

# Herglotz-NET: Implicit Neural Representation of Spherical Data with Harmonic Positional Encoding

Théo Hanon<sup>1</sup>, Nicolas Mil-Homens Cavaco<sup>1</sup>, John Kiely<sup>2</sup>, Laurent Jacques<sup>1</sup>

<sup>1</sup>: INMA/ICTEAM, ULouvain, Belgium; <sup>2</sup>: Yale University, CT, USA.

**Abstract**—Representing and processing data in spherical domains presents unique challenges, primarily due to the curvature of the domain, which complicates the application of classical Euclidean techniques. Implicit neural representations (INRs) have emerged as a promising alternative for high-fidelity data representation; however, to effectively handle spherical domains, these methods must be adapted to the inherent geometry of the sphere to maintain both accuracy and stability. In this context, we propose Herglotz-NET (HNET), a novel INR architecture that employs a harmonic positional encoding based on complex Herglotz mappings. This encoding yields a well-posed representation on the sphere with interpretable and robust spectral properties. Moreover, we present a unified expressivity analysis showing that any spherical-based INR satisfying a mild condition exhibits a predictable spectral expansion that scales with network depth. Our results establish HNET as a scalable and flexible framework for accurate modeling of spherical data.

**Index Terms**—Herglotz, spherical harmonics, spectral analysis, implicit neural representation.

## I. INTRODUCTION

Over the past few years, implicit neural representations (INRs), such as SIREN [1] and NeRF [2], have attracted significant attention as a powerful alternative to classical, explicitly discretized representations of continuous signals. In these models, neural networks learn the mapping from spatial coordinates to signal values, yielding differentiable and memory-efficient representations. These architectures have found applications in areas such as continuous 3D scene reconstruction [2], data compression [3], as well as in solving partial differential equations and various inverse problems [1, 4]. Despite their success in Euclidean settings, extending INRs to spherical domains poses additional challenges as they are prone to generate artifacts, especially near the poles. To enhance the capabilities of INRs in representing signals on the sphere, [5] has proposed to incorporate spherical harmonics (SH) into the positional encoding (PE) of INRs (cf. Sec. II). However, their approach, SPH-SIREN, requires closed-form evaluations (or recursive computation) of a potentially large list of spherical harmonics, which can be computationally intensive and prone to numerical instability at high orders.

E-mails: theo.hanon@student.uclouvain.be, {nicolas.mil-homens, laurent.jacques}@uclouvain.be, john.kiely@yale.edu. Computational resources have been provided by the supercomputing facilities of UCLouvain (CISM) and the Consortium des Equipements de Calcul Intensif en Fédération Wallonie Bruxelles (CECI) funded by the Fond de la Recherche Scientifique de Belgique (FRS-FNRS). Part of this work is funded by the Belgian F.R.S.-FNRS (PDR project QuadSense; T.0160.24).

In this work, we propose Herglotz-NET (HNET), a new implicit neural representation for spherical data, as an alternative to SPH-SIREN [5] that does not require explicit evaluations of spherical harmonics (cf. Sec. III). In particular, we propose to replace the Euclidean periodic functions used in INRs by the exponentiation of an Herglotz mapping [6, Sec. VII.7], a *harmonic* PE. In Sec. IV, we generalize part of the spectral analysis of Euclidean INR in [7] to Spherical INRs, providing insights on the expressive power of both HNET and SPH-SIREN [5] in function of their depth. Finally, in Sec. V, we demonstrate in two experiments—a super-resolution application and the estimation of a continuous spherical Laplacian map—that, despite the simplicity of its PE, HNET’s performance is on a par with that of SPH-SIREN, both schemes being superior to non-spherical SIREN for the Laplacian estimation.

