Spectral Gating via Damped Oscillations for Adaptive Implicit Neural Representations

Alex Costanzino[1], Pierluigi Zama Ramirez[⋆2], Giuseppe Lisanti[1], and Luigi Di Stefano[1]

[1] CVLab, University of Bologna
[2] Ca' Foscari University of Venice

Abstract. Implicit Neural Representations (INRs) have been proven successful in encoding continuous signals through coordinate-based networks, yet facing a spectral dilemma: periodic activations capture fine details but act as all-pass filters that memorise noise, while spatially compact activations regularise effectively but suffer from low-frequency bias. Existing attempts to resolve this trade-off introduce computational overhead or tuning frailty. We propose to model each neuron's activation as the steady-state response of a sinusoidally-forced damped harmonic oscillator, whose amplitude naturally governs the network's spectral selectivity during training. By jointly optimising the oscillator parameters alongside the network weights, our method adapts to the target signal's spectral content without explicit regularisation. Initialised in the stop-band, the network exhibits a coarse-to-fine learning curriculum that progressively expands its spectral gate, capturing low-frequency structures first and high-frequency details only when justified by the reconstruction objective. Comprehensive experiments show that our approach consistently achieves state-of-the-art or competitive results against established INRs, while requiring no task-specific tuning of any hyperparameters. Project Page available at `https://alex-costanzino.github.io/fdho/`.

Keywords: Implicit Neural Representations · Spectral Gating

1 Introduction

Implicit Neural Representations (INRs) have emerged as a powerful paradigm for encoding continuous signals, such as audio, images, shapes, and physical fields, into the weights of coordinate-based neural networks [13, 16, 21]. Their appeal lies in being resolution-independent, differentiable, and compact, making them natural fits for tasks ranging from novel view synthesis [4] to solving partial differential equations [21]. At the heart of any INR lies the choice of the activation function, which governs the network's spectral expressivity, *i.e.*, its ability to faithfully represent both smooth global structure and fine-grained local details. This choice confronts practitioners with what we term the spectral dilemma. Periodic activations, exemplified by SIREN [21], endow the network

⋆ Partially supported by the IRIDE-B program, Ca' Foscari University of Venice.

with the capacity to represent arbitrarily high frequencies. However, their infinite bandwidth makes them susceptible to noise memorisation, treating random perturbations as legitimate signal content. Non-periodic alternatives, such as Gaussian activations [17], provide spatial compactness and inherent regularisation but introduce a low-frequency bias that limits their ability to capture oscillatory or fine-grained content. Intermediate approaches exist – Gabor wavelets in WIRE [18] offer joint space-frequency localisation, FINER [12] introduces variable-periodic functions, and Fourier Reparametrisation [20] addresses spectral bias through weight decomposition – but each comes with its own trade-offs in computational cost, sensitivity to hyperparameters, or restricted representational capacity. More recently, MIRE [9] proposes learning per-layer activations selected from a dictionary of predefined functions, and IGA-INR [19] adjusts gradient dynamics through inductive gradient correction, both methods highlighting that the interaction between activation design and training dynamics is central to INR performance. Despite this progress, no existing method provides a mechanism to adapt the bandwidth robustly with respect to noise.

Inspired by classical mechanics, we model each neuron as the steady-state response of a sinusoidally-forced damped harmonic oscillator. The resulting activation takes the form of an amplitude-modulated sinusoidal carrier. Such modulation depends on three physical quantities: the forcing frequency ω, the natural frequency ω_n, and the damping factor ξ. In particular, the amplitude is the transfer function of a second-order linear system, a well-characterised band-pass filter whose centre frequency, bandwidth, and roll-off are determined by ω_n and ξ. We refer to this amplitude-modulated bandwidth control as *spectral gating*: the oscillator parameters determine a passband, and only frequency content from the forcing within this band receives non-negligible amplitude gain. When these parameters are made learnable, the network acquires the ability to adapt its bandwidth during training. In particular, with our formulation, the optimisation cost to admit higher frequencies is lower for coherent signal content than for random noise (Sec. 3.3, Corollary 1). Moreover, a coarse-to-fine learning curriculum emerges as we initialise the oscillator parameters in the stopband, so that the network begins with a narrow bandwidth. As training progresses and the loss landscape demands richer frequency content, the bandwidth widens, producing a natural coarse-to-fine learning progression without explicit scheduling (Sec. 3.4). Finally, the damping factor ξ provides a continuous, learnable axis connecting distinct spectral regimes: as $\xi \to 0$, the activation recovers the unregulated bandwidth of a pure SIREN-style sinusoid, as it increases, the function smoothly transitions toward a more spatially compact representation. Hence, rather than committing to a fixed point in the INR's design space, the network can learn where to operate along this axis.

We validate these claims through extensive experiments and comparisons with established INRs, achieving state-of-the-art or competitive results across all evaluated settings, while requiring no task-specific tuning of any hyperparameters.

2 Related Work

Activation Design for INRs. The choice of activation function is central to INR performance. ReLU-based networks are known to suffer from spectral bias toward low frequencies [23]. SIREN [21] addressed this with sinusoidal activations, enabling high-frequency representation but acting as an all-pass filter with no intrinsic spatial localisation or frequency regularisation, which leaves it prone to noise memorisation and sensitive to initialisation. These limitations motivated a line of activation-level remedies: Gaussian activations [17] provide smoothness and an implicit low-frequency bias, while WIRE [18] combines spatial and frequency localisation through Gabor wavelets. FINER [12] and MIRE [9] instead retain the sinusoidal form but add per-layer spectral control through variable-periodic activations and learned function selection, respectively. These methods address the spectral dilemma through activation shape, but none provide a formal mechanism explaining why certain frequencies are favoured or suppressed during training. Our approach differs by deriving the activation from a physical system whose transfer function yields provable spectral gating properties.

Spectral Bias and Training Dynamics. A parallel line of work addresses spectral limitations not through the activation itself but through the training procedure or input encoding. Fourier feature mappings [23] lift coordinates into a high-dimensional sinusoidal basis, enabling ReLU-based networks to learn high frequencies. Their NTK [8] analysis shows that the feature bandwidth directly controls the kernel's spectral support. Multiplicative filter networks [7] achieve frequency control through element-wise products of Gabor filters across layers. BACON [11] enforces a band-limited hierarchy via additive multi-scale outputs. TUNER [14] instead keeps the sinusoidal activation fixed and controls the spectrum at initialisation, bounding the network's frequency support to stabilise training and prevent overfitting. Fourier Reparameterisation [20] decomposes network weights in the frequency domain to mitigate spectral bias, and IGA-INR [19] corrects gradient dynamics through inductive gradient adjustment. Our work is complementary: rather than modifying the input encoding, weight parameterisation, or gradient update rule, we embed spectral control directly into the activation's differentiable structure, where it arises as an implicit consequence of optimisation.

