)$ determining the contribution of each component.

$$\mathcal{L}_{\mathrm{AA}} = \sum_{\omega > \omega_c} w(\omega) |\widehat{r}_\theta(\omega)|^2, \qquad \mathcal{L} = \mathcal{L}_{\mathrm{PINN}} + \lambda_{\mathrm{AA}} \mathcal{L}_{\mathrm{AA}}. \tag{5}$$

The objective combines the pointwise PINN loss with a frequency-weighted residual penalty. Raw HFRE is therefore a diagnostic of the band targeted by training; solution error and dense residual are reported separately. At the continuous equation level, a smooth exact solution has zero PDE residual and hence zero residual spectral penalty. The regularizer acts on the differential residual rather than directly on the solution representation. Fourier features and the spectral residual penalty address these different parts of the approximation process. The relation to residual-gradient regularization is discussed in Section 2.3.

For a sufficiently smooth periodic scalar residual on Ω, consider an orthonormal Fourier expansion $r = \sum_{\boldsymbol{\kappa}} \widehat{r}_{\boldsymbol{\kappa}} \phi_{\boldsymbol{\kappa}}$, with physical wavevectors satisfying $\nabla \phi_{\boldsymbol{\kappa}} = i\boldsymbol{\kappa} \phi_{\boldsymbol{\kappa}}$. Parseval's identity gives

$$\|\nabla r\|^2_{L^2(\Omega)} = \sum_{\boldsymbol{\kappa}} \|\boldsymbol{\kappa}\|^2 |\widehat{r}_{\boldsymbol{\kappa}}|^2.$$

The gradient penalty thus weights the full residual spectrum by squared frequency, whereas Eq. 5 restricts the penalty to a selected band. The band and its weights determine which residual components enter the additional objective. The local spectral score in Section 3.3 uses residual-frequency information for a different operation: choosing new collocation points.

3.3 Spectrum-Aware Adaptive Sampling

The sampler combines pointwise residual magnitude with local spectral energy. For a candidate $z \in \mathcal{Q}$, write $a(z) = |r_\theta(z)|$ and $b(z) = E^{\mathrm{local}}_{\mathrm{HF}}(z)$. Each signal is min–max normalized over the candidate pool before the weighted score is formed:

$$s(z) = \alpha \widetilde{a}(z) + \beta \widetilde{b}(z), \qquad \mathcal{C}_{\mathrm{new}} = \text{top-k}_{z \in \mathcal{Q}}\, s(z). \tag{6}$$

The selected points are added through $\mathcal{C} \leftarrow \mathcal{C} \cup \mathcal{C}_{\mathrm{new}}$. Local spectral energy is obtained by evaluating the residual on a small regular window centered at a candidate, applying a multidimensional FFT, and summing energy above the radial cutoff.

The pointwise term ranks candidates by the residual already observed at their coordinates. The spectral term uses the residual in a surrounding window, so two points with similar residual magnitudes can receive different scores. A large residual remains influential even when its local high-frequency contribution is small. The coefficients α and β control the relative weights of the two normalized signals, and k determines the number of points added at each update. The candidate-pool size and addition schedule are listed in Table 1.

The spectral penalty and the sampling score play different roles. The penalty changes the objective evaluated during an optimization step, whereas the score changes the locations included in subsequent pointwise training. Section 4.4 compares each component with their combination and with the final refinement stage.

3.4 Training Profiles and HFRE Guard

We report a full-core profile and a reduced-probe, fixed-epoch profile. Both use the architecture and 3000 Adam-stage epochs in Table 1. The full-core profile uses 32 probe points per coordinate and a candidate pool of 4096, adding 128 points every 150 epochs. The reduced-probe profile uses 24 probe points per coordinate and 3072 candidates, adding 96 points every 100 epochs. Its final L-BFGS refinement is shortened or omitted. The label "equal-time" in the existing displays refers to this reduced-probe profile; the comparison fixes the stated settings rather than a measured common wall-clock duration.

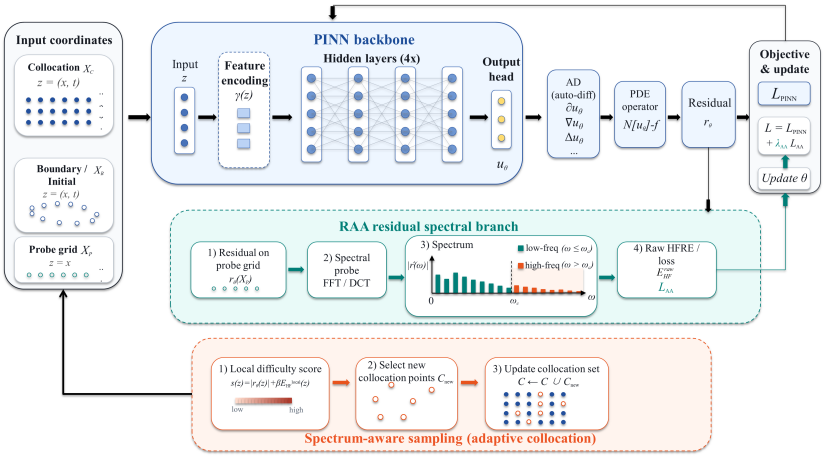


Figure 2: RAA-PINN training procedure. The residual spectral branch supplies the training penalty, while residual magnitude and local spectral energy determine the adaptive collocation score.

For the guarded comparisons, an RAA candidate is eligible when its raw HFRE is below the reference value from the applicable non-RAA methods for the same task. Supplementary Table S3 reports 3, 19, 1, 8, and 9 eligible candidates for AC, CD-const, CD-rot, CD-vort, and Kovasznay, respectively. These are counts of eligible candidates, rather than the total number of runs. The guard specifies this eligibility condition; the retained summaries do not specify the final within-set selection rule or checkpoint correspondence across metrics. Figure 4 is therefore presented as a selected-candidate comparison. Appendix G records the available configuration and reporting information.

4 Experiments

The experiments examine residual spectral energy, solution error, and the effects of the training components. Field comparisons and FEM reference evaluations provide spatial views of the prediction errors.