## II. RELATED WORK

Typically, an implicit neural representation (INR)  $f_\Theta : \mathcal{E} \rightarrow \mathbb{K}$  approximates a function  $f : \mathcal{E} \rightarrow \mathbb{K}$  (with  $\mathbb{K} = \mathbb{R}$  or  $\mathbb{C}$ , and  $\mathcal{E} \subset \mathbb{R}^d$ ) from its sampling  $\mathcal{D} = \{(\mathbf{x}_i, y_i = f(\mathbf{x}_i))\}_{i=1}^N$  on a set of  $N$  points  $\{\mathbf{x}_i\}_{i=1}^N \subset \mathcal{E}$ . This is done by *training* the set of parameters  $\Theta$  on  $\mathcal{D}$  so that  $f_\Theta$  interpolates  $f$ , i.e.,  $f_\Theta(\mathbf{x}_i) \approx y_i$ . An INR combines a positional encoding (PE) composed of  $n$  neurons

$$\psi_{\sin}(\mathbf{x}) = \left( \sin(\boldsymbol{\omega}_i^\top \mathbf{x} + b_i) \right)_{i=1}^n, \quad (1)$$

which extracts  $n$  Fourier features from the input coordinates  $\mathbf{x} \in \mathcal{E}$ , with a multilayer perceptron (MLP) whose activations are typically ReLU or sine wave. Frequencies  $\boldsymbol{\omega}_i$  and bias  $b_i$  in the PE can be either considered as trainable weights or generated randomly, leading then to the so-called Random Fourier Features [8]. This periodic positional encoding layer of the INR network is key for accurate representations of a signal [9] as it enables the representation of high-frequency signal features. The final architecture of an INR reads

$$f_\Theta(\mathbf{x}) := \mathbf{W}^{(Q)}(\phi^{(Q-1)} \circ \dots \circ \phi^{(1)})(\psi(\mathbf{x})), \quad (2)$$

with the PE  $\psi = \psi_{\sin}$ ,  $\phi^{(q)}(\cdot) := \sigma^{(q)}(\mathbf{W}^{(q)} \cdot + \mathbf{b}^{(q)})$ ,  $q = 1, \dots, Q-1$ , and where each layer  $q$  is characterized by a weight matrix  $\mathbf{W}^{(q)} \in \mathbb{R}^{d_q \times d_{q-1}}$ , a bias vector  $\mathbf{b}^{(q)} \in \mathbb{R}^{d_q}$ , and an activation function  $\sigma^{(q)} : \mathbb{R} \rightarrow \mathbb{R}$  applied element-wise.

Moreover, these architectures benefit from a comprehensive understanding of their underlying spectral properties [7], namely that the frequency support of the initial positional

encoding entirely governs the expressive capacity of the model. In addition, assuming the activation function admits a Taylor series expansion, each layer in the network expands the frequency spectrum by the order of the activation function's expansion. For networks employing sinusoidal activation functions (e.g., SIRENs [1]) this expansion is theoretically infinite due to the infinite Taylor series, while in practice most of the coefficients of the polynomial expansion decay at a factorial rate so that some terms become negligible up to some order.

As they do not consider the spherical geometry, conventional INRs, such as SIRENs, can generate singular spherical functions prone to sampling artifacts, particularly at the poles. In [5], the authors addressed these limitations by introducing SPH-SIREN; a SIREN network integrating spherical harmonics (SH)  $Y_{\ell m}(\theta, \varphi)$  with order  $\ell \geq 0$  and moment  $|m| \leq \ell$  [6] into its PE, i.e., they set the PE  $\psi = \psi_{\text{SH}}$  in (2) with

$$\mathbf{x}(\theta, \varphi) \in \mathbb{S}^2 \mapsto \psi_{\text{SH}}(\mathbf{x}) = \{Y_{\ell m}(\theta, \varphi) : |m| \leq \ell\}_{\ell=0}^{L_0}.$$

Our investigation revealed that, for orders above  $L_0 > 25$ , accurate closed-form expressions of SH in SPH-SIREN become impractical, while possible recursive estimation methods require high-precision arithmetic, limiting scalability.

### III. AN ALTERNATIVE SPHERICAL INR ON $\mathbb{S}^2$

We propose an alternative spherical INR relying on the Herglotz formalism [6, Sec. VII.7]. Specifically, given a complex vector  $\mathbf{a} \in \mathbb{C}^3$  satisfying the Herglotz condition  $\mathbf{a}^\top \mathbf{a} = 0$  (e.g., the vector  $(1, i, 0)$ ) and any twice complex-differentiable function  $f : \mathbb{C} \rightarrow \mathbb{R}$ , the function  $g(\mathbf{x}) := f(\mathbf{a}^\top \mathbf{x})$ , composing  $f$  with the Herglotz mapping  $\mathbf{x} \mapsto \mathbf{a}^\top \mathbf{x}$ , is a harmonic function, i.e., its Laplacian  $\Delta$  vanishes everywhere,  $\Delta g(\mathbf{x}) = 0$ .

