

# A Structured Dictionary Perspective on Implicit Neural Representations

Gizem Yüce\*  
ETH Zurich

gyuece@ethz.ch

Guillermo Ortiz-Jiménez\* Beril Besbinar Pascal Frossard  
École Polytechnique Fédérale de Lausanne (EPFL)

{name.surname}@epfl.ch

## Abstract

Implicit neural representations (INRs) have recently emerged as a promising alternative to classical discretized representations of signals. Nevertheless, despite their practical success, we still do not understand how INRs represent signals. We propose a novel unified perspective to theoretically analyse INRs. Leveraging results from harmonic analysis and deep learning theory, we show that most INR families are analogous to structured signal dictionaries whose atoms are integer harmonics of the set of initial mapping frequencies. This structure allows INRs to express signals with an exponentially increasing frequency support using a number of parameters that only grows linearly with depth. We also explore the inductive bias of INRs exploiting recent results about the empirical neural tangent kernel (NTK). Specifically, we show that the eigenfunctions of the NTK can be seen as dictionary atoms whose inner product with the target signal determines the final performance of their reconstruction. In this regard, we reveal that meta-learning has a reshaping effect on the NTK analogous to dictionary learning, building dictionary atoms as a combination of the examples seen during meta-training. Our results permit to design and tune novel INR architectures, but can also be of interest for the wider deep learning theory community.

## 1. Introduction

Implicit neural representations (INRs) have recently emerged as a powerful alternative to classical, discretized, representations of multimedia signals [9, 13, 30, 31, 39, 47, 48, 52, 53]. In contrast to traditional methods, INRs parameterize the continuous mapping between coordinates and signal values using neural networks. This allows for an efficient and compact representation of signals that can be easily integrated into modern differentiable learning pipelines.

The recent success of INRs in many applications, such as surface representation [47], volume rendering [27, 31, 40,

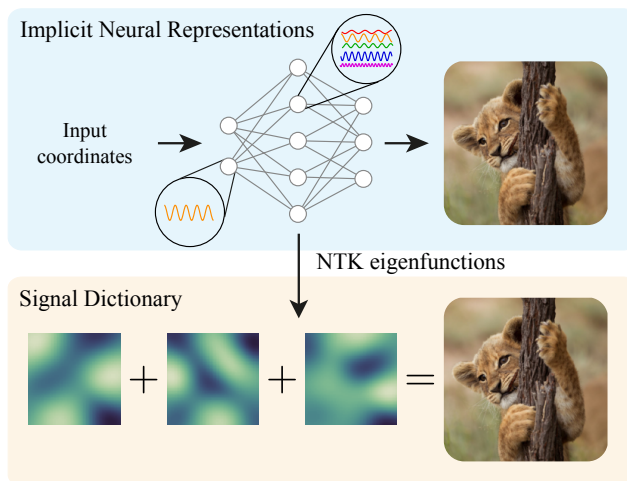


Figure 1. Conceptual illustration of our main theoretical contributions: i) Each layer of an INR increases the frequency support of the representation by splitting a signal into higher order harmonics. ii) INRs can be interpreted as signal dictionaries whose atoms are the eigenfunctions of their NTK at initialization.

50] or generative modelling [7, 14] can be largely attributed to the development of new periodic representations that can circumvent the spectral bias of standard neural networks. Indeed, there is ample evidence that the use of periodic representations [22, 31, 48, 52] can mitigate the bias of standard architectures towards low frequency [43].

Nevertheless, even if INRs have become widely adopted in practice, the theoretical understanding of their principles and properties is rather limited. For example, there is no clear criterion to select between different INR families, their parameters are mostly based on heuristics, and their limitations are not well understood. These shortcomings are slowing down further research developments. In this work, we therefore take a step back and focus on understanding the mechanisms behind the success of modern INRs, but also their failure modes, in order to develop more informed design strategies. We provide a unified perspective with the aim to answer the following important questions:

1. *What is the expressive power of INRs?*

\*The first two authors contributed equally to this work.

## 2. How does initialization affect their inductive bias?

Specifically, we first leverage results from harmonic analysis and deep learning theory, and we discover that the expressive power of most INRs is equivalent to that of a structured signal dictionary whose atoms are integer harmonics of the frequencies that define their initial input mapping (see Fig. 1). This unifies many INR architectures under a single perspective, and can serve to understand them better and mitigate some of their common problems.

Then, we delve deeper on the inductive bias of INRs. We build upon the foundational work in [52], and exploit recent results in deep learning theory to give a new unifying framework to analyse the inductive bias of any INR architecture in terms of its empirical neural tangent kernel (NTK) [20]. In particular, we reveal the existence of a close analogy between the eigenfunctions of the empirical NTK and the atoms of a signal dictionary, and show that the difficulty of learning a signal with an INR is intimately connected to how efficiently it can be encoded by this dictionary.

Finally, we use our novel perspective to explain the role of meta-learning in improving the performance of INRs. INRs are known to be notoriously inefficient, requiring long training times, and a large sample exposure to achieve good results, especially in 3D settings [17, 18, 45]. However, recent works have shown that using meta-learning algorithms to initialize INRs can greatly improve their speed of convergence and sample complexity [46, 51]. In this work, we show that meta-learning works as a dictionary learning algorithm, transforming the NTK of an INR into a rich signal dictionary whose atoms are formed by combinations of the examples seen during meta-training. This increases the representation efficiency of the target signals by the NTK [37], thus improving performance and training speed.

In summary, the main contributions of our work are:

- We provide a unified perspective to theoretically analyze the expressive power and inductive bias of INRs.
- We show that the frequency support of INRs grows exponentially with depth, as each layer splits its input into higher order harmonics, demonstrating their efficiency in representing wide spectrum signals.
- We use this theory to explain the main failure modes of INRs: imperfect recovery and aliasing.
- We show that the inductive bias of INRs can be characterized by the ability of their empirical NTKs to encode different target signals efficiently.
- Finally, we discover that meta-learning greatly increases the encoding efficiency of the NTK by constructing a rich signal dictionary using different combinations of the meta-training tasks.

