

EVOS: Efficient Implicit Neural Training via EVolutionary Selector

Weixiang Zhang¹, Shuzhao Xie¹, Chengwei Ren¹, Siyi Xie², Chen Tang³,
Shijia Ge¹, Mingzi Wang¹, Zhi Wang^{1*}

¹Tsinghua Shenzhen International Graduate School, Tsinghua University

²Xi'an Jiaotong University ³The Chinese University of Hong Kong

Abstract

We propose *EVolutionary Selector (EVOS)*, an efficient training paradigm for accelerating *Implicit Neural Representation (INR)*. Unlike conventional INR training that feeds all samples through the neural network in each iteration, our approach restricts training to strategically selected points, reducing computational overhead by eliminating redundant forward passes. Specifically, we treat each sample as an individual in an evolutionary process, where only those fittest ones survive and merit inclusion in training, adaptively evolving with the neural network dynamics. While this is conceptually similar to *Evolutionary Algorithms*, their distinct objectives (selection for acceleration vs. iterative solution optimization) require a fundamental redefinition of evolutionary mechanisms for our context. In response, we design sparse fitness evaluation, frequency-guided crossover, and augmented unbiased mutation to comprise EVOS. These components respectively guide sample selection with reduced computational cost, enhance performance through frequency-domain balance, and mitigate selection bias from cached evaluation. Extensive experiments demonstrate that our method achieves approximately 48%-66% reduction in training time while ensuring superior convergence without additional cost, establishing state-of-the-art acceleration among recent sampling-based strategies. Our code is available at [this link](#).

1. Introduction

Implicit Neural Representation (INR) is an emerging paradigm [12, 43, 46] that leverages Multilayer Perceptrons (MLPs) to map spatial coordinates (e.g., (x, y) for images) to corresponding attributes (e.g., RGB values of pixels), enabling continuous representation of diverse natural signals, including sketches [2], audio [45], images [43], 3D shapes [31] and climate data [16]. Compared to conven-

tional discrete representations, INRs offer significant advantages in memory efficiency and inverse problem solving capability, demonstrating great potential in data compression [44], image enhancement [8], inverse imaging [17], novel view synthesis [27], and physics simulation [34].

However, the encoding process of INRs remains computationally intensive due to dense iterations required for optimization, presenting a significant barrier to widespread adoption. Furthermore, spectral bias [33], where neural networks inherently prioritize low-frequency over high-frequency components, impedes fine detail fitting and introduces additional training overhead. In response, various acceleration approaches have been proposed, each with inherent limitations: partition-based methods [21, 25, 38] increase architectural complexity through multiple networks, explicit representation collaborations [5, 28, 49] compromise memory efficiency, and meta-learning approaches [9, 47] demand substantial homogeneous data for pre-training.

Unlike these approaches, recent sampling-based methods [13, 53, 55] propose to accelerate training through sparsified forward passes by selectively sampling important coordinates (samples). While this sampling paradigm offers potential cost-free acceleration, existing methods are limited by either relying solely on static signal structures without considering network dynamics [13, 53] or employing computationally expensive greedy algorithms for coordinate selection [53], thus constraining their efficiency gains.

To address these limitations, we propose *EVolutionary Selector (EVOS)*, a sampling-based method to accelerate INR training, inspired by principles of biological evolution. The basic motivation behind EVOS is that, the selection of optimal coordinate subsets conceptually parallels *natural selection*, where only the fittest individuals survive and merit computational resources, adaptively evolving with network dynamics via *crossover* and *mutation* operations. This paradigm shares principles with *Evolutionary Algorithms (EA)* [15], yet serves a fundamentally different objective. While EA maintains a population of search points to iteratively refine candidate solutions toward op-

*Corresponding author.

timal outcomes, our approach focuses on efficient subset selection rather than iterative solution optimization. Despite this distinction, evolutionary principles offer compelling advantages for complex [52] and dynamic optimization problems [29], making them well-suited for addressing the non-convex nature of neural optimization [10]. This motivates us to preserve the core evolutionary concept while redefining key components (*fitness evaluation*, *crossover*, and *mutation*) to align with INR acceleration goals.

Specifically, given a signal composed of coordinates \mathbf{x} and corresponding attributes \mathbf{y} , the objective is to optimize an MLP $F_\theta(\mathbf{x})$ to approximate \mathbf{y} with minimal loss. At each iteration t , EVOS treats a subset of coordinates $\mathbf{x}'_t \subset \mathbf{x}$ as *survivors*, which is selected through *fitness evaluation* from the full coordinates \mathbf{x} . Subsequently, *offspring* \mathbf{z}_t are generated from these *survivors* through *crossover* and *mutation* operations, and then fed into $F_\theta(\mathbf{z}_t)$ to achieve sparsified forward passes ($|\mathbf{z}_t| \leq |\mathbf{x}|$). The network parameters θ are then updated through backpropagation. Through reducing the number of forward passes from $|\mathbf{x}|$ to $|\mathbf{z}_t|$ in each iteration, computational acceleration is achieved.

We redefine *fitness evaluation*, *crossover*, and *mutation* as follows: **(1) Sparse Fitness Evaluation** (Sec. 3.2). Computing fitness values in our context requires network forward passes for each coordinate, making conventional per-iteration fitness evaluation computationally prohibitive. To address this, we perform sparse evaluations at strategic intervals and cache results for intermediate steps, significantly reducing computational overhead. **(2) Frequency-Guided Crossover** (Sec. 3.3). To address spectral bias that impedes high-frequency component fitting, we select coordinates from both low and high-frequency perspectives. These serve as parents for crossover operations, generating offspring that balance different frequency preferences and enhance performance with negligible computational cost. **(3) Augmented Unbiased Mutation** (Sec. 3.4). Continuous reliance on cached evaluation results can introduce training bias and degrade performance [58]. We mitigate this through stochastic incorporation of non-surviving points into the offspring set, introducing controlled uncertainty in each iteration through mutation-like operations.

