
Neural Implicit Dictionary Learning via Mixture-of-Expert Training

Peihao Wang¹ Zhiwen Fan¹ Tianlong Chen¹ Zhangyang Wang¹

Abstract

Representing visual signals by coordinate-based deep fully-connected networks has been shown advantageous in fitting complex details and solving inverse problems than discrete grid-based representation. However, acquiring such a continuous Implicit Neural Representation (INR) requires tedious per-scene training on tons of signal measurements, which limits its practicality. In this paper, we present a generic INR framework that achieves both data and training efficiency by learning a Neural Implicit Dictionary (NID) from a data collection and representing INR as a functional combination of basis sampled from the dictionary. Our NID assembles a group of coordinate-based subnetworks which are tuned to span the desired function space. After training, one can instantly and robustly acquire an unseen scene representation by solving the coding coefficients. To parallelly optimize a large group of networks, we borrow the idea from Mixture-of-Expert (MoE) to design and train our network with a sparse gating mechanism. Our experiments show that, NID can improve reconstruction of 2D images or 3D scenes by 2 orders of magnitude faster with up to 98% less input data. We further demonstrate various applications of NID in image inpainting and occlusion removal, which are considered to be challenging with vanilla INR. Our codes are available in <https://github.com/VITA-Group/Neural-Implicit-Dict>.

1. Introduction

Implicit Neural Representations (INRs) have recently demonstrated remarkable performance in representing multimedia signals in computer vision and graphics (Park et al., 2019; Mescheder et al., 2019; Saito et al., 2019; Chen et al.,

2021c; Sitzmann et al., 2020b; Tancik et al., 2020; Mildenhall et al., 2020). In contrast to classical discrete representations, where real-world signals are sampled and vectorized before processing, INR directly parameterizes the continuous mapping between coordinates and signal values using deep fully-connected networks (also known as multi-layer perceptron or MLP). This continuous parameterization enables to represent more complex and flexible scenes without being limited by grid extents and resolution in a more compact and memory efficient way.

However, one significant drawback of this approach is that acquiring an INR usually requires a *tedious per-scene* training of neural networks on *dense* measurements, which limits the practicality. Yu et al. (2021); Wang et al. (2021); Chen et al. (2021a) generalizes Neural Radiance Field (NeRF) (Mildenhall et al., 2020) across various scenes by projecting image features to a 3D volumetric proxy and then rendering feature volume to generate novel views. To speed up INR training, Sitzmann et al. (2020a); Tancik et al. (2021) apply meta-learning algorithms to learn the initial weight parameters for the MLP based on the underlying class of signals being represented. However, this line of works are either hard to be extended beyond NeRF scenario or incapable of producing high-fidelity results with insufficient supervision.

In this paper, we design a unified INR framework that simultaneously achieves optimization and data efficiency. We think of reconstructing an INR from few-shot measurements as solving an underdetermined system. Inspired by compressed sensing techniques (Donoho, 2006), we represent every neural implicit function as a linear combination of a function basis sampled from an over-complete *Neural Implicit Dictionary* (NID). Unlike conventional basis representation as a wide matrix, an NID is parameterized by a group of small neural networks that acts as continuous function basis spanning the entire target function space. The NID is shared across different scenes while the sparse codes are specified by each scene. We first acquire the NID “offline” by jointly optimizing it with per-scene coding across a class of instances in a training set. When transferring to unseen scenarios, we re-use the NID and only solves the scene specific coding coefficients “online”.

To effectively scale to thousands of subnetworks inside our dictionary, we employ the Mixture-of-Expert (MoE) training

¹Department of Electrical and Computer Engineering, University of Texas at Austin. Correspondence to: Zhangyang Wang <atlaswang@utexas.edu>.

for NID learning (Shazeer et al., 2017). We model each function basis in our dictionary as an expert subnetwork and the coding coefficients as its gating state. During each feed-forward, we utilize a routing module to generate sparsely coded gates, i.e., activating a handful of basis experts and linearly combining their responses. Training with MoE also “kills two birds with one stone” by constructing transferable dictionaries and avoiding extra computational overheads.

Our contributions can be summarized as follows:

- We propose a novel data-driven framework to learn a *Neural Implicit Dictionary (NID)* that can transfer across scenes, to both accelerate per-scene neural encoding and boost their performance.
- NID is parameterized by a group of small neural networks that acts as continuous function basis to span the neural implicit function space. The dictionary learning is efficiently accomplished via MoE training.
- We conduct extensive experiments to validate the effectiveness of NID. For training efficiency, we show that our approach is able to achieve 100× faster convergence speed for image regression task. For data efficiency, our NID can reconstruct signed distance function with 98% less point samples, and optimize a CT image with 90% fewer views. We also demonstrate more practical applications for NID, including image inpainting, medical image recovery, and transient object detection for surveillance videos.

2. Preliminaries

Compressed Sensing in Inverse Imaging. Compressed sensing and dictionary learning are widely applied in inverse imaging problems (Lustig et al., 2008; Metzler et al., 2016; Fan et al., 2018). In classical signal processing, signals are discretized and represented by vectors. A common goal is to reconstruct signals (or digital images) $\mathbf{x} \in \mathbb{R}^N$ from M measurements $\mathbf{y} \in \mathbb{R}^M$, which are formed by linearly transforming the underlying signals plus noise: $\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$. However, \mathbf{A} is often highly ill-posed, i.e., number of measurements is much smaller than the number of unknowns ($M \ll N$), which makes this inverse problem rather challenging. Compressed sensing (Candès et al., 2006; Donoho, 2006) provides an efficient approach to solve this underdetermined linear system by assuming signals $\mathbf{x} \in \mathbb{R}^N$ are compressible and representing it in terms of few vectors inside a group of spanning vectors $\boldsymbol{\Psi} = [\boldsymbol{\psi}_1 \ \cdots \ \boldsymbol{\psi}_K] \in \mathbb{R}^{N \times K}$. Then we can reconstruct \mathbf{x} through the following optimization objective:

$$\arg \min_{\boldsymbol{\alpha}} \|\boldsymbol{\alpha}\|_0 \text{ subject to } \|\mathbf{y} - \mathbf{A}\boldsymbol{\Psi}\boldsymbol{\alpha}\|_2 < \varepsilon \quad (1)$$

where $\boldsymbol{\alpha} \in \mathbb{R}^K$ is known as the sparse code coefficient, and $\|\boldsymbol{\eta}\|_2 \leq \varepsilon$ is a bound on the noise level. One often replaces the ℓ_0 semi-norm with ℓ_1 to obtain a convex objective. The spanning vectors $\boldsymbol{\Psi}$ can be chosen from orthonormal bases or, more often than not, over-complete dictionaries ($N \ll K$) (Kreutz-Delgado et al., 2003; Tošić & Frossard, 2011; Aharon et al., 2006; Chen & Needell, 2016). Rather than a bunch of spanning vectors, Chan et al. (2015); Tariyal et al. (2016); Pappayan et al. (2017) proposed hierarchical dictionary implemented by neural network layers.