Overall, we believe that our findings can impact the future research in INRs and their applications, and contribute

to speeding up the development of new principled algorithms in the field. It gives a fresh perspective to understand and alleviate the drawbacks of the current architectures, as well as new intuitions to design better INR algorithms. Finally, our analysis on the effect of meta-learning can also be of broader interest for the deep learning theory community<sup>1</sup>.

## 2. Implicit Neural Representations

The goal of an implicit neural representation is to encode a continuous target signal  $g : \mathbb{R}^D \rightarrow \mathbb{R}^C$  using a neural network  $f_{\theta} : \mathbb{R}^D \rightarrow \mathbb{R}^C$ , parameterized by a set of weights  $\theta \in \mathbb{R}^N$ , by representing the mapping between input coordinates  $\mathbf{r} \in \mathbb{R}^D$ , e.g., pixels, and signal values  $g(\mathbf{r}) \in \mathbb{R}^C$ , e.g., RGB colors. This is achieved minimizing a distortion measure, like mean-squared error, during training using some form of (stochastic) gradient descent (SGD).

The continuous parameterization of INRs allows to store signals at a constant memory cost regardless of the spatial resolution, which makes INRs stand out for reconstructing high-dimensional signals, such as videos or 3D scenes. The main challenge for INRs, though, is to reconstruct high frequency details present in most multimedia signals, e.g., textures in images. Classical neural network architectures are well-known for their strong spectral bias towards lower frequencies [43], and this has made them traditionally useless for implicit representation tasks. Recently, however, few works [47, 52] have come up with different solutions to circumvent the spectral bias of neural networks, allowing faster convergence and greater fidelity of INRs.

In what follows, we provide an overview of the main solutions under a unified *architecture* formulation. Specifically, we note that most INR architectures can be decomposed into a mapping function  $\gamma : \mathbb{R}^D \rightarrow \mathbb{R}^T$  followed by a multilayer perceptron (MLP), with weights  $\mathbf{W}^{(\ell)} \in \mathbb{R}^{F_{\ell-1} \times F_{\ell}}$ , bias  $\mathbf{b}^{(\ell)} \in \mathbb{R}^{F_{\ell}}$ , and activation function  $\rho^{(\ell)} : \mathbb{R} \rightarrow \mathbb{R}$ , applied elementwise; at each layer  $\ell = 1, \dots, L - 1$ . That is, if we denote by  $\mathbf{z}^{(\ell)}$  each layers post activation, most INR architectures compute

$$\begin{aligned} \mathbf{z}^{(0)} &= \gamma(\mathbf{r}), \\ \mathbf{z}^{(\ell)} &= \rho^{(\ell)} \left( \mathbf{W}^{(\ell)} \mathbf{z}^{(\ell-1)} + \mathbf{b}^{(\ell)} \right), \quad \ell = 1, \dots, L - 1 \\ f_{\theta}(\mathbf{r}) &= \mathbf{W}^{(L)} \mathbf{z}^{(L-1)} + \mathbf{b}^{(L)}. \end{aligned} \tag{1}$$

We now examine the two most popular INR architectures:

**Fourier feature networks (FFNs)** In [52], Tancik *et al.* proposed to use a Fourier mapping  $\gamma(\mathbf{r}) = \sin(\mathbf{\Omega}\mathbf{r} + \phi)$ , with parameters  $\mathbf{\Omega} \in \mathbb{R}^{T \times D}$  and  $\phi \in \mathbb{R}^T$  followed by an MLP with  $\rho^{(\ell)} = \text{ReLU}$ . Specifically, they showed that initializing  $\Omega_{i,j} \sim \mathcal{N}(0, \sigma^2)$  with random Fourier features [44] can modulate the spectral bias of an FFN, with

<sup>1</sup>Code to reproduce this work: [github.com/gortizji/inr.dictionaries](https://github.com/gortizji/inr.dictionaries)

larger values of  $\sigma$  biasing these networks towards higher frequencies. Alternative formulations with deterministic initialization, commonly used for neural rendering algorithms [31] can be considered as a special category of these networks where the frequencies in  $\Omega$  are taken to be powers of 2 and the frequencies in  $\phi$  alternate between  $\{0, \pi/2\}$ .

**Sinusoidal representation networks (SIRENs)** In [47], Sitzmann *et al.* proposed to use MLP with sinusoidal activations, i.e.,  $\rho^{(\ell)} = \sin$ , where the first layer post activation,  $z^{(0)} = \sin(\omega_0(\mathbf{W}^{(0)}\mathbf{r} + \mathbf{b}^{(0)}))$  can be interpreted as  $\gamma(\mathbf{r}) = \sin(\Omega\mathbf{r} + \phi)$ . They showed that, by rescaling the parameters at initialization by the constant factor  $\omega_0$ , SIRENs can also modulate the spectral bias, with larger  $\omega_0$  biasing these networks towards higher frequencies.

Nonetheless, despite the ample empirical evidence that shows that these architectures are effective at representing natural images or other visual signals, there is little theoretical understanding of how they do so. Moreover, since the design of each of these networks is guided by very different principles, the sheer diversity in the structure of these architectures makes their analysis very involved.

In the next sections, we provide a unified perspective to analyze the expressive power and inductive bias of INRs and show that all modern INRs are intrinsically guided by the same fundamental principles, which let them express a wide range of signals. However, it also makes them prone to the same type of failure modes. Our novel framework can be used to design new principled solutions to address these shortcomings, but also simplify the tuning of current INRs.

### 3. Expressive Power of INRs

We now provide an integrated analysis of the expressive power of INRs. To that end, we will follow the formulation in Eq. (1), where, to simplify our derivations, we will restrict ourselves to polynomial activation functions, i.e., nonlinearities of the form  $\rho(x) = \sum_{k=0}^K \alpha_k x^k$ . Note that this is a very mild assumption, as all analytic activation functions, e.g., sinusoids, can be approximated using polynomials with a naïve Taylor expansion; and that even the non-differentiable ReLUs can be effectively approximated using Chebyshev polynomials [28]. Note, also, that the sequence of coefficients of the polynomial expansion of most activation functions used in practice decays very rapidly [28].

