Factor Fields: A Unified Framework for Neural Fields and Beyond

Anpei Chen[1,4]   Zexiang Xu[2]   Xinyue Wei[3]   Siyu Tang[1]
Hao Su[3]   Andreas Geiger[4,5]

[1] ETH Zürich   [2] Adobe Research   [3] University of California, San Diego
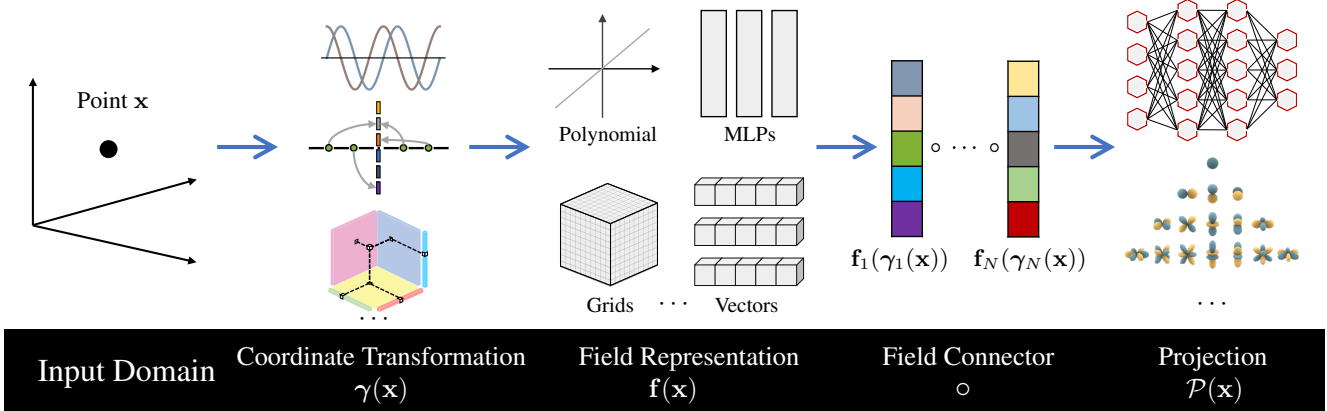[4] University of Tübingen   [5] Tübingen AI Center

Figure 1: **Factor Fields** is a framework which unifies many recently proposed neural field representations (e.g., NeRF, EG3D, Instant-NGP, TensoRF) and allows for the creation of powerful new ones such as the proposed **Dictionary Factorization**. In particular, Factor Fields decomposes a signal into N factors $\mathbf{f}_1$ to $\mathbf{f}_N$ (top-center), each of which is represented by one out of many different field representations operating on coordinate transformations $\boldsymbol{\gamma}_1$ to $\boldsymbol{\gamma}_N$. The resulting product field is passed to a projection function (e.g., MLP) which maps it to the signal's output.

Abstract

We present Factor Fields, a novel framework for modeling and representing signals. Factor Fields decomposes a signal into a product of factors, each represented by a classical or neural field representation which operates on transformed input coordinates. This decomposition results in a unified framework that accommodates several recent signal representations including NeRF, Plenoxels, EG3D, Instant-NGP, and TensoRF. Additionally, our framework allows for the creation of powerful new signal representations, such as the "Dictionary Field" (DiF) which is a second contribution of this paper. Our experiments show that DiF leads to improvements in approximation quality, compactness, and training time when compared to previous fast reconstruction methods. Experimentally, our representation achieves better image approximation quality on 2D image regression tasks, higher geometric quality when reconstructing 3D signed distance fields, and higher compactness for radiance field reconstruction tasks. Furthermore, DiF enables generalization to unseen images/3D scenes by sharing bases across signals during training which greatly benefits use cases such as image regression from sparse observations and few-shot radiance field reconstruction. Our code is available at https://apchenstu.github.io/FactorFields/.

1. Introduction

Effectively representing multi-dimensional digital content – like 2D images or 3D geometry and appearance – is critical for computer graphics and vision applications. These digital signals are traditionally represented discretely as pixels, voxels, textures, or polygons. Recently, significant headway has been made in developing advanced neural representations [36, 51, 37, 10, 49], which demonstrated superiority in modeling accuracy and efficiency over traditional representations for different image synthesis and scene reconstruction applications.

In order to gain a better understanding of existing representations, make comparisons across their design principles, and create powerful new representations, we propose *Factor Fields*, a novel mathematical framework that unifies many previous neural representations for multi-dimensional signals. This framework offers a simple formulation for modeling and representing signals.