Implicit Neural Representation. Implicit Neural Representation (INR) in computer vision and graphics replaces traditional discrete representations of multimedia objects with continuous functions parameterized by multilayer perceptrons (MLP) (Tancik et al., 2020; Sitzmann et al., 2020b). Since this representation is amenable to gradient-based optimization, prior works managed to apply coordinate-based MLPs to many inverse problems in computational photography (Park et al., 2019; Mescheder et al., 2019; Mildenhall et al., 2020; Chen et al., 2021c;b; Sitzmann et al., 2021; Fan et al., 2022; Attal et al., 2021b; Shen et al., 2021) and scientific computing (Han et al., 2018; Li et al., 2020; Zhong et al., 2021). Formally, we denote an INR inside a function space \mathcal{F} by $f_{\theta} : \mathbb{R}^m \rightarrow \mathbb{R}$, which continuously maps m -dimension spatio-temporal coordinates (say (x, y) with $m = 2$ for images) to the value space (say pixel intensity). Consider a functional $\mathcal{R} : \mathcal{F} \times \Omega \rightarrow \mathbb{R}$, we intend to find the network weights θ^* such that:

$$\mathcal{R}(f_{\theta^*} | \omega) = 0, \text{ for every } \omega \in \Omega \quad (2)$$

where Ω records the measurement settings. For instance, in computed tomography (CT), \mathcal{R} is called the volumetric projection integral and Ω specifies the ray parameterization and corresponding colors. When solving ordinal differential equations, \mathcal{R} takes form of $\rho(\mathbf{x}, f, \nabla f, \nabla^2 f, \dots)$ if $\mathbf{x} \in \Omega \setminus \partial\Omega$, while $\mathcal{R} = f(\mathbf{x}) - C$ for some constant C if $\mathbf{x} \in \partial\Omega$, given a compact set Ω and operator $\rho(\cdot)$ which combines derivatives of f (Sitzmann et al., 2020b).

Mixture-of-Expert Training. Shazeer et al. (2017) proposed outrageously wide neural networks with dynamic routing to achieve larger model capacity and higher data parallel. Their approach is to introduce an Mixture-of-Expert (MoE) layer with a number of expert subnetworks and train a gating network to select a sparse combination of the experts to process each input. Let us denote by $G(\mathbf{x})$ and $E_i(\mathbf{x})$ the output of the gating network and the output of the i -th expert network for a given input \mathbf{x} . The output of the MoE module can be written as:

$$\mathbf{y} = \sum_{i=1}^n G(\mathbf{x})_i E_i(\mathbf{x}), \quad (3)$$

where n is the number of experts and $\|G(\mathbf{x})\|_0 = k$. In Shazeer et al. (2017), computation is saved based on the sparsity of $G(\mathbf{x})$. The common sparsification strategy is called noisy top- k gating, which can be formulated as:

$$G(\mathbf{x}) = \text{Normalize}(\text{TopK}(H(\mathbf{x}), k)), \quad (4)$$

$$\text{TopK}(\mathbf{x}, k)_i = \begin{cases} x_i & \text{if } x_i \text{ is in top } k \text{ elements} \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where $H(\mathbf{x})$ synthesizes raw gating activations, $\text{TopK}(\cdot)$ masks out $n - k$ smallest elements, and $\text{Normalize}(\cdot)$ scales the magnitude of remaining weights to a constant, which can be chosen from softmax or ℓ_p -norm normalization.

3. Neural Implicit Dictionary Learning

As we discussed before, inverse imaging problems are often ill-posed and it is also true for Implicit Neural Representation (INR). Moreover, training an INR network is also time-consuming. How to kill two birds with one stone by efficiently and robustly acquiring an INR from few-shot observations remains uninvestigated. In this section, we answer this question by presenting our approach *Neural Implicit Dictionary* (NID), which are learned from data collections *a priori* and can be re-used to quickly fit an INR. We will first reinterpret two-layer SIREN (Sitzmann et al., 2020b) and point out the limitation of current design. Then we will elaborate on our proposed models and the techniques to improve its generalizability and stability.

3.1. Motivation by Two-Layer SIREN

Common INR architectures are pure Multi-Layer Perceptrons (MLP) with periodic activation functions. Fourier Feature Mapping (FFM) (Tancik et al., 2020) places a sinusoidal transformation after the first linear layer, while Sinusoidal Representation Network (SIREN) (Sitzmann et al., 2020b) replaces every nonlinear activation with a sinusoidal function. For the sake of simplicity, we only consider two-layer INR architectures to unify the formulation of FFM and SIREN. To be consistent with the notation in Section 2, let us denote INR by function $f: \mathbb{R}^m \rightarrow \mathbb{R}$, which can be formulated as below:

$$\gamma(\mathbf{x}) = [\sin(\mathbf{w}_1^T \mathbf{x} + b_1) \quad \cdots \quad \sin(\mathbf{w}_n^T \mathbf{x} + b_n)]^T, \quad (6)$$

$$f(\mathbf{x}) = \boldsymbol{\alpha}^T \gamma(\mathbf{x}) + c, \quad (7)$$

where $\mathbf{w}_i \in \mathbb{R}^m, b_i \in \mathbb{R}, \forall i \in [n]$ and $\boldsymbol{\alpha} \in \mathbb{R}^n, c \in \mathbb{R}$ are all network parameters, and mapping $\gamma(\cdot)$ (cf. Equation 6) is called positional embedding (Mildenhall et al., 2020;




Figure 1. Illustration of our NID pipeline. The blue experts are activated while grey ones are ignored.

Zhong et al., 2021). After simply rewriting, we can obtain:

$$f(\mathbf{x}) = \sum_{i=1}^n \alpha_i \sin(\mathbf{w}_i^T \mathbf{x} + b_i) + c \quad (8)$$

$$\approx \int_{\mathbb{R}^m} \frac{\alpha(\mathbf{w})}{\sqrt{\pi}} \sin\left(\mathbf{w}^T \mathbf{x} + \frac{\pi}{4}\right) d\mathbf{w}, \quad (9)$$

from which we discover Equations 6-7 can be considered as an approximation of inverse Hartley (Fourier) transform (cf. Equation 9). The weights of the first SIREN layer sample frequency bands on the Fourier domain, and passing coordinates through sinusoidal activation functions maps spatial positions onto cosine-sine wavelets. Then training a two-layer SIREN amounts to finding the optimal frequency supports and fitting the coefficients in Hartley transform.

