NASA-INR: NON-LINEAR ADAPTIVE SPECTRALLY-SPANNING ACTIVATION FOR IMPLICIT NEURAL REPRESENTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Implicit neural representations (INRs) have emerged as a powerful approach to model complex data offering continuous, resolution independent, and differentiable representations. However, existing methods do not explicitly address symmetry induced spectral attenuation without relying on complex modulation. To counteract this effect, we introduce a non-linear adaptive spectral spanning activation(NASA). We also introduce a non-linear phase expansion (NOPE) for activations in INR. Using harmonic analysis and Chebyshev polynomials, we showcase that the resulting activations contain both odd and even polynomial components, leading to a reduced spectral attenuation in the activation. We test where NASA stands among other existing activations. Furthermore, we investigate how the addition of NOPE to existing activations and adaptive forms of existing activations with tunable parameters compare with their non-modified forms. Finally, the resulting models are further evaluated on image reconstruction, inpainting, super-resolution, and denoising tasks against their classic counterparts. Through our work, nearest SOTA methods were surpassed by a mean increment of 2.57 dB in image reconstruction, 0.61 dB in super resolution, 0.26 dB in image inpainting 0.26. Overall, the study investigates phase asymmetry as a possible design dimension for improving the spectral representation capabilities of implicit neural networks.

1 INTRODUCTION

Implicit neural representations (INRs) use multilayer perceptrons to model signals as continuous functions,and they offer advantages such as resolution independence and memory-efficient representation. This approach has been applied to audio, images, 3D objects, and high-dimensional signals (Essakine et al., 2025). Despite these advantages, INRs remain affected by spectral bias, a frequency dependent learning imbalance that favours smooth structures over fine-scale variations (Rahaman et al., 2019; Jayasundara & Patel, 2026). In image representation, this can weaken the recovery of edges, textures, and other fine details. One well-known approach to addressing this limitation is to design activation functions with greater spectral expressivity.

This direction has evolved from SIREN's sinusoidal formulation toward spatially localized, adaptive, and learnable activation designs (Sitzmann et al., 2020; Jayasundara & Patel, 2026). Beyond the choice of activation family, the symmetry of an activation also influences its spectral contributions. Specifically, odd and even symmetric activations have alternating zero Chebyshev coefficients, which have been linked to attenuation of the post-activation spectrum (Thennakoon et al., 2026). Complex sinusoidal modulation has been proposed to mitigate this attenuation, but it introduces complex-valued intermediate representations (Thennakoon et al., 2026). This motivates investigating whether symmetry breaking can instead be achieved through a direct, real-valued reparameterization applicable across different activation functions.

Building on harmonic analyses of nonlinear networks and implicit neural representations (Mehmeti-Göpel et al., 2021; Yüce et al., 2022), as well as the symmetry-induced spectral attenuation identified by COSMO-INR (Thennakoon et al., 2026), we investigate whether activation symmetry can be broken without complex-valued modulation. As a result, we introduce NOPE-INR, a non-linear

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

phase expansion for sinusoidal activations in INR models. This is designed to mitigate the effects of the odd-even frequency bias while remaining compatible with different INR activations.

Another comparatively underexplored direction in INR design is layer wise adaptive activations. Instead of using one activation configuration throughout the network, each hidden layer can learn its own scale or frequency parameters (Jagtap et al., 2020). This motivated us in formulating the Non-linear Adaptive Sepctrally-spanning Activation (NASA) which is layer wise adaptive while fully utilizing our non-linear phase expantion.

Our main contributions are as follows:

1. Introducing a new activation function NASA, which can adapt it's frequency response without using complex variables which allows models using NASA to match to the image spectrum.

2. An exploration of how layer-wise adaptivity effects INR performance by utilizing layer-wise adaptive versions of classic activations and formulating NASA as a layer wise adaptive activation.

3. Exploring how the non-linear phase expansion (NOPE) used in NASA effects SOTA activation functions.

4. Through a series of benchmark tests, including image representation and computer vision inverse tasks such as denoising, super-resolution and image inpainting showcasing the improvements in performance of using NASA in INR tasks and improvements to some SOTA activations by incorporating a NOPE.

2 Related Work

Implicit neural representations. Recent INR research evaluates methods not only in terms of representation quality, but also their spectral expressivity, computational efficiency, and robustness across tasks (Essakine et al., 2025). The existing methods for INR can generally be classified according to whether they modify positional encoding, activation-function design, or network and training parameterization, with some approaches combining multiple strategies (Essakine et al., 2025; Jayasundara & Patel, 2026). Among these directions, activation functions are particularly important because they strongly influence the frequency content learned by an INR (Jayasundara & Patel, 2026).

Effect of activation functions. Early coordinate-based MLPs commonly used ReLU activations, whose piecewise-linear structure limits fine-detail and derivative representation (Sitzmann et al., 2020). Such networks also tend to learn low-frequency components before high-frequency ones (Rahaman et al., 2019). SIREN addressed this limitation by introducing periodic sinusoidal activations, enabling INRs to represent complex signals and their derivatives more effectively (Sitzmann et al., 2020). However, sinusoidal activations are not spatially localized and can produce ringing artifacts around sharp features. WIRE reduces these artifacts using Gabor-wavelet activations that provide spatial and frequency localization (Saragadam et al., 2023). Beyond these approaches, more recent methods explore flexible and adaptive activation designs. FINER uses a variable periodic activation (Liu et al., 2024). MIRE selects a matched activation for each layer from a dictionary (Jayasundara et al., 2025), and SL2A-INR parameterizes a learnable activation using Chebyshev polynomials (Rezaeian et al., 2025). More recently, MSA introduced parallel sine activations with different frequency scales to improve multiscale representation (Han et al., 2026). These developments demonstrate that activation design strongly influences the spatial-frequency behavior of INRs, motivating a closer examination of spectral bias.

Layer-wise adaptive activations. (Jagtap et al., 2020) introduced layer-wise locally adaptive activation functions (L-LAAF) by introducing a scalable parameter inside the function itself and optimizing it using a stochastic gradient descent algorithm in deep and physics-informed neural networks. Later, (Jagtap et al., 2022) introduced the Rowdy activation function, which includes trainable parameters and can be used in any type of neural network. In addition, Rowdy captures high-frequency components in the target function by introducing high-frequency sinusoidal fluctuations in the activation functions, proposing a potential method to mitigate the spectral bias problem in neural networks. Furthermore,(Morsali et al., 2025) used the idea of layer-wise locally adaptive activation functions directly in INRs, where they used an activation parameterized in the form of a

Fourier series where the amplitude, frequency and phase parameters are learned dynamically during training, enabling the network to adapt its activation function to the specific characteristics of the signal being modeled.

Spectral bias analysis. Spectral bias describes the tendency of neural networks to learn low-frequency components more readily than high-frequency details (Rahaman et al., 2019). Harmonic analyses show that nonlinear activations can expand spectral support by generating higher-order harmonics, whose magnitudes depend on the activation's polynomial coefficients (Mehmeti-Göpel et al., 2021; Yüce et al., 2022). Building on this analysis, COSMO-INR showed that odd and even symmetric activations have Chebyshev coefficients that alternate between zero, leading to spectral attenuation. In order to overcome this limitation, it introduced complex sinusoidal modulation,

$$g(x) = \phi(x)e^{j\zeta x},$$

to retain nonzero Chebyshev coefficients at both odd and even orders (Thennakoon et al., 2026). Despite its improved spectral support, COSMO-INR produces complex-valued intermediate representations and normalizes activation inputs and outputs onto the unit circle while preserving phase (Thennakoon et al., 2026). Consequently, magnitude variation is not propagated through these normalized representations. In its reported efficiency comparison, the complex-sinusoidal activation and prior embedder also introduce runtime overhead that is not reflected by linear-layer FLOP counts (Thennakoon et al., 2026). Alternatively, I-INR recovers high-frequency details through iterative refinement while adding only 0.8–1.6% additional FLOPs over the corresponding baseline (Haider et al., 2026).