We conduct extensive experiments to verify the performance (Sec. 4.2), compatibility (Sec. 4.3 & Sec. 4.4) and effectiveness (Sec. 4.5) of our method. With the integration of EVOS, the time cost of INRs’ training can be reduced by 48%-66% without compromising reconstruction quality. Compared to existing sampling-based methods, including recent approaches such as Soft Mining [18] and INT [53], EVOS achieves state-of-the-art training efficiency. A key insight emerging from our study is that strategic sparsification of training samples not only reduces computational cost but also consistently enhances training performance, challenging the conventional paradigm of exhaustive data

utilization for signal fitting tasks [12, 43]. To summarize, our contributions can be outlined as follows:

- Inspired by principles of biological evolution, we propose EVOLUTIONARY Selector (EVOS), a sampling-based method that accelerates INR training while preserving reconstruction fidelity.
- We reformulate key notions in conventional evolutionary computation to meet our acceleration objectives, including sparse fitness evaluation, frequency-guided crossover, and augmented unbiased mutation.
- We conduct extensive experiments to verify the performance and compatibility of EVOS, demonstrating that EVOS achieves state-of-the-art efficiency improvement among recent sampling-based acceleration techniques.

2. Related Work

2.1. Implicit Neural Representations

Implicit Neural Representations (INRs) [12, 43, 46] continuously represent natural signals by training coordinate-based MLPs, whose structures are meticulously designed to achieve high representation capability, including position encoding [46] and frequency-guided activation functions [23, 35, 39, 43]. This representation paradigm has demonstrated broad applicability across diverse modalities, including images [11, 43, 49], videos [7, 19, 24], shapes [4, 20, 31], 3D scenes [5, 27, 28], and unconventional data formats [2, 6, 16]. Leveraging its memory efficiency and inherent capability in solving inverse problems, INRs have expanded beyond signal representation to applications in data compression [44], generation [30], novel view synthesis [27], inverse imaging [17], and partial differential equation solving [34]. Despite these advantages, the encoding process of INRs remains constrained by intensive computational requirements and spectral bias [33], which impedes high-fidelity detail reconstruction and limits their widespread adoption.

2.2. Acceleration for INR Training

Numerous approaches have been proposed to address the computational intensity of INR training, which can be categorized into four main streams: **(1) Partition-based methods** decompose the input signal into smaller components, employing multiple compact MLPs for local representation. These include approaches based on regular blocks [37], adaptive blocks [25], Voronoi diagrams [36], Laplacian pyramids [38], and segmentation maps [21]. **(2) Prior-driven methods** leverage knowledge distilled from large-scale datasets to expedite convergence. This category encompasses meta-learning approaches for parameter initialization [9, 47] and hyper-network-based weight generation [43, 45]. **(3) Explicit caching methods** trade spatial complexity for temporal efficiency by maintaining en-

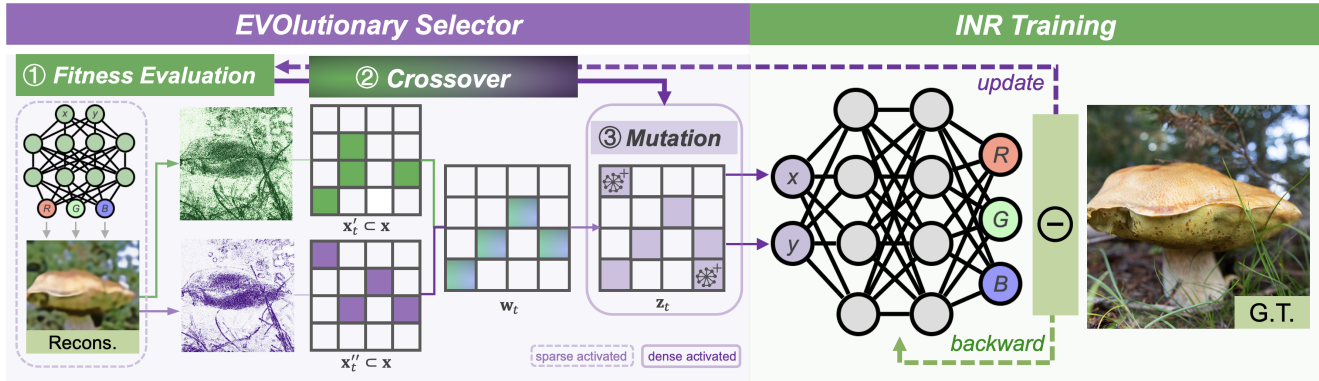


Figure 1. **Overview of EVOS framework.** The proposed method aims to optimize a MLP for implicit signal representation via evolutionary selector. The process comprises three key components: (1) (Sparse) Fitness Evaluation (Sec. 3.2) for efficiently guiding coordinate selection, (2) (Frequency-Guided) Crossover (Sec. 3.3) for improving performance by balancing frequency domain preferences, and (3) (Augmented Unbiased) Mutation (Sec. 3.4) for mitigating selection bias in each iteration. The selected coordinates from this evolutionary process are then fed into the network, enabling sparsified forward passes and reduced computational costs.

coded features in explicit data structures, such as hash maps [28, 49], point cloud [50], low-rank tensors [5], and tree structures [22, 51]. (4) Alternative approaches explore diverse strategies to enhance training efficiency, including modulator networks [26], reparameterized training [42], data transformation [40, 56], batch normalization [3], and gradient adjustment [41].

Our method aligns with recent sampling-based acceleration approaches. Several methods have emerged in this direction: EGRA [13], Expansive Supervision [55], and Soft Mining [18], tailored for Neural Radiance Field [27] (NeRF) training, demonstrate acceleration through selective *Ray Sampling* based on edge detection, frequency priors, and Monte-Carlo methods, respectively. Most relevant to our approach, INT [53] addresses general INR acceleration by reformulating the learning process as a nonparametric teaching problem and employing greedy functional algorithms for coordinate selection. EVOS distinctively leverages evolutionary principles for coordinate selection, introducing novel mechanisms for fitness evaluation, crossover, and mutation that achieve superior performance compared to existing sampling-based methods.

3. Method

3.1. Formulation and Overview

Consider a natural signal $S = (\mathbf{x}, \mathbf{y})$, where $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^n$ represent coordinates and their corresponding attributes in m and n dimensions, respectively. Our target is to learn an INR by overfitting a MLP $F_\theta(\mathbf{x}) : \mathbf{x} \in \mathbb{R}^m \mapsto \mathbf{y} \in \mathbb{R}^n$ to represent signal S . Under the EVOS training paradigm, the optimization of F_θ is guided by a sparse forward pass $L(F_\theta(\mathbf{z}_t), \mathbf{y}_z)$ rather than the conventional $L(F_\theta(\mathbf{x}), \mathbf{y})$, where $\mathbf{z}_t \subset \mathbf{x}$ represents the selected

coordinate subset at iteration t , and $\mathbf{y}_z \subset \mathbf{y}$ denotes their corresponding ground truth values. Here, $L(\cdot)$ represents the loss function. The coordinate subset \mathbf{z}_t is determined through three key steps:

1. The *survivors* $\mathbf{x}'_t \subset \mathbf{x}$ are selected through Sparse Fitness Evaluation (Sec.3.2), efficiently identifying the fittest candidates at a coarse level.
2. The *survivors* \mathbf{x}'_t undergo Frequency-Guided Crossover (Sec.3.3) to generate *offspring* \mathbf{w}_t , where parent selection is guided by frequency preferences to balance different spectral components.
3. The *offspring* \mathbf{w}_t are further refined through Augmented Unbiased Mutation (Sec.3.4) to obtain the final subset \mathbf{z}_t , incorporating stochastic perturbations to mitigate selection bias.