Although trigonometric polynomials are dense in continuous function space, cosine-sine waves may not be always desirable as approximating functions at arbitrary precision with finite neurons can be infeasible. In fact, some other bases, such as Gegenbauer basis (Feng & Varshney, 2021) and Plücker embedding (Attal et al., 2021a), have been proven useful in different tasks. However, we argue that since handcrafted bases are agnostic to data distribution, they cannot express intrinsic information about data, thus may generalize poorly across various scenes. This causes per-scene training to re-select the frequency supports and re-fit the Fourier coefficients. Moreover, when observations are scarce, sinusoidal basis can also result in severe over-fitting in reconstruction (Sutherland & Schneider, 2015).

3.2. Learning Implicit Function Basis

Having reasoned why current INR architectures generalize badly and demand tons of measurements, we intend to introduce the philosophy of sparse dictionary representation (Kreutz-Delgado et al., 2003; Tošić & Frossard, 2011; Aharon et al., 2006) into INR. A dictionary contains a group of over-complete basis that spans the signal space. In contrast to handcrafted bases or wavelets, dictionary are usu-

ally learned from a data collection. Since it is aware of the distribution of the underlying signals to be represented, expressing signals using dictionary enjoys higher sparsity, robustness and generalization power.

Even though dictionary learning algorithms are well established in Aharon et al. (2006), it is far from trivial to design dictionaries amenable to INR on the continuous domain. Formally, we want to obtain a set of continuous maps: $b_i : \mathbb{R}^m \rightarrow \mathbb{R}, \forall i \in [n]$ such that for every signal $f : \mathbb{R}^m \rightarrow \mathbb{R}$ inside our target signal space \mathcal{F} , there exists a sparse coding $\alpha \in \mathbb{R}^n$ that can express the signal:

$$f(\mathbf{x}) = \alpha_1 b_1(\mathbf{x}) + \dots + \alpha_n b_n(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^m, \quad (10)$$

where n is the size of the dictionary, and α satisfies $\|\alpha\|_0 \leq k$ for some sparsity $k \ll n$. We parameterize each component in the dictionary with small coordinate-based networks by $b_{\theta_1}, \dots, b_{\theta_n}$, where θ_i denotes the network weights of the i -th element. We call this group of function basis *Neural Implicit Dictionary (NID)*.

We adopt an end-to-end optimization scheme to learn the NID. During training stage, we jointly optimize the sub-networks inside NID and the sparse coding assigned with each instance. Suppose we own a data collection with measurements captured from T multimedia instances to be represented (say T images or geometries of objects): $\mathcal{D} = \{\Omega^{(i)} \in \mathbb{R}^{t_i \times m}, \mathbf{Y}^{(i)} \in \mathbb{R}^{t_i}\}_{i=1}^T$, where $\Omega^{(i)}$ is the observation parameters (say coordinates on 2D lattice for images), m is the dimension of such parameters, $\mathbf{Y}^{(i)}$ are measured observations (say corresponding RGB colors), t_i denotes the number of observations for i -th instance. Then we optimize the following objective on the training dataset:

$$\begin{aligned} \arg \min_{\substack{\theta_1, \dots, \theta_n \\ \alpha^{(1)}, \dots, \alpha^{(T)}}} & \sum_{i=1}^T \sum_{j=1}^{t_i} \mathcal{L} \left(\mathcal{R}(f^{(i)} | \Omega_j^{(i)}), \mathbf{Y}_j^{(i)} \right) \quad (11) \\ & + \lambda \mathcal{P} \left(\alpha^{(1)}, \dots, \alpha^{(T)} \right), \\ \text{subject to } & f^{(i)}(\mathbf{x}) = \sum_{j=1}^n \alpha_j^{(i)} b_{\theta_j}(\mathbf{x}) \quad \forall \mathbf{x} \in \mathbb{R}^m, \end{aligned}$$

where $f^{(i)} \in \mathcal{F}$ is the INR of the i -th instance, $\mathcal{R}(f|\omega) : \mathcal{F} \times \Omega \rightarrow \mathbb{R}$ is a functional measuring function f with respect to a group of parameters ω . $\mathcal{L}(\cdot)$ is the loss function dependent of downstream tasks. $\mathcal{P}(\cdot)$ places a regularization onto the sparse coding, $\lambda = 0.01$ is fixed in our experiments. Besides sparsity penalty, we also consider some joint prior distributions among all codings, which will be discussed in Section 3.3. When transferring to unseen scenes, we fix NID basis $\{b_{\theta_i}\}_{i=1}^N$ and only compute the corresponding sparse coding to minimize the objective in Equation 11.

3.3. Training Thousands of Subnetworks with Mixture-of-Expert Layer

Directly invoking thousands of networks causes inefficiency and redundancy due to sample dependent sparsity. Moreover, this brute force computational strategy fails to properly utilize the advantage of modern computing architectures in parallelism. As we introduced in Section 2, Mixture-of-Expert (MoE) training system (Shazeer et al., 2017; He et al., 2021) provides a conditional computation mechanism that achieves stable and parallel training on a outrageously large networks. We notice that MoE layer and NID share the intrinsic similarity in the underlying running paradigm. Therefore, we propose to leverage an MoE layer to represent an NID accommodating thousands of implicit function basis. Specifically, each element in NID is an expert network in MoE layer, and the sparse coding encodes the gating states. Below we elaborate on the implementation details of the MoE based NID layer part by part:

Expert Networks. Each expert network is a small SIREN (Sitzmann et al., 2020b) or FFM (Tancik et al., 2020) network. To downsize the whole MoE layer, we share the positional embedding and the first 4 layers among all expert networks. Then we append two independent layers for each expert. We note this design can make two experts share the early-stage features and adjust their coherence.

Gating Networks. The generated gating is used as the sparse coding of an INR instance. We provide two alternatives to obtain the gating values: 1) We employ an encoder network as the gating function to map the (partial) observed measurements to the pre-sparsified weights. For grid-like modality, we utilize convolutional neural networks (CNN) (He et al., 2016; Liu et al., 2018; Gordon et al., 2019). For unstructured point modality, we adopt set encoders (Zaheer et al., 2017; Qi et al., 2017a;b). 2) We can also leverage a lookup table (Bojanowski et al., 2017) where each scene is assigned with a trainable embedding jointly optimized with expert networks. After computing the raw gating weights, we recall the method in Equation 3 to sparsify gates. Different from Shazeer et al. (2017), we do not perform softmax normalization to gating logits. Instead, we sort gating weights with respect to their absolute values, and normalize the weights by its ℓ_2 norm. Comparing aforementioned two gating functions, encoder-based gating networks benefit in parameter saving and instant inference without need of re-fitting sparse coding. However, headless embeddings demonstrate more strength in training efficiency and achieve better convergence.