Interestingly, when restricted on the 2-sphere  $\mathbb{S}^2$ , the function  $\mathbf{x} \mapsto \exp(\mathbf{a}^\top \mathbf{x})$  is a generating function of all spherical harmonics [6]. In fact, given some  $\omega_0 > 0$  and defining  $\mathbf{a}_{\Re}, \mathbf{a}_{\Im} \in \mathbb{R}^3$  the real and imaginary parts of  $\mathbf{a}$ , on  $\mathbb{S}^2$ , the function

$$g(\mathbf{x}) = \exp(\omega_0(\mathbf{a}^\top \mathbf{x})), \quad \mathbf{x} \in \mathbb{S}^2, \quad (3)$$

is centered on  $\mathbf{a}_{\Re}/\|\mathbf{a}_{\Re}\|$  with oscillations locally pointing in the direction of  $\mathbf{a}_{\Im}$ , and a frequency increasing with  $\omega_0$  (see Fig. 1). This can be seen by first noting that  $\mathbf{a}_{\Re}^\top \mathbf{a}_{\Im} = 0$  from the Herglotz condition. Moreover, since  $\mathbf{a}^\top \mathbf{x} = \mathbf{a}_{\Re}^\top \mathbf{x} + i\mathbf{a}_{\Im}^\top \mathbf{x}$ ,  $g(\mathbf{x}) = |g(\mathbf{x})| \exp(i\omega_0(\mathbf{a}_{\Im}^\top \mathbf{x}))$ , with a maximal amplitude  $|g(\mathbf{x})| = \exp(\omega_0(\mathbf{a}_{\Re}^\top \mathbf{x}))$  on  $\mathbf{x} \propto \mathbf{a}_{\Re}$ .

**Herglotz-NET:** From the above analysis, we propose an alternative INR for spherical data that preserves SIREN's flexibility and scalability while providing a well-defined framework on  $\mathbb{S}^2$ . Specifically, we define a Herglotz-NET (or HNET) by adapting the INR positional encoding  $\psi$  in (2) using the mapping (3). Given  $n$  Herglotz vectors  $(\mathbf{a}_1, \dots, \mathbf{a}_n)^\top =: \mathbf{A} \in \mathbb{C}^{n \times 3}$ , with  $\mathbf{a}_i^\top \mathbf{a}_i = 0$  for all  $i \in [n]$ ,  $n$  trainable weights  $(w_1, \dots, w_n) =: \mathbf{w}$  and biases  $(b_1, \dots, b_n) =: \mathbf{b}$ , the PE is composed of  $n$  neurons

$$\psi_{\mathbf{A}, \mathbf{w}, \mathbf{b}}(\mathbf{x}) := \{\psi_{\mathbf{a}_i, w_i, b_i}(\mathbf{x}) = e^{\omega_0(w_i \mathbf{a}_i^\top \mathbf{x} + b_i)}\}_{i=1}^n \quad (4)$$




Fig. 1. Representation of the real part of the Herglotz atom  $g(\mathbf{x})$  for  $\mathbf{a} = (1, 1, 1)/\sqrt{3} + i(1, -1, 0)/\sqrt{2}$ , and different values of  $\omega_0$ . The function  $g$  is centered on  $\mathbf{a}_{\Re} = (1, 1, 1)/\sqrt{3}$  with oscillations locally oriented along the direction  $\mathbf{a}_{\Im} = (1, -1, 0)/\sqrt{2}$  and frequency proportional to  $\omega_0$ .




Fig. 2. The Herglotz-NET architecture using the  $n$ -neuron PE defined by (4).

encoding the location  $\mathbf{x}(\theta, \varphi) \in \mathbb{S}^2$ , with  $(\theta, \varphi) \in [0, \pi] \times [0, 2\pi)$ . Consequently, each neuron corresponds to a harmonic function if  $\mathbf{x} \in \mathbb{R}^3$  if not restricted to  $\mathbb{S}^2$ . As for SIREN in (2), we add scaling factor  $\omega_0 \in \mathbb{R}$  into the PE. This can drastically affect training—yielding faster convergence and enhanced stability when properly tuned. The PE is subsequently fed to an MLP with periodic activations  $\sigma^{(q)}(\cdot) = \sin^{(q)}(\cdot)$  (see Fig. 2).

