Hybrid Neural Representations for Spherical Data

Hyomin Kim [1] **Yunhui Jang** [1] **Jaeho Lee** [1] **Sungsoo Ahn** [1]

Abstract

In this paper, we study hybrid neural representations for spherical data, a domain of increasing relevance in scientific research. In particular, our work focuses on weather and climate data as well as cosmic microwave background (CMB) data. Although previous studies have delved into coordinate-based neural representations for spherical signals, they often fail to capture the intricate details of highly nonlinear signals. To address this limitation, we introduce a novel approach named Hybrid Neural Representations for Spherical data (HNeR-S). Our main idea is to use spherical feature-grids to obtain positional features which are combined with a multi-layer perceptron to predict the target signal. We consider feature-grids with equirectangular and hierarchical equal area isolatitude pixelization structures that align with weather data and CMB data, respectively. We extensively verify the effectiveness of our HNeR-S for regression, super-resolution, temporal interpolation, and compression tasks.

1. Introduction

Coordinate-based neural representations (CNRs) form a family of techniques to parameterize the target signal using a neural network that maps input coordinates to signal values (Xie et al., 2021). This results in a continuous function that seamlessly interpolates discretized coordinates of the target signal. As a case in point, when applied to image representation, CNRs trained to fit RGB values on 2D-pixel coordinates can also predict the continuum of coordinates between pixels. This capability has been proven to be useful in representing, interpolating, and increasing the signal resolutions of various modalities, including images (Chen et al., 2021), signed distance functions (Park et al., 2019), and radiance fields (Mildenhall et al., 2021) as special cases.

[1]Pohang University of Science and Technology. Correspondence to: Hyomin Kim <hyomin126@postech.ac.kr>, Sungsoo Ahn <sungsoo.ahn@postech.ac.kr>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024. Copyright 2024 by the author(s).

However, CNRs have been less investigated for spherical signals despite the important scientific applications, e.g., global weather and climate modeling (Hersbach et al., 2018) or the analysis of cosmic microwave background (CMB) radiation (de Bernardis et al., 2000). This is noteworthy given the necessity for interpolating and increasing the resolution of spherical signals in these applications. For example, spatial or temporal interpolation is required for building a fine-grained weather forecasting system (Lam et al., 2023; Bi et al., 2023). Increasing the resolution of CMB data may discover new insights about the early universe (Bennett et al., 1996; Spergel et al., 2003; Ade et al., 2014).

To this end, recent works (Esteves et al., 2022; Grattarola & Vandergheynst, 2022; Koestler et al., 2022; Schwarz et al., 2023; Huang & Hoefler, 2023) have investigated CNRs based on coordinate-based multi-layer perceptrons (MLPs). To featurize the spherical coordinates, they use the sinusoids from the Cartesian coordinate system (Schwarz et al., 2023; Huang & Hoefler, 2023), spectral embeddings of the discretized manifold (Grattarola & Vandergheynst, 2022; Koestler et al., 2022), or the spherical harmonics with fixed frequency components (Esteves et al., 2022). However, these methods suffer from the inherent limitation of coordinate-based MLPs: lack of expressive power to approximate highly nonlinear signals to their finest details.

For Euclidean data, recent works (Chabra et al., 2020; Müller et al., 2022; Martel et al., 2021; Liu et al., 2020; Peng et al., 2021) have resolved this issue by using hybrid neural representations that combine feature-grid structures with the MLP. Their key idea is to decompose the input space into grid structures and assign learnable parameters as positional features for each point in the grid. Then the positional features can largely fluctuate accordingly with the unconstrained point-wise parameters. However, extending the existing hybrid neural representation to spherical data is non-trivial since they rely on grid structures and interpolation algorithms specific to Euclidean domains, e.g., sparse voxel grids and Euclidean distance-based interpolation, that are highly redundant or ill-defined for spherical data.

Contribution. In this work, we develop a novel hybrid neural representation for spherical data (HNeR-S) that combines spherical feature-grids with MLP. Our main idea is to use a hierarchy of coarse- to fine-grained spherical feature-

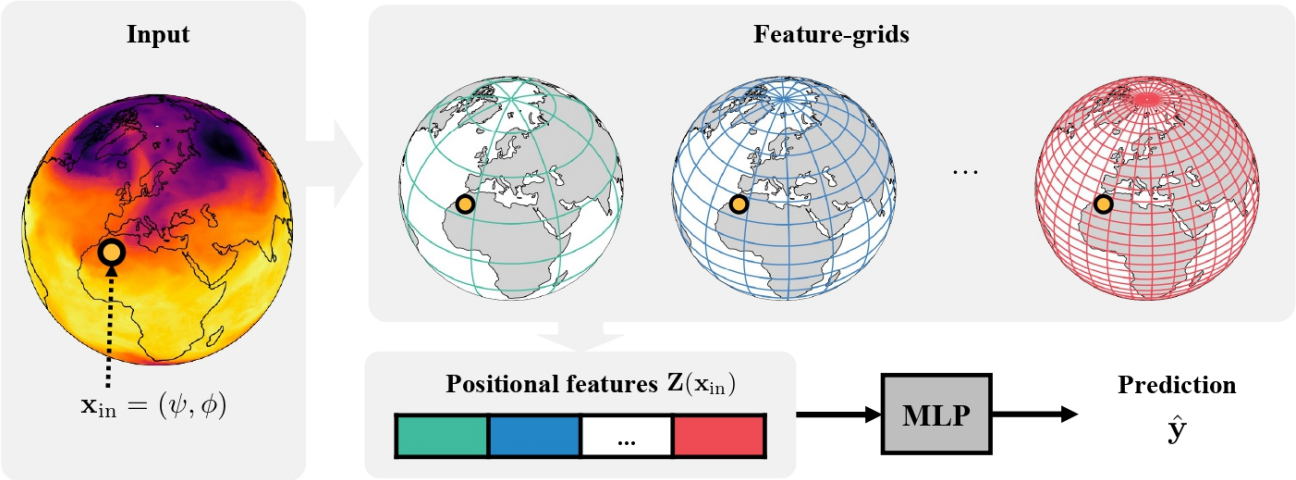

Figure 1. **Overview of hybrid neural representation for spherical data (HNeR-S).** HNeR-S considers an input point as the spherical coordinate $\mathbf{x}_{\text{in}} = (\psi, \phi)$, a pair of latitude $\psi \in [-\frac{1}{2}\pi, \frac{1}{2}\pi]$ and longitude $\phi \in [0, 2\pi)$. Then the model interpolates the neighborhood feature-grid parameters and constructs the positional features $\mathbf{Z}(\mathbf{x}_{\text{in}})$. The MLP predicts the target signal values from the positional features.

grids to generate positional features of a spherical coordinate at multiple frequencies. We focus on scientific data: weather and climate data as well as CMB data. We provide an overview of HNeR-S in Figure 1.

To compute the positional feature for an input point, we assign parameters to the grid point in the spherical feature-grid and apply a spherical interpolation algorithm to output the feature for the input position. In particular, we choose to align the feature-grid structure with the inherent data structure. This leads to designing our method with (1) an equirectangular feature-grid for the weather and climate data and (2) a hierarchical equal area isolatitude pixelization (HEALPix) feature-grid for the CMB data.

We thoroughly evaluate our algorithm on weather and climate (Hersbach et al., 2018) and CMB (Ade et al., 2014) data for super-resolution, regression, temporal interpolation, and compression tasks. Furthermore, we implement and compare with five CNR baselines (Mildenhall et al., 2021; Sitzmann et al., 2020; Esteves et al., 2022; Saragadam et al., 2023; Tancik et al., 2020) applicable to spherical data.

To summarize, our contributions are as follows:

- We propose a novel hybrid neural representation for spherical data (HNeR-S), which tailors the existing hybrid neural representations for spherical data using spherical feature-grids and interpolation algorithms.
- We empirically evaluate our algorithms on climate and weather data as well as CMB data. The results demonstrate our superiority over five INR baselines on super-resolution, regression, temporal interpolation, and compression tasks.

2. Related Work

Coordinate-based neural representations (CNRs). CNRs aim to learn a mapping from coordinates to the signal values via neural networks. The choice of positional encodings and activation functions are important for CNRs to capture high-frequency information of the underlying signal (Rahaman et al., 2019). To this end, sinusoidal positional encoding (Tancik et al., 2020) and activation function (Sitzmann et al., 2020) are proposed for Euclidean data, e.g., images and videos. Saragadam et al. (2023) proposed the Gabor wavelet activations to improve the performance and the robustness of sinusoidal encodings for the image domain. Tancik et al. (2020) analyzed the benefits of positional encoding to learn high-frequency functions via neural tangent kernel theory (Jacot et al., 2018).