Theoretical analysis of INRs. Theoretical understanding of INRs remains limited: [23] analyses the NTK of Fourier features to explain their accelerated convergence; [25] shows that SIREN layers compose via Bessel functions; [14] develops an amplitude–phase expansion of sinusoidal MLPs, characterising how layer composition generates new frequencies and bounding the network's representable spectrum; [17] derives Lipschitz bounds and NTK eigenspectra for Gaussian activations; [18] analyses Gabor wavelet support to motivate WIRE's frequency-space trade-off; and [21] establishes the distribution of activations and gradients through depth for SIREN. However, these analyses characterise representational properties – what each activation can express – rather than optimisation dynamics – how gradient descent navigates the parameter space. In particular, while TUNER governs the spectrum statically by bounding frequency support at initialisation, none of these works provides a gradient-level mechanism for noise rejection or adaptive bandwidth control. Conversely, our contribution features an explicit gating condition that allows us to reject noise while fitting genuine signal components as training progresses.

3 Formulation

We derive our activation function from the steady-state analysis of a forced damped harmonic oscillator (Sec. 3.1), establish its formal properties (Sec. 3.2, proofs in Sec. A.1), analyse how its learnable parameters create implicit spectral regularisation through gradient dynamics (Sec. 3.3, proofs in Sec. A.2), characterise the multi-layer frequency composition via Bessel-function expansions (Sec. 3.4 and Sec. A.3), and discuss its neural tangent kernel structure (Sec. A.4).

3.1 Activation Definition

Consider a damped harmonic oscillator with natural frequency ω_n and damping factor ξ, driven by a sinusoidal forcing at frequency ω. Its motion equation is $\ddot{x} + 2\xi\omega_n\dot{x} + \omega_n^2 x = \omega_n^2 \sin(\omega t)$, and the steady-state solution has the form $x_s(t) = A(\omega, \omega_n, \xi)\sin(\omega t + \varphi)$, with the amplitude given by:

$$A(\omega, \omega_n, \xi) = \frac{\omega_n^2}{\sqrt{(\omega_n^2 - \omega^2)^2 + (2\xi\omega_n\omega)^2}}. \tag{1}$$

Definition 1. *Given the set of learnable parameters $\theta = (\omega, \omega_n, \xi, \varphi)$, with $\omega; \omega_n; \xi > 0$, our proposed activation function, $\sigma\colon \mathbb{R} \to \mathbb{R}$, is defined as:*

$$\sigma(z; \theta) = A(\omega, \omega_n, \xi)\sin(\omega z + \varphi). \tag{2}$$

Each hidden layer uses a parameter set θ shared across all neurons. Hence, the activation is applied element-wise to the input $z = \mathbf{w}^\top\mathbf{x} + b$, yielding $\sigma(\mathbf{W}\mathbf{x} + \mathbf{b}; \theta)$. Parameters θ and weights $(\mathbf{W}, \mathbf{b})$ are jointly optimised via gradient descent.

It is worth highlighting that for a given set of layer parameters, the amplitude gain A is a scalar constant, not a spatially varying envelope. Hence, the network does not filter different frequency components at inference time; rather, the landscape of $A(\omega, \omega_n, \xi)$ shapes the gradient flow during training, determining how readily the network can amplify or suppress content at different frequencies as the optimisation evolves. The inductive bias, therefore, operates through optimisation dynamics, not through the functional form of a single forward pass.

One might consider achieving similar adaptivity by simply making the amplitude, frequency, and phase of SIREN's sinusoidal activation independently learnable. However, when these quantities are decoupled, nothing prevents the network from assigning a large amplitude to any frequency, including those driven by noise. The oscillator formulation is fundamentally different, as the amplitude $A(\omega, \omega_n, \xi)$ is not a free parameter but is constrained by the relationship between forcing frequency, natural frequency, and damping factor. This physical coupling ensures that the optimisation cost to admit higher frequencies is lower for coherent signal content than for random noise, as we formalise in Sec. 3.3, and that the learned parameters converge to stable equilibria rather than drifting indefinitely, as we verify empirically by comparing our proposal to learnable SIREN activations in Sec. C.1.

3.2 Basic Properties

Before analysing the training dynamics of our proposed activation, dubbed *Forced Damped Harmonic Oscillator* (FDHO), we establish three properties: smoothness, Lipschitz continuity, and output boundedness. Smoothness in both z and θ guarantees that all parameters can be jointly optimised via gradient-based methods; Lipschitz continuity bounds the magnitudes of gradients flowing through the network during backpropagation; and the output bound quantifies how the parameters control the activation's dynamic range. All proofs are in Sec. A.1.

Proposition 1 (Smoothness). *For $\omega_n > 0$ and $\xi > 0$, the activation function $\sigma(z; \theta)$ is C^∞ with respect to both z and all components of θ.*

Remark 1. At $\xi = 0$ the denominator of $A(\omega, \omega_n, \xi)$ vanishes when $\omega = \omega_n$ (undamped resonance). In practice, we constrain $\xi \in (0, 1)$ via sigmoid reparameterisation, which is physically motivated as all real oscillatory systems exhibit non-zero damping.

Proposition 2 (Lipschitz Bound). *For fixed θ, the activation is Lipschitz in z with a constant $L_z = |\omega| \, A(\omega, \omega_n, \xi)$.*

Proposition 3 (Output Bound). *$|\sigma(z; \theta)| \leq A(\omega, \omega_n, \xi)$ uniformly in z for all $z \in \mathbb{R}$. The peak amplitude $A_{peak} = \left(2\xi\sqrt{1 - \xi^2}\right)^{-1}$ is attained at the resonant frequency $\omega_r = \omega_n\sqrt{1 - 2\xi^2}$ for $\xi < 1/\sqrt{2}$. For $\xi \geq 1/\sqrt{2}$ the amplitude decreases monotonically from $A(0) = 1$.*

Proposition 2 reveals a key structural property: representing high-frequency content requires *both* a large forcing frequency ω *and* a non-negligible amplitude A at that frequency, since the Lipschitz constant is their product. This multiplicative coupling is the source of the implicit regularisation that we formalise next.

3.3 Spectral Gating

The properties above confirm that our activation is well-behaved for optimisation; however, they do not yet explain why it should favour signal over noise. We now turn to the core theoretical question: *how do the gradients of the reconstruction loss with respect to the parameters behave during training?* By decomposing the gradient with respect to the damping factor ξ, following the principle of analysing implicit biases of gradient descent [3, 10], we obtain a particularly interpretable form of spectral regularisation, requiring no explicit penalty in the loss. The gradient naturally separates into: (i) a term that pushes ξ upward, narrowing the bandwidth unconditionally; (ii) a term that pushes ξ downward, widening the bandwidth, but only when the target contains coherent energy at the carrier frequency. We refer to these competing dynamics as *closing* and *opening* the spectral gate, respectively. Theorems and corollaries are proven in Sec. A.2.

