

Research Article

Implicit Branch Selection in Physics-Informed Neural Networks for an Underdetermined Exterior Laplace Problem: Potential Flow Around a Circular Cylinder with Weak Far-Field Regularization

Y. Jangid^{1,2}, P. Sathya², Der-Liang Young^{2,3}

1. Department of Civil Engineering, MBM University, Jodhpur, India; 2. Department of Civil Engineering, National Taiwan University, Taiwan, Province of China; 3. Core Tech System Co. Ltd, Moldex3D, Taiwan

Physics-informed neural networks (PINNs) have emerged as a promising framework for solving partial differential equations by embedding physical laws directly into the learning process. This study investigates the behaviour of a PINN applied to the classical exterior potential-flow problem around a circular cylinder, focusing on convergence under weak soft-penalty regularization rather than on explicit far-field boundary conditions. The no-penetration condition on the cylinder surface is the primary boundary constraint, while the physics-informed loss enforces the Laplace equation for the stream function. A weak far-field soft penalty provides qualitative free-stream information without constituting a hard far-field boundary condition when applied at radial distances greater than five-cylinder radii and weighted one-tenth of the solid boundary loss.

It is important to establish that the exterior Laplace problem with only a solid wall boundary condition is mathematically non-unique; many harmonic functions, including the trivial zero solution, satisfy both the governing equation and the no-penetration condition. This work does not seek to resolve that non-uniqueness analytically or enforce mathematical uniqueness. Instead, it examines the empirical behaviour of the PINN under the given formulation. The results show convergence toward the classical uniform-flow solution branch, which is interpreted as an implicit branch selection influenced by the combined effects of global residual minimization, boundary anchoring, the soft far-field penalty, and the network's bias toward smooth, low-frequency functions.

A parametric study spanning 33 training runs across 2000 to 10000 epochs, and 5000 to 60000 collocation points, reveals a highly non-monotonic optimization landscape. A recurrent failure mode is identified in which the absolute error converges to approximately 5.4, confirmed to equal the maximum magnitude of the analytical stream function over the evaluation domain, indicating collapse toward the trivial zero solution. A stability-optimal epoch regime near 8000 epochs is identified, producing zero failures across all tested collocation densities. The best-performing configuration achieves close agreement with the analytical solution for the stream function, velocity field, and pressure coefficient distribution.

Corresponding authors: D.L. Young, dlyoung@ntu.edu.tw; P. Sathya, sathyap@ntu.edu.tw; Y. Jangid, Yashs99834@gmail.com

1. Introduction

The simulation of fluid flow around solid bodies is a fundamental problem in fluid mechanics with broad applications in aerodynamics, hydrodynamics, and engineering design ^{[1][2][3]}. Among classical benchmark problems, the two-dimensional incompressible and irrotational flow around a circular cylinder plays a central role due to its analytical tractability and well-understood physical characteristics ^{[4][5][6]}. The problem serves as a standard test case for assessing numerical methods and data-driven approaches for solving partial differential equations in exterior domains.

It is important to clarify the mathematical context of this study before presenting the formulation. The exterior Laplace problem subject only to a solid wall boundary condition is not uniquely determined. The governing equation admits infinitely many harmonic solutions satisfying the no-penetration condition, including higher angular-mode functions of the form $\psi_n(r, \theta) = A_n(r^n - a^{2n}r^{-n}) \sin(n\theta)$ for any integer $n \geq 1$, as well as the trivial zero solution $\psi = 0$. Classical theory resolves this non-uniqueness by imposing an explicit far-field condition. The present work does not impose such a condition as a hard constraint. Instead, it demonstrates empirically that a specific PINN formulation, including a weak soft far-field penalty, can converge to the physically relevant $n = 1$ solution branch under appropriate training conditions. This convergence is interpreted as architecture- and optimizer-dependent implicit branch selection, not as mathematical recovery of the unique solution from the governing equation and wall boundary alone.

Traditional numerical methods for solving exterior flow problems, such as finite difference, finite volume, and finite element methods, typically require the introduction of artificial far-field boundaries to truncate the unbounded physical domain. Appropriate boundary conditions must then be imposed at these fictitious boundaries to approximate the behaviour at infinity. While effective, this approach introduces additional modelling assumptions and may influence the accuracy and stability of the numerical solution, particularly when the far-field behaviour is not known *a priori*.

Physics-informed neural networks (PINNs) have recently attracted significant attention as an alternative framework for solving partial differential equations [7][8]. By embedding governing equations and boundary conditions directly into the loss function, PINNs enable the approximation of solution fields without the need for mesh generation or labelled training data. Since their introduction, PINNs have been successfully applied to a wide range of problems involving elliptic, parabolic, and hyperbolic equations [9][10].

Despite these advances, most existing PINN studies addressing exterior flow problems continue to rely on artificial far-field boundary conditions to ensure well-posedness and numerical stability. The extent to which PINNs can infer correct far-field behaviour solely from the governing equations and solid boundary information remains an open question. Understanding this capability is particularly important for unbounded-domain problems, where the introduction of fictitious boundaries may be undesirable or impractical. In this work, we investigate the performance of physics-informed neural networks for the classical exterior potential flow around a circular cylinder, with a specific focus on replacing explicit hard far-field boundary conditions with weak soft far-field regularization. The stream function formulation is adopted, and the Laplace equation governing incompressible, irrotational flow is enforced throughout the domain. The primary prescribed boundary condition is the no-penetration condition on the surface of the cylinder, supplemented by a weak soft far-field penalty as described in Section 3.5.

The main contributions of this study are presented as follows. First, a precise characterization of the PINN loss formulation, explicitly distinguishing the primary solid boundary condition from the weak soft far-field penalty and clarifying the role of each in promoting implicit branch selection. Second, empirical demonstration across 33 independent training runs that a stability-optimal epoch regime near 8000 epochs produces consistent convergence to the physical solution branch across collocation densities ranging from 5000 to 60000 points. Third, identification and confirmation of the recurrent failure mode: the recurring absolute error of approximately 5.4 is shown to correspond to the maximum magnitude of the analytical stream function over the evaluation domain, confirming failure as collapse

toward the trivial zero solution rather than convergence to an alternative harmonic branch. Fourth, a systematic parametric study documenting the non-monotonic relationship between training epochs, collocation density, and solution accuracy. Fifth, a clear discussion of limitations, including the absence of strict uniqueness guarantees and sensitivity to hyperparameter selection.

2. Governing Equations and Classical Solution

In this study, we consider the two-dimensional incompressible and irrotational flow of an inviscid fluid around a circular cylinder of radius a placed in a uniform free stream of velocity U , as shown in Fig. 1. Under these assumptions, the flow can be described using potential flow theory, for which the velocity field is derived from a scalar potential or, equivalently, from a stream function formulation.




Figure 1. Schematic of the exterior potential flow problem around a circular cylinder of radius a in a uniform free stream of velocity U . The governing equation is enforced in the exterior domain, while the no-penetration condition $\psi = 0$ is prescribed on the cylinder surface. No explicit hard far-field boundary condition is imposed; instead, weak soft far-field regularization is incorporated into the loss formulation.

2.1. Stream Function Formulation

For incompressible flow, the velocity field $\mathbf{u} = (u, v)$ satisfies the continuity equation:

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

In two dimensions, this condition is automatically satisfied by introducing a stream function $\psi(x, y)$ such that

$$u = \frac{\partial \psi}{\partial y}, v = -\frac{\partial \psi}{\partial x}. \quad (2)$$

For irrotational flow, the vorticity vanishes, leading to the condition:

$$\nabla \times \mathbf{u} = 0. \quad (3)$$

Substituting the stream function definition into the vorticity-free condition yields the governing equation

$$\nabla^2 \psi = 0 \quad (4)$$

which is the Laplace equation. Therefore, the stream function for incompressible, irrotational flow satisfies a second-order elliptic partial differential equation throughout the flow domain.

2.2. Boundary Conditions for Exterior Cylinder Flow

The physical domain of interest is the exterior region $\Omega = \mathbb{R}^2 \setminus \{r \leq a\}$, where $r = \sqrt{x^2 + y^2}$. On the surface of the cylinder, the no-penetration condition requires that the normal component of velocity vanishes. In terms of the stream function, this condition implies that ψ is constant along the solid boundary. Without loss of generality, the stream function on the cylinder surface is set to zero:

$$\psi(r=a)=0. \quad (5)$$

In classical formulations, an additional far-field condition is imposed to enforce uniform flow behaviour as $r \rightarrow \infty$, typically of the form

$$\psi \sim Uy \text{ as } r \rightarrow \infty. \quad (6)$$

In the present work, no explicit hard far-field boundary condition is prescribed. Instead, weak soft far-field regularization is incorporated into the composite loss function, as described in Section 3.6. This regularization introduces qualitative free-stream information without enforcing a strict asymptotic boundary constraint. The governing equation is enforced throughout a finite rectangular sampling domain representing the exterior domain, and the no-penetration condition on the cylinder surface is

the only boundary condition applied. The trained solution in this domain is utilized to approximate the external flow field without imposing far-field requirements at the domain boundaries. The ability of the physics-informed neural network to converge toward the physically relevant asymptotic behaviour under this weakly constrained formulation is a central focus of this study.