CNRs for non-Euclidean data. Recently, researchers have also considered generalizing CNRs to non-Euclidean data (Esteves et al., 2022; Grattarola & Vandergheynst, 2022; Koestler et al., 2022; Schwarz et al., 2023; Huang & Hoefler, 2023). A simple approach is to lift the two-dimensional sphere into three-dimensional Euclidean space so that the Cartesian coordinates of the sphere can be used as inputs to the existing three-dimensional CNRs (Schwarz et al., 2023; Huang & Hoefler, 2023). To incorporate the geometry of the data, researchers considered using the eigenfunctions of the Laplace-Beltrami operator for general manifolds (Grattarola & Vandergheynst, 2022; Koestler et al., 2022). For spherical data, this corresponds to using spherical harmonics to featurize the spherical coordinates (Esteves et al., 2022).

Learning weather and climate data. Recently, weather and climate data have gained popularity as an application of machine learning (Yuval & O'Gorman, 2020), which

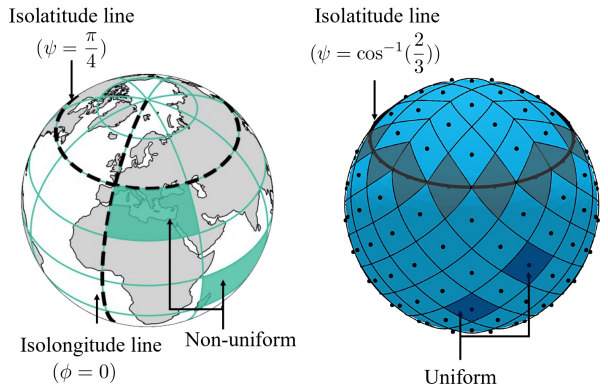

Figure 2. **Overview of equirectangular and HEALPix grid.** Note that the uniformity of feature-grid refers to the consistency in the area covered by each unit cell within the grid.

is attributed to its significant impact on climate prediction, mitigation, and adaptation. Employing neural networks for weather and climate data tasks has shown impressive results across a variety of tasks, e.g., super-resolution (Wang et al., 2021; Yang et al., 2022), temporal modeling (Wang et al., 2018; Stengel et al., 2020), and compression (Dupont et al., 2022; Huang & Hoefler, 2023). Neural network architectures like spherical convolutional neural networks (Cohen et al., 2018) and spherical Fourier neural operators (Bonev et al., 2023) have been tailored for this domain.

Learning cosmic microwave background (CMB) data. CMB data encapsulates the thermal radiation left over from the Big Bang, the origin of the universe. Machine learning has been applied to predict the posterior distribution of the cosmological parameters (Hortúa et al., 2020) as well as signal recovery and dust cleaning (Caldeira et al., 2019; Wang et al., 2022; Casas et al., 2022). A potential application of machine learning for CMB data is to increase the resolution of the data, which has mainly been progressed from the physical development of new sensors (Bennett et al., 1996; Spergel et al., 2003; Ade et al., 2014).

3. Method

In this section, we introduce our hybrid neural representations for spherical data (HNeR-S). We first elaborate on the general grid-agnostic framework (Section 3.1) and provide detailed descriptions for the equirectangular and HEALPix feature-grids (Sections 3.2 and 3.3). We provide a brief overview of each grid in Figure 2.

3.1. General Framework

We formulate our problem as learning a mapping from spherical coordinates to the target signal. To this end, our HNeR-S employs a hierarchy of spherical feature-grids and a bilinear spherical interpolation algorithm to output the positional feature of the input coordinate. Then HNeR-S applies a multi-layer perceptron (MLP) on the positional feature to predict the target signal. We provide an overview of HNeR-S in Figure 1.

Problem definition. To be specific, we consider learning some ground-truth spherical signal $f(\mathbf{x}_{\text{in}}) \in \mathbb{R}^{D_{\text{out}}}$, where the input coordinate $\mathbf{x}_{\text{in}} = (\psi, \phi)$ is a pair of latitude $\psi \in [-\frac{1}{2}\pi, \frac{1}{2}\pi]$ and longitude $\phi \in [0, 2\pi)$. To approximate this signal, we train a neural network f_θ on a dataset of tuples $(\mathbf{x}_{\text{in}}, \mathbf{y})$ with ground truth label $\mathbf{y} = f(\mathbf{x}_{\text{in}})$.

Overall architecture. Our HNeR-S makes a prediction $\hat{\mathbf{y}} = f_\theta(\mathbf{x}_{\text{in}})$ from an input point $\mathbf{x}_{\text{in}}$ as follows:

$$f_\theta(\mathbf{x}_{\text{in}}) = \text{MLP}(\mathbf{Z}(\mathbf{x}_{\text{in}})),$$

where $\text{MLP}(\cdot)$ denotes the MLP and $\mathbf{Z}(\mathbf{x}_{\text{in}}) \in \mathbb{R}^{D \times L}$ denotes the concatenation of D-dimensional positional features constructed from L levels of feature-grids:

$$\mathbf{Z}(\mathbf{x}_{\text{in}}) = \text{Concat}(\mathbf{z}^{(1)}(\mathbf{x}_{\text{in}}), \ldots, \mathbf{z}^{(L)}(\mathbf{x}_{\text{in}})),$$

where $\text{Concat}(\cdot)$ denotes concatenation of elements and the positional feature $\mathbf{z}^{(\ell)}(\mathbf{x}_{\text{in}}) \in \mathbb{R}^D$ is constructed from the ℓ-th level feature-grid for $\ell = 1, \ldots, L$. We assign a higher level ℓ for a finer grid.

Positional feature. The ℓ-th positional feature $\mathbf{z}^{(\ell)}(\mathbf{x}_{\text{in}})$ is constructed by interpolating the positional features assigned to grid points $\mathcal{V}^{(\ell)}$ in the ℓ-th feature-grid. We let $\mathbf{x}_i$ denote the coordinate of a grid point $i \in \mathcal{V}^{(\ell)}$. Different grid types, e.g., equirectangular and HEALPix grids, correspond to a different organization of position $\mathbf{x}_i^{(\ell)}$ for each grid point i. We also let $\mathbf{z}_i^{(\ell)} \in \mathbb{R}^D$ denote the associated positional feature, which is also a parameter to train for our framework.

Then given an input point with position $\mathbf{x}_{\text{in}}$, the HNeR-S constructs the ℓ-th positional feature $\mathbf{z}^{(\ell)}(\mathbf{x}_{\text{in}})$ by interpolation of neighboring points $\mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}}) \subseteq \mathcal{V}^{(\ell)}$ in the feature-grid:

$$\mathbf{z}^{(\ell)}(\mathbf{x}_{\text{in}}) = \text{Interpolate}(\mathbf{x}_{\text{in}}, \{(\mathbf{x}_i^{(\ell)}, \mathbf{z}_i^{(\ell)}) : i \in \mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}})\}).$$

The neighborhood structure of the equirectangular and HEALPix grid allows our HNeR-S to use bilinear spherical interpolation based on the latitude and the longitude.

Neighborhood structure. We consider neighborhood functions that output $\mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}}) \subseteq \mathcal{V}^{(\ell)}$ as a set of four points, denoted by $\mathbf{x}_{i,j} = (\psi_{i,j}, \phi_{i,j})$ associated with positional feature $\mathbf{z}_{i,j}^{(\ell)}$ for $i, j \in \{1, 2\}$. Importantly, we consider grids that allow the construction of the neighborhood as pairs of isolatitude points and let $\psi_1 = \psi_{1,1} = \psi_{1,2}, \psi_2 = \psi_{2,1} = \psi_{2,2}$ without loss of generality.

