

Grounding and Enhancing Grid-based Models for Neural Fields

Zelin Zhao¹, Fenglei Fan^{*2}, Wenlong Liao^{1,3}, and Junchi Yan¹

¹Department of CSE & MoE Key Lab of AI, Shanghai Jiao Tong University

²Department of Mathematics, The Chinese University of Hong Kong

³Cowa Tech. Ltd.

Abstract

Many contemporary studies utilize grid-based models for neural field representation, but a systematic analysis of grid-based models is still missing, hindering the improvement of those models. Therefore, this paper introduces a theoretical framework for grid-based models. This framework points out that these models' approximation and generalization behaviors are determined by grid tangent kernels (GTK), which are intrinsic properties of grid-based models. The proposed framework facilitates a consistent and systematic analysis of diverse grid-based models. Furthermore, the introduced framework motivates the development of a novel grid-based model named the Multiplicative Fourier Adaptive Grid (MulFAGrid). The numerical analysis demonstrates that MulFAGrid exhibits a lower generalization bound than its predecessors, indicating its robust generalization performance. Empirical studies reveal that MulFAGrid achieves state-of-the-art performance in various tasks, including 2D image fitting, 3D signed distance field (SDF) reconstruction, and novel view synthesis, demonstrating superior representation ability. The project website is available at [this link](#).

1. Introduction

Neural fields [29, 56] are coordinate-based networks representing a field, a continuous parameterization representing a physical quantity of an object or a scene. These fields have demonstrated significant success in image regression [46], view synthesis [29], and 3D model reconstruction [51]. Recent studies find that neural field techniques can be applied to visual computing problems and beyond [56]. Therefore, this field is poised to have a topographic impact on computer vision and machine learning.

Recent empirical studies [9, 18, 43, 54, 58] have sub-

stantiated that grid-based models, parameterized by grid feature tensors and operating on grids, can achieve a maximum two-orders-of-magnitude speed-up compared to MLP-based neural fields [29, 31, 43], all while upholding high-fidelity representation quality. Grid-based models can be either regular (based on regular grids [43, 44, 58]) or irregular (based on point cloud [18] or mesh [33, 51]). However, to the best of our knowledge, there is no known theory to analyze the learning behaviors of grid-based models systematically. Consequently, the empirical success of grid-based models not only lacks a theoretical foundation but also reveals a deficiency in effective principles for designing enhanced grid-based models.

To ground and enhance grid-based models, we propose a theory inspired by neural tangent kernels [6, 17, 53]. Our theory aims to capture the optimization characteristics and generalization performance of grid-based models. Different from neural tangent kernels (NTKs) [17] that study behaviors of MLPs, we target grid-based models and propose tangent kernels for them named grid tangent kernels (GTKs). GTKs are defined as the covariance between model gradients with respect to their parameters at two different input data. They describe precisely how changes in grid-based models' parameters impact their predictions during training. We show that GTKs of grid-based models remain unchanged during training so that a grid-based model behavior can be understood as a linear kernelized model when the kernel used in the GTK remains unchanged. Therefore, the GTK is particularly useful for understanding the connection between grid-based model architectures and their training dynamics. Furthermore, we derive a generalization bound based on Rademacher complexity [3], which measures how well a trained grid-based model performs on unseen data. This property helps researchers interested in neural fields understand the factors influencing generalization in grid-based models. Our numerical studies reveal that the performance of previous grid-based models can be escalated further by designing better structures toward a suitable GTK

*Corresponding author.

spectrum and a better generalization bound.

Guided by the theory of GTK, we propose a new grid-based model called the Multiplicative Fourier Adaptive Grid (MulFAGrid), which leverages multiplicative filters [11] to model nodal functions with constructed Fourier features, and then the extracted nodal features are normalized via a node-wise normalization function to form kernel functions. Finally, the features are gathered by element-wise multiplication between nodal features and extracted kernel features. An adaptive learning technique is adopted to optimize kernel features and grid features of MulFAGrid jointly. Our model supports both regular and irregular grids, depending on the choice of the index function. The technical comparison between our model and related methods is presented in Table 1, where we highlight our method is a general-purpose grid-based model, supporting a wide range of applications such as 2D image fitting [46], 3D SDF reconstruction [42] and view synthesis [29].

We methodologically analyze the characteristics of state-of-the-art grid-based models [9, 31, 54] and ours via the proposed GTK, and our mathematical findings indicate that the spectrum of MulFAGrid’s GTK is wider in the high-frequency domain, leading to better learning efficiency in learning high-frequency components. Furthermore, our numerical study in 2D toy examples shows that MulFAGrid has a tighter generalization bound in most regions of the 2D data plane than InstantNGP [31], NFFB [54] and NeuRBF [9]. Finally, the visualization of image regression experiments verifies the connection between the GTK analysis and predictive performance.

We then conduct systematic experiments on three fundamental tasks utilizing neural fields: 2D image fitting, 3D reconstruction with the signed distance field (SDF), and novel view synthesis via neural radiance fields (NeRF). In 2D image fitting and 3D SDF reconstruction, MulFAGrid achieves competitive performance with other grid-based models InstantNGP [31], NFFB [54] and NeuRBF [9]. In neural radiance fields, we experiment on five benchmarks ranging from bounded scenes and unbounded scenarios, and we find that MulFAGrid outperforms previous grid-based models by a notable margin. Moreover, despite the above experiments based on regular grids, we also experiment with non-regular grids (point clouds) and observe that MulFAGrid is competitive compared to the strong baseline 3DGS [18]. However, our rendering speed is much lower than 3DGS [18]. **Our contributions are:**

1. We present a theory based on tangent kernels for grid-based models. The proposed theory highlights that the grid tangent kernel (GTK) of grid-based models stays unchanged during training, and the GTK can characterize the generalization bound of grid-based models.
2. We propose a grid-based model based on multiplicative filters and Fourier features, and we propose an adaptive

learning approach to optimize the kernel features and grid features jointly.

3. We conduct GTK analysis for several state-of-the-art grid-based models, and the numerical studies show that MulFAGrid has a better generalization bound than other grid-based models.
4. Our empirical results show that MulFAGrid achieves state-of-the-art performances in 2D image fitting, 3D SDF reconstruction, and NeRF reconstruction, compared to grid-based models and other methods.