Now, without loss of generality, let  $D = 1$  and consider what happens when a single-frequency mapping, i.e.  $\gamma(r) = e^{j\omega r}$ , goes through such a polynomial activation: The output of the activation consists of a linear combination of the integer harmonics of the input frequency, i.e.,

$$\rho(\gamma(r)) = \rho(e^{j\omega r}) = \sum_{k=0}^K \alpha_k e^{jk\omega r}. \quad (2)$$

This harmonic expansion is precisely the mechanism that

controls the frequency representation in INRs. More generally, the mapping  $\gamma(\mathbf{r})$  acts as a collection of single frequency basis, whose spectral support is expanded after each non-linear activation into a collection of higher order harmonics. This particular structure is shared among all FFNs and SIRENs and it gives rise to the following result regarding their expressive power, i.e. the class of functions that can be represented with these architectures.

**Theorem 1.** *Let  $f_{\theta} : \mathbb{R}^D \rightarrow \mathbb{R}$  be an INR of the form of Eq. (1) with  $\rho^{(\ell)}(z) = \sum_{k=0}^K \alpha_k z^k$  for  $\ell > 1$ . Furthermore, let  $\Omega = [\Omega_0, \dots, \Omega_{T-1}]^T \in \mathbb{R}^{T \times D}$  and  $\phi \in \mathbb{R}^T$  denote the matrix of frequencies and vector of phases, respectively, used to map the input coordinate  $\mathbf{r} \in \mathbb{R}^D$  to  $\gamma(\mathbf{r}) = \sin(\Omega\mathbf{r} + \phi)$ . This architecture can only represent functions of the form*

$$f_{\theta}(r) = \sum_{\omega' \in \mathcal{H}(\Omega)} c_{\omega'} \sin(\langle \omega', \mathbf{r} \rangle + \phi_{\omega'}), \quad (3)$$

where

$$\mathcal{H}(\Omega) \subseteq \left\{ \sum_{t=0}^{T-1} s_t \Omega_t \mid s_t \in \mathbb{Z} \wedge \sum_{t=0}^{T-1} |s_t| \leq K^{L-1} \right\}. \quad (4)$$

*Proof.* See Appendix.  $\square$

Thm. 1 shows that the expressive power of FFNs and SIRENs is restricted to functions that can be expressed as a linear combination of certain harmonics of the feature mapping  $\gamma(\mathbf{r})$ . That is, these architectures have the same expressive power as a structured signal dictionary whose atoms are sinusoids with frequencies equal to sums and differences of the integer harmonics of the mapping frequencies<sup>2</sup>. Interestingly, an analogous result was also proven for the Multiplicative Filter Networks (MFNs) [15], a proof-of-concept architecture based on a multiplicative connection between layers instead of the usual compositional structure of MLPs. In particular, it can be shown that MFNs, although very different in structure, are also only able to express linear combinations of certain harmonics of their sinusoidal filters [15], which means that they have the same expressive power as FFNs and SIRENs.

Besides this unification, Thm. 1 also highlights that the way all these architectures encode different signals is very similar. Indeed, instead of representing a signal by directly learning the coefficients of the linear combination, which would require to store  $\mathcal{O}(TK^L)$  coefficients  $c_{\omega'}$ ; the multi-layer structure of all INRs imposes a certain low rank structure over the coefficients – akin to the sparsity assumption in classical dictionaries [54] – which can greatly save on memory as it only requires to store  $\mathcal{O}(T^2L)$  parameters. This is better understood through an illustrative example<sup>3</sup>.

<sup>2</sup>We will refer to these components as the harmonics of  $\gamma(\mathbf{r})$ .

<sup>3</sup>Similar examples for other architectures can be found in the Appendix.

**Example.** Let  $f_\theta$  be a three-layer SIREN defined as

$$f_\theta(r) = \mathbf{w}^{(2)\top} \sin\left(\mathbf{W}^{(1)} \sin(\Omega r)\right), \quad (5)$$

where  $\Omega \in \mathbb{R}^T$ ,  $\mathbf{W}^{(1)} \in \mathbb{R}^{F \times T}$ , and  $\mathbf{w}^{(2)} \in \mathbb{R}^F$ . The output of this network can equivalently be represented as

$$f_\theta(r) = \sum_{m=0}^{F-1} \sum_{s_1, \dots, s_T = -\infty}^{\infty} c_{m, s_1, \dots, s_T} \sin\left(\sum_{t=0}^{T-1} s_t \omega_t r\right), \quad (6)$$

where

$$c_{m, s_1, \dots, s_T} = \left( \prod_{t=0}^{T-1} J_{s_t} \left( W_{m,t}^{(1)} \right) \right) w_m^{(2)}, \quad (7)$$

and  $J_s$  denotes the Bessel function of first kind of order  $s$ .

*Proof.* See Appendix  $\square$

As we can see, the harmonic expansion introduced by the nested sinusoids of this simple SIREN can be developed into a signal with a very large bandwidth. Indeed, the few coefficients of this network are enough to represent a signal supported by an infinite number of frequency harmonics.

On the other hand, note that composing sinusoids is a common operation in communication theory as it defines the basis of frequency modulation (FM) technology [42]. Interestingly, drawing analogies between FM signals and SIRENs is a good source of inspiration to intuitively understand how these networks modulate their spectral bias: Recall that for FM signals, such as  $\sin(\beta \sin(\omega_0 r))$ , the parameter  $\beta$  controls the bandwidth of the modulation, which is generally limited by the decreasing nature of the Bessel coefficients  $J_n(\beta)$  in  $n$ . Increasing  $\beta$  has the effect of expanding the spectral support of the modulation, as the arguments of the Bessel functions increase.

The analogous phenomenon can be observed in Eq. (6) for this simple SIREN, but can be extended to more general architectures. In general, we see that due to the decreasing nature of the Bessel functions  $J_{s_t}(W_{m,t}^{(1)})$ , the high order harmonics in Eq. (6) tend to have smaller weights than the lower ones. This specific parameterization acts as an implicit bias mechanism, which focuses most of the energy of the output signal in a narrow band around the input frequencies  $\Omega$ . Nevertheless, we can also see that increasing the scale of the coefficients in the inner layer, e.g.,  $\mathbf{W}^{(1)}$ , makes the coefficients of higher order terms in Eq. (7) larger, thus increasing the power of the higher order harmonics, and allowing the network to learn a wider range of frequencies.