In contrast, our method directly applies the real-valued symmetry-breaking phase reparameterization

$$u \mapsto r_{k,\eta}(u)$$

to each activation input, as formally defined in Section 3.2. This mapping introduces input-dependent local frequencies without complex-valued modulation.

3 METHADOLOGY

An implicit neural representation (INR) behaves as a map from a discrete signal $g(X)$ to its corresponding continuous signal $f_\theta(x)$. Here, the continuous signal is originally unknown and is learned from the sampled data space of the discrete signal. Such a network can be expressed as, $f_\theta : \mathbb{R}^{D_i} \to \mathbb{R}^{D_o}$ where θ denotes the network parameters and D_i and D_o denotes input and output dimensions. The network is trained by minimizing the mean square error, which is given by,

$$L_{\text{MSE}} = \mathbb{E}_{x \in X} \left[\| f_\theta(x) - g(X) \|^2 \right] \tag{1}$$

Activation function $f_\theta(x)$ plays a vital role in capturing the signal spectrum. In this section, we build up on prior INR work to introduce novel methods to improve spectral and signal representation capacity of INRs.

3.1 HARMONIC DISTORTION IN ACTIVATION FUNCTIONS

A nonlinear activation modifies its input spectrum by generating higher-order harmonics. Harmonic-distortion analyses show that different polynomial orders contribute to the post-activation spectrum, with each contribution weighted by the corresponding coefficient (Mehmeti-Göpel et al., 2021; Yüce et al., 2022).

Because many practical activations are not finite polynomials, we follow Thennakoon et al. (2026) and represent an activation over $u \in [-1, 1]$ using the Chebyshev expansion

$$\phi(u) = \sum_{n=0}^{N} a_n T_n(u), \tag{2}$$

where T_n is the Chebyshev polynomial of degree n. For an even activation, the odd-order Chebyshev coefficients are zero, whereas for an odd activation, the even-order coefficients are zero:

$$\phi \text{ even} \implies a_{2q+1} = 0, \qquad \phi \text{ odd} \implies a_{2q} = 0. \tag{3}$$

Therefore, even and odd symmetric activations lack one parity class of polynomial contributions, which can attenuate their post-activation spectra (Thennakoon et al., 2026). This motivates us to investigate a non-linear phase expansion for activations that allows representations to include both odd and even order Chebyshev coefficients.

3.2 Non-linear Phase Expansion for INR

To address the issue of spectral bias arising from spectral attenuation due to the symmetry of activation functions, we introduce a non-linear phase expansion for activations in INRs, referred to as NOPE-INR. This phase expansion can be expressed as follows.

Let

$$u = \mathbf{w}^\top \mathbf{x} + b, \tag{4}$$

where $\mathbf{w}$ and b denote the weights and bias of the neural layer, respectively. We define the non-linear phase expansion as

$$\boxed{r_{k,\eta}(u) = u + ku^2 + \frac{1+\eta}{3}k^2u^3} \tag{5}$$

For example, the reparameterized SIREN activation function would be

$$\sin(\omega_0 u) \longrightarrow \sin\left[\omega_0\left(u + ku^2 + \frac{1+\eta}{3}k^2u^3\right)\right].$$

The effect of adding NOPE to the SIREN and FINER activation functions can be seen seen in Figure 1 showcasing the Chebyshev coefficient energy of each activation. WIRE is used as a control to see how complex functions which are already asymmetric will be affected by the phase expansion. The choice of the cubic term and the coefficient $(1+\eta)/3$ is discussed in Appendix A.1.

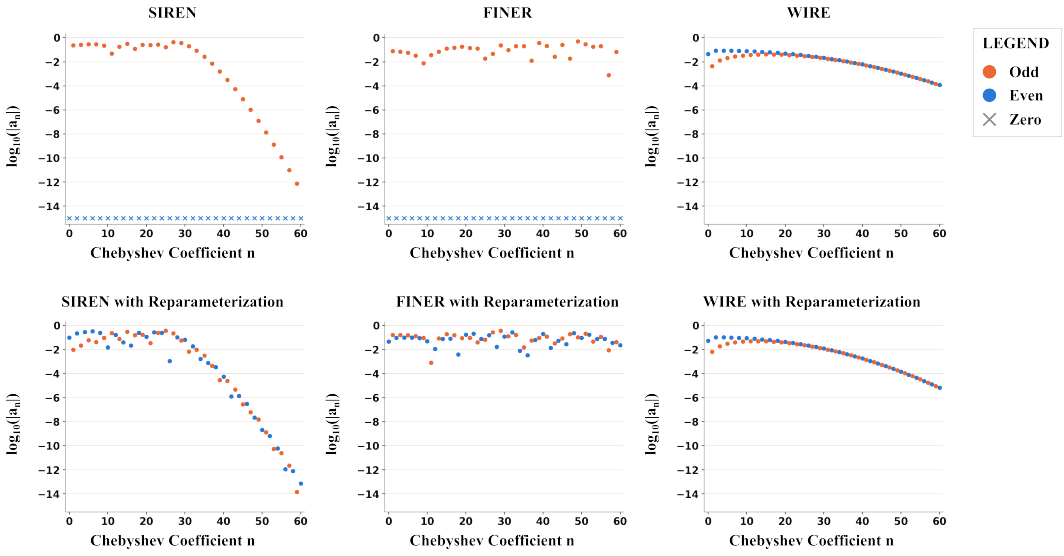


Figure 1: Chebyshev component analysis for SIREN, FINER and WIRE before and after reparameterization

3.3 Non-linear Adaptive Variable-spectral Activation function

The results of adding NOPE to classic activations in image reconstruction tasks are shown Table 2. It's clear from these results that mitigating the odd-even spectral bias dose lead to better performance in the SIREN activation function.

However, these gains do not stay consistent across all activation families and images. This can be attributed to these activations not being designed to fully utilize a non-linear phase expansion.

Therefore, taking inspiration from previous works such as Thennakoon et al. (2026) we formulate a non-linear spectrally-spanning activation function named NASA. The formulation of NASA is shown in Equations 6 - 8.

Let's define a tanh window as,

$$w_{\tanh}(u) = \frac{1}{2}\left[\tanh\left(\frac{u+T/2}{\tau}\right) - \tanh\left(\frac{u-T/2}{\tau}\right)\right], \tag{6}$$

and the asymmetric phase as

$$\phi(u) = (1+c)u + ku^2 + \frac{1+\eta}{3}k^2u^3. \tag{7}$$

We define the NASA function as

$$\boxed{f(u) = w_{\tanh}(u)\sin\left(\omega_0\phi(u)\right)} \tag{8}$$

For the work presented in this paper, $\omega_0 = 30$, $\tau = 0.1$, $\eta = 0.05$, and $k = 0.5$ were used. The c and T variables can shift the functions center frequency and bandwidth as demonstrated in Figure 7. An explantion for the denser Chebychev plot is provided in Appendix refA.2

Work such as Kazerouni et al. (2024), Thennakoon et al. (2026) and Morsali et al. (2025) had demonstrated that having trainable parameters in the activation function significantly improves INR performance. However the use of the ResNet-32 encoder in Kazerouni et al. (2024), Thennakoon et al. (2026) had made the INR models much larger compared to INR models using classic activation.

