



A Hybrid Legendre Polynomial–Physics-Informed Neural Network for Predicting Solutions of Laplace's Equation in Spherical Coordinates

Eman Mohamed Ahmed¹

¹Department of Mathematics, Faculty of Education – Al-Ajaylat, University of Zawia

Email: e.mmohammed@zu.edu.ly

<https://orcid.org/0009-0008-8257-5612>

Received: 29/6/2026 Revised: 20/7/2026 Accepted: 28/7/2026 Publication: 29/8/2026

Abstract

Boundary value problems for Laplace's equation on spherical and shell-shaped domains arise in potential theory, heat conduction, electrostatics, and gravitation, with classical solutions expressed as series in spherical harmonics. Standard physics-informed neural networks (PINNs) represent such solutions with a generic, domain-agnostic multilayer perceptron that ignores the exact angular separability of the Laplacian on spherical domains. This paper examines whether embedding the known Legendre and spherical-harmonic structure of the operator into a physics-informed model changes accuracy, computational cost, and interpretability relative to a standard PINN. We propose a hybrid framework in which a compact network learns only the radial expansion coefficients $\hat{a}_{\ell m}(r)$ of a truncated spherical-harmonic series, while the angular dependence is supplied analytically through associated Legendre polynomials; this reduces the PDE residual to an exact, one-dimensional radial ODE residual per mode, removing the need to differentiate with respect to the angular coordinates. A composite loss combining PDE-residual, boundary-condition, and data terms is minimized with gradient descent. On a reproducible spherical-shell Laplace benchmark with a known closed-form multipole solution, the proposed model is trained and evaluated against a matched-budget standard PINN using RMSE, MAE, relative L_2 error, and R^2 , with ablations over degree, collocation-point count, and network capacity. The two models achieve comparable accuracy, with relative L_2 error below 4%; the proposed model trains in roughly half the time and is markedly more accurate at low capacity, while the standard PINN is marginally more accurate at larger capacity. The learned coefficients give a direct, per-degree interpretation unavailable from a standard PINN.

Index Terms— Legendre polynomials; spherical harmonics; Laplace's equation; boundary value problems; partial differential equations; physics-informed neural networks; numerical approximation; computational mathematics.

المستخلص

تعد مسائل القيم الحدية لمعادلة لابلاس في المجالات الكروية والقشرية من المسائل الأساسية في الرياضيات التطبيقية ونظرية المعادلات التفاضلية الجزئية، وتمثل حلولها التحليلية تقليدياً باستخدام التوافقيات الكروية المبنية على دوال لجندر المرتبطة، وذلك عبر فصل المتغيرات في الإحداثيات الكروية. تهدف هذه الورقة إلى دراسة أثر دمج هذا التمثيل الرياضي التحليلي المعروف مباشرة داخل نموذج تقريبي عددي قائم على الشبكات العصبية المستتيرة بالمعادلات الحاكمة - (Physics-Informed Neural Networks, PINNs)، بدلاً من الاعتماد على شبكة عامة لا تستغل البنية الزاوية التحليلية للمؤثر اللابلاسي. تقترح الورقة إطاراً هجيناً تتعلم فيه شبكة عصبية مدمجة فقط المعاملات الشعاعية $\hat{a}_{\ell m}(r)$ ضمن متسلسلة كروية توافقية مقطوعة عند أقصى درجة L ، بينما يشق الاعتماد الزاوي تحليلياً من التوافقيات الكروية $Y_{\ell}^m(\theta, \phi)$ ، بما يختزل معادلة لابلاس إلى معادلة تفاضلية عادية شعاعية أحادية البعد لكل نمط. جرب النموذج المقترح وقورن بشبكة PINNs قياسية على مسألة قيمة حدية مرجعية على قشرة كروية ذات حل تحليلي معروف، باستخدام مقاييس الخطأ التربيعي الجذري المتوسط (RMSE) والخطأ المطلق المتوسط (MAE) والخطأ النسبي L_2 ومعامل التحديد R^2 ، إلى جانب دراسات حذف لأثر أقصى درجة كروية توافقية وعدد نقاط التجميع وسعة الشبكة. أظهرت النتائج الفعلية دقة تنبؤية متقاربة بين النموذجين، إذ كانت الشبكة القياسية أدق بقليل في الإعداد الأساسي، بينما تفوق النموذج المقترح بوضوح في زمن التدريب وفي حالات

السعة الحاسوبية المنخفضة. وتوفر المعاملات الشعاعية المتعلمة تفسيرًا عدديًا مباشرًا لإسهام كل نمط كروي توافقي في الحل. الكلمات المفتاحية: حدوديات لجندر؛ التوافقيات الكروية؛ معادلة لابلاس؛ الشبكات العصبية؛ المعادلات التفاضلية الجزئية؛ التقريب العددي؛ الرياضيات التطبيقية.

I. INTRODUCTION

Boundary value problems (BVPs) governed by Laplace's equation on spherical and shell-shaped domains are among the most thoroughly studied problems in applied mathematics and the theory of partial differential equations (PDEs). Classical potential theory expresses their solutions as series in spherical harmonics built from associated Legendre polynomials, and this representation underlies applications ranging from electrostatics and heat conduction to gravitational potential modeling and geodesy. The analytical machinery — separation of variables, Legendre's differential equation, and the orthogonality of spherical harmonics on the sphere — is exact, well understood, and computationally cheap once the expansion coefficients are known from boundary data.

In parallel, physics-informed neural networks (PINNs) have emerged as a mesh-free approach to approximating solutions of partial differential equations (PDEs) by embedding the governing equation directly into the training loss of a neural network, rather than requiring it to fit only labeled solution data [1]. A standard PINN represents the unknown field with a generic multilayer perceptron (MLP) that maps spatial coordinates directly to a scalar (or vector) output, and is trained by minimizing a composite loss built from the PDE residual, boundary and initial conditions, and, when available, sparse observational data. Because the MLP architecture is domain-agnostic, it does not exploit any closed-form structure the operator may have on a particular geometry; recent surveys document rapid growth of physics-informed and scientific machine learning methods across many equation families and domains [2], [3], but comparatively few studies examine what is gained or lost when a PINN's function representation is deliberately re-parameterized to match a known analytic separation of variables on a curved domain.

Spherical domains are a natural test case for this question, because Laplace's equation separates exactly in spherical coordinates: the angular dependence of every harmonic mode is captured exactly by a spherical harmonic $Y_\ell^m(\theta, \phi)$, and the entire coordinate dependence of the solution reduces, mode by mode, to a single radial function. A generic PINN must instead learn this angular structure implicitly and approximately from automatic differentiation of a Cartesian- or spherical-coordinate MLP, at the cost of extra representational capacity, extra derivative computations at every training step, and reduced interpretability of what the network has learned.