The fact that all modern INRs encode information in a similar way can explain why all these architectures are as powerful, in practice. However, it may also explain why they all suffer from the same failure modes. In Sec. 4, we study these in more detail.

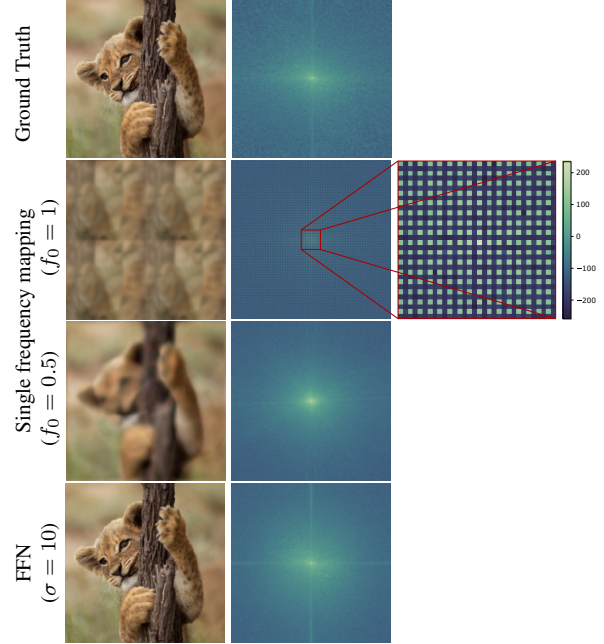


Figure 2. **Left** Image reconstruction with different mappings of the input coordinates. **Right:** Magnitude of the DFT of the reconstruction. The **FFN** uses random Fourier encodings as defined in Sec. 2, and the **single frequency mappings** correspond to  $\gamma(r) = [\cos(2\pi f_0 r), \sin(2\pi f_0 r)]^\top$ .

## 4. Failure modes of INRs

We now move on to study of the main failure modes of INRs. In particular, we will see how the specific harmonic expansion from Thm. 1 can sometimes lead to very recognizable artifacts in the learned reconstructions. Specifically, imperfect signal recovery and aliasing.

### 4.1. Imperfect recovery

One of the main consequences of Thm. 1 is that the set of frequencies that define the base embedding  $\gamma(r)$  completely determines the frequency support of the reconstruction  $f_\theta(r)$ . In this sense, it is fundamental to guarantee that the set  $\mathcal{H}(\Omega)$  permits to properly cover the spectrum of  $g(r)$ . When this is not the case, the reconstructed representations can exhibit severe artifacts in the spatial domain stemming from an incorrect choice of fundamental frequencies determined by the INR mapping in Eq. (1).

Let us illustrate this phenomenon for FFNs<sup>4</sup>, but note that other types of architectures, such as SIRENs, also can suffer from spatial artifacts<sup>5</sup>. To that end, let us take the extreme case of an FFN  $f_\theta : \mathbb{R}^2 \rightarrow \mathbb{R}^3$ , with a deterministic single frequency Fourier encoding  $\gamma(r) = [\sin(2\pi f_0 r), \cos(2\pi f_0 r)]^\top$ , reconstructing an image  $f : [-1, 1]^2 \rightarrow \mathbb{R}^3$ , from samples in a grid of  $N \times N$ .

<sup>4</sup>The details of all our experiments can be found in the Appendix.

<sup>5</sup>We replicate our experiments for other networks in the Appendix.

Now, note that, in light of Thm. 1, this network can only represent signals with a frequency support in  $\mathcal{H}(\Omega) \subseteq \{2k \cdot \pi f_0 | k \in \mathbb{Z}\}$ , i.e., containing only even multiples of  $\pi f_0$ . This means that if we choose  $f_0 = 1$ , the discrete Fourier transform (DFT) of the reconstruction will only have non-zero coefficients at frequencies corresponding to  $2k \cdot 2\pi/N$ , for  $k = 0, \dots, \lfloor (N-1)/2 \rfloor$ . This frequency covering is certainly not enough to completely represent images, as it misses all odd multiples of  $2\pi/N$ .

As shown in Fig. 2, reconstructing an image with such network produces severe artifacts. The learned representation with  $f_0 = 1$  is highly distorted. That is, we see multiple displaced versions of the target image imposed over each other. The nature of this artifact is much more clear when we inspect the DFT of the reconstruction, which is supported on a perfect grid in the spectral domain, missing all the values of the spectrum at the odd coefficients.

Strikingly, setting  $f_0 = 0.5$  is enough to completely get rid of this type of artifact. Indeed, when  $f = 0.5$  the set  $\mathcal{H}(\Omega) \subseteq \{\pi k | k \in \mathbb{Z}\}$ , which means that the DFT of the reconstruction can have energy in all spectral coefficients. Nonetheless, we also observe that the resulting image is quite blurry. As we will see, this is due to the fast decay of the polynomial coefficients in Eq. (2) for most activation functions, including ReLUs [28], which causes the weights of the high frequency harmonics in Eq. (3) to be very small. This phenomenon can be greatly alleviated, however, by increasing the frequency cover of the initial mapping  $\gamma(\mathbf{r}) = \sin(\Omega \mathbf{r} + \phi)$  and sampling  $\Omega \in \mathbb{R}^{D \times T}$  using  $\Omega_{i,j} \sim \mathcal{N}(0, \sigma^2)$ . Indeed, using a large  $T$  with a large  $\sigma$  can reduce the probability of having a limited representation of the frequency spectrum of the target signal. Nevertheless, as we will see in Sec. 4.2, setting  $\sigma$  too large can introduce other problems.

## 4.2. Aliasing

It has been empirically shown that INRs with high fundamental frequencies in  $\gamma(\mathbf{r})$  converge faster, and achieve higher performances in the training set [47, 52]; even for targets with high frequency details. Nevertheless, it has also been reported that initializing these frequencies too high leads to poor performance outside the exact support of the training set, and produces aliasing artifacts [5]. To the best of our knowledge, this behavior is still poorly understood.