Patch-wise Dictionary. It is implausible to construct an over-complete dictionary to represent entire signals. We adopt the walkround in (Reiser et al., 2021; Turki et al.,

Table 1. Performance of NID compared with FFM, SIREN, and Meta on CelebA dataset. \uparrow the higher the better, \downarrow the lower the better. The unit of # Params is megabytes, FLOPs is in gigabytes, and throughput is in #images/s.

Methods	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)	# Params	FLOPs	Throughput
FFM (Tancik et al., 2020)	22.60	0.636	0.244	147.8	20.87	0.479
SIREN (Sitzmann et al., 2020b)	26.11	0.758	0.379	66.56	4.217	0.540
Meta + 5 steps (Tancik et al., 2021)	23.92	0.583	0.322	66.69	4.217	0.536
Meta + 10 steps (Tancik et al., 2021)	29.64	0.651	0.182	66.69	4.217	0.536
NID + init. ($k = 128$)	28.75	0.892	0.061	8.972	23.30	30.37
NID + 5 steps ($k = 128$)	33.57	0.941	0.027	8.972	23.30	30.37
NID + 10 steps ($k = 128$)	35.10	0.954	0.021	8.972	23.30	30.37
NID + init. ($k = 256$)	30.26	0.919	0.045	8.972	29.55	21.23
NID + 5 steps ($k = 256$)	35.09	0.960	0.019	8.972	29.55	21.23
NID + 10 steps ($k = 256$)	37.75	0.971	0.012	8.972	29.55	21.23

2021) by partitioning the coordinate space into regular and overlapped patches, and assign separate NID to each block. We implement this by setting up multiple MoE layers and dispatch the coordinate inputs to corresponding MoE with respect to the region where they are located.

Utilization Balancing and Warm-Up. It was observed that gating network tends to converge to a self-reinforcing imbalanced state, where it always produces large weights for the same few experts (Shazeer et al., 2017). To tackle this problem, we pose a regularization on the Coefficient of Variation (CV) of the sparse codings following Bengio et al. (2015); Shazeer et al. (2017). The CV penalty is defined as:

$$\mathcal{P}_{CV}(\alpha^{(1)}, \dots, \alpha^{(T)}) = \frac{\text{Var}(\bar{\alpha})}{(\sum_{i=1}^n \bar{\alpha}_i/n)^2}, \quad (12)$$

$$\text{where } \bar{\alpha} = \sum_{i=1}^T \alpha^{(i)}. \quad (13)$$

Evaluating this regularization over the whole training set is infeasible. Instead we estimate and minimize this loss per batch. We also find hard sparsification will stop gradient back-propagation, which leads to stationary gating states equal to the initial stage. To address this side-effect, we first abandon hard thresholding and train the MoE layer with ℓ_1 penalty $\mathcal{P}_{\ell_1} = \sum_{i=1}^T \|\alpha^{(i)}\|_1$ on codings for several epochs, and enable sparsification afterwards.

4. Experiments and Applications

In this section, we demonstrate the promise of NID by showing several applications in scene representation.

4.1. Instant Image Regression

A prototypical example of INR is to regress a 2D image with an MLP which takes in coordinates on 2D lattice and is supervised with RGB colors. Given a $D \times D$ image

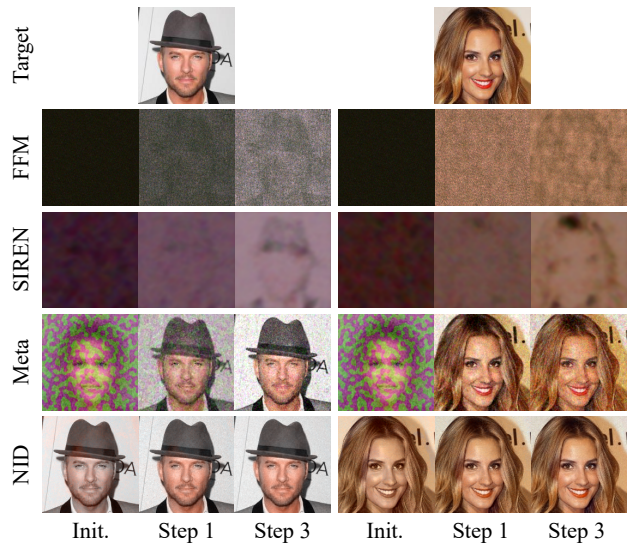


Figure 2. A closer look at the early training stages of FFM, SIREN, Meta, and NID, respectively.

$\mathbf{Y} \in \mathbb{R}^{D \times D \times 3}$, our goal is to approximate the mapping $f: \mathbb{R}^2 \mapsto \mathbb{R}^3$ by optimizing $\|f(i, j) - \mathbf{Y}_{ij}\|_2$ for every $(i, j) \in [0, D]^2$, where $f_\theta = \sum_i \alpha_i b_{\theta_i}$. In conventional training scheme, each image is encoded into a dedicated network after thousands of iterations. Instead, we intend to use NID to instantly acquire such INR without training or with only few steps of gradient descent.

Experimental Settings. We choose to train our NID on CelebA face dataset (Liu et al., 2015), where each image is cropped to 178×178 . Our NID contains 4096 experts, each of which share a 4-layer backbone with 256 hidden dimension and own a separate 32-dimension output layer. We adopt 4 residual convolutional blocks (He et al., 2016) as the gating network. During training, the gating network is tuned with the dictionary. NID is warmed up within 10 epochs and then start to only keep top 128 experts for each input for 5000 epochs. At the inference stage, we




Figure 3. Visualization of foreground-background decomposition results for surveillance video via principal component pursuit with NID.

let gating network directly output the sparse coding of the test image. To further improve the precision, we utilize the output as the initialization, and then use gradient descent to further optimize the sparse coding with the dictionary fixed. We contrast our methods to FFM (Tancik et al., 2020), SIREN (Sitzmann et al., 2020b) and Meta (Tancik et al., 2021). In Table 1, we demonstrate the overall PSRN, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) of these four models on test set (with 500 images) under the limited training step setting, where FFM and SIREN are only trained for 100 steps. We also present the inference time metrics in Table 1, including the number of parameters to represent 500 images, FLOPs to render a single image, and measured throughput of images rendered per second. In Figure 2, we zoom into the initialization and early training stages of each model.