Theorem 1 (Amplitude is Monotone in ξ). *For $\omega; \omega_n; \xi > 0$ then $\frac{\partial A}{\partial \xi} < 0$.*

Theorem 1 establishes that increasing ξ always decreases the amplitude. The only route to amplifying the activation's output, thereby increasing the network's representational capacity at the carrier frequency, is to decrease ξ. Since gradient descent updates $\xi \leftarrow \xi - \eta \frac{\partial \mathcal{L}}{\partial \xi}$, the spectral gate opens only when $\frac{\partial \mathcal{L}}{\partial \xi} > 0$. To characterise when this condition holds, let us consider the mean squared error loss $\mathcal{L} = \frac{1}{N} \sum_{i=1}^{N} \big(\sigma(z_i; \theta) - y_i\big)^2$ for a single neuron with target values $\{y_i\}$.

Theorem 2 (Decomposition of the Damping Gradient). *The gradient of $\mathcal{L}$ with respect to ξ decomposes as:*

$$\frac{\partial \mathcal{L}}{\partial \xi} = \underbrace{\frac{\partial A}{\partial \xi} \frac{2A}{N} \sum_i \sin^2(\omega z_i + \varphi)}_{T_{\text{self}}} + \underbrace{\left(-\frac{\partial A}{\partial \xi}\right) \frac{2}{N} \sum_i y_i \sin(\omega z_i + \varphi)}_{T_{\text{signal}}}. \tag{3}$$

The self-interaction term satisfies $T_{\text{self}} = \frac{\partial A}{\partial \xi} \cdot \frac{2A}{N} \sum_i \sin^2(\omega z_i + \varphi)$. For dense, uniform sampling[3] of z, the sum converges to $\frac{1}{2}$, yielding $T_{\text{self}} \approx \frac{\partial A}{\partial \xi} A < 0$ by Theorem 1. This term always pushes ξ to increase, acting as a regularising force that tends to close the spectral gate. The signal term satisfies $T_{\text{signal}} = -2 \frac{\partial A}{\partial \xi} \langle y, \sin(\omega \cdot + \varphi) \rangle_N$, where $\langle \cdot, \cdot \rangle_N$ denotes the empirical inner product. Since $-\frac{\partial A}{\partial \xi} > 0$ by Theorem 1, we have $\operatorname{sign}(T_{\text{signal}}) = \operatorname{sign}\big(\langle y, \sin(\omega \cdot + \varphi) \rangle_N\big)$. That is, T_{signal} is positive precisely when the target has energy at the carrier frequency.

Corollary 1 (Spectral Gating Conditions). *The spectral gate opens ($\frac{\partial \mathcal{L}}{\partial \xi} > 0$) if and only if $T_{\text{signal}} > |T_{\text{self}}|$, which requires:*

$$\langle y, \sin(\omega \cdot + \varphi) \rangle_N > \frac{A}{2}. \tag{4}$$

Three cases arise:

(a) ***Coherent signal.*** *If the target contains a Fourier component $A_y \sin(\omega z + \psi)$ at the carrier frequency ω, then $\langle y, \sin(\omega \cdot + \varphi) \rangle_N \approx \frac{A_y}{2} \cos(\varphi - \psi)$. The gate opens when $A_y \cos(\varphi - \psi) > A$,* i.e., *when the neuron is underfitting the corresponding frequency component. As A increases toward A_y, the inequality saturates, and ξ reaches equilibrium;*

[3] We analyse the population gradient obtained in the limit $N \to \infty$, where empirical inner products converge to their expectations.

(b) ***Random noise.*** *If y is i.i.d. with zero mean and variance σ_y^2, then $\langle y, \sin(\omega \cdot + \varphi)\rangle_N = O_p(N^{-1/2})$ by the central limit theorem, while $A/2 = O(1)$. For any reasonably sized sample, $|T_{\text{self}}|$ dominates, and the gate remains closed;*

(c) ***Mismatched frequency.*** *If $y = A' \sin(\omega' z)$ with $\omega' \neq \omega$, the inner product $\langle y, \sin(\omega \cdot + \varphi)\rangle_N \to 0$ as $N \to \infty$ by the orthogonality of sinusoids at distinct frequencies. The self-term dominates, correctly signalling that this neuron should not fit the mismatched component.*

Corollary 1, therefore, establishes that the optimisation problem is inherently biased to fit the coherent components of the signal and reject random noise.

Remark 2 (Multi-neuron layers). In a full layer with weight vectors $\{\mathbf{w}_j\}$ and biases $\{b_j\}$, each neuron receives $z_j = \mathbf{w}_j^\top \mathbf{x} + b_j$, creating different effective frequencies $\omega \|\mathbf{w}_j\|$ along different directions. The oscillator parameters $(\omega, \omega_n, \xi, \varphi)$, shared across the layer, control the envelope of representable frequencies, while the weight distribution determines how individual neurons tile the frequency space within that envelope.

3.4 Multi-Layer Frequency Composition

Cascading single-neuron oscillator layers produce a frequency-modulated signal [6] whose harmonic content is governed by the Jacobi-Anger expansion [2,24]. The effective modulation index at layer ℓ is $\beta_\ell = \omega_\ell w_\ell \prod_{k=1}^{\ell-1} A_k$, which controls the number of non-negligible harmonics, which are approximately $\lfloor \beta \rfloor + 1$, since $|J_n(\beta)| \ll 1$ for $|n| > \beta + O(\beta^{1/3})$ [1,15]. Since each A_k decreases with ξ_k (Theorem 1), the damping factors collectively gate the network's harmonic bandwidth: larger ξ_k narrows the spectrum, while smaller ξ_k widens it. The product structure ensures geometric decay of β_ℓ at initialisation ($\beta_\ell \propto A_0^{\ell-1}$ with $A_0 < 1$), so that the deepest layers initially contribute near-zero frequency support and bandwidth expands progressively as the individual A_k grow during training, producing a *coarse-to-fine learning curriculum* without explicit scheduling. This extends the observation by [25] that composing sinusoidal layers in SIREN produces Bessel-governed harmonics. The key difference is that our oscillator amplitude A_k makes the modulation index learnable. Full formulation and derivations are provided in Sec. A.3 of the Supplemental Material.

4 Experiments

We evaluate our method across several tasks organised by observation model: signal representation (Sec. 4.1), recovery from corrupted observations (Sec. 4.2), and recovery from incomplete observations (Sec. 4.3). More tasks are reported in Sec. D of the Supplemental Materials. We compare against eight established INRs: SIREN [21], Gauss [17], WIRE [18], FINER [12], BACON [11], MFN [7], FR [20], and Fourier Features [23]. For each competitor, we adopt the task-specific hyperparameters and architecture configuration recommended by the original authors when available. When no configuration is provided for a given task, we manually tune the architecture and hyperparameters to obtain competitive results. In contrast, our method uses the same architecture and hyperparameters across all tasks, with no task-specific tuning, directly verifying the claim that the spectral gate adapts to the signal's complexity without manual intervention. All these details are in the Supplemental Materials.

To ensure a fair and reproducible comparison, we have re-implemented all methods and tasks within a unified codebase under identical training conditions. We will publicly release this library, dubbed `INRs Playground`, as part of our contributions. We run each experiment 10 times and average the results. For image-related tasks, we provide qualitative results of their best run in Fig. 1.