In our formulation of NASA we let the T and C values tune through the back propagation pipeline which enable the models to achieve the frequency domain freedom while keeping the number of parameters to a minimum.

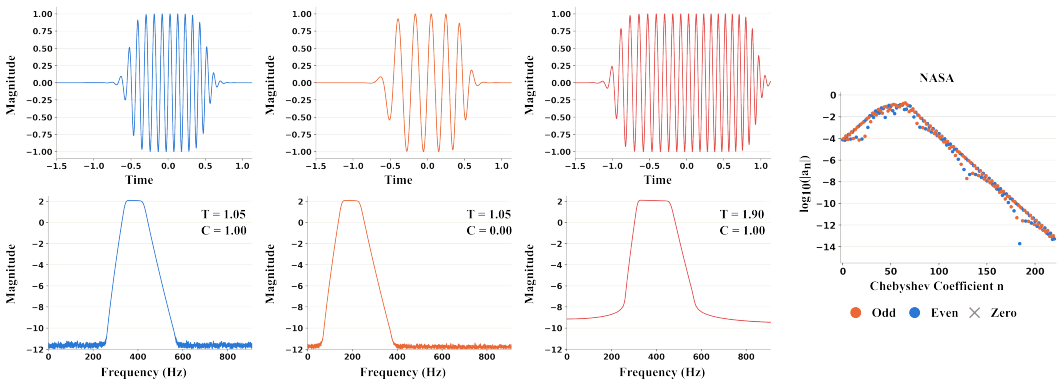


Figure 2: Effect of varying c and T on NASA function.The parameter c shift the center frequency of the response, while T to control the bandwidth. The corresponding chebyshev plot illustrates the chebyshev coefficient distribution for the NASA function.

3.4 Locally layer-wise adaptive activations

NASA activation function is an adaptive activation, which we propose as a potential activation that can improve spectral support. Inspired by the work of (Morsali et al., 2025), to compare the newly formulated activation function with classic activations we have introduced adaptive forms of the classic activations. This approach led us to explore the field of layer-wise adaptive activation functions.

Layer wise adaptive activation functions include a trainable parameter in each layer which is dynamically adjusted during the training process. Locally layer wise adaptive activations (L-LAAF)

carry the trainable parameters inside the function itself(Jagtap et al., 2020). To investigate the performance of INRs employing L-LAAF, we have used adaptive forms of classic activations. Angular frequency of SIREN(Sitzmann et al., 2020) and FINER(Liu et al., 2024) activation functions, angular frequency and spacial decay parameter of WIRE(Saragadam et al., 2023) activation function are the selected activation functions and the corresponding locally adaptive parameters which are adjusted through back propagation.

Table 1: Activation functions and their locally adaptive variables.

Activation	**Equation**	**Locally adaptive variable(s)**
SIREN	$\sin(\omega_a u)$	ω_a
FINER	$\sin(\omega_a(\lvert u\rvert + 1)u)$	ω_a
WIRE	$\exp(j\omega_a u)\ \exp\left[-\sigma_a^2(u^2+v^2)\right]$	$\omega_a,\ \sigma_a$
NASA	$w_{\tanh}(u)\ \sin(\omega_0\phi(u))$	$T,\ c$

4 Results and Discussion

24 lossless images from the Kodak dataset Eastman Kodak Company were used to evaluate the image representation task where the images were trained at native resolution, while images from the DIV2K dataset Agustsson & Timofte (2017) were used to evaluate the super-resolution, inpainting and denoising tasks. The experiments were performed using an NVIDIA RTX 6000 Ada Generation GPU with 48 GB All experiments were based on a PyTorch framework and used Adam as the optimiser. The benchmark test was performed against SIREN(Sitzmann et al., 2020),WIRE(Saragadam et al., 2023), FINER(Liu et al., 2024), their above mentioned variants and INCODE(Kazerouni et al., 2024) .

4.1 Neural Tangent Kernel Analysis

The Neural Tangent Kernel (NTK) models the learning process of a neural network as a kernel regression(Jacot et al., 2018). Utiliizing NTK, the convergence and the ability to represent a target signal can be predicted by analysing the prominence of the main diagonal of the NTK matrix and the spread of its eigenvalues. High prominence of the main diagonal indicates better shift variance and better convergence while large eigenvalues guarantees faster learning of high frequency components (Bai et al., 2023)(Tancik et al., 2020).

We follow the method of (Liu et al., 2024) to calculate the NTK of a simplified network written as $f(x;\theta) = \sum_{k=1}^{n} c_k \sigma(w_k x + b_k)$, employing the NASA activation function. By definition, the NTK of NASA is given by

$$\begin{aligned} K(x_i, x_j) &= \mathbb{E}_\theta \left[\langle \nabla_\theta f(x_i;\theta), \nabla_\theta f(x_j;\theta)\rangle\right] \\ &= (x_i^\top x_j + 1)\, \mathbb{E}_\theta \sum_{r=1}^{n} a_r^2 \\ &\quad \times \Big[w'(g_{ri}) \sin(\omega_0 \Phi(g_{ri})) + \omega_0 w(g_{ri}) \Phi'(g_{ri}) \cos(\omega_0 \Phi(g_{ri}))\Big] \\ &\quad \times \Big[w'(g_{rj}) \sin(\omega_0 \Phi(g_{rj})) + \omega_0 w(g_{rj}) \Phi'(g_{rj}) \cos(\omega_0 \Phi(g_{rj}))\Big]. \end{aligned} \tag{9}$$

Here, $g_{ri} = \mathbf{w}_r^\top \mathbf{x}_i + b_r$, while $w(g_{ri})$ and $\Phi(g_{ri})$ denote the tanh window $w_{\tanh}(u)$ and nonlinear phase $\Phi(u)$ defined in Eqs. (6) and (7), respectively. Their derivatives with respect to the pre-activation are $\Phi'(g_{ri})$ and $w'(g_{ri})$.

Higher values along the main diagonal and very small values as non-diagonal elements of the NTK matrix suggests better shift invariance which improves the learning capability of the INR due to the minimum co-dependance between two inputs points(Jacot et al., 2018)(Tancik et al., 2020).

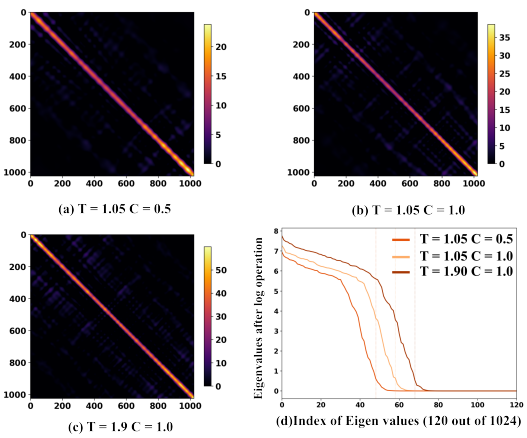


Figure 3: Visualization of the NTKs and the corresponding eigenvalues in NASA. (a)-(c) visualize NTKs when T and c are initialized at 1.05 and 0.5, 1.05 and 1.0, 1.9 and 1.0 respectively. (d) plots the corresponding eigenvalues after log operation. Because the eigenvalues are spread through a large range, they are processed by a log function for visualization.