Our framework decomposes a signal by factorizing it into multiple factor fields $(\mathbf{f}_1, \ldots, \mathbf{f}_N)$ operating on suitably chosen coordinate transformations $(\boldsymbol{\gamma}_1, \ldots, \boldsymbol{\gamma}_N)$ as illustrated in Fig. 1. More specifically, each factor field decodes multi-channel features at any spatial location of a coordinate-transformed signal domain. The target signal is then regressed from the factor product via a learned projection function (e.g., MLP).

Our framework accommodates most previous neural representations. Many of them can be represented in our framework as a single factor with a domain transformation – for example, the MLP network as a factor with a positional encoding transformation in NeRF [36], the tabular encoding as a factor with a hash transformation in Instant-NGP [37], and the feature grid as a factor with identity transformation in DVGO [51] and Plenoxels [16]. Recently, TensoRF [10] introduced a tensor factorization-based representation, which can be seen as a representation of two Vector-Matrix or three CANDECOMP-PARAFAC decomposition factors with axis-aligned orthogonal 2D and 1D projections as transformations. The potential of a multi-factor representation has been demonstrated by TensoRF, which has led to superior quality and efficiency on radiance field reconstruction and rendering, while being limited to orthogonal transformations.

This motivates us to generalize previous classic neural representations via a single unified framework which enables easy and flexible combinations of previous neural fields and transformation functions, yielding novel representation designs. As an example, we present *Dictionary Field (DiF)*, a two-factor representation that is composed of (1) a basis function factor with periodic transformation to model the commonalities of patterns that are shared across the entire signal domain and (2) a coefficient field factor with identity transformation to express localized spatially-varying features in the signal. The combination of both factors allows for an efficient representation of the global and local properties of the signal. Note that, most previous single-factor representations can be seen as using only one of such functions – either a basis function, like NeRF and Instant-NGP, or a coefficient function, like DVGO and Plenoxels. In DiF, jointly modeling two factors (basis and coefficients) leads to superior quality over previous methods like Instant-NGP and enables compact and fast reconstruction, as we demonstrate on various downstream tasks.

As DiF is a member of our general Factor Fields family,

we conduct a rich set of ablation experiments over the choice of basis/coefficient functions and basis transformations. We evaluate DiF against various variants and baselines on three classical signal representation tasks: 2D image regression, 3D SDF geometry reconstruction, and radiance field reconstruction for novel view synthesis. We demonstrate that our factorized DiF representation is able to achieve state-of-the-art reconstruction results that are better or on par with previous methods, while achieving superior modeling efficiency. For instance, compared to Instant-NGP our method leads to better reconstruction and rendering quality, while effectively *halving* the total number of model parameters (capacity) for SDF and radiance field reconstruction, demonstrating its superior accuracy and efficiency.

Moreover, in contrast to recent neural representations that are designed for purely per-scene optimization, our factorized representation framework is able to learn basis functions across different scenes. As shown in preliminary experiments, this enables learning across-scene bases from multiple 2D images or 3D radiance fields, leading to signal representations that generalize and hence improve reconstruction results from sparse observations such as in the few-shot radiance reconstruction setting. In summary,

- We introduce a common framework *Factor Fields* that encompasses many recent radiance field / signal representations and enables new models from the Factor Fields family.

- We propose *DiF*, a new member of the Factor Fields family representation that factorizes a signal into coefficient and basis factors which allows for exploiting similar signatures spatially and across scales.

- Our model can be trained jointly on multiple signals, recovering general basis functions that allow for reconstructing parts of a signal from sparse or weak observations.

- We present thorough experiments and ablation studies that demonstrate improved performance (accuracy, runtime, memory), and shed light on the performance improvements of DiF vs. other models in the Factor Fields family.