Initialisation We initialise ω_0 and $\omega_{n,0}$ such that the forcing frequency lies below the natural frequency (*e.g.*, $\omega_0 = 40, \omega_{n,0} = 45$), placing the system away from resonance with a suppressed initial amplitude. Both frequencies are reparameterised via softplus to ensure positivity. We initialise the damping factor to $\xi_0 = 1/\sqrt{2}$, corresponding to a Butterworth (maximally-flat) magnitude response [5], which provides moderate amplitude without a resonance peak. We initialise φ to the physical phase lag of the oscillator at the initial parameters: $\varphi_0 = -\arctan\big(\frac{2\xi_0\omega_{n,0}\omega_0}{\omega_{n,0}^2-\omega_0^2}\big)$. This ensures physical consistency at the start of training; the phase is then free to evolve. Finally, we follow the SIREN initialisation scheme [21] but scaled by the initial natural frequency: the first layer composed of m neurons draws $w \sim \mathcal{U}(-1/m, 1/m)$ and hidden layers draw $w \sim \mathcal{U}\big(-\sqrt{6/m}/\omega_n, \sqrt{6/m}/\omega_n\big)$, which maintains approximate unit variance of pre-activations across layers.

4.1 Signal Representation

We first evaluate the activation's ability to faithfully encode continuous signals under direct, complete supervision, *i.e.*, fitting a function $y = f(\mathbf{x})$ from dense samples. This setting isolates expressivity: every sample point is observed without noise or missing data, so performance differences between activations reflect their capacity to represent diverse spectral content. In terms of our theoretical framework, these experiments validate the properties established in Sec. 3.2. The network must have sufficient spectral bandwidth to match the target and provide baseline conditions under which the spectral gate should open during training.

1D Signal Fitting. We fit a one-dimensional composite signal containing multiple frequency components. This controlled setting allows direct inspection of how the learned FDHO parameters align with the target's spectral content, providing the most transparent validation of the spectral gating mechanism described in Sec. 3.3. Results are reported in Tab. 1. We evaluate on two signal types: a **Square Wave** (rich in odd harmonics) and a **Chirp** (continuously sweeping frequency), at increasing fundamental frequencies, and report both final and peak PSNR over training to distinguish converged reconstruction quality from the best performance achieved at any point during optimisation.

Table 1: Signal Fitting. Best results in **bold**, runner-ups underlined.

INR	Square Wave						Chirp					
	100 Hz		250 Hz		500 Hz		250 Hz		500 Hz		1000 Hz	
	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR
FDHO	**150.76 ± 1.51**	**160.40 ± 7.32**	**135.95 ± 1.39**	**142.70 ± 1.22**	**126.83 ± 1.17**	**135.38 ± 0.69**	**117.69 ± 15.63**	120.97 ± 17.13	**122.38 ± 16.85**	**124.98 ± 18.65**	**86.39 ± 4.10**	94.15 ± 2.60
SIREN [21]	116.68 ± 1.46	148.39 ± 2.32	74.98 ± 3.06	136.40 ± 0.33	64.49 ± 39.68	125.86 ± 0.80	70.43 ± 9.88	126.55 ± 2.20	57.65 ± 34.99	120.78 ± 3.15	61.05 ± 6.62	105.05 ± 3.27
Gauss [17]	123.83 ± 5.02	132.48 ± 0.16	91.60 ± 42.04	125.54 ± 0.07	26.63 ± 8.40	42.08 ± 4.32	104.32 ± 16.93	113.31 ± 15.97	49.10 ± 14.06	66.77 ± 4.61	16.70 ± 1.65	17.96 ± 1.61
WIRE [18]	95.64 ± 27.03	140.20 ± 1.98	54.83 ± 8.10	88.38 ± 14.79	34.41 ± 6.80	45.35 ± 3.07	96.81 ± 21.76	132.48 ± 3.72	51.09 ± 10.94	74.48 ± 5.82	30.92 ± 7.67	34.94 ± 6.57
BACON [11]	29.03 ± 0.00	29.03 ± 0.00	33.01 ± 0.00	33.01 ± 0.00	36.02 ± 0.00	36.02 ± 0.00	75.03 ± 2.71	107.59 ± 1.87	79.23 ± 0.73	109.08 ± 3.24	73.28 ± 0.96	78.69 ± 2.56
FINER [12]	107.64 ± 25.40	107.64 ± 25.40	101.85 ± 10.61	132.31 ± 0.81	53.70 ± 33.03	119.07 ± 0.07	82.68 ± 22.96	136.74 ± 1.10	65.76 ± 6.98	123.49 ± 0.63	80.50 ± 30.38	**112.57 ± 1.36**
MFN [7]	94.72 ± 13.13	126.92 ± 8.56	52.23 ± 4.76	63.77 ± 2.22	6.03 ± 0.00	6.04 ± 0.00	56.23 ± 4.40	81.54 ± 2.57	14.60 ± 0.21	14.62 ± 0.22	10.90 ± 0.14	10.91 ± 0.12
Fourier [23]	150.33 ± 6.87	155.47 ± 2.46	54.75 ± 8.71	125.31 ± 7.14	24.75 ± 13.57	34.00 ± 2.47	104.82 ± 22.65	**137.95 ± 1.28**	43.50 ± 9.43	60.16 ± 4.76	19.64 ± 1.53	22.31 ± 1.62
FR [20]	6.02 ± 0.00	6.03 ± 0.00	6.02 ± 0.00	6.02 ± 0.00	6.02 ± 0.00	6.02 ± 0.00	9.40 ± 0.36	10.48 ± 0.20	9.47 ± 0.25	9.75 ± 0.11	9.24 ± 0.16	9.36 ± 0.12

The FDHO activation achieves the highest final PSNR across all six settings, demonstrating consistently superior converged reconstruction quality. It also attains the highest peak PSNR in four of six conditions, with the two exceptions being the **Chirp** at 250 Hz (where Fourier Features reaches 137.95 dB vs. our 120.97 dB) and at 1000 Hz (where FINER reaches 112.57 dB vs. our 94.15 dB). The fact that FDHO leads on final PSNR even in these two settings suggests that the spectral gating mechanism provides more stable long-term convergence: once the gate opens for a coherent frequency component, the equilibrium predicted by Corollary 1(a) prevents further drift, whereas methods without this mechanism may overshoot during training and settle at a lower endpoint. The most revealing pattern emerges at higher frequencies. Non-periodic and spatially compact activations, such as Gauss, WIRE, and MFN, degrade severely as frequency increases: on the 500 Hz **Square Wave**, Gauss drops to 42 dB, and MFN collapses to 6 dB, consistent with their documented low-frequency bias. SIREN shows a characteristic pattern: modest final PSNR but competitive peak PSNR, reflecting its all-pass nature: it can match the signal's frequency content but lacks the spectral gate to stabilise convergence, resulting in a gap between peak and final performance. FDHO combines the strengths of both regimes: strong converged quality (highest final PSNR) with stable training dynamics (smallest gap between peak and final PSNR), consistent with the spectral gating mechanism and the coarse-to-fine curriculum of Sec. 3.4, preventing subsequent drift. The variance across runs also deserves attention. Several INRs show high instability at elevated frequencies: SIREN at 500 Hz **Square Wave** (±39.68 dB), FINER at 1000 Hz **Chirp** (±30.38 dB), and Gauss at 250 Hz **Square Wave** (±42.04 dB), whereas FDHO exhibits a contained variance across all conditions, suggesting that the spectral gating mechanism provides a more robust optimisation landscape.