According to Figures A.3(a)-(c), the stronger main diagonals in the calculated matrices show a strong shift invariance directly resulting in better overall convergence of the model network using NASA activation.

Figure A.3 representing the eigenvalue distribution shows that increased values for T and c significantly expand the outward spread of the curves, producing a greater quantity of large eigenvalues before the exponential decay occurs. Maintaining larger eigenvalues reduces the effect of the spectral bias issue, leading to faster convergence for high-frequency signal components.

4.2 IMAGE RECONSTRUCTION

Table 2: PSNR (dB, single seed 42) on selected nine Kodak images, grouped by activation family. For each image, the highest, second-highest, and third-highest PSNR values are highlighted in light red, light orange, and light yellow, respectively, and shown in bold.

Activation	**kodim01**	**kodim03**	**kodim05**	**kodim07**	**kodim10**	**kodim13**	**kodim15**	**kodim17**	**kodim20**
Siren	32.488	36.591	34.497	37.643	36.312	30.175	34.281	37.042	33.491
Siren-Adaptive	32.425	36.618	34.570	37.581	36.317	30.188	34.201	37.056	33.419
Siren-NOPE	33.073	37.558	34.532	38.769	36.907	30.629	34.607	37.350	34.275
Finer	**35.976**	**39.782**	**36.968**	**40.190**	**39.033**	33.662	36.953	**39.642**	36.925
Finer-Adaptive	35.943	39.778	**36.882**	**40.228**	39.018	**33.703**	36.937	**39.585**	36.885
Finer-NOPE	34.694	**40.018**	35.414	39.579	38.379	32.443	**36.957**	39.021	**37.026**
Wire	24.589	32.304	25.375	31.739	31.237	22.249	30.766	30.290	30.072
Wire-Adaptive	22.934	30.142	22.590	27.380	27.945	20.808	27.973	27.885	27.444
Wire-NOPE	22.966	29.970	22.520	27.436	27.727	20.861	27.893	27.838	27.426
Incode	**36.512**	39.738	35.490	39.633	**39.309**	**33.942**	**37.667**	39.384	**37.807**
NASA	**38.296**	**42.319**	**37.100**	**42.008**	**42.059**	**35.151**	**41.116**	**41.940**	**43.168**

With respect to state of the art (SOTA) models, NASA on Kodak dataset achieves the highest performance in image reconstruction task with an average PSNR of 40.21 dB, outperforming the nearest SOTA method INCODE which remains largely above NASA in terms of parameter cost. It can also be noted that adding NOPE to SIREN improves image reconstruction throughout the tested data set. In contrast WIRE shows a clear decline in performance upon addition of NOPE and adaptive parameters. This observation can be justified by the fact that WIRE is an asymmetric function with no significant effect on its spectral scope after addition of NOPE. Furthermore, the high sensitivity of WIRE performance to function parameters can explain the reduced performance in the adaptive modification.

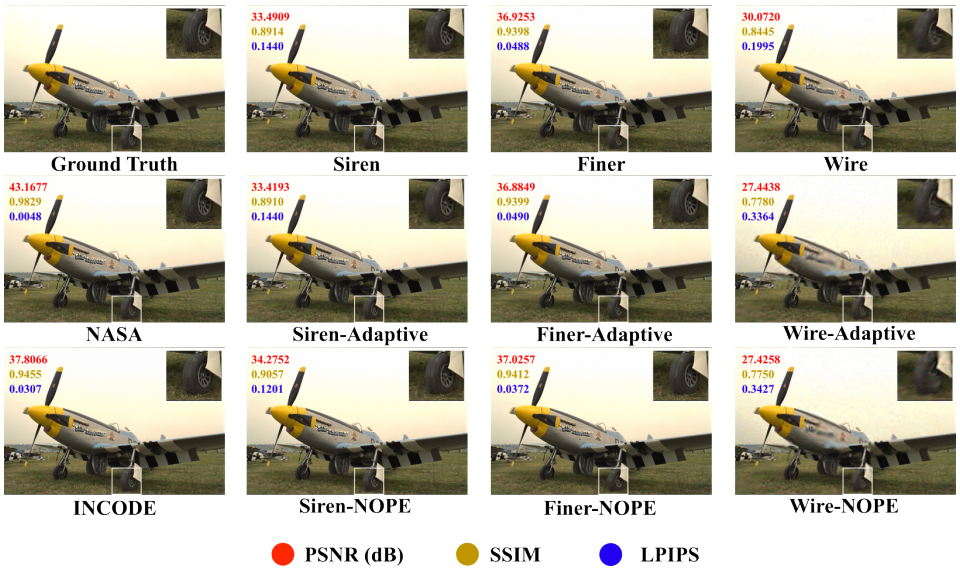


Figure 4: Image representation task compared with SOTA methods for Kodak image 20

4.3 INPAINTING

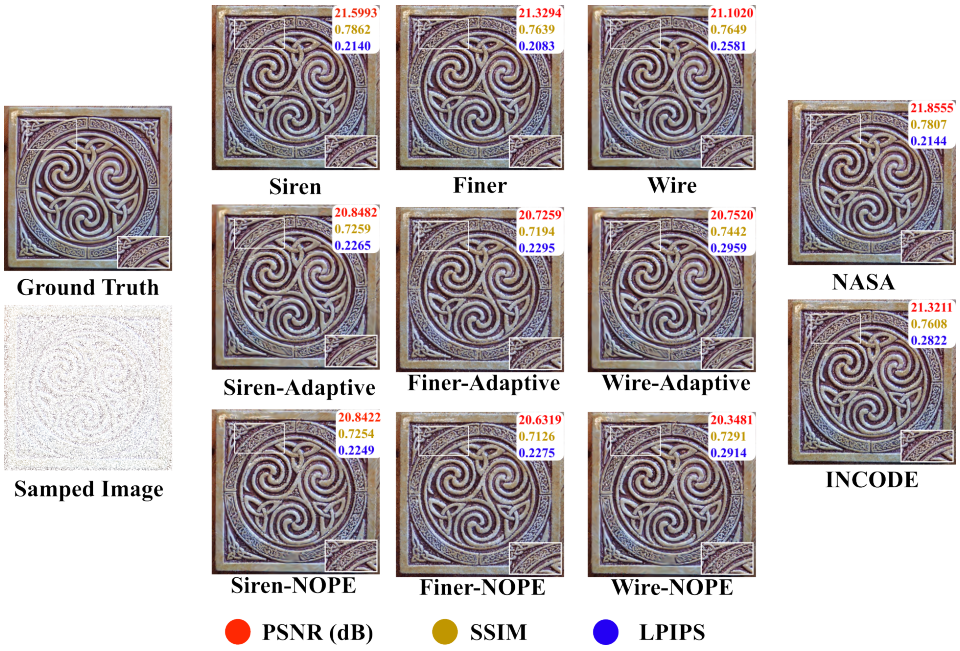


Figure 5: Image inpainting task for the Celtic spiral knots image compared to SOTA methods.

The inpainting task was done utilising a Celtic spiral knots image ($572 \times 582 \times 3$) of which 20% of the pixels from the original image was used with a binary mask applied to identify missing pixels. The model was tasked to reconstruct the full image

As shown in Figure.5, Although The modified activations with adaptive parameters and adding NOPE to classic activations did not show an improvement in metric values, the performance of NASA marginally exceeds SOTA methods in single image inpainting.