Finally, the network parameters are optimized by gradients of $L(F_\theta(\mathbf{z}_t), \mathbf{y}_z)$. By reducing the number of forward passes from $|\mathbf{x}|$ to $|\mathbf{z}_t|$ in each iteration, EVOS achieves an acceleration factor of $I(\frac{|\mathbf{x}|-|\mathbf{z}_t|}{|\mathbf{x}|})$, where $I(\cdot)$ represents the ratio of forward pass time to total computation time. The overview of EVOS is illustrated in Fig. 1.

3.2. Sparse Fitness Evaluation

In conventional evolutionary algorithms, the fitness function evaluates the quality of each individual in the population, iteratively refining it until convergence to an optimal solution. Adapting this concept to EVOS, we define the fitness function as the distance between reconstructed and ground truth values, prioritizing coordinates with larger discrepancies for selection. Formally, the fitness function is expressed as $f(\mathbf{x}) = \mathcal{D}(F_\theta(\mathbf{x}), \mathbf{y})$, where $\mathcal{D}(\cdot)$ denotes the distance metric. The survivors \mathbf{x}'_t are selected through $\mathbf{x}'_t = \arg \max_{\mathbf{x}'_t \subset \mathbf{x}} (f(\mathbf{x}))$.

However, since implementing the fitness function $f(\mathbf{x})$

requires computing all coordinates, performing such resource-intensive computations at each iteration would impose significant overhead, contradicting our efficiency objectives. To address this, we introduce sparse fitness evaluation, which performs assessment only at key iterations while caching the results for intermediate steps. The sparse fitness function at iteration t is formulated as:

$$f_t(\mathbf{x}) = \Gamma(t)(\mathcal{D}(F_\theta(\mathbf{x}), \mathbf{y})) + (1 - \Gamma(t))\hat{f}(\mathbf{x}), \quad (1)$$

where $\hat{f}(\mathbf{x})$ represents the cached evaluation from the last key iteration and $\Gamma(t)$ is an indicator function that equals 1 at key iterations and 0 otherwise. When $\Gamma(t) = 1$, the cached evaluation is updated: $\hat{f}(\mathbf{x}) := f_t(\mathbf{x})$.

Notably, the frequency of key iterations in $\Gamma(t)$ presents a trade-off: more frequent evaluations offer finer-grained guidance but increase computational overhead. Our empirical analysis reveals that the distribution changes in $\mathcal{D}(F_\theta(\mathbf{x}), \mathbf{y})$ exhibit a gradual linear increase across iterations. Based on this observation, we implement a linearly increasing schedule for key iterations to balance performance and efficiency:

$$\Gamma(t) = \begin{cases} 1, & \text{if } t \bmod (\tau - \frac{\theta t}{T}) = 1 \\ 0, & \text{otherwise,} \end{cases} \quad (2)$$

where T denotes the total number of iterations, and the interval between key iterations is $\tau - \frac{\theta t}{T}$, controlled by coefficients τ and θ . This interval decreases as training progresses, resulting in more frequent fitness evaluations.

3.3. Frequency-Guided Crossover

After fitness evaluation, the survivors \mathbf{x}' are selected using the fitness function $f(\mathbf{x}) = \mathcal{D}(F_\theta(\mathbf{x}), \mathbf{y})$. While defining $\mathcal{D}(\cdot)$ as the square error $\|F_\theta(\mathbf{x}) - \mathbf{y}\|^2$ presents a straightforward measure of reconstruction quality, this coordinate-wise approach neglects spatial correlations, resulting in oversmoothed reconstructions and insufficient emphasis on high-frequency details. Given that INRs inherently exhibit spectral bias toward low-frequency components, this simple square error-based selection further exacerbates the challenge of learning high-frequency features.

To address this limitation, we propose a dual-perspective coordinate selection strategy that considers both low and high-frequency components. These two complementary subsets serve as parents in a crossover operation, producing offspring that balance different frequency preferences. Specifically, we employ the Laplacian operator L_{lap} , a differential operator based on the gradient divergence, to measure high-frequency components, while retaining square error for low-frequency assessment. At iteration t , given coordinates \mathbf{x}_t , we compute two sets of parents: \mathbf{x}'_t (low-frequency preference) and \mathbf{x}''_t (high-frequency preference)

by selecting survivors using fitness functions $f_t^{low}(\mathbf{x})$ and $f_t^{high}(\mathbf{x})$, respectively:

$$\begin{aligned} \mathbf{x}'_t &= \arg \max_{\{\mathbf{x}'\}_k \subseteq \{\mathbf{x}\}_N} \underbrace{(\|F_\theta(\mathbf{x}) - \mathbf{y}\|^2)}_{f_t^{low}(\mathbf{x})}, \\ \mathbf{x}''_t &= \arg \max_{\{\mathbf{x}''\}_k \subseteq \{\mathbf{x}\}_N} \underbrace{(\|L_{lap}(F_\theta(\mathbf{x})) - L_{lap}(\mathbf{y})\|^2)}_{f_t^{high}(\mathbf{x})}, \end{aligned} \quad (3)$$

where N denotes the total number of coordinates and k represents the number of selected survivors. This selection process integrates with the sparse fitness evaluation mechanism described in Eq. 1.