Audio Fitting. We encode raw audio waveforms as coordinate-based networks mapping time to amplitude, following the protocol of [21]. Audio signals are characterised by rich harmonic structure spanning several orders of magnitude in frequency, making them a demanding test of spectral expressivity. Results are reported in Tab. 2. We evaluate on two clips: **Bach** (tonal, harmonically structured) and **Counting** (speech, broadband and transient-rich), again reporting both final and peak PSNR.

FDHO achieves the highest final PSNR on both signals (56.23 dB on **Bach**, 43.17 dB on **Counting**), confirming that the spectral gating mechanism yields

Table 2: Audio Fitting. Best results in **bold**, runner-ups underlined.

Signal	Metric	FDHO	SIREN	Gauss	WIRE	BACON	FINER	MFN	Fourier	FR
Bach	Final PSNR	**56.23 ± 0.21**	52.02 ± 2.14	19.72 ± 0.10	19.80 ± 0.06	19.50 ± 0.00	38.69 ± 1.64	7.38 ± 0.25	50.75 ± 1.14	17.98 ± 1.20
	Peak PSNR	56.32 ± 0.19	**58.47 ± 0.16**	20.25 ± 0.15	20.80 ± 0.17	19.56 ± 0.01	41.97 ± 0.19	7.38 ± 0.25	52.36 ± 0.64	20.63 ± 0.05
Counting	Final PSNR	**43.17 ± 2.45**	41.64 ± 2.84	24.43 ± 0.01	24.83 ± 0.19	24.43 ± 0.00	34.02 ± 0.60	6.84 ± 0.25	36.20 ± 0.65	24.29 ± 0.18
	Peak PSNR	**45.83 ± 4.20**	45.22 ± 0.66	24.72 ± 0.03	25.11 ± 0.23	24.43 ± 0.00	36.20 ± 1.17	6.84 ± 0.25	36.55 ± 0.40	24.64 ± 0.04

Table 3: Image Fitting. Best results in **bold**, runner-ups underlined.

INR	**Tiger**		**Tiles**		**Bikers**		**Butterfly**		**Knot**	
	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR	Final PSNR	Peak PSNR
FDHO	**63.79 ± 0.23**	**63.79 ± 0.23**	52.24 ± 5.82	57.09 ± 1.31	**57.06 ± 0.97**	**57.25 ± 0.78**	**57.85 ± 0.71**	**57.96 ± 0.61**	**59.21 ± 7.84**	**64.05 ± 3.68**
SIREN [21]	53.69 ± 1.77	58.52 ± 0.25	52.05 ± 1.76	**58.36 ± 0.21**	43.11 ± 11.9	55.92 ± 0.21	48.81 ± 1.51	55.83 ± 0.48	51.90 ± 0.53	59.72 ± 0.78
Gauss [17]	44.10 ± 0.62	46.46 ± 0.21	42.34 ± 1.09	44.77 ± 0.19	41.45 ± 0.74	42.59 ± 0.08	43.19 ± 0.62	45.01 ± 0.44	43.81 ± 1.19	47.08 ± 0.24
WIRE [18]	38.64 ± 0.19	39.20 ± 0.06	40.92 ± 0.45	42.06 ± 0.21	36.71 ± 0.15	36.99 ± 0.26	36.78 ± 0.33	37.10 ± 0.26	40.32 ± 0.27	40.80 ± 0.11
BACON [11]	29.88 ± 0.03	29.90 ± 0.04	32.20 ± 0.04	32.30 ± 0.01	27.45 ± 0.05	27.45 ± 0.05	25.01 ± 0.10	25.01 ± 0.11	31.22 ± 0.03	31.25 ± 0.02
FINER [12]	56.58 ± 0.12	57.09 ± 0.14	**54.09 ± 0.97**	55.31 ± 0.27	47.46 ± 3.13	53.53 ± 0.17	53.65 ± 0.86	55.10 ± 0.08	53.39 ± 1.88	59.72 ± 0.20
MFN [7]	46.97 ± 0.99	47.01 ± 0.97	41.92 ± 3.32	42.59 ± 2.60	38.11 ± 0.53	38.11 ± 0.53	44.12 ± 0.74	44.25 ± 0.72	48.71 ± 1.95	50.59 ± 1.02
Fourier [23]	51.79 ± 0.66	52.77 ± 0.26	46.29 ± 1.45	50.79 ± 0.16	45.35 ± 1.09	47.80 ± 0.12	45.30 ± 3.40	51.58 ± 0.40	51.49 ± 0.52	53.02 ± 0.56
FR [20]	23.01 ± 0.22	23.01 ± 0.22	19.79 ± 0.02	19.79 ± 0.02	19.89 ± 0.05	19.89 ± 0.05	18.68 ± 0.03	18.68 ± 0.03	17.35 ± 0.12	17.35 ± 0.12

consistently superior converged representations even on spectrally demanding targets. On peak PSNR, FDHO ranks first on **Counting** (45.83 dB vs. SIREN's 45.22 dB) and second on **Bach** (56.32 dB vs. SIREN's 58.47 dB). The peak gap on **Bach** is informative: SIREN's all-pass, purely periodic activation is well matched to the sustained harmonic structure of instrumental music, allowing it to transiently achieve high fidelity, but it does not maintain this performance at convergence (final PSNR drops to 52.02 dB, over 4 dB below our activation). This pattern is consistent with SIREN's lack of a stabilising mechanism: without the equilibrium condition of Corollary 1(a), the network continues to adjust frequencies after they have been adequately captured, leading to a lower endpoint despite a higher peak. The broader comparison reveals a stark divide. Nonperiodic activations, such as Gauss, WIRE, BACON, and FR, plateau around 20 to 25 dB on both clips, and MFN fails entirely (below 8 dB), demonstrating that spatially compact activations lack the spectral bandwidth required for audiofrequency content. FINER provides a meaningful improvement (42 dB peak on **Bach**, 36 dB on **Counting**) but remains well below FDHO and SIREN, indicating that variable-periodic modulation alone is insufficient without a mechanism to coordinate spectral bandwidth across the full frequency range. Fourier Features performs competitively on **Bach** (52.36 dB peak) but drops on **Counting** (36.55 dB), suggesting that its fixed random frequency basis struggles with the non-stationary spectral content of speech.

Image Fitting. We fit natural images by learning a mapping from pixel coordinates to RGB values, following the standard protocol established by [21] and adopted across the INR literature [9, 12, 17, 18, 20]. Natural images exhibit a broad, spatially varying frequency spectrum, ranging from smooth regions to fine textures and sharp edges, requiring the activation to accommodate a wide range of spectral content. Results on five diverse images are reported in Tab. 3.