2.3. Velocity Potential Formulation

For irrotational flow, the velocity field can also be expressed in terms of a scalar velocity potential $\phi(x, y)$ as

$$u = \frac{\partial \phi}{\partial x}, v = \frac{\partial \phi}{\partial y}. \quad (7)$$

Incompressibility of continuity equation of Eq. (1) implies that ϕ satisfies the same Laplace equation as ψ in Eq. (4) throughout the exterior domain Ω .

$$\nabla^2 \phi = 0 \quad (8)$$

Hence, ψ and ϕ are harmonic functions representing the same physical flow field. For the exterior potential flow around a circular cylinder of radius a in a uniform free stream of velocity U , the classical analytical expression for the velocity potential in polar coordinates (r, θ) is given by ^[11],

$$\phi(r, \theta) = Ur \left(1 + \frac{a^2}{r^2} \right) \cos(\theta). \quad (9)$$

This solution satisfies the Laplace equation in the exterior domain and represents uniform flow perturbed by the presence of the solid cylinder. The corresponding stream function, given in Section 2.2, is orthogonal to the velocity potential, and the two functions together form a harmonic conjugate pair. The equipotential lines $\phi = \text{constant}$ are everywhere orthogonal to the streamlines $\psi = \text{constant}$. This orthogonality property provides a useful physical interpretation of the flow field and serves as an additional qualitative validation of numerical solutions. In the present work, equipotential lines are used in conjunction with streamlines to assess the physical consistency of the PINN-predicted flow field.

2.4. Classical Analytical Solution

For reference and validation, the analytical solution for the potential flow around a circular cylinder in a uniform stream is briefly summarized. In the polar coordinates (r, θ) , the classical stream function is given by ^[11]

$$\psi(r,\theta)=Ur\left(1-\frac{a^2}{r^2}\right)\sin(\theta). \quad (10)$$

This solution satisfies the Laplace equation in the exterior domain and the no-penetration condition $\psi(a, \theta) = 0$. The corresponding velocity components in the polar coordinates are

$$v_r = \frac{1}{r} \frac{\partial \psi}{\partial \theta} = U \left(1 - \frac{a^2}{r^2}\right) \cos(\theta), \quad (11)$$

$$v_\theta = -\frac{\partial \psi}{\partial r} = -U \left(1 + \frac{a^2}{r^2}\right) \sin(\theta) \quad (12)$$

On the surface of the cylinder ($r = a$), the tangential velocity reduces to

$$v_\theta(r = a) = -2U \sin(\theta), \quad (13)$$

indicating that the maximum velocity occurs at the top and bottom of the cylinder ($\theta = \pm\pi/2$).

2.5. Pressure Coefficient

The pressure distribution on the cylinder surface is obtained from Bernoulli's equation. Neglecting elevation changes and assuming steady inviscid flow, the pressure coefficient C_p is defined as ^[11],

$$C_p = \frac{p - p_\infty}{\frac{1}{2}\rho U^2}, \quad (14)$$

where p_∞ is the pressure in the uniform free stream. Using the surface velocity expression, the analytical pressure coefficient for potential flow around a circular cylinder is

$$C_p(\theta) = 1 - 4 \sin^2(\theta). \quad (15)$$

This pressure distribution serves as an important benchmark for assessing the accuracy of the velocity gradients predicted by the PINN, as correct recovery of C_p requires accurate first-order derivatives of the stream function along the solid boundary.

3. Physics-Informed Neural Network Formulation

In this section, the physics-informed neural network (PINN) formulation used to approximate the exterior potential flow around a circular cylinder is described. The network architecture, the construction of the physics-informed loss function, and the treatment of boundary conditions are presented in a manner consistent with the governing equations introduced in Section 2.




Figure 2. Schematic illustration of the physics-informed neural network (PINN) used to approximate the stream function ψ . The network takes the spatial coordinates (x, y) as inputs and outputs the predicted stream function ψ . The composite loss function consists of the governing equation residual enforcing the Laplace equation, the solid boundary condition enforcing the no-penetration condition on the cylinder surface, and a weak soft far-field regularization term.

3.1. Neural Network Architecture

The stream function $\psi(x, y)$ is approximated by a fully connected feedforward neural network $\psi_\theta(x, y)$, where θ denotes the collection of trainable parameters (weights and biases). Therefore, the network provides a parametric representation of the stream function, written as $\psi(x, y) \approx \psi_\theta(x, y)$, mapping the spatial coordinates (x, y) in a two-dimensional domain to a scalar output representing the stream function value at that location.

The architecture consists of three hidden layers, each containing 128 neurons, Fig.2. The SiLU (Sigmoid Linear Unit) activation function is used throughout, defined as $\text{SiLU}(z) = z \cdot \sigma(z)$ where σ is the standard sigmoid function. SiLU is selected because it is smooth and infinitely differentiable, which is essential for the stable computation of the second-order Laplacian residual through automatic differentiation. Unlike ReLU, SiLU does not produce zero gradients for negative inputs, which improves gradient flow during training of deep networks solving second-order PDEs.

All network weights are initialized using Xavier normal initialization ^[12], and all biases are initialized to zero. In Xavier initialization, weights are sampled from a normal distribution with variance $2/(n_{in} + n_{out})$, where n_{in} and n_{out} are the numbers of input and output neurons for each layer. This initialization

ensures that signal variance is approximately preserved across layers and prevents vanishing or exploding gradients in early training.

Parameter	Value	Justification
Hidden layers	3	Sufficient depth for smooth harmonic approximation
Neurons per layer	128	Balances expressivity and training speed
Activation function	SiLU	Smooth, differentiable, good gradient flow
Weight initialization	Xavier normal	Preserves activation variance across layers
Bias initialization	Zero	Standard practice
Optimizer	Adam	Adaptive learning rate, robust to noise
Initial learning rate	0.001	Standard Adam default
Loss weight — PDE residual (λ_{PDE})	1.0	Base reference weight
Loss weight — boundary condition (λ_{BC})	50.0	Enforces no-penetration strongly
Loss weight — far-field soft penalty (λ_{FF})	0.1	Weak qualitative guidance only
Far-field penalty activation radius	$r > 5a$	Beyond this, the flow approaches uniform
Sampling domain (x)	$[-3, 7]$ m	Asymmetric to cover the downstream region
Sampling domain (y)	$[-3, 3]$ m	Symmetrical about the centreline
Boundary points on a cylinder	200	Uniform circumferential distribution

Parameter	Value	Justification
Cylinder radius a	1.0 m	Unit radius for non-dimensionalisation
Free-stream velocity U	1.0 m/s	Unit velocity for non-dimensionalisation
Precision	float32	Single precision throughout
Framework	PyTorch with autograd	Automatic differentiation for all gradients
Random seed	Independent random initialization for the parametric study; fixed seeds (1–5) used for the seed robustness study	See Section 6.2

Table 1. Complete implementation parameters for reproducibility.

Additional runs were also conducted using near-zero weight initialization, obtained by scaling the Xavier initialization by a factor of 0.01, to investigate the sensitivity of the optimization process to parameter initialization. These experiments consistently converged to the zero-solution failure mode, producing maximum absolute errors of approximately 5.4–5.9. This behaviour indicates that initialization significantly influences the optimization trajectory and the solution branch selected during training. Since the dominant components of the loss function—the Laplacian residual and the solid boundary condition—are also satisfied by the trivial solution $\psi = 0$ on the bounded computational domain, insufficient guidance from the remaining loss terms may cause optimization to collapse toward this non-physical branch. Therefore, all results reported in this work employ the standard Xavier normal initialization.

3.2. Physics-Informed Residual of the Governing Equation

The governing equation for incompressible, irrotational flow is the Laplace equation, which is given in Eq. (4). Within the PINN framework, this equation is enforced by minimizing the residual of the Laplacian

evaluated using automatic differentiation. Given a set of interior collocation points $\{(x_i, y_i)\}_{i=1}^{N_f}$ in the exterior domain, the PDE residual is defined as

$$\mathcal{R}_{\text{PDE}}(x_i, y_i) = \frac{\partial^2 \psi_\theta}{\partial x^2}(x_i, y_i) + \frac{\partial^2 \psi_\theta}{\partial y^2}(x_i, y_i). \quad (16)$$

The corresponding physics-informed loss term associated with the governing equation is

$$L_{\text{PDE}} = \frac{1}{N_f} \sum_{i=1}^{N_f} |\mathcal{R}_{\text{PDE}}(x_i, y_i)|^2. \quad (17)$$

This term enforces the harmonicity of the learned stream function throughout the exterior domain.