**HNET initialization:** We initialize all the weights and biases of HNET using the scheme proposed in SIREN [1]. Regarding the PE, we select  $n$  i.i.d. *unit* random Herglotz vectors  $\mathbf{a}_i \sim \text{i.i.d. } \mathbf{a}$ , with  $2\mathbf{a}_{\Re}$  picked uniformly at random on  $\mathbb{S}^2$  (i.e.,  $\|\mathbf{a}_{\Re}\| = 1/2$ ) and  $2\mathbf{a}_{\Im}$  picked uniformly at random on the circle  $\mathbb{S}^1$  in the plane orthogonal to  $2\mathbf{a}_{\Re}$  (with also  $\|\mathbf{a}_{\Im}\| = 1/2$ ). As one can show that the distribution of  $\mathbf{a}^\top \mathbf{x} \in \mathbb{C}$  is unchanged if  $\mathbf{x} \rightarrow \mathbf{R}\mathbf{x}$ , with  $\mathbf{R} \in SO(3)$  an arbitrary rotation, no neurons  $e^{\omega_0(w_i \mathbf{a}_i^\top \mathbf{x})}$  favors, in expectation, a specific point of  $\mathbb{S}^2$  or a specific direction of oscillation.

**Spectral analysis:** One can easily understand the spectral properties of the function generated by each neuron of the PE, i.e., its spectrum in the spherical harmonics basis  $Y_{\ell m}(\theta, \varphi)$ . Indeed, neglecting the biases, a Taylor series expansion gives

$$\exp(\omega_0 w_i (\mathbf{a}_i^\top \mathbf{x})) = \sum_{\ell \geq 0} \frac{\omega_0^\ell}{\ell!} w_i^\ell (\mathbf{a}_i^\top \mathbf{x})^\ell, \quad (5)$$

where each term  $(\mathbf{a}_i^\top \mathbf{x})^\ell$  activates a unique spectral component at order  $\ell$ , as stated in Lemma 1 below.

**Lemma 1.** *Given a vector  $\mathbf{a} \in \mathbb{C}^3$ , with  $\mathbf{a}^\top \mathbf{a} = 0$  and  $\ell \in \mathbb{N}$  and  $\|\mathbf{a}\|_2 = 1$ , the function  $h_\ell(\mathbf{x}) = (\mathbf{a}^\top \mathbf{x})^\ell$  is harmonic in  $\mathbb{R}^3$  and its restriction to all  $\mathbf{x} = \mathbf{x}(\theta, \varphi) \in \mathbb{S}^2$  reads*

$$h_\ell(\mathbf{x}) = \sum_{|m| \leq \ell} c_{\ell m}(\mathbf{a}) Y_{\ell m}(\theta, \varphi),$$

with  $c_{\ell m}(\mathbf{a}) \in \mathbb{C}$  such that  $\sum_{|m| \leq \ell} |c_{\ell m}|^2 \leq 4\pi(2\ell + 1)$ .

*Proof.* Given  $\mathbf{x} = r\mathbf{u} \in \mathbb{R}^3$  with  $r = \|\mathbf{x}\|$  and  $\mathbf{u} \in \mathbb{S}^2$ , since  $h_\ell$  is homogeneous of degree  $\ell$  (i.e.,  $h_\ell(\mathbf{x}) = r^\ell h_\ell(\mathbf{u})$ ), its Laplacian in  $\mathbb{R}^3$  reads

$$\Delta h_\ell(\mathbf{x}) = r^{\ell-2} (\ell(\ell+1)h_\ell(\mathbf{u}) + \Delta_{\mathbb{S}^2} h_\ell(\mathbf{u})).$$




Fig. 3. (left) Earth representation on the fine grid ( $L = 300$ ). (right) Difference between this representation and HNET upsampled on the same grid.

