Spectral Periodic Networks for Neural Rendering

HALLISON PAZ, IMPA, Brazil
TIAGO NOVELLO, IMPA, Brazil
LUIZ VELHO, IMPA, Brazil

Additional Key Words and Phrases: Periodic Functions; Fourier Series; Texture Mapping; Anti-aliasing

ACM Reference Format:
Hallison Paz, Tiago Novello, and Luiz Velho. 2024. Spectral Periodic Networks for Neural Rendering . *ACM Trans. Graph.* 1, 1 (April 2024), 5 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

One of the fundamental aspects of neural rendering is a scene representation. Different methods and techniques rely on properties of this underlying model to effectively accomplish synthesis and analysis tasks. In that sense, the scene representation should be tailored to the application at hand.

The scope of this work is to develop an *implicit neural representation* (INR) for describing periodic signals in neural rendering. Thus, it aims to encode attribute functions through a periodic neural network $f : \mathbb{R}^n \to \mathcal{A}$, where $\mathcal{A}$ is the attribute space, that can be trained based on data samples. In particular, without loss of generality, this paper concentrates in the case of 2-dimensional functions of visual attributes. More specifically, we explore two important cases in neural rendering: (a) the representation of seamless tileable texture materials and (b) the visualization of 360 degrees spherical panoramas.

We explore sinusoidal neural networks to represent periodic attribute functions. Our approach leverages the Fourier series by initializing the first layer of a sinusoidal network with integer frequencies with a period P. We prove that the compositions of sinusoidal layers generate only integer frequencies with period P. As a result, our network learns a continuous and periodic signal, enabling direct evaluation at any spatial coordinate without the need for interpolation. Our neural implicit representation is compact and suitable for various neural rendering scenarios that deals with periodic signals, such as tileable texture materials and omnidirectional images.

2 SEAMLESS TILEABLE MATERIALS

We find that initializing the first layer of a sinusoidal MLP, f, with periodic neurons having periods (P_1, P_2), ensures that the network itself is periodic with the same periods. Mathematically, this can be expressed as: $f(x_1, x_2) = f(x_1 + P_1, x_2 + P_2)$ For simplicity, consider f as a 1D function; the general case extends similarly.

Authors' addresses: Hallison Paz, IMPA, Rio de Janeiro, RJ, Brazil, hallpaz@impa.br; Tiago Novello, IMPA, Rio de Janeiro, Brazil, tiago.novello@impa.br; Luiz Velho, IMPA, Rio de Janeiro, RJ, Brazil, lvelho@impa.br.

Manuscript submitted to ACM

Assuming the first layer of f is periodic with a period P, its frequencies can be represented as: $\omega_j = k_j \frac{2\pi}{P}$ where $k_j \in \mathbb{Z}$. Furthermore, we observe that a hidden neuron, $h(x) = \sin\left(\sum_i^n a_j \sin(\omega_j x + \varphi_i)\right)$, with weights a_i, can be expanded as a sum of harmonics with period P. This is supported by the identity [Novello 2022]:

$$h(x) = \sum_{\mathbf{l} \in \mathbb{Z}^n} \alpha_{\mathbf{l}}(a) \sin\Big(\langle \mathbf{l}, \omega x + \varphi \rangle + b\Big). \tag{1}$$

Here, $\alpha_{\mathbf{l}}(a)$ is a term dependent solely on a. Each frequency term in this expansion, $\langle \mathbf{l}, \omega \rangle$, can be expressed as $\frac{2\pi}{P} \langle \mathbf{l}, \mathbf{k} \rangle$, confirming that h is periodic with period P. A similar expansion was also presented in [Yüce et al. 2022]. We generalize this finding to sinusoidal MLPs with two hidden layers through induction.

Theorem 2.1. *If the first layer of a sinusoidal MLP f is periodic with period P, then f is also periodic with period P.*

Theorem 2.1 establishes that sinusoidal MLPs can effectively represent periodic functions. We introduce the term *spectral periodic INR* to denote a sinusoidal MLP with a periodic first layer. Utilizing a spectral periodic INR allows for the representation of any tileable image across the entire coordinate space. By training on a small tile and sampling in an expansive domain, the image can be represented comprehensively. Figure 1 illustrates two sinusoidal MLPs trained to depict an image confined within the set $[-1, 1]^2$ and sampled across $[-2, 2]^2$. The network on the left employs our initialization method, whereas the one on the right uses Siren's initialization [Sitzmann et al. 2020]. Note that outside the training interval, Siren exhibits noise while our method accurately replicates the pattern.