This paper investigates whether re-introducing the known Legendre/spherical-harmonic separation structure into a physics-informed model while keeping the physics-informed training paradigm intact changes accuracy, computational cost, and interpretability relative to a standard PINN, for boundary value problems governed by Laplace's equation on a spherical shell. The central research question is not whether Laplace's equation can be solved analytically (it can); it is whether coupling the analytic angular structure with a learned radial representation yields a more accurate or more efficient physics-informed solver than treating

the whole field as an unstructured learned function, and whether the resulting per-degree coefficients offer diagnostic value that a standard PINN does not provide.

The specific contributions of this paper are as follows.

1) A hybrid mathematical-computational framework, the Legendre–Spherical-Harmonic PINN, in which a compact neural network learns only the radial expansion coefficients $\hat{a}_{\ell m}(r)$ of a truncated spherical-harmonic series, while the angular part is represented analytically through associated Legendre polynomials; the predicted field is reconstructed as $\hat{u}(r, \theta, \varphi) = \sum_{\ell m} \hat{a}_{\ell m}(r) Y_{\ell}^m(\theta, \varphi)$.

2) A physics-informed loss in which the Laplacian PDE residual is reduced — exactly, using the eigenvalue property of spherical harmonics under the angular Laplacian — to a one-dimensional radial ODE residual per mode, avoiding automatic differentiation with respect to θ and φ altogether and removing the coordinate-singularity issues that arise near the poles of a spherical coordinate system.

3) A reproducible numerical study: a spherical-shell Laplace boundary value problem with a known closed-form multipole solution is used to train and evaluate the proposed model against a matched-budget standard PINN baseline, with a fixed random seed, an independent held-out test set, and quantitative metrics (RMSE, MAE, relative L_2 error, and R^2). Ablation studies quantify the effect of the maximum spherical-harmonic degree L , the number of PDE collocation points, and network capacity on both models.

4) An up-to-date, verifiable survey of related work (2023–2026) on physics-informed neural networks, neural methods for PDEs and functions on the sphere, and Legendre/spectral-basis neural networks, together with a comparison table that situates the present study relative to the closest prior work and states the specific, defensible research gap it addresses.

The results in Section VI show that the two models are competitive rather than one uniformly dominating the other at the scale tested: the proposed model trains in roughly half the wall-clock time of the standard PINN at matched iteration budgets and is markedly more accurate in the low-capacity regime, while the standard PINN is marginally more accurate once both models are given enough capacity to fit the benchmark well. The proposed architecture's principal advantages are therefore framed as computational efficiency, sample efficiency at low capacity, and the interpretability of explicit per-degree coefficients, rather than a uniform accuracy improvement.

The remainder of the paper is organized as follows. Section II reviews the mathematical background: the Laplacian in spherical coordinates, Legendre and associated Legendre polynomials, and spherical harmonics. Section III surveys related work and states the research gap. Section IV presents the proposed Legendre–Spherical-Harmonic PINN, including its loss function. Section V describes the experimental setup: the benchmark problem, data generation, network architectures, and evaluation metrics. Section VI reports results, comparisons, and ablations. Section VII concludes and outlines future work.

II. MATHEMATICAL BACKGROUND

This section summarizes, in self-contained form, the classical mathematical results used to construct the proposed hybrid model: the Laplacian operator in spherical coordinates, separation of variables, Legendre and associated Legendre polynomials, and spherical harmonics. The material follows standard references [4], [5] and is consistent with the derivations collected in [6]; no text, structure, problem statement, or numerical result from [6] is reproduced beyond these standard identities, which are used here only as background for an independent research problem, methodology, and set of experiments.

A. Laplace's Equation and Spherical Coordinates

Let u be a twice continuously differentiable scalar field. Laplace's equation states that u is harmonic:

$$\nabla^2 u = 0 \quad (1)$$

Spherical coordinates (r, θ, ϕ) , with $r \geq 0$ the radial distance, $\theta \in [0, \pi]$ the polar (zenith) angle, and $\phi \in [0, 2\pi)$ the azimuthal angle, relate to Cartesian coordinates through $x = r \sin \theta \cos \phi$, $y = r \sin \theta \sin \phi$, $z = r \cos \theta$. In these coordinates the Laplacian operator takes the form

$$\nabla^2 u = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial u}{\partial r} \right) + \frac{1}{r^2 \sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial u}{\partial \theta} \right) + \frac{1}{r^2 \sin^2 \theta} \frac{\partial^2 u}{\partial \phi^2} \quad (2)$$

so that Laplace's equation $\nabla^2 u = 0$ on a ball or spherical shell reads

$$\frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial u}{\partial r} \right) + \frac{1}{r^2 \sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial u}{\partial \theta} \right) + \frac{1}{r^2 \sin^2 \theta} \frac{\partial^2 u}{\partial \phi^2} = 0. \quad (3)$$

B. Separation of Variables

Seeking product solutions $u(r, \theta, \phi) = R(r) \Theta(\theta) \Phi(\phi)$ and substituting into (3), the standard separation-of-variables argument [4] shows that the angular part must satisfy an eigenvalue problem with a real separation constant, conventionally written $\ell(\ell+1)$ with ℓ a non-negative integer for solutions that remain finite and single-valued on the whole sphere, while the radial factor obeys the Euler–Cauchy equation

$$r^2 R'' + 2rR' - \ell(\ell+1)R = 0, \quad (4)$$

whose general solution is

$$R(r) = C_1 r^\ell + C_2 r^{-(\ell+1)} \quad (5)$$

The two independent radial behaviors r^ℓ and $r^{-(\ell+1)}$ are, respectively, the regular (interior) and singular (exterior) solutions classically associated with multipole expansions; both are retained on a spherical shell that excludes the origin and infinity, as in the boundary value problem studied in Section V.

C. Legendre's Equation and Legendre Polynomials

With $x = \cos \theta$, the polar factor of the angular eigenproblem with azimuthal symmetry (order $m = 0$) reduces to Legendre's differential equation

$$(1 - x^2)y'' - 2xy' + \ell(\ell + 1)y = 0, \quad (6)$$

whose polynomial solutions that remain bounded on $[-1, 1]$ are the Legendre polynomials $P_\ell(x)$, obtainable from Rodrigues' formula

$$P_\ell(x) = \frac{1}{2^\ell \ell!} \frac{d^\ell}{dx^\ell} ((x^2 - 1)^\ell) \quad (7)$$

and satisfy the orthogonality relation on $[-1, 1]$,

$$\int_{-1}^1 P_\ell(x) P_k(x) dx = \frac{2}{2\ell + 1} \delta_{\ell k} \quad (8)$$