3.3. Boundary Condition Enforcement on the Cylinder Surface

On the surface of the circular cylinder, the no-penetration condition requires that the stream function be constant. As discussed in Section 2, this constant is set to zero without loss of generality in Eq. (5),

Let $\{(x_j^b, y_j^b)\}_{j=1}^{N_b}$ denote a set of collocation points distributed along the cylinder surface. The boundary loss term is defined as

$$L_{\text{BC}} = \frac{1}{N_b} \sum_{j=1}^{N_b} |\psi_\theta(x_j^b, y_j^b)|^2. \quad (18)$$

While no hard far-field boundary condition is imposed in the classical sense, a weak soft penalty is additionally applied at large radial distances, as described in Section 3.5. This penalty provides qualitative free-stream information without constituting an explicit constraint at a defined outer boundary. This formulation allows direct assessment of whether the governing equation alone is sufficient to guide the network toward the correct asymptotic behaviour.

3.4. Composite Loss Function

The total loss function minimized during training is a weighted sum of five terms:

$$L = L_{\text{PDE}} + \lambda_{\text{BC}} \cdot L_{\text{BC}} + \lambda_{\text{FF}} \cdot L_{\text{FF}} + L_{\text{sym}} + L_{\text{Cp}} \quad (19)$$

where L_{PDE} is the Laplacian residual loss enforcing the governing equation, L_{BC} is the solid boundary condition loss enforcing the no-penetration condition, L_{FF} is the weak far-field soft penalty described in Section 3.5, L_{sym} is a symmetry regularization term, and L_{Cp} is a pressure coefficient supervision term at the cylinder surface. The weights $\lambda_{\text{BC}} = 50$ and $\lambda_{\text{FF}} = 0.1$ are reported in Table 1.

The boundary loss is assigned a large weight of 50 to ensure strong enforcement of the no-penetration condition and to prevent boundary violations from propagating into the interior solution. The far-field penalty is intentionally assigned a small weight of 0.1, making it the weakest component of the total loss. Its role is qualitative: it provides information about the free-stream velocity scale U and discourages solutions that grow unboundedly, without constituting a hard boundary constraint.

The balance between these loss terms is critical for achieving stable convergence to the physical solution branch, as demonstrated by the parametric study in Section 6.

3.5. Loss Function Evaluated at the Analytical Solution

To provide insight into the structure of the loss landscape, the contribution of each loss term is evaluated at the known analytical solution $\psi_{\text{exact}}(r, \theta) = Ur(1 - a^2/r^2) \sin(\theta)$. Since ψ_{exact} satisfies the Laplace equation exactly, the PDE residual loss evaluates to:

$$L_{PDE}[\psi_{\text{exact}}] = (1/N_f) \sum |\Delta \psi_{\text{exact}}(x_i, y_i)|^2 = 0 \quad (20)$$

The boundary condition loss also evaluates to zero, since $\psi_{\text{exact}}(r=a, \theta) = 0$ exactly:

$$L_{BC}[\psi_{\text{exact}}] = (1/N_b) \sum |\psi_{\text{exact}}(x_j^b, y_j^b)|^2 = 0 \quad (21)$$

The far-field soft penalty evaluates to zero in the limit $r \rightarrow \infty$, since $\psi_{\text{exact}} \rightarrow Uy$ as $r \rightarrow \infty$. On the finite training domain at $r > 5a$, the deviation $\psi_{\text{exact}} - Uy$ is small but nonzero due to the cylinder perturbation term, giving $L_{FF}[\psi_{\text{exact}}] \approx \varepsilon$ where ε is small.

The pressure coefficient loss evaluates to zero exactly, since $C_{p^{\text{exact}}}(\theta) = 1 - 4\sin^2(\theta)$ and the velocity components derived from ψ_{exact} reproduce this distribution.

Therefore, the analytical solution achieves near-zero total loss. However, as noted above, $\psi = 0$ also achieves zero loss on the two dominant terms L_{PDE} and L_{BC} . The analytical solution is distinguished from $\psi = 0$ only by the far-field penalty and pressure coefficient terms. This confirms that these additional terms are essential for recovering the physical branch and motivates their inclusion in the composite loss.

3.6. Far-Field Soft Penalty

A soft penalty is applied at collocation points where $r > 5a$. At these locations, the network output is penalized for deviating from the far-field behaviour $\psi \approx Uy$, with a weight of 0.1 relative to the PDE residual. This term is explicitly distinguished from a classical hard far-field boundary condition. It does

not enforce a Dirichlet constraint at a specific outer boundary and is not sufficient alone to guarantee convergence to the physical solution branch. Its role is qualitative: it introduces the free-stream velocity scale U into the learning problem, which is otherwise unspecified by the solid wall condition alone, and it discourages solutions that grow unboundedly at large distances. This weak penalty, combined with the symmetry constraint, the pressure coefficient loss, and the spectral bias of the network, collectively promotes selection of the physically relevant branch during optimization.

3.7. Velocity and Pressure Recovery via Automatic Differentiation

Once the stream function approximation ψ_θ is obtained, the velocity components are computed directly through automatic differentiation,

$$u_\theta = \frac{\partial \psi_\theta}{\partial y}, v_\theta = -\frac{\partial \psi_\theta}{\partial x}. \quad (22)$$

These derivatives are evaluated exactly with respect to the neural-network surrogate function through automatic differentiation.

$$C_p = 1 - \frac{u_\theta^2 + v_\theta^2}{U^2}. \quad (23)$$

Accurate prediction of the pressure coefficient provides a stringent test of the PINN's ability to recover physically meaningful velocity gradients near the solid boundary.

4. Exterior Flow with Minimal Far-Field Information

In this section, the numerical treatment of the exterior flow problem without far-field boundary conditions is described. Emphasis is placed on the sampling strategy, domain representation, and training considerations required to ensure stability and accuracy when the computational domain is unbounded and only the solid boundary condition is prescribed.

4.1. Computational Domain and Sampling Strategy

The physical problem under consideration is defined on an unbounded exterior domain. In practice, training points must be sampled within a finite computational region while avoiding the introduction of artificial boundary constraints. To this end, a rectangular domain sufficiently large to capture the dominant flow features is selected for sampling, while no boundary conditions are imposed on the outer edges of the domain.

Interior collocation points are distributed throughout the exterior region, excluding the interior of the cylinder. To improve resolution in regions of strong gradients, particularly near the cylinder surface, a non-uniform sampling strategy is adopted in which a higher density of points is placed close to the solid boundary. This targeted sampling ensures accurate enforcement of the governing Laplace equation in regions that have a significant influence on the global solution.

Boundary collocation points are distributed uniformly along the circumference of the cylinder to enforce the no-penetration condition. No explicit hard boundary constraints are imposed on the outer edges of the sampled domain. Instead, weak soft regularization terms are included within the composite loss function to promote physically consistent far field and symmetric behaviour. The distribution of interior and boundary collocation points used for training is shown in Fig. 3.

4.2. Training with Weak Far-Field Regularization

The replacement of hard far-field boundary conditions with a weak soft penalty presents a challenging learning problem, as the exterior Laplace equation admits infinitely many harmonic solutions and the penalty alone is insufficient to guarantee convergence to the physical branch.

During training, the physics-informed loss enforces harmonicity across the entire sampled domain, while the solid boundary condition anchors the solution on the cylinder surface. The network is therefore required to infer the appropriate decay and symmetry properties of the flow solely from the governing equation and boundary information at the solid surface.

This formulation provides a stringent test of the expressive and generalization capabilities of PINNs, particularly in unbounded-domain settings where explicit asymptotic information is not provided.

4.3. Numerical Stability and Loss Weighting

To ensure numerical stability and accurate enforcement of the solid boundary condition, the boundary loss term is assigned a relatively large weight compared to the governing equation residual. This weighting prioritizes satisfaction of the no-penetration condition and prevents boundary violations from propagating into the interior solution.




Figure 3. Training Point Distribution

The balance between the PDE residual loss and the boundary loss is critical. Insufficient boundary weighting may lead to spurious penetration through the cylinder surface, while excessive weighting can reduce the network's ability to minimize the Laplace residual in the exterior domain. The selected loss weights were determined empirically to achieve stable convergence and low global error.