By definition of  $\mathbf{a}$ ,  $h_\ell$  is harmonic (i.e.,  $\Delta h_\ell(\mathbf{x}) = 0$ ), and  $\Delta_{\mathbb{S}^2} h_\ell(\mathbf{u}) = -\ell(\ell+1)h_\ell(\mathbf{u})$ , and  $h_\ell(\mathbf{u})$  can thus be expanded in the basis  $\{Y_{\ell m}\}_{m=-\ell}^{\ell}$  of spherical harmonics of degree  $\ell$ . Moreover, since  $\|\mathbf{a}\| = 1$ ,  $|h_\ell(\mathbf{x})| \leq 1$  for  $\mathbf{x} \in \mathbb{S}^2$ , and  $\sum_{|m| \leq \ell} |c_{\ell m}|^2 \leq 4\pi(2\ell+1)$  using Parseval's theorem [6].  $\square$

From this lemma, we conclude that each neuron  $\psi_{\mathbf{a}_i, w_i, b_i}(\mathbf{x})$  in (4), has bounded SH coefficients, i.e.,  $|\hat{\psi}_{\ell m}| \leq \frac{\omega_0}{\ell!} w_i^\ell |c_{\ell m}| \|\mathbf{a}_i\|^\ell$ , with  $\hat{\psi}_{\ell m} := \langle \psi_{\mathbf{a}_i, w_i, b_i}, Y_{\ell m} \rangle$ . Since  $\ell! \geq e(\ell/e)^\ell$ , the spectrum of each neuron, computed with  $S_\psi(\ell) := (2\ell+1)^{-1} \sum_{|m| \leq \ell} |\hat{\psi}_{\ell m}|^2$ , is bounded by  $S_\psi(\ell) \leq \bar{S}_\psi(\ell) := \frac{4\pi}{e^2} (e\omega_0 w_i \|\mathbf{a}_i\|)^{2\ell}$ , with  $\bar{S}_\psi(\ell)$  reaching a maximum around  $\ell \approx e\omega_0 w_i \|\mathbf{a}_i\|$ , before to quickly decay to 0 when  $\ell$  increases.

This shows that each neuron spectrum is essentially band-limited and characterized by a dominant order around  $\ell \propto \omega_0 w_i \|\mathbf{a}_i\|$ , which is consistent with the frequency of the oscillatory pattern in Fig. 1. The fixed  $\omega_0 \in \mathbb{R}$  and the trainable  $w_i \in \mathbb{C}$  then allow the model to adapt itself to various spectrum configurations.

#### IV. SPECTRUM EXPANSION OF SPHERICAL INRS

In this section, we build on the work of Yüce et al. [7] by examining the spectral expansion of spherical-based INRs. Our goal is to shed light on the empirical performance of both HNET and SPH-SIREN [5] and to offer insights on the underlying properties of these models.

Our development relies on the key assumption that the HNET exponential functions in (4) and the subsequent MLP activations in both HNET and SPH-SIREN can be accurately approximated by truncated Taylor series. In particular, we approximate the exponential as  $\exp(z) \approx \sum_{\ell=0}^{L_0} \beta_\ell z^\ell$ , for a certain degree  $L_0$ , and represent the MLP activation functions as  $\sigma(z) \approx \sum_{k=1}^K \alpha_k z^k$ . This assumption is motivated by the observation that many popular activation functions have rapidly decaying polynomial expansion coefficients—a property evident in  $\exp(z)$  and, by extension,  $\sin(z)$ .

This assumption allows us to reformulate the output of the first layer following the PE  $\psi(\mathbf{x})$  as a finite spectral expansion. In particular, for the  $j$ -th neuron we obtain

$$[\mathbf{W}^{(1)}\psi(\mathbf{x}) + \mathbf{b}^{(1)}]_j = \sum_{\ell=0}^{L_0} \sum_{|m| \leq \ell} c_{j\ell m} Y_{\ell m}(\theta, \varphi) \quad (6)$$

for some  $c_{j\ell m} \in \mathbb{C}$ . Notably, Eq. (6) is exact for SPH-SIREN, whereas it is only approximately fulfilled for HNET, given

Lemma 1. The following theorem formalizes the expansion property and delineates the expressivity limits of these architectures.