Furthermore, the neighbors bound the input point by a latitude interval $\Psi = [\psi_{1,1}, \psi_{1,2}]$ and two longitude intervals Φ_1, Φ_2. When a neighbor of the input point does not

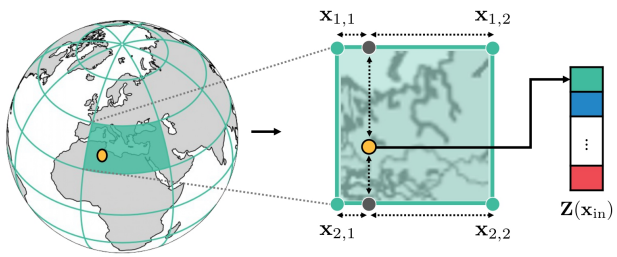

Figure 3. **Neighborhood structure of the equirectangular grid.** The yellow point indicates the input point $\mathbf{x}_{\text{in}}$ and green points indicate the neighborhood grid points $\mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}})$.

overlap with the prime meridian ($\phi = 0$), the longitude intervals are defined similarly, e.g., $\Phi_1 = [\phi_{1,1}, \phi_{1,2}]$. Otherwise, the longitude intervals are set by $\Phi_1 = [\phi_{1,1}, 2\pi] \cup [0, \phi_{1,2}]$ and Φ_2 is set similarly.

Bilinear spherical interpolation. Given the neighbor as pairs of isolatitude points, our framework employs bilinear interpolation using the longitude and the latitude to compute the associated feature. To be specific, given an input point $\mathbf{x}_{\text{in}} = (\psi, \phi)$, we compute $\mathbf{z}^{(\ell)}(\psi, \phi)$ via two steps: (1) computing $\mathbf{z}^{(\ell)}(\psi_1, \phi)$ and $\mathbf{z}^{(\ell)}(\psi_2, \phi)$ via isolatitude interpolation and (2) computing $\mathbf{z}^{(\ell)}(\psi, \phi)$ via interpolation along the longitude between features computed in (1).

Precisely, the interpolation along the isolatitude neighbors $\mathbf{x}_{1,1} = (\psi_1, \phi_{1,1}), \mathbf{x}_{1,2} = (\psi_1, \phi_{1,2})$ is defined as follows:

$$\mathbf{z}^{(\ell)}(\psi_1, \phi) = \lambda \mathbf{z}^{(\ell)}_{1,1} + (1 - \lambda)\mathbf{z}^{(\ell)}_{1,2},$$

where $\mathbf{z}^{(\ell)}_{1,1} = \mathbf{z}^{(\ell)}(\psi_1, \phi_{1,1})$ and $\mathbf{z}^{(\ell)}_{1,2} = \mathbf{z}^{(\ell)}(\psi_1, \phi_{1,2})$ are learnable parameters. For the interpolation weight λ, the algorithm sets $\lambda = \frac{d(\phi, \phi_{1,1})}{d(\phi_{1,1}, \phi_{1,2})}$ where $d(\cdot, \cdot)$ is the difference in the longitude for isolatitude points, e.g., $d(\phi_{1,1}, \phi_{1,2}) = \phi_{1,2} - \phi_{1,1}$ if the interval Φ_1 does not overlap with the prime meridian. One can also compute $\mathbf{z}^{(\ell)}(\psi_2, \phi)$ similarly using $\mathbf{z}^{(\ell)}_{2,1}$ and $\mathbf{z}^{(\ell)}_{2,2}$.

Finally, the algorithm computes the final positional feature by interpolation along the latitude:

$$\mathbf{z}^{(\ell)}(\psi, \phi) = \mu \mathbf{z}^{(\ell)}(\psi_1, \phi) + (1 - \mu)\mathbf{z}^{(\ell)}(\psi_2, \phi),$$

where the interpolation weight is set to $\mu = \frac{\psi - \psi_1}{\psi_2 - \psi_1}$.

3.2. Equirectangular Feature-Grid

We use HNeR-S based on the equirectangular feature-grid for weather and climate data since the data is also distributed along an equirectangular grid. It is based on a projection that maps meridians (isolongitude lines) to vertical straight lines of constant spacing and circles of latitude (isolatitude lines) to horizontal straight lines of constant spacing. We index each point in the ℓ-th grid by a tuple of integers (n, m)

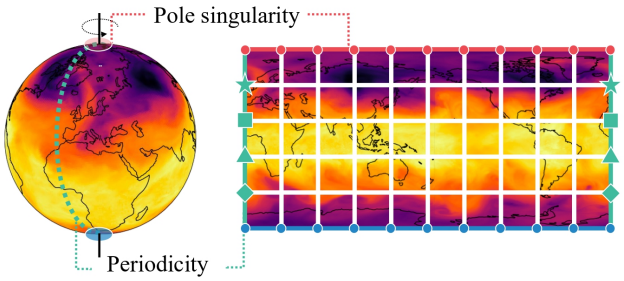

Figure 4. **Pole singularity and periodicity of equirectangular grid.** Different spherical coordinates can indicate the same point on a sphere, i.e., points at the North Pole ($\psi = -\frac{\pi}{2}$), South Pole ($\psi = \frac{\pi}{2}$), and the prime meridian ($\phi = 0$ and $\phi = 2\pi$). Our HNeR-S avoided assigning different parameters for such spherical coordinates. Points with the same marker and color share the associated parameters.

with the position $\mathbf{x}^{(\ell)}_{n,m}$ described as follows:

$$\mathbf{x}^{(\ell)}_{n,m} = (\psi^{(\ell)}_n, \phi^{(\ell)}_m) = \left(\pi\left(\frac{n}{N^{(\ell)}_{\text{lat}}} - \frac{1}{2}\right), 2\pi\frac{m}{N^{(\ell)}_{\text{lon}}}\right).$$

The non-negative integers $n \leq N^{(\ell)}_{\text{lat}}$ and $m \leq N^{(\ell)}_{\text{lon}}$ index the point-wise latitude and longitude, respectively. The integers $N^{(\ell)}_{\text{lat}}, N^{(\ell)}_{\text{lon}}$ decide the latitude-wise and longitude-wise resolutions of the feature-grid. We choose them by:

$$N^{(\ell)}_{\text{lat}} = \lfloor \gamma^{\ell-1} \cdot N^{(1)}_{\text{lat}} \rfloor, \quad N^{(\ell)}_{\text{lon}} = \lfloor \gamma^{\ell-1} \cdot N^{(1)}_{\text{lon}} \rfloor,$$

where $\gamma > 1$ is the scaling factor, $\lfloor \cdot \rfloor$ is the floor function, and $N^{(1)}_{\text{lat}}, N^{(1)}_{\text{lon}}$ are resolutions of the coarsest grid ($\ell = 1$).

Neighborhood structure. Given an input point $\mathbf{x}_{\text{in}} = (\psi, \phi)$ as an input of the CNR, the neighborhood forms a unit cell of the feature-grid as follows:

$$\mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}}) = \{(n, m), (n, m + 1), (n + 1, m), (n + 1, m + 1)\}.$$

Our framework determines the index (n, m) to ensure the position $\mathbf{x}_{\text{in}}$ is inside the rectangle:

$$(n, m) = \left(\left\lfloor \left(\frac{\psi}{\pi} + \frac{1}{2}\right) N^{(\ell)}_{\text{lat}} \right\rfloor, \left\lfloor \frac{\phi}{2\pi} N^{(\ell)}_{\text{lon}} \right\rfloor\right).$$

We describe the neighborhood structure in Figure 3.

Non-uniformity. Since the equirectangular projection does not preserve area, the resulting equirectangular grid is irregular and the cell-wise areas are non-uniformly distributed. To incorporate this during training and evaluation, one needs to reweight each data point by the corresponding cell area (Huang & Hoefler, 2023). The non-uniformity of the equirectangular grid is illustrated in Figure 2.

Pole singularity. It is important to note that the equirectangular grid assigns multiple points on a single pole, i.e.,

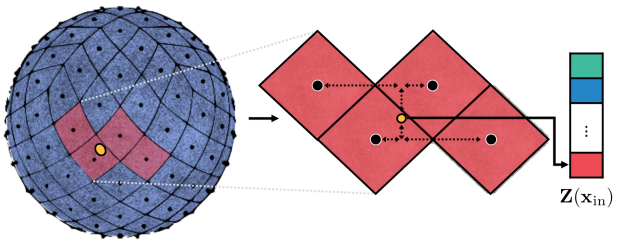

Figure 5. **Neighborhood structure of the HEALPix grid.** The yellow point indicates the input point $\mathbf{x}$ and black points indicate the neighborhood grid points $\mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}})$.

$\mathbf{x}_{0,m} = \mathbf{x}_{0,m'}$ for any $m, m' < N_{\text{lon}}$. To incorporate this fact, we merge the associated parameters into a single parameter for both training and evaluation. We also remark on how the interpolation of our algorithm is well-defined even for the input points near the poles. We illustrate the pole singularity in Figure 4.

Periodicity. The equirectangular projection of maps inherently introduces discontinuity, where the grid points sharing identical coordinates find themselves at the opposite ends of the map, on the left and right sides of the prime meridian, as described in Figure 4. To address this issue, we merge the parameters corresponding to grid points at the same latitude from both ends of the map into a unified parameter.