D. Associated Legendre Functions

When the azimuthal order $m \neq 0$ is retained, the polar factor instead satisfies the associated Legendre equation

$$(1 - x^2)y'' - 2xy' + \left[\ell(\ell + 1) - \frac{m^2}{(1 - x^2)} \right] y = 0, \quad (9)$$

whose regular solutions, the associated Legendre functions P_ℓ^m , are generated from the Legendre polynomials by

$$P_\ell^m(x) = (-1)^m (1 - x^2)^{\frac{m}{2}} \frac{d^m}{dx^m} P_\ell(x), \quad (10)$$

with the sign convention $(-1)^m$ following the standard Condon–Shortley phase [4]. The associated Legendre functions satisfy the orthogonality relation

$$\int_{-1}^1 P_\ell^m(x) P_k^m(x) dx = \frac{2}{2\ell + 1} \frac{(\ell + m)!}{(\ell - m)!} \delta_{\ell k} \quad (11)$$

E. Spherical Harmonics

Combining the polar factor $P_\ell^m(\cos \theta)$ with the azimuthal factor $e^{im\phi}$ (the bounded, single-valued solution of the Φ -equation) defines the complex spherical harmonics

$$Y_\ell^m(\theta, \phi) = N_{\ell m} P_\ell^m(\cos \theta) e^{im\phi}, \quad (12)$$

with normalization constant

$$N_{\ell m} = \sqrt{\frac{(2\ell + 1)(\ell - m)!}{4\pi(\ell + m)!}} \quad (13)$$

chosen so that the Y_ℓ^m form an orthonormal set on the unit sphere with respect to the solid-angle measure $d\Omega = \sin \theta \, d\theta \, d\phi$:

$$\int_0^{2\pi} \int_0^\pi Y_\ell^m(\theta, \phi) Y_{\ell'}^{m'*}(\theta, \phi) \sin \theta \, d\theta \, d\phi = \delta_{\ell\ell'} \delta_{mm'} \quad (14)$$

A further identity used throughout Sections IV and V is that the Y_ℓ^m are eigenfunctions of the angular part of the Laplacian, with eigenvalue $\ell(\ell+1)$:

$$-\left[\frac{1}{\sin \theta} \frac{\partial}{\partial \theta} \left(\sin \theta \frac{\partial Y}{\partial \theta} \right) + \frac{1}{\sin^2 \theta} \frac{\partial^2 Y}{\partial \phi^2} \right] = \ell(\ell + 1)Y, \quad (15)$$

where Y is shorthand for $Y_\ell^m(\theta, \phi)$. Equation (15) is precisely what allows the three-dimensional Laplacian to reduce to an exact one-dimensional radial operator once a field is expanded in the Y_ℓ^m basis, a fact exploited directly by the proposed architecture in Section IV.

F. General Separable Solution

Combining the radial solution (5) with the angular eigenfunctions (12), the general axisymmetric-free (multipole) solution of Laplace's equation on a spherical shell is the series

$$\sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} (A_{\ell m} r^\ell + B_{\ell m} r^{-(\ell+1)}) Y_\ell^m(\theta, \phi) \quad (16)$$

with constants $A_{\ell m}$, $B_{\ell m}$ fixed by boundary conditions. Equation (16) is the classical starting point for potential-theory boundary value problems [5] and is the basis of both the neural representation proposed in Section IV and the exact analytical benchmark solution used in Section V.

III. RELATED WORK

This section surveys work published between 2023 and 2026 on physics-informed neural networks, neural methods for PDEs and functions defined on the sphere, and Legendre/spectral-basis neural networks, and positions the present study relative to it. Foundational, pre-2023 references are cited only where a specific method or definition originates there.

A. Physics-Informed Neural Networks

PINNs were introduced by Raissi et al. as a general framework for embedding a PDE residual, evaluated by automatic differentiation, directly into the training loss of a neural network that approximates the PDE solution [1]. Subsequent work has extended PINNs to a very wide range of equations, loss formulations, and training strategies; recent surveys catalogue this growth and its open challenges, including training pathologies caused by stiff or multi-scale residual landscapes, spectral bias, and the difficulty of enforcing hard constraints exactly [2], [3]. None of these general surveys is specific to spherical geometries or to spectral/Legendre-basis representations, which is the gap this paper addresses.

B. Neural Methods for PDEs and Functions on the Sphere

A smaller body of work targets PDEs or functions defined intrinsically on the sphere. Lei et al. propose physics-informed convolutional neural networks for solving PDEs on the sphere and provide an approximation-theoretic analysis based on spherical harmonic decompositions of the network's error, together with numerical experiments [7]. Bonev et al. introduce Spherical Fourier Neural Operators (SFNOs), which perform spectral convolutions using the spherical harmonic transform instead of the planar Fourier transform used by standard Fourier Neural Operators, and demonstrate stable long-horizon autoregressive rollouts for global atmospheric dynamics [8]. Zhong and Meidani propose a physics-informed, geometry-aware neural operator (PI-GANO) that conditions a physics-informed operator network on an encoded domain geometry so that a single trained model generalizes across varying geometries and PDE parameters without labeled solution data [9]. These three studies share the present paper's physics-informed or spectral-operator orientation, but none of them explicitly parameterizes a physics-informed loss through per-degree associated-Legendre radial coefficients for a spherical boundary value problem; SFNO and PI-GANO instead learn convolutional or operator-network weights on top of a spherical or geometry-encoded latent representation.

A separate line of work uses spherical harmonics as a feature representation for supervised learning rather than for solving a physics-informed BVP. Rußwurm et al. combine spherical harmonics with sinusoidal representation networks to build a location encoder for geographic coordinates on the globe, improving downstream geospatial prediction tasks relative to naive rectangular encodings [10]. Hendriks et al. parameterize the diffusion MRI signal on each voxel's angular sphere with a neural spherical-harmonics representation to obtain a spatially and angularly coherent, denoised signal reconstruction [11]. Both works demonstrate the practical value of coupling neural networks with an explicit spherical-harmonic basis, but in a supervised representation-learning setting with no PDE residual or physics-informed loss, which is a different problem class from the boundary value problem addressed here.

C. Legendre and Spectral-Basis Neural Networks

Closer to the radial-coefficient idea used in this paper, Ali proposes Legendre spectral neural networks that use Legendre polynomials as basis functions inside a physics-informed scheme for solving elliptic PDEs, reporting improved accuracy and convergence relative to standard PINNs on one- and two-dimensional elliptic benchmarks [12]. That work uses a one- or two-dimensional Cartesian Legendre expansion and does not address the three-dimensional spherical geometry, the coupled Legendre–spherical-harmonic angular structure, or the degree/order (ℓ , m) coefficient interpretation central to the present study. In a related but distinct direction, Hao et al. develop a neural-network method for the integral fractional Laplacian, a nonlocal generalization of the classical Laplace operator, using a different (non-spectral) network parameterization tailored to the fractional-order kernel [13]. Table I summarizes these and the other most closely related studies.