4.4 SUPER RESOLUTION

Table 3: Super-resolution PSNR (dB) at factors of 2x, 4x and 6x, grouped by activation family. For each upscale factor, the highest, second-highest, and third-highest PSNR values are highlighted in light red, light orange, and light yellow, respectively, and shown in bold.

Factor	Siren			Finer			Wire			Incode	NASA
	Base	Adap	NOPE	Base	Adap	NOPE	Base	Adap	NOPE		
2×	32.066	32.140	32.155	**32.871**	32.791	**32.917**	31.066	31.043	31.143	30.273	**33.286**
4×	29.710	29.726	29.691	**29.816**	**29.843**	29.796	26.759	27.032	26.791	28.056	**30.355**
6×	**27.332**	**27.331**	27.318	26.872	26.869	26.834	24.067	24.262	24.072	26.403	**27.742**

An image of resolution 2040×1356 was downsampled by factors of 1/2, 1/4 and 1/6 and was utilized to reconstruct the image at the original resolution using the continous nature of INRs.

The results of super resolution experiments given in Table.3 clearly demonstrate how NASA surpasses the SOTA methods. However, No significant effect is observed when adding NOPE or adaptively modifying the classic activations.

4.5 DENOISING

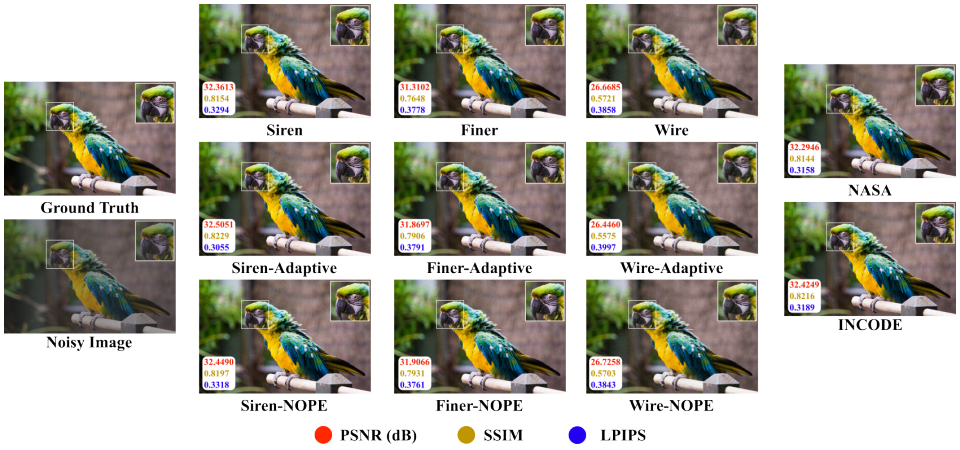


Figure 6: Image denoising task compared with SOTA methods.

The image for the denoising task was prepared by adding photon noise to the ground truth applying independant Poisson random variables to each pixel. The mean photon count and the readout counts were set to 30 and 2 respectively.

In the image denoising task, the addition of adaptive parameters and NOPE has increased performance in symmetric SIREN and FINER activations. Nevertheless, performance of NASA in denoising remains comparable with the SOTA methods with no significant improvement.

5 CONCLUTION

In this work we have proposed and verified NASA, a non-linear adaptive activation that can span the frequency domain without the use of complex variables. Through this we address the effect of the odd-even frequency bias of INRs established in prior work. Furthermore we explore the effects of adding the non-linear phase expansion used in NASA to SOTA activations, as well as exploring the effects of using layer-wise adaptive activations for INRs. Through benchmark tests we demonstrate that NASA out performs SOTA activations in image fitting, inpainting and super resoltion tasks.

6 Additional Information

AI use statement

In this work, we used generative AI tools for grammar checking and correction, LaTeX code generation, and identification of relevant literature. We did not use generative AI tools for research tasks, including the development of theoretical models or conceptual frameworks, formulation of mathematical claims, writing of mathematical proofs, refinement of hypotheses, design of research methodology, implementation of methods. Other required-disclosure tasks, including synthetic dataset generation, translation, dataset cleaning or reformatting, and qualitative or thematic data analysis, were not applicable to this work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

References

Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, July 2017.

Jinshuai Bai, Gui-Rong Liu, Ashish Gupta, Laith Alzubaidi, Xi-Qiao Feng, and YuanTong Gu. Physics-informed radial basis network (pirbn): A local approximating neural network for solving nonlinear partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 415:116290, 2023.

Eastman Kodak Company. Kodak lossless true color image suite. Kodak Lossless True Color Image Suite. Accessed: 2026.

Amer Essakine, Yanqi Cheng, Chun-Wun Cheng, Lipei Zhang, Zhongying Deng, Lei Zhu, Carola-Bibiane Schönlieb, and Angelica I. Aviles-Rivero. Where do we stand with implicit neural representations? a technical and performance survey. *Transactions on Machine Learning Research*, pp. 1–25, 2025. URL https://openreview.net/forum?id=QTsJXSvAI2.

Ali Haider, Muhammad Salman Ali, Maryam Qamar, Tahir Khalil, Soo Ye Kim, Jihyong Oh, Enzo Tartaglione, and Sung-Ho Bae. I-INR: Iterative implicit neural representations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(6):4520–4528, 2026. doi: 10.1609/aaai.v40i6.42451.

Jufeng Han, Shu Wei, Min Wu, Lina Yu, Weijun Li, Linjun Sun, Hong Qin, and Yan Pang. Implicit neural representation with multi-scale sine activation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(26):21567–21575, 2026. doi: 10.1609/aaai.v40i26.39305.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

Ameya D. Jagtap, Kenji Kawaguchi, and George Em Karniadakis. Locally adaptive activation functions with slope recovery for deep and physics-informed neural networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 476(2239):20200334, 2020. doi: 10.1098/rspa.2020.0334. URL https://doi.org/10.1098/rspa.2020.0334.

Ameya D Jagtap, Yeonjong Shin, Kenji Kawaguchi, and George Em Karniadakis. Deep kronecker neural networks: A general framework for neural networks with adaptive activation functions. *Neurocomputing*, 468:165–180, 2022.

Dhananjaya Jayasundara and Vishal M. Patel. Implicit neural representations: A signal processing perspective. *arXiv preprint arXiv:2604.15047*, 2026. doi: 10.48550/arXiv.2604.15047. URL https://arxiv.org/abs/2604.15047.

Dhananjaya Jayasundara, Heng Zhao, Demetrio Labate, and Vishal M. Patel. MIRE: Matched implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8279–8288, 2025. doi: 10.1109/CVPR52734.2025.00775.

Amirhossein Kazerouni, Reza Azad, Alireza Hosseini, Dorit Merhof, and Ulas Bagci. Incode: Implicit neural conditioning with prior knowledge embeddings. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 1287–1296. IEEE, 2024.

Zhen Liu, Hao Zhu, Qi Zhang, Jingde Fu, Weibing Deng, Zhan Ma, Yanwen Guo, and Xun Cao. FINER: Flexible spectral-bias tuning in implicit NEural representation by variable-periodic activation functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2713–2722, June 2024. doi: 10.1109/CVPR52733.2024.00262. URL https://openaccess.thecvf.com/content/CVPR2024/html/Liu_FINER_Flexible_Spectral-bias_Tuning_in_Implicit_NEural_Representation_by_Variable-periodic_CVPR_2024_paper.html.