2. Factor Fields

We seek to compactly represent a continuous Q-dimensional signal $\mathbf{s} : \mathbb{R}^D \rightarrow \mathbb{R}^Q$ on a D-dimensional domain. We assume that signals are not random, but structured and hence share similar signatures within the same signal (spatially and across different scales) as well as between different signals. In the following, we develop our factor fields model step-by-step, starting from a standard basis expansion.
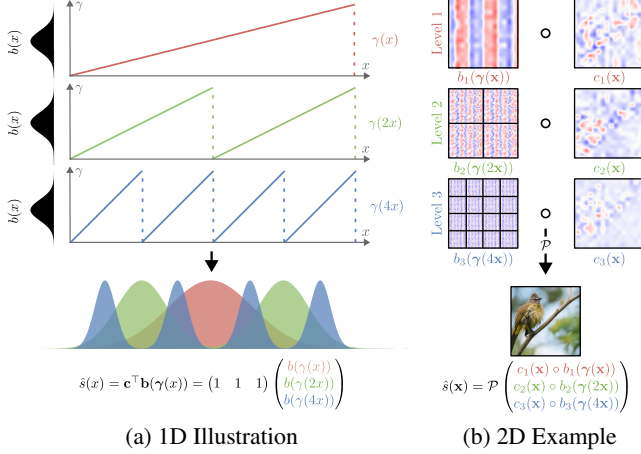
(a) 1D Illustration

$$\hat{s}(x) = \mathbf{c}^\top \mathbf{b}(\gamma(x)) = (1 \quad 1 \quad 1) \begin{pmatrix} b(\gamma(x)) \\ b(\gamma(2x)) \\ b(\gamma(4x)) \end{pmatrix}$$

(b) 2D Example

$$\hat{s}(\mathbf{x}) = \mathcal{P} \begin{pmatrix} c_1(\mathbf{x}) \circ b_1(\gamma(\mathbf{x})) \\ c_2(\mathbf{x}) \circ b_2(\gamma(2\mathbf{x})) \\ c_3(\mathbf{x}) \circ b_3(\gamma(4\mathbf{x})) \end{pmatrix}$$

Figure 2: **Local Basis.** (a) Choosing a (periodic) coordinate transformation $\boldsymbol{\gamma}(x)$ allows for applying the same basis function $b(x)$ at multiple spatial locations and scales. For clarity, we have chosen constant coefficients $\mathbf{c} = \mathbf{1}$. (b) Composing multiple bases at different spatial resolutions with their respective coefficients yields a powerful representation for signal $s(\mathbf{x})$. In practice, we use multiple bases and coefficient fields at each resolution.

Dictionary Field (DiF): Let us first consider a 1D signal $s(\mathbf{x}) : \mathbb{R}^D \to \mathbb{R}$. Using basis expansion, we decompose $s(\mathbf{x})$ into a set of coefficients $\mathbf{c} = (c_1, \ldots, c_K)^\top$ with $c_k \in \mathbb{R}$ and basis functions $\mathbf{b}(\mathbf{x}) = (b_1(\mathbf{x}), \ldots, b_K(\mathbf{x}))^\top$ with $b_k : \mathbb{R}^D \to \mathbb{R}$:

$$\hat{s}(\mathbf{x}) = \mathbf{c}^\top \mathbf{b}(\mathbf{x}) \tag{1}$$

Note that we denote $s(\mathbf{x})$ as the true signal and $\hat{s}(\mathbf{x})$ as its approximation.

Representing the signal $s(\mathbf{x})$ using a global set of basis functions is inefficient as information cannot be shared spatially. We hence generalize the above formulation by (i) exploiting a spatially varying coefficient field $\mathbf{c}(\mathbf{x}) = (c_1(\mathbf{x}), \ldots, c_K(\mathbf{x}))^\top$ with $c_k : \mathbb{R}^D \to \mathbb{R}$ and (ii) transforming the coordinates of the basis functions via a coordinate transformation function $\boldsymbol{\gamma} : \mathbb{R}^D \to \mathbb{R}^B$:

$$\hat{s}(\mathbf{x}) = \mathbf{c}(\mathbf{x})^\top \mathbf{b}(\boldsymbol{\gamma}(\mathbf{x})) \tag{2}$$

When choosing $\boldsymbol{\gamma}$ to be a periodic function, this formulation allows us to apply the same basis at multiple spatial locations and optionally also at multiple different scales while varying the coefficients $\mathbf{c}$, as illustrated in Fig. 2. Note that in general B does not need to match D, and hence the domain of the basis functions also changes accordingly: $b_k : \mathbb{R}^B \to \mathbb{R}$. Furthermore, we obtain Eq. (1) as a special case when setting $\mathbf{c}(\mathbf{x}) = \mathbf{c}$ and $\boldsymbol{\gamma}(\mathbf{x}) = \mathbf{x}$.

So far, we have considered a 1D signal $s(\mathbf{x})$. However, many signals have more than a single dimension (e.g., 3 in the case of RGB images or 4 in the case of radiance fields).