4.4. Evaluation Metrics

The performance of the PINN is assessed using both local and global error measures. The absolute error in the stream function is evaluated pointwise by comparison with the analytical solution,

$$\varepsilon_{\text{abs}}(x, y) = | \psi_{\theta}(x, y) - \psi_{\text{exact}}(x, y) |. \quad (24)$$

To quantify global accuracy, the relative L_2 error over the sampled domain is computed as

$$\varepsilon_{L_2} = \frac{\| \psi_{\theta} - \psi_{\text{exact}} \|_2}{\| \psi_{\text{exact}} \|_2}. \quad (25)$$

In addition to stream function errors, the pressure coefficient distribution on the cylinder surface is compared against the analytical solution. Agreement in C_p serves as a sensitive indicator of the accuracy

of the predicted velocity gradients near the solid boundary.

4.5. Visualization and Diagnostic Analysis

To provide qualitative insight into the learned solution, streamlines, equipotential lines, and velocity-derived quantities are visualized across the exterior domain. Error contours and local L_2 error distributions are examined to identify regions of higher deviation and to assess whether errors exhibit physically meaningful spatial patterns.

Together, these diagnostics allow a comprehensive evaluation of the PINN's ability to recover the correct exterior flow behaviour without explicit far-field boundary enforcement.

5. Numerical Examples and Discussion

Example 1. Exterior Potential Flow Around a Circular Cylinder with Far-Field Condition

In this example, the classical exterior potential flow around a circular cylinder of radius a placed in a uniform free stream of velocity U is considered as a baseline validation problem for the proposed physics-informed neural network (PINN) framework. The objective of this example is to verify that the implemented network architecture, loss formulation, and automatic differentiation procedures can accurately reproduce the known analytical solution under a conventionally constrained formulation.

The stream function ψ satisfies the Laplace equation throughout the exterior domain Ω , Eq. (4), subject to the no-penetration condition on the cylinder surface Eq. (5) together with the classical far-field condition enforcing uniform flow behaviour Eq. (6):

The PINN is trained using interior collocation points distributed throughout the exterior region, boundary points located uniformly along the cylinder surface, and additional outer-region sampling points used to impose the far-field condition. The purpose of this example is not to investigate implicit branch selection or optimization instability, but rather to establish a reliable reference solution under standard boundary enforcement conditions.

The predicted streamline structure reproduces the characteristic features of classical potential flow around a circular cylinder, including symmetric streamlines, stagnation points at the front and rear of the cylinder, and smooth acceleration of the flow around the upper and lower surfaces, as shown in Fig. 4. The corresponding pressure coefficient distribution also matches the analytical solution closely, indicating that the network accurately captures the velocity gradients near the solid boundary.

This baseline validation confirms that the implemented PINN framework can successfully recover the expected harmonic flow structure when both the solid boundary condition and conventional far-field guidance are included during training. Establishing this reference case is important because it demonstrates that the numerical framework itself can reproduce the analytical solution before investigating the more challenging weakly constrained formulation considered in Example 2.

The following example then examines whether physically meaningful exterior-flow solutions can still emerge when explicit hard far-field boundary conditions are removed and replaced by weak soft far-field regularization within the composite loss formulation.

Example 2. Exterior Potential Flow Using Only the Solid Boundary Condition

In this example, the same exterior flow problem is solved using only the solid boundary condition on the cylinder surface. The governing equation remains Eq. (4), with the no-penetration condition Eq. (5), while no far-field boundary condition is imposed.

Interior collocation points are distributed in the exterior domain, with increased density near the cylinder surface to accurately resolve high-gradient regions. Boundary points are placed uniformly along the cylinder circumference. No explicit hard far-field boundary condition is imposed. Instead, weak soft regularization terms are incorporated into the composite loss function, as described in Section 3.

Despite the absence of far-field enforcement, the PINN successfully recovers the global flow structure. The predicted streamlines exhibit correct stagnation points and symmetry, and the solution approaches uniform flow behaviour at large distances from the cylinder. The corresponding absolute error and local normalized L2 error distributions are shown in Fig. 5.

Example 3. Comparison of Boundary-Enforced and Solid-Boundary-Only PINNs

In this example, a direct comparison is performed between the PINN solutions obtained in Examples 1 and 2. The comparison focuses on global accuracy, local error behaviour, and pressure coefficient recovery. While the inclusion of a far-field boundary condition leads to slightly reduced error near the cylinder surface, the solid-boundary-only formulation maintains comparable accuracy throughout most of the exterior domain. The pressure coefficient distribution predicted using only the solid boundary condition remains in close agreement with the analytical solution, and near-boundary velocity gradients are accurately captured.




Figure 4. Comparison of streamlines for the exterior potential flow around a circular cylinder. (a) An analytical solution obtained from classical potential flow theory. (b) Streamlines predicted by the physics-informed neural network using only the solid boundary condition, without imposing any explicit hard far-field boundary condition. Solid lines represent streamlines, while dashed red lines denote equipotential lines.




Figure 5. Absolute error (left) and local normalized L2 error (right) distributions between the PINN-predicted stream function and the analytical solution. Both metrics remain low across most of the domain, with localized increases in the downstream region where velocity gradients are stronger.




Figure 6. Pressure coefficient distribution along the surface of the circular cylinder. The PINN prediction is compared with the analytical solution from classical potential flow theory. Close agreement is observed over the entire surface, demonstrating that the PINN accurately captures boundary pressure variations without requiring any far-field boundary condition.

The comparison highlights the robustness of the physics-informed formulation and confirms that under the proposed weak soft far-field regularization strategy, accurate solutions can be recovered under appropriate training conditions for this class of exterior potential flow problems. Figure 6 compares the pressure coefficient predicted by the PINN with the analytical solution from classical potential flow theory. The PINN accurately reproduces the characteristic pressure distribution, with correct stagnation pressure at the front and rear of the cylinder and maximum suction at the top and bottom. The close agreement over the entire circumference indicates that the network successfully learns the velocity gradients on the solid boundary, despite the absence of any far-field boundary condition.

To provide a quantitative summary of the solution accuracy, the key error metrics are reported in Table 2. The table lists the maximum absolute error of the stream function across the computational domain. Despite the absence of any far-field boundary condition, the solid-boundary-only PINN achieves error levels comparable to the boundary-enforced case. These results demonstrate that the proposed PINN formulation, including weak soft far-field regularization, can recover accurate exterior potential flow solutions under appropriate training conditions.

5.1. Outer Ring Velocity Evaluation

To assess the spatial accuracy of the velocity field beyond the near-cylinder region, the PINN velocity components were evaluated on three concentric rings at radii $r = 2a$, $4a$, and $8a$, using the best-performing model from the seed robustness study (seed 1, maximum absolute error 0.1015). At each ring, the root mean squared error (RMSE) of the x-velocity component u , y-velocity component v , and speed

magnitude $|V|$ was computed against the analytical solution over 360 uniformly spaced angular positions.

The corresponding velocity RMSE values are summarized in Table 2. Velocity errors increase systematically with radial distance, consistent with the expected physical behaviour: the cylinder perturbation decays as $1/r^2$ in potential flow theory, meaning that the flow approaches uniform far-field behaviour most rapidly near the symmetry axes and more slowly at intermediate angles. The PINN captures this structure, with the smallest errors at $r = 2a$, where the flow gradients are strongest, and the network has the most training data, and the largest errors at $r = 8a$, which approaches the boundary of the sampled training domain.

The speed magnitude RMSE at $r = 8a$ is 0.1202 m/s, compared to the free-stream velocity of 1.0 m/s, representing a relative error of approximately 12%. While this indicates that far-field accuracy degrades at large distances from the cylinder, the error remains bounded, and the solution does not diverge. This is consistent with the soft far-field penalty preventing unbounded growth, even though it does not enforce exact far-field conditions.

Radius	u RMSE (m/s)	v RMSE (m/s)	Speed RMSE (m/s)
$r = 2a$	0.0123	0.0075	0.0128
$r = 4a$	0.0419	0.0337	0.0420
$r = 8a$	0.1179	0.1015	0.1202

Table 2. RMSE of PINN-predicted velocity components compared with the analytical solution at concentric rings of increasing radius (best model, seed 1). Errors increase with distance from the cylinder, consistent with the decay of the cylinder perturbation and the reduced training point density at large radii.

6. Observations and Parametric Sensitivity Analysis

To characterize the conditions under which implicit branch selection toward the physical solution occurs, a systematic parametric study was conducted spanning 33 independent training runs. Training epochs were varied across the set {2000, 4000, 6000, 7000, 7500, 8000, 9000, 10000} and collocation point counts across {5000, 10000, 20000, 30000, 40000, 50000, 60000}. Each run used a fresh random

initialization of the network weights and a freshly sampled set of collocation points. No warm starting was employed.

Out of the 33 runs, 11 (33.3%) converged to the failure mode identified in Section 5 as the trivial zero solution. The remaining 22 runs converged to the physical solution branch with varying accuracy. The best maximum absolute error achieved was 0.0648, obtained at 8000 epochs with both 10,000 and 60,000 collocation points, demonstrating that within the stability-optimal epoch regime, collocation density has limited impact on peak accuracy.