4.1 Experimental Setup

The five main cases comprise one space–time Allen–Cahn problem, three two-dimensional steady convection–diffusion problems with different velocity fields, and two-dimensional steady Kovasznay flow at $\mathrm{Re} = 40$. The coordinate pair for Allen–Cahn is (x, t); the scalar convection–diffusion cases use (x, y). Table 1 specifies the equations and their domains. The common network has five hidden layers of width 96, and the principal result tables summarize seeds $\{0, 1, 2\}$.

The comparison includes PINN, Fourier features, RAR, RAD, GradPenalty, gPINN, Causal PINN, and PINNacle-like sampling, where applicable. Raw HFRE is compared within each task. Relative L_2 error describes discrepancy from the reference solution, dense residual is the squared-residual mean in Eq. (3), and the field summaries report mean error, upper-tail error, and pointwise comparisons. Training profiles and candidate eligibility are described in Section 3.4. Detailed result tables appear in Appendix A.

4.2 Main Results under Strong Baselines

Table 2 reports raw HFRE for the reduced-probe, fixed-epoch profile. RAR has the lowest non-RAA values on AC, CD-rot, and CD-vort, while RAD is the closest comparator on CD-const and Kovasznay. Relative to these methods, the selected RAA entries are lower by approximately 4.3%, 3.5%, 5.2%, 33.3%, and 16.6%, respectively. The larger differences occur on CD-vort and Kovasznay, while the remaining cases have smaller margins.

Item	Setting
Allen–Cahn	$u_t - 0.01u_{xx} + u^3 - u = f$, $(x,t) \in [0,1] \times [0,1]$.
Convection–diffusion	$-0.01\Delta u + \boldsymbol{\beta} \cdot \nabla u = f$ on $[0,1]^2$. $\boldsymbol{\beta}_c = (1, 0.35)$; $\boldsymbol{\beta}_r = (-(y-0.5), x-0.5)$; $\boldsymbol{\beta}_v = (\sin \pi x \cos \pi y, -\cos \pi x \sin \pi y)$.
Kovasznay	Steady incompressible flow at $\mathrm{Re} = 40$ on $[-0.5, 1] \times [-0.5, 1.5]$.
Reference data	Forcing and boundary/initial data follow the manufactured or analytic solutions.
Network	Five hidden layers of width 96; common initial task samples.
Metrics	Raw HFRE; relative L_2; worst-region error; dense residual.
Full-core	3000 Adam-stage epochs; initial $N_c/N_b = 1024/256$; 32 probe points per coordinate; 4096 candidates; add 128 points every 150 epochs.
Reduced-probe	Same network, Adam-stage epochs, and initial N_c/N_b; 24 probe points per coordinate; 3072 candidates; add 96 points every 100 epochs; fast/prefiltered probes and short or omitted L-BFGS refinement.
Baselines	PINN, Fourier, RAR, RAD, GradPenalty, gPINN, Causal PINN, PINNacle-like.
Repeats	Seeds $\{0, 1, 2\}$; seed-aggregated values in the original result tables.

Table 1: Experimental settings. The reduced-probe label refers to the fixed-epoch profile originally labelled equal-time. Candidate eligibility is defined in Section 3.4.

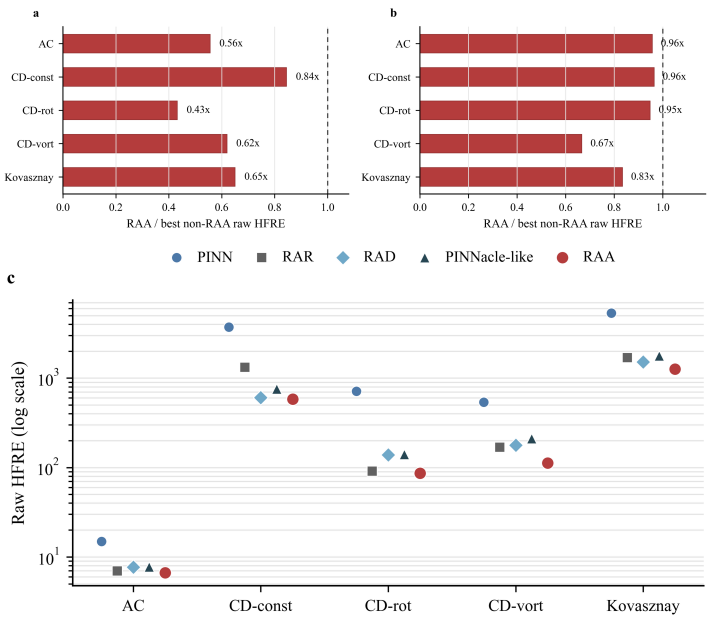


Figure 3: Raw HFRE under the full-core and reduced-probe profiles. Panels (a,b) show ratios to the lowest applicable non-RAA value within each task; panel (c) displays the reduced-probe comparison for the methods shown. The original label "equal-time" denotes the fixed-epoch profile described in Section 3.4. Table 2 gives the complete reduced-probe baseline matrix.

The full-core results in Table S1 give raw HFRE values of 3.28 for RAA and 5.89 for RAR on AC, and 60.4 versus 140 on CD-rot. These correspond to reductions of approximately 44.3% and 56.9%, respectively. The margins on these two tasks are larger than those in the reduced-probe comparison. Table 1 lists the probing and refinement settings associated with each profile.

Figure 3 presents the task-wise ratios and a logarithmic view of the displayed methods. The full-core baseline values are given separately in Table S1. Comparing methods within each task preserves the task-specific scale of raw spectral energy. RAR and RAD remain the closest competitors across most cases, with their ordering changing across the convection–diffusion fields. The lower-HFRE rankings summarize the selected entries under these training profiles; conventional solution errors are examined next.

Method	AC	CD-const	CD-rot	CD-vort	Kov.	Wins
PINN	14.920	3733.3	714.2	539.1	5345.3	—
Fourier	1.19×10^{10}	3.58×10^{10}	3.30×10^{10}	3.34×10^{10}	3.98×10^{10}	—
RAR	6.965	1320.5	91.079	168.1	1696.9	—
RAD	7.684	603.9	138.8	177.7	1511.8	—
GradPenalty	21164.0	38979.8	42393.1	4184.3	19583.0	—
gPINN	38735.0	54023.8	62307.5	5427.0	29719.6	—
Causal PINN	1730.5	—	—	—	—	—
PINNacle-like	7.740	758.6	140.7	210.3	1770.7	—
RAA (selected)	6.664	582.5	86.336	112.1	1260.9	5/5

Table 2: Raw HFRE for the reduced-probe, fixed-epoch profile. Lower values indicate lower residual-band energy. Candidate reporting is described in Section 3.4.