Christian H. X. Ali Mehmeti-Göpel, David Hartmann, and Michael Wand. Ringing ReLUs: Harmonic distortion analysis of nonlinear feedforward networks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=TaYhv-q1Xit.

Alireza Morsali, MohammadJavad Vaez, Mohammadhossein Soltani, Amirhossein Kazerouni, Babak Taati, and Morteza Mohammad-Noori. A unified theory of sinusoidal activation families for implicit neural representations. *arXiv preprint arXiv:2502.00869*, 2025.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 5301–5310. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/rahaman19a.html.

Reza Rezaeian, Moein Heidari, Reza Azad, Dorit Merhof, Hamid Soltanian-Zadeh, and Ilker Hacihaliloglu. SL2A-INR: Single-layer learnable activation for implicit neural representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 26065–26074, 2025.

Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G. Baraniuk. WIRE: Wavelet implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18507–18516, June 2023. doi: 10.1109/CVPR52729.2023.01775. URL https://openaccess.thecvf.com/content/CVPR2023/html/Saragadam_WIRE_Wavelet_Implicit_Neural_Representations_CVPR_2023_paper.html.

Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 7462–7473. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper/2020/hash/53c04118df112c13a8c34b38343b9c10-Abstract.html.

Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33:7537–7547, 2020.

Pandula Thennakoon, Avishka Ranasinghe, Mario De Silva, Buwaneka Epakanda, Roshan Godaliyadda, Mervyn Parakrama Bandara Ekanayake, and Vijitha Herath. COSMO-INR: Complex sinusoidal modulation for implicit neural representations. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=yGJrvSU6wK.

Gizem Yüce, Guillermo Ortiz-Jiménez, Beril Besbinar, and Pascal Frossard. A structured dictionary perspective on implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19206–19216, 2022. doi: 10.1109/CVPR52688.2022.01863.

A Appendix

A.1 Phase Folding and Cubic Stabilization

The nonlinear phase used by NASA is

$$\Phi(u) = (1+c)u + ku^2 + \frac{1+\eta}{3}k^2u^3, \tag{10}$$

where k controls the chirp rate, c controls the linear phase gain, and $\eta > 0$ is a fixed cubic stabilization parameter.

Phase folding in a purely quadratic map. To first illustrate the folding problem, consider the quadratic phase map

$$\Phi_{\text{quad}}(u) = (1+c)u + ku^2. \tag{11}$$

Its derivative is

$$\Phi'_{\text{quad}}(u) = (1+c) + 2ku. \tag{12}$$

A phase map is injective, and therefore non-folding, whenever its derivative remains strictly positive over its domain. On the normalized interval $u \in [-1, 1]$,

$$\min_{u\in[-1,1]} \Phi'_{\text{quad}}(u) = (1+c) - 2|k|. \tag{13}$$

Hence the quadratic phase is guaranteed to remain monotonic only if

$$\boxed{|k| < \frac{1+c}{2}.} \tag{14}$$

For the special case $c = 0$, this reduces to

$$|k| < \frac{1}{2}. \tag{15}$$

Once this bound is violated, the derivative vanishes at

$$u^\star = -\frac{1+c}{2k}, \tag{16}$$

provided that $u^\star$ lies inside the active support. Beyond this turning point, $\Phi'_{\text{quad}}(u)$ changes sign and the phase map is no longer one-to-one. We refer to this phenomenon as *phase folding*.

Quadratic–cubic stabilization. NASA augments the quadratic phase with the cubic term in Eq. equation 10. Differentiating gives

$$\boxed{\Phi'(u) = (1+c) + 2ku + (1+\eta)k^2u^2.} \tag{17}$$

For $k \neq 0$, Eq. equation 17 is a quadratic polynomial in u with positive leading coefficient $(1+\eta)k^2$. Its discriminant is

$$\Delta = (2k)^2 - 4(1+\eta)k^2(1+c) \tag{18}$$

$$= 4k^2\left[1 - (1+\eta)(1+c)\right] \tag{19}$$

$$= -4k^2\left[\eta + c(1+\eta)\right]. \tag{20}$$

For the NASA parameter range $c \geq 0$ and any $\eta > 0$,

$$\Delta < 0 \qquad (k \neq 0). \tag{21}$$

Since the leading coefficient is positive and the quadratic has no real roots,

$$\boxed{\Phi'(u) > 0, \qquad \forall u \in \mathbb{R}.} \tag{22}$$

Thus the quadratic–cubic phase is globally monotonic and cannot fold, irrespective of the magnitude of k.

The location of the minimum derivative follows from the vertex of the quadratic:

$$u^\star = -\frac{1}{(1+\eta)k}. \tag{23}$$

Substitution into Eq. equation 17 gives

$$\Phi'(u^\star) = (1+c) - \frac{1}{1+\eta} \tag{24}$$

$$= \boxed{\frac{\eta + c(1+\eta)}{1+\eta}}. \tag{25}$$

Consequently, even at the most conservative setting $c = 0$,

$$\boxed{\Phi'_{\min} = \frac{\eta}{1+\eta} > 0.} \tag{26}$$

For the fixed value $\eta = 0.05$ used in our experiments,

$$\Phi'_{\min} = \frac{0.05}{1.05} \approx 0.04762, \tag{27}$$

which bounds the phase slope strictly away from zero for every chirp rate k.

If $\eta = 0$ and $c = 0$, the derivative instead becomes

$$\Phi'(u) = 1 + 2ku + k^2u^2 = (1+ku)^2, \tag{28}$$

which never becomes negative but can touch zero at $u = -1/k$. The strictly positive choice $\eta > 0$ therefore removes this near-singular case and guarantees a finite positive minimum slope.

Hence, the cubic term does not remove the symmetry-breaking quadratic component ku^2; rather, it stabilizes the phase map by adding sufficient positive curvature to prevent the local phase slope from reaching or crossing zero.

A.2 Dense Chebyshev Representation of NASA

The NASA activation is

$$\rho(u) = w(u)\sin(\omega_0\Phi(u)), \tag{29}$$

where

$$w(u) = \frac{1}{2}\left[\tanh\left(\frac{u+T/2}{\tau}\right) - \tanh\left(\frac{u-T/2}{\tau}\right)\right], \tag{30}$$

and

$$\Phi(u) = (1+c)u + ku^2 + \frac{1+\eta}{3}k^2u^3. \tag{31}$$

We show that, for nonzero chirp rate k, NASA generically contains nonzero Chebyshev coefficients in both the even and odd parity subsequences, and that both subsequences are infinitely supported.