FDHO achieves the highest final and peak PSNR on four of five images (**Tiger**, **Bikers**, **Butterfly**, **Knot**), with margins of up to 10 dB over the clos-

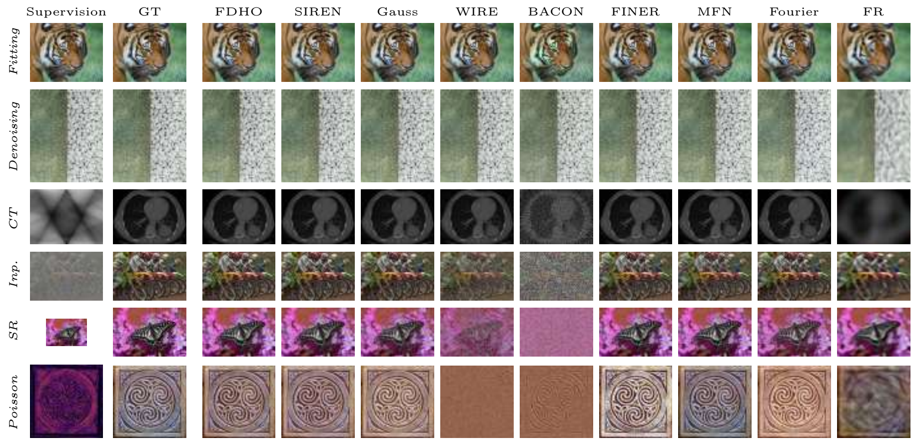

Fig. 1: Qualitative results. Please zoom to appreciate the differences.

est competitor on final quality (**Bikers**: 57.06 dB vs. FINER's 47.46 dB). The sole exception is **Tiles**, where FINER reaches 54.09 dB vs. FDHO's 52.24 dB on final PSNR, and SIREN reaches 58.36 dB vs. FDHO's 57.09 dB on peak PSNR. **Tiles** is a highly regular, repetitive texture, precisely the setting where FINER's variable-periodic activations are best suited, as the spectral content is concentrated at a few dominant spatial frequencies with little need for adaptive bandwidth control. The most evident pattern is the gap between the final and peak PSNR for SIREN. On **Bikers**, SIREN peaks at 55.92 dB but converges to only 43.11 dB, a degradation of nearly 13 dB; similar drops appear on **Tiger** (-4.83 dB), **Butterfly** (-7.02 dB), and **Tiles** (-6.31 dB). This is the overshoot-then-decay behaviour expected from an all-pass activation: SIREN reaches the target frequencies but, lacking a stabilising mechanism, continues to adjust its parameters and drifts away from optimal solutions. FDHO shows no such degradation: on **Tiger**, final and peak PSNR are identical (63.79 dB), and across all images the gap never exceeds 5 dB (with the highest being on **Knot**), consistent with the equilibrium condition of Corollary 1(a) that stabilises ξ once the neuron adequately represents its frequency component. Among the remaining INRs, FINER and SIREN are the strongest overall competitors. Gauss, WIRE, and MFN cluster between 37 and 50 dB, reflecting the low-frequency ceiling of non-periodic and spatially compact activations on images with significant high-frequency content. FR and BACON remain below 32 dB across all images. The variance of FDHO is moderate across conditions, while several INRs exhibit notable instability (*i.e.*, on **Bikers**: SIREN ± 11.9 dB vs. FDHO ± 0.97 dB).

4.2 Signal Recovery from Corrupted Observations

We next evaluate settings where the observed data is degraded, following the model $y = f(\mathbf{x}) + \varepsilon$, where ε represents observation noise. These experiments directly test the noise rejection mechanism formalised in Sec. 3.3. According to Corollary 1, the spectral gate should remain closed for frequency components driven by incoherent noise while opening for those supported by genuine signal content, validating that the implicit regularisation arising from the T_{self} / T_{signal} decomposition (Theorem 2) translates into practical robustness.

Image Denoising. We fit a noisy image (with $\sigma = 0.1$) directly, without any explicit denoising loss or regularisation, relying entirely on the INR's inductive bias to separate signal from noise. This is the purest test of Corollary 1(b): the noise component ε is i.i.d. and produces $O_p(N^{-1/2})$ correlation with the carrier, so the gate should remain closed for noise-driven frequencies while opening for the underlying clean signal. We follow the evaluation protocol described in [12,18] and report only the final PSNR against the clean ground truth in Tab. 4, since peak PSNR is not meaningful when the training target itself is corrupted.

FDHO achieves the highest PSNR on all five images. This is a direct empirical confirmation of the noise rejection mechanism: without any architectural modification, early stopping, or regularisation term, the spectral gate autonomously prevents the network from fitting the noise, as it adapts the bandwidth to the signal content rather than imposing a fixed spectral ceiling. The runner-up across all images is MFN, which is notable, as in the clean image fitting experiments (Tab. 3), MFN ranked among the weaker methods. Its relative strength here stems from its multiplicative structure imposing a strong low-frequency bias that acts as implicit regularisation in the presence of noise, though at the cost of severely limited representational capacity. SIREN, which led or rivalled FDHO on peak PSNR in clean fitting, drops to 26.2 dB here, nearly identical to Gauss and WIRE, confirming that its all-pass nature offers no protection against noise memorisation. FINER similarly falls to the level of the non-periodic activations, indicating that its variable-periodic mechanism does not provide noise rejection. FR performs worst overall (below 23 dB on four of five images), suggesting that its Fourier reparametrisation amplifies rather than suppresses noise-driven frequency content. The variance of our activation is slightly higher than that of the other baselines (e.g., ± 1.25 dB on **Tiger** vs. ± 0.02 dB for SIREN), which reflects the fact that FDHO is actively making non-trivial spectral decisions, *i.e.*, determining which frequencies to fit and which to suppress, whereas the baselines converge to a similar worse solution regardless of the run.

Computed Tomography (CT) Reconstruction. We reconstruct a 2D image from a sparse set of Radon projections, and observations are corrupted by measurement noise. This task combines aspects of both indirect and incomplete observation models: the signal is accessed only indirectly through line integrals, and the limited number of projection angles makes the inverse problem severely ill-posed. The INR must leverage its inductive bias to regularise the reconstruction, making this a joint test of noise robustness and the coarse-to-fine learning curriculum described in Sec. A.3. We follow the setup of [18, 22] and report results in the rightmost column of Tab. 4.

Table 4: Image Denoising and CT Scan Reconstruction. Best results in **bold**, runner-ups <u>underlined</u>.