3.3. HEALPix Grid

We consider feature-grids with hierarchical equal area isolatitude pixelation (HEALPix) structure for the CMB data (Gorski et al., 2005), since the data points are distributed according to the HEALPix grid structure. Notably, the data points are located in the center of the HEALPix grid cells. The grid structure uniformly partitions the sphere and the data points are uniformly distributed over the sphere.

In particular, the HEALPix center points are located on N_{line} lines of constant latitude, and each line (or azimuth) is uniformly divided by the points. The lines with latitude ψ satisfying $|\cos\psi| \leq \frac{2}{3}$ are divided into the same number of pixels N_{eq}. The remaining lines are located within the polar cap ($|\cos\psi| > \frac{2}{3}$) and contain a varying number of pixels, increasing from line to line with increasing distance from the poles by one pixel within each quadrant. We elaborate on HEALPix structure at Appendix A.

The resolution of the grid is expressed by the parameter N_{side}, which denotes the number of divisions along the side of a base resolution pixel that is needed to reach a desired high-resolution partition. One can observe that $N_{\text{line}} = 4N_{\text{side}} - 1$ and $N_{\text{eq}} = 4N_{\text{side}}$. This results in $N_{\text{pix}}^{(\ell)} = 12(N_{\text{side}}^{(\ell)})^2$ pixels and the corresponding center points. The resolution increases by the power of two for each level, i.e., we set $N_{\text{side}}^{(\ell)} = 2^{\ell-1}$ for the ℓ-th grid.

Neighborhood structure. We choose the neighborhood $\mathcal{N}^{(\ell)}(\mathbf{x}_{\text{in}})$ in the HEALPix grid as the center of pixels that are adjacent to the pixel that includes the point $\mathbf{x}_{\text{in}}$. Furthermore, while a pixel in the HEALPix grid is adjacent to at most eight pixels, we choose the neighborhood as the center of the closest pixels on the two lines above and below the location. Such neighborhood allows the application of bilinear interpolation as described in Section 3.1. We illustrate the HEALPix neighborhood structure in Figure 5.

4. Experiments

In this section, we assess the performance of our two hierarchical hybrid neural representations designed for spherical data: the equirectangular grid and the HEALPix grid. To this end, we first validate our framework for climate data regarding four tasks: regression, super-resolution, temporal interpolation, and compression (Section 4.1). We also explore two tasks on CMB data with the HEALPix grid regarding super-resolution and regression (Section 4.2). We will release our code upon acceptance.

Datasets. Mainly, we utilized climate and CMB data for our experiment. Weather and climate data is gathered from ECMWF reanalysis 5th generation (Hersbach et al., 2018, ERA5) archive where data can be directly downloaded by using Climate Data Storage API (Buontempo et al., 2020). To be specific, we gathered geopotential and temperature data in the year 2000 for super-resolution, regression, and temporal interpolation. For the compression task, we gathered geopotential data in the year 2016 following Huang & Hoefler (2023). The detailed description of datasets is in Table 2.

Next, we use the CMB temperature data from Planck Public Data Release 1 (PR1) Mission Science Maps data at NASA/IPAC Infrared Science Archive (IRSA)[1]. We use the version where missing values are determined by spectral matching independent component analysis (Delabrouille et al., 2003, SMICA).

Baselines. On one hand, for super-resolution, regression, and temporal interpolation tasks, we compare our methods to five different encoding schemes that are applicable for learning spherical signals: logarithmically spaced sinusoidal positional encoding (Mildenhall et al., 2021, ReLU+P.E.), sinusoidal representation networks (Sitzmann et al., 2020, SIREN), wavelet implicit neural representation (Saragadam et al., 2023, WIRE), spherical harmonics implicit neural representation (Esteves et al., 2022, SHINR), Fourier feature networks (Tancik et al., 2020, FFN), and HashGrid (Müller et al., 2022, HashGrid).

[1] https://irsa.ipac.caltech.edu/data/Planck/release_1/all-sky-maps/

Table 1. **Results of the regression and super-resolution.** For the weather dataset including geopotential (Geo.) and temperature (Temp.) data, evaluation is based on weighted PSNR, while for the CMB dataset, it is assessed using standard PSNR. The best results are highlighted in **bold**. Regression and $\times 2$, $\times 4$ super-resolution tasks are included. Additionally, weighted RMSE results are provided at Appendix B.

		ReLU + P.E.	SIREN	SHINR	WIRE	FFN	HashGrid	**Ours**
Geo.	Reg.	59.79 ± 1.87	58.27 ± 1.21	54.26 ± 0.88	65.24 ± 1.59	58.96 ± 3.81	46.54 ± 0.39	**66.83 ± 4.93**
	$\times 2$	59.09 ± 1.54	55.04 ± 2.05	53.21 ± 1.10	54.33 ± 1.04	58.60 ± 0.40	36.30 ± 0.30	**67.51 ± 1.81**
	$\times 4$	55.15 ± 1.62	51.84 ± 1.81	52.35 ± 2.20	51.80 ± 1.05	53.00 ± 2.38	28.03 ± 0.17	**59.72 ± 2.37**
Temp.	Reg.	54.56 ± 0.50	52.65 ± 0.80	43.67 ± 2.80	56.18 ± 1.02	54.96 ± 0.86	43.95 ± 1.07	**60.87 ± 0.77**
	$\times 2$	51.19 ± 1.77	51.87 ± 1.67	42.54 ± 0.66	51.12 ± 0.53	55.29 ± 2.35	28.03 ± 0.17	**57.24 ± 0.96**
	$\times 4$	49.45 ± 0.05	46.76 ± 1.87	42.28 ± 0.87	47.05 ± 1.37	49.76 ± 0.76	25.77 ± 0.21	**51.31 ± 0.15**
CMB	Reg.	23.60 ± 0.05	28.07 ± 0.20	22.59 ± 0.01	29.91 ± 0.07	27.98 ± 0.04	25.69 ± 0.08	**38.08± 0.09**
	$\times 2$	22.21 ± 0.01	23.32 ± 0.07	22.24 ± 0.01	26.24 ± 0.37	24.22±0.17	23.00 ± 0.38	**28.84 ± 1.40**
	$\times 4$	22.19 ± 0.01	23.27 ± 0.06	22.24 ± 0.00	25.65 ± 0.17	23.91 ± 0.16	21.77 ± 0.01	**26.66 ± 0.36**

Table 2. **Details on the datasets used for weather and climate (ERA5) and CMB (PR1) data.**

	ERA5	PR1
Grid	Equirectangular	HEALPix
Target	Geopotential, Temperature	CMB Temperature
Spatial Res.	0.25°, 0.50°, 1.00°	5 arcmin
Temporal Res.	Daily, Weekly	-

On the other hand, for the compression task, we compare our methods to the state-of-the-art compression scheme without neural network (Liang et al., 2022) and the recently proposed coordinate-based neural network for compression (Huang & Hoefler, 2023).

Experimental setting. We conduct all experiments using a single RTX 3090 GPU. For the regression, super-resolution, and temporal interpolation experiments, we employ a consistent architecture, consisting of a 4-layer multi-layer perceptron (MLP) with 256 units in each hidden layer. Detailed settings for the compression experiment are provided in Section 4.1.

4.1. Weather and Climate Data

In this section, we evaluate our model on weather and climate data for $\times 2$ and $\times 4$ super-resolution, regression, temporal interpolation, and compression. The climate data is often represented using an equirectangular grid (i.e., a regular latitude-longitude grid), where data distribution is concentrated around the poles. To address this unique challenge, we employ our hierarchical equirectangular feature-grid. In addition, every task on an equirectangular grid is evaluated by latitude-weighted metric, i.e., weighted RMSE, MAE, and PSNR (Huang & Hoefler, 2023). This assigns smaller weights to the poles, compensating for the irregularity of data sampled from the equirectangular grid.

Regression. We first evaluate our framework for the regression task where we split the train, valid, and test dataset in the portion of 6:2:2. We report the results in Table 1. One can observe how our method consistently demonstrates superiority in PSNR performance and faster convergence over the baselines. We further visualize our results in Figure 7. The figure illustrates our method's accuracy in regression tasks, showcasing minimal errors across diverse geographical regions.

In Figure 7, one can observe how our algorithm demonstrates significant improvement over the baselines. The improvement is more pronounced for predicting the temperature, where our HNeR-S achieves a 60.87 PSNR score whereas the second-best baseline (WIRE) achieves 56.18.

