
GeoWavelets: Spherical Wavelets for Fair Implicit Representations of Earth Data

Daniel Cai*

Department of Computer Science
Brown University
daniel_cai@brown.edu

Randall Balestrierio†

Department of Computer Science
Brown University
randall_balestrierio@brown.edu

Abstract

In the realm of geographic data, implicit neural representations (INRs) often incorporate location embeddings to enhance performance in downstream tasks. Recent advancements in this field have introduced location encodings based on Fourier domain decomposition. However, given the localized nature of many geographic signals, these approaches may introduce significant biases against certain geographic subgroups. To this end, we propose an alternative encoding mechanism, lifting the theoretical guarantees of spherical wavelets. Leveraging the FAIR-Earth dataset (1), we demonstrate that our novel encoding successfully mitigates biases in localized regions while simultaneously maintaining competitive global performance. Our approach represents a significant step forward in creating more equitable and accurate INRs for geographic data.

1 Introduction

In recent years, implicit neural representations (INRs) have emerged as a novel approach to geospatial modeling, offering a continuous and compact representation of complex signals (4). As INRs are being deployed increasingly for various geospatial tasks, it is crucial to understand and emphasize fairness of solutions beyond average-case performance. This of course is practically motivated: for tasks like natural disaster risk assessment where consequences are severe, there is a natural emphasis on improving worst-case performance, rather than average-case metrics (10). In particular, the multi-scale nature of geographic phenomena necessitates approaches that reconcile global patterns with localized ones.

Recent work in the INR domain (13) centers on advancing a certain portion of the pipeline—input embeddings—via decomposition onto the harmonic domain. While these Fourier-based approaches have shown efficacy in capturing global geographic patterns, it is well-established both theoretically and empirically that the global support of Fourier bases introduces biases against localized signals, compromising more fine-scaled geographic representations (7; 1).

To address these limitations, we propose a novel encoding mechanism grounded in existing research into spherical wavelets. Then, leveraging the Fairness Assessment for Implicit Representations of Earth Data (FAIR-Earth) for fine-grained evaluation, we show our spherical wavelet encodings reconcile global signals with local variations on the sphere, reducing bias in localized regions significantly compared to spherical harmonic encodings (1). Moreover, we qualitatively evaluate the corrections achieved with spherical wavelets, and confirm that models can concurrently maintain competitive global performance.

*Equal contribution

†Equal contribution

2 Background

Departing from the spherical harmonic encodings proposed in (13), we implement an encoding based on existing literature in spherical wavelets by (5) and (12). The foundation of our approach lies in the fundamental properties developed from one-dimensional wavelets, which have revolutionized the task of multiresolution decomposition of signals.

In the context of geospatial tasks, wavelets possess several key characteristics for natural data (14). Unlike Fourier transforms, wavelets provide excellent localization in both time/space and frequency domains, allowing for efficient representation of local features and discontinuities in signals (7). Similarly, wavelets facilitate the decomposition of signals at multiple scales, enabling the representation of both fine details and broad structures. Finally, the existence of various wavelet families (e.g., Haar, Daubechies, Mexican Hat), allows for flexibility in choosing appropriate bases (11).

Beyond the one-dimension case, there exist many extensions of wavelet analysis to two-dimensional and spherical domains. In particular, (12) introduces two different schemes for constructing spherical wavelets: via direct stereographic projection of Euclidean wavelets, and via direct construction in harmonic space. The former method is well-defined only in the real space, whereas the latter is only well-defined in harmonic space; following the convention in (13), we later construct encodings based on projection for ease of analysis in the real domain.

2.1 FAIR-Earth

To compare performance across different subgroups, including localized features such as islands and fine groups like coastlines, we leverage the FAIR-Earth dataset, which offers both high-resolution and subgroup monitoring (1). In contrast to the data initially used in (13), FAIR-Earth offers high-resolution granularity along multiple modalities, motivated by pressing representation tasks in climate and carbon dioxide monitoring. For the purpose of analysis, we examine the land-sea data, which employs a uniform $0.1^\circ \times 0.1^\circ$ gridding of the globe, yielding a consistent 1800×3600 map size. A full dataset description and set of visualizations are available in the Table 2.