Subsequently, the offspring \mathbf{w}_t is generated through crossover operation between \mathbf{x}'_t and \mathbf{x}''_t :

$$\mathbf{w}_t = \mathbf{x}'_t \odot \mathbf{x}''_t = \{\mathbf{x}'_t \cap \mathbf{x}''_t\} \cup \Psi(\mathbf{x}'_t, \mathbf{x}''_t), \quad (4)$$

where $\{\mathbf{x}'_t \cap \mathbf{x}''_t\}$ represents coordinates of cross-frequency importance preserved in the offspring \mathbf{w}_t . The balancer function $\Psi(\mathbf{x}'_t, \mathbf{x}''_t)$ augments the offspring by selecting additional individuals from \mathbf{x}'_t and \mathbf{x}''_t based on the current reconstruction quality across frequency components. This function is formulated as $\Psi(\mathbf{x}'_t, \mathbf{x}''_t) = \mathbf{v}'_t \cup \mathbf{v}''_t$, where:

$$\begin{aligned} \mathbf{v}'_t &\sim \mathcal{U}_{pl}(\mathbf{x}'_t \setminus \mathbf{x}''_t), \\ \mathbf{v}''_t &\sim \mathcal{U}_{(1-p)l}(\mathbf{x}''_t \setminus \mathbf{x}'_t). \end{aligned} \quad (5)$$

Here, $p = \frac{f_t^{low}(\mathbf{x})}{(f_t^{low}(\mathbf{x}) + f_t^{high}(\mathbf{x}))}$, $l = k - |\mathbf{x}'_t \cap \mathbf{x}''_t|$ denotes the number of samples, and \mathcal{U} represents the uniform distribution. When low-frequency reconstruction exhibits superior quality ($p < 0.5$), $\Psi(\mathbf{x}'_t, \mathbf{x}''_t)$ prioritizes coordinates from the high-frequency preference set \mathbf{x}''_t , and vice versa.

Cross-Frequency Supervision. To leverage the Laplacian gradient computed in Eq. 3, we introduce cross-frequency loss to complement the crossover operation. For the mutated offspring \mathbf{z}_t (Sec. 3.4), we formulate the cross-frequency loss as:

$$\begin{aligned} L_t^c &= \lambda_l L_t^l + \lambda_h L_t^h, \\ L_t^l &= \|F_\theta(\mathbf{z}_t) - \mathbf{y}_z\|^2, \\ L_t^h &= \|L_{lap}[F_\theta(\mathbf{z}_t) + \hat{F}_\theta(\mathbf{x} \setminus \mathbf{z}_t)] - L_{lap}(\mathbf{y})\|^2. \end{aligned} \quad (6)$$

Since the $L_{lap}(\cdot)$ operator requires a full volume of attributes for its spatial convolution filter, the sparsified coordinates \mathbf{z}_t cannot be directly applied. We address this by incorporating cached evaluation results $\hat{F}_\theta(\mathbf{x} \setminus \mathbf{z}_t)$ from the last key iteration. λ_l and λ_h denote the scaling factors.

3.4. Augmented Unbiased Mutation

While EVOS's sparse fitness evaluation (Sec. 3.2) significantly reduce computational cost, they introduce a potential

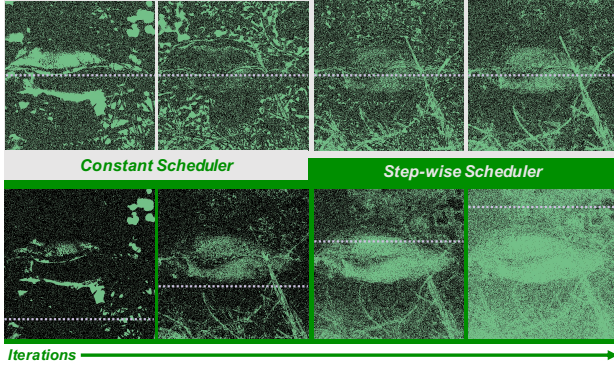


Figure 2. Visualization of Constant (Top) and Step-wise (Bottom) schedulers. Constant maintains fixed selection ratio throughout iterations, while Step-wise implements progressive increase.

drawback: the network receives deterministic offspring coordinates \mathbf{w}_t during consecutive iterations, potentially inducing selection bias and degrading performance [58]. To address this, we incorporate uncertainty through mutation, stochastically integrating a small portion of non-selected individuals $\{\mathbf{x} \setminus \mathbf{w}_t\}$ into the offspring. This strategy maintains population diversity and mitigates training bias.

At iteration t , given offspring set \mathbf{w}_t , we generate mutated coordinates \mathbf{z}_t through the augmented unbiased mutation operator $\mathcal{M}(\cdot)$:

$$\mathbf{z}_t = \mathcal{M}(\mathbf{w}_t) = \mathbf{w}_t \cup \mathbf{m}_t, \quad \mathbf{m}_t \sim \mathcal{U}_{\alpha k}(\{\mathbf{x} \setminus \mathbf{w}_t\}), \quad (7)$$

where \mathbf{m}_t denotes mutated coordinates of length αk uniformly sampled from the non-selected set $\{\mathbf{x} \setminus \mathbf{w}_t\}$, and α represents the mutation ratio controlling uncertainty. The final number of mutated offspring is $|\mathbf{z}_t| = (1 + \alpha)k < N$. Given the negligible computational cost of uniform sampling, Eq. 7 is executed in each iteration.

Finally, we utilize the mutated offspring \mathbf{z}_t as the fittest subset of coordinates for iteration t in training network F_θ . The network parameters θ are optimized according to:

$$\theta := \theta - \eta \nabla L_t^c(F_\theta(\mathbf{z}_t), \mathbf{y}_z), \quad (8)$$

where η denotes the learning rate and L_t^c represents the *cross-frequency loss* detailed in Eq. 6.

3.5. Discussions

Scheduler for Selection Size. For determining the selection intensity q ($q = |\mathbf{z}_t|$), we implement two scheduling mechanisms: (1) Constant scheduler: maintains a fixed selection size $q = \beta N$ throughout training, where $\beta \in (0, 1)$ and N denotes the total number of coordinates \mathbf{x} , ensuring consistent memory efficiency. (2) Step-wise scheduler: progressively increases q following a step function during training, which achieves optimal performance as demonstrated in INT [53]. We adopt this scheduler with default

parameters with INT (q increasing from $0.2N$ to N). Both schedulers are implemented by adjusting k (Eq. 3) according to $q = (1 + \alpha)k$. Fig. 2 illustrates the temporal dynamics of both scheduling mechanisms.

Complexity. Our method achieves acceleration without introducing additional storage overhead. Selection operations in EVOS account for merely 1.34% of training time, which is negligible considering the method’s 48%-66% reduction in total training duration.

4. Experiments

4.1. Implement Details

Experimental Settings. We set the total number of iterations T to 5000, with hyperparameters $\tau = 100$, $\theta = 0.01$, $\alpha = 0.5$, $\lambda_l = 1$, $\lambda_h = 1e-5$, and $\beta = 0.5$. Following [23, 46], we evaluated our method on processed DIV2K datasets [1] using 3×256 MLP with SIREN [43] architecture unless otherwise stated. Additional experiments across various modalities and datasets are provided in the *supplementary materials*. All experiments in this section were performed using the PyTorch framework [32] on 4 NVIDIA RTX 3090 GPUs with 24.58 GB VRAM each.