2. Related work

Neural Radiance Fields (NeRF). Neural radiance fields (NeRF) [29] proposes representing colors and densities via MLPs and learning the implicit 3D representations via differentiable volumetric rendering. After that, NeRF-based approaches dominate novel view synthesis [4, 5, 8, 26, 60]. NeRF has been applied to a wide range of topics, such as generation [32], surface reconstruction [51], and SLAM [39]. Previous works improve the efficiency of NeRFs in various aspects [16, 31, 35, 36]. FastNeRF [16] and Instant-NGP [31] apply advanced caching techniques to speed up the training of NeRFs. KiloNeRF [35] proposes decomposing the high-capacity MLP into thousands of small MLPs. Grid-based approaches [43, 44, 58] engage a lot of researchers [7, 24, 47, 50] because they are simple and fast.

Unbounded Scene Reconstruction. When depth cameras or multi-view stereos are available, one can use structure from motion (SfM) to reconstruct 3D scenes [1, 13–15, 49]. Recent researchers are interested in learning unbounded or large-scale radiance fields [5, 26, 26, 27, 37, 47, 48, 52, 55, 59]. Mip-NeRF-360 [5] improves the parameterization and efficiency of Mip-NeRF [4]. 3DGS [18] proposes to represent scenes with 3D Gaussians that achieve strong performance in free-view synthesis.

Neural Tangent Kernels (NTKs). Our analysis is closely related to NTK theories [3, 6, 10, 17, 22]. The neural tangent kernel [17, 22] finds that a wide network of any depth evolves as a linear model under gradient descent. Arora et al. [3] describe the generalization ability by the eigenvalues of the kernel. The closest work to ours is Matthew et al. [46], which shows that the Fourier feature mapping helps learn the high-frequency component of low-level representations.

Fourier Features. Random Fourier features (RFF) [34] can accelerate the training of kernel machines and are beneficial to large-scale classification and regression tasks. Fourier neural operators (FNO [21, 46]) demonstrate strong performance in the partial derivative equation (PDE) domain. Fourier feature mappings are critical in the original paper of NeRF [29]. Fourier PlenOctrees [50] demonstrates the effectiveness of Fourier feature learning in dy-

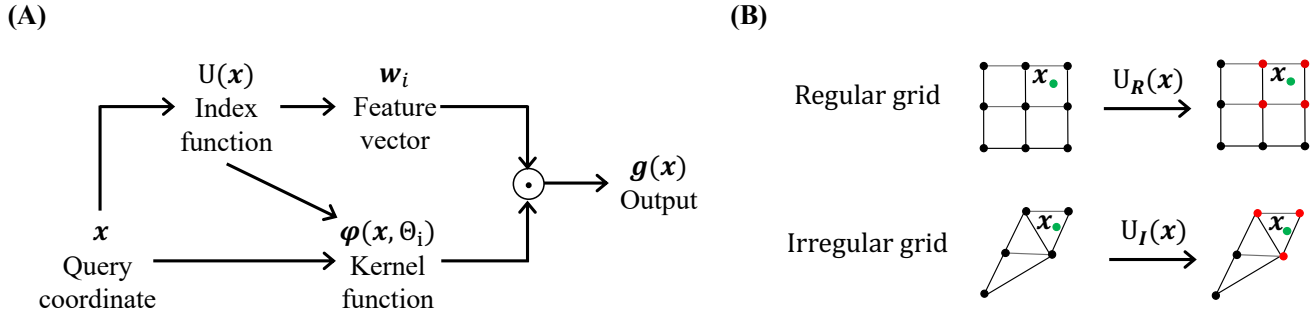


Figure 1. Formulations for grid-based models. (A) A grid-based model takes a query coordinate \mathbf{x} as the input, which is sent to an index function U to acquire a set of feature vectors \mathbf{w} from the grid. Then, the model outputs a weighted average of the kernel function φ and the feature vectors \mathbf{w} . (B) Our formulation supports grid-based models using a regular grid or an irregular grid, depending on the index function. The query coordinate is shown in green, and the queried points are in red. Please refer to Sec. 3.1.1 for more details.

Method	Publication venue	3D SDF reconstruction	2D image fitting	Novel view synthesis	Theoretical results	Grid-based models	Multiplicative filters	Support irregular grids
BACON [25]	CVPR22	✓	✓	✓	✓	✗	✗	✗
DVGO [43]	CVPR22	✗	✗	✓	✗	✓	✗	✗
Plenoxels [58]	CVPR22	✗	✗	✓	✗	✓	✗	✗
MINER [40]	ECCV22	✓	✓	✓	✗	✗	✗	✗
PNF [57]	NeurIPS22	✓	✓	✓	✓	✗	✗	✗
InstantNGP [31]	SIGGRAPH22	✓	✓	✓	✗	✓	✗	✗
3DGS [18]	SIGGRAPH23	✗	✗	✓	✗	✓	✗	✓
NFFB [54]	ICCV23	✓	✓	✓	✗	✓	✗	✗
NeuRBF [9]	ICCV23	✓	✓	✓	✗	✓	✗	✓
MulFAGrid (ours)	CVPR24	✓	✓	✓	✓	✓	✓	✓

Table 1. A detailed technical comparison among other recent neural field models (since 2022) and MulFAGrid.

dynamic scenes. NeuRBF [9] proposes a sinusoidal composition technique to fuse features of different frequencies.

3. Methodology

3.1. Understanding grid-based models

3.1.1 Formulations

Grid-based models play an important role in state-of-the-art neural-field methods because of their efficiency and scalable representation power [9, 18, 31, 43, 54, 58]. As shown in Figure 1, we define a grid-based model $g(\mathbf{x}, \mathbf{w})$ as a machine-learning model with the weighted-average form:

Definition 1. Given the input query coordinate \mathbf{x} , and φ as the kernel function of the grid-based model parameterized by Θ (Θ is an empty set if φ has no parameters), \mathbf{w}_i is the weight vector associated with the node i , and $U(\mathbf{x})$ is an index function which returns a set of indices given the location \mathbf{x} , a grid-based model is defined as a four-element tuple $\langle \varphi, \Theta, U, \mathbf{w} \rangle$ with the following computation:

$$g(\mathbf{x}, \mathbf{w}) = \sum_{i \in U(\mathbf{x})} \varphi(\mathbf{x}, \Theta_i) \mathbf{w}_i. \quad (1)$$

For example, in the simplest case of regular grids, $U(\mathbf{x})$ returns the eight nearest grid points around \mathbf{x} , φ is the bilinear interpolation with interpolation weights Θ , and \mathbf{w} denotes features stored in grid points. As shown in Figure 1 (B), the formulation in Equation (1) supports both *regular grid-based models* and *irregular grid-based models*, where the main difference between them is the index function U . The regular grid-based model uniformly discretizes the coordinate domain with equal intervals, which makes the index function fast. On the other hand, the irregular grid-based model leverages the geometric prior from point-cloud [18] or meshes [51] and does not regularly discretize the coordinate domain.