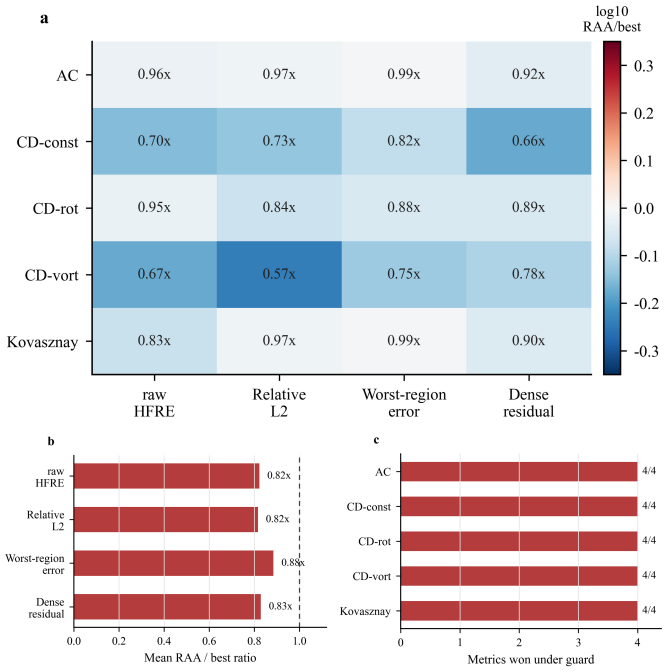


Figure 4: HFRE-guarded candidate summaries. Each ratio compares a selected RAA result with the applicable non-RAA reference for that task and metric. Ratios below one indicate a lower reported value. Panel (c) counts metrics below one, rather than statistically significant wins. Candidate eligibility and reporting scope are described in Section 3.4 and Appendix G.

4.3 Guarded Auxiliary Metrics and Field-Level Evidence

Figure 4 summarizes the HFRE-guarded comparisons in Section 3.4. For CD-vort, the ratios are 0.573 for relative L_2 error, 0.747 for worst-region error, and 0.779 for dense residual. The relative L_2 ratios for AC and Kovasznay are 0.966 and 0.968, respectively. The solution-error margins vary across tasks even when the HFRE ranking is unchanged.

Table S2 gives absolute values for the separate full-core summary. Its RAA relative L_2 entries are 9.15×10^{-4}, 0.00442, 0.00880, 0.00862, and 0.00342 for AC, CD-const, CD-rot, CD-vort, and Kovasznay; the corresponding non-RAA references are 9.98×10^{-4}, 0.00575, 0.0119, 0.00999, and 0.00684. The RAA dense residuals are 6.97×10^{-6}, 3.78×10^{-4}, 3.64×10^{-5}, 5.50×10^{-5}, and 2.95×10^{-4}. These numbers give the error scale of the full-core summary, while Figure 4 presents the guarded ratio series.

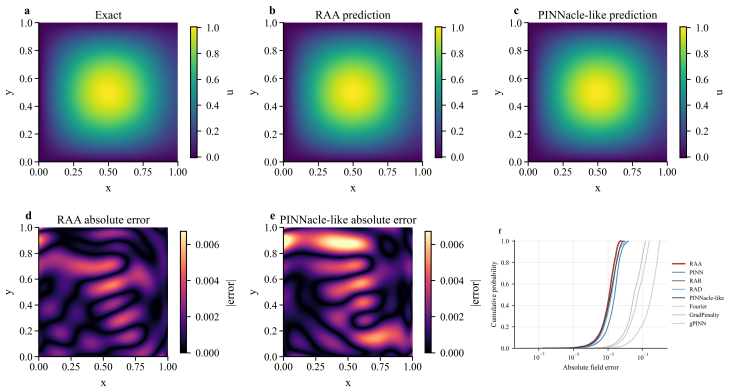


Figure 5: CD-const predictions and field errors. The panels show the reference field, RAA and PINNacle-like predictions, absolute-error maps, and cumulative distributions of absolute error for the displayed methods.

Figure 5 compares the CD-const fields, absolute-error maps, and error CDFs. Table S4 gives PINNacle-like/RAA ratios of 1.31 for mean error and 1.43 for P95 error. Expressed relative to PINNacle-like, these correspond to reductions of approximately 23.7% and 30.1%. Against RAR and RAD, the mean- and P95-error ratios of 1.43 and 1.63 correspond to reductions of 30.1% and 38.7%. The RAA pointwise win fractions are 54% against PINNacle-like and 59% against RAR and RAD. The error ratios describe the magnitude of the differences, whereas the win fraction counts the locations with lower error. Together, they distinguish changes in error magnitude from the spatial extent of improvement.

4.4 Ablation Study

Table 3 compares six configurations. A0 is PINN; A1 adds RAR; A2 uses the spectral loss; A3 uses RAA sampling; A4 combines the spectral and sampling terms; and A5 includes final refinement. Before refinement, A4 has lower raw HFRE than A1 on AC, CD-const, CD-vort, and Kovasznay by approximately 25.7%, 43.2%, 3.0%, and 18.8%, respectively. On CD-rot, A4 is 10.6% higher. This comparison shows where the combined configuration differs from magnitude-based refinement.

The individual additions also have different effects. Sampling alone changes the CD-const value from 718.4 in A0 to 292.5 in A3. Adding the spectral term changes A3 to A4 and lowers the value on four tasks, whereas CD-rot increases from 38.35 to 41.30. A5 has the lowest value in the ablation series on every task. Relative to A4, final refinement reduces HFRE by approximately 53.7%, 75.0%, 78.6%, 69.6%, and 64.0% across the five cases. Both the preceding training configuration and the refinement stage contribute to the reported final values. The spectral contribution varies across tasks, and final refinement produces a substantial additional decrease.

The spectral-only A2 configuration lowers raw HFRE from 5.549 to 4.145 on AC and from 1740.0 to 1556.7 on Kovasznay, while CD-vort increases from 155.7 to 169.9.

Appendix D reports sampling-rule, parameter, precision, and sensitivity checks over the tested settings.