TABLE I. Comparison of closely related studies (2023–2026)

Year	Study	Math./AI Approach	PDE / Problem Type	Domain	Gap Relative to This Study
2025	Lei et al. [7]	Physics-informed CNN + spherical-harmonic error analysis	General PDEs	Unit sphere (2-D manifold)	No explicit per-degree Legendre radial-coefficient network; convolutional, not coefficient-based
2023	Bonev et al. [8]	Spherical Fourier Neural Operator (spectral convolution)	Atmospheric dynamics PDEs	Sphere (global weather)	Data-driven operator learning, not physics-informed residual minimization; no radial 3-D BVP
2025	Zhong & Meidani [9]	Physics-informed geometry-aware neural operator	Parametric elliptic PDEs	Varying 2-D/3-D geometries	Geometry-conditioned operator network; no spherical-harmonic angular decomposition
2024	Rußwurm et al. [10]	Spherical harmonics + SIREN location encoder	Supervised geospatial regression/classification	Sphere (Earth surface)	Representation learning, no PDE residual or BVP
2023	Hendriks et al. [11]	Neural spherical-harmonic signal model	Diffusion MRI signal reconstruction	Angular sphere per voxel	Signal-fitting objective, not a physics-informed PDE loss
2025	Ali [12]	Legendre spectral physics-informed NN	Elliptic PDEs	1-D / 2-D Cartesian intervals	No 3-D spherical geometry or spherical-harmonic angular coupling
2023	Hao et al. [13]	Neural network for fractional Laplacian	Integral fractional Laplace equation	Bounded Euclidean domains	Different (nonlocal, fractional-order) operator; no spectral radial-coefficient PINN
2019	Raissi et al. [1]	Original physics-informed neural network	General PDEs (forward/inverse)	General Euclidean domains	Domain-agnostic MLP; no analytic angular structure exploited
<i>Proposed study (this paper) — shown for reference, not a previously published work</i>					
—	Proposed study	Legendre–spherical-harmonic radial-coefficient PINN	Laplace equation (BVP)	Spherical shell	Combines exact angular eigenstructure with a learned radial PINN and systematic ablations

Table entries above the divider are ordered by first appearance in the text; see the References section for full citations. The final, shaded row is the study proposed in this paper and is listed for direct comparison only.

D. Research Gap and Positioning

The literature reviewed above shows active, recent interest in three overlapping directions: (i) physics-informed neural solvers in general [1]–[3]; (ii) neural or operator-learning methods that operate natively on the sphere, either physics-informed [7], [9] or purely data-driven and spectral [8]; and (iii) spherical-harmonic or Legendre feature representations used for supervised or signal-reconstruction tasks [10]–[12]. To the best of the authors' knowledge from this review, no publicly available, peer-reviewed or preprint study combines an explicit associated-Legendre/spherical-harmonic radial-coefficient parameterization with a physics-informed loss for the Laplace equation on a three-dimensional spherical boundary value problem, together with a controlled, matched-budget comparison to a standard PINN and ablations over spherical-harmonic degree, collocation density, and network capacity. This is

stated as a specific, narrowly scoped gap relative to the reviewed literature, not as a claim that the general idea of combining spectral bases with neural PDE solvers is unprecedented — Ali's Legendre spectral neural networks [12] and Lei et al.'s spherical-harmonic error analysis of physics-informed CNNs [7] are close in spirit and are treated as the nearest prior art throughout this paper.

Accordingly, the goal of this study is not to claim a categorically new capability, but to empirically examine, under a controlled and reproducible protocol, whether introducing this specific Legendre–spherical-harmonic structure into a physics-informed model changes the accuracy, computational efficiency, and interpretability of the learned solution to a spherical Laplace BVP relative to an architecture-matched standard PINN.

IV. PROPOSED LEGENDRE–SPHERICAL HARMONIC PINN

This section presents the proposed hybrid model. The central design decision is to change what the neural network is asked to learn: instead of mapping spatial coordinates directly to the scalar field u , the network maps the radial coordinate r to a vector of spherical-harmonic expansion coefficients, and the field itself is reconstructed analytically from those coefficients and the closed-form spherical-harmonic basis. Fig. 1 summarizes the resulting pipeline, from the spherical coordinate domain through to error evaluation.

A. Representation

Let L denote a chosen maximum spherical-harmonic degree. The proposed model approximates the solution of Laplace's equation by the truncated expansion

$$\hat{u}(r, \theta, \phi; w) = \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} \hat{a}_{\ell m}(r; w) Y_{\ell}^m(\theta, \phi), \quad (17)$$

where w denotes the trainable parameters of a single small feed-forward neural network $N_w: \mathbb{R} \rightarrow \mathbb{R}^{(L+1)^2}$, whose scalar input is r and whose output vector collects every coefficient $\hat{a}_{\ell m}(r; w)$ for $0 \leq \ell \leq L$ and $-\ell \leq m \leq \ell$; the basis functions $Y_{\ell}^m(\theta, \phi)$ are the fixed, closed-form spherical harmonics of (12), not learned. In this sense the network is not asked to approximate a general trivariate function of (r, θ, ϕ) ; it is asked to approximate only $(L+1)^2$ univariate functions of r , with all angular dependence supplied analytically. Once trained, the predicted field at any (r, θ, ϕ) is obtained by evaluating $N_w(r)$ and forming the weighted sum in (17); the coefficients $\hat{a}_{\ell m}(r; w)$ are directly interpretable as the (learned) radial profile of the energy carried by spherical-harmonic mode (ℓ, m) , analogous to a physics-informed multipole spectrum.

For the real-valued numerical implementation in Sections V and VI, the complex basis $\{Y_{\ell}^m\}$ of (12) is replaced by an equivalent real orthonormal basis $\{Y_{\ell 0}, Y_{\ell m}^c, Y_{\ell m}^s\}$ obtained from the standard cosine/sine combination of $Y_{\ell}^{\pm m}$; this is a unitary change of basis that spans the same function space for each degree ℓ and preserves the orthonormality and eigenvalue properties of Sections II-E and IV-B, while allowing $\hat{a}_{\ell m}(r; w)$ and the network output to

remain real-valued throughout. The angular eigenvalue property that (18)-(19) below rely on holds identically for this real basis.

B. Physics-Informed Loss

Substituting the expansion (17) into the Laplacian (2) and using linearity together with the radial-derivative rule and the spherical-harmonic eigenvalue property (15) - which holds exactly for the closed-form Y_ℓ^m regardless of what $\hat{a}_{\ell m}(r;w)$ is - the three-dimensional PDE residual collapses to a sum of one-dimensional radial residuals, one per mode:

$$R_{\ell m}(r; w) = \hat{a}_{\ell m}''(r; w) + \frac{2}{r} \hat{a}_{\ell m}'(r; w) - \frac{\ell(\ell + 1)}{r^2} \hat{a}_{\ell m}(r; w), \quad (18)$$

so that the pointwise PDE residual of the reconstructed field is the exact, closed-form combination

$$R(r, \theta, \phi; w) = \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} R_{\ell m}(r; w) Y_\ell^m(\theta, \phi) \quad (19)$$

Equations (18)-(19) are the key computational advantage of the proposed architecture: because the angular part is handled analytically rather than through automatic differentiation, computing the PDE residual requires only first and second derivatives of each $\hat{a}_{\ell m}$ with respect to the single scalar r , never derivatives with respect to θ or ϕ , and is therefore free of the coordinate-singularity behavior that a generic PINN can encounter near the poles $\theta \rightarrow 0, \pi$ of a spherical coordinate system.