3.1.2 The grid tangent kernel (GTK) theory

To analyze existing grid-based models [9, 43, 54, 56, 58] for neural fields, we propose a theory for grid-based models inspired by the neural tangent kernel (NTK) [17]. Based on the formulation of grid-based models, we introduce a new term called *Grid Tangent Kernel (GTK)* to denote the tangent kernel of a grid-based model, which is crucial in understanding training and generalization performance of

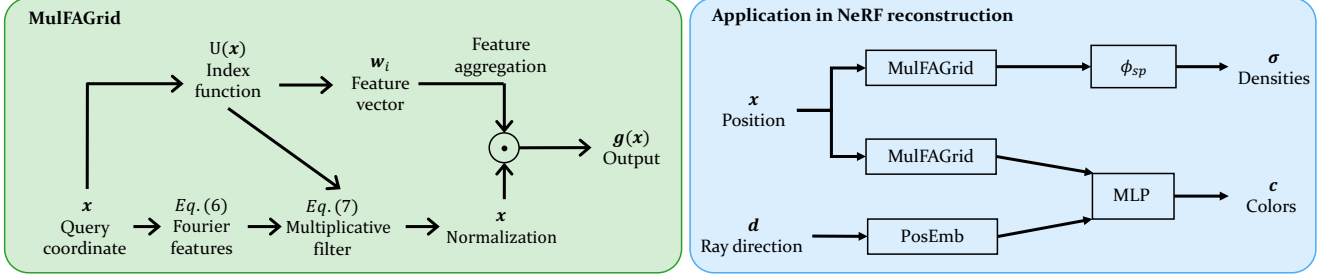


Figure 2. **(Left)** The diagram of MulFAGrid. The input query coordinate is passed to the multiplicative filter to produce Fourier features and then sent to the normalization layer to compute the aggregation weights. See Section 3.2 for details. **(Right)** The full architecture for neural radiance fields (NeRF). We obtain the densities σ via a MulFAGrid and the activation ϕ_{sp} . For the colors c , we encode the position x via the MulFAGrid with an MLP to post-process the features. After that, we combine the queried spatial features with ray direction information to get color predictions. Please refer to Section 4.4 for detailed explanations.

grid-based models.

Definition 2. Let \mathbf{X} be a collection of input data where \mathbf{X}_i is the i -th data, and $\mathbf{w}(t)$ is the weight at the training time t . The Grid Tangent Kernel (GTK) of the grid-based model g is defined as an $n \times n$ positive semidefinite matrix $\mathbf{G}_g(t)$ whose (i, j) -th element is:

$$[\mathbf{G}_g(t)]_{i,j} = \left\langle \frac{\partial g(\mathbf{X}_i, \mathbf{w}(t))}{\partial \mathbf{w}}, \frac{\partial g(\mathbf{X}_j, \mathbf{w}(t))}{\partial \mathbf{w}} \right\rangle. \quad (2)$$

With the GTK, we further introduce some theoretical results based on the GTK, while we put the proofs to those results in the Supplementary Section 7 due to space constraints. Our analysis follows a supervised regression setup [17, 46], where the task is to regress target labels \mathbf{Y} given the input data \mathbf{X} . Firstly, we introduce the following theorem characterizing the training dynamics:

Theorem 1. Let $\mathbf{O}(t) = (g(\mathbf{X}_i, \mathbf{w}(t)))_{1 \leq i \leq n}$ be the outputs of a grid-based model g where $\mathbf{X} = (\mathbf{X}_i)_{1 \leq i \leq n}$ is the input data at time t , and $\mathbf{Y} = (\mathbf{Y}_i)_{1 \leq i \leq n}$ is the corresponding label. Then $\mathbf{O}(t)$ follows this evolution:

$$\frac{d\mathbf{O}(t)}{dt} = -\mathbf{G}_g(t) \cdot (\mathbf{O}(t) - \mathbf{Y}). \quad (3)$$

According to the above theorem, GTK reflects model training dynamics. Therefore, analyzing GTK is critical to understanding the training behaviors of the grid-based model. Beyond this result, we further uncover a conservation property of the GTK of grid-based models:

Theorem 2. The GTK of a grid-based model g , denoted by \mathbf{G}_g , stays stationary during training. Formally, this property can be written as:

$$\mathbf{G}_g(t) = \mathbf{G}_g(0), \quad (4)$$

where $\mathbf{G}_g(0)$ is the initial GTK of the grid-based model. This property holds for any size of the grid-based model.

The above theorem shows that the GTK is a fundamental property of the grid-based model, which does not evolve through time. This means the evolution of the output $\mathbf{O}(t)$ is characterized by Equation (3), which is an ordinary differentiable equation (ODE). Therefore, grid-based models can be understood as simple linear models [3]. Besides training performance, another important property of a machine-learning model is the generalization gap. We further derive the following generalization bound based on the Rademacher complexity [3]:

Theorem 3. Given a probability $\delta_p \in (0, 1)$, suppose the dataset $S = (\mathbf{X}, \mathbf{Y})$ contains n i.i.d. samples from a distribution where $n \gg \log \frac{2}{\delta_p}$ and the minimum eigenvalue of the GTK, denoted by \mathbf{G} , is at least a constant λ_0 : $\lambda_{\min}(\mathbf{G}) \geq \lambda_0$. For any grid-based model g that is optimized by gradient descent with a learning rate η , and for any loss function $\mathcal{L} : \mathbb{R} \times \mathbb{R} \rightarrow [0, 1]$, which is 1-Lipschitz in the first argument, we define the population loss as $\mathcal{L}_{\mathcal{D}}(t) = \mathbb{E}_{(\mathbf{X}_i, \mathbf{Y}_i) \sim \mathcal{D}} [\mathcal{L}(g(\mathbf{X}_i, \mathbf{w}(t)), \mathbf{Y}_i)]$. Then, with probability at least $1 - \delta_p$, a randomly initialized grid-based model trained by gradient descent for $t \geq \Omega\left(\frac{1}{\eta \lambda_0} \log \frac{n}{\delta_p}\right)$ iterations has a generalization bound:

$$L_{\mathcal{D}}(t) \leq \sqrt{\frac{2\mathbf{Y}^\top \mathbf{G}^{-1} \mathbf{Y}}{n}} + O\left(\sqrt{\frac{\log \frac{2}{\delta_p}}{n}}\right). \quad (5)$$

From this theorem, we can find out that the dominating term affecting generalization for a given dataset with a fixed number of data n is $\Delta = \mathbf{Y}^\top \mathbf{G}^{-1} \mathbf{Y}$. If Δ is tighter, the generalization gap between training and testing data is smaller. This theorem provides another useful tool to assess grid-based models.