We generalize our model to Q-dimensional signals $\mathbf{f}(\mathbf{x})$ by introducing a projection function $\mathcal{P} : \mathbb{R}^K \to \mathbb{R}^Q$ and replacing the inner product with the element-wise/Hadamard product (denoted by $\circ$ in the following):

$$\hat{\mathbf{s}}(\mathbf{x}) = \mathcal{P}\left(\mathbf{c}(\mathbf{x}) \circ \mathbf{b}(\boldsymbol{\gamma}(\mathbf{x}))\right) \tag{3}$$

We refer to Eq. (3) as **Dictionary Field (DiF)**. Note that in contrast to the scalar product $\mathbf{c}^\top \mathbf{b}$ in Eq. (2), the output of $\mathbf{c} \circ \mathbf{b}$ is a K-dimensional vector which comprises the individual coefficient-basis products as input to the projection function $\mathcal{P}$ which itself can be either linear or non-linear. In the linear case, we have $\mathcal{P}(\mathbf{x}) = \mathbf{A}\mathbf{x}$ with $\mathbf{A} \in \mathbb{R}^{Q \times K}$. Furthermore, note that for $Q = 1$ and $\mathbf{A} = (1, \ldots, 1)$ we recover Eq. (2) as a special case. The projection operator $\mathcal{P}$ can also be utilized to model the volumetric rendering operation when reconstructing a 3D radiance field from 2D image observations as discussed in Section 2.3.

Factor Fields: To allow for more than 2 factors, we generalize Eq. (3) to our full Factor Fields framework by replacing the coefficients $\mathbf{c}(\mathbf{x})$ and basis $\mathbf{b}(\mathbf{x})$ with a set of factor fields $\{\mathbf{f}_i(\mathbf{x})\}_{i=1}^N$:

$$\hat{\mathbf{s}}(\mathbf{x}) = \mathcal{P}\left(\prod_{i=1}^N \mathbf{f}_i\left(\boldsymbol{\gamma}_i(\mathbf{x})\right)\right) \tag{4}$$

Here, $\prod$ denotes the element-wise product of a sequence of factors. Note that in this general form, each factor $\mathbf{f}_i : \mathbb{R}^{F_i} \to \mathbb{R}^K$ may be equipped with its own coordinate transformation $\boldsymbol{\gamma}_i : \mathbb{R}^D \to \mathbb{R}^{F_i}$.

We obtain **DiF** in Eq. (3) as a special case of our **Factor Fields** framework in Eq. (4) by setting $\mathbf{f}_1(\mathbf{x}) = \mathbf{c}(\mathbf{x})$, $\boldsymbol{\gamma}_1(\mathbf{x}) = \mathbf{x}$, $\mathbf{f}_2(\mathbf{x}) = \mathbf{b}(\mathbf{x})$ and $\boldsymbol{\gamma}_2(\mathbf{x}) = \boldsymbol{\gamma}(\mathbf{x})$ with $N = 2$. Besides DiF, the Factor Fields framework generalizes many recently proposed radiance field representations in one unified model as we will discuss in Section 3.

In our formulation, $\{\boldsymbol{\gamma}_i\}$ are considered deterministic functions while $\mathcal{P}$ and $\{\mathbf{f}_i\}$ are parametric mappings (e.g., polynomials, multi-layer perceptrons or 3D feature grids) whose parameters (collectively named θ below) are optimized. The parameters θ can be optimized either for a single signal or jointly for multiple signals. When optimizing for multiple signals jointly, we share the parameters of the projection function and basis factors (but not the parameters of the coefficient factors) across signals.

2.1. Factor Fields $\mathbf{f}_i$

For modeling the factor fields $\mathbf{f}_i : \mathbb{R}^{F_i} \to \mathbb{R}^K$, we consider various different representations in our Factor Fields framework as illustrated in Fig. 1 (bottom-left). In particular, we consider polynomials, MLPs, 2D and 3D feature grids and 1D feature vectors.

MLPs have been proposed as signal representations in Occupancy Networks [35], DeepSDF [40] and NeRF [36].
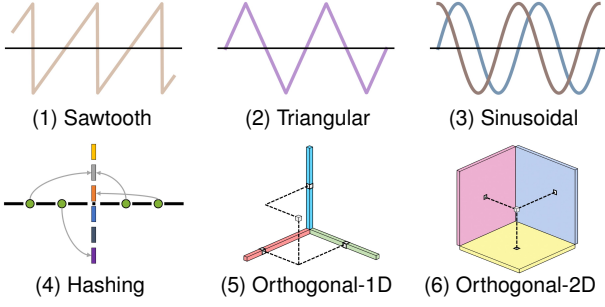
Figure 3: **Coordinate Transformations.** We show various periodic (top) and non-periodic (bottom) coordinate transformations γ used in our framework.