INR	Tiger	Tiles	Bikers	Butterfly	Knot	CT Scan
FDHO	**30.81 ± 1.25**	**28.57 ± 0.29**	**29.35 ± 0.32**	**31.86 ± 0.69**	**28.67 ± 0.83**	**39.97 ± 0.25**
SIREN [21]	26.22 ± 0.02	26.21 ± 0.02	26.26 ± 0.08	26.40 ± 0.03	26.21 ± 0.01	29.81 ± 1.81
Gauss [17]	26.44 ± 0.03	26.46 ± 0.03	26.77 ± 0.06	27.05 ± 0.06	26.25 ± 0.03	31.22 ± 0.06
WIRE [18]	26.36 ± 0.02	26.69 ± 0.02	26.72 ± 0.02	26.69 ± 0.05	26.29 ± 0.03	27.15 ± 0.19
BACON [11]	25.87 ± 0.02	27.30 ± 0.06	25.48 ± 0.04	23.60 ± 0.04	26.01 ± 0.01	11.11 ± 0.12
FINER [12]	26.23 ± 0.01	26.20 ± 0.03	26.39 ± 0.01	26.54 ± 0.02	26.21 ± 0.01	28.42 ± 6.06
MFN [7]	<u>27.65 ± 0.19</u>	<u>27.62 ± 0.04</u>	<u>28.21 ± 0.10</u>	<u>28.18 ± 0.69</u>	<u>26.89 ± 0.10</u>	<u>37.19 ± 0.07</u>
Fourier [23]	26.59 ± 0.01	26.13 ± 0.44	26.85 ± 0.10	27.36 ± 0.06	26.08 ± 0.25	34.92 ± 0.19
FR [20]	22.87 ± 0.15	19.77 ± 0.02	19.85 ± 0.07	18.62 ± 0.04	17.29 ± 0.16	22.84 ± 0.62

FDHO achieves 39.97 dB, leading the second-best method (MFN, 37.19 dB) by nearly 2.8 dB, the largest absolute margin among all experiments in this section. This result is particularly significant because CT reconstruction compounds two challenges that the spectral gate is designed to address: noise rejection (the projections are noisy) and progressive frequency recovery (sparse angular sampling provides incomplete spectral coverage, the network must reconstruct missing frequency content in a principled order). The ranking mirrors the denoising results: MFN is the runner-up, benefiting from its conservative spectral bias, followed by Fourier Features (34.92 dB) and Gauss (31.22 dB). SIREN drops to 29.81 dB, confirming that its all-pass nature is a liability when the observations are both indirect and noisy: without spectral gating, SIREN fits noise in the sinogram, which propagates into streak artefacts in the reconstructed image (Fig. 1). FINER shows instability in this setting (±6.06 dB), suggesting that its variable-periodic activations are sensitive to the conditioning of the Radon inverse problem. BACON (11.11 dB) and FR (22.84 dB) fail to produce useful reconstructions, indicating that their spectral structures are incompatible with the frequency-domain characteristics of tomographic data.

4.3 Signal Recovery from Incomplete Observations

Finally, we evaluate settings where the signal is observed only partially or indirectly, following the model $y = \Phi(f(\mathbf{x}))$ with Φ a non-invertible forward operator that discards information. Unlike the corrupted setting, the challenge here is not noise but ill-posedness: infinitely many signals are consistent with the observations, and the network must rely on its inductive bias to select a plausible reconstruction. These experiments test whether the coarse-to-fine learning curriculum (Sec. 3.4 and Sec. A.3) and the progressive bandwidth expansion (Sec. A.4) provide useful structural priors for underdetermined inverse problems.

Image Inpainting. We reconstruct an image from a subset of observed pixels, with the forward operator Φ being a binary mask that retains only a fraction (20%) of the pixel locations. Thus, the network must interpolate coherently

Table 5: Image Inpainting and Super-resolution. Best results in **bold**, runner-ups underlined.

INR	Image Inpainting					Image Super-resolution				
	Tiger	Tiles	Bikers	Butterfly	Knot	Tiger	Tiles	Bikers	Butterfly	Knot
FDHO	26.86 ± 0.39	20.47 ± 0.52	23.05 ± 0.10	**24.65 ± 0.10**	**21.11 ± 0.09**	**23.92 ± 0.06**	19.60 ± 0.81	**21.45 ± 0.02**	**20.81 ± 0.04**	**17.84 ± 0.03**
SIREN [21]	**27.35 ± 0.09**	20.70 ± 0.06	**23.33 ± 0.08**	24.15 ± 0.17	20.91 ± 0.04	23.59 ± 0.02	20.27 ± 0.28	20.69 ± 0.04	20.24 ± 0.26	17.40 ± 0.17
Gauss [17]	24.69 ± 0.04	20.66 ± 0.06	22.38 ± 0.03	22.97 ± 0.07	19.74 ± 0.01	20.93 ± 0.04	19.41 ± 0.03	19.79 ± 0.03	18.97 ± 0.02	16.62 ± 0.03
WIRE [18]	17.32 ± 0.14	16.86 ± 0.21	17.51 ± 0.10	15.66 ± 0.21	15.87 ± 0.10	15.47 ± 0.11	15.39 ± 0.11	15.72 ± 0.10	13.37 ± 0.12	14.47 ± 0.07
BACON [11]	8.21 ± 0.05	9.42 ± 0.07	7.65 ± 0.03	6.92 ± 0.02	9.69 ± 0.06	5.50 ± 0.12	8.01 ± 3.68	11.90 ± 0.09	9.00 ± 0.06	5.72 ± 0.19
FINER [12]	26.78 ± 0.10	20.90 ± 0.03	23.00 ± 0.07	24.61 ± 0.08	20.90 ± 0.01	23.57 ± 0.05	20.52 ± 0.14	21.34 ± 0.07	20.67 ± 0.00	17.76 ± 0.02
MFN [7]	24.65 ± 0.04	19.42 ± 0.14	20.74 ± 0.06	22.83 ± 0.13	19.33 ± 0.12	23.34 ± 0.10	20.56 ± 0.07	21.00 ± 0.07	20.42 ± 0.04	17.31 ± 0.04
Fourier [23]	25.25 ± 0.10	**20.87 ± 0.03**	22.34 ± 0.17	23.33 ± 0.11	20.12 ± 0.18	21.28 ± 0.25	19.52 ± 0.23	20.27 ± 0.14	18.87 ± 0.13	16.76 ± 0.09
FR [20]	22.47 ± 0.18	19.46 ± 0.02	19.54 ± 0.04	18.40 ± 0.03	16.73 ± 0.08	21.76 ± 0.10	19.48 ± 0.02	19.33 ± 0.04	18.13 ± 0.02	16.18 ± 0.07

across missing regions. We follow the protocol described in [12, 18] and report results in Tab. 5.

Inpainting is the most competitive setting across baselines, with the top three methods, FDHO, SIREN, and FINER, separated by fractions of a decibel on most images. FDHO ranks first on **Butterfly** (24.65 dB) and **Knot** (21.11 dB), and second on **Tiger** (26.86 dB, behind SIREN's 27.35 dB) and **Bikers** (23.05 dB, behind SIREN's 23.33 dB). On **Tiles**, Fourier Features leads (20.87 dB), with SIREN second (20.70 dB) and FDHO at 20.47 dB. Unlike denoising and super-resolution, where the network must generalise beyond corrupted or downsampled observations, inpainting provides clean pixel values at the observed locations. The missing data is spatially, not spectrally, incomplete. The task, therefore, rewards full spectral bandwidth at the observed coordinates, which is exactly SIREN's strength. The noise rejection is less relevant here as the observations are uncorrupted, and the coarse-to-fine curriculum offers a smaller benefit when the low-resolution structure is already fully determined by the observed pixels scattered across the image. The fact that FDHO nonetheless matches or exceeds SIREN on two of five images, while remaining within 0.5 dB on two others, indicates that the spectral gate does not significantly penalise the network when noise rejection is unnecessary. FINER again emerges as a strong competitor, trailing our activation on **Butterfly** and matching it closely elsewhere, consistent with its strong performance in the clean fitting regime. WIRE, BACON, and FR underperform, confirming that spatial localisation without sufficient spectral bandwidth is inadequate for coherent interpolation across missing regions.