4.5 External FEM/COMSOL Validation

We compare the neural predictions with FEM reference fields for a two-dimensional CD-const problem and a three-dimensional manufactured convection–diffusion problem. The three-dimensional equation is

$$-\varepsilon\,(u_{xx} + u_{yy} + u_{zz}) + \boldsymbol{\beta} \cdot \nabla u = f(x, y, z). \tag{7}$$

Predictions and the reference field are evaluated on a common volume grid, and the displayed slices visualize the three-dimensional field.

ID	Variant	Spectral	RAA sampling	Final refine.	AC	CD-c	CD-r	CD-v	Kov.
A0	PINN	no	no	no	5.549	718.4	168.7	155.7	1740.0
A1	PINN + RAR	no	no	no	2.540	495.9	37.35	64.86	251.8
A2	Spectral	yes	no	no	4.145	673.9	130.9	169.9	1556.7
A3	RAA sampling	no	yes	no	2.233	292.5	38.35	63.85	303.5
A4	Spectral + sampling	yes	yes	no	1.887	281.9	41.30	62.90	204.4
A5	**Full RAA-PINN**	**yes**	**yes**	**yes**	**0.873**	**70.45**	**8.837**	**19.15**	**73.49**

Table 3: Component configurations and raw HFRE. The sampling column indicates RAA sampling; A1 instead uses RAR. A5 includes final refinement. CD-c/r/v denote CD-const/rot/vort, and Kov. denotes Kovasznay.

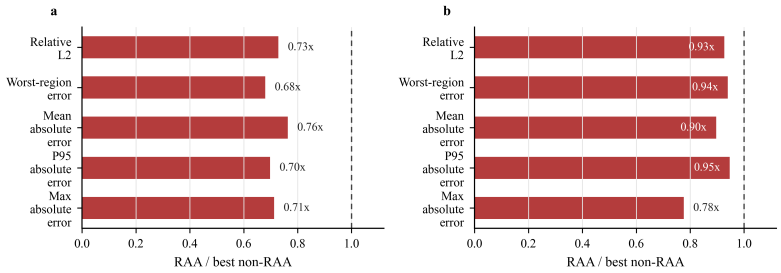


Figure 6: Discrepancies against FEM reference fields in the two-dimensional and three-dimensional manufactured cases. Ratios compare RAA with the lowest applicable non-RAA value for each metric. Absolute metric values are given in Tables S7 and S8.

In the two-dimensional comparison, Table S7 gives relative L_2 discrepancies of 0.00354 for RAA and 0.00486 for PINNacle-like sampling, a difference of approximately 27.2%. The mean absolute discrepancies are 0.00142 and 0.00187, and the maximum absolute discrepancies are 0.00609 and 0.00855. Figure 6 summarizes the error ratios, while Appendix E provides the field maps and CDF.

The three-dimensional comparison in Table S8 gives relative L_2 discrepancies of 0.0458 for RAA and 0.0494 for PINNacle-like, a difference of 7.3%. Maximum absolute discrepancy decreases from 0.0781 to 0.0607, or approximately 22.3%. The difference is larger for the maximum local error than for the global relative L_2 measure. These errors are measured against the FEM reference fields used in the two cases.

The P95 absolute discrepancies provide an additional comparison of the error distributions. In two dimensions, RAA reduces this value from 0.00509 to 0.00355 relative to PINNacle-like, a reduction of approximately 30.3%. In three dimensions, the corresponding values are 0.0367 and 0.0347, a reduction of 5.4%. The three-dimensional P95 change is smaller than the 22.3% reduction in maximum absolute discrepancy, distinguishing the change at the 95th percentile from that at the largest observed error.

5 Discussion and Conclusion

RAA-PINN uses residual-frequency information in both a spectral training penalty and an adaptive collocation score. The penalty acts on a selected residual band, while the score combines pointwise residual magnitude with local spectral energy to choose collocation points.

Across five benchmarks, the combined configuration has lower HFRE than RAR on four cases before refinement, while CD-rot favors RAR; final refinement gives the lowest ablation HFRE. Solution-error summaries and the two FEM-reference comparisons complement the spectral evaluation with field-level discrepancies.

Reproducibility Statement

The full-core and guarded benchmark tables, together with the field diagnostics, are documented in Appendices A–C. Ablation, sampling-sensitivity, precision, and two- and three-dimensional FEM/COMSOL comparisons are provided in Appendices D–E, with the Kovasznay diagnostics in Appendix F. Appendix G records the implementation settings and reporting limits. The anonymous code package supports a CPU smoke-training run and post-run metric aggregation and plotting from summary outputs; paper-scale checkpoints, raw training data, and licensed COMSOL resources are not included. The package README gives the tested commands, inputs, outputs, and external-resource requirements.

AI Use Statement

OpenAI Codex was used solely to assist with LaTeX typesetting and document formatting.

References

Ali Baheri. What does a neural pde solver really learn? a residual-spectrum diagnostic. In *AI and Partial Differential Equations Workshop at ICLR*, 2026. URL https://openreview.net/forum?id=e49Aef9YtL. Workshop paper.

Atilim Gunes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, 18(153):1–43, 2018. URL https://www.jmlr.org/papers/v18/17-468.html.

Weinan E and Bing Yu. The deep ritz method: A deep learning-based numerical algorithm for solving variational problems. *Communications in Mathematics and Statistics*, 6:1–12, 2018. doi: 10.1007/s40304-018-0127-z.

Bengt Fornberg. *A Practical Guide to Pseudospectral Methods*. Cambridge University Press, 1998. doi: 10.1017/CBO9780511626357.

Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018. doi: 10.1073/pnas.1718942115.

Thomas J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Dover Publications, 2000. ISBN 978-0-486-41181-1.

George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021. doi: 10.1038/s42254-021-00314-5.

Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*, 24(89):1–97, 2023. URL https://www.jmlr.org/papers/v24/21-1524.html.

Isaac E. Lagaris, Aristidis Likas, and Dimitrios I. Fotiadis. Artificial neural networks for solving ordinary and partial differential equations. *IEEE Transactions on Neural Networks*, 9(5):987–1000, 1998. doi: 10.1109/72.712178.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=c8P9NQVtmnO.