**Theorem 1.** Let  $f_\Theta : \mathbb{S}^2 \rightarrow \mathbb{C}$  be an INR defined as in (2) with PE  $\psi(\mathbf{x})$  satisfying (6),  $Q \geq 1$  layers and polynomial activation functions  $\sigma^{(q)}(z) = \sum_{k=0}^K \alpha_k z^k$ . Then, there exist coefficients  $\kappa_{\ell m} \in \mathbb{C}$  such that

$$f_\Theta(\theta, \varphi) = \sum_{\ell=0}^{K^{Q-1}L_0} \sum_{|m| \leq \ell} \kappa_{\ell m} Y_{\ell m}(\theta, \varphi) \quad (7)$$

*Proof.* See Appendix.  $\square$

With Theorem 1 in mind, we see that both HNET and SIREN rely on the frequency content of their positional encodings to determine their expressive power. In HNET, the learnable parameters  $w_i \in \mathbb{C}$  enables the network to adaptively select from a wide range of frequencies, which makes the model capable of representing an unbounded spectrum.

#### V. NUMERICAL EXPERIMENTS

In this section, we present a comprehensive evaluation of our proposed HNET and compare its performance with two baselines: a conventional SIREN [1] and its spherical variant SPH-SIREN [5]. Our experiments use Earth data from EARTH2014 representation available in PySHTools datasets.

**Data & Preprocessing:** The training data consists on the evaluation of the EARTH2014 spherical harmonic representation on a Gauss-Legendre (GL) grid [10] with  $L = 150$ . For testing, a finer GL grid with  $L = 300$  is used. Prior to model training, the training data is standardized.

**Model Architectures:** We compare three models. First, HNET with a PE based on 50 Herglotz mappings, followed by 3 hidden layers with 100 features each, and uses a scaling parameter  $\omega_0 = 10$  (#param = 30601). Secondly, SIREN with a positional encoding with 100 neurons, 3 hidden layers with 100 features each, and a first layer scaling parameter  $\omega_0 = 10$  (#param = 30700). Lastly, SPH-SIREN with a positional encoding containing the spherical harmonics up to  $L = 10$  (resulting in 110 neurons), 2 hidden layers with 100 features each and we incorporate a modified activation for the first hidden layer given by  $\sin(\omega_0 x)$  with  $\omega_0 = 3$  (#param = 22401).

**Training:** Each model was trained for 2000 epochs with a batch size of 2048, using the Adam optimizer [11] with a fixed learning rate of  $10^{-3}$ .

##### A. Super-Resolution

This experiment is conducted in a super-resolution context. After training the model on a coarse Gauss-Legendre grid [10] (with  $L = 150$ ), we evaluate its generalization capability on a finer grid ( $L = 300$ ). Fig. 3 illustrates the residuals between the HNET output and the ground truth on the  $L = 300$  grid. Quantitatively, HNET, SIREN, SPH-SIREN achieve SNR values of 22.26 dB, 22.25dB and 22.29dB, respectively. Visual inspection reveals that the residual patterns of all three models

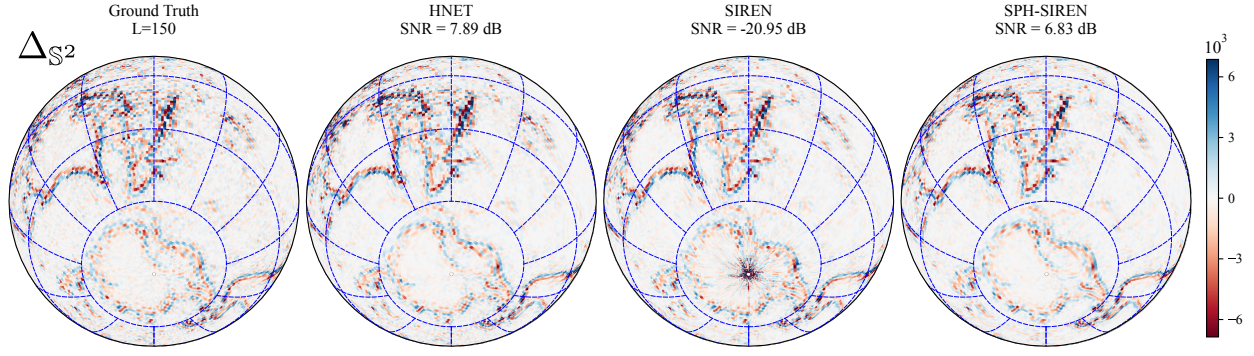


Fig. 4. Comparison of the Laplacian of the Earth representation computed with spherical harmonics up to  $L = 150$  (left) and the Laplacian of the models HNET, SIREN, and SPH-SIREN (from left to right). The ground truth Laplacian is computed using the spherical harmonic transform, while the model Laplacians are obtained via automatic differentiation of the continuous functions.

are nearly identical. Overall, all the three architectures perform equivalently on this task.