3 Spherical Wavelets for Geographic Location Encoding

3.1 Spherical Morlet Wavelet

As introduced in (12), the inverse stereographic projection of an admissible *Euclidean* wavelet yields an admissible *spherical* wavelet; that is, the projection satisfies the necessary zero-mean condition (5)

$$C_\psi \equiv \int_{S^2} d\mu(\theta, \varphi) \frac{\psi(\theta, \varphi)}{1 + \cos \theta} = 0$$

Thus, applying the inverse projection onto the 2D Morlet wavelet with width factor and wave number $w, k = 1$ yields the admissible spherical Morlet mother wavelet:

$$\psi_M(\theta, \phi) = [\Pi^{-1}\psi_{\mathbb{R}^2}](\theta, \phi) = \frac{e^{i \tan(\theta/2) \cos(\phi)} e^{-(1/2) \tan^2(\theta/2)}}{1 + \cos \theta}$$

3.2 Lifting Scheme

From an admissible mother wavelet, we can now construct a wavelet basis from affine transformation on the sphere. Analogous to time-frequency localization of Euclidean wavelets, spherical wavelets provide localization via rotation and dilation. The rotation operator $\mathcal{R}(\rho) \equiv \mathcal{R}(\alpha, \beta, \gamma)$ and dilation operator $\mathcal{D}(a)$ both transport a function $f \in L^2(\mathbb{S}^2, d\mu(\theta, \phi))$; refer to (12) for rigid formulation. Applying these transformations, we can then strictly define a set of orthogonal basis functions

$$\{\psi_{a,\rho} \equiv \mathcal{R}(\rho)\mathcal{D}(a)\psi_M\}$$

Finally, to discretize the reconstruction formula, we follow convention in wavelet literature. First, we split ρ into a uniform grid $\mathcal{C}_N = \{\frac{2\pi n_\alpha}{N}, \frac{2\pi n_\beta}{N}, \frac{2\pi n_\gamma}{N}, 0 \leq n_\alpha, n_\beta, n_\gamma \leq N\}$, where N denotes the

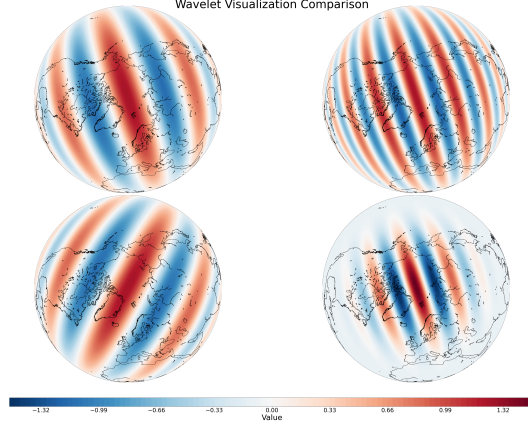


Figure 1: (a) Top Left: Untransformed (b) Top Right: Higher Wave Number (c) Bottom Left: Rotated Wavelet (d) Bottom Right: Dilated Wavelet

granularity of rotations. Then, we take $a \in \{2^{\frac{i}{Q}}, 0 < i \leq M\}$, where $Q \leq 8$ is a fixed value (11). This yields an embedding size on the order of $\mathcal{O}(N^3M)$, and the discretized approximation

$$f(\theta, \phi)_{N,M} \equiv \sum_{\rho \in \mathcal{C}_N} \sum_{i=1}^M w_{a,\rho} \psi_{a,\rho}(\theta, \phi)$$

4 Experiments

Localized Biases are Largely Corrected With Spherical Wavelets To examine the biases presented in (1), we begin by analyzing the trends in training and evaluation, based on a uniform set of hyperparameters. In particular, we leverage FAIR-Earth to compare algorithm performance between local signals and global signals. As a baseline, we compare to results presented in (1), which indicated a striking negative trend between land loss and island loss ($R^2 = 0.59$). To control for confounding factors, we replicate their experimental design (including network design, hyperparameters, etc.), with the sole difference coming in the form of spherical wavelet embeddings.