Results. Results in Table 1 show that NID ($k = 256$) can achieve best performance among all compared models even without subsequent optimization steps. A relative sparser NID ($k = 128$) can also surpass both FFM and SIREN (trained with 100 steps) with the initially inferred coding. Compared with meta-learning based method, our model can outperform them by a significant margin (≥ 5 dB) within the same optimization steps. We note that since NID only further tunes the coding vector, both computation and convergence speed are much faster than meta-learning approaches which fine-tune parameters of the whole network. Figure 2 illustrates that the initial sparse coding inferred from the gating network is enough to produce high-accuracy reconstructed images. With 3 more gradient descent steps (which usually takes 5 seconds), it can reach the quality of well-tuned per-scene training INR (which takes 10 minutes). We argue that although meta learning is able to find a reasonable start point, but the subsequent optimization is sensitive to saddle points where the represented images are fuzzy and noisy. In regard to model efficiency, our NID is 8 times more compact than single-MLP representation, as NID shares dictionary among all samples and only needs to additionally record an small gating network. Moreover, our MoE implementation results in a significant throughput gain, as it makes inference highly parallelable. We point



Figure 4. Qualitative results of inpainting image from corruptions with NID.

out that meta-learning can only provide an initialization. To represent all test images, one has to save all dense parameters separately. Horizontally compared, denser NID is more expressive than sparser one though sacrificing efficiency.

4.2. Facial Image Inpainting.

Image inpainting recovers images corrupted by occlusion. Previous works (Liu et al., 2018; Yu et al., 2019) only establish algorithms based on discrete representation. In this section, we demonstrate image inpainting directly on continuous INR. Given a corrupted image $\mathbf{Y} \in \mathbb{R}^{D \times D \times 3}$, we remove outliers by projecting \mathbf{Y} onto some low-dimension linear (function) subspace spanned by components in a dictionary. We achieve this by trying to represent the corrupted image as a linear combination of a pre-trained NID, while simultaneously enforcing the sparsity of this combination. Specifically, we fix the dictionary in Equation 11 and choose ℓ_1 norm as the loss function \mathcal{L} (Candès et al., 2011), where we assume noises are sparsely distributed on images.

Experimental Settings. We corrupt images by randomly pasting a 48×48 color patch. To recover images, we borrow the dictionary trained on CelebA dataset from Section 4.1. However, we do not leverage the gating network to synthesize the sparse coding. Instead, we directly optimize a randomly initialized coding to minimize Equation 11. Our baseline includes SIREN and Meta (Tancik et al., 2021). We change their loss function to ℓ_1 norm to keep consistent. To inpaint with Meta, we start from its learned initialization, and optimize two steps towards the objective.

Results. The inpainting results are presented in Figure 4. Our findings are 1) SIREN overfits all given signals as it does not rely on any image priors. 2) Meta-learning based approach implicitly poses a prior by initializing the networks around a desirable optimum. However, our experiment shows that the learned initialization is ad-hoc to a certain data distribution. When noises are added, Meta turns unstable and converges to a trivial solution. 3) Our NID displays stronger robustness by accurately locating and removing the occlusion pattern.

4.3. Self-Supervised Surveillance Video Analysis

In this section, we establish a self-supervision algorithm that can decompose foreground and background for surveillance videos based on NID. Given a set of video frames $\{\mathbf{Y}^{(t)} \in \mathbb{R}^{D \times D \times 3}\}_{t=1}^T$, our goal is to find a continuous mapping $f(x, y, t)$ representing the clip that can be decomposed to: $f(x, y, t) = f_X(x, y, t) + f_E(x, y, t)$, where f_X is the background and f_E are transient noises (e.g., pedestrians). We borrow the idea from Robust Principal Component Analysis (RPCA) (Candès et al., 2011; Ji et al., 2010) where background is assumed to be “low-rank” and noises are assumed to be sparse. Despite well-established for discrete representation, modeling “low-rank” in continuous domain remains elusive. We achieve this by assuming $f_X(x, y, t)$ at each time stamp are largely represented by the same group of experts, i.e., the non-zero elements in the sparse codings concentrate to several points, and the coding weights follow a decay distribution. Mathematically, we first rewrite f by decoupling spatial coordinates and time: $f(x, y, t) = \sum_i \alpha_i(t) b_{\theta_i}(x, y)$, where every time slice shares a same dictionary, and sparse coding $\alpha_i(t)$ depends on the timestamp. Then we minimize:

$$\arg \min_{\theta_1, \dots, \theta_n, \alpha(t)} \sum_{t=1}^T \sum_{(x,y)} \left\| \sum_{i=1}^n \alpha_i(t) b_{\theta_i}(x, y) - \mathbf{Y}_{xy}^{(t)} \right\|_1 \quad (14)$$

$$+ \lambda \sum_{t=1}^T \sum_{i=1}^n \frac{|\alpha_i(t)|}{\exp(-\beta i)},$$

where the second term penalize the sparsity of $\alpha(t)$ according to an exponentially increasing curve (controlled by β), which implies the larger i is, the more sparsity is enforced. As a consequence, every time slice are largely approximated by the first few components in NID, which simulates the nature of “low-rank” representation for continuous functions.

Results. We test the above algorithm on BMC-Real dataset (Vacavant et al., 2012). In our implementation, $\alpha(t)$ is also parameterized by another MLP, and we choose $\beta = 0.5$. Our qualitative results are presented in Figure 3. We verify that our algorithm can decompose the background and foreground correctly by imitating the behavior of RPCA.




Figure 5. Qualitative results of CT reconstruction from sparse measurements.

This application further demonstrates the potential of our NID in combining with subspace learning techniques.

Table 2. Quantitative results of CT reconstruction compared with FFM, SIREN, and Meta. (PSNR in dB)

Methods	128 views		16 views		8 views	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
FFM (Tancik et al., 2020)	22.81	0.845	15.22	0.122	13.58	0.095
SIREN (Sitzmann et al., 2020b)	24.32	0.891	18.48	0.510	17.26	0.483
Meta (Tancik et al., 2021)	32.70	0.948	21.39	0.822	18.28	0.574
NID ($k = 128$)	36.56	0.939	24.48	0.818	16.24	0.619
NID ($k = 256$)	37.49	0.944	26.32	0.829	16.77	0.636

4.4. Computed Tomography Reconstruction

Computed tomography (CT) is a widely used medical imaging technique that captures projective measurements of the volumetric density of body tissue. This imaging formation can be formulated as below:

$$Y(r, \phi) = \int_{\mathbb{R}^2} f(x, y) \delta(r - x \cos \phi - y \sin \phi) dx dy,$$

where r is the location on the image plane, ϕ is the viewing angle, and $\delta(\cdot)$ is known as Dirac delta function. Due to limited number of measurements, reconstructing f through inverting this integral is often ill-posed. We propose to shrink the solution space by using NID as a regularization.