3.2. MulFAGrid

We propose a new grid-based model called multiplicative Fourier adaptive grid (MulFAGrid), where our numerical study in Section 4.1 suggests that it has a faster convergence

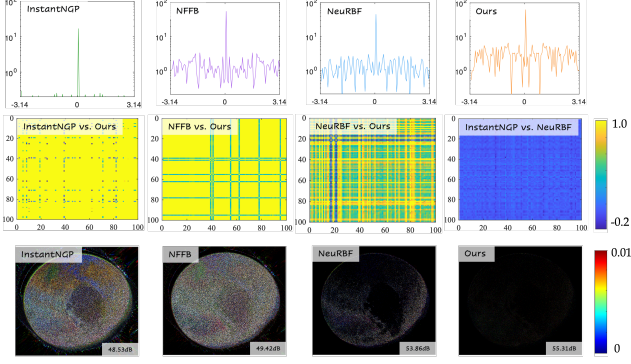


Figure 3. Analysis of grid-based models (InstantNGP [31], NFFB [54], NeuRBF [9], and ours) based on grid tangent kernels (GTKs) and image regression results. **(Top)** Visualizations of the GTK Fourier spectrum. MulFAGrid has a wide spectrum, especially in the high-frequency domain, leading to faster convergence for high-frequency components [46]. **(Mid)** Comparisons between generalization bounds of pairs of methods. In this experiment, we construct a dataset, which only contains two data points with labels $\mathbf{Y} = (\mathbf{Y}_1, \mathbf{Y}_2)$, shown in the x-axis and y-axis correspondingly. MulFAGrid has a tighter (lower) generalization bound for most values of \mathbf{Y}_1 and \mathbf{Y}_2 . These findings help explain why MulFAGrid demonstrates better representation ability than other grid-based models. **(Bot)** Error maps of the fitted images in comparison with ground truth ones.

speed and a tighter GTK-based generalization bound than state-of-the-art grid-based models. Our model is called an adaptive grid-based model because the kernel function φ is learned through data. We illustrate the design of MulFAGrid in Figure 2.

Fourier features. We construct a set of Fourier features $\mathbf{z}^{(1)}$ [29, 46] based on the input query location \mathbf{x} :

$$\gamma(\mathbf{x}, j) = \begin{cases} \sin(2^{\lfloor j/2 \rfloor} \pi \mathbf{x}), & \text{if } j \bmod 2 = 0, \\ \cos(2^{\lfloor j/2 \rfloor} \pi \mathbf{x}), & \text{else.} \end{cases} \quad (6a)$$

$$\mathbf{z}^{(1)} = \gamma(\mathbf{x}, 1) \oplus \gamma(\mathbf{x}, 2) \oplus \dots \oplus \gamma(\mathbf{x}, d_f), \quad (6b)$$

where d_f is a hyperparameter of the dimension of the constructed Fourier features, and \oplus denotes the concatenation operator. Our theory indicates that the input Fourier features are important in narrowing the generalization bound.

Multiplicative filters. We adopt multiplicative filters [11] to apply non-linearity to the constructed Fourier features and inform the model with the node index $i \in U(\mathbf{x})$. The output of multiplicative filters is denoted by $\tilde{\varphi}(\mathbf{x}, \Theta)$:

$$\begin{aligned} s(i; \theta^{(k)}) &= \sin(\omega^{(k)} i + \phi^{(k)}), \\ \mathbf{z}^{(k+1)} &= (W^{(k)} \mathbf{z}^{(k)} + b^{(k)}) \circ s(i; \theta^{(k+1)}), \\ \tilde{\varphi}(\mathbf{x}, \Theta_i) &= W^{(n_m)} \mathbf{z}^{(n_m)} + b^{(n_m)}. \end{aligned} \quad (7)$$

Here, $k = 1, 2, \dots, n_m - 1$ is the index of multiplicative filters, n_m is the total number of filters, s is the applied sinusoidal filter, $\Theta_i = \{\omega^{(k)}, \phi^{(k)}, W^{(k)}, b^{(k)}\}$ are parameters in the filters, and \circ denotes the element-wise multiplication.

Normalization and feature aggregation. We apply a node-wise normalization layer to acquire the final kernel:

$$\varphi(\mathbf{x}, \Theta_i) = \frac{\tilde{\varphi}(\mathbf{x}, \Theta_i)}{\sum_{i \in U(\mathbf{x})} \tilde{\varphi}(\mathbf{x}, \Theta_i)}. \quad (8)$$

With this kernel function, we produce the final representation of grid-based models via a feature aggregation procedure according to Equation (1).

Adaptive learning. We propose an adaptive learning method to adapt the kernel function to the data. In theory, kernel parameters Θ and features of the grid-based model \mathbf{w} can be iteratively optimized, similar to the expectation-maximization algorithm [9, 30]. Learning of the kernel parameters Θ contributes to a better GTK, which further narrows the generalization bound. The feature learning stage finds the best weight \mathbf{w} while the GTK is unchanged during this procedure (according to Theorem 2). While in experiments, we jointly optimize Θ and \mathbf{w} together for simplicity, and we find this work well in practice (see Section 4.5).

4. Experimental results

We first provide a numerical study based on the grid tangent kernel (GTK) to analyze the performance of several grid-based models and ours. Then, we evaluate our method and baselines on various applications using neural fields. For each application, we first describe the dataset and baselines and then provide qualitative and quantitative results. Choices of baselines for each task follow previous works [9, 54], and we wish to note that some baselines do not support all applications (see Table 1).

4.1. Numerical study based on the GTK

In this section, we provide various studies based on the GTK to analyze various grid-based models.