### B. Laplacian Reconstruction

In this experiment, we evaluate the ability of the models to accurately reconstruct the spherical Laplacian on the training grid (i.e.,  $L = 150$ ); for consistency, we employ the same trained models from the previous experiment.

Notably, while the super-resolution results indicated that all models perform equivalently, the Laplacian reconstruction reveals distinct differences. As shown in Fig. 4, SIREN—although visually accurate in reconstructing the Laplacian away from the poles—suffers from significant artifacts near the poles that hinder its SNR. In contrast, both HNET and SPH-SIREN exhibit no visible artifacts at the poles and faithfully reconstruct the Laplacian.

In summary, although SIREN remains an effective fitting tool on spherical domains, it loses the well-posed differentiability that is crucial for applications such as PDE solving and inverse problems. In contrast, both HNET and SPH-SIREN maintain well-posed differentiability, making them better suited for such tasks. Our results demonstrate that HNET effectively combines SIREN’s fitting speed and capacity with the well-posedness inherent to SPH-SIREN on spherical domain.

## VI. CONCLUSION

In this paper, we developed a novel implicit neural representation, HNET, adapted to spherical data. By leveraging the Herglotz formalism to construct a harmonic positional encoding, HNET circumvents the need for explicit spherical harmonic evaluations. Our experiments have shown that HNET provides performances that are on par with the SPH-SIREN architecture that explicitly use spherical harmonics. Moreover, we demonstrated that, assuming a polynomial approximation of the activation functions, the spectral expressivity of both HNET and SPH-SIREN scales with the network depth, enabling the model to capture a rich (possibly) infinite set of frequencies on the sphere.

## APPENDIX : PROOF OF THEOREM 1

The proof of Theorem 1 requires the following lemma. We also use the handy notation  $\sum_{(\ell,m)}^{[L']} \equiv \sum_{\ell=0}^{L'} \sum_{|m| \leq \ell}$  (for  $L' \in \mathbb{N}$ ).

**Lemma 2.** *Given a complex value polynomial  $p_K(z) = \sum_{k=1}^K a_k z^k$  of degree  $K \geq 0$  and an order  $L$ , for any spherical function  $f(\theta, \varphi) := \sum_{\ell=0}^L \sum_{|m| \leq \ell} c_{\ell m} Y_{\ell m}(\theta, \varphi)$  (with  $c_{\ell m} \in \mathbb{C}$ ) with bandlimit  $L$ ,  $p_K \circ f$  has a bandlimit  $KL$ , i.e., there exist coefficients  $\kappa_{\ell m} \in \mathbb{C}$  such that*

$$p_K(f(\theta, \varphi)) = \sum_{\ell=0}^{KL} \sum_{|m| \leq \ell} \kappa_{\ell m} Y_{\ell m}(\theta, \varphi).$$

*Proof.* By induction, we start by observing that  $K = 0$  is trivial since  $Y_{0,0}(\theta, \varphi)$  is constant. Next, for general  $K > 0$ , the proof only needs to consider monomials of degree  $K + 1$  since  $p_{K+1}(z) = a_{K+1}z^{K+1} + p_K(z)$ . We can develop

$$\begin{aligned} f(\theta, \varphi)^{K+1} &= \sum_{(\ell,m)}^{[KL]} \sum_{(\ell',m')}^{[L]} \kappa_{\ell m} c_{\ell' m'} Y_{\ell m}(\theta, \varphi) Y_{\ell' m'}(\theta, \varphi) \\ &= \sum_{(\ell,m)}^{[KL]} \sum_{(\ell',m')}^{[L]} \kappa_{\ell m} c_{\ell' m'} \sum_{(k,p)}^{[\ell+\ell']} \alpha_{kp} Y_{kp}(\theta, \varphi) \\ &= \sum_{(\ell,m)}^{[(K+1)L]} \hat{\kappa}_{\ell m} Y_{\ell m}(\theta, \varphi), \end{aligned}$$

where the second line uses the contraction rule of SHs [12, Sec. 5.8], the coefficients  $\alpha_{kp} \in \mathbb{C}$  involve Wigner 3- $j$  symbols, and  $\hat{\kappa}_{\ell m} \in \mathbb{C}$ .  $\square$