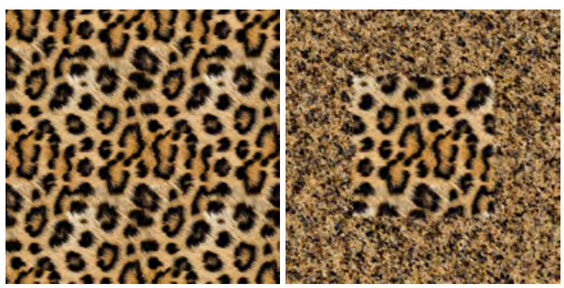

Fig. 1. Comparing our method (left) with Siren (right) on extrapolation.

This behavior is also observed across multiple levels of multiresolution when using a multistage network as MR-Net [Paz et al. 2022, 2023]. Figure 2 shows 3 levels of detail of an M-Net model trained in $[-1, 1]^2$ and sampled in $[-2, 2]^2$.

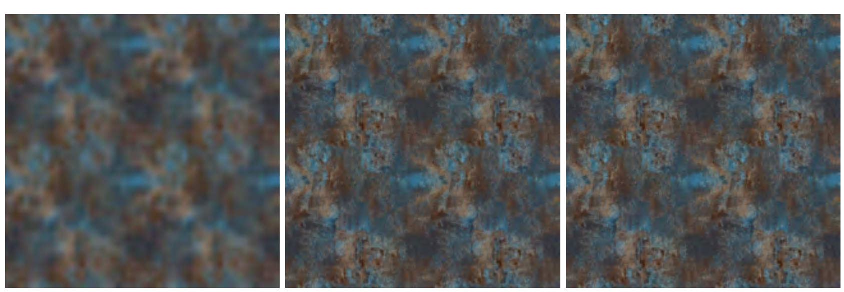

Fig. 2. Reconstructed multiresolution levels extrapolation. Top left: level 2; bottom left: level 4; right: level 6.

Some patterns are not tileable but could be transformed into seamless materials through operations on the INR f. We propose a loss function $\mathscr{L}$ based on Poisson equation to learn a seamless representation based on the sample image $\not{f}$. Since the network is periodic over the whole domain, we can interchange the boundary and interior equations in the image domain Ω. Specifically, we define a function $\lambda : \Omega \to [0, 1]$ indicating that the gradient of ∇f should match the ground truth gradient V if it is close to one, and $\not{f} = f$, otherwise. Thus, $\mathscr{L}$ is given by:

$$\mathscr{L}(\theta) = \int_\Omega \lambda \|\nabla f - V\|^2 dx + \int_\Omega (1 - \lambda)\left(\not{f} - f\right)^2 dx. \tag{2}$$

Using this approach, the boundary is reconstructed primarily using gradient information while the interior propagates the pixels colors through the image. Figure 3 shows the result on a jeans fabric and a patch of a wall.

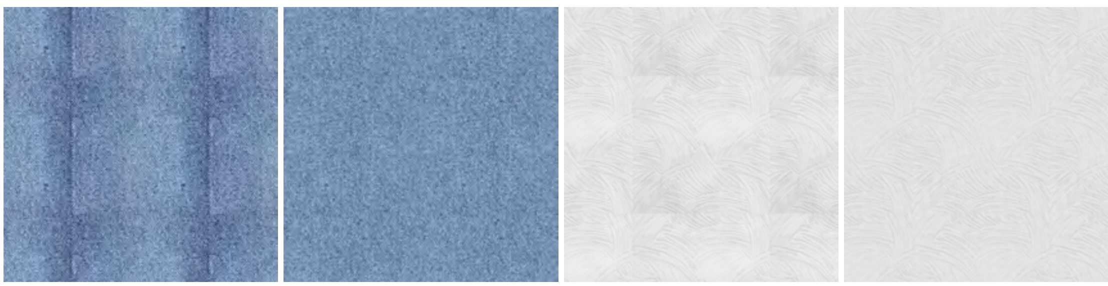

Fig. 3. Left: concatenation of tiles of jeans fabric. Right: seamless reconstruction of the pattern.

3 360 DEGREES PANORAMAS AND NEURAL RENDERING

We present the initial experiments of incorporating our periodic INR into rendering applications. For the underlying INR we adopt the *equirectangular projection mapping* that transforms spherical coordinates, (θ, ϕ), $(\theta \in [-\pi, \pi]; \phi \in [-\pi/2, \pi/2])$ into planar coordinates, (x, y). Giving this parametrization we train our periodic INR by a multiresolution metric-aware adaptive sampling in order to compensate the mapping distortion. This effectively takes into account the stretching of $1/\cos\theta$ along the vertical direction.

Figure 4 demonstrates the adaptive sampling pattern and the reconstruction of a panorama in a INR. Note that, since the image sampling wasn't done in a resolution compatible with the sampling pattern, there are noticeable artifacts around the poles (top and bottom lines of the image). We argue that this could be solved by training a multiresolution INR such as [Paz et al. 2022], due to the compatible filtering.