Baselines and data. We compare ours against several grid-based models: InstantNGP [31], NFFB [54], and NeuRBF [9]. Each grid-based model is trained until convergence, and we ensure our model does not contain more parameters than baselines. The numerical analysis is conducted in a 2D image dataset composed of the ultra-high-resolution image ‘‘Pluto’’ [9].

Spectrum analysis. Theorem 2 indicates that the GTK reflects a property of a grid-based model, which makes it possible to understand the training of grid-based models in functional space instead of parametric space [3]. Studies in the NTK domain [3, 46] show that the spectrum of the NTK reflects their convergence and generalization performance corresponding to different frequencies. Similarly, we conduct a spectrum analysis [10, 46] of grid-based models. We

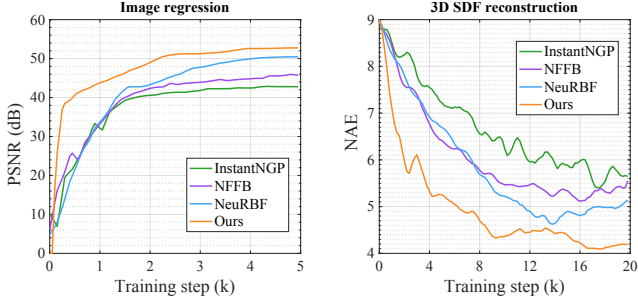


Figure 4. Comparison curves of several grid-based models: InstantNGP [31], NFFB [54], NeuRBF [9], and MulFAGrid. (Left) Training curves of the image regression task on the Kodak dataset [12]. (Right) The evolution of the normal angular error (NAE) through training of the 3D SDF reconstruction task [9].

	Steps	Time↓	# Params↓	PSNR↑
BACON [25]	5k	85.2s	268K	38.51
PNF [57]	5k	480.9s	287K	38.99
InstantNGP [31]	35k	1.9m	511K	39.14
MINER [40]	35k	14.2m	415K	39.25
NFFB [54]	5k	39.4s	154K	45.28
NeuRBF [9]	5k	28.5s	128K	54.84
Ours	5k	29.4s	119K	56.19

Table 2. 2D image fitting results on the validation set of DIV2K dataset [2]. Images are center cropped and down-sampled to $256 \times 256 \times 3$ following the practice of BACON [25] and NeuRBF [9].

	Steps	# Params↓	IoU↑	NAE↓
NGLOD [45]	245k	78.84M	0.9963	6.14
InstantNGP [31]	20k	950K	0.9994	5.70
NFFB [54]	20k	1.4M	0.9994	5.23
NeuRBF [9]	20k	856K	0.9995	4.93
Ours	5k	823K	0.9995	4.51

Table 3. 3D signed distance field (SDF) reconstruction results on the sampled 3D models dataset [9].

sample 100 data points from the training image to compute GTKs of different methods, where the ground-truth is the color value, and the calculation follows the GTK’s definition Equation (2). The value of GTK is normalized to the $[0, 1]$ range, and the Fourier spectrum of the GTK is shown in Figure 3. We find that InstantNGP [31] has a peak spectrum at the low frequency, corresponding to a very “narrow” kernel [17]. The drawback of a too-narrow kernel is overfitting to the training points [46]. Recent grid-based models such as NFFB [54] and NeuRBF [9] can mitigate such artifacts by introducing hierarchical structures [54] or radial basis functions [9]. Besides, we find that MulFAGrid has a wider GTK with higher values in high-frequency regions

so that it can learn better details. Therefore, MulFAGrid can balance better between “underfitting” and “overfitting” extremes [38].

Generalization bound comparisons. In Theorem 3, we show a theoretical result that provides an upper bound on the generalization error of a grid-based model. The GTK-based generalization bound can measure the generalization ability of a grid-based model without actually training it. As shown in Equation (5), the dominating term affecting generalization for a given dataset with a fixed number of data is $\Delta = \mathbf{Y}^\top \mathbf{G}^{-1} \mathbf{Y}$. We visualize the difference in generalization bounds between two grid-based models. Each plot characterizes the generalization bound difference $\Delta_A - \Delta_B$ between two grid-based models, A and B . We consider a two-point dataset where the ground true values of this dataset are $\mathbf{Y} = (\mathbf{Y}_1, \mathbf{Y}_2)$. We consider all possible values of the vector \mathbf{Y} and visualize the generalization bound difference $\Delta_A - \Delta_B$ in a 2D plane, where $\Delta_A - \Delta_B > 0$ means the generalization bound of the model B is better than that of model A . The result is shown in the middle of Figure 3. We observe that MulFAGrid has tighter generalization bounds than InstantNGP [31] and NFFB [54] in almost all possible data combinations. Meanwhile, the generalization bound of MulFAGrid is tighter than NeuRBF [9] in most regions. These facts indicate that MulFAGrid has better generalization performance than the compared grid-based models. We further show the image regression results at the bottom of Figure 3, which verifies that MulFAGrid has quicker optimization performance in modeling 2D neural fields than other grid-based models.

4.2. 2D image fitting

We then evaluate the task of fitting large-scale 2D images. The task goal is to fit an image based on spatial coordinates. The model is trained via the mean squared error, and all grid-based methods are based on regular grids in this task.

Baselines. Despite the mentioned grid-based models before (InstantNGP [31], NFFB [54] and NeuRBF [9]), we also compare to other recent works MINER [40], BACON [25] and PNF [57].

Dataset. Following previous work [9], we use the val split of the DIV2K dataset [2] consisting of 100 natural images of 2K resolution with a diversity of contents. Besides, we incorporate the Kodak dataset [12] whose images are sampled by the previous work NeuRBF [9].

Results. We present the learning curve of 2D image fitting, comparing to other grid-based models, in Figure 4. We observe that our method reaches a high PSNR at a faster speed than the state-of-the-art NeuRBF [9]. The quantitative results comparing a wide range of baselines are presented in Table 2. Our model achieves a better PSNR with fewer parameters, demonstrating the ability to represent target signals precisely.