Thm. 1 can, however, shed new light on this phenomenon. To that end, it is useful to see INRs as digital-to-analog converters (DAC), since INRs do little more than reconstruct a continuous signal from a set of discrete training samples. Classical sampling theory [36] guarantees that one can reconstruct a bandlimited signal from its samples provided the sampling frequency is above the Nyquist rate. Nevertheless, it also states that without this prior knowledge, the problem of reconstructing a continuous signal




Figure 3. Magnitude of the spectrum of  $g(r) = \sin(2\pi \cdot 23r)$  and its SIREN reconstruction trained at  $f_s = 128$  Hz. **Top row** shows  $\omega_0 = 300$ , and **bottom row**  $\omega_0 = 30$ . On the **left** the signals are sampled at  $f_s = 128$  Hz and on the **right** at  $f_s = 256$  Hz.

from its samples is, in general, ill-posed – there are many continuous functions that can lead to the exact same samples. Since INRs do not have an explicit knowledge of the bandwidth of the target, only their implicit bias can determine which of all these functions they reconstruct.

When the implicit bias does not match the nature of the signals, this can lead to reconstruction artifacts. Take for instance the problem of reconstructing a single-frequency signal  $g(r) = \sin(2\pi \cdot 23r)$  using a SIREN ( $\omega_0 = 300$  rad/s) trained on 128 evenly spaced samples in the range  $[0, 1]$ , i.e., sampled with a frequency of  $f_s = 128$  Hz. As we can see in Fig. 3, the discrete-time Fourier transform of the reconstruction at the training points perfectly matches the target signal, i.e., the training loss is zero. Surprisingly, though, if one reconstructs the signal on a finer grid, e.g.,  $f_s = 256$  Hz, which contains coordinates not seen during training, one can see that the spectrum of the reconstruction has an additional peak at 105 Hz that is not present in the target signal. That is, the implicit bias of the network has “chosen” to reconstruct the signal using an aliased higher frequency component, as it had no way to discard this feasible solution. Interestingly, if one initializes the SIREN using  $\omega_0 = 30$  rad/s, instead, this aliased copy disappears.

Thm. 1 gives the key to understand this behaviour. Specifically, note that most non-linearities used in INRs, e.g., ReLU or sin, can be effectively approximated by polynomials of small order, or with rapidly decaying coefficients. As a result, even if the frequency support of the INRs can include harmonics of very high frequencies, theoretically, those components tend to be weighted with much smaller coefficients in practice. Increasing the value of the fundamental frequencies does help to include higher frequency components without relying in very high order harmonics. However, it does so, at the cost of introducing high frequency components with large weights in Eq. (3), thus increasing the chances of yielding aliased reconstructions.

Reconstructing signals at low sampling rates makes the aliased high frequency components in Eq. (3) indistinguishable from lower frequency components. As we have seen this phenomenon stems from the underspecification [11] of the reconstruction of the reconstruction problem in INRs, which can yield aliasing artifacts when testing at higher sampling rates. Solving this issues is crucial in application where a certain degree of generalization is required from the INRs. Applications such as super-resolution [8, 21] or scene reconstruction [47] cannot rely on pure overfitting, and require INRs to generalize outside of their training support. Overall, we hope that our new insights can support the design of a new generation of INR architectures and algorithms that can mitigate this underspecification.

## 5. Inductive bias of INRs

All our results, so far, have only dealt with expressive power, i.e., the type of functions that can be represented by INRs. However, even if a network can express a signal, it does not mean that it can learn to represent it efficiently. MLPs, for instance, are widely known to be universal function approximators [10], but still they have a hard time learning to high frequency functions [43]. To the best of our knowledge, the inductive bias of INRs is a largely unexplored topic. Besides the fact that INRs can circumvent the spectral bias [47, 52], little is known of how different design choices influence the learnability of different signals.

In what follows, we will try to narrow this knowledge gap, as we will leverage recent results from deep learning theory to shed new light on the inductive bias of INRs, and how their initialization has a crucial role on what they learn.

### 5.1. Overview of NTK theory

Studying the inductive bias of deep learning is hard. This is mostly due to the non-linear nature of the mapping between parameters and functions specified by neural networks. Recent studies, however, have started arguing that studying learnability approximately is much more tractable. Notably, the neural tangent kernel (NTK) framework [20] proposes to approximate any neural network by its first order Taylor decomposition around the initialization  $\theta_0$ , i.e.,

$$f_{\theta}(\mathbf{r}) \approx f_{\theta_0}(\mathbf{r}) + (\theta - \theta_0)^{\top} \nabla_{\theta} f_{\theta_0}(\mathbf{r}), \quad (8)$$

since using this approximation, the network is reduced to a simple linear predictor defined by the kernel

$$\Theta(\mathbf{r}_1, \mathbf{r}_2) = \langle \nabla_{\theta} f_{\theta_0}(\mathbf{r}_1), \nabla_{\theta} f_{\theta_0}(\mathbf{r}_2) \rangle. \quad (9)$$

Remarkably, while the understanding of deep learning is still in its infancy, the learning theory of kernels is much more developed [49]. Specifically, it can be shown that using the kernel in Eq. (9), the sample complexity, and optimization difficulty, of learning a target function  $g$  grows proportionally to its kernel norm [6], i.e.,




Figure 4. Average energy concentration of 100 validation images from CelebA on subspaces spanned by the eigenfunctions of the empirical NTK associated to eigenvalues greater than a given threshold. Legend shows the average test PSNR after training to reconstruct those images from 50% randomly selected pixels.

$$\|g\|_{\Theta}^2 = \sum_{i=0}^{\infty} \frac{1}{\lambda_i} |\langle \phi_i, g \rangle|^2, \quad (10)$$

where  $\langle \phi_i, g \rangle = \mathbb{E}_{\mathbf{r}}[\phi_i(\mathbf{r})g(\mathbf{r})]$ , and  $\{\lambda_i, \phi_i\}_{i=0}^{\infty}$  denote the eigenvalue, eigenfunction pairs of the kernel given by its Mercer’s decomposition, i.e.,  $\Theta(\mathbf{r}_1, \mathbf{r}_2) = \sum_{i=0}^{\infty} \lambda_i \phi_i(\mathbf{r}_1) \phi_i(\mathbf{r}_2)$ . That is, those targets that are more concentrated in the span of the eigenfunctions associated with the largest eigenvalues of the kernel are easier to learn.