Experimental Settings. We conduct experiments on Shepp-Logan phantoms dataset (Shepp & Logan, 1974) with 2048 randomly generated 128×128 CTs. We first

directly train an NID over 1k CT images, during which the total number of experts is 1024, and each CT selects 128/256 experts. In CT scenario, a look-up table is chosen as our gating network. Afterwards, we randomly sample 128 viewing angles, and synthesize 2D integral projections of a bundle of 128 parallel rays from these angles as the measurement. To testify the effectiveness of our method under limited number of observations, we downsample 128 views by 12.5%(16) and 6.25%(8) respectively. Again, we choose FFM (Tancik et al., 2020), SIREN (Sitzmann et al., 2020b), and Meta (Tancik et al., 2021) as our baselines.

Results. The quantitative results are listed in Table 2. We observe that our NID consistently leads two metrics in the table. When sampled views are sufficient, NID achieves the highest PSNR, while when views are reduced, our NID takes advantage in SSIM. We also plot the qualitative results in Figure 5. We find that our NID can regularize the reconstructed results to be smooth and shape-consistent, which leads to less missing wedge artifacts.

4.5. Shape Representation from Point Clouds

Recent works (Park et al., 2019; Sitzmann et al., 2020a;b; Gropp et al., 2020) convert point clouds to continuous surface representation through directly regressing a Signed Distance Function (SDF) parameterized by MLPs. Suppose $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ is our target SDF, given a set of points $\Omega \subset \mathbb{R}^3$, we fit f by solving a integral equation of the form below (Park et al., 2019):

$$\arg \min_f \int_{\mathbf{x} \in \Omega} |f(\mathbf{x})| d\mathbf{x} + \int_{\mathbf{x} \in \mathbb{R}^3 \setminus \Omega} |f(\mathbf{x}) - d(\mathbf{x}, \Omega)| d\mathbf{x}, \quad (15)$$

where $d(\mathbf{x}, \Omega)$ denotes the signed shortest distance from point \mathbf{x} to point set Ω . During optimization, we evaluate the first integral via sampling inside the given point cloud and the second term via uniformly sampling over the whole space. Tackling this integral with sparsely sampled points around the surface is challenging (Park et al., 2019). Similarly, we introduce NID to learn *a priori* SDF basis from data and then leverage it to regularize the solution.

Table 3. Quantitative results of SDF reconstruction compared with SIREN, DeepSDF, MetaSDF. CD is short for Chamfer Distance (magnified by 10^3), NC means Normal Consistency. \uparrow the higher the better, \downarrow the lower the better.

Methods	500k points		50k points		10k points	
	CD(\downarrow)	NC(\uparrow)	CD(\downarrow)	NC(\uparrow)	CD(\downarrow)	NC(\uparrow)
SIREN (Sitzmann et al., 2020b)	0.051	0.962	0.163	0.801	1.304	0.169
IGR (Gropp et al., 2020)	0.062	0.927	0.170	0.812	0.961	0.676
DeepSDF (Park et al., 2019)	0.059	0.925	0.121	0.856	2.751	0.194
MetaSDF (Sitzmann et al., 2020a)	0.067	0.884	0.097	0.878	0.132	0.755
ConvONet (Peng et al., 2020)	0.052	0.938	0.082	0.914	0.133	0.845
NID ($k = 128$)	0.058	0.940	0.067	0.948	0.093	0.921
NID ($k = 256$)	0.053	0.956	0.063	0.952	0.088	0.945

Experimental Settings. Our experiments about SDF are conducted on ShapeNet (Chang et al., 2015) datasets, from which we pick the *chair* category for demonstration. To guarantee meshes are watertight, we run the toolkit provided by Huang et al. (2018) to convert the whole dataset. We split the *chair* category following Choy et al. (2016), and fit our NID over the training set. The total number of experts is 4096, and after 20 warm-up epochs, only 128/256 experts will be preserved for each sample. We choose lookup table as our gating network. During inference time, we sample 500k, 50k and 10k point clouds, respectively, from the test surfaces. Then we optimize objective in Equation 15 to obtain the regressed SDF with f represented by our NID. In addition to SIREN and IGR (Gropp et al., 2020), We choose DeepSDF (Park et al., 2019), MetaSDF (Sitzmann et al., 2020a), and ConvONet (Peng et al., 2020) as our baselines. Our evaluation metrics are Chamfer distance (the average minimal pairwise distance) and normal consistency (the angle between corresponding normals).

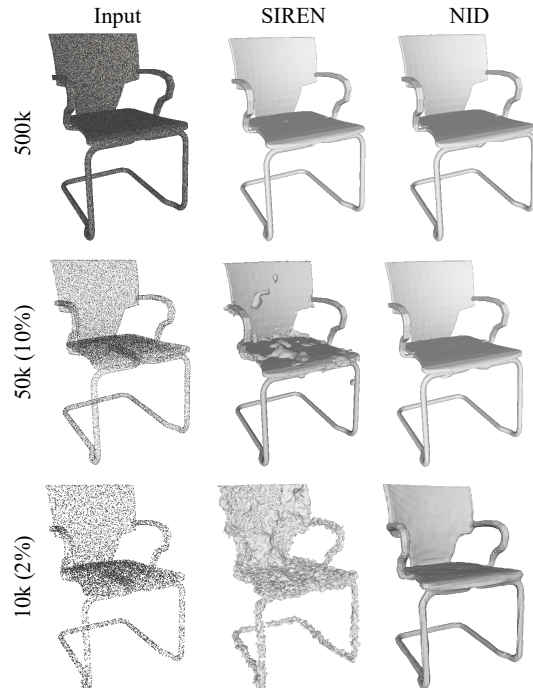


Figure 6. Qualitative results of SDF reconstruction from sparse point clouds.

Results. We put our numerical results in Table 3, from which we can summarize that our NID is more robust to smaller number of points. As the performance of other methods drops quickly, the CD metric of NID stays below 0.1 and NC keeps above 0.9. We also provide qualitative illustration in Figure 6. We conclude that thanks to the constraint of our NID, the SDF will not collapse at some point where observations are missing. DeepSDF and ConvONet reply on latent feature space to decode geometries, which shows the potential in regularizing geometries. However, the superiority of

our model suggests our dictionary based representation is advantageous over conditional implicit representation.