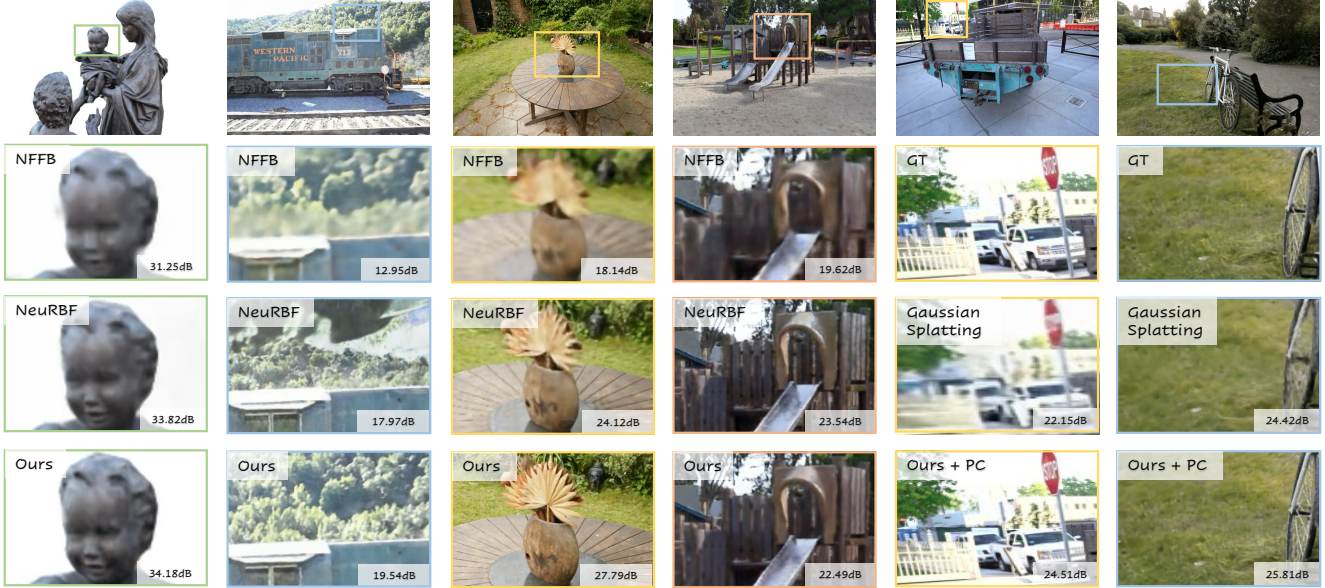


Figure 5. Rendering results on various scenes from SyntheticNeRF [29], Tanks&Temples [20] and Mip-NeRF-360 [5]. We visualize comparison results against some grid-based models NFFB [54], NeuRBF [9], and the strong baseline 3DGS [18] based on point cloud initialized from structure-from-motion (SfM).

Benchmark	SyntheticNeRF [29]			LLFF [28]			Tanks&Temples [20]			Mip-NeRF-360 [5]			SFMB [47]		
Method Metric	Time	#Params	PSNR	Time	#Params	PSNR	Time	#Params	PSNR	Time	#Params	PSNR	Time	#Params	PSNR
InstantNGP [31]	3.80m	12.2M	32.08	26.7m	13.2M	25.28	12.4m	11.3M	19.45	20.1m	12.4M	23.14	32.5m	10.9M	25.52
NFFB [54]	36.9m	18.5M	32.04	35.2m	16.3M	21.25	20.1m	15.9M	18.99	25.4m	19.4M	25.21	35.2m	16.2M	25.11
NeuRBF [9]	33.6m	17.7M	34.62	31.1m	18.7M	27.05	19.6m	16.0M	20.12	21.4m	15.2M	26.12	40.1m	9.24M	25.21
DVGOv2 [44]	10.9m	5.20M	32.80	10.9m	6.90M	26.34	21.4m	5.70M	20.10	20.3m	5.64M	25.42	37.8m	7.68M	26.42
Plenoxels [58]	11.4m	194M	31.71	24.5m	500M	26.29	20.8m	523M	20.40	21.9m	511M	20.59	39.4m	524M	25.98
Ours	10.3m	6.14M	34.68	24.2m	5.10M	27.22	19.2m	5.32M	20.85	20.4m	5.43M	28.98	33.1m	7.21M	29.15
3DGS [18]	10.7m	14.9M	33.15	26.5m	15.2M	27.25	27.0m	15.3M	22.22	41.2m	15.9M	27.21	35.1m	7.34M	29.44
Ours (w/ PC)	10.9m	13.3M	34.69	26.2m	16.1M	27.50	26.5m	16.2M	22.45	42.1m	16.5M	29.11	36.1m	7.34M	30.12

Table 4. Comparison results on five NeRF benchmarks. We report training time, the number of parameters, and PSNR. Reported numbers are averaged across all scenes in each benchmark. The bottom two models are based on point cloud initializations, while the rest do not use such information.

4.3. 3D signed distance fields reconstruction

We further validate our framework in reconstructing 3D models represented in 3D signed distance fields (SDF). In this application, all grid-based models use regular grids.

Baselines. We compare our method against NGLoD [45], targeted at learning SDF representations with an octree-based feature volume. We also compare InstantNGP [31], NFFB [54], and NeuRBF [9], where the hyper-parameters of those methods follow NeuRBF [9].

Dataset. We use ten commonly used 3D models to benchmark the performances of 3D SDF reconstruction models. Those models are sampled by NeuRBF [9] and originated from public repositories [23].

Results. We present the quantitative result of 3D SDF models in Table 3 and learning curves are presented in the

right of Figure 4. Our model has achieved strong performance in the normal angular error (NAE) with fewer training parameters. Visualization results are provided in Supplementary Section 9.

4.4. Novel view synthesis

Model details. The overall architecture for the view synthesis task is shown at the right of Figure 2. We model the density σ via one MulFAGrid:

$$\sigma = \phi_{sp}(\text{MulFAGrid}^{(1)}(\mathbf{x})), \quad (9a)$$

where ϕ_{sp} is the softplus activation [43]:

$$\phi_{sp}(\mathbf{x}) = \log(1 + \exp(\mathbf{x} + \eta_b)). \quad (9b)$$

The hyperparameter $\eta_b = 1 \times e^{-3}$ controls the empirical bias, which penalizes the number of non-zero elements in

grid tensors to alleviate overfitting [43]. We then model the color \mathbf{c} via two terms considering position-dependent colors and view-dependent colors,

$$\mathbf{c} = \text{MLP}(\text{MulFAGrid}^{(2)}(\mathbf{x}) \oplus \text{PosEmb}(\mathbf{d})), \quad (9c)$$

where $\text{MulFAGrid}^{(2)}$ denotes another MulFAGrid , and PosEmb refers to the positional embedding [43, 44]. After getting colors $(\mathbf{c}_i)_{1 \leq i \leq N}$ and densities $(\sigma_i)_{1 \leq i \leq N}$ for N segments, differentiable volumetric rendering [29, 43] is adopted to get the rendered pixel color.