Eq. (8) holds with equality only if the neural network  $f_{\theta}$  is infinitely wide and has a specific structure [1, 20]. For the finite-size neural networks used in practice, it only provides a rough approximation. Fortunately, recent studies have shown that even if finite-size neural networks and their kernel approximations do not have exactly the same dynamics, their sample complexity when learning a target  $g$  scales in both cases with its kernel norm [37], which makes Eq. (10) a good proxy for learnability in deep learning.

### 5.2. NTK eigenfunctions as dictionary atoms

The fact that the empirical NTK can approximately capture learnability in deep learning leads to a new interpretation of INRs: we can view INRs as signal dictionaries whose atoms are given by the eigenfunctions of the NTK at initialization. In this view, the study of the inductive bias of an INR is equivalent to the study of the representation capabilities of its NTK dictionary, in the sense that the functions that can be efficiently encoded by this dictionary are the ones that will be easier to learn.

The simplicity of this analogy allows us to investigate phenomena that appear complex otherwise. For example, we can use this perspective to constructively characterize the effect of the parameter  $\omega_0$  in the inductive bias of a SIREN, and compare different networks, or initializations. To that end, we measure the average energy concentration<sup>6</sup> of  $N = 100$  validation images  $\{g_n\}_{n=1}^N$  from the CelebA

<sup>6</sup>Details of the experiments can be found in the Appendix.

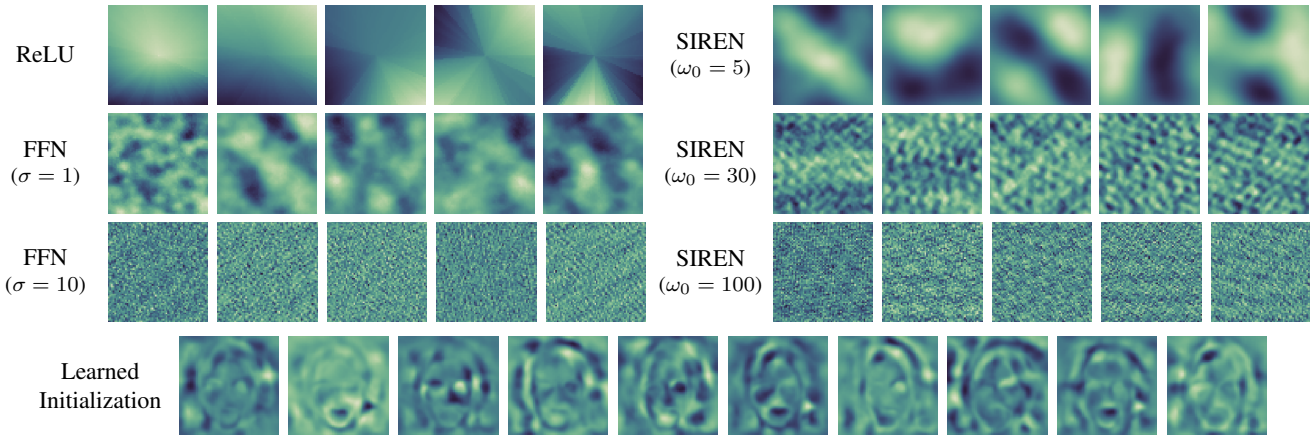


Figure 5. First eigenfunctions of the empirical NTK of different INRs at initialization. The first six architectures are initialized as described in Sec. 2. The **learned initialization** row shows the eigenfunctions of a SIREN initialized after meta-learning on 1,000 training images from the CelebA dataset [25] following the procedure described in [51]. Details of this experiment can be found in the Appendix.

dataset [25] on the span of the eigenfunctions of the NTK associated to eigenvalues greater than a given threshold, i.e.,

$$\mathcal{E}(\lambda) = \frac{1}{N} \sum_{n=1}^N \sum_{\lambda_i/\lambda_0 \geq \lambda} \frac{|\langle \phi_i, g_n \rangle|^2}{|\langle g_n, g_n \rangle|^2}. \quad (11)$$

This metric is intimately connected to the kernel norm in Eq. (10), and it can give us a convenient perspective of the region of the NTK spectrum that will represent an image. The results of this procedure applied to different networks are shown in Fig. 4. Remarkably, for very low values of  $\omega_0$ , most of the energy of these images is concentrated on the eigenfunctions corresponding to small eigenvalues. However, as we increase  $\omega_0$ , the energy concentration gets more skewed towards the eigenfunctions associated with large eigenvalues. Interestingly, after some point ( $\omega_0 > 40$ ), the energy profile starts receding to the right, again.

Comparing the energy profiles with the generalization performance of these networks, we observe a clear pattern: the more energy is concentrated on the eigenfunctions associated with larger eigenvalues, the better the test peak signal-to-noise ratio (PSNR)<sup>7</sup>. To understand this phenomenon, we can inspect the eigenfunctions of the NTK. As it is shown in Fig. 5, the eigenfunctions of the SIRENs with larger  $\omega_0$  have higher frequency content. This means that increasing  $\omega_0$  can have a positive effect in generalization as it yields a dictionary that better spans the medium-high frequency spectrum of natural images. Increasing  $\omega_0$  too much, on the other hand, yields atoms with an overly high frequency content that cannot span the space of natural images efficiently, which explains their poor reconstruction performance of these networks.

Overall, we see how interpreting learnability as encoding efficiency of the NTK dictionary is a powerful analogy that

<sup>7</sup>Correlations with other training metrics are shown in the Appendix.

can explain diverse phenomena, and lets us study under a single framework all sorts of INR questions, including those which might not be readily understood from Thm. 1. This is a very powerful tool that we further exploit in Sec. 5.3 to provide novel insights on the role of meta-learning in INRs.

### 5.3. Meta-learning as dictionary learning

Prior work has shown that a correct initialization is key to ensure a good performance for INRs [47, 52]. In this sense, recent studies [46, 51] have shown that the use of learned initialization, such as the ones obtained from meta-learning algorithms [16], can significantly boost the performance of INRs. Indeed, initializing with meta-learned weights is one of the most effective remedies against the slow speed of convergence, and high sample complexity of INRs. However, while there has been recently great progress in understanding traditional forms of deep learning, the role of meta-learning on the inductive bias of deep neural networks remains largely overlooked. Interestingly, we now show how using the connections between INRs and signal dictionaries can help us understand meta-learning in general.