Super-resolution. Next, we evaluate our method for the super-resolution task. To this end, we split the low-resolution data points into train and validation sets with a ratio of 8:2 and evaluate them on high-resolution grids. Precisely, we train the models on the resolution of 0.50° and 1.00° and evaluate them on the resolution of 0.25° for the $\times 2$ and $\times 4$ setting, respectively.

We report the results in Table 1. Here, one can observe how our algorithm demonstrates significant improvements over the baselines. We also visualize the results in Figure 6, which demonstrate our method's ability to reconstruct high-resolution details accurately, improving clarity and definition in climate data visualizations.

Temporal interpolation. Finally, we evaluate the ability of our framework to interpolate between snapshots of the weather at different timesteps. This task employs two temporal resolutions: daily and weekly, with a total of 30 data snapshots sampled. Since this task requires the CNRs to

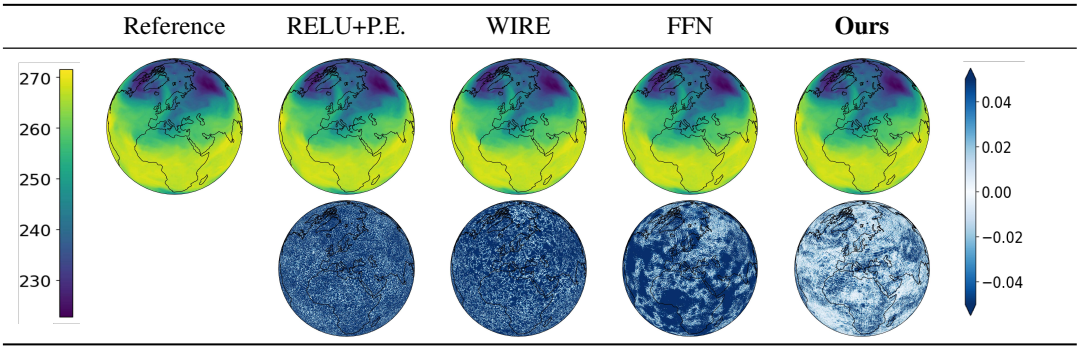

Figure 6. **Prediction (top) and error maps (bottom) for $\times 2$ super-resolution in climate and weather temperature data.**

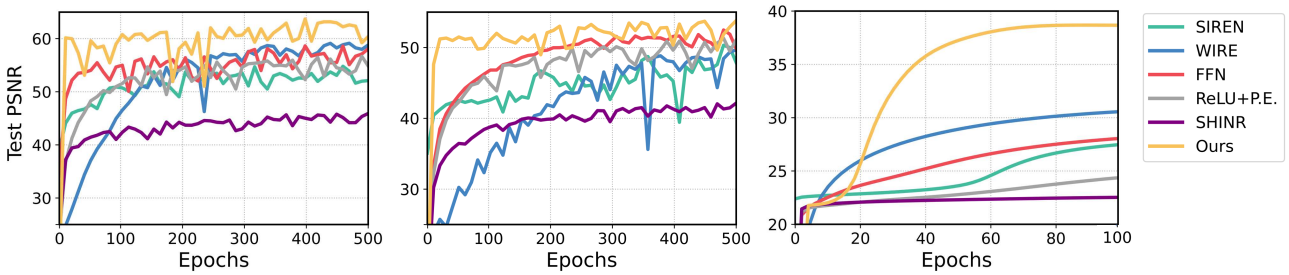

Figure 7. **Test PSNR curves.** These plots show the test PSNR curve for ($\times 4$) super-resolution (left) and regression (middle) for the temperature of weather and climate data and regression for CMB data (right). Our model converges to the best PSNR in the fastest manner.

Table 3. **Performance for temporal interpolation on weather and climate data.** The evaluation metric is weighted PSNR and the best results are highlighted in **bold**.

	Geopotential		Temperature	
	Daily	Weekly	Daily	Weekly
ReLU+P.E.	28.43	23.92	27.12	24.35
SIREN	**28.80**	23.16	27.26	23.34
SHINR	28.46	22.75	27.27	22.90
WIRE	28.66	22.99	27.07	23.18
FFN	28.53	23.54	26.65	23.73
HNeR-S (ours)	28.79	**23.97**	**27.55**	**24.50**

Table 4. **Results of compression on weather and climate data.** The best results of spatial compression models are highlighted in **bold**. †FFN-T and SZ3 results are from Huang & Hoefler (2023).

	$0.031 \leq$ BPP ≤ 0.080			$0.099 \leq$ BPP ≤ 0.119		
	BPP (↓)	WRMSE (↓)	WMAE (↓)	BPP (↓)	WRMSE (↓)	WMAE (↓)
Spatiotemporal Compression						
SZ3†	0.080	698.5	574.9	0.115	493.9	408.8
FFN-T†	0.031	143.1	101.2	0.111	73.6	53.0
Spatial Compression						
SIREN	0.063	669.5	462.1	0.119	616.0	404.8
WIRE	0.069	2617.2	1433.3	0.099	688.7	478.6
HNeR-S (ours)	0.057	**144.8**	**99.4**	0.114	**124.0**	**50.3**

condition on the timestep of the snapshot, we modify our algorithm and the baselines by incorporating the timestep t as a concatenation of the positional feature, which serves as input for the MLP.

The outcomes, detailed in Table 3 and Figure 7, highlight our method's exceptional PSNR performance. The visualizations in Figure 6 reveal our model's capability to accurately interpolate temporal changes.

Compression. As an additional experiment we apply HNeR-S for compression, driven by the demand to store daily-growing high-resolution climate and weather datasets (Huang & Hoefler, 2023). To this end, we combine our method with the existing HNeR-based compression scheme (Ladune et al., 2023) to apply image-wise compression, i.e., we train one neural network for the compression of each image. As a non-machine learning baseline, we consider a modular composable SZ3 (Liang et al., 2022) framework. We also compare with the Fourier feature network with temporal encoding (Huang & Hoefler, 2023, FFN-T) proposed for compression of climate and weather data. We note that FFN-T is a spatiotemporal compressor, i.e., it takes both spatial and temporal coordinate that trains a series of snapshots,

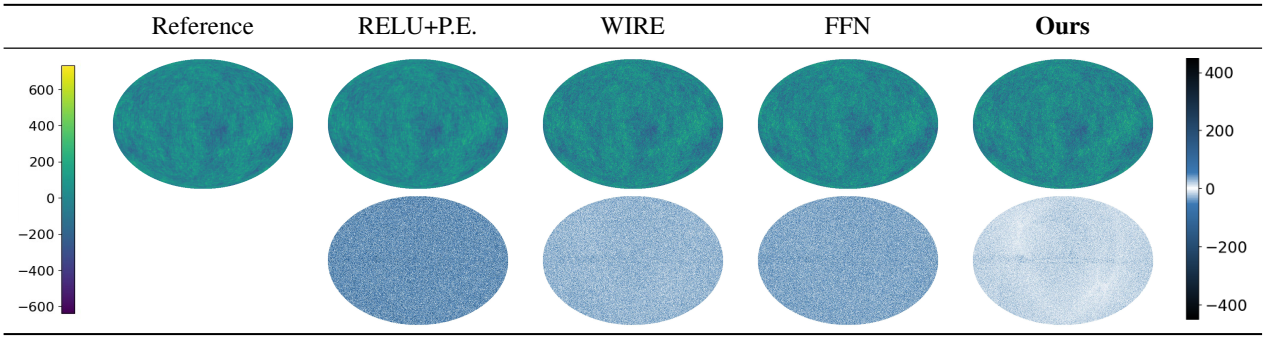

Figure 8. **Prediction (top) and error maps (bottom) for regression in CMB data.**

while ours is a spatial compressor, i.e., it trains a network for each image. Additionally, we consider two implicit neural network baselines including WIRE (Saragadam et al., 2023) and SIREN (Sitzmann et al., 2020) for fair comparison on spatial compression.