Given N_f collocation points $\{(r_i, \theta_i, \phi_i)\}$ sampled in the interior of the domain, N_b boundary points $\{(r_{bj}, \theta_{bj}, \phi_{bj})\}$ on the Dirichlet boundary, and (optionally) N_d labeled reference points $\{(r_{dk}, \theta_{dk}, \phi_{dk}), u_{dk}\}$ - used to test the framework's ability to assimilate sparse data alongside the physics constraint - the model is trained by minimizing a composite loss with three terms. The PDE-residual loss is

$$L_{pde} = \frac{1}{N_f} \sum_{i=1}^{N_f} (R(r_i, \theta_i, \phi_i; w))^2 \quad (20)$$

The boundary-condition loss enforces the (Dirichlet) boundary data $g(\theta, \phi)$:

$$L_{bc} = \frac{1}{N_b} \sum_{j=1}^{N_b} (\hat{u}(r_{bj}, \theta_{bj}, \phi_{bj}; w) - g(\theta_{bj}, \phi_{bj}))^2 \quad (21)$$

and, when reference values are available, the data loss anchors the prediction to labeled points:

$$L_{data} = \frac{1}{N_d} \sum_{k=1}^{N_d} (\hat{u}(r_{dk}, \theta_{dk}, \phi_{dk}; w) - u_{dk})^2 \quad (22)$$

and the three terms are combined into the total training objective

$$L(w) = \lambda_{pde} L_{pde} + \lambda_{bc} L_{bc} + \lambda_{data} L_{data}, \quad (23)$$

with non-negative weights λ_{pde} , λ_{bc} , λ_{data} set once and held fixed for both models compared in this study (Section V-D). $L(w)$ is minimized with a first-order stochastic gradient method; Section V details the optimizer, sampling scheme, and training budget.

C. Baseline for Comparison

To isolate the effect of the Legendre–spherical-harmonic representation itself, the proposed model is compared against a standard PINN that is physics-informed in exactly the same sense - trained with a PDE-residual, boundary, and data loss of the same functional form as (20)-(22) - but whose function representation is a single generic MLP mapping Cartesian coordinates directly to the scalar field, $M_w: \mathbb{R}^3 \rightarrow \mathbb{R}$, $\hat{u}(x,y,z;w) = M_w(x,y,z)$, with the Laplacian evaluated by automatic differentiation in Cartesian coordinates, $\nabla^2 \hat{u} = \frac{\partial^2 \hat{u}}{\partial x^2} + \frac{\partial^2 \hat{u}}{\partial y^2} + \frac{\partial^2 \hat{u}}{\partial z^2}$. This is the natural, structure-agnostic baseline: it uses the same training paradigm, loss weights, optimizer, and (as closely as architecture differences allow) parameter budget as the proposed model, differing only in whether the spherical-harmonic angular structure of Section II is built into the representation or left for the network to infer implicitly.




Fig. 1. Proposed Legendre–Spherical–Harmonic PINN framework.

V. METHODOLOGY AND EXPERIMENTAL SETUP

This section specifies the benchmark boundary value problem, the data-generation procedure, the two model architectures, the training protocol, the evaluation metrics, and the ablation protocol used to produce the results reported in Section VI, with sufficient detail for the experiments to be reproduced.

A. Benchmark Boundary Value Problem

The benchmark domain is a spherical shell $\Omega = \{(r, \theta, \phi): r_{in} \leq r \leq r_{out}, 0 \leq \theta \leq \pi, 0 \leq \phi < 2\pi\}$ with $r_{in} = 0.5$ and $r_{out} = 1.0$, so that both radial solution branches r^ℓ and $r^{-(\ell+1)}$ in (5) are regular on Ω and the boundary value problem has a unique classical solution. An exact reference solution is constructed directly from the general series (16), restricted to a fixed finite set S of five (ℓ, m) modes spanning degrees $\ell = 0$ through $\ell = 3$ so that the maximum-degree ablation (Section VI-C) has a genuine, non-trivial truncation effect for $L < 3$:

$$u_{exact}(r, \theta, \phi) = \sum_{(\ell, m) \in S} (A_{\ell m} r^\ell + B_{\ell m} r^{-(\ell+1)}) Y_\ell^m(\theta, \phi), \quad (24)$$

with S and the coefficients $(A_{\ell m}, B_{\ell m})$ listed in Table II. The Dirichlet boundary data used for training and evaluation are obtained directly by evaluating (24) at $r = r_{in}$ and $r = r_{out}$, i.e., $g(\theta, \phi) = u_{exact}(r_{in} \text{ or } r_{out}, \theta, \phi)$.

TABLE II. Modes and coefficients defining the exact benchmark solution (24)

ℓ	m	$A_{\ell m}$	$B_{\ell m}$
0	0	1.00	0.05
1	0	0.80	0.02
1	1	0.50	0.01
2	1	0.30	0.01
3	-1	0.20	0.005

B. Data Generation

PDE collocation points $\{(r_i, \theta_i, \phi_i)\}$ are sampled independently with $r_i \sim \text{Uniform}(r_{in}, r_{out})$, $\cos \theta_i \sim \text{Uniform}(-1, 1)$ (the correct measure for a uniform angular distribution on the sphere, avoiding pole clustering), and $\phi_i \sim \text{Uniform}(0, 2\pi)$. Boundary points are generated the same way with r fixed at r_{in} for half the points and r_{out} for the other half. A small set of interior data points is generated identically to the collocation points, independently reseeded, and labeled with u_{exact} from (24), to test the data-assimilation term (21). Finally, an independent held-out test set of $N_{test} = 2000$ points is drawn from a separate random-number stream (seed 12345, disjoint from every training seed) and used only for evaluation; it is never seen during training or model selection.

C. Network Architectures

The standard PINN baseline (Section IV-C) is a fully connected MLP with Cartesian input (x, y, z) , tanh hidden activations, and a single linear scalar output $\hat{u}$. The proposed model is a fully connected MLP with scalar input r , tanh hidden activations, and a linear output layer of

width $(L+1)^2$, producing every $\hat{a}_{\ell m}(\mathbf{r})$ simultaneously; both networks share the same hidden width and depth in every experiment reported in Section VI unless that quantity is itself the ablated variable. Table III lists the two architectures' input/output dimensions and how their parameter count is determined by hidden width H and depth (number of hidden layers) n_h .