A.2.1 Parity decomposition of the phase

Write the phase as the sum of its odd and even components,

$$\Phi(u) = \Phi_{\mathrm{o}}(u) + \Phi_{\mathrm{e}}(u), \tag{32}$$

where

$$\Phi_{\mathrm{o}}(u) = (1+c)u + \frac{1+\eta}{3}k^2u^3, \tag{33}$$

$$\Phi_{\mathrm{e}}(u) = ku^2. \tag{34}$$

Clearly,

$$\Phi_{\mathrm{o}}(-u) = -\Phi_{\mathrm{o}}(u), \qquad \Phi_{\mathrm{e}}(-u) = \Phi_{\mathrm{e}}(u). \tag{35}$$

The tanh window is itself even. Indeed, using the oddness of $\tanh(\cdot)$,

$$w(-u) = w(u). \tag{36}$$

Applying

$$\sin(A + B) = \sin A \cos B + \cos A \sin B \tag{37}$$

to Eq. equation 29 gives

$$\begin{aligned}\rho(u) = w(u) \sin(\omega_0 \Phi_{\mathrm{o}}(u)) \cos(\omega_0 \Phi_{\mathrm{e}}(u)) \\ + w(u) \cos(\omega_0 \Phi_{\mathrm{o}}(u)) \sin(\omega_0 \Phi_{\mathrm{e}}(u)) \,.\end{aligned} \tag{38}$$

Define

$$\rho_{\mathrm{o}}(u) = w(u) \sin(\omega_0 \Phi_{\mathrm{o}}(u)) \cos(\omega_0 \Phi_{\mathrm{e}}(u)) \,, \tag{39}$$

$$\rho_{\mathrm{e}}(u) = w(u) \cos(\omega_0 \Phi_{\mathrm{o}}(u)) \sin(\omega_0 \Phi_{\mathrm{e}}(u)) \,. \tag{40}$$

Since w is even, $\sin(\omega_0 \Phi_{\mathrm{o}})$ is odd, $\cos(\omega_0 \Phi_{\mathrm{o}})$ is even, $\cos(\omega_0 \Phi_{\mathrm{e}})$ is even, and $\sin(\omega_0 \Phi_{\mathrm{e}})$ is even, we obtain

$$\rho_{\mathrm{o}}(-u) = -\rho_{\mathrm{o}}(u), \tag{41}$$

and

$$\rho_{\mathrm{e}}(-u) = \rho_{\mathrm{e}}(u). \tag{42}$$

Therefore

$$\boxed{\rho(u) = \rho_{\mathrm{o}}(u) + \rho_{\mathrm{e}}(u)} \tag{43}$$

contains both odd and even components whenever $k \neq 0$.

A.2.2 Connection to the Chebyshev basis

On $u \in [-1, 1]$, expand the activation in Chebyshev polynomials of the first kind:

$$\rho(u) = \sum_{n=0}^{\infty} a_n T_n(u), \qquad T_n(u) = \cos(n \arccos u) \,. \tag{44}$$

The coefficients are

$$a_n = \frac{2 - \delta_{n0}}{\pi} \int_0^{\pi} \rho(\cos\theta) \cos(n\theta) \, d\theta. \tag{45}$$

Chebyshev polynomials satisfy

$$T_n(-u) = (-1)^n T_n(u). \tag{46}$$

Hence T_{2m} is even and T_{2m+1} is odd. Consequently,

$$\rho_{\mathrm{e}}(u) = \sum_{m=0}^{\infty} a_{2m} T_{2m}(u), \tag{47}$$

$$\rho_{\mathrm{o}}(u) = \sum_{m=0}^{\infty} a_{2m+1} T_{2m+1}(u). \tag{48}$$

For a purely odd sinusoidal activation, all even coefficients vanish:

$$a_{2m} = 0. \tag{49}$$

NASA does not obey this constraint because $\rho_{\mathrm{e}} \not\equiv 0$ whenever $k \neq 0$. Therefore there exists nonzero energy in the even Chebyshev subsequence:

$$\boxed{E_{\mathrm{even}} = \sum_{n \text{ even}} |a_n|^2 > 0.} \tag{50}$$

Likewise, because $\rho_{\mathrm{o}} \not\equiv 0$,

$$\boxed{E_{\mathrm{odd}} = \sum_{n \text{ odd}} |a_n|^2 > 0.} \tag{51}$$

A.2.3 WHY THE REPRESENTATION IS DENSE

A finite Chebyshev expansion

$$\sum_{n=0}^{N} a_n T_n(u) \tag{52}$$

is an ordinary polynomial of degree at most N. Conversely, any non-polynomial analytic function must have infinitely many nonzero Chebyshev coefficients.

For finite T and $\tau > 0$, the window $w(u)$ is analytic on the real interval $[-1, 1]$. Furthermore,

$$\sin\left[\omega_0 \left((1+c)u + ku^2 + \frac{1+\eta}{3}k^2u^3\right)\right] \tag{53}$$

is a non-polynomial analytic function for nonzero ω_0. Therefore NASA is non-polynomial and its Chebyshev expansion cannot terminate at finite order.

More strongly, for $k \neq 0$, both parity components in Eq. equation 38 are nonzero, analytic, and non-polynomial under nondegenerate parameter choices. Hence neither can possess a finite Chebyshev expansion. It follows that

$$\boxed{\#\{m : a_{2m} \neq 0\} = \infty, \qquad \#\{m : a_{2m+1} \neq 0\} = \infty.} \tag{54}$$

Thus NASA possesses infinitely supported Chebyshev components in both parity classes. This is the sense in which we refer to the NASA Chebyshev representation as *dense*: symmetry no longer forces either the even or odd coefficient subsequence to vanish.

This result should not be interpreted as requiring $a_n \neq 0$ for every individual integer n; isolated coefficients may vanish because of parameter-dependent cancellation. Rather, the activation is generically non-sparse across both parity classes.

A.2.4 SMALL-CHIRP INTERPRETATION

The emergence of the even coefficients can also be seen directly for small $|k|$. Since

$$\Phi_{\mathrm{e}}(u) = ku^2, \tag{55}$$

we have

$$\sin(\omega_0 k u^2) = \omega_0 k u^2 + \mathcal{O}(k^3). \tag{56}$$

Therefore the even component satisfies

$$\rho_{\mathrm{e}}(u) = \omega_0 k\, w(u) u^2 \cos(\omega_0(1+c)u) + \mathcal{O}(k^3). \tag{57}$$

The leading perturbation is explicitly even in u and is linear in the chirp rate k. Hence arbitrarily small nonzero chirp immediately injects even-parity Chebyshev content into the otherwise predominantly odd sinusoidal carrier.

For numerical characterization, the parity energies may be summarized by

$$R_{\mathrm{even}} = \frac{E_{\mathrm{even}}}{E_{\mathrm{even}} + E_{\mathrm{odd}} + \varepsilon}. \tag{58}$$

A purely odd sinusoidal activation has $R_{\mathrm{even}} \approx 0$, whereas a nonzero value directly measures the parity-breaking Chebyshev content introduced by the nonlinear phase.

A.3 NEURAL TANGENT KERNEL OF NASA

Consider a scalar-output, one-hidden-layer coordinate network

$$f_\theta(\mathbf{x}) = \frac{1}{\sqrt{n}} \sum_{r=1}^{n} a_r \rho(g_r(\mathbf{x})), \tag{59}$$

where

$$g_r(\mathbf{x}) = \mathbf{w}_r^\top \mathbf{x} + b_r, \tag{60}$$

and $\rho(\cdot)$ is the NASA activation.