We report the results in Table 4. We compare the trade-off between bits per pixel (BPP) and the quality of the compressed data measured in weighted RMSE and weighted MAE.[2] Here, one can observe that our method shows superior results to spatial compression baselines. Our method achieves comparable results with FFN-T, which is a spatial-temporal compression algorithm that is not directly comparable to ours. In particular, FFN-T takes advantage of sharing parameters across different time steps, while our scheme can easily incorporate the daily growth of the dataset.[3]

4.2. Cosmic Microwave Background Data

Here, we evaluate our model for the CMB data, which is stored as a HEALPix grid where the distribution of the data is uniform across the grids. For every CMB experiment, we used a resolution of $N_{\text{side}} = 2048$, which is the highest resolution that the Planck satellite provides. One should note that we used the highest resolution of the CMB data (i.e., N_{side} = 2048) where the data is often downgraded or only utilizes part of the map due to its extensive data size (Krachmalnicoff & Tomasi, 2019; Montefalcone et al., 2021). Meanwhile, our method utilizes the full resolution and highlights the scalability and ability to learn intricate signals of the CMB data.

Super-resolution. Similar to the regression task, we can apply super-resolution for CMB data. However, the experiment setting is slightly adjusted for CMB data since the resolution of HEALPix data is defined by N_{side}. Here, the $\times 2$ setting involves training on $N_{\text{side}} = 1024$ to predict $N_{\text{side}} = 2048$, and the $\times 4$ setting trains the model on $N_{\text{side}} = 512$ to predict $N_{\text{side}} = 2048$.

Note that training on high-resolution $N_{\text{side}} = 2048$ data is challenging due to its large size, a hurdle our CNR-based approach overcomes, unlike prior works that often use downgraded data such as $N_{\text{side}} = 64$ (Defferrard et al., 2020). As can be seen in the Table 1 our approach outperforms the other baselines. Despite the highly intricate and noisy pattern of the CMB data, our model significantly outperforms the baselines. This implies that our model can learn complex real-world data as well.

We provide experimental results in Table 1 and Figure 8. We observe that our method consistently shows superior PSNR results. Moreover, in Figure 8, our model converges into desirable PSNR in the early stages of the training.

Regression. Similar to the regression in climate data, we predict the unknown signal from the constructed hybrid neural representation. We demonstrate the results in Table 1. One can observe that our method highly outperforms compared to the baselines. As depicted in Figure 8, both the prediction and error map indicate that our method can predict the target signal within a small error. Moreover, according to Figure 7, our model converges to the promising accuracy by showing a large performance gap compared to the other baselines.

5. Conclusion

In this paper, we presented two new hybrid neural representations tailored for spherical data, utilizing hierarchical spherical feature-grids to generate positional encodings through spherical interpolation algorithms. Our extensive experimental evaluations across multiple datasets and tasks have verified the effectiveness of our method. An interesting avenue for future work will be the application of adaptive grid (Martel et al., 2021; Martinez-Castellanos et al., 2022) which will lead to a more compact feature-grid. Also, there remains an exploration for interpolation methods tailored for spherical grids, i.e., geodesic-aware interpolation.

[2]Weights are assigned based on varying cell areas as in training.

[3]Our scheme can train a small network only for the new data, while FFN-T requires re-training on the updated dataset.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. RS-2019-II191906, Artificial Intelligence Graduate School Program(POSTECH)), the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. 2022R1C1C1013366), Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education(2022R1A6A1A0305295413), and the National Research Foundation of Korea(NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00213710, Neural network optimization with minimal optimization costs).

We would like to express our sincere gratitude to Langwen Huang for providing the detailed results of their work, which served as a foundational baseline method for our research. Additionally, our thanks extend to Jihoon Tack and Doyup Lee for their constructive feedback that greatly contributed to the improvement of this paper.

Impact Statement

Our framework can advance machine learning for applications like climatology and cosmology. Machine learning for weather and climate forecasting can benefit our society via improved resilience to extreme weather events, agriculture productivity, and water resource management. Furthermore, machine learning for cosmology can accelerate scientific discovery, in particular our understanding of dark matter and dark energy.

References

Ade, P. A., Aghanim, N., Alves, M., Armitage-Caplan, C., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Aussel, H., Baccigalupi, C., et al. Planck 2013 results. I. Overview of products and scientific results. *Astronomy & Astrophysics*, 571:A1, 2014.

Bennett, C. L., Banday, A. J., Górski, K. M., Hinshaw, G., Jackson, P., Keegstra, P., Kogut, A., Smoot, G. F., Wilkinson, D. T., and Wright, E. L. Four-year COBE DMR cosmic microwave background observations: maps and basic results. *The Astrophysical Journal*, 464(1):L1, 1996.

Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X., and Tian, Q. Accurate medium-range global weather forecasting with 3D neural networks. *Nature*, 619(7970):533–538, 2023. doi: 10.1038/s41586-023-06185-3. URL https://doi.org/10.1038/s41586-023-06185-3.

Bonev, B., Kurth, T., Hundt, C., Pathak, J., Baust, M., Kashinath, K., and Anandkumar, A. Spherical Fourier neural operators: Learning stable dynamics on the sphere. In *International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=TwsJ9IOZDx.

Buontempo, C., Hutjes, R., Beavis, P., Berckmans, J., Cagnazzo, C., Vamborg, F., Thépaut, J.-N., Bergeron, C., Almond, S., Amici, A., et al. Fostering the development of climate services through Copernicus Climate Change Service (C3S) for agriculture applications. *Weather and Climate Extremes*, 27:100226, 2020.

Caldeira, J., Wu, W. K., Nord, B., Avestruz, C., Trivedi, S., and Story, K. T. DeepCMB: Lensing reconstruction of the cosmic microwave background with deep neural networks. *Astronomy and Computing*, 28:100307, 2019.

Casas, J., Bonavera, L., González-Nuevo, J., Baccigalupi, C., Cueli, M., Crespo, D., Goitia, E., Santos, J., Sánchez, M., and de Cos, F. CENN: A fully convolutional neural network for CMB recovery in realistic microwave sky simulations. *Astronomy & Astrophysics*, 666:A89, 2022.

Chabra, R., Lenssen, J. E., Ilg, E., Schmidt, T., Straub, J., Lovegrove, S., and Newcombe, R. Deep local shapes: Learning local SDF priors for detailed 3D reconstruction. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pp. 608–625. Springer, 2020.

Chen, Y., Liu, S., and Wang, X. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8628–8638, 2021.

Cohen, T. S., Geiger, M., Köhler, J., and Welling, M. Spherical CNNs. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=Hkbd5xZRb.

de Bernardis, P., Ade, P. A., Bock, J. J., Bond, J., Borrill, J., Boscaleri, A., Coble, K., Crill, B., De Gasperis, G., Farese, P., et al. A flat universe from high-resolution maps of the cosmic microwave background radiation. *Nature*, 404(6781):955–959, 2000.

Defferrard, M., Milani, M., Gusset, F., and Perraudin, N. DeepSphere: A graph-based spherical CNN. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=B1e3OlStPB.

Delabrouille, J., Cardoso, J.-F., and Patanchon, G. Multidetector multicomponent spectral matching and applications for cosmic microwave background data analysis. *Monthly Notices of the Royal Astronomical Society*, 346 (4):1089–1102, 2003.

Dupont, E., Loya, H., Alizadeh, M., Golinski, A., Teh, Y. W., and Doucet, A. COIN++: Neural compression across modalities. *Transactions on Machine Learning Research*, 2022. ISSN 2835-8856. URL https://openreview.net/forum?id=NXB0rEM2Tq.

Esteves, C., Lu, T., Suhail, M., fan Chen, Y., and Makadia, A. Generalized Fourier features for coordinate-based learning of functions on manifolds, 2022. URL https://openreview.net/forum?id=g6UqpVislvH.

Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, 2005.

Grattarola, D. and Vandergheynst, P. Generalised implicit neural representations. *Advances in Neural Information Processing Systems*, 35:30446–30458, 2022.

Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., et al. ERA5 hourly data on single levels from 1979 to present. *Copernicus climate change service (c3s) climate data store (cds)*, 10(10.24381), 2018.

Hortúa, H. J., Volpi, R., Marinelli, D., and Malagò, L. Parameter estimation for the cosmic microwave background with Bayesian neural networks. *Physical Review D*, 102 (10):103509, 2020.

Huang, L. and Hoefler, T. Compressing multidimensional weather and climate data into neural networks. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=Y5SEe3dfniJ.

Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

Kim, H., Bauer, M., Theis, L., Schwarz, J. R., and Dupont, E. C3: High-performance and low-complexity neural compression from a single image or video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

Kingma, D. and Ba, J. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, San Diega, CA, USA, 2015.

Koestler, L., Grittner, D., Moeller, M., Cremers, D., and Lähner, Z. Intrinsic neural fields: Learning functions on manifolds. In *European Conference on Computer Vision*, pp. 622–639. Springer, 2022.