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

Lu Lu, Xuhui Meng, Zhiping Mao, and George Em Karniadakis. DeepXDE: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021b. doi: 10.1137/19M1274067.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

Pratik Rathore, Weimu Lei, Zachary Frangella, Lu Lu, and Madeleine Udell. Challenges in training PINNs: A loss landscape perspective. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 2024. URL https://proceedings.mlr.press/v235/rathore24a.html.

Justin Sirignano and Konstantinos Spiliopoulos. DGM: A deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018. doi: 10.1016/j.jcp.2018.08.029.

Kejun Tang, Jiayu Zhai, Xiaoliang Wan, and Chao Yang. Adversarial adaptive sampling: Unify PINN and optimal transport for the approximation of PDEs. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=7QI7tVrh2c.

Lloyd N. Trefethen. *Spectral Methods in MATLAB*. SIAM, 2000. doi: 10.1137/1.9780898719598.

Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021. doi: 10.1137/20M1318043.

Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022. doi: 10.1016/j.jcp.2021.110768.

Sifan Wang, Shyam Sankaran, and Paris Perdikaris. Respecting causality for training physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 421:116813, 2024. doi: 10.1016/j.cma.2024.116813.

Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023. doi: 10.1016/j.cma.2022.115671.

Haixu Wu, Huakun Luo, Yuezhou Ma, Jianmin Wang, and Mingsheng Long. RoPINN: Region optimized physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/c745bfa5b50544882938ff4f89ff26ac-Abstract-Conference.html.

Jeremy Yu, Lu Lu, Xuhui Meng, and George Em Karniadakis. Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. *Computer Methods in Applied Mechanics and Engineering*, 393:114823, 2022. doi: 10.1016/j.cma.2022.114823.

Method	AC	CD-const	CD-rot	CD-vort	Kov.	Wins
PINN	13.0	3.22×10^3	610	444	4.00×10^3	—
Fourier	1.15×10^{10}	3.36×10^{10}	3.10×10^{10}	3.16×10^{10}	3.35×10^{10}	—
RAR	5.89	1.53×10^3	140	248	1.37×10^3	—
RAD	5.89	1.53×10^3	140	248	1.37×10^3	—
GradPenalty	2.49×10^4	4.42×10^4	4.23×10^4	4.08×10^3	1.18×10^4	—
gPINN	4.56×10^4	5.93×10^4	6.19×10^4	5.94×10^3	1.44×10^4	—
Causal PINN	1.51×10^3	—	—	—	—	—
PINNacle-like	26.9	849	164	203	1.84×10^3	—
RAA (selected)	3.28	717	60.4	126	888	5/5

Table S1: Full-core raw-HFRE baseline summary. Lower values indicate lower residual-band energy. This table and the guarded metric summaries are distinct reported result series.

Metric	Task	RAA	Non-RAA reference	Ratio
Raw HFRE	AC	3.62	5.89	0.615
Raw HFRE	CD-const	717	770	0.931
Raw HFRE	CD-rot	78.7	140	0.563
Raw HFRE	CD-vort	153	204	0.752
Raw HFRE	Kovasznay	124	1.37×10^3	0.0908
Relative L_2	AC	9.15×10^{-4}	9.98×10^{-4}	0.916
Relative L_2	CD-const	0.00442	0.00575	0.769
Relative L_2	CD-rot	0.00880	0.0119	0.738
Relative L_2	CD-vort	0.00862	0.00999	0.862
Relative L_2	Kovasznay	0.00342	0.00684	0.500
Worst-region	AC	0.00160	0.00164	0.976
Worst-region	CD-const	0.00491	0.00702	0.699
Worst-region	CD-rot	0.00936	0.0133	0.706
Worst-region	CD-vort	0.0110	0.0144	0.766
Worst-region	Kovasznay	0.00649	0.0150	0.431
Dense residual	AC	6.97×10^{-6}	8.88×10^{-6}	0.785
Dense residual	CD-const	3.78×10^{-4}	5.26×10^{-4}	0.718
Dense residual	CD-rot	3.64×10^{-5}	6.01×10^{-5}	0.606
Dense residual	CD-vort	5.50×10^{-5}	7.00×10^{-5}	0.785
Dense residual	Kovasznay	2.95×10^{-4}	3.96×10^{-4}	0.745

Table S2: Full-core selected-candidate metric summary. Reference values and ratios are retained as reported in the corresponding result series. Ratios compare RAA with the applicable non-RAA reference for each metric; the reference is not assumed to be the same method for every row.

A Complete Baseline Tables

This appendix contains the full-core baseline table, selected-candidate metric summaries, field-error comparisons, sensitivity and precision checks, and two- and three-dimensional FEM-reference results. Appendix G gives the recorded training settings and reporting details.

Table S1 gives the full-core raw-HFRE baseline summary, and the main text gives the reduced-probe comparison. Tables S2 and S3 contain the selected-candidate metric summaries. Table S2 includes the absolute error and residual values; Table S3 gives the guarded ratios and counts of eligible candidates. The eligibility condition is defined in Section 3.4 of the main text.

Task	Raw HFRE	Relative L_2	Worst	Dense	Eligible candidates
AC	0.957	0.966	0.986	0.919	3
CD-const	0.704	0.732	0.820	0.664	19
CD-rot	0.948	0.840	0.884	0.886	1
CD-vort	0.667	0.573	0.747	0.779	8
Kovasznay	0.834	0.968	0.985	0.895	9

Table S3: HFRE-guarded ratio summary for the reduced-probe profile. Ratios below one indicate lower reported values for RAA. The final column counts eligible candidates.