Evaluation Mechanisms. We evaluated reconstruction quality using PSNR, SSIM [48], and LPIPS [54] metrics. The efficiency improvements were assessed through two mechanisms: (1) comparing time required to achieve target reconstruction quality, and (2) comparing performance and computational cost under fixed iterations. The first directly reveals efficiency gains across different strategies, while the second compares the per-iteration performance between full data utilization and partial data sampling.

4.2. Comparison with State-of-the-arts Strategies

Settings. We compared our method against Uniform Sampling (Uniform.), EGRA [13], Expansive Supervision (Expan.) [55], Soft Mining [18], INT [53], and conventional full-coordinate training (Standard) commonly used in INR encoding [23, 35, 39, 43, 46]. For INT, which addresses the identical application scenario as ours, we strictly followed to its official implementation and settings. Although EGRA and Expan. were originally tailored for NeRF training, their straightforward mechanisms enabled direct adaptation to our experiments. For Soft Mining, we modified parameters to accommodate general INR training, though it demonstrated limited effectiveness in 2D image fitting tasks. Through empirical tuning, optimal performance was achieved with $\alpha_{soft} = 0.5$ and disabled fuzzy indexing (detailed in supplementary materials). Apart from these necessary adjustments, all experimental settings remained consistent across methods.

Quantitative Results. Tables 1 and 2 present quantitative comparisons under fixed iterations and target PSNR

Strategies	1k Iterations			2k Iterations			5k Iterations			Time↓ (sec)
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	
Standard	31.06	0.899	0.123	34.34	0.944	0.042	37.10	0.964	0.021	180.45
Uniform.	30.35	0.885	0.149	33.33	0.932	0.058	36.11	0.956	0.026	94.13
EGRA [13]	30.37	0.885	0.150	33.40	0.932	0.057	36.22	0.957	0.025	103.31
Expan. [55]	30.95	0.893	0.121	33.70	0.933	0.049	36.28	0.955	0.023	96.39
Soft Mining [18]	31.49	0.903	0.124	33.52	0.931	0.072	35.31	0.948	0.045	100.35
INT [53] (incre.)	31.67	0.901	0.117	31.95	0.894	0.096	34.62	0.923	0.059	110.55
INT [53] (dense.)	31.67	0.901	0.117	34.53	0.938	0.052	37.21	0.959	0.027	161.89
EVOS (w/o CFS.)	31.64	0.899	0.107	34.73	0.940	0.040	37.49	0.960	0.019	95.43
EVOS (proposed)	32.71	0.912	0.082	35.89	0.951	0.031	37.81	0.962	0.018	97.39
<u>EGRA [13]</u>	29.59	0.867	0.187	33.01	0.927	0.068	37.08	0.963	0.020	105.71
<u>Expansive Sup. [55]</u>	30.58	0.883	0.138	33.55	0.930	0.058	36.95	0.960	0.026	103.09
<u>INT [53] (incre.)[†]</u>	31.52	0.880	0.137	30.59	0.863	0.138	37.17	0.962	0.020	124.35
<u>EVOS (w/o CFS.)</u>	30.99	0.873	0.124	34.49	0.930	0.050	37.81	0.965	0.019	105.82
<u>EVOS (proposed)</u>	31.69	0.882	0.106	35.61	0.940	0.041	38.43	0.968	0.016	108.60

[†] denotes the best-performing variant reported in INT [53].

Table 1. Comparison of sampling strategies under fixed iterations. Strategies without underlines employ constant scheduler ($\beta = 0.5$), while underlined strategies implement step-wise scheduler. **Forest** : the best performance; **Mint** : exceeds standard training.

PSNR	Stand.	EGRA	Expan.	Soft.	INT [†]	EVOS
25 dB	9.04	6.19	5.08	4.99	6.15	4.73
30 dB	30.75	19.63	16.42	14.14	18.68	12.48
35 dB	88.64	67.15	59.41	84.81	75.35	29.73

Table 2. Comparison of computational time (seconds) required to achieve target PSNR across different strategies.

benchmarks, respectively. In Table 1, INT (incre.) and INT (dense.) denote variants optimized for minimal time cost and maximal reconstruction quality, respectively. For fair comparison independent of loss function variations, we include EVOS (w/o CFS.), which implements vanilla L2 loss without Cross-Frequency Supervision (Eq. 6). Our method demonstrates consistent superior efficiency across all comparative strategies. Specifically, under fixed iterations, our approach achieves 39.82% time reduction while improving PSNR by 1.33 dB compared to standard training (final row of Table 1). Furthermore, when targeting PSNR values of 25, 30, and 35 dB, our method reduces computational time by 47.68%, 59.41%, and 66.46% respectively (Table 2), outperforming all alternative approaches.

Visualization Results. Fig. 3 presents reconstructed images after 60 seconds of training using different acceleration methods. Under identical time constraints, our method achieves superior reconstruction quality, surpassing alternative approaches by 2.31 to 6.44 dB. Detailed examination reveals that our method particularly excels in preserving high-fidelity details, as evidenced by the wolf’s eye and

fur regions, producing results closest to ground truth.

Temporal Evolution of Training. The temporal evolution of reconstruction quality is illustrated in Fig. 4, which plots PSNR and SSIM metrics against training time. For precise evaluation, additional full-coordinate forward passes were performed after each optimization step to compute accurate metrics, with these evaluation passes excluded from the total training time. Our method demonstrates both faster convergence and superior performance compared to alternative approaches, particularly in PSNR metrics. Moreover, EVOS exhibits more stable performance gains than INT, maintaining its advantage throughout the training process.

Analysis of EGRA and Expan. While EGRA and Expan. leverage edge detection and frequency priors for coordinate sampling, accelerating training in the early stages, their overreliance on static image information without considering dynamic network changes results in suboptimal long-term performance.

Analysis of Soft Mining. Despite empirical parameter tuning, Soft Mining exhibits deteriorating reconstruction quality, particularly in later training stages. This limitation likely stems from sampling space incompatibility. While Soft Mining effectively samples approximately 5% of coordinates from vast 3D radiance field data in NeRF training, conventional INR training typically processes all available data in each iteration due to relatively small data volume. This fundamental difference in sampling requirements explains its reduced effectiveness in non-NeRF INR applications. Supporting experimental evidence is provided in the supplementary material.