To do so, we follow the same experimental protocol as in Sec. 5.2, where instead of computing the eigenfunctions of the NTK at a random initialization point, we linearize the INRs using Eq. (8) at the meta-learned weights, after pre-training on 1,000 training images from CelebA using model agnostic meta-learning (MAML) [16, 51].

As it is shown in Fig. 4, the meta-learned weights yield an eigenstructure that concentrates most of the energy of the target images on a subspace spanned by the eigenfunctions of the NTK with the largest eigenvalues, with almost no energy concentrated on the eigenfunctions corresponding to smaller eigenvalues. Therefore, training this INR starting from the meta-learned weights, results in a very fast speed of convergence and superior generalization capacity.

As it happened with the role of  $\omega_0$  in Sec. 5.2, visually inspecting the eigenfunctions of the NTK can help to build an intuition around this phenomenon. In this regard, recall that the CelebA dataset consists of a collection of face images. Strikingly, as illustrated in Fig. 5, the first eigenfunctions of the meta-learned NTK also look like faces. Clearly, meta-learning has reshaped the NTK so that the eigenfunctions have a large correlation with the target images.

To the best of our knowledge, we are the first to report the NTK reshaping behavior of meta-learning, which cannot be obviously explained by first order approximation theories (cf. Eq. (8)). This result is remarkable for deep learning theory, as it helps us understand the high-order dynamics of the NTK during training, which remains one of the main open questions of the field. Prior work had observed that standard training procedures change the first few eigenfunctions of the NTK so that they look like the target task [4, 23, 37, 38], but our observations in Fig. 4 and Fig. 5 go one step further, and show that meta-learning has the potential to reshape a much larger space of the NTK dictionary by combining many tasks together, thus increasing the capacity of the NTK to efficiently encode a full meta-distribution of signals<sup>8</sup>. In this sense, we believe that drawing parallels between classical dictionary learning algorithms [54] and meta-learning can be a strong abstraction which can simplify the complexity of this problem, thus leading to a promising avenue for future research. Delving deeper in this connection will not only improve our understanding of meta-learning as a whole, but it can also provide new insights for the design of more efficient INRs by leveraging data to construct richer dictionaries.

## 6. Related work

INRs are a very active research field in computer vision, as they have become integral parts of many applications such as volume reconstruction [30, 39], scene rendering [31, 34, 48], texture synthesis [19, 35], generative modelling [7, 9, 33], or compression [13]. Recent architectural advances have focused mostly on improving the inference and training cost [12, 24, 32, 41, 46, 51] of INRs, as well as on mitigating aliasing and improving generalization [5, 29].

The theory behind INRs has attracted much less attention, however. Similar to our work, Fathony *et al.* studied the expressive power of INRs, but their results only apply to their proposed multiplicative filter network architecture [15]. Zheng *et al.* [55], on the other hand, studied the trade-off between the rank and distance-preserving properties of different activation functions on INRs. Most notably, however, Tancik *et al.* [52] showed that precoding the input of an infinitely wide ReLU-network with random Fourier features [44] is equivalent to using a tunable

<sup>8</sup>In the Appendix we provide a more detailed experimental discussion.

shift-invariant kernel method. This gives a static intuition of how randomly initialized FFNs circumvent the spectral bias [43]. Our work goes one step further, and builds upon recent empirical results [37] to extend this NTK analysis to finite networks with arbitrary weights and activations, e.g., meta-learned SIRENs. This allows us to investigate dynamical aspects of INRs such as the role of pre-training.

Interestingly, Kopitkov and Indelman [23] also used the visualization of the eigenfunctions of the NTK during training to understand other high-order training effects, such as the increase of alignment of the NTK with the target signal [4, 23, 37, 38]. Our experiments use a similar approach to show the complex dictionary learning behaviour of MAML [16] in the NTK, which to the best of our knowledge is the first time this has been reported in the literature.

Connected to Thm. 1, other works have also used a similar harmonic expansion to analyze certain effects in deep learning, such as the increase in roughness of the loss landscape with respect to the weights for deeper layers [28], or how skip-connections can avoid shattered gradients [3].

Finally, we note that most of our work draws inspirations from the classical signal processing literature [36]. Some of our derivations are intimately connected to standard techniques in communications [42], and most of our analogies are founded on the field of signal representation [26] and dictionary design [54]. Moving forward, delving deeper on these connections will be a fruitful avenue for future work.

## 7. Conclusion

In this paper, we have analyzed the expressive power and inductive bias of modern INRs from a unified perspective. We have shown that the expressive power of a large class of INRs with sinusoidal encodings is given by the space of linear combinations of the integer harmonics of their input mapping. This allows INRs to encode signals with an exponentially large frequency support using a few coefficients, but also cause them to suffer from imperfect signal recovery or aliasing. We have also seen that the inductive bias of INRs is captured by the ability of the empirical NTK to encode signals efficiently, and we have revealed that meta-learning can modify the NTK and increase this efficiency.

A natural future extension would be to generalize Thm. 1 to input mappings beyond sinusoids [5, 15] or include normalization layers [2]. Similarly, one could also study the effect of out-of-distribution data on the alignment with the NTK after meta-training.

Finally, it is important to note that our insights should be readily extensible to higher dimensional settings, although most of our practical results were performed using one or two-dimensional signals. In this sense, designing methods to visualize the eigenfunctions of the NTK in higher dimensions would clearly help to inform practitioners about the inductive bias of different INRs.