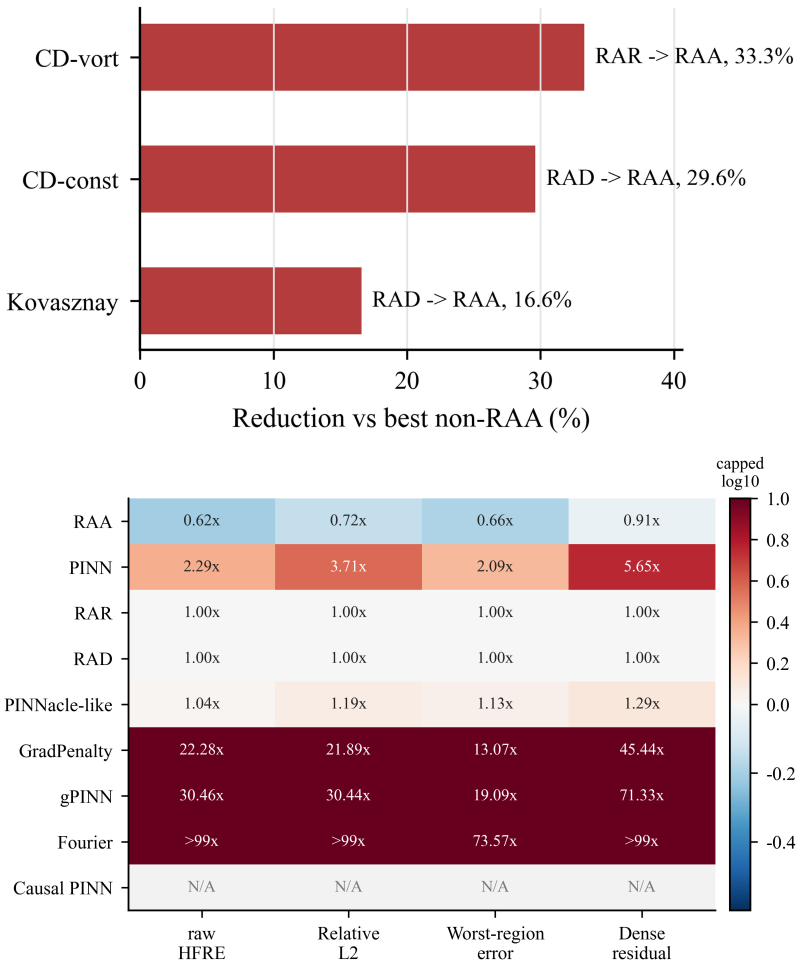


Figure S1: Additional selected-result context for the main comparisons.

B Additional Main-Result Visual Context

Figure S1 retains an additional selected-result display. Its ratios are not a direct transcription of Table S3; the unresolved correspondence is recorded in Appendix G. Quantitative comparisons in the main text use the explicitly identified tables and figures.

C Complete Field-Level Results

Figure S2 compares CD-const predictions, absolute-error maps, and error CDFs. Table S4 reports error ratios relative to RAA and the fraction of evaluation locations at which RAA has lower point-

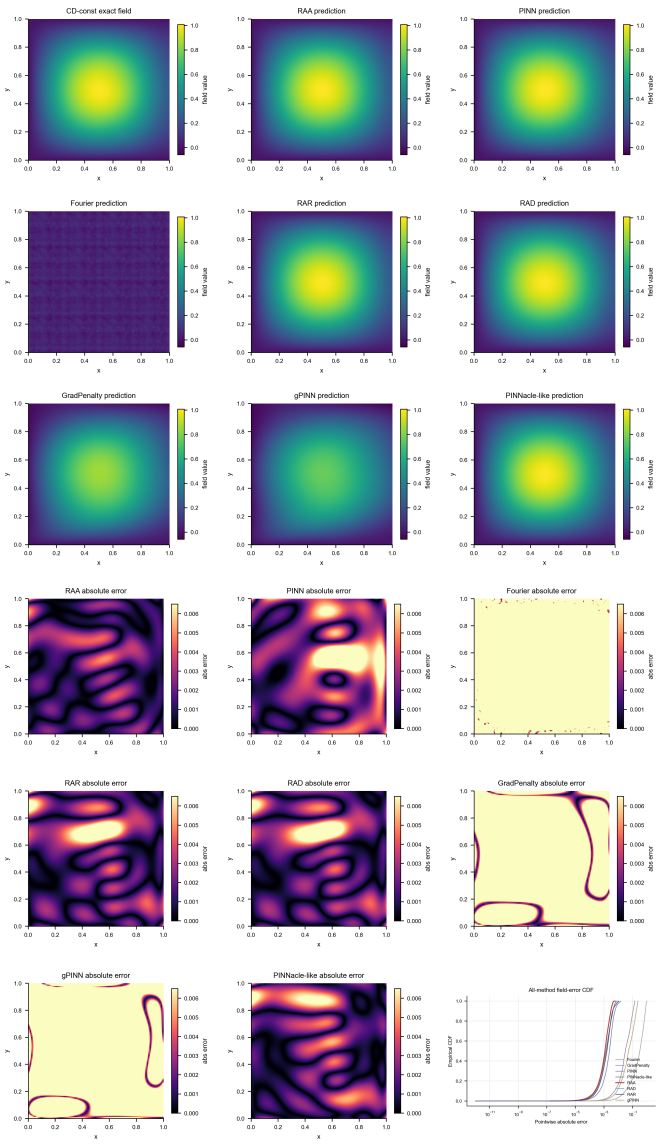


Figure S2: Complete CD-const field-level prediction and error results.

wise error. PINNacle-like has mean- and P95-error ratios of 1.31 and 1.43, while the RAA pointwise win fraction is 0.54. Against RAR and RAD, the corresponding error ratios are 1.43 and 1.63 and the pointwise win fraction is 0.59. These statistics describe the distribution and spatial extent of the differences.

D Ablation, Sensitivity, and Precision Robustness

Table S5 lists the tested spectral weights, cutoff ratios, probe sizes, numbers of added points, and floating-point precisions. Figure S3 shows the corresponding sensitivity results, and Table S6 summarizes the comparisons. The mixed sampler has lower HFRE in the nine listed comparisons with the single-signal rules. The default spectral weight matches the lowest tested value on CD-const and Kovasznay and lies 2.04% above the lowest tested value on CD-vort.

The cutoff and probe-size panels characterize their respective sweeps. Their minima refer to the spectral statistic under the displayed settings. Changing the cutoff also changes which frequencies

Baseline	Mean error over RAA	P95 error over RAA	RAA pointwise win fraction
PINN	2.06	2.17	0.73
Fourier	261	247	1.00
RAR	1.43	1.63	0.59
RAD	1.43	1.63	0.59
GradPenalty	32.5	35.5	0.98
gPINN	54.5	59.6	0.99
PINNacle-like	1.31	1.43	0.54

Table S4: CD-const field-error summary. The first two ratios divide baseline error by RAA error, so values above one favor RAA. The final column gives the fraction of evaluated locations at which RAA has lower pointwise error.