TABLE III. Architecture summary

Model	Input	Output dim.	Hidden layers	Activation
Standard PINN	$(x, y, z) \in \mathbb{R}^3$	1 ($\hat{u}$)	n_h layers, width H	tanh
Proposed model	$\mathbf{r} \in \mathbb{R}$	$(L+1)^2$ ($\hat{a}_{\ell m}$)	n_h layers, width H	tanh

Every dense layer of input size a and output size b contributes exactly $ab + b$ parameters; exact total parameter counts for every (H, n_h, L) configuration used in Section VI are reported alongside each result (Tables IV and VII).

D. Training Protocol

Both models minimize the composite loss (23) with $\lambda_{pde} = 1$, $\lambda_{bc} = 5$, and $\lambda_{data} = 5$ (the boundary and data terms are up-weighted because they involve far fewer points than the PDE-residual term and would otherwise be dominated by it). Parameters are initialized with a Glorot-uniform-scaled Gaussian and optimized with Adam using a fixed learning rate (2×10^{-3} for the standard PINN, 5×10^{-3} for the proposed model, chosen by a small preliminary search so that neither model is handicapped by an obviously poor learning rate) and no learning-rate schedule. All baseline results in Section VI-B use random seed 42, $N_f = 500$ collocation points, $N_b = 150$ boundary points, $N_d = 30$ data points, hidden width $H = 24$, depth $n_h = 2$, and 3000 training iterations; ablation runs (Section VI-C-VI-E) use a reduced iteration budget (1500-2000 iterations) to keep the total computational cost of the ablation grid tractable, as stated explicitly with each result.

The automatic differentiation required for both the PDE residual (2)-(3) (Cartesian second derivatives for the baseline; the radial second derivative in (18) for the proposed model) and for backpropagating the resulting loss to every network weight was implemented with a purpose-built reverse-mode automatic-differentiation engine in NumPy, together with a forward-mode (dual-number) recursion for the required derivatives, since the compute environment used for these experiments had no outbound network access to install a GPU-oriented autodiff framework. The engine's gradients were validated against central finite differences prior to any reported experiment, for both first- and second-order derivatives and for backpropagation through loss terms that themselves depend on second derivatives, with maximum absolute error below 10^{-5} in every check.

E. Evaluation Metrics

Let $\hat{u}_p$, for $p = 1, \dots, N$, denote model predictions on the held-out test set and u_p the corresponding exact values from (24), with $\bar{u}$ the mean of $\{u_p\}$. Four metrics are reported: the root-mean-square error,

$$RMSE = \sqrt{\frac{1}{N} \sum_{p=1}^N (\hat{u}_p - u_p)^2} \quad (25)$$

the mean absolute error,

$$MAE = \frac{1}{N} \sum_{p=1}^N [|\hat{u}_p - u_p|] \quad (26)$$

the relative L_2 error,

$$Rel. L_2 = \frac{\sqrt{\sum_{p=1}^N (\hat{u}_p - u_p)^2}}{\sqrt{\sum_{p=1}^N (u_p)^2}} \quad (27)$$

and the coefficient of determination,

$$R^2 = 1 - \frac{\sum_{p=1}^N (\hat{u}_p - u_p)^2}{\sum_{p=1}^N (u_p - \bar{u})^2} \quad (28)$$

Lower RMSE, MAE, and relative L_2 error indicate better agreement with the exact solution; R^2 closer to 1 indicates that the model explains a larger share of the variance of the exact solution over the test set.

F. Ablation Protocol

Three ablations are reported in Section VI, each varying one factor while holding the others at the baseline configuration of Section V-D (with the reduced iteration budget noted above). Ablation 1 varies the maximum spherical-harmonic degree $L \in \{1, 2, 3\}$ of the proposed model only, since L is a hyperparameter specific to that architecture; because the exact solution (24) contains an $\ell = 3$ mode, $L < 3$ introduces a genuine representational truncation. Ablation 2 varies the number of PDE collocation points $N_f \in \{100, 300, 600, 1000\}$ for both models. Ablation 3 varies network capacity, jointly changing hidden width and depth $(H, n_h) \in \{(8,1), (16,2), (32,3)\}$ for both models. All ablation runs use the same fixed seed (42) and the same independent test set as the baseline comparison.

VI. RESULTS AND DISCUSSION

All results below were produced by the code and protocol of Section V, with the fixed random seed stated in each case.

A. Baseline Comparison

Table IV reports RMSE, MAE, relative L_2 error, and R^2 on the independent $N = 2000$ held-out test set, together with parameter count and total wall-clock training time, for the

standard PINN and the proposed model trained under the baseline protocol of Section V-D ($H = 24$, $n_h = 2$, $N_f = 500$, 3000 iterations, seed 42).

TABLE IV. Baseline comparison: Standard PINN vs. proposed model

Model	RMSE	MAE	Rel. L_2 (%)	R^2	Params.	Train time (s)
Standard PINN	0.01359	0.00892	3.67	0.9962	721	42.2
Proposed model	0.01412	0.00983	3.81	0.9959	1048	21.9

Both models fit the benchmark closely (relative L_2 error below 4%, R^2 above 0.995). The standard PINN is marginally more accurate at this specific, matched-capacity configuration, while the proposed model trains in 21.9 s versus 42.2 s for the standard PINN - roughly 52% of the wall-clock time - despite having more parameters (1048 vs. 721). This speed advantage follows directly from the architectural design of Section IV: the standard PINN's Cartesian Laplacian requires three separate second-derivative sweeps of the network (one per Cartesian direction) at every collocation point, whereas the proposed model differentiates a much narrower (single-input) coefficient network only with respect to the scalar r , and the angular part of every derivative is supplied analytically through (15) at no additional automatic-differentiation cost. Fig. 2 shows the total training loss versus iteration for both models on a logarithmic scale; both converge over roughly two orders of magnitude, with intermittent loss spikes for both models that are a known symptom of Adam optimization on physics-informed objectives whose residual involves second derivatives, rather than a divergence, since the loss recovers within a few dozen iterations each time and the overall trend is monotonically decreasing.




Fig. 2. Training convergence: total loss (log scale) vs. training iteration, standard PINN vs. proposed model (baseline configuration, seed 42).




Fig. 3. Actual vs. predicted field values on the $N = 2000$ held-out test set for both models; the dashed line is the identity $y = x$.

Fig. 3 plots predicted versus exact field values on the held-out test set for both models; both scatter plots lie tightly along the identity line with no visible systematic bias, consistent with the R^2 values in Table IV.