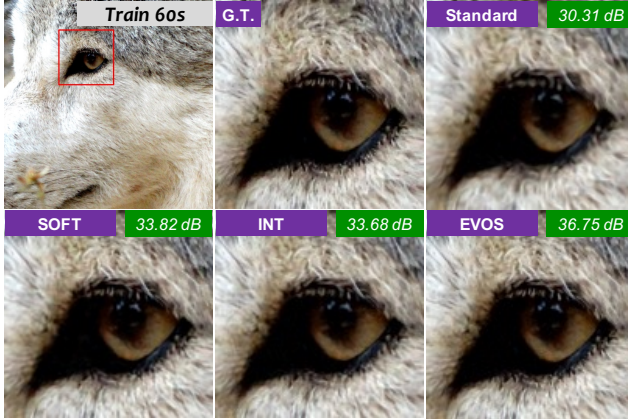


Figure 3. Visual comparison of sampling-based acceleration methods for INR training. All methods are evaluated under identical conditions with a fixed training duration of 60 seconds.

Analysis of INT. INT, which closely aligns with our application context, employs a greedy algorithm to select high-priority coordinates. While its best-performing variant INT(incre.)[†] demonstrates improvements over standard training (31% time savings, 0.07 dB PSNR gain), it exhibits unstable reconstruction quality, as shown in Fig. 4. The performance degradation in early stages stems from abrupt reductions in MLP evaluation frequency inherent to its greedy algorithm. EVOS circumvents this limitation through our sparse fitness evaluation design. Moreover, our crossover and mutation mechanisms achieve superior performance compared to INT’s greedy selection approach.

Sample Weighting ALSO Benefits “Overfitting”. An intuitive assumption suggests that sparsifying training samples would compromise per-iteration reconstruction quality while reducing computational cost, due to incomplete data utilization. However, the results of EVOS (w/o CFS) in Table 1 challenge this intuition, demonstrating that our strategy achieves superior fitting quality compared to full-data training under identical iteration counts. This phenomenon can be understood through the lens of sample weighting [14, 57], a technique that improves model generality by adjusting sample observation frequencies during training. Indeed, EVOS can be viewed as a specialized form of sample weighting that reweights signal coordinates during signal fitting, implicitly regularizing the loss function through selective sampling. Notably, this finding extends the benefits of sample weighting beyond traditional model generalization to signal fitting tasks (inherently an overfitting scenario without test set validation), revealing an intriguing contradiction that merits further investigation.

4.3. Compatibility for Different Backbones

To evaluate EVOS’s compatibility, we conducted experiments across multiple INR architectures: PEMLP [46],




Figure 4. Comparison of reconstruction quality (PSNR & SSIM) across different sampling-based methods over training time.

Backbones	PSNR			SSIM	Time
	(1k) \uparrow	(2k) \uparrow	(5k) \uparrow	(5k) \uparrow	(sec) \downarrow
PEMLP [46]	25.99	27.61	29.43	0.841	148.84
PEMLP+EVOS	26.10	28.00	30.04	0.842	83.18
SIREN [43]	30.37	33.35	36.24	0.960	177.26
SIREN+EVOS	32.12	35.10	36.84	0.961	100.43
WIRE [39]	28.96	31.17	33.72	0.918	553.31
WIRE+EVOS	28.59	32.91	34.56	0.927	291.30
GAUSS [35]	30.38	32.83	35.48	0.945	244.91
GAUSS+EVOS	30.21	34.36	35.48	0.941	133.28
FINER [23]	36.93	38.67	40.72	0.978	260.81
FINER+EVOS	36.49	39.54	40.61	0.976	123.60

Table 3. Quantitative result across different backbones.

WIRE [39], GAUSS [35], and FINER [23]. These architectures represent diverse approaches to enhancing INR expressiveness through innovations in frequency encoding and activation functions. Following [23], we implemented a consistent learning rate scheduler across all backbones while maintaining their official initial learning rates. The quantitative results in Table 3 demonstrate EVOS’s consistent acceleration capabilities across different architectures, showing substantial efficiency gains and performance improvements. In most cases (highlighted in green), EVOS achieves simultaneous improvements in both efficiency and reconstruction quality. Fig. 5 illustrates the reconstruction quality across different backbones after 120 seconds of training. EVOS integration consistently enhances detail preservation across all architectures, particularly in the fine features of the parrot’s eye region.

4.4. Compatibility for Different Network Sizes

To evaluate EVOS’s compatibility with varying network architectures, we tested its performance across MLP configurations ranging from 1×64 to 3×256 . Table 4 presents primary results, demonstrating EVOS’s effectiveness in improving training efficiency across diverse architectures,




Figure 5. Visualization of performance improvements across different backbone architectures when integrated EVOS. All experiments are conducted under consistent conditions with a fixed training duration of 120 seconds.

Sizes	PSNR				Time (sec) ↓
	(1k) ↑	(2k) ↑	(5k) ↑	(10k) ↑	
1×64	20.60	21.49	22.31	22.58	47.71
1×64 (EVOS)	20.65	21.53	22.26	22.54	37.28
2×64	22.57	23.71	24.57	24.79	72.19
2×64 (EVOS)	22.72	23.85	24.56	24.77	47.18
2×128	25.29	26.97	28.33	28.82	114.04
2×128 (EVOS)	25.68	27.29	28.40	28.83	74.59
2×256	28.25	30.30	32.27	33.30	262.46
2×256 (EVOS)	29.05	31.09	32.51	33.49	144.98
3×256	31.06	34.34	37.10	38.60	366.30
3×256 (EVOS)	32.71	36.09	37.81	38.71	254.33

Table 4. Quantitative result across different network sizes.

with efficiency gains becoming more pronounced as network size increases. From these observations, we draw two key conclusions: **(1)** EVOS demonstrates strong generalization across network architectures, enabling efficient training acceleration for varying INR expressiveness requirements. **(2)** The effectiveness of EVOS increases with network size, suggesting that larger architectures exhibit greater redundancy which our method effectively addresses.

4.5. Ablation Study

Settings. We conducted comprehensive ablation studies to evaluate the quantitative contributions of EVOS components: Sparse Fitness Evaluation (Eval.), Frequency-Guided Crossover (Cross.), and Augmented Unbiased Mutation (Mutat.). Beyond component-level analysis, we examined key design parameters including the interval parameter τ and linear increasing coefficient θ in Eq. 2, Cross-Frequency Supervision (CFS, Eq. 6), and mutation ratio α in Eq. 7. All experiments followed the settings in Sec. 4.2, varying only the parameter under investigation.