6.1. Threshold Behaviour in Epoch–Training Point Space

The parametric study reveals a highly irregular stability structure in the epoch–training-point parameter space. Although the results appear to exhibit an approximate threshold-like boundary separating stable and unstable configurations, several exceptions indicate that the stability boundary is not strictly deterministic.

Configurations with both high epoch counts and large collocation densities frequently converge to a recurrent high-error failure mode (absolute error ≈ 5.4 – 5.9), corresponding to collapse toward the zero-solution attractor. Parameter combinations slightly below this approximate instability region provide smaller errors and physically meaningful flow fields, even when neighbouring configurations vary greatly.

These observations demonstrate that the relationship between training duration and collocation density is strongly non-monotonic. Increasing either parameter does not guarantee improved accuracy or convergence stability. Instead, the optimization process evolves within a fragmented multi-attractor landscape in which small hyperparameter changes can produce abrupt transitions between successful convergence and catastrophic failure.

Despite this irregular behaviour, one consistent pattern emerges clearly from the results: configurations trained near 8000 epochs repeatedly produce the most accurate and stable solutions across a broad range of collocation densities. Importantly, accurate convergence is observed even at relatively low training-point counts such as 5000 points, indicating that optimization stability depends more strongly on the training regime than on collocation density alone.

This behaviour suggests that identifying stability-optimal epoch regimes may be more important than simply increasing computational scale. The results further indicate that PINN convergence in exterior-

domain problems is governed primarily by optimization dynamics rather than by training-point density alone.

6.2. Seed Robustness Study

To investigate the sensitivity of training outcomes to network initialization, five independent training runs were conducted at 8000 epochs and 40000 collocation points, each using a fixed random seed (seeds 1 through 5) with Xavier normal weight initialization. The results are presented in Table 3 and Fig. 7.

Of the five runs, two (seeds 1 and 2) converged to the physical solution branch with maximum absolute errors of 0.102 and 0.196, respectively, Fig. 8. Three runs (seeds 3, 4, 5) collapsed to the zero-solution attractor with errors near 5.9, Fig. 9 consistent with the failure mode identified in Section 5. The 40% convergence rate at this hyperparameter configuration confirms that initialization plays a critical role in determining which solution branch the optimizer enters.

These observations can be understood by examining the structure of the composite loss function. The two dominant loss components—the Laplacian residual and the solid boundary condition—are both satisfied by the trivial solution $\psi = 0$ on the bounded computational domain. Consequently, the optimization landscape contains a strong low-loss basin associated with this non-physical solution. Recovery of the physically meaningful exterior-flow branch, therefore, depends on the additional guidance provided by the weak far-field penalty, symmetry regularization, and pressure-coefficient supervision. If these auxiliary terms fail to provide sufficient gradient information during the early stages of optimization, the network may instead converge toward the trivial branch.

Importantly, the failure pattern across seeds is not random noise — the three failed seeds all converged to the same error magnitude (≈ 5.9), which slightly exceeds the expected $\max|\psi_{\text{exact}}| \approx 5.4$ over the evaluation domain. The small excess is attributed to the finite training domain and the fact that the network output does not perfectly collapse to $\psi = 0$ but to a near-zero state. The consistent failure value confirms the structural nature of the zero-solution attractor.

6.3. Hyperparameter Sensitivity and Zero-Solution Collapse

The observed behaviour indicates a highly non-convex optimization landscape that is strongly sensitive to initialization and stochastic training effects. Small changes in epoch count can alter the optimization trajectory sufficiently to move the network between distinct solution attractors. As a result,

neighbouring hyperparameter configurations may converge either to the physically relevant harmonic branch or to the recurrent zero-solution attractor.

A recurrent failure pattern was observed throughout the parametric study, where multiple training runs converged to a nearly identical high-error state with maximum absolute error approximately equal to 5.4–5.9. The consistency of this error magnitude across different epochs and collocation-point combinations indicates that the instability is not random numerical noise, but rather convergence toward a structurally stable incorrect solution branch. When compared to the analytical solution, this error magnitude roughly matches the maximum absolute value of the analytical stream function over the computational domain, showing that the failed solutions compress toward a near-zero stream-function field.

This behaviour is directly related to the structure of the loss function. The dominant components of the loss — namely the Laplacian residual and the solid boundary condition — are both satisfied by the trivial solution $\psi = 0$. Consequently, unless the additional loss terms provide sufficient gradient information during optimization, the network may converge toward this zero-solution attractor instead of the physical exterior-flow solution branch.

The repeated appearance of this failure mode demonstrates that the optimization landscape contains multiple stable attractors. Some training configurations converge to the physically correct harmonic branch, while others collapse into the trivial near-zero state. This observation highlights the importance of loss balancing, initialization strategy, and hyperparameter selection in exterior-domain PINN problems. Although the present mechanism is specific to the loss formulation considered here.

Seed	Max Abs Error (ψ)	Rel L2 Error	Status
1	0.1015	1.92×10^{-2}	Converged
2	0.1959	2.17×10^{-2}	Converged
3	5.8939	1.446	Failed (zero soln.)
4	5.9092	1.428	Failed (zero soln.)
5	5.8949	1.438	Failed (zero soln.)
Mean (converged)	0.149 ± 0.047	—	2/5 converged

Table 3. Seed robustness study at 8000 epochs, 40000 training points. Each run uses Xavier normal initialization with a fixed random seed. Two of five runs converged to the physical solution branch, while three collapsed to the zero-solution attractor. The consistent failure error near 5.9 confirms a structural attractor rather than random failure.

6.4. Optimization Stability and Convergence Behaviour

The parametric study demonstrates that convergence of the PINN solution is highly sensitive to optimization dynamics. Increasing the number of epochs or collocation points does not always improve accuracy, as training can shift abruptly between successful convergence and failure.

A distinct stability-optimal region is observed near 8000 training epochs, where the network repeatedly converges to the physically relevant harmonic solution across a broad range of collocation densities. Outside this region, the probability of collapse toward the zero-solution attractor increases substantially.




Figure 7. Seed robustness study results. Green bars indicate convergence to the physical solution branch; red bars indicate failure to the zero-solution attractor. The consistent failure error near 5.9 confirms a structural attractor.

The results indicate that convergence reliability is controlled more by optimization dynamics than by computational scale alone. Under favourable conditions, the network converges toward the physical exterior-flow solution. Under unfavourable conditions, the optimization collapses toward a near-zero harmonic field that satisfies the dominant loss terms but not the physical solution.

6.5. Accuracy–Cost Trade-Off

The parametric study reveals an important trade-off between solution accuracy and computational expense. The highest-accuracy solutions (maximum absolute error = 0.0648) were obtained near 8000 epochs, using 10000–60000 collocation points, but required approximately 12–16 hours of training depending on hardware conditions.

Such computational requirements may be acceptable for benchmark validation studies, but may be impractical for rapid engineering evaluation.




Figure 8. Converged solution (seed 1). Top row: exact streamlines (left) and PINN-predicted streamlines with equipotential lines (right). Bottom row: absolute error (left) and local L2 error (right) distributions.




Figure 9. Failed solution (seed 3). The PINN converges to a near-zero stream function, corresponding to the trivial zero-solution attractor. Streamlines are physically inconsistent, and errors are near the maximum possible value ($\approx \max |\psi_{\text{exact}}|$).

A more computationally efficient configuration was identified at 4000 epochs with 20000 collocation points, which achieved a maximum absolute error of 0.162 while requiring substantially lower computational time and maintain a stable convergence and physically reasonable results.

These observations indicate that practical PINN deployment for exterior-domain problems requires balancing computational cost against the desired level of solution accuracy, as larger training configurations do not always yield proportionally better performance.

6.6. Implications for Future PINN Development

The observed instability patterns highlight the strong sensitivity of PINNs for exterior-domain problems to initialization and hyperparameters, indicating that conventional fully connected architectures remain vulnerable to ensuring robust convergence.

Future work should focus on stability-enhanced training strategies such as adaptive collocation, dynamic loss balancing, gradient normalization, and residual neural network architectures. Deep residual learning introduced by He et al.^[13] has demonstrated significantly improved optimization stability in deep neural networks and may help reduce collapse toward non-physical solution branches in PINN training. These methods may improve convergence reliability and reduce the probability of collapse toward non-physical solution branches. The results also emphasize the importance of reproducibility using multi-seed testing and statistical evaluation rather than relying on single-run results alone.

Overall, the present findings suggest that successful application of PINNs to unbounded exterior-domain problems depends not only on the governing equations and boundary conditions, but also on the stability characteristics of the underlying optimization process.