While MLPs excel in compactness and induce a useful smoothness bias, they are slow to evaluate and hence increase training and inference time. To address this, DVGO [51] proposes a 3D voxel grid representation for radiance fields. While voxel grids are fast to optimize, they increase memory significantly and do not easily scale to higher dimensions. To better capture the sparsity in the signal, Instant-NGP [37] proposes a hash function in combination with 1D feature vectors instead of a dense voxel grid, and TensoRF [10] decomposes the signal into matrix and vector products. Our *Factor Fields* framework allows any of the above representations to model any factor $\mathbf{f}_i$. As we will see in Section 3, many existing models are hence special cases of our framework.

2.2. Coordinate Transformations γ_i

The coordinates input to each factor field $\mathbf{f}_i$ are transformed by a coordinate transformation function $\gamma_i : \mathbb{R}^D \to \mathbb{R}^{F_i}$.

Coefficients: When the factor field $\mathbf{f}_i$ represents coefficients, we use the identity $\gamma_i(\mathbf{x}) = \mathbf{x}$ for the corresponding coordinate transformation since coefficients vary freely over the signal domain.

Local Basis: The coordinate transformation γ_i enables the application of the same basis function $\mathbf{f}_i$ at *multiple locations* as illustrated in Fig. 2. In this paper, we consider sawtooth, triangular, sinusoidal (as in NeRF [36]), hashing (as in Instant-NGP [37]) and orthogonal (as in TensoRF [10]) transformations in our framework, see Fig. 3.

Multi-scale Basis: The coordinate transformation γ_i also allows for applying the same basis $\mathbf{f}_i$ at *multiple spatial resolutions* of the signal by transforming the coordinates $\mathbf{x}$ with (periodic) transformations of different frequencies as illustrated in Fig. 2. This is crucial as signals typically carry both high and low frequencies, and we seek to exploit our basis representation across the full spectrum to model fine details of the signal as well as smooth signal components.

Specifically, we model the target signal with a set of multi-scale (field of view) basis functions. We arrange the basis into L levels where each level covers a different scale. Let $[\mathbf{u}, \mathbf{v}]$ denote the bounding box of the signal along one dimension. The corresponding scale is given by $(\mathbf{v} - \mathbf{u})/f_l$ where f_l is the frequency at level l. A large scale basis (e.g., level 1) has a low frequency and covers a large region of the target signal while a small scale basis (e.g., level L) has a large frequency f_L covering a small region of the target signal.

We implement our multi-scale representation (PR) by multiplying the scene coordinate $\mathbf{x}$ with the level frequency f_l before feeding it to the coordinate transformation function γ_i and then concatenating the results across the different scale $l = 1, \ldots, L$:

$$\gamma_i^{\text{PR}}(\mathbf{x}) = (\gamma_i(\mathbf{x}\, f_1), \ldots, \gamma_i(\mathbf{x}\, f_L)) \qquad (5)$$

Here, γ_i is any of the coordinate transformations in Fig. 3, and γ_{PR} is the final coordinate transform of our multi-scale representation.

As illustrated in Fig. 2 (b), when considering one coefficient factor $\mathbf{f}_1(\mathbf{x}) = \mathbf{c}(\mathbf{x})$ and one basis factor $\mathbf{f}_2(\mathbf{x}) = \mathbf{b}(\mathbf{x})$ with coordinate transformation $\gamma_2^{\text{PR}}(\mathbf{x})$ results in the target signal $\mathbf{s}(\mathbf{x})$ being decomposed as the product of spatial varying coefficient maps and multi-level basis maps which comprise repeated local basis functions.

2.3. Projection $\mathcal{P}$

To represent multi-dimensional signals, we introduced a projection function $\mathcal{P} : \mathbb{R}^K \to \mathbb{R}^Q$ that maps from the K-dimensional Hadamard product $\prod_i \mathbf{f}_i$ to the Q-dimensional target signal. We distinguish two cases in our framework: The case where direct observations from the target signal are available (e.g., pixels of an RGB image) and the indirect case where observations are projections of the target signal (e.g., pixels rendered from a radiance field).