Krachmalnicoff, N. and Tomasi, M. Convolutional neural networks on the HEALPix sphere: A pixel-based algorithm and its application to CMB data analysis. *Astronomy & Astrophysics*, 628:A129, 2019.

Ladune, T., Philippe, P., Henry, F., Clare, G., and Leguay, T. COOL-CHIC: Coordinate-based low complexity hierarchical image codec. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13515–13522, 2023.

Lam, R., Sanchez-Gonzalez, A., Willson, M., Wirnsberger, P., Fortunato, M., Alet, F., Ravuri, S., Ewalds, T., Eaton-Rosen, Z., Hu, W., Merose, A., Hoyer, S., Holland, G., Vinyals, O., Stott, J., Pritzel, A., Mohamed, S., and Battaglia, P. Learning skillful medium-range global weather forecasting. *Science*, 2023.

Liang, X., Zhao, K., Di, S., Li, S., Underwood, R., Gok, A. M., Tian, J., Deng, J., Calhoun, J. C., Tao, D., et al. SZ3: A modular framework for composing prediction-based error-bounded lossy compressors. *IEEE Transactions on Big Data*, 9(2):485–498, 2022.

Liu, L., Gu, J., Zaw Lin, K., Chua, T.-S., and Theobalt, C. Neural sparse voxel fields. *Advances in Neural Information Processing Systems*, 33:15651–15663, 2020.

Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.

Martel, J. N., Lindell, D. B., Lin, C. Z., Chan, E. R., Monteiro, M., and Wetzstein, G. ACORN: Adaptive coordinate networks for neural representation. *ACM Trans. Graph. (SIGGRAPH)*, 2021.

Martinez-Castellanos, I., Singer, L. P., Burns, E., Tak, D., Joens, A., Racusin, J. L., and Perkins, J. S. Multiresolution HEALPix maps for multiwavelength and multimessenger astronomy. *The Astronomical Journal*, 163(6):259, 2022.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

Montefalcone, G., Abitbol, M. H., Kodwani, D., and Grumitt, R. Inpainting CMB maps using partial convolutional neural networks. *Journal of Cosmology and Astroparticle Physics*, 2021(03):055, 2021.

Müller, T., Evans, A., Schied, C., and Keller, A. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4): 1–15, 2022.

Park, J. J., Florence, P., Straub, J., Newcombe, R., and Lovegrove, S. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 165–174, 2019.

Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., and Zhou, X. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9054–9063, 2021.

Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., and Courville, A. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301–5310. PMLR, 2019.

Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., and Baraniuk, R. G. WIRE: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18507–18516, 2023.

Schwarz, J. R., Tack, J., Teh, Y. W., Lee, J., and Shin, J. Modality-agnostic variational compression of implicit neural representations. In *International Conference on Machine Learning*, 2023.

Sitzmann, V., Martel, J., Bergman, A., Lindell, D., and Wetzstein, G. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33:7462–7473, 2020.

Spergel, D. N., Verde, L., Peiris, H. V., Komatsu, E., Nolta, M., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., et al. First-year Wilkinson microwave anisotropy probe (WMAP) observations: determination of cosmological parameters. *The Astrophysical Journal Supplement Series*, 148(1):175, 2003.

Stengel, K., Glaws, A., Hettinger, D., and King, R. N. Adversarial super-resolution of climatological wind and solar data. *Proceedings of the National Academy of Sciences*, 117(29):16805–16815, 2020.

Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33: 7537–7547, 2020.

Wang, G.-J., Shi, H.-L., Yan, Y.-P., Xia, J.-Q., Zhao, Y.-Y., Li, S.-Y., and Li, J.-F. Recovering the CMB signal with machine learning. *The Astrophysical Journal Supplement Series*, 260(1):13, 2022.

Wang, J., Liu, Z., Foster, I., Chang, W., Kettimuthu, R., and Kotamarthi, V. R. Fast and accurate learned multiresolution dynamical downscaling for precipitation. *Geoscientific Model Development*, 14(10):6355–6372, 2021.

Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., and Change Loy, C. ESRGAN: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pp. 0–0, 2018.

Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., and Sridhar, S. Neural fields in visual computing and beyond. *arXiv preprint 2111.11426*, 2021.

Yang, H., Yan, J., Liu, Y., and Song, Z. Statistical downscaling of numerical weather prediction based on convolutional neural networks. *Global Energy Interconnection*, 5(2):217–225, 2022.

Yuval, J. and O'Gorman, P. A. Stable machine-learning parameterization of subgrid processes for climate modeling at a range of resolutions. *Nature communications*, 11(1): 3295, 2020.

A. About HEALPix

HEALPix (Hierarchical equal area isolatitude pixelation) is an algorithm for equally partitioning the sphere. This partitioning is done by an iterative subdivision process that starts with twelve base pixels (partitions) that are recursively divided into smaller pixels, creating a hierarchical structure that allows for different levels of resolution.

As a result, a pixel is the smallest unit that comprises the HEALPix grid. The centers of all pixels align with the isolatitude lines, the horizontal lines sharing the same latitude. The total number of pixels determines the resolution of the grid, where the base resolution (i.e., the coarsest resolution) consists of twelve pixels with three isolatitude lines. As the number of pixels increases, so does the resolution of the grid. Specifically, the resolution increases by the number of divisions along the side of the base resolution pixel N_{side}. It is defined by $N_{\text{side}} = 2^{\ell-1}$, for $\ell \in \{1, \dots, L\}$ where ℓ indicates the level of the grid. In addition, the number of the isolatitude lines is given by $N_{\text{line}} = 4 \times N_{\text{side}} - 1$. For example, the base resolution has $N_{\text{side}} = 1$ and $N_{\text{line}} = 3$.

The arrangement principle of the pixels varies according to their latitudinal position. Considering latitude $\psi \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, pixels on the isolatitude lines satisfying $|\cos\psi| \leq \frac{2}{3}$ (i.e., near-by the equator) consist of the same number of pixels, defined by $N_{\text{eq}} = 4 \times N_{\text{side}}$. For example, the base resolution has $N_{\text{eq}} = 4 \times 1 = 4$. On the other hand, for the pixels on the isolatitude line with $|\cos\psi| > \frac{2}{3}$ (i.e., near-by the pole), the number of pixels increases by line moving away from the pole. In detail, there are always four pixels surrounding the North Pole. Starting from these four pixels, the number of the pixels on the next closest line to the North Pole increases by one for each upper pixel. This principle is identically applied to the region surrounding the South Pole.

B. Weighted RMSE results

Table 5. **Results on the real-world datasets.** Both weather datasets (Geo. and Temp.) and CMB dataset are assessed by weighted RMSE. Each method is evaluated for $\times 2$ and $\times 4$ super-resolution tasks and regression (Reg.) task. The units for geopotential, temperature, and cosmic microwave background are m^2/s^2, K, and μK_{CMB}, respectively.

		ReLU + P.E.	SIREN	SHINR	WIRE	FFN	HashGrid	**Ours**
Geo.	$\times 2$	12.83 ± 2.30	21.81 ± 5.20	25.11 ± 3.19	21.98 ± 2.72	13.44 ± 0.61	175.09 ± 6.04	$\mathbf{4.89 \pm 1.05}$
	$\times 4$	20.21 ± 3.55	36.06 ± 10.68	28.18 ± 6.73	29.56 ± 3.67	26.27 ± 7.59	454.04 ± 9.05	$\mathbf{12.10 \pm 3.11}$
	Reg.	11.88 ± 2.39	14.05 ± 2.03	22.22 ± 2.30	6.33 ± 1.14	13.71 ± 5.70	53.89 ± 2.43	$\mathbf{5.82 \pm 3.41}$
Temp.	$\times 2$	0.13 ± 0.03	0.15 ± 0.03	0.36 ± 0.03	0.14 ± 0.01	0.09 ± 0.02	1.00 ± 0.03	$\mathbf{0.07 \pm 0.01}$
	$\times 4$	0.16 ± 0.00	0.44 ± 0.46	0.38 ± 0.04	0.22 ± 0.04	0.16 ± 0.01	2.52 ± 0.06	$\mathbf{0.13 \pm 0.00}$
	Reg.	0.09 ± 0.01	0.11 ± 0.01	0.33 ± 0.11	0.08 ± 0.01	0.09 ± 0.01	0.31 ± 0.04	$\mathbf{0.04 \pm 0.00}$
CMB	$\times 2$	103.21 ± 0.18	90.80 ± 0.77	102.78 ± 0.06	64.94 ± 2.78	81.91 ± 1.59	94.27 ± 4.09	$\mathbf{48.49 \pm 7.93}$
	$\times 4$	103.36 ± 0.17	91.31 ± 0.61	102.83 ± 0.06	69.41 ± 1.32	84.83 ± 1.52	108.56 ± 0.08	$\mathbf{61.82 \pm 2.56}$
	Reg.	87.94 ± 0.51	52.53 ± 1.23	98.78 ± 0.13	42.58 ± 0.45	53.10 ± 0.24	69.08 ± 0.64	$\mathbf{16.60 \pm 0.18}$

C. Architecture Details

Equirectangular feature-grid. For the equirectangular feature-grid we adjust the level of the hierarchical map and the base resolution of the feature-grid, by considering the resolution of the training data. For example, comparing the task of $\times 2$ and $\times 4$ super-resolution, we assign the higher level and base resolution for the $\times 2$ task. Mainly, we used upscaling factor $\gamma = 1.5$. We initialized the parameter for the feature-grid with uniform initialization $\mathcal{U}(-1\text{e}{-4}, 1\text{e}{-4})$.