Baselines. Besides previously mentioned baselines InstantNGP [31], NFFB [54] and NeuRBF [9], we also compare to other grid-based NeRF methods DVGOv2 [43, 44] and Plenoxels [58]. We adopt regular grid-based models in this comparison, which is called ‘‘Ours’’. We also provide a comparison with 3DGS [18], a recent very strong method in neural fields. Note that 3DGS [18] requires external structure-from-motion to construct its point cloud, so when comparing to 3DGS [18], we also leverage the same point cloud (PC) as initialization, which means irregular grids are adopted in this case. Our method is called ‘‘Ours (w/ PC)’’ in this case.

Benchmarks. We evaluate related methods on bounded NeRF benchmarks: LLFF [28] and SyntheticNeRF [29]. Besides, we also adopt two unbounded NeRF benchmarks: Tanks&Temples [20] and Mip-NeRF-360 [5]. Furthermore, we also conduct experiments on San Francisco Mission Bay (SFMB) provided by Block-NeRF [47]. Dataset details are presented in Supplementary Section 8.

Results. The quantitative comparison results on the five benchmarks are shown in Table 4. Our model consistently improves over NeuRBF [9] for all the scenes, while MulFAGrid keeps a good balance between training speed and effectiveness. Furthermore, our method achieves competitive performance compared to the strong baseline of 3DGS [18] when the point cloud initialization is derived by running the SfM initialization module of 3DGS [18]. However, the rendering speed of 3DGS [18] is still much faster than ours, partially due to their advanced CUDA implementation.

4.5. Ablation studies

We verify our design choices via ablation studies on the 2D image fitting dataset and the 3D SDF reconstruction dataset. Despite the mentioned changes, architectures and hyperparameters are kept the same. We train all models for 5k steps, and results are presented in Table 5. Numbers of NeuRBF [9] and our full model are presented in the last two rows for reference.

We validate the effectiveness of adaptive learning in the first row, and we use a simple interpolation kernel to replace the learned kernel. The performance is significantly downgraded in this case, revealing the necessity of learning the kernel function φ . Then, we remove the normalization

	2D image fitting		3D SDF reconstruction	
	PSNR \uparrow	SSIM \uparrow	IoU \uparrow	NAE \downarrow
Learned kernel \rightarrow interpolation	48.12	0.9930	0.9991	5.24
No normalization	56.01	0.9980	0.9995	4.81
Fourier features \rightarrow MSC	55.93	0.9978	0.9995	4.59
Fourier features \rightarrow SIREN	51.54	0.9958	0.9993	5.04
Joint learning \rightarrow decoupled	55.12	0.9980	0.9995	4.88
NeuRBF [9]	54.84	0.9975	0.9995	4.93
Ours full	56.19	0.9983	0.9995	4.51

Table 5. Ablation studies on 2D image fitting and 3D SDF reconstruction. We use the same validation dataset as that of NeuRBF [9]. Please refer to Section 4.5 for more details.

function (Equation (8)), and we find that the performance of the model slightly decreases. We replace the adopted Fourier features (Equation (6)) with the multi-frequency sinusoidal composition (MSC) proposed by NeuRBF [9], where it shows slightly worse performance than our full model. In the fourth row, we replace the Fourier features with SIREN [42] features, which demonstrates even worse performance. Lastly, we replace the joint optimization scheme of MulFAGrid with the decoupled learning method [9], and we find that it has worse performance, which demonstrates the benefit of jointly learning the kernel parameters with grid features.

5. Conclusion

We have proposed the grid tangent kernel (GTK) theory, which grounded grid-based models and helped analyze their optimization and generalization performance. We conducted a fine-grained analysis inspired by GTK to explain the behaviors of grid-based models. Then, we proposed MulFAGrid, a grid-based model for general neural field modeling. We proposed an adaptive learning scheme for MulFAGrid by jointly optimizing the kernel parameters and grid features. MulFAGrid balanced between ‘‘underfitting’’ and ‘‘overfitting’’ extremes and made more precise predictions. Meanwhile, MulFAGrid supported both regular grids and irregular grids. Experimental results on 2D image regression, 3D SDF reconstruction, and novel view synthesis demonstrated that MulFAGrid achieved state-of-the-art performance compared to various grid-based models. We hope our theoretical findings can provide insight for designing better grid-based models.

6. Acknowledgements

Dr. Fenglei Fan would like to acknowledge that this paper was supported by the Direct Grant for Research from the Chinese University of Hong Kong and ITS/173/22FP from the Innovation and Technology Fund of Hong Kong. This work was also partly supported by NSFC (62222607) and the Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102).