Direct Observations: In the simplest case, the projection function realizes a learnable linear mapping $\mathcal{P}(\mathbf{x}) = \mathbf{A}\mathbf{x}$ with parameters $\mathbf{A} \in \mathbb{R}^{Q \times K}$ to map the K-dimensional Hadamard product $\prod_i \mathbf{f}_i$ to the Q-dimensional signal. However, a more flexible model is attained if $\mathcal{P}$ is represented by a shallow non-linear multi-layer perceptron (MLP) which is the default setting in all of our experiments.

Indirect Observations: In some cases, we only have access to *indirect* observations of the signal. For example, when optimizing neural radiance fields, we typically only observe 2D images instead of the 4D signal (density and radiance). In this case, extend $\mathcal{P}$ to also include the differentiable volumetric rendering process. More concretely, we first apply a multi-layer perceptron to map the view direction $\mathbf{d} \in \mathbb{R}^3$ and the product features $\prod_i \mathbf{f}_i$ at a particular location $\mathbf{x} \in \mathbb{R}^3$ to a color value $\mathbf{c} \in \mathbb{R}^3$ and a volume density $\sigma \in \mathbb{R}$.

Next, we follow Mildenhall et al. [36] and approximate the intractable volumetric projection integral using numerical integration. More formally, let $\{(\mathbf{c}_i, \sigma_i)\}_{i=1}^N$ denote the color and volume density values of N random samples along a camera ray. The RGB color value $\mathbf{c}$ at the corresponding pixel is obtained using alpha composition

$$\mathbf{c}_r = \sum_{i=1}^N T_i\,\alpha_i\,\mathbf{c}_i \quad T_i = \prod_{j=1}^{i-1}(1-\alpha_j) \quad \alpha_i = 1 - \exp\left(-\sigma_i\delta_i\right) \tag{6}$$

where T_i and α_i denote the transmittance and alpha value of sample i and $\delta_i = \|\mathbf{x}_{i+1} - \mathbf{x}_i\|_2$ is the distance between neighboring samples. The *composition* of the learned MLP and the volume rendering function in Eq. (6) constitute the projection function $\mathcal{P}$.

2.4. Space Contraction

We normalize the input coordinates $\mathbf{x} \in \mathbb{R}^D$ to $[0,1]$ before passing them to the coordinate transformations $\gamma_i(\mathbf{x})$ by applying a simple space contraction function to $\mathbf{x}$. We distinguish two settings:

For **bounded signals** with D-dimensional bounding box $[\mathbf{u}, \mathbf{v}]$ (where $\mathbf{u}, \mathbf{v} \in \mathbb{R}^D$), we utilize a simple linear mapping to normalize all coordinates to the range $[0,1]$:

$$\mathrm{contract}(\mathbf{x}) = \frac{\mathbf{x} - \mathbf{u}}{\mathbf{v} - \mathbf{u}} \tag{7}$$

For **unbounded signals** (e.g., an outdoor radiance field), we adopt Mip-NeRF 360's [3] space contraction function:

$$\mathrm{contract}(\mathbf{x}) = \begin{cases} \mathbf{x} & \|\mathbf{x}\|_2 \le 1 \\ \left(2 - \frac{1}{\|\mathbf{x}\|_2}\right)\left(\frac{\mathbf{x}}{\|\mathbf{x}\|_2}\right) & \|\mathbf{x}\|_2 > 1 \end{cases} \tag{8}$$

2.5. Optimization

Given samples $\{(\mathbf{x}, \mathbf{s}(\mathbf{x}))\}$ from the signal, we minimize

$$\underset{\theta}{\mathrm{argmin}}\ \mathbb{E}_\mathbf{x}\left[\|\mathbf{s}(\mathbf{x}) - \hat{\mathbf{s}}_\theta(\mathbf{x})\|_2 + \Psi(\theta)\right] \tag{9}$$

where $\Psi(\theta)$ is a regularizer on the model parameters. We optimize this objective using stochastic gradient descent.

Sparsity Regularization: While using the ℓ_0 norm for sparse coefficients is desirable, this leads to a difficult optimization problem. Instead, we use a simpler strategy which we found to work surprisingly well. We regularize our objective by randomly dropping a subset of the K features of our model by setting them to zero with probability μ. This forces the signal to be represented with random combinations of features at every iteration, encouraging sparsity and preventing co-adaptation of features. We implement this dropout

---

In our implementation, we slightly modify Eq. (8) to map coordinates to a unit ball centered at 0.5 which avoids negative coordinates when indexing feature grids.

regularizer using a random binary vector which we multiply element-wise with the factor product: $\circ \prod_i \mathbf{f}_i$.