7. Visualization of Error Behaviour and Stability Regions

To better understand the numerical behaviour discussed in the previous section, three additional figures are presented to illustrate the observed stability patterns, failure modes, and convergence characteristics of the proposed PINN framework.

7.1. Recurrent Failure Mode (Absolute Error ≈ 5.4)

Figure 10 shows a representative solution corresponding to the recurrent high-error state (absolute error ≈ 5.4), which appeared across multiple epoch–training point combinations. The failed solutions yield physically inconsistent pressure distributions and deviate significantly from the analytical streamline structure, indicating convergence to a non-physical solution branch rather than a slightly inaccurate approximation.

The recurrence of almost identical error values and solution structures indicates that the optimization process converges to a stable but inaccurate attractor. This supports the hypothesis that the training instability is structural rather than stochastic. Such behaviour has been reported in PINN literature, where an imbalance in residual terms or gradient dominance leads to convergence toward trivial or unphysical solutions.

7.2. Epoch–Training Point Stability Map (Absolute Error Marked)

Figure 11 presents the parametric stability map obtained by systematically varying epochs ranging from 2000 to 10000 and training points from 10000 to 60000, with each plotted point labelled by its

corresponding absolute error.




Figure 10. Representative solution corresponding to the recurrent failure mode (absolute error ≈ 5.4). The near-identical structure observed across multiple configurations indicates convergence to a stable but incorrect optimization basin.

The stability map reveals a highly fragmented optimization structure. Several configurations with simultaneously large epoch count and collocation densities collapse toward the recurrent high-error attractor, while neighbouring configurations converge successfully to the physical solution branch. This behaviour confirms that convergence is strongly non-monotonic and highly sensitive to hyperparameter selection.

The most accurate solutions are consistently obtained near 8000 epochs, where stable convergence occurs across multiple collocation densities. In contrast, several configurations outside this regime fail despite using substantially larger computational resources, as shown in Fig. 12. These observations demonstrate that optimization stability plays a more important role than computational scale alone in determining PINN solution quality for exterior-domain problems.

7.3. L2 Error Distribution and Spatial Convergence Behaviour

Figure 13 illustrates the spatial distribution of the local L2 error across the computational domain. The following patterns are observed:

- The highest errors concentrate near the upper and lower far-field regions.
- Error magnitude decays smoothly away from the cylinder boundary.
- The downstream region shows moderate accumulation of error.
- Near the solid boundary, the model maintains strong enforcement of the no-penetration condition.

This distribution indicates that the PINN accurately captures near-field behaviour while exhibiting reduced accuracy in the far-field region. The global relative L2 error values further confirm the convergence trends observed in the absolute error map.




Figure 11. Stability map showing the relationship between training epochs and the number of training points. Each marker is labelled with the corresponding absolute error. A threshold-like region separates stable and unstable regimes.




Figure 12. Three-dimensional landscape of the absolute error as a function of training epochs and the number of training points. The surface illustrates the highly irregular structure of the optimization landscape, revealing both stable low-error regions and unstable regions associated with the recurrent failure mode.

7.4. Structure of the Optimization Landscape

The three-dimensional error landscapes as shown in Fig. 12 and 14 respectively depict a highly fragmented optimization space with isolated low-error regions surrounded by abrupt instability zones. This leads to a high sensitivity to hyperparameters, wherein slight variations in collocation density or epochs can result in either correct convergence or collapse to a non-physical attractor. Therefore, increased computational effort does not guarantee improved performance because solution quality is strongly determined by the interaction between loss balancing, initialization, optimizer dynamics, and collocation sampling.




Figure 13. Spatial distribution of the local L2 error in the exterior domain. Higher errors are observed in the far-field region, while strong accuracy is maintained near the cylinder boundary.




Figure 14. Three-dimensional visualization of the relative L2 error as a function of training epochs and training points. The landscape highlights the complex interaction between sampling density and training duration, demonstrating that optimal solutions occur only within specific regions of the parameter space.

The fragmented nature of the error landscape also explains the strong sensitivity observed throughout the study. Since stable low-error regions occupy only limited portions of the parameter space, successful

convergence requires selecting hyperparameter combinations that remain within these narrow stability regimes.

7.5. Emergence of a Stability-Optimal Epoch Regime

Despite the highly irregular optimization behaviour observed throughout the parametric study, one consistent pattern emerges clearly from the results: the existence of a stability-optimal epoch regime near 8000 epochs. Across multiple collocation densities, training within this region repeatedly produced the most accurate and reliable solutions obtained in the study.

Notably, stable convergence in this regime was observed even with relatively low collocation densities (5000 training points), while some higher-density cases outside it failed. This indicates that convergence is more strongly governed by training dynamics than collocation density alone. The results suggest distinct training phases: at lower epoch counts lead to an underdeveloped solution structure, while excessively large epoch counts may induce instability and convergence to non-physical attractors. The 8000-epoch regime appears to provide an effective balance between these competing effects, allowing stable convergence toward the physical branch before instability becomes dominant.

Overall, the results emphasize that finding stability-optimal training strategies is more crucial than just extending training time, as optimization stability significantly influences the quality of solutions in exterior-domain PINN problems.




Figure 15. Absolute error as a function of training epochs. The error exhibits highly irregular and non-monotonic behaviour, with abrupt transitions between low-error and failure regimes. A distinct stability-optimal region is observed near 8000 epochs.

7.6. Training Points versus Absolute Error

The three-dimensional error landscapes in Fig. 14 show that maximum absolute error and local L2 error have fragmented distributions, with isolated low-error regions surrounded by abrupt instability zones. This fragmented structure explains why nearby hyperparameter combinations can produce dramatically different outcomes despite only minor numerical changes.

These observations support the interpretation that exterior-domain PINN optimization is fundamentally multi-attractor in nature. The network does not converge toward a unique analytical solution through classical mathematical guarantees but instead evolves toward different solution branches depending on initialization, optimization trajectory, and hyperparameter selection. Under favourable conditions, in Fig.15 the physically relevant harmonic branch can be recovered consistently; under unfavourable conditions, in Fig. 16 the optimization collapses into stable but non-physical attractors.

8. Theoretical Considerations on Uniqueness in Unbounded Domains

The classical exterior potential flow problem governed by the Laplace equation presents a subtle theoretical challenge in the absence of far-field boundary conditions. In this section, we discuss the mathematical background of uniqueness in unbounded domains and provide theoretical insight into why the physics-informed neural network (PINN) formulation employed in this work can recover the physically correct solution without explicit far-field enforcement.




Figure 16. Absolute error as a function of the number of training points. The absence of monotonic convergence highlights the sensitivity of PINN optimization to sampling density and the existence of discrete stability regions.

8.1. Non-Uniqueness of the Laplace Equation in Exterior Domains

Consider the two-dimensional Laplace equation Eq. (4), where Ω denotes the exterior domain outside a circular cylinder of radius R . On the solid boundary $\partial\Omega_s$, the no-penetration condition is imposed,

$$\psi = 0 \text{ on } \partial\Omega_s. \quad (26)$$

In classical potential theory, this boundary value problem is not uniquely determined unless an additional condition is imposed at infinity. The reason is that the Laplace equation admits infinitely many harmonic solutions in an unbounded domain.

More precisely, if ψ_1 is a harmonic solution satisfying the solid boundary condition, then

$$\psi = \psi_1 + h, \quad (27)$$

where h is any harmonic function vanishing on the solid boundary, and it is also a solution without an asymptotic constraint, such as

$$\psi \sim Ur \sin(\theta) \text{ as } r \rightarrow \infty, \quad (28)$$

The solution space is not restricted to the physically meaningful uniform flow behaviour. In traditional numerical methods (finite element, finite difference, boundary element methods), this ambiguity is resolved by introducing artificial outer boundaries and prescribing far-field conditions. These additional constraints enforce asymptotic behaviour and ensure uniqueness.

The central question addressed in this work is therefore: Why does the PINN formulation recover the correct exterior potential flow solution without explicit far-field boundary enforcement?

8.2. Loss Function, Bounded Domain, and the Zero-Solution Attractor

It is important to recognize that the PINN optimization is performed on a finite, bounded computational domain rather than on the mathematically unbounded exterior domain. Consequently, the optimization landscape differs from the classical exterior boundary-value problem and may contain additional low-loss solution branches that do not correspond to the desired physical solution.

$$L_{PDE} + \lambda_{BC} \cdot L_{BC} = (1/N_f) \Sigma |\Delta \psi_\theta|^2 + (\lambda_B C / N_b) \Sigma |\psi_\theta|^2 \quad (29)$$

is a discretized approximation of the squared L^2 norm of the Laplacian plus the squared $L^2(\partial\Omega_s)$ norm of the boundary trace. On a bounded Lipschitz domain, this functional is bounded below by the squared $H^{1/2}(\partial\Omega_s)$ norm of the boundary trace via standard elliptic regularity. Importantly, this means that $\psi = 0$ is not merely a trivial harmonic function satisfying the boundary condition — it is the unique global minimizer of this two-term loss on the bounded domain. The theoretical minimum of $L_{PDE} + \lambda_{BC} \cdot L_{BC}$ is zero, achieved exactly by $\psi = 0$.