B. Qualitative Error Analysis

Fig. 4 shows the exact field, both models' predictions, and the resulting absolute-error maps over the full angular domain (θ, ϕ) at the mid-shell radius $r = 0.75$. Both models reproduce the two-lobe angular pattern of the exact solution (set by the $\ell = 1$ and $\ell = 2$ modes in Table II) closely; the largest absolute errors for both models occur near $\theta \approx 20^\circ$ - 40° and $\theta \approx 140^\circ$ - 160° , i.e., away from the equator, where the higher-degree ($\ell = 2, 3$) angular structure is most pronounced and hardest to fit exactly with a finite training budget. The spatial error patterns of the two models differ in detail (the proposed model's largest error region is more concentrated near $\phi \approx 60^\circ$ and $\phi \approx 280^\circ$) but are of comparable magnitude, consistent with the aggregate metrics of Table IV.

A Hybrid Legendre Polynomial–Physics-Informed Neural Network for Predicting Solutions of Laplace's Equation in Spherical Coordinates




Fig. 4. Angular maps at fixed mid-shell radius $r = 0.75$: exact field, standard-PINN prediction, proposed-model prediction (top row), and absolute-error maps for both models (bottom row).

C. Ablation: Maximum Spherical-Harmonic Degree L

Table V reports the effect of the maximum degree L of the proposed model's spherical-harmonic truncation on test-set accuracy, with all other settings fixed at the reduced ablation budget of Section V-F ($N_f = 500$, 2000 iterations, seed 42). Because the exact benchmark solution (24) contains an $\ell = 3$ mode (Table II), any $L < 3$ makes the proposed representation structurally incapable of reproducing that mode exactly, regardless of training.

TABLE V. Ablation on maximum spherical-harmonic degree L (proposed model)

L	RMSE	MAE	Rel. L_2 (%)	R^2	Train time (s)
1	0.06973	0.05765	18.84	0.9000	12.4
2	0.03700	0.02992	10.00	0.9718	13.7
3	0.01568	0.01154	4.24	0.9949	15.1

The trend is clean and monotonic: relative L_2 error falls from 18.84% at $L = 1$ to 10.00% at $L = 2$ to 4.24% at $L = 3$, and R^2 rises from 0.900 to 0.972 to 0.995. This is the expected signature of a representational truncation error dominating the total error at small L , and it directly confirms that the proposed architecture's accuracy is governed, in an interpretable and predictable way, by how many spherical-harmonic degrees are retained - a diagnostic that is not directly available from a standard PINN's undifferentiated scalar output.

D. Ablation: Number of Collocation Points

Table VI reports the effect of the number of PDE collocation points N_f on test accuracy for both models, with the reduced ablation training budget (1500 iterations, seed 42) and all other settings fixed at baseline.

TABLE VI. Ablation on number of PDE collocation points N_f

N_f	Std. RMSE	Std. Rel. L_2 (%)	Std. R^2	Prop. RMSE	Prop. Rel. L_2 (%)	Prop. R^2
100	0.02258	6.10	0.9895	0.02045	5.53	0.9914
300	0.01641	4.44	0.9945	0.02431	6.57	0.9879
600	0.01921	5.19	0.9924	0.02161	5.84	0.9904
1000	0.02157	5.83	0.9904	0.02085	5.63	0.9911

At the reduced, fixed 1500-iteration budget used for this ablation, neither model shows a clean monotonic improvement with more collocation points; relative L_2 error stays within a comparatively narrow 4.4%-6.6% band across a full order-of-magnitude range of N_f (100 to 1000) for both models. This indicates that, for this smooth benchmark problem, the bottleneck at a fixed training budget is optimization progress rather than collocation density - beyond a few hundred points, adding more collocation points increases the per-iteration computational cost (Table VI training times, not reproduced here for brevity, scale roughly linearly with N_f) without a correspondingly large accuracy gain within a fixed number of gradient steps. This should be read as a property of the fixed-iteration-budget protocol used for this ablation, not as evidence that collocation density is unimportant in general; the baseline comparison in Table IV, which uses a larger training budget, achieves lower error than any configuration in Table VI.

E. Ablation: Network Capacity

Table VII reports the effect of jointly varying hidden width H and depth n_h for both models, with $N_f = 500$, 1500 iterations, seed 42.

TABLE VII. Ablation on network capacity (hidden width H , depth n_h)

(H, n_h)	Std. Rel. L_2 (%)	Std. R^2	Std. Params.	Prop. Rel. L_2 (%)	Prop. R^2	Prop. Params.
(8, 1)	13.08	0.9518	41	9.88	0.9725	160
(16, 2)	6.56	0.9879	353	8.46	0.9798	576
(32, 3)	4.02	0.9955	2273	4.05	0.9954	2704

This ablation shows the clearest qualitative difference between the two architectures. In the smallest configuration, $(H, n_h) = (8, 1)$ with only 41-160 parameters, the proposed model is markedly more accurate than the standard PINN (9.88% vs. 13.08% relative L_2 error; R^2 0.973 vs. 0.952): with very little capacity to spend, encoding the correct angular structure analytically leaves the small coefficient network free to spend its entire capacity on the one-dimensional radial profiles, while the standard PINN must spend part of its limited capacity discovering the angular structure from scratch. At the intermediate configuration (16, 2) the standard PINN is more accurate (6.56% vs. 8.46%), and by the largest configuration (32, 3) the two models are essentially tied (4.02% vs. 4.05%, R^2 0.9955 vs. 0.9954). Taken together with the collocation-density ablation, this suggests that the proposed architecture's inductive bias is most valuable in the low-capacity, low-data regime, and that its advantage narrows or disappears once the standard PINN is given enough capacity and training budget to discover the angular structure on its own - a nuanced finding rather than a uniform superiority claim, and one that is consistent with the general expectation that structural priors help most when the model or the data are limited.

F. Discussion and Limitations

Across all experiments, the proposed Legendre-spherical-harmonic PINN is competitive with a matched-budget standard PINN on the spherical-shell Laplace benchmark, and offers three concrete, measured advantages: (i) faster training at matched iteration counts, because the analytic angular reduction (18)-(19) removes the need to automatically differentiate with respect to θ and φ ; (ii) higher accuracy in the low-capacity regime (Table VII); and (iii) directly interpretable per-degree radial coefficients $\hat{a}_{\ell m}(r)$, which a scalar-output standard PINN does not expose. Its main measured limitation relative to the standard PINN is that, once both models have sufficient capacity and training budget, the standard PINN is marginally more accurate on this particular smooth, low-degree benchmark (Table IV), and the proposed model's accuracy is capped by the chosen truncation degree L (Table V) in a way the standard PINN's is not.