Upon initial analysis, we see in Fig. 2 and Table 1 the glaring disappearance of this trend; that is, there no longer appears to be a tradeoff between global (land) and localized (island) performance. In fact, for finer resolutions, the correlation is entirely insignificant ($R^2 = 0.04$), and even positive for training resolution of 15000. Compared to the results in (1), we see that when optimized for total loss, our spherical wavelet encoding largely corrects the biases against localized regions.

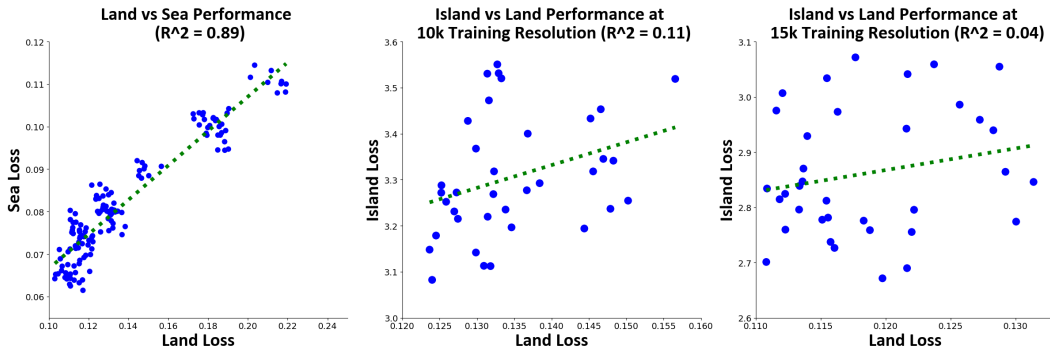


Figure 2: Correlation Analysis Between Local and Global Signals. Contrary to spherical harmonic encodings, there exists insignificant correlation between local and global signals.

Table 1: Comparison of Correlation Between Land and Island Performance By Encoding and Training Resolution

Training Samples	Pearson Correlation (R)	
	Spherical Harmonics	Spherical Wavelets
5000	-0.75	-0.44
10000	-0.77	-0.23
15000	-0.71	0.03
20000	-0.64	-0.21

Spherical Wavelets are Competitive with Existing Methods To better understand the behavior of spherical wavelet encodings, we select a pair of spherical harmonic and spherical wavelet model results, controlled for all other model parameters. In Fig. 3 we note a primary artifact of spherical harmonics encodings that spherical wavelet encodings address: blending of landmasses. Specifically in localized areas, spherical wavelet encodings tend to produce sharper, more fine boundaries compared to spherical harmonic encodings. Moreover, we note that the global performance of spherical wavelets ($BCE_Loss = 0.096$) is on-par with that of spherical harmonics ($BCE_Loss = 0.099$), as indicated by the visual similarities in their learned representations.