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. [2](#)
- [2] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2017. [6](#)
- [3] Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pages 322–332. PMLR, 2019. [1](#), [2](#), [4](#), [5](#), [3](#)
- [4] 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*, pages 5855–5864, 2021. [2](#), [5](#)
- [5] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. [2](#), [7](#), [8](#), [5](#)
- [6] Alberto Bietti and Julien Mairal. On the inductive bias of neural tangent kernels. *Advances in Neural Information Processing Systems*, 32, 2019. [1](#), [2](#)
- [7] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [8] Guojin Chen, Zehua Pei, Haoyu Yang, Yuzhe Ma, Bei Yu, and Martin Wong. Physics-informed optical kernel regression using complex-valued neural fields. In *Proceedings of the 60th ACM/IEEE Design Automation Conference*, 2023. [2](#)
- [9] Zhang Chen, Zhong Li, Liangchen Song, Lele Chen, Jingyi Yu, Junsong Yuan, and Yi Xu. Neurbf: A neural fields representation with adaptive radial basis functions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4182–4194, 2023. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [10] Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR, 2019. [2](#), [5](#)
- [11] Rizal Fathony, Anit Kumar Sahu, Devin Willmott, and J Zico Kolter. Multiplicative filter networks. In *International Conference on Learning Representations*, 2020. [2](#), [5](#)
- [12] Rich Franzen. Kodak lossless true color image suite. *source: http://r0k.us/graphics/kodak*, 4(2):9, 1999. [6](#)
- [13] Christian Früh and Avidesh Zakhor. An automated method for large-scale, ground-based city model acquisition. *International Journal of Computer Vision*, 60(1):5–24, 2004. [2](#)
- [14] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1362–1376, 2009.
- [15] Yasutaka Furukawa, Brian Curless, Steven M Seitz, and Richard Szeliski. Towards internet-scale multi-view stereo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1434–1441. IEEE, 2010. [2](#)
- [16] Stephan J. Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fast-nerf: High-fidelity neural rendering at 200fps. <https://arxiv.org/abs/2103.10380>, 2021. [2](#)
- [17] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in Neural Information Processing Systems*, 31, 2018. [1](#), [2](#), [3](#), [4](#), [6](#)
- [18] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023. [1](#), [2](#), [3](#), [7](#), [8](#), [4](#)
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [20] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. [7](#), [8](#), [5](#)
- [21] Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error bounds for fourier neural operators. *arXiv preprint arXiv:2107.07562*, 2021. [2](#)
- [22] Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [23] Marc Levoy, Kari Pulli, Brian Curless, Szymon Rusinkiewicz, David Koller, Lucas Pereira, Matt Ginzton, Sean Anderson, James Davis, Jeremy Ginsberg, et al. The digital michelangelo project: 3d scanning of large statues. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, pages 131–144, 2000. [7](#)
- [24] Ruofan Liang, Hongyi Sun, and Nandita Vijaykumar. Co-ordx: Accelerating implicit neural representation with a split mlp architecture. *arXiv preprint arXiv:2201.12425*, 2022. [2](#)
- [25] David B. Lindell, Dave Van Veen, Jeong Joon Park, and Gordon Wetzstein. Bacon: Band-limited coordinate networks for multiscale scene representation. In *CVPR*, 2022. [3](#), [6](#)
- [26] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM 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*, pages 7210–7219, 2021. [2](#)
- [27] Zhenxing Mi and Dan Xu. Switch-nerf: Learning scene decomposition with mixture of experts for large-scale neural radiance fields. In *International Conference on Learning Representations (ICLR)*, 2023. [2](#)
- [28] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and

- Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019. 7, 8
- [29] 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 *European Conference on Computer Vision*, pages 405–421. Springer, 2020. 1, 2, 5, 7, 8
- [30] Todd K Moon. The expectation-maximization algorithm. *IEEE Signal Processing Magazine*, 13(6):47–60, 1996. 5
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *CoRR*, abs/2201.05989, 2022. 1, 2, 3, 5, 6, 7, 8
- [32] 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*, pages 11453–11464, 2021. 2
- [33] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9054–9063, 2021. 1
- [34] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances in Neural Information Processing Systems*, 20, 2007. 2
- [35] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps, 2021. 2
- [36] Christian Reiser, Richard Szeliski, Dor Verbin, Pratul P Srinivasan, Ben Mildenhall, Andreas Geiger, Jonathan T Barron, and Peter Hedman. Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *arXiv preprint arXiv:2302.12249*, 2023. 2
- [37] Konstantinos Rematas, Andrew Liu, Pratul P Srinivasan, Jonathan T Barron, Andrea Tagliasacchi, Thomas Funkhouser, and Vittorio Ferrari. Urban radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12932–12942, 2022. 2
- [38] Daniel A Roberts, Sho Yaida, and Boris Hanin. *The principles of deep learning theory*. Cambridge University Press Cambridge, MA, USA, 2022. 6, 1, 2
- [39] Antoni Rosinol, John J Leonard, and Luca Carlone. Nerfslam: Real-time dense monocular slam with neural radiance fields. *arXiv preprint arXiv:2210.13641*, 2022. 2
- [40] Vishwanath Saragadam, Jasper Tan, Guha Balakrishnan, Richard G. Baraniuk, and Ashok Veeraraghavan. MINER: multiscale implicit neural representations. *CoRR*, abs/2202.03532, 2022. 3, 6
- [41] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 5
- [42] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *CoRR*, abs/2006.09661, 2020. 2, 8
- [43] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 1, 2, 3, 7, 8
- [44] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Improved direct voxel grid optimization for radiance fields reconstruction. *arXiv preprint arXiv:2206.05085*, 2022. 1, 2, 7, 8, 5
- [45] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11353–11362, Los Alamitos, CA, USA, 2021. IEEE Computer Society. 6, 7
- [46] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. 1, 2, 4, 5, 6
- [47] Matthew Tancik, Vincent Casser, Xinchen Yan, Sabeek Pradhan, Ben Mildenhall, Pratul Srinivasan, Jonathan T. Barron, and Henrik Kretzschmar. Block-NeRF: Scalable large scene neural view synthesis. *arXiv preprint arXiv:2202.05263*, 2022. 2, 7, 8, 4, 5
- [48] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12922–12931, 2022. 2
- [49] Michael Waechter, Nils Moehrl, and Michael Goesele. Let there be color! large-scale texturing of 3d reconstructions. In *European Conference on Computer Vision*, pages 836–850. Springer, 2014. 2
- [50] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Jingyi Yu, and Lan Xu. Fourier plenotrees for dynamic radiance field rendering in real-time. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13524–13534, 2022. 2
- [51] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. 1, 2, 3
- [52] Peng Wang, Yuan Liu, Zhaoxi Chen, Lingjie Liu, Ziwei Liu, Taku Komura, Christian Theobalt, and Wenping Wang. F2-nerf: Fast neural radiance field training with free camera trajectories. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4150–4159, 2023. 2
- [53] Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why pinns fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022. 1
- [54] Zhijie Wu, Yuhe Jin, and Kwang Moo Yi. Neural fourier filter bank. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14153–14163, 2023. 1, 2, 3, 5, 6, 7, 8

- [55] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. Citynerf: Building nerf at city scale. *arXiv preprint arXiv:2112.05504*, 2021. [2](#)
- [56] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *CoRR*, abs/2111.11426, 2021. [1](#), [3](#)
- [57] Guandao Yang, Sagie Benaim, Varun Jampani, Kyle Genova, Jonathan Barron, Thomas Funkhouser, Bharath Hariharan, and Serge Belongie. Polynomial neural fields for subband decomposition and manipulation. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022. [3](#), [6](#)
- [58] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. *arXiv preprint arXiv:2112.05131*, 2021. [1](#), [2](#), [3](#), [7](#), [8](#)
- [59] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. [2](#), [4](#)
- [60] Zelin Zhao and Jiaya Jia. End-to-end view synthesis via nerf attention, 2022. [2](#)