Initialization: During all our experiments, we initialize the basis factors using the discrete cosine transform (DCT), while initializing the parameters of the coefficient factors and projection MLP randomly. We experimentally found this to improve the quality of the solution as illustrated in our ablation study in Table 3a to Table 3e.

Multiple Signals: When optimizing for multiple signals jointly, we share the parameters of the projection function and basis factors (but not the parameters of the coefficient factors) across signals. As evidenced by our experiments in Section 4.3, sharing bases across different signals while encouraging sparse coefficients improves generalization and enables reconstruction from sparse observations.

3. Factor Fields As A Common Framework

Advanced neural representations have emerged as a promising replacement for traditional representations and been applied to improve the reconstruction quality and efficiency in various graphics and vision applications, such as novel view synthesis [65, 32, 54, 2, 36, 31, 10, 58, 55], generative models [47, 9, 8, 17], 3D surface reconstruction [39, 60, 56, 62, 25], image processing [13], graphics asset modeling [45, 27, 66], inverse rendering [5, 4, 63, 6, 7, 64], dynamic scene modeling [44, 30, 41, 29], and scene understanding [42, 34] amongst others.

Inspired by classical factorization and learning techniques, like sparse coding [57, 59, 18] and principal component analysis (PCA) [46, 33], we propose a novel neural factorization-based framework for neural representations. Our Factor Fields framework unifies many recently published neural representations and enables the instantiation of new models in the Factor Fields family, which, as we will see, exhibit desirable properties in terms of approximation quality, model compactness, optimization time and generalization capabilities. In this section, we will discuss the relationship to prior work in more detail. A systematic performance comparison of the various Factor Field model instantiations is provided in Section 4.4.

Occupancy Networks, IMNet and DeepSDF: [35, 14, 40] represent the surface implicitly as the continuous decision boundary of an MLP classifier or by regressing a signed distance value. The vanilla MLP representation provides a continuous implicit 3D mapping, allowing for the extraction of 3D meshes at any resolution. This setting corresponds to our Factor Fields model when using a single factor (i.e., $N = 1$) with $\gamma_1(\mathbf{x}) = \mathbf{x}$, $\mathbf{f}_1(\mathbf{x}) = \mathbf{x}$, $\mathcal{P}(\mathbf{x}) = \mathrm{MLP}(\mathbf{x})$, thus $\hat{\mathbf{s}}(\mathbf{x}) = \mathrm{MLP}(\mathbf{x})$. While this representation is able to generate high-quality meshes, it fails to model high-frequency signals, such as images due to the implicit smoothness bias

of MLPs.

NeRF: [36] proposes to represent a radiance field via an MLP in Fourier space by encoding the spatial coordinates with a set of sinusoidal functions. This corresponds to our Factor Fields setting when using a single factor with $\gamma_1(\mathbf{x}) = (\sin(\mathbf{x}f_1), \cos(\mathbf{x}f_1), \ldots, \sin(\mathbf{x}f_L), \cos(\mathbf{x}f_L))$, $\mathbf{f}_1(\mathbf{x}) = \mathbf{x}$ and $\mathcal{P}(\mathbf{x}) = \text{MLP}(\mathbf{x})$. Here, the coordinate transformation $\gamma_1(\mathbf{x})$ is a sinusoidal mapping as shown in Fig. 3 (3), which enables high frequencies.

Plenoxels: [16] use sparse voxel grids to represent 3D scenes, allowing for direct optimization without neural networks, resulting in fast training. This corresponds to our Factor Fields framework when setting $N = 1$, $\gamma_1(\mathbf{x}) = \mathbf{x}$, $\mathbf{f}_1(\mathbf{x}) = \text{3D-Grid}(\mathbf{x})$, and $\mathcal{P}(\mathbf{x}) = \mathbf{x}$ for the density field while $\mathcal{P}(\mathbf{x}) = \text{SH}(\mathbf{x})$ (Spherical Harmonics) for the radiance field. In related work, DVGO [51] proposes a similar design, but replaces the sparse 3D grids with dense grids and uses a tiny MLP as the projection function $\mathcal{P}$. While dense grid modeling is simple and leads to fast feature queries, it requires high spatial resolution (and hence large memory) to represent fine details. Moreover, optimization is more easily affected by local minima compared to MLP representations that benefit from their inductive smoothness bias.