Image Super-resolution. We recover a higher-resolution (512 px) image from a lower-resolution input (128 px), where Φ represents spatial downsampling. The network is supervised only on the low-resolution grid and must hallucinate plausible high-frequency content at the fine scale. This tests whether the coarse-to-fine curriculum produces a natural frequency hierarchy: the network should first lock onto the observed low-frequency content and then extend to higher frequencies guided by the signal's spectral structure. We follow the protocol described in [12, 18, 20] and report results in Tab. 5.

FDHO achieves the highest PSNR on four of five images (**Tiger**, **Bikers**, **Butterfly**, **Knot**), with FINER as the consistent runner-up on those four. The margins are tighter than in clean fitting or denoising, as expected, since super-resolution is fundamentally constrained by the information available in the low-resolution input. On **Tiles**, FDHO (19.60 dB) falls behind MFN (20.56 dB) and FINER (20.52 dB), echoing the results in clean fitting, where FINER also led: the highly periodic texture of **Tiles** favours activations whose spectral support aligns with its regular structure, reducing the advantage of adaptive bandwidth control. An interesting shift from the fitting and denoising experiments is MFN's performance. MFN was among the weakest methods on clean fitting (Tab. 3) but the second-best on denoising (Tab. 4), and here it is competitive with FINER and SIREN across all images. Super-resolution shares a key property with denoising: the network must generalise beyond the observed data, and methods with conservative spectral behaviour avoid hallucinating high-frequency artefacts at unsupervised locations. FDHO achieves this conservatism through the spectral gating. The coarse-to-fine curriculum ensures that low-frequency structure is captured first and higher frequencies are admitted only when justified, while MFN achieves it through a fixed low-frequency ceiling that, as the clean fitting results show, limits its capacity when the full spectrum must be represented. WIRE and BACON further degrade in the super-resolution setting, suggesting that their spatial localisation, while beneficial for fitting observed data, does not extrapolate well to unobserved coordinates. SIREN remains competitive on **Tiger** (23.59 dB, second-best) but falls behind on the remaining images, consistent with its lack of a mechanism to prioritise the observed low-frequency content before extending to higher frequencies.

5 Discussion

We have presented FDHO, a physics-grounded activation function for INRs derived from the steady-state response of a damped harmonic oscillator. The core insight is that coupling amplitude to frequency through a physical transfer function creates an implicit spectral regulariser: the gradient of the reconstruction loss with respect to the damping factor decomposes into a self-closing force and a signal-driven opening force, yielding a precise gating condition that distinguishes coherent signal from noise without any explicit penalty in the loss.

This mechanism yields state-of-the-art or competitive results across all the considered tasks without the need to tune either the architecture or any other hyperparameters, whereas each competitor INR needs specific tuning, and performs well in some settings but degrades significantly in others. Notably, the cases where FDHO turns out less beneficial – clean, spatially incomplete observations and highly regular textures – are those where Corollary 1 predicts the least advantage, confirming that the theory correctly characterises the method's behaviour.

Extending the formulation to per-neuron parameters and formalising the multi-layer gradient dynamics are the main directions for future work.

References

1. Abramowitz, M., Stegun, I.A.: Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. Dover Publications (1964)
2. Arfken, G.B., Weber, H.J., Harris, F.E.: Mathematical Methods for Physicists. Academic Press, 7th edn. (2012)
3. Arora, S., Cohen, N., Hu, W., Luo, Y.: Implicit regularization in deep matrix factorization. In: Advances in Neural Information Processing Systems (NeurIPS). vol. 32 (2019)
4. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: ICCV (2021)
5. Butterworth, S.: On the theory of filter amplifiers. Wireless Engineer **7**(6), 536–541 (1930)
6. Carson, J.R.: Notes on the theory of modulation. Proceedings of the Institute of Radio Engineers **10**(1), 57–64 (1922)
7. Fathony, R., Sahu, A.K., Willmott, D., Kolter, J.Z.: Multiplicative filter networks. In: ICLR (2021)
8. Jacot, A., Gabriel, F., Hongler, C.: Neural tangent kernel: Convergence and generalization in neural networks. In: NeurIPS (2018)
9. Jayasundara, D., Zhao, H., Labate, D., Patel, V.M.: Mire: Matched implicit neural representations. In: CVPR (2024)
10. Li, Z., Luo, Y., Lyu, K.: Towards resolving the implicit bias of gradient descent for matrix factorisation: Greedy low-rank learning. In: ICLR (2021)
11. Lindell, D.B., Van Veen, D., Park, J.J., Wetzstein, G.: Bacon: Band-limited coordinate networks for multiscale scene representation. In: CVPR (2022)
12. Liu, Z., Zhu, H., Zhang, Q., Fu, J., Deng, W., Ma, Z., Guo, Y., Cao, X.: Finer: Flexible spectral-bias tuning in implicit neural representation by variable-periodic activation functions. In: CVPR (2024)
13. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
14. Novello, T., Aldana, D., Araujo, A., Velho, L.: Tuning the frequencies: Robust training for sinusoidal neural networks. In: CVPR (2025)
15. Olver, F.W.J., Lozier, D.W., Boisvert, R.F., Clark, C.W.: NIST Handbook of Mathematical Functions. Cambridge University Press (2010)
16. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
17. Ramasinghe, S., Lucey, S.: Beyond periodicity: towards a unifying framework for activations in coordinate-mlps. In: ECCV (2022)
18. Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., Baraniuk, R.G.: Wire: Wavelet implicit neural representations. In: CVPR (2023)
19. Shi, K., Chen, H., Zhang, L., Gu, S.: Inductive gradient adjustment for spectral bias in implicit neural representations. In: ICML (2025)
20. Shi, K., Zhou, X., Gu, S.: Improved implicit neural representation with fourier reparameterized training. In: CVPR (2024)
21. Sitzmann, V., Martel, J.N., Bergman, A.W., Lindell, D.B., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: NeurIPS (2020)
22. Sun, Y., Liu, J., Xie, M., Wohlberg, B., Kamilov, U.S.: Coil: Coordinate-based internal learning for tomographic imaging. IEEE Transactions on Computational Imaging (2021)

23. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: NeurIPS (2020)
24. Watson, G.N.: A Treatise on the Theory of Bessel Functions. Cambridge University Press, 2nd edn. (1944)
25. Yüce, G., Ortiz-Jiménez, G., Bekerš, B., Frossard, P.: A structured dictionary perspective on implicit neural representations. In: CVPR (2022)