This is a precise mathematical explanation for the recurrent failure mode observed in 11 of 33 training runs. When the additional loss terms (the soft far-field penalty L_{FF} , the symmetry loss L_{sym} , and the pressure coefficient loss L_{Cp}) are not sufficient to override the gradient signal from $L_{PDE} + \lambda_{BC} \cdot L_{BC}$,

the optimizer collapses toward $\psi = 0$. This collapse arises naturally from the structure of the loss landscape.

The physically relevant solution $\psi = Ur (1 - a^2/r^2) \sin(\theta)$ is NOT the global minimizer of the two-term loss. It is selected only when the additional loss terms — particularly L_{FF} and L_{Cp} — provide sufficient gradient signal to guide the optimizer away from the zero-solution basin. This confirms that the soft far-field penalty and pressure coefficient supervision are not merely optional enhancements: they are mathematically essential for recovering the physical branch on a bounded training domain.

This understanding also clarifies the role of network initialization. Near-zero weight initialization places the initial network function close to $\psi \approx 0$, which is the global attractor of the dominant loss terms. Starting from this position makes it difficult for the optimizer to escape toward the physical branch. Standard Xavier initialization, which produces nonzero initial outputs, reduces the probability of being captured by the zero-solution basin during early training.

8.3. Role of Neural Network Inductive Bias

Another important factor in solution selection is the inductive bias of neural networks. Fully connected feedforward networks with smooth activation functions (e.g., SiLU) possess spectral bias, meaning they tend to approximate smooth, low-frequency functions before representing highly oscillatory behaviour, as shown in prior work [14]. Rahaman et al. [14] demonstrated through theoretical analysis and experiments that gradient-based training of neural networks preferentially fits low-frequency components of target functions before high-frequency ones. This frequency-ordered learning is a consequence of the gradient structure of mean squared error losses combined with the smooth activation functions used in standard architectures.

In the context of the present problem, this spectral bias has a concrete implication: among all harmonic functions satisfying the solid boundary condition, the physical $n=1$ mode $\psi \propto \sin(\theta)$ has the simplest angular structure, while higher-mode solutions $\psi_n \propto \sin(n\theta)$ oscillate n times faster in the azimuthal direction. The spectral bias of the network therefore makes the $n=1$ mode more easily representable and more likely to be selected during training, provided the other loss terms do not strongly favour a competing branch.

In the context of harmonic problems, it is observed that this bias behaviour favours:

- Smooth global solutions,

- Symmetric flow structures,
- Low-complexity harmonic fields.

Among the infinitely many harmonic functions satisfying the solid boundary condition, the physically correct exterior flow solution is the simplest global harmonic extension consistent with uniform flow symmetry. Thus, the neural network architecture itself implicitly promotes the selection of the physically meaningful branch of the solution space.

8.4. Effect of Finite Sampling in an Unbounded Domain

Although the physical domain is unbounded, training is performed over a large but finite computational region. Importantly, no boundary conditions are imposed on the outer edges of this region. However, the PDE residual is enforced throughout the sampled area. Harmonic functions exhibiting large growth over the sampled region generally produce larger residual and regularization penalties during optimization and are therefore less likely to be selected during training. As a result, such functions are disfavoured during optimization.

Therefore, the combination of:

- Interior residual minimization,
- Solid boundary anchoring,
- Finite-domain sampling,
- Neural network smoothness bias,

restricts the effective solution space and promotes convergence toward the physically correct asymptotic behaviour.

8.5. Comparison with Classical Uniqueness Requirements

Classically, uniqueness in exterior Laplace problems is ensured by imposing either:

1. A Dirichlet condition at infinity,
2. A Neumann condition at infinity,
3. A prescribed growth rate condition.

In the present PINN formulation, none of these conditions is explicitly enforced. Instead, convergence toward the physical branch occurs through optimization dynamics under favourable conditions, without

constituting mathematical uniqueness.

This distinction highlights an important conceptual difference:

- Traditional methods impose uniqueness analytically.
- PINNs promote convergence toward the physical branch through optimization dynamics, without constituting mathematical uniqueness.

The results presented in Sections 5–7 demonstrate that, for the exterior potential flow around a circular cylinder, under the specific training conditions described, this combination of factors is sufficient to select the physical solution branch, velocity field, and pressure coefficient distribution.

8.6. Limitations and Scope of Validity

This implicit branch selection mechanism may not generalize to all exterior problems. The success observed in this study is influenced by several favourable properties of the test case:

- The governing equation is linear.
- The analytical solution is smooth.
- The flow exhibits strong symmetry.
- The harmonic solution grows at most linearly at infinity.

For nonlinear or more complex external PDEs, additional constraints or architectural modification may be required to assure stability and uniqueness. Overall, the findings show that PINNs can recover the physically correct branch in classical linear exterior flows with a weak far-field penalty. However, this should be interpreted as architecture- and optimizer-dependent implicit branch selection, not as mathematical inference of far-field behaviour from the governing equation and wall condition alone.

8.7. Summary of Theoretical Insights

The theoretical discussion presented above supports the numerical findings of this study. The following conclusions are drawn:

1. The exterior Laplace equation with only a solid wall boundary condition is mathematically non-unique. Infinitely many harmonic functions, including the trivial zero solution, satisfy both the governing equation and the no-penetration condition.

2. The composite PINN loss function does not resolve this non-uniqueness analytically. The Laplacian residual loss is satisfied by any harmonic function; the far-field soft penalty provides only weak discouragement of non-physical solutions.
3. Convergence toward the physical $n=1$ branch, when observed, results from the combined effect of the soft far-field penalty, the symmetry and pressure coefficient loss terms, and the spectral bias of the network architecture toward smooth low-frequency functions.
4. The existence of a zero-solution failure mode in 11 of 33 training runs confirms that the loss landscape contains multiple stable attractors. The failure mode is identified as collapse to $\psi \approx 0$, whose maximum magnitude over the evaluation domain equals the recurring error value of 5.4.
5. These findings should be interpreted as evidence of architecture- and optimizer-dependent implicit branch selection in an underdetermined exterior problem, not as a demonstration that far-field behaviour is uniquely determined by the PDE and wall boundary condition alone.

9. Discussion on Stability-Enhanced Residual Architectures

Recent advances in physics-informed neural networks have emphasized the importance of architectural modifications to improve training stability, gradient flow, and convergence reliability. Residual-based neural network architectures have been shown to significantly enhance the robustness of deep learning models applied to partial differential equations. In this section, we briefly discuss how residual-based architectures—specifically the Squared Residual Network and the Power-Enhanced Residual Network—relate to the instability phenomena observed in the present study.

9.1. Residual Networks for Stable Weight Updating

Residual networks (ResNets) introduce skip connections that allow information to propagate directly across layers. Instead of learning a mapping $\mathcal{F}(x)$, the network learns a residual correction:

$$x^{(l+1)} = x^{(l)} + \mathcal{F}^{(l)}(x^{(l)}). \quad (30)$$

This structure improves gradient flow during backpropagation and mitigates the vanishing and exploding gradient problems common in deep architectures.

In the context of PDE learning, stable weight updating is particularly critical because:

- Higher-order derivatives (such as those required for the Laplace equation) amplify gradient instability.
- Loss functions combine multiple competing terms (PDE residual, boundary enforcement, etc.).

- Optimization landscapes are highly non-convex.

The recent works by [15] introduced residual-based modifications specifically tailored to improve weight updating stability in PINNs.

9.2. Squared Residual Network

The Squared Residual Network proposed in [15], Engineering Analysis with Boundary Elements] incorporates squared residual connections to enhance training stability. This architecture was shown to:

- Improve convergence reliability,
- Reduce sensitivity to hyperparameter variation,
- Stabilize gradient propagation in deep networks,
- Produce more consistent PDE solutions.

The recurrent failure mode observed in the present study (absolute error ≈ 5.4) suggests that the plain fully connected network used here may suffer from unstable weight updating or gradient imbalance. Given that the Laplace residual requires computation of second-order derivatives via automatic differentiation, the optimization process is inherently sensitive to weight scaling. Residual-based architectures may reduce the likelihood of collapse into the high-error basin identified in Section 7.