ConvONet and EG3D: [43, 8] use a tri-plane representation to model 3D scenes by applying an orthogonal coordinate transformation to spatial points within a bounded scene, and then representing each point as the concatenation of features queried from a set of 2D feature maps. This representation allows for aggregating 3D features using only 2D convolution, which significantly reduces memory footprint compared to standard 3D grids. The setting can be viewed as an instance of our Factor Fields framework, with $N = 1$, $\gamma_1(\mathbf{x}) = \text{Orthogonal-2D}(\mathbf{x})$, $\mathbf{f}_1(\mathbf{x}) = \text{2D-Maps}(\mathbf{x})$ and $\mathcal{P}(\mathbf{x}) = \text{MLP}(\mathbf{x})$. However, while the axis-aligned transformation allows for dimension reduction and feature sharing along the axis, it can be challenging to handle complicated structures due to the axis-aligned bias of this representation.

Instant-NGP: [37] exploits a multi-level hash grid to efficiently model internal features of target signals by hashing spatial locations to 1D feature vectors. This approach corresponds to our Factor Fields framework when using $N = 1$, $\gamma_1(\mathbf{x}) = \text{Hashing}(\mathbf{x})$, $\mathbf{f}_1(\mathbf{x}) = \text{Vectors}(\mathbf{x})$ and $\mathcal{P}(\mathbf{x}) = \text{MLP}(\mathbf{x})$ with $L = 16$ scales. However, the multi-level hash mappings can result in dense collisions at fine scales, and the one-to-many mapping forces the model to distribute its capacity bias towards densely observed regions and noise in areas with fewer observations. The concurrent work VQAD [52] introduces a hierarchical vector-quantized auto-decoder (VQ-AD) representation that learns an index table as the coordinate transformation function which allows

for higher compression rates.

TensoRF: [10] factorizes the radiance fields into the products of vectors and matrices (TensoRF-VM) or multiple vectors (TensoRF-CP), achieving efficient feature queries at low memory footprint. This setting instantiates our Factor Fields framework for $N = 2$, $\gamma_1(\mathbf{x}) = \text{Orthogonal-1D}(\mathbf{x})$, $\mathbf{f}_1(\mathbf{x}) = \text{Vectors}(\mathbf{x})$, $\gamma_2(\mathbf{x}) = \text{Orthogonal-2D}(\mathbf{x})$, $\mathbf{f}_2(\mathbf{x}) = \text{2D-Maps}(\mathbf{x})$ for VM decomposition and $N = 3$, $\gamma_i(\mathbf{x}) = \text{Orthogonal-1D}(\mathbf{x})$, $\mathbf{f}_i(\mathbf{x}) = \text{Vectors}(\mathbf{x})$ for CP decomposition. Moreover, TensoRF uses both SH and MLP models for the projection $\mathcal{P}$. Similar to ConvONet and EG3D, TensoRF is sensitive to the orientation of the coordinate system due to the use of an orthogonal coordinate transformation function. Note that, with the exception of TensoRF [10], all of the above representations factorize the signal using a single factor field, that is $N = 1$. As we will show in Table 3a to Table 3d, using multiple factor fields (i.e., $N > 1$) provides stronger model capacity.

ArXiv Preprints: The field of neural representation learning is advancing fast and many novel representations have been published as preprints on ArXiv recently. We now briefly discuss the most related ones and their relationship to our work: *Phasorial Embedding Fields (PREF)* [19] proposes to represent a target signal with a set of phasor volumes and then transforms them into the spatial domain with an inverse fast Fourier Transform (iFFT) for compact representation and efficient scene editing. This method shares a similar idea with DVGO and extends the projection function $\mathcal{P}$ in our formulation with an iFFT function. *Tensor4D* [48] extends the triplane representation to 4D human reconstruction by using 3 triplane Factor Fields and indexing plane features with 3 orthogonal coordinate transformations (i.e., $\text{Orthogonal-2D}(\mathbf{x}) = (xy, xt, yt)$, $\text{Orthogonal-2D}(\mathbf{x}) = (xz, xt, zt)$, $\text{Orthogonal-2D}(\mathbf{x}) = (yz, yt, zt)$). *NeRF-Player* [50] represents dynamic scenes via deformation, newness and decomposition fields, using multiple factors similar