Results and Analysis. Table 5 presents the result of ablation experiment. The absence of Eval. (Sec. 3.2) results in

Settings	PSNR		SSIM (5k) ↑	Time (sec) ↓
	(1k) ↑	(5k) ↑		
EVOS	32.71	37.81	0.962	97.39
w/o Eval. (Sec. 3.2)	20.80	35.41	0.952	93.41
$\theta = 0$ (Eq. 2)	32.69	37.54	0.961	95.26
$\tau = 1$ (Eq. 2)	31.57	35.51	0.943	142.39
$\tau = 10$ (Eq. 2)	32.42	36.37	0.953	105.23
w/o Cross. (Sec. 3.3)	31.67	37.04	0.956	94.35
w/o CFS. (Eq. 6)	31.64	37.49	0.960	95.43
w/o Mutat. (Sec. 3.4)	31.18	34.94	0.921	97.13
$\alpha = 0.1$ (Eq. 7)	32.33	37.11	0.953	96.64
$\alpha = 1.0$ (Eq. 7)	32.10	37.14	0.958	101.51

Table 5. Ablation study. Coral: the worst performance.

unguided selection, significantly degrading early-stage performance. Without Mutat. (Sec. 3.4), selection bias from cached selections substantially impairs final performance. Components from Cross. (Sec. 3.3), CFS Eq. 6, and the linear increasing indicator (Eq. 2) enhance performance with minimal computational overhead. Interestingly, increasing evaluation frequency (reducing τ to 1 or 10) does not improve performance despite higher computational costs. We suspect that overly frequent evaluation induces excessive training instability, making it difficult for INRs to consistently capture signal features during the training process.

5. Conclusion

In this paper, we propose EVOLutionary Selector (EVOS) to accelerate INR training. Extensive experiments demonstrate EVOS’s superior performance, surpassing existing sampling-based acceleration methods. Notably, we discover that strategic sparsification of training samples not only reduces computational cost but also improves per-iteration performance, suggesting that sample weighting techniques can enhance optimization in signal fitting tasks beyond their traditional role in model generalization.

Acknowledgments

This work is supported in part by National Key Research and Development Project of China (Grant No. 2023YFF0905502), National Natural Science Foundation of China (Grant No. 92467204 and 62472249), and Shenzhen Science and Technology Program (Grant No. JCYJ20220818101014030 and KJZD20240903102300001). We thank anonymous reviewers for their valuable advice.

References

- [1] 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 workshops*, pages 126–135, 2017. 5
- [2] Hmrishav Bandyopadhyay, Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, Tao Xiang, Timothy M. Hospedales, and Yi-Zhe Song. Sketchinr: A first look into sketches as implicit neural representations. *CoRR*, abs/2403.09344, 2024. 1, 2
- [3] Zhicheng Cai, Hao Zhu, Qiu Shen, Xinran Wang, and Xun Cao. Batch normalization alleviates the spectral bias in coordinate networks. In *CVPR*, pages 25160–25171. IEEE, 2024. 3
- [4] Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 608–625. Springer, 2020. 2
- [5] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, pages 333–350. Springer, 2022. 1, 2, 3
- [6] Chia-Hao Chen, Ying-Tian Liu, Zhifei Zhang, Yuan-Chen Guo, and Song-Hai Zhang. Joint implicit neural representation for high-fidelity and compact vector fonts. In *ICCV*, pages 5515–5525. IEEE, 2023. 2
- [7] Hao Chen, Bo He, Hanyu Wang, Yixuan Ren, Ser-Nam Lim, and Abhinav Shrivastava. Nerv: Neural representations for videos. In *NeurIPS*, pages 21557–21568, 2021. 2
- [8] Xiang Chen, Jinshan Pan, and Jiangxin Dong. Bidirectional multi-scale implicit neural representations for image deraining. In *CVPR*, pages 25627–25636. IEEE, 2024. 1
- [9] Yinbo Chen and Xiaolong Wang. Transformers as meta-learners for implicit neural representations. In *ECCV (17)*, pages 170–187. Springer, 2022. 1, 2
- [10] Yann N. Dauphin, Razvan Pascanu, Çağlar Gülçehre, KyungHyun Cho, Surya Ganguli, and Yoshua Bengio. Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. In *NIPS*, pages 2933–2941, 2014. 2
- [11] Emilien Dupont, Adam Goliński, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. Coin: Compression with implicit neural representations. *arXiv preprint arXiv:2103.03123*, 2021. 2
- [12] Emilien Dupont, Hyunjik Kim, S. M. Ali Eslami, Danilo Jimenez Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you can treat it like one. In *ICML*, pages 5694–5725. PMLR, 2022. 1, 2
- [13] Zhenbiao Gai, Zhenyang Liu, Min Tan, Jiajun Ding, Jun Yu, Mingzhao Tong, and Junqing Yuan. Egra-nerf: Edge-guided ray allocation for neural radiance fields. *Image and Vision Computing*, 134:104670, 2023. 1, 3, 5, 6
- [14] Yue He, Pengfei Tian, Renzhe Xu, Xinwei Shen, Xingxuan Zhang, and Peng Cui. Model-agnostic random weighting for out-of-distribution generalization. In *KDD*, pages 1050–1061. ACM, 2024. 7
- [15] John H Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press, 1992. 1
- [16] Langwen Huang and Torsten Hoefler. Compressing multidimensional weather and climate data into neural networks. In *ICLR*. OpenReview.net, 2023. 1, 2
- [17] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, and Qing Wang. Inverting the imaging process by learning an implicit camera model. In *CVPR*, pages 21456–21465. IEEE, 2023. 1, 2
- [18] Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Accelerating neural field training via soft mining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20071–20080, 2024. 2, 3, 5, 6
- [19] Subin Kim, Sihyun Yu, Jaeho Lee, and Jinwoo Shin. Scalable neural video representations with learnable positional features. In *NeurIPS*, 2022. 2
- [20] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H. Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *CVPR*, pages 8456–8465. IEEE, 2023. 2
- [21] Ke Liu, Feng Liu, Haishuai Wang, Ning Ma, Jiajun Bu, and Bo Han. Partition speeds up learning implicit neural representations based on exponential-increase hypothesis. In *ICCV*, pages 5451–5460. IEEE, 2023. 1, 2
- [22] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *NeurIPS*, 2020. 3
- [23] 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. *CoRR*, abs/2312.02434, 2023. 2, 5, 7
- [24] Shishira R. Maiya, Sharath Girish, Max Ehrlich, Hanyu Wang, Kwot Sin Lee, Patrick Poirson, Pengxiang Wu, Chen Wang, and Abhinav Shrivastava. NIRVANA: neural implicit representations of videos with adaptive networks and autoregressive patch-wise modeling. In *CVPR*, pages 14378–14387. IEEE, 2023. 2
- [25] Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: adaptive coordinate networks for neural scene representation. *ACM Trans. Graph.*, 40(4):58:1–58:13, 2021. 1, 2