To prove Theorem 1, we first consider a bias-free scenario ( $\mathbf{b}^{(q)} = \mathbf{0}$  for all  $q = 1, \dots, Q - 1$ ). The  $i$ -th component of the pre-activation input of the second layer reads

$$\begin{aligned} z_i^{(2)} &= \sum_{j=1}^{d^{(2)}} W_{ij}^{(2)} \sigma^{(1)}([\mathbf{W}^{(1)} \psi(\mathbf{x}) + \mathbf{b}^{(1)}]_j) \\ &= \sum_{j=1}^{d^{(2)}} W_{ij}^{(2)} \sigma^{(1)}\left(\sum_{(\ell,m)}^{[L_0]} c_{j\ell m} Y_{\ell m}(\theta, \varphi)\right) \\ &= \sum_{(\ell,m)}^{[KL_0]} \left(\sum_{j=1}^{d^{(2)}} W_{ij}^{(2)} \kappa_{j\ell m}\right) Y_{\ell m}(\theta, \varphi) \\ &= \sum_{(\ell,m)}^{[KL_0]} \hat{\kappa}_{i\ell m} Y_{\ell m}(\theta, \varphi), \end{aligned}$$

with  $\hat{\kappa}_{i\ell m} \in \mathbb{C}$ . The second line comes from (6), and the third one from Lemma 2 for some  $\kappa_{j\ell m} \in \mathbb{C}$ . The result of Theorem 1 follows by applying Lemma 2 for each layer, that is,  $(Q - 1)$  times. Adding biases shifts the constant term of the polynomial activation function. Consequently, biases do not alter the above reasoning, which concludes the proof.



## REFERENCES

- [1] V. Sitzmann et al. “Implicit neural representations with periodic activation functions”. In: *Proceedings of the 34th International Conference on Neural Information Processing Systems*. NIPS ’20. Vancouver, BC, Canada: Curran Associates Inc., 2020.
- [2] B. Mildenhall et al. “NeRF: representing scenes as neural radiance fields for view synthesis”. In: *Commun. ACM* 65.1 (Dec. 2021), pp. 99–106.
- [3] E. Dupont et al. “COIN: COmpression with Implicit Neural representations”. In: *Neural Compression: From Information Theory to Applications–Workshop@ ICLR 2021*. 2021.
- [4] R. Liu et al. “Recovery of continuous 3d refractive index maps from discrete intensity-only measurements using neural fields”. In: *Nature Machine Intelligence* 4.9 (2022), pp. 781–791.
- [5] M. Rußwurm et al. “Geographic Location Encoding with Spherical Harmonics and Sinusoidal Representation Networks”. In: *Proceedings of the International Conference on Learning Representations (ICLR)*. 2024.
- [6] R. Courant and D. Hilbert. *Methods of Mathematical Physics*. Vol. 1. New York: Wiley, 1962.
- [7] G. Yüce et al. “A Structured Dictionary Perspective on Implicit Neural Representations”. In: *arXiv preprint arXiv:2112.01917* (2022).
- [8] A. Rahimi and B. Recht. “Random Features for Large-Scale Kernel Machines”. In: *Neural Information Processing Systems*. 2007.
- [9] M. Tancik et al. “Fourier features let networks learn high frequency functions in low dimensional domains”. In: *Advances in neural information processing systems* 33 (2020), pp. 7537–7547.
- [10] M. A. Price and J. D. McEwen. “Differentiable and accelerated spherical harmonic and Wigner transforms”. In: *Journal of Computational Physics* 510 (2024), p. 113109. eprint: arXiv:2311.14670.
- [11] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization”. In: *CoRR* abs/1412.6980 (2014).
- [12] V. Devanathan. *Angular Momentum Techniques in Quantum Mechanics*. Fundamental Theories of Physics 108. ISBN:978-0-306-47123-0. Dordrecht: Kluwer Academic Publishers, 2002.