## References

- [1] Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, Ruslan Salakhutdinov, and Ruosong Wang. On exact computation with an infinitely wide neural net. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 6
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *stat*, 1050:21, 2016. 8
- [3] David Balduzzi, Marcus Frean, Lennox Leary, JP Lewis, Kurt Wan-Duo Ma, and Brian McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In *International Conference on Machine Learning (ICML)*, 2017. 8
- [4] Aristide Baratin, Thomas George, César Laurent, R. Devon Hjelm, Guillaume Lajoie, Pascal Vincent, and Simon Lacoste-Julien. Implicit regularization via neural feature alignment. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021. 8
- [5] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 5, 8
- [6] Peter L. Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. In *Computational Learning Theory (COLT)*, 2001. 6
- [7] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 8
- [8] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
- [9] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 8
- [10] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989. 6
- [11] Alexander D’Amour, Katherine A. Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D. Hoffman, Farhad Hormozdiari, Neil Houlsby, Shaobo Hou, Ghassen Jerfel, Alan Karthikesalingam, Mario Lucic, Yi-An Ma, Cory Y. McLean, Diana Mincu, Akinori Mitani, Andrea Montanari, Zachary Nado, Vivek Natarajan, Christopher Nielson, Thomas F. Osborne, Rajiv Raman, Kim Ramasamy, Rory Sayres, Jessica Schrouff, Martin Seneviratne, Shannon Sequeira, Harini Suresh, Victor Veitch, Max Vladymyrov, Xuezhi Wang, Kellie Webster, Steve Yadlowsky, Taedong Yun, Xiaohua Zhai, and D. Sculley. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395*, 2020. 6
- [12] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021. 8
- [13] Emilien Dupont, Adam Golinski, Milad Alizadeh, Yee Whye Teh, and Arnaud Doucet. COIN: CCompression with implicit neural representations. In *Neural Compression: From Information Theory to Applications ICLRW*, 2021. 1, 8
- [14] Emilien Dupont, Yee Whye Teh, and Arnaud Doucet. Generative models as distributions of functions. *arXiv preprint arXiv:2102.04776*, 2021. 1
- [15] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations (ICLR)*, 2021. 3, 8
- [16] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning (ICML)*, 2017. 7, 8
- [17] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. *arXiv preprint arXiv:2103.10380*, 2021. 2
- [18] Peter Hedman, Pratul P Srinivasan, Ben Mildenhall, Jonathan T Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *arXiv preprint arXiv:2103.14645*, 2021. 2
- [19] Philipp Henzler, Niloy J Mitra, and Tobias Ritschel. Learning a neural 3d texture space from 2d exemplars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 8
- [20] Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, 2018. 2, 6
- [21] Jaechang Kim, Yunjoo Lee, Seunghoon Hong, and Jungseul Ok. Learning continuous representation of audio for arbitrary scale super resolution. *arXiv preprint arXiv:2111.00195*, 2021. 6
- [22] Sylwester Kłoczek, Łukasz Maziarka, Maciej Wołczyk, Jacek Tabor, Jakub Nowak, and Marek Śmieja. Hypernetwork functional image representation. In *International Conference on Artificial Neural Networks (ICANN)*, 2019. 1
- [23] Dmitry Kopitkov and Vadim Indelman. Neural spectrum alignment: Empirical study. In *International Conference on Artificial Neural Networks (ICANN)*, 2020. 8
- [24] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 8
- [25] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015. 7
- [26] Stéphane Mallat. *A wavelet tour of signal processing*. Elsevier, 1999. 8
- [27] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [28] Christian H.X. Ali Mehmeti-Göpel, David Hartmann, and Michael Wand. Ringing relus: Harmonic distortion analysis of nonlinear feedforward networks. In *International Conference on Learning Representations (ICLR)*, 2021. 3, 5, 8
- [29] Ishit Mehta, Michaël Gharbi, Connelly Barnes, Eli Shechtman, Ravi Ramamoorthi, and Manmohan Chandraker. Modulated periodic activations for generalizable local functional representations. *arXiv preprint arXiv:2104.03960*, 2021. 8
- [30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 8
- [31] 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 *Proceedings of European Conference on Computer Vision (ECCV)*, 2020. 1, 3, 8
- [32] Thomas Neff, Pascal Stadlbauer, Mathias Parger, Andreas Kurz, Joerg H. Mueller, Chakravarty R. Alla Chaitanya, Anton S. Kaplanyan, and Markus Steinberger. DONeRF: Towards Real-Time Rendering of Compact Neural Radiance Fields using Depth Oracle Networks. *Computer Graphics Forum*, 2021. 8
- [33] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 8
- [34] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 8
- [35] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 8
- [36] Alan V Oppenheim. *Discrete-time signal processing*. Pearson Education India, 1999. 5, 8
- [37] Guillermo Ortiz-Jiménez, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. What can linearized neural networks actually say about generalization? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 2, 6, 8
- [38] Jonas Paccolat, Leonardo Petrini, Mario Geiger, Kevin Tyloo, and Matthieu Wyart. Geometric compression of invariant manifolds in neural networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2021. 8
- [39] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 8
- [40] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 1
- [41] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021. 8
- [42] John G. Proakis and Masoud Salehi. *Fundamentals of Communication Systems*. Pearson Education Limited, 2014. 4, 8
- [43] Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning (ICML)*, 2019. 1, 2, 6, 8
- [44] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2008. 2, 8
- [45] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps, 2021. 2
- [46] Vincent Sitzmann, Eric R Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. *arXiv preprint arXiv:2006.09662*, 2020. 2, 7, 8
- [47] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 3, 5, 6, 7
- [48] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618*, 2019. 1, 8
- [49] Alexander J Smola and Bernhard Schölkopf. On a kernel-based method for pattern recognition, regression, approximation, and operator inversion. *Algorithmica*, 22(1):211–231, 1998. 6
- [50] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [51] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 7, 8
- [52] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1, 2, 5, 6, 7, 8

- [53] A. Tewari, O. Fried, J. Thies, V. Sitzmann, S. Lombardi, Z. Xu, T. Simon, M. Nießner, E. Tretschk, L. Liu, B. Mildenhall, P. Srinivasan, R. Pandey, S. Orts-Escolano, S. Fanello, M. Guo, G. Wetzstein, J.-Y. Zhu, C. Theobalt, M. Agrawala, D. B Goldman, and M. Zollhöfer. Advances in neural rendering. In *ACM SIGGRAPH 2021 Courses*, SIGGRAPH '21, 2021. 1
- [54] Ivana Tošić and Pascal Frossard. Dictionary learning. *IEEE Signal Processing Magazine*, 28(2):27–38, 2011. 3, 8
- [55] Jianqiao Zheng, Sameera Ramasinghe, and Simon Lucey. Rethinking positional encoding. *arXiv preprint arXiv:2107.02561*, 2021. 8