5. Related Work

Generalizable Implicit Neural Representations. Implicit Neural Representation (INR) (Tancik et al., 2020; Sitzmann et al., 2020b) notoriously suffers from the limited cross-scene generalization capability. Tancik et al. (2021); Sitzmann et al. (2020a) propose meta-learning based algorithms to better initialize INR weights for fast convergence. Chen et al. (2021c); Park et al. (2019); Chabra et al. (2020); Chibane et al. (2020); Jang & Agapito (2021); Martin-Brualla et al. (2021); Rematas et al. (2021) introduce learnable latent embeddings to encode scene specific information and condition the INR on the latent code for generalizable representation. In Sitzmann et al. (2020b), the authors further utilize a hyper-network (Ha et al., 2016) to predict INR weights directly from inputs. Compared with conditional fields or hyper-network based methods, sparse coding based NID, with just one last layer, can achieve faster adaptation. The dictionary representation simplifies the mapping between latent spaces to a sparse linear combination over the additive basis, which can be manipulated more interpretably and also contributes to transferability. Last but not least, it is known that imposing sparsity can help overcome noise in ill-posed inverse problems (Donoho, 2006; Candès et al., 2011).

Mixture of Experts (MoE). Mixture of Experts (Jacobs et al., 1991; Jordan & Jacobs, 1994; Chen et al., 1999; Yuksel et al., 2012; Roller et al., 2021) perform conditional computations composed of a group of parallel sub-models (a.k.a. experts) according to a routing policies (Dua et al., 2021; Roller et al., 2021). Recent advances (Shazeer et al., 2017; Lepikhin et al., 2020; Fedus et al., 2021) improve MoE by adopting a sparse-gating strategy, which only activates a minority of experts by selecting top candidates according to the scores given by the gating networks. This brings massive advantages in model capacity, training time, and achieved performance (Shazeer et al., 2017). Fedus et al. (2021) even built language models with trillions of parameters. To stabilize the training, Hansen (1999); Lepikhin et al. (2020); Fedus et al. (2021) investigated auxiliary loading loss to balance the selection of experts. Alternatively, Lewis et al. (2021); Clark et al. (2022) encourage a balanced routing by solving a linear assignment problem.

6. Conclusion

We propose Neural Implicit Dictionary (NID) learned from data collection to represent the signals as a sparse combination of the function basis inside. Unlike tradition dictionary, our NID contains continuous function basis, which

are parameterized by subnetworks. To train thousands of networks efficiently, we employ Mixture-of-Expert training strategy. Our NID enjoys higher compactness, robustness, and generalization. Our experiments demonstrate promising applications of NID in instant regression, image inpainting, video decomposition, and reconstruction from sparse observations. Our future work may bring in subspace learning theories to analyze NID.

Acknowledgement

Z. W. is in part supported by a US Army Research Office Young Investigator Award (W911NF2010240).

References

- Aharon, M., Elad, M., and Bruckstein, A. K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on signal processing*, 54(11):4311–4322, 2006.
- Attal, B., Huang, J.-B., Zollhoefer, M., Kopf, J., and Kim, C. Learning neural light fields with ray-space embedding networks. *arXiv preprint arXiv:2112.01523*, 2021a.
- Attal, B., Laidlaw, E., Gokaslan, A., Kim, C., Richardt, C., Tompkin, J., and O’Toole, M. Törf: Time-of-flight radiance fields for dynamic scene view synthesis. *Advances in neural information processing systems*, 34, 2021b.
- Bengio, E., Bacon, P.-L., Pineau, J., and Precup, D. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015.
- Bojanowski, P., Joulin, A., Lopez-Paz, D., and Szlam, A. Optimizing the latent space of generative networks. *arXiv preprint arXiv:1707.05776*, 2017.
- Candès, E. J., Romberg, J., and Tao, T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006.
- Candès, E. J., Li, X., Ma, Y., and Wright, J. Robust principal component analysis? *Journal of the ACM (JACM)*, 58(3): 1–37, 2011.
- Chabra, R., Lenssen, J. E., Ilg, E., Schmidt, T., Straub, J., Lovegrove, S., and Newcombe, R. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *European Conference on Computer Vision*, pp. 608–625. Springer, 2020.
- Chan, T.-H., Jia, K., Gao, S., Lu, J., Zeng, Z., and Ma, Y. Pcanet: A simple deep learning baseline for image classification? *IEEE transactions on image processing*, 24(12):5017–5032, 2015.

- Chang, A. X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., and Yu, F. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015.
- Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., and Su, H. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. *arXiv preprint arXiv:2103.15595*, 2021a.
- Chen, G. and Needell, D. Compressed sensing and dictionary learning. *Finite Frame Theory: A Complete Introduction to Overcompleteness*, 73:201, 2016.
- Chen, H., He, B., Wang, H., Ren, Y., Lim, S. N., and Shrivastava, A. Nerv: Neural representations for videos. *Advances in Neural Information Processing Systems*, 34, 2021b.
- Chen, K., Xu, L., and Chi, H. Improved learning algorithms for mixture of experts in multiclass classification. *Neural networks*, 12(9):1229–1252, 1999.
- Chen, Y., Liu, S., and Wang, X. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8628–8638, 2021c.
- Chibane, J., Alldieck, T., and Pons-Moll, G. Implicit functions in feature space for 3d shape reconstruction and completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6970–6981, 2020.
- Choy, C. B., Xu, D., Gwak, J., Chen, K., and Savarese, S. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- Clark, A., Casas, D. d. l., Guy, A., Mensch, A., Paganini, M., Hoffmann, J., Damoc, B., Hechtman, B., Cai, T., Borgeaud, S., et al. Unified scaling laws for routed language models. *arXiv preprint arXiv:2202.01169*, 2022.
- Donoho, D. L. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- Dua, D., Bhosale, S., Goswami, V., Cross, J., Lewis, M., and Fan, A. Tricks for training sparse translation models. *arXiv preprint arXiv:2110.08246*, 2021.
- Fan, Z., Sun, L., Ding, X., Huang, Y., Cai, C., and Paisley, J. A segmentation-aware deep fusion network for compressed sensing mri. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 55–70, 2018.
- Fan, Z., Jiang, Y., Wang, P., Gong, X., Xu, D., and Wang, Z. Unified implicit neural stylization. *arXiv preprint arXiv:2204.01943*, 2022.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.
- Feng, B. Y. and Varshney, A. Signet: Efficient neural representation for light fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14224–14233, 2021.
- Gordon, J., Bruinsma, W. P., Foong, A. Y., Requeima, J., Dubois, Y., and Turner, R. E. Convolutional conditional neural processes. *arXiv preprint arXiv:1910.13556*, 2019.
- Gropp, A., Yariv, L., Haim, N., Atzmon, M., and Lipman, Y. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020.
- Ha, D., Dai, A., and Le, Q. V. Hypernetworks. *arXiv preprint arXiv:1609.09106*, 2016.
- Han, J., Jentzen, A., and Weinan, E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- Hansen, J. V. Combining predictors: comparison of five meta machine learning methods. *Information Sciences*, 119(1-2):91–105, 1999.
- He, J., Qiu, J., Zeng, A., Yang, Z., Zhai, J., and Tang, J. Fastmoe: A fast mixture-of-expert training system. *arXiv preprint arXiv:2103.13262*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016.
- Huang, J., Su, H., and Guibas, L. Robust watertight manifold surface generation method for shapenet models. *arXiv preprint arXiv:1802.01698*, 2018.
- Jacobs, R. A., Jordan, M. I., Nowlan, S. J., and Hinton, G. E. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- Jang, W. and Agapito, L. Codenerf: Disentangled neural radiance fields for object categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12949–12958, 2021.
- Ji, H., Liu, C., Shen, Z., and Xu, Y. Robust video denoising using low rank matrix completion. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 1791–1798. IEEE, 2010.