Quantity	Tasks	Tested values
λ_{AA}	CD-const, CD-vort	$\{0, 2.6, 3.9, 5.2, 6.5, 7.8, 10.4\} \times 10^{-7}$
λ_{AA}	Kovasznay	$\{0, 4, 6, 8, 10, 12, 16\} \times 10^{-7}$
Cutoff ratio	CD-const, CD-vort, Kovasznay	$\{0.45, 0.50, 0.55, 0.60, 0.65\}$
Probe points	CD-const, CD-vort, Kovasznay	$\{8, 16, 24, 32\}$
Adaptive additions	CD-const, CD-vort, Kovasznay	$\{64, 96, 128, 192\}$
Precision	Poisson, CD-const, Kovasznay	FP32 and FP64

Table S5: Tested settings for the sensitivity and precision comparisons. The spectral weights are absolute values in the reported scale; a unit value on a relative-scale plot does not denote an absolute weight of one.

Comparison	Observation	Numerical summary
A0–A5	A5 has the lowest raw HFRE on five cases	Maximum A5 / best A0–A4 $= 0.463$
Sampling rules	Mixed rule is lower in the nine listed comparisons	Maximum mixed / single-signal $= 0.940$
Spectral weight	Default / lowest tested HFRE: CD-const, CD-vort, Kovasznay	$1.000,\ 1.020,\ 1.000$
Cutoff sweep	Selected / lowest tested statistic in the displayed sweep	1.000 for all three tasks
Floating-point precision	RAA is lower in the six listed comparisons	Maximum RAA / reference $= 0.316$

Table S6: Summary of the component, sampling, parameter, and precision comparisons. The listed numerical summaries describe their respective experiments.

enter the tail sum; a cutoff sweep is therefore distinct from comparing all trained models with a common fixed-band evaluator. Likewise, changing the training probe size is distinct from refining the evaluation grid of one frozen model. The precision checks compare FP32 and FP64 on Poisson, CD-const, and Kovasznay.

E External FEM/COMSOL Validation Details

Figure S4 and Table S7 compare the two-dimensional CD-const predictions with the FEM reference field. RAA has a relative L_2 discrepancy of 0.00354, a mean absolute discrepancy of 0.00142, and a maximum absolute discrepancy of 0.00609. The corresponding PINNacle-like values are 0.00486, 0.00187, and 0.00855.

Figure S5 and Table S8 give the three-dimensional convection–diffusion comparison. The metrics use the common volume grid, and the displayed slices visualize the volume solution. RAA has

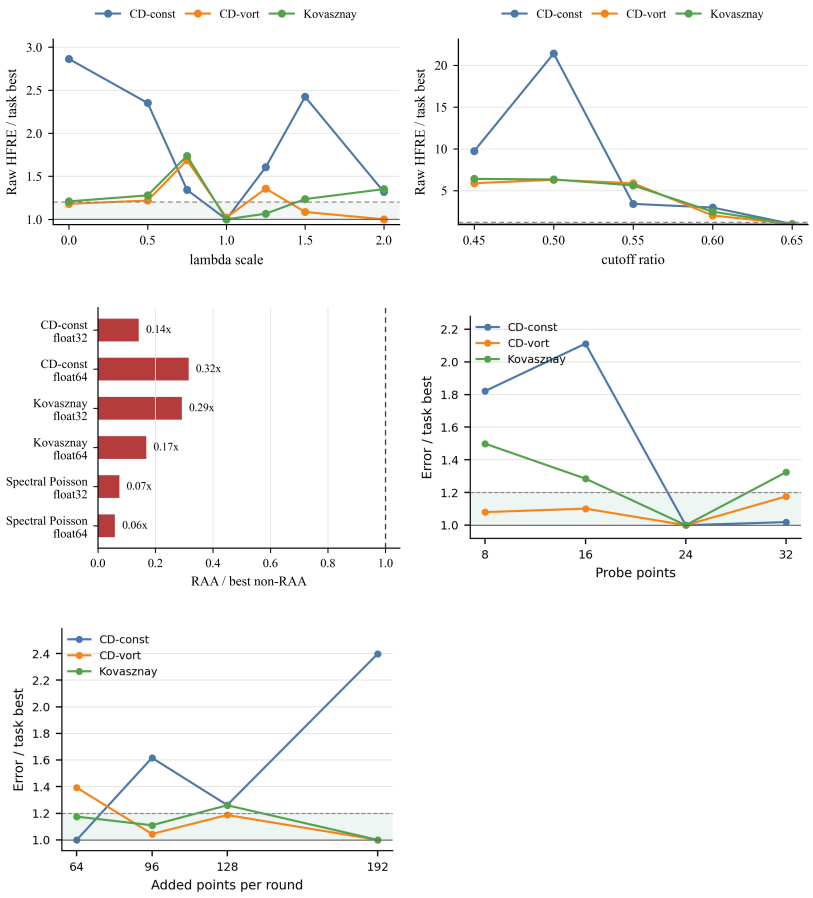


Figure S3: Sensitivity and precision robustness results.

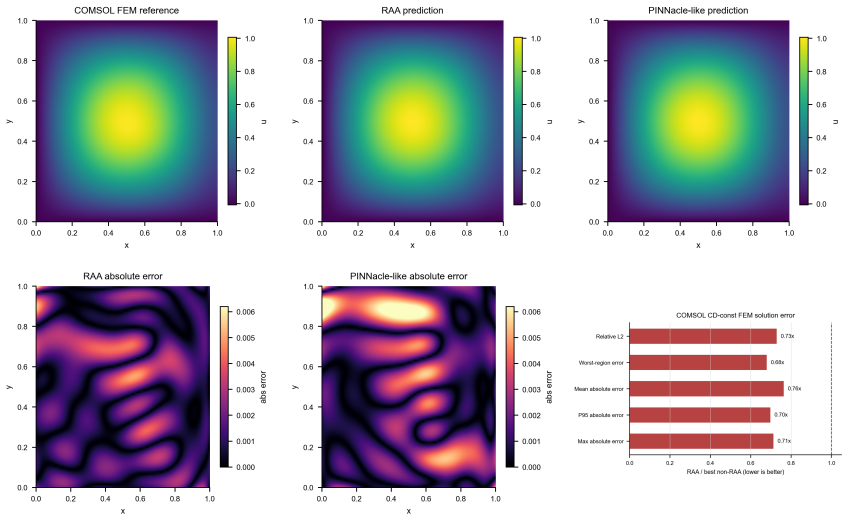


Figure S4: Complete two-dimensional COMSOL/FEM validation results.