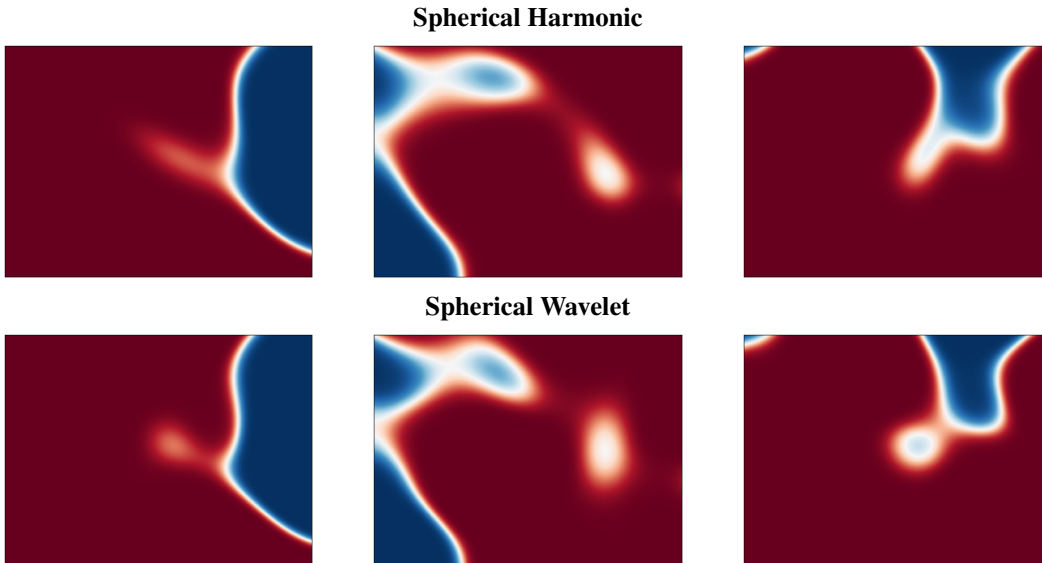


Figure 3: Comparisons of Spherical Harmonic and Spherical Wavelet Behavior in Localized Regions. Full figures available in Fig. 9 and Fig. 10

5 Conclusion

Our construction, experiments, and analyses lay the groundwork for a novel location encoding mechanism based on spherical wavelets. In comparison to spherical harmonic encodings, our encodings produce sharper and more well-defined boundaries in the land-sea classification task. In the process, we leverage the FAIR-Earth dataset (1) to show general corrections to biases against localized regions, a crucial improvement to building fairer INRs of geographic data. While our current study was limited to a subset of land-sea boundary signals, there exist diverse Earth signals that may benefit from a similar nuanced, motivated approach to encoding.

References

- [1] A. Anonymous and B. Anonymous. No location left behind: Introducing the fairness assessment for implicit representations of earth data. *Under Review*, 2024.