- [26] Ishit Mehta, Michaël Gharbi, Connelly Barnes, Eli Shechtman, Ravi Ramamoorthi, and Manmohan Chandraker. Modulated periodic activations for generalizable local functional representations. In *ICCV*, pages 14194–14203. IEEE, 2021. 3
- [27] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV (1)*, pages 405–421. Springer, 2020. 1, 2, 3
- [28] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 1, 2, 3
- [29] Trung Thanh Nguyen, Shengxiang Yang, and Juergen Branke. Evolutionary dynamic optimization: A survey of the state of the art. *Swarm and Evolutionary Computation*, 6:1–24, 2012. 2
- [30] Dogyun Park, Sihyeon Kim, Sojin Lee, and Hyunwoo J. Kim. DDMI: domain-agnostic latent diffusion models for synthesizing high-quality implicit neural representations. *CoRR*, abs/2401.12517, 2024. 2
- [31] Jeong Joon Park, Peter R. Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, pages 165–174. Computer Vision Foundation / IEEE, 2019. 1, 2
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8024–8035, 2019. 5
- [33] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred A. Hamprecht, Yoshua Bengio, and Aaron C. Courville. On the spectral bias of neural networks. In *ICML*, pages 5301–5310. PMLR, 2019. 1, 2
- [34] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, 378:686–707, 2019. 1, 2
- [35] Sameera Ramasinghe and Simon Lucey. Beyond periodicity: Towards a unifying framework for activations in coordinate-mlps. In *ECCV (33)*, pages 142–158. Springer, 2022. 2, 5, 7
- [36] Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. In *CVPR*, pages 14153–14161. Computer Vision Foundation / IEEE, 2021. 2
- [37] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *ICCV*, pages 14315–14325. IEEE, 2021. 2
- [38] Vishwanath Saragadam, Jasper Tan, Guha Balakrishnan, Richard G. Baraniuk, and Ashok Veeraraghavan. MINER: multiscale implicit neural representation. In *ECCV (23)*, pages 318–333. Springer, 2022. 1, 2
- [39] Vishwanath Saragadam, Daniel LeJeune, Jasper Tan, Guha Balakrishnan, Ashok Veeraraghavan, and Richard G. Baraniuk. WIRE: wavelet implicit neural representations. In *CVPR*, pages 18507–18516. IEEE, 2023. 2, 5, 7
- [40] Junwon Seo, Sangyoon Lee, Kwang In Kim, and Jaeho Lee. In search of a data transformation that accelerates neural field training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [41] Kexuan Shi, Hai Chen, Leheng Zhang, and Shuhang Gu. Inductive gradient adjustment for spectral bias in implicit neural representations, 2024. 3
- [42] Kexuan Shi, Xingyu Zhou, and Shuhang Gu. Improved implicit neural representation with fourier bases reparameterized training. *CoRR*, abs/2401.07402, 2024. 3
- [43] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NeurIPS*, 2020. 1, 2, 5, 7
- [44] Yannick Strümpfer, Janis Postels, Ren Yang, Luc Van Gool, and Federico Tombari. Implicit neural representations for image compression. In *ECCV (26)*, pages 74–91. Springer, 2022. 1, 2
- [45] Filip Szatkowski, Karol J. Piczak, Przemyslaw Spurek, Jacek Tabor, and Tomasz Trzcinski. Hypernetworks build implicit neural representations of sounds. In *ECML/PKDD (4)*, pages 661–676. Springer, 2023. 1, 2
- [46] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *NeurIPS*, 2020. 1, 2, 5, 7
- [47] Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P. Srinivasan, Jonathan T. Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations. In *CVPR*, pages 2846–2855. Computer Vision Foundation / IEEE, 2021. 1, 2
- [48] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4): 600–612, 2004. 5
- [49] Shaowen Xie, Hao Zhu, Zhen Liu, Qi Zhang, You Zhou, Xun Cao, and Zhan Ma. DINER: disorder-invariant implicit neural representation. In *CVPR*, pages 6143–6152. IEEE, 2023. 1, 2, 3
- [50] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *CVPR*, pages 5428–5438. IEEE, 2022. 3
- [51] Runzhao Yang. TINC: tree-structured implicit neural compression. In *CVPR*, pages 18517–18526. IEEE, 2023. 3
- [52] Zhi-Hui Zhan, Lin Shi, Kay Chen Tan, and Jun Zhang. A survey on evolutionary computation for complex continuous optimization. *Artificial Intelligence Review*, 55(1):59–110, 2022. 2

- [53] Chen Zhang, Steven Tin Sui Luo, Jason Chun Lok Li, Yik-Chung Wu, and Ngai Wong. Nonparametric teaching of implicit neural representations. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. [1](#), [2](#), [3](#), [5](#), [6](#)
- [54] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595. Computer Vision Foundation / IEEE Computer Society, 2018. [5](#)
- [55] Weixiang Zhang, Shuzhao Xie, Shijia Ge, Wei Yao, Chen Tang, and Zhi Wang. Expansive supervision for neural radiance field. *arXiv preprint arXiv:2409.08056*, 2024. [1](#), [3](#), [5](#), [6](#)
- [56] Weixiang Zhang, Shuzhao Xie, Chengwei Ren, Shijia Ge, Mingzi Wang, and Zhi Wang. Enhancing implicit neural representations via symmetric power transformation, 2024. [3](#)
- [57] Xingxuan Zhang, Peng Cui, Renzhe Xu, Linjun Zhou, Yue He, and Zheyang Shen. Deep stable learning for out-of-distribution generalization. In *CVPR*, pages 5372–5382. Computer Vision Foundation / IEEE, 2021. [7](#)
- [58] Tianle Zhong, Jiechen Zhao, Xindi Guo, Qiang Su, and Geoffrey Fox. RINAS: training with dataset shuffling can be general and fast. *CoRR*, abs/2312.02368, 2023. [2](#), [5](#)