9.3. Power-Enhanced Residual Network

In a related study published in Applied Mathematics and Computation [16], a Power-Enhanced Residual Network was introduced to further stabilize learning of smooth and non-smooth functions in PDE contexts. The key idea of the Power-Enhanced architecture is to incorporate power terms into residual elements, effectively modifying the gradient dynamics and improving convergence behaviour. Numerical experiments in that work demonstrated:

- Superior accuracy compared to plain MLPs,
- Improved robustness under hyperparameter variation,
- Faster convergence in inverse and forward PDE problems.

The extreme sensitivity observed in the epoch-error and training-point-error curves of the present study suggests that the plain PINN architecture is operating near instability margins. The existence of discrete stability regimes and abrupt transitions between low-error and catastrophic states indicates that

gradient flow control plays a central role in solution quality. Incorporating power-enhanced residual connections may mitigate these abrupt transitions by smoothing the optimization landscape.

9.4. Interpretation of the 5.4 Failure Mode considering Residual Architectures

The repeated convergence to an absolute error of approximately 5.4 across multiple parameter combinations strongly suggests that the optimization process collapses into a structurally stable but physically incorrect attractor. This behaviour is consistent with known PINN training instabilities, including:

- Imbalance between loss components,
- Gradient domination by high-frequency modes,
- Poor weight scaling across layers,
- Vanishing or exploding gradients during backpropagation.

Residual-based architectures address precisely these issues by:

- Facilitating stable weight updating,
- Preserving gradient magnitude across layers,
- Improving training smoothness,
- Reducing sensitivity to hyperparameter fluctuations.

Therefore, it is reasonable to hypothesize that replacing the plain feedforward architecture used in this study with a Squared ResNet or Power-Enhanced Residual Network may reduce or eliminate the observed catastrophic failure regime.

9.5. Scope for Future Extension

While the present work focuses on demonstrating that exterior potential flow can be recovered without explicit far-field boundary conditions, future research should investigate the use of residual-based architectures to further improve stability.

Potential extensions include:

1. Implementing a Squared Residual PINN for the exterior Laplace problem.
2. Comparing convergence reliability between plain MLP and residual architectures.
3. Studying whether the 8000-epoch stability regime persists under residual-based training.

4. Evaluating whether the 5.4 failure attractor disappears entirely.

Such investigations would provide deeper insight into the relationship between network architecture and implicit uniqueness selection in unbounded-domain PINNs.

9.6. Concluding Remarks on Architectural Considerations

The numerical evidence presented in Sections 6 and 7 indicates that training stability plays a central role in the success of the solid-boundary-only formulation. Although the plain network architecture is sufficient to recover the correct solution under favourable hyperparameter combinations, its sensitivity highlights the need for stability-aware designs.

Residual-based architectures, as demonstrated in recent literature, offer a promising pathway toward more robust and less hyperparameter-sensitive PINN formulations. Integrating such architectures with the unbounded-domain formulation proposed in this work represents a natural and valuable direction for future investigation.

10. Conclusions

This study investigated the behaviour of physics-informed neural networks (PINNs) applied to the classical exterior potential flow around a circular cylinder under a weakly constrained formulation in which explicit hard far-field boundary conditions were replaced by weak soft far-field regularization. The governing Laplace equation was enforced throughout the exterior domain, while the no-penetration condition on the cylinder surface served as the primary physical boundary constraint.

The numerical results demonstrate that the proposed PINN formulation can recover the physically relevant exterior flow solution with good agreement to the analytical stream function, velocity field, and pressure coefficient distribution under appropriate training conditions. The predicted streamlines reproduced the expected stagnation points, global symmetry, and near-boundary flow behaviour associated with classical potential flow theory.

A major outcome of the study is the identification of implicit branch selection behaviour within the PINN optimization process. Although the exterior Laplace problem with only a solid-wall boundary condition is mathematically non-unique, the network frequently converged toward the physically meaningful uniform-flow solution branch when the combined effects of the governing-equation residual, weak soft

far-field regularization, symmetry enforcement, pressure-coefficient supervision, and neural-network spectral bias acted favourably during training.

At the same time, the parametric study revealed that the optimization landscape is highly sensitive and strongly multi-stable. Abrupt transitions between successful convergence and catastrophic failure were observed across nearby hyperparameter combinations. A recurrent failure mode with absolute error near 5.4–5.9 was consistently identified and shown to correspond to collapse toward the trivial near-zero solution, confirming the existence of competing stable attractors within the loss landscape.

The study further demonstrated that convergence reliability depends more strongly on optimization stability than on computational scale alone. Increasing the number of training points or epochs did not guarantee improved accuracy. Instead, a stability-optimal regime emerged near 8000 epochs, where accurate convergence was repeatedly achieved across multiple collocation densities. This observation suggests that careful control of optimization dynamics is essential for reliable PINN performance in exterior-domain problems.

The results also highlight important limitations of the present formulation. The observed convergence behaviour does not constitute a mathematical proof of uniqueness or analytical recovery of far-field behaviour from the governing equation and wall condition alone. The successful solutions obtained in this work should instead be interpreted as optimizer- and architecture-dependent convergence toward the physical branch under favourable training conditions. Furthermore, the study considered only a single linear benchmark problem involving smooth incompressible potential flow around a circular cylinder. The extent to which similar behaviour persists for more complex geometries, nonlinear PDEs, turbulent flows, or three-dimensional exterior problems remains uncertain.

Nevertheless, the findings provide useful insight into the behaviour of PINNs in underdetermined exterior-domain problems and demonstrate that physically meaningful solution branches can emerge even under weak asymptotic constraints. The instability patterns identified in this work also emphasize the importance of stability-aware network design, reproducibility analysis, and robust optimization strategies in future PINN research.

Potential directions for future work include the application of residual-based architectures, adaptive collocation strategies, dynamic loss-balancing techniques, and extension of the framework to nonlinear and three-dimensional exterior-flow problems. Further theoretical investigation into optimization-driven branch selection and stability mechanisms in PINNs may also provide a deeper understanding of how neural-network architectures interact with underdetermined PDE systems.

Acknowledgements

The first author acknowledges the support and academic environment provided by the Department of Civil Engineering, National Taiwan University, which facilitated the completion of this research.

References

1. [△]Milne-Thomson LM (1968). *Theoretical Hydrodynamics*. 5th ed. Macmillan.
2. [△]Batchelor GK (2000). *An Introduction to Fluid Dynamics*. Cambridge University Press.
3. [△]Anderson JD (2010). *Fundamentals of Aerodynamics*. 5th ed. McGraw-Hill.
4. [△]Streeter VL (1948). *Fluid Dynamics*. McGraw-Hill.
5. [△]Currie IG (1974). *Fundamental Mechanics of Fluids*. McGraw-Hill.
6. [△]Kirchhoff RH (1985). *Potential Flows, Computer Graphic Solutions*. Marcel Dekker.
7. [△]Raissi M, Perdikaris P, Karniadakis GE (2019). "Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations." *J Comput Phys*. 378:686–707.
8. [△]Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L (2021). "Physics-Informed Machine Learning." *Nat Rev Phys*. 3:422–440.
9. [△]Lagaris IE, Likas A, Fotiadis DI (1998). "Artificial Neural Networks for Solving Ordinary and Partial Differential Equations." *IEEE Trans Neural Netw*. 9(5):987–1000.
10. [△]Sirignano J, Spiliopoulos K (2018). "DGM: A Deep Learning Algorithm for Solving Partial Differential Equations." *J Comput Phys*. 375:1339–1364.
11. [△][△]Munson BR, Young DF, Okiishi TH (2006). *Fundamentals of Fluid Mechanics*. 5th ed. Wiley.
12. [△]Glorot X, Bengio Y (2010). "Understanding the Difficulty of Training Deep Feedforward Neural Networks." In: *Proc Thirteenth Int Conf Artif Intell Stat (AISTATS)*. 249–256.
13. [△]He K, Zhang X, Ren S, Sun J (2016). "Deep Residual Learning for Image Recognition." In: *Proc IEEE Conf Comput Vis Pattern Recognit (CVPR)*.
14. [△][△]Rahaman N, Baratin A, Arpit D, Draxler F, Lin M, Hamprecht F, Bengio Y, Courville A (2019). "On the Spectral Bias of Neural Networks." In: *Proc 36th Int Conf Mach Learn (ICML 2019)*. PMLR 97:5301–5310.
15. [△][△]Noorizadegan A, Cavoretto R, Young DL, Chen CS (2024a). "Stable Weight Updating: A Key to Reliable PDE Solutions Using Deep Learning." *Eng Anal Bound Elem*. 168:105933.

16. [△]Noorizadegan A, Young DL, Hon YC, Chen CS (2024b). "Power-Enhanced Residual Network for Function Approximation and Physics-Informed Inverse Problems." *Appl Math Comput.* **480**:128910.

Declarations

Funding: National Science and Technology Council of Taiwan under grant number 112-2221-E-002-097-MY3.

Potential competing interests: No potential competing interests to declare.