Method	Relative L_2	Worst	Mean abs.	P95 abs.	Max abs.
RAA	0.00354	0.00411	0.00142	0.00355	0.00609
PINN	0.00763	0.0106	0.00294	0.00769	0.0154
Fourier	0.925	0.918	0.372	0.877	1.02
RAR	0.00544	0.00793	0.00203	0.00580	0.0110
RAD	0.00544	0.00793	0.00203	0.00580	0.0110
GradPenalty	0.120	0.136	0.0462	0.126	0.147
gPINN	0.203	0.226	0.0777	0.212	0.241
PINNacle-like	0.00486	0.00604	0.00187	0.00509	0.00855

Table S7: Discrepancies against the two-dimensional FEM reference field. Lower values indicate closer agreement with the reference.

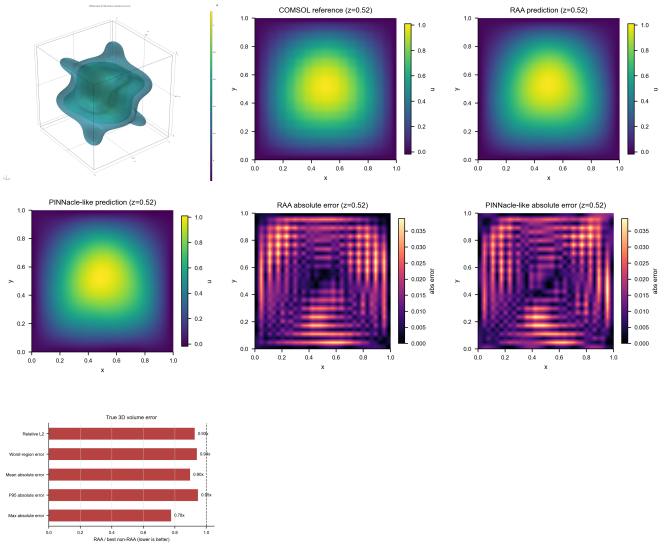


Figure S5: Complete three-dimensional COMSOL/FEM validation results.

Method	Relative L_2	Worst	Mean abs.	P95 abs.	Max abs.
RAA	0.0458	0.0407	0.0120	0.0347	0.0607
PINN	0.0505	0.0446	0.0136	0.0371	0.105
Fourier	1.00	0.872	0.251	0.764	1.06
RAR	0.191	0.126	0.0594	0.113	0.165
RAD	0.0517	0.0455	0.0139	0.0382	0.0917
GradPenalty	0.274	0.240	0.0727	0.203	0.345
gPINN	0.470	0.386	0.130	0.337	0.512
PINNacle-like	0.0494	0.0433	0.0133	0.0367	0.0781

Table S8: Discrepancies against the three-dimensional FEM reference field. Lower values indicate closer agreement with the reference.

relative L_2 discrepancy 0.0458 and maximum absolute discrepancy 0.0607, compared with 0.0494 and 0.0781 for PINNacle-like. Appendix G lists the available reference information.

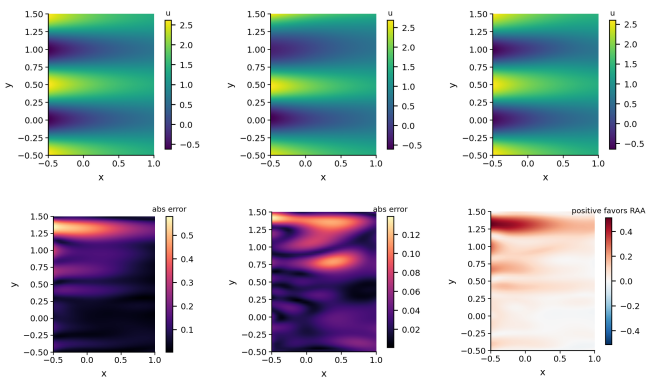


Figure S6: Kovasznay field diagnostics.

F Representative Kovasznay Diagnostics

Kovasznay flow couples the two velocity components u, v with pressure p. At $\mathrm{Re} = 40$, the standard steady incompressible residual components are

$$r_u = uu_x + vu_y + p_x - \mathrm{Re}^{-1}(u_{xx} + u_{yy}),$$
$$r_v = uv_x + vv_y + p_y - \mathrm{Re}^{-1}(v_{xx} + v_{yy}),$$
$$r_c = u_x + v_y.$$

Figure S6 compares the displayed velocity predictions and error fields for RAR and RAA. These component equations specify the mathematical problem; the component scaling used in a numerical loss or spectral aggregate is a separate implementation choice.

G Reproducibility

The principal tables summarize three seeds, $\{0, 1, 2\}$. The recorded architecture has five hidden layers of width 96, with 1024 initial collocation points and 256 boundary points. The two profiles use 3000 Adam-stage epochs. Their probe sizes, candidate pools, and point-addition schedules are listed in Table 1. The reported computing environment is a Kaggle notebook using an NVIDIA Tesla T4 and PyTorch 2.10.0+cu128. Local RTX 3050 Ti and CPU checks are excluded from the principal tables.

The retained tables contain seed aggregates and selected-candidate summaries. They do not specify the final within-set candidate choice, checkpoint correspondence across metrics, the aggregation operation, or per-seed dispersion. The main baseline table, guarded summaries, and ablation table are discussed separately because their run-level correspondence has not been recovered. In particular, the guarded CD-const HFRE ratio is not obtained by dividing its entries in the main baseline table. No measured common wall-clock duration is established by the retained configuration summaries.

The descriptions specify local FFT-based probing and the tested sensitivity ranges. They leave the global transform normalization, per-run cutoff and spectral weights, local window geometry, coupled-residual scaling, and worst-region construction unresolved. The FEM summaries provide discrepancy values but not the mesh, element order, interpolation details, or a reference discretization-error estimate. The cutoff/probe sensitivity plots and reference-field comparisons are interpreted with these reporting limits. No formal significance test is reported.

The Parseval relation underlying the frequency-weighting interpretation is given in Section 3.2 of the main text. The appendix retains the reporting limits above rather than duplicating that derivation.