- [2] H. E. Beck, E. F. Wood, M. Pan, C. K. Fisher, D. G. Miralles, A. I. J. M. van Dijk, T. R. McVicar, and R. F. Adler. Mswep v2 global 3-hourly 0.1° precipitation: Methodology and quantitative assessment. *Bulletin of the American Meteorological Society*, 100:473–500, 2019. doi: 10.1175/BAMS-D-17-0138.1.
- [3] Center for International Earth Science Information Network - CIESIN - Columbia University. Gridded population of the world, version 4 (gpwv4): Population density adjusted to match 2015 revision un wpp country totals, revision 11. Technical report, NASA Socioeconomic Data and Applications Center (SEDAC), Palisades, New York, 2018. Accessed 29 August, 2024.
- [4] Elijah Cole, Grant Van Horn, Christian Lange, Alexander Shepard, Patrick Leary, Pietro Perona, Scott Loarie, and Oisín Mac Aodha. Spatial implicit neural representations for global-scale species mapping. In *International Conference on Machine Learning*, pages 6320–6342. PMLR, 2023.
- [5] Laurent Demanet and Pierre Vandergheynst. Gabor wavelets on the sphere. In Michael A. Unser, Akram Aldroubi, and Andrew F. Laine, editors, *Wavelets: Applications in Signal and Image Processing X*, volume 5207 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 208–215, November 2003. doi: 10.1117/12.506436.
- [6] X. Dou, J. Hong, P. Ciais, et al. Near-real-time global gridded daily co2 emissions 2021. *Scientific Data*, 10:69, 2023. doi: 10.1038/s41597-023-01963-0.
- [7] Phil PG Dyke and PP Dyke. *An introduction to Laplace transforms and Fourier series*. Springer, 2001.
- [8] G. Huffman, D. Bolvin, D. Braithwaite, K. Hsu, R. Joyce, and P. Xie. Integrated multi-satellite retrievals for gpm (imerg), version 4.4. Technical report, NASA’s Precipitation Processing Center, 2014. URL <ftp://arthurhou.pps.eosdis.nasa.gov/gpmdata/>. Accessed 29 August, 2024.
- [9] D. N. Karger, O. Conrad, J. Böhrer, T. Kawohl, H. Kreft, R. W. Soria-Auza, N. E. Zimmermann, P. Linder, and M. Kessler. Climatologies at high resolution for the earth land surface areas. *Scientific Data*, 4:170122, 2017. doi: 10.1038/sdata.2017.122.
- [10] Luke Kemp, Chi Xu, Joanna Depledge, Kristie L Ebi, Goodwin Gibbins, Timothy A Kohler, Johan Rockström, Marten Scheffer, Hans Joachim Schellnhuber, Will Steffen, and Timothy M Lenton. Climate endgame: Exploring catastrophic climate change scenarios. *Proceedings of the National Academy of Sciences*, 119(34):e2108146119, 2022. doi: 10.1073/pnas.2108146119. PMID: 35914185; PMCID: PMC9407216.
- [11] Stephane Mallat. *A wavelet tour of signal processing*. Academic Press, 1999.
- [12] Jason D. McEwen, Michael P. Hobson, Daniel J. Mortlock, and Anthony N. Lasenby. Fast directional continuous spherical wavelet transform algorithms. *IEEE Transactions on Signal Processing*, 55(2):520–529, February 2007. ISSN 1053-587X. doi: 10.1109/tsp.2006.887148. URL <http://dx.doi.org/10.1109/TSP.2006.887148>.
- [13] Marc Rußwurm, Konstantin Klemmer, Esther Rolf, Robin Zbinden, and Devis Tuia. Geographic location encoding with spherical harmonics and sinusoidal representation networks, 2024. URL <https://arxiv.org/abs/2310.06743>.
- [14] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G. Baraniuk. Wire: Wavelet implicit neural representations, 2023. URL <https://arxiv.org/abs/2301.05187>.

A Appendix

A.1 Dataset Info

Table 2: FAIR-Earth Components

Category	Description	Misc.	Source
Land-Sea	Binary and continuous data, with data on percent water surface coverage for every grid point.	Additional metadata available for island and coastline labeling.	(8)
Population	Population distribution data for regions across the Earth, based on primary sources and interpolation.	Errors in Egypt and Greenland are smoothed via nearest-neighbor interpolation.	(3)
Precipitation	Global precipitation patterns and measurements.	Includes time-slice data from the year 2018 with monthly resolution.	(2)
Temperature	Temperature data and trends across different regions.	Includes time-slice data from the year 2018 with monthly resolution.	(9)
Carbon Dioxide Emissions	Data on CO2 emissions from various sources globally.	–	(6)

A.2 Extended Figures

A.2.1 Dataset Figures

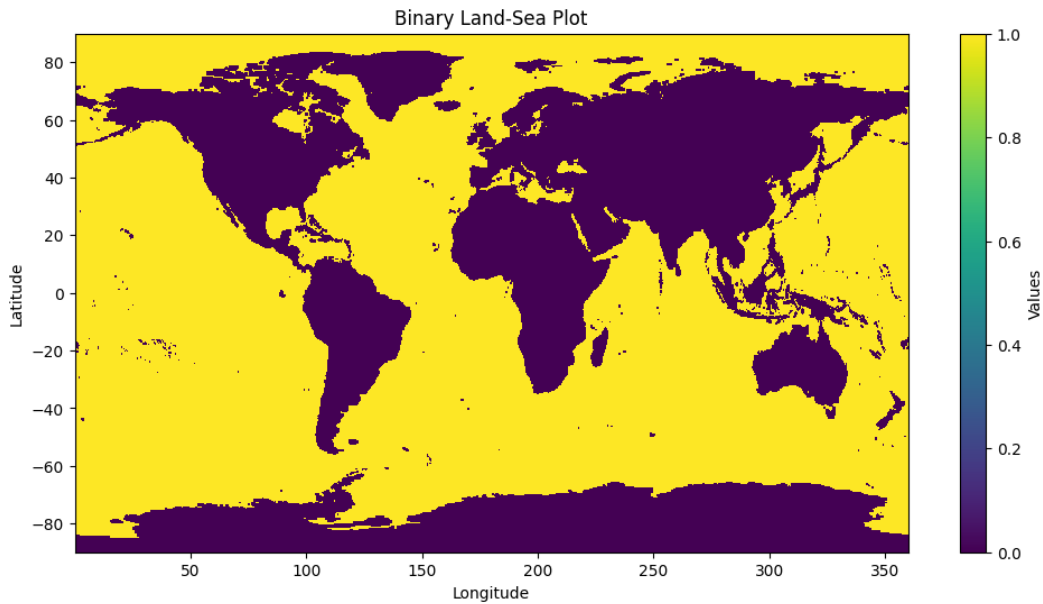


Figure 4: Binary Land-Sea

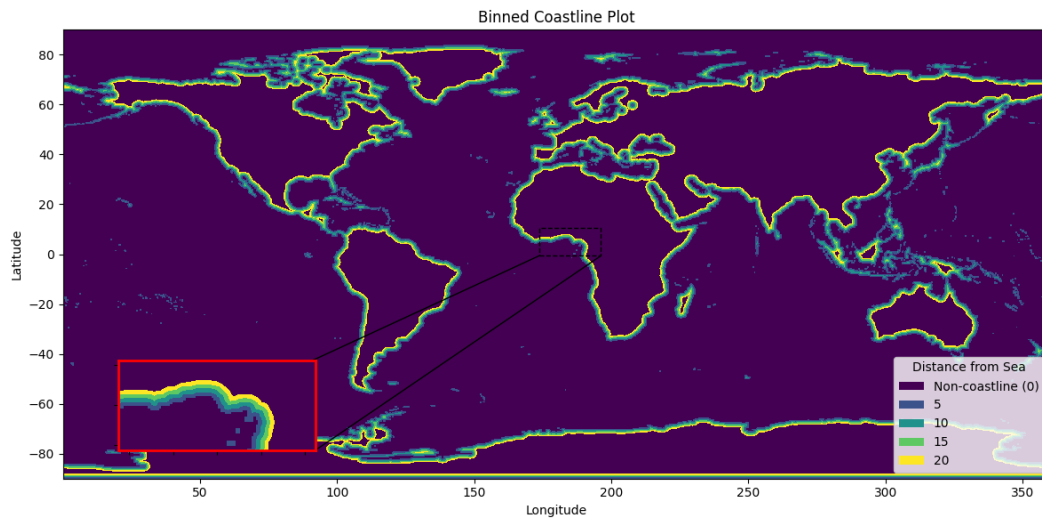


Figure 5: Coastline

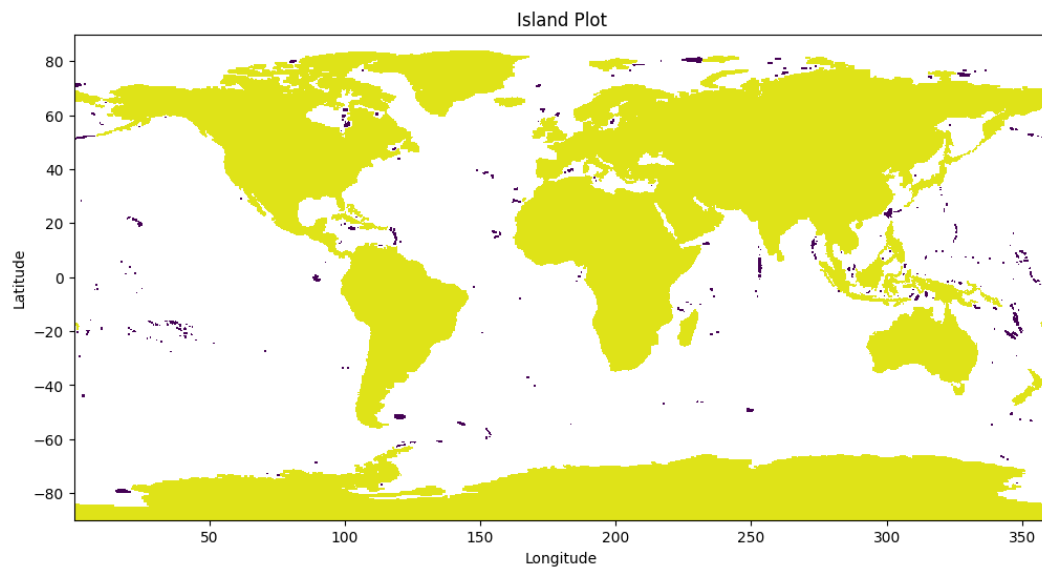


Figure 6: Islands (in Dark Purple)

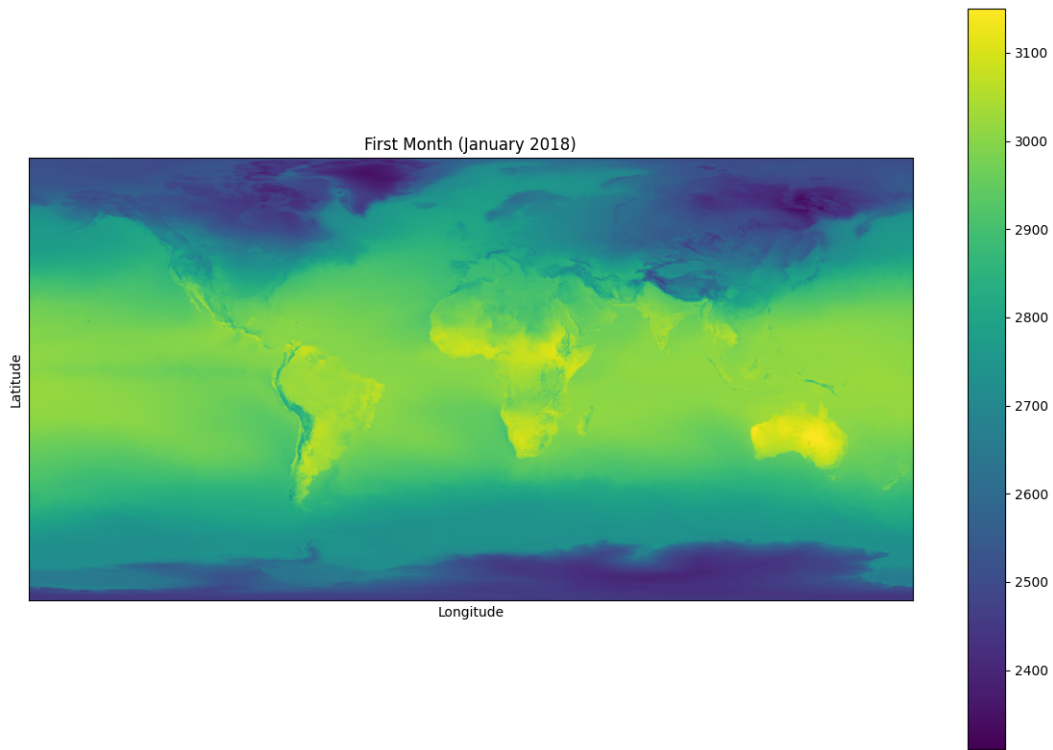


Figure 7: Air Surface Temperature Plot (Jan. 2018)




Figure 8: Cumulative Precipitation Plot (Jan. 2018)

A.2.2 Misc. Figures




Figure 9: Spherical Harmonics Encoding Prediction Plot




Figure 10: Spherical Wavelets Encoding Prediction Plot