- Jordan, M. I. and Jacobs, R. A. Hierarchical mixtures of experts and the em algorithm. *Neural computation*, 6(2): 181–214, 1994.
- Kreutz-Delgado, K., Murray, J. F., Rao, B. D., Engan, K., Lee, T.-W., and Sejnowski, T. J. Dictionary learning algorithms for sparse representation. *Neural computation*, 15(2):349–396, 2003.
- Lepikhin, D., Lee, H., Xu, Y., Chen, D., Firat, O., Huang, Y., Krikun, M., Shazeer, N., and Chen, Z. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.
- Lewis, M., Bhosale, S., Dettmers, T., Goyal, N., and Zettlemoyer, L. Base layers: Simplifying training of large, sparse models. In *International Conference on Machine Learning*, pp. 6265–6274. PMLR, 2021.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- Liu, G., Reda, F. A., Shih, K. J., Wang, T.-C., Tao, A., and Catanzaro, B. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 85–100, 2018.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008.
- Martin-Brualla, R., Radwan, N., Sajjadi, M. S., Barron, J. T., Dosovitskiy, A., and Duckworth, D. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7210–7219, 2021.
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., and Geiger, A. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4460–4470, 2019.
- Metzler, C. A., Maleki, A., and Baraniuk, R. G. From denoising to compressed sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pp. 405–421. Springer, 2020.
- Papayan, V., Romano, Y., and Elad, M. Convolutional neural networks analyzed via convolutional sparse coding. *The Journal of Machine Learning Research*, 18(1):2887–2938, 2017.
- Park, J. J., Florence, P., Straub, J., Newcombe, R., and Lovegrove, S. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 165–174, 2019.
- Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., and Geiger, A. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pp. 523–540. Springer, 2020.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 652–660, 2017a.
- Qi, C. R., Yi, L., Su, H., and Guibas, L. J. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017b.
- Reiser, C., Peng, S., Liao, Y., and Geiger, A. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. *arXiv preprint arXiv:2103.13744*, 2021.
- Rematas, K., Martin-Brualla, R., and Ferrari, V. Sharf: Shape-conditioned radiance fields from a single view. *arXiv preprint arXiv:2102.08860*, 2021.
- Roller, S., Sukhbaatar, S., Szlam, A., and Weston, J. E. Hash layers for large sparse models. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=lMgDDWb1ULW>.
- Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., and Li, H. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2304–2314, 2019.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.
- Shen, S., Wang, Z., Liu, P., Pan, Z., Li, R., Gao, T., Li, S., and Yu, J. Non-line-of-sight imaging via neural transient fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

- Shepp, L. A. and Logan, B. F. The fourier reconstruction of a head section. *IEEE Transactions on nuclear science*, 21(3):21–43, 1974.
- Sitzmann, V., Chan, E. R., Tucker, R., Snavely, N., and Wetzstein, G. Metasdf: Meta-learning signed distance functions. *arXiv preprint arXiv:2006.09662*, 2020a.
- Sitzmann, V., Martel, J., Bergman, A., Lindell, D., and Wetzstein, G. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Sitzmann, V., Rezkikov, S., Freeman, W. T., Tenenbaum, J. B., and Durand, F. Light field networks: Neural scene representations with single-evaluation rendering. *arXiv preprint arXiv:2106.02634*, 2021.
- Sutherland, D. J. and Schneider, J. On the error of random fourier features. *arXiv preprint arXiv:1506.02785*, 2015.
- Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *arXiv preprint arXiv:2006.10739*, 2020.
- Tancik, M., Mildenhall, B., Wang, T., Schmidt, D., Srinivasan, P. P., Barron, J. T., and Ng, R. Learned initializations for optimizing coordinate-based neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2846–2855, 2021.
- Tariyal, S., Majumdar, A., Singh, R., and Vatsa, M. Deep dictionary learning. *IEEE Access*, 4:10096–10109, 2016.
- Tošić, I. and Frossard, P. Dictionary learning. *IEEE Signal Processing Magazine*, 28(2):27–38, 2011.
- Turki, H., Ramanan, D., and Satyanarayanan, M. Meganerf: Scalable construction of large-scale nerfs for virtual fly-throughs. *arXiv preprint arXiv:2112.10703*, 2021.
- Vacavant, A., Chateau, T., Wilhelm, A., and Lequievre, L. A benchmark dataset for outdoor foreground/background extraction. In *Asian Conference on Computer Vision*, pp. 291–300. Springer, 2012.
- Wang, Q., Wang, Z., Genova, K., Srinivasan, P. P., Zhou, H., Barron, J. T., Martin-Brualla, R., Snavely, N., and Funkhouser, T. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4690–4699, 2021.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Yu, A., Ye, V., Tancik, M., and Kanazawa, A. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4578–4587, 2021.
- Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., and Huang, T. S. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4471–4480, 2019.
- Yuksel, S. E., Wilson, J. N., and Gader, P. D. Twenty years of mixture of experts. *IEEE Transactions on Neural Networks and Learning Systems*, 23(8):1177–1193, 2012. doi: 10.1109/TNNLS.2012.2200299.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R., and Smola, A. Deep sets. *arXiv preprint arXiv:1703.06114*, 2017.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.
- Zhong, E. D., Bepler, T., Berger, B., and Davis, J. H. Cryodrnn: reconstruction of heterogeneous cryo-em structures using neural networks. *Nature Methods*, 18(2):176–185, 2021.