Several limitations bound the scope of these conclusions. First, the benchmark solution (24) is smooth and uses only five low-degree modes; results on solutions with sharper angular features, higher-degree content, or non-smooth boundary data may differ and are not tested here. Second, both models were trained on a single CPU-class environment with a custom NumPy automatic-differentiation engine rather than a GPU-accelerated framework (Section V-D), so absolute wall-clock numbers should be read as relative comparisons between the two models under identical implementation and hardware conditions, not as an estimate of performance on production PINN toolchains. Third, network sizes (tens to a few thousand parameters) and collocation counts (hundreds to low thousands of points) are modest by the standards of recent large-scale physics-informed and operator-learning studies [7], [8]; whether the trends reported here persist at that scale is an open question. Fourth, only Laplace's equation (a linear, source-free elliptic PDE) is studied; extending the radial-ODE reduction of Section IV-B to Poisson's equation, the Helmholtz equation, or other operators that do not separate as cleanly in spherical coordinates would require additional derivation and is left to future work (Section VII).

VII. CONCLUSION AND FUTURE WORK

This paper studied whether embedding the exact Legendre/spherical-harmonic angular structure of Laplace's equation into a physics-informed neural network changes accuracy, computational cost, and interpretability relative to a standard, domain-agnostic PINN. The proposed Legendre-Spherical Harmonic PINN reparameterizes the physics-informed representation so that a compact neural network learns only radial expansion coefficients $\hat{a}_{\ell m}(r)$, while the angular dependence is supplied analytically by the closed-form spherical harmonics; this reduces the three-dimensional PDE residual to an exact, per-mode, one-dimensional radial ODE residual, removing the need to differentiate with respect to the angular coordinates and the associated coordinate-singularity concerns near the poles.

On a reproducible spherical-shell Laplace boundary value problem with a known closed-form multipole solution, the proposed model and a matched-budget standard PINN were trained and evaluated with a fixed random seed and an independent held-out test set, using RMSE, MAE, relative L_2 error, and R^2 . Both models achieved relative L_2 error below 4% at

the baseline configuration. The proposed model trained in roughly half the wall-clock time of the standard PINN at matched iteration counts, and was markedly more accurate in the low-capacity regime (Table VII), while the standard PINN was marginally more accurate once both models were given sufficient capacity. Ablation studies showed a clean, monotonic accuracy improvement in the proposed model as the maximum retained spherical-harmonic degree L increased toward the true generating degree of the benchmark solution (Table V), directly confirming that the truncation degree governs the hybrid model's accuracy in a predictable, interpretable way not directly available from a standard PINN's scalar output.

These findings support a specific, bounded claim rather than a general superiority claim: for Laplace boundary value problems on spherical domains, coupling a physics-informed loss with an explicit Legendre/spherical-harmonic radial-coefficient representation is a competitive and, in the low-capacity and low-training-budget regime, favorable alternative to a standard PINN, with concrete benefits in training efficiency and coefficient interpretability. The results reported here were obtained on a single benchmark family of modest scale, as discussed in Section VI-F, and should be read accordingly.

Future work includes: (i) extending the radial-ODE reduction of Section IV-B from Laplace's equation to the Poisson and Helmholtz equations, and to operators that do not separate exactly in spherical coordinates, where an approximate or hybrid analytic-numeric angular treatment would be required; (ii) using a genuinely complex-valued coefficient network to avoid the real-basis restriction adopted in Section IV-A and directly test whether it affects accuracy; (iii) an adaptive or learned degree-selection scheme that grows L during training instead of fixing it in advance; (iv) repeating the comparisons of Section VI on GPU-accelerated PINN toolchains and at larger network and collocation-point scale, and against the spherical convolutional and operator-learning baselines of [7] and [8] directly, rather than only against a standard MLP-based PINN; (v) extending the benchmark family beyond smooth, low-degree multipole solutions to cases with sharper angular features or non-smooth boundary data; and (vi) uncertainty quantification for the learned radial coefficients, to accompany point predictions with calibrated error estimates.

REFERENCES

- [1] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686-707, 2019, doi: 10.1016/j.jcp.2018.10.045.
- [2] Z. Ren, S. Zhou, D. Liu, and Q. Liu, "Physics-informed neural networks: A review of methodological evolution, theoretical foundations, and interdisciplinary frontiers toward next-generation scientific computing," *Appl. Sci.*, vol. 15, no. 14, art. 8092, 2025, doi: 10.3390/app15148092.
- [3] C. Meng, S. Griesemer, D. Cao, S. Seo, and Y. Liu, "When physics meets machine learning: A survey of physics-informed machine learning," *Mach. Learn. Comput. Sci. Eng.*, vol. 1, art. 20, 2025, doi: 10.1007/s44379-025-00016-0.
- [4] G. B. Arfken, H. J. Weber, and F. E. Harris, *Mathematical Methods for Physicists*, 7th ed. Waltham, MA, USA: Academic Press, 2013.
- [5] J. D. Jackson, *Classical Electrodynamics*, 3rd ed. New York, NY, USA: Wiley, 1998.

- [6] I. M. Mohamed, "Legendre polynomials and spherical harmonics with their applications," M.Sc. thesis, Basic Sciences School, Dept. of Mathematics, Applied Mathematics Div., The Libyan Academy, Tripoli, Libya, 2019.
- [7] G. Lei, Z. Lei, L. Shi, C. Zeng, and D.-X. Zhou, "Solving PDEs on spheres with physics-informed convolutional neural networks," *Appl. Comput. Harmon. Anal.*, vol. 74, art. 101714, 2025, doi: 10.1016/j.acha.2024.101714.
- [8] B. Bonev, T. Kurth, C. Hundt, J. Pathak, M. Baust, K. Kashinath, and A. Anandkumar, "Spherical Fourier neural operators: Learning stable dynamics on the sphere," in *Proc. 40th Int. Conf. Machine Learning (ICML)*, Honolulu, HI, USA, 2023, arXiv:2306.03838.
- [9] W. Zhong and H. Meidani, "Physics-informed geometry-aware neural operator," *Comput. Methods Appl. Mech. Eng.*, vol. 434, art. 117540, 2025, doi: 10.1016/j.cma.2024.117540.
- [10] M. Rußwurm, K. Klemmer, E. Rolf, R. Zbinden, and D. Tuia, "Geographic location encoding with spherical harmonics and sinusoidal representation networks," in *Proc. Int. Conf. Learning Representations (ICLR)*, Vienna, Austria, 2024, arXiv:2310.06743.
- [11] T. Hendriks, A. Vilanova, and M. Chamberland, "Neural spherical harmonics for structurally coherent continuous representation of diffusion MRI signal," in *Proc. Computational Diffusion MRI Workshop (MICCAI)*, 2023, arXiv:2308.08210.
- [12] I. Ali, "Advanced machine learning technique for solving elliptic partial differential equations using Legendre spectral neural networks," *Electron. Res. Arch.*, vol. 33, no. 2, pp. 826-848, 2025, doi: 10.3934/era.2025037.
- [13] Z. Hao, M. Park, and Z. Cai, "Neural network method for integral fractional Laplace equations," *East Asian J. Appl. Math.*, vol. 13, no. 1, pp. 95-118, 2023, doi: 10.4208/eajam.010122.210722.