HEALPix feature-grid. Similarly to the equirectangular feature-grid, we adjust the feature-grid level by the resolution of the training data. The choice of the level is quite intuitive. In the case of super-resolution and regression tasks, mostly the optimal choice was to choose the level as $\ell_{\text{grid}} = \ell_{\text{data}} - 1$. This choice was beneficial since it enables overall parameters to be trained while capturing the fine details. The parameter of the feature-grid is initialized by normal initialization $\mathcal{N}(0, 1\text{e}{-4})$.

D. Baseline Details

ReLU+P.E. (Mildenhall et al., 2021) is a positional encoding defined as:

$$\gamma(x) = \left(\sin\left(2^0\pi x\right), \cos\left(2^0\pi x\right), \cdots, \sin\left(2^{L-1}\pi x\right), \cos\left(2^{L-1}\pi x\right)\right)$$

which separately encodes each dimension of the input point $\mathbf{x} \in \mathbb{R}^n$ with sinusoidals in a axis-aligned and logarithmically spaced manner. In this study, we adjust the level of the frequency L with $\{2, 5, 7\}$.

SIREN (Sitzmann et al., 2020) is a sinusoidal activation $\sin(\omega_0 W\mathbf{x} + b)$ where the frequency of a sinusoidal function is composed by linear weight and scaled by ω_0. Here, ω_0 is set as a hyperparameter and we adjust its range within $\{10, 20, 30\}$.

FFN (Tancik et al., 2020) is a random fourier feature positional encoding $\gamma(\mathbf{x}) = ((\cos(2\pi\mathbf{B}\mathbf{x}), \sin(2\pi\mathbf{B}\mathbf{x}))^T$ where $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{B} \in \mathbb{R}^{m\times n}$ is a sampled frequency from $\mathcal{N}(0, \sigma^2)$. In our study, we adjust σ in the range of $\{2, 5, 10\}$.

WIRE (Saragadam et al., 2023) is a Gabor wavelet based activation $\exp(j\omega_0\mathbf{x})\exp(-|s_0\mathbf{x}|^2)$, which assigns the scale (ω_0) and width (s_0) as hyper parameters. In our study, we adjust both scale and width parameters in the range of $\{1, 10\}$.

SHINR (Esteves et al., 2022) is a spherical harmonics based positional encoding $\gamma(\mathbf{x}) = \bigoplus_{l=0}^{L}\bigoplus_{m=-l}^{l} Y_m^l(\theta, \psi)$, which takes spherical coordinate (θ, ψ) and L is a level of the encoding. In our experiment, we adjust the level L in the range of $\{2, 3, 4\}$.

E. Experimental Details

E.1. Hyperparameters.

In this section, we present the spectrum of hyperparameters applied to our approach. Without loss of generality, the hyperparameters selected for weather and climate datasets correspond to those used for the equirectangular feature-grid, while those chosen for cosmic microwave background data pertain to the HEALPix feature-grid. The common settings across the tasks are 4 hidden layers, 256 hidden dimensions, ReLU activation function, and AdamW optimizer (Loshchilov & Hutter, 2019) for the MLP.

Table 6. Hyperparameter table for super-resolution task for geopotential, temperature and CMB data.

INR training	**Geo. SR (×2)**	**Geo. SR (×4)**	**Temp. SR (×2)**	**Temp. (×4)**	**CMB (×2)**	**CMB (×4)**
Level (L)	$\{8, 6\}$	$\{8, 6\}$	$\{8, 6\}$	$\{8, 6\}$	$\{11, 10\}$	$\{10, 9\}$
Parameter dim. (d)	$\{2, 3\}$	$\{2, 3\}$	$\{2, 3\}$	$\{2, 3\}$	$\{2, 3\}$	$\{2, 3\}$
Scaling factor (γ)	1.5	1.5	1.5	1.5	–	–
Base resolution	16	16	16	16	–	–

Table 7. Hyperparameter table for regression and temporal task for geopotential, temperature and CMB data. **1 day** stands for daily and **1 week** stands for weekly interpolation task.

INR training	**Geo. (Reg.)**	**Temp. (Reg.)**	**CMB (Reg.)**	**Geo. (1 day)**	**Temp. (1 day)**	**Geo. (1 week)**	**Temp. (1 week)**
Level (L)	$\{8, 6\}$	$\{8, 6\}$	$\{11, 12\}$	$\{7, 8\}$	$\{6, 8\}$	$\{8\}$	$\{6, 8\}$
Parameter dim. (d)	$\{2, 3\}$	$\{2, 3\}$	$\{2, 3\}$	$\{6, 8\}$	$\{2, 3\}$	$\{8, 9\}$	$\{2, 3\}$
Scaling factor (γ)	1.5	1.5	–	1.5	1.5	$\{1.1, 1.2\}$	1.5
Base resolution	16	16	–	16	4	4	16

E.2. Training Details.

We train our model using the weighted RMSE (Huang & Hoefler, 2023) and AdamW optimizer (Loshchilov & Hutter, 2019) with a learning rate of 1e-5. For the compression task, we employ the weighted RMSE and use Adam optimizer (Kingma & Ba, 2015) with a learning rate of 1e-2. We utilize the C3 (Kim et al., 2024) framework for compression task which also utilizes the learnable positional encoding to encode the data. Specifically, we apply our equirectangular grid feature instead of the original learnable positional grid which is targeted for a single or a series of 2D image data.

E.3. Task Details.

Super-resolution. The study applies two settings for enhancing the resolution of CMB, weather, and climate data: $\times 2$ and $\times 4$. In the $\times 2$ setting, the model is trained using 361×720 training points, each representing a 0.50° degree. For the $\times 4$ setting, training involves 181×360 points, with each point representing a 1.00° degree. The model's performance is

assessed on its ability to predict data at a finer resolution of $0.25°$, across a total of 721×1440 points. In the context of CMB super-resolution, the $\times 2$ setting involves training on data downsampled to $N_{\text{side}} = 1024$, yielding approximately 10M training points at a resolution of about $N\mu$min. For the $\times 4$ setting, the training data is further downsampled to $N_{\text{side}} = 512$, with around 30K points. The prediction task targets data with $N_{\text{side}} = 2048$, encompassing about 50M points at the full resolution provided by the Planck satellite.

Regression. The regression task involves dividing data points into training and testing sets in an 8:2 ratio. Specifically, for weather and climate data, the split is conducted on geopotential and temperature datasets that total 721×1440 points, resulting in approximately 83K training points and 20K test points. For CMB data, the split includes 50M points, allocating 40M for training and 10M for testing.

Temporal interpolation. This task employs two temporal resolutions: daily and weekly, with a total of 30 data snapshots sampled. Here, snapshots indicate the data that represents the captured moment. The spacing between snapshots corresponds to the chosen temporal resolution. Training data is constructed by selecting the first snapshot, skipping the next, and then selecting the third snapshot. Conversely, test data is compiled by skipping the first snapshot, selecting the second, and continuing this pattern. This method ensures a consistent and desired time step between training and test datasets.

Compression. The two settings of the compression task are divided by the bits per pixel (BPP). The BPP is computed by dividing the model size by the number of pixels. On one hand, for the spatiotemporal compression task, the number of pixels is the multiplication of the number of time indices, pressure levels, and the number of latitude and longitude levels. In our dataset, they are 366, 11, 361, and 720, respectively, following Huang & Hoefler (2023). On the other hand, for the spatial compression task, the number of pixels is the multiplication of the number of latitude and longitude levels.