The neural tangent kernel between two coordinates $\mathbf{x}_i$ and $\mathbf{x}_j$ is defined as

$$\boxed{K(\mathbf{x}_i, \mathbf{x}_j) = \langle \nabla_\theta f_\theta(\mathbf{x}_i), \nabla_\theta f_\theta(\mathbf{x}_j) \rangle\,.} \tag{61}$$

A.3.1 NASA DERIVATIVE

The activation is

$$\rho(u) = w(u)\sin(\omega_0\Phi(u))\,, \tag{62}$$

with

$$w(u) = \frac{1}{2}\left[\tanh\left(\frac{u+T/2}{\tau}\right) - \tanh\left(\frac{u-T/2}{\tau}\right)\right], \tag{63}$$

and

$$\Phi(u) = (1+c)u + ku^2 + \frac{1+\eta}{3}k^2u^3. \tag{64}$$

The window derivative is

$$\boxed{w'(u) = \frac{1}{2\tau}\left[\operatorname{sech}^2\left(\frac{u+T/2}{\tau}\right) - \operatorname{sech}^2\left(\frac{u-T/2}{\tau}\right)\right].} \tag{65}$$

The phase derivative is

$$\boxed{\Phi'(u) = (1+c) + 2ku + (1+\eta)k^2u^2.} \tag{66}$$

Using the product and chain rules,

$$\begin{aligned}\rho'(u) = {} & w'(u)\sin(\omega_0\Phi(u)) \\ & + \omega_0 w(u)\Phi'(u)\cos(\omega_0\Phi(u))\,.\end{aligned} \tag{67}$$

For compactness, define

$$\boxed{\Psi(u) = w'(u)\sin(\omega_0\Phi(u)) + \omega_0 w(u)\Phi'(u)\cos(\omega_0\Phi(u))\,.} \tag{68}$$

Thus

$$\rho'(u) = \Psi(u). \tag{69}$$

A.3.2 PARAMETER GRADIENTS

Let

$$g_{ri} = g_r(\mathbf{x}_i), \qquad g_{rj} = g_r(\mathbf{x}_j). \tag{70}$$

The gradient with respect to the output coefficient is

$$\frac{\partial f_\theta(\mathbf{x}_i)}{\partial a_r} = \frac{1}{\sqrt{n}}\rho(g_{ri}). \tag{71}$$

The gradient with respect to the hidden weight vector is

$$\frac{\partial f_\theta(\mathbf{x}_i)}{\partial \mathbf{w}_r} = \frac{1}{\sqrt{n}}a_r\Psi(g_{ri})\mathbf{x}_i. \tag{72}$$

The gradient with respect to the hidden bias is

$$\frac{\partial f_\theta(\mathbf{x}_i)}{\partial b_r} = \frac{1}{\sqrt{n}}a_r\Psi(g_{ri}). \tag{73}$$

Therefore,

$$\begin{aligned}&\left\langle \frac{\partial f(\mathbf{x}_i)}{\partial \mathbf{w}_r}, \frac{\partial f(\mathbf{x}_j)}{\partial \mathbf{w}_r} \right\rangle + \frac{\partial f(\mathbf{x}_i)}{\partial b_r}\frac{\partial f(\mathbf{x}_j)}{\partial b_r} \\ &\quad = \frac{1}{n}a_r^2\left(\mathbf{x}_i^\top\mathbf{x}_j + 1\right)\Psi(g_{ri})\Psi(g_{rj}).\end{aligned} \tag{74}$$

The output-weight contribution is

$$\frac{1}{n}\rho(g_{ri})\rho(g_{rj}). \tag{75}$$

Consequently, the complete finite-width NTK of the one-hidden-layer network is

$$\boxed{K_{\mathrm{BC}}(\mathbf{x}_i,\mathbf{x}_j)=\frac{1}{n}\sum_{r=1}^{n}\Big[\rho(g_{ri})\rho(g_{rj}) + a_r^2\left(\mathbf{x}_i^\top\mathbf{x}_j+1\right)\Psi(g_{ri})\Psi(g_{rj})\Big].} \tag{76}$$

Taking an expectation over random initialization gives

$$\boxed{\mathcal{K}_{\mathrm{BC}}(\mathbf{x}_i,\mathbf{x}_j)=\mathbb{E}_\theta\left[\frac{1}{n}\sum_{r=1}^{n}\Big\{\rho(g_{ri})\rho(g_{rj}) + a_r^2(\mathbf{x}_i^\top\mathbf{x}_j+1)\Psi(g_{ri})\Psi(g_{rj})\Big\}.\right.} \tag{77}$$

If only the hidden affine parameters $\{\mathbf{w}_r, b_r\}$ are included in the kernel, as in a FINER-style hidden-layer NTK expression, the first term is omitted and the kernel reduces to

$$\boxed{\mathcal{K}_{\mathrm{BC}}^{(w,b)}(\mathbf{x}_i,\mathbf{x}_j)=(\mathbf{x}_i^\top\mathbf{x}_j+1)\,\mathbb{E}_\theta\left[\frac{1}{n}\sum_{r=1}^{n}a_r^2\Psi(g_{ri})\Psi(g_{rj})\right].} \tag{78}$$

Substituting Eq. equation 68 explicitly gives

$$\begin{aligned}\mathcal{K}_{\mathrm{BC}}^{(w,b)}(\mathbf{x}_i,\mathbf{x}_j)=(\mathbf{x}_i^\top\mathbf{x}_j+1)\mathbb{E}_\theta\Bigg[\frac{1}{n}\sum_{r=1}^{n}a_r^2 &\times\Big\{w'(g_{ri})\sin[\omega_0\Phi(g_{ri})]+\omega_0 w(g_{ri})\Phi'(g_{ri})\cos[\omega_0\Phi(g_{ri})]\Big\}\\ &\times\Big\{w'(g_{rj})\sin[\omega_0\Phi(g_{rj})]+\omega_0 w(g_{rj})\Phi'(g_{rj})\cos[\omega_0\Phi(g_{rj})]\Big\}\Bigg].\end{aligned} \tag{79}$$

Equation equation 79 shows that the NASA tangent kernel contains four interactions: window–window, window–phase, phase–window, and phase–phase contributions.

A.4 Additional results

Figure 5 presents the results of the 1/6× downsampled for the super-resolution task. From this figure, it is evident that when compared with methods such as SIREN, WIRE, FINER , their adaptive versions, and NOPE-added versions have been surpased by NASA.

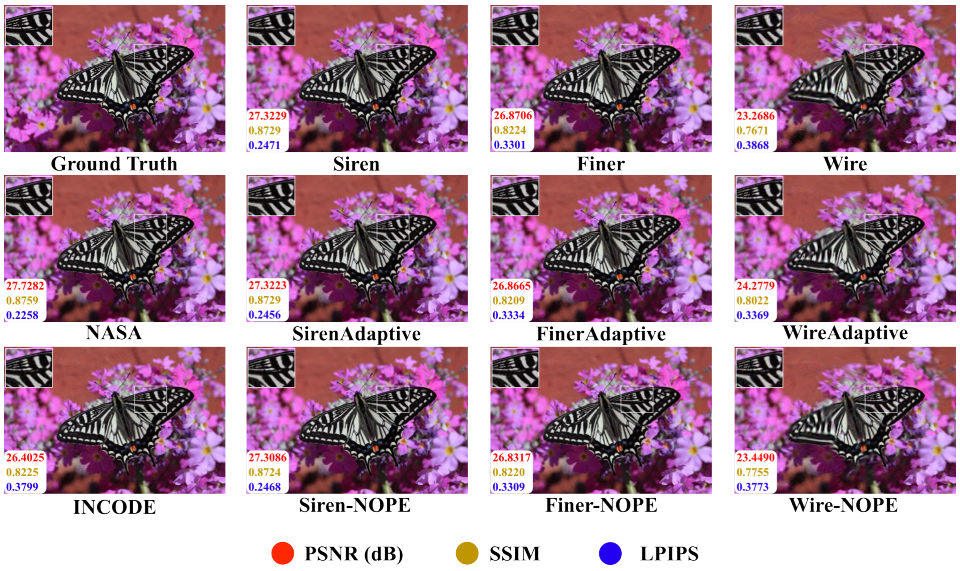


Figure 7: Super resolution task compared with SOTA methods.