In our first rendering experiment, we use the PyTorch3D [Ravi et al. 2020] differentiable renderer to visualize a scene with a single sphere parameterized by latitude and longitude, setting the camera in the center of the sphere, and using the panorama as a texture. Figure 5 shows an inner view of the projected equirectangular panorama and an outside view of the scene projected on the sphere.

We propose to use INR of panoramas to exploit the decomposition of a 3D scene into foreground and background using NeRF [Mildenhall et al. 2020]. This is an important case in the representation of NeRF for unbounded scenes. An interesting approach would be to extend DirectVoxGO++ [Perazzo et al. 2022] to incorporate our INR model.

The key idea of the original method is to decompose the 3D scene into two regions divided by a sphere of radius r. The points inside of the sphere are reconstructed as a regular 3D radiance field, whereas the points outside the

Fig. 4. Sampling pattern for equirectangular panorama (top) and its trained reconstruction using a sinusoidal network (bottom).

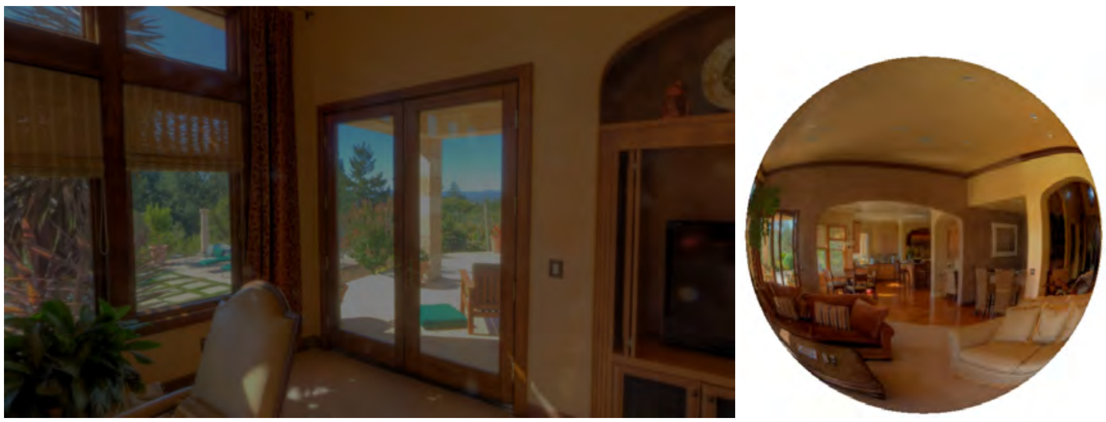

Fig. 5. Rendering of a viewpoint of a panorama (left) and visualization on the sphere (right).

sphere are represented as a spherical panorama. This employs the compactification of the background region using the following mapping: $T(x, y, z, \frac{1}{r}) = (\frac{x}{r}, \frac{y}{r}, \frac{z}{r})$

In our extension, the background would be represented as a periodic INR. Combining our method with [Paz et al. 2022] it would be possible to achieve a compact and multiresolution representation of the background, suitable for anti-aliasing. Figure 6 exemplifies the decomposition of the scene reconstruction into foreground and background.

Fig. 6. Decomposition of a NeRF scene into foreground / background. Composed image (left) and only the foreground (right). Images from [Perazzo et al. 2022].

Note that these ideas, including periodic INRs, could also be applied to cylindrical panoramas.

4 CONCLUSION AND FUTURE WORK

Lore Ipsum

REFERENCES

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.

Tiago Novello. 2022. Understanding Sinusoidal Neural Networks. *arXiv preprint arXiv:2212.01833* (2022).

Hallison Paz, Tiago Novello, Vinicius Silva, Guilherme Schardong, Luiz Schirmer, Fabio Chagas, Helio Lopes, and Luiz Velho. 2022. Multiresolution Neural Networks for Imaging. In *2022 35th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, Vol. 1. 174–179. https://doi.org/10.1109/SIBGRAPI55357.2022.9991765

Hallison Paz, Daniel Perazzo, Tiago Novello, Guilherme Schardong, Luiz Schirmer, Vinicius da Silva, Daniel Yukimura, Fabio Chagas, Helio Lopes, and Luiz Velho. 2023. MR-Net: Multiresolution sinusoidal neural networks. *Computers & Graphics* (2023).

Daniel Perazzo, Joao Paulo Lima, Luiz Velho, and Veronica Teichrieb. 2022. DirectVoxGO++: Fast Neural Radiance Fields for Object Reconstruction. In *Proceedings of SIBGRAPI*.

Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. 2020. Accelerating 3D Deep Learning with PyTorch3D. *arXiv:2007.08501* (2020).

Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. 2020. Implicit Neural Representations with Periodic Activation Functions. In *Proc. NeurIPS*.

Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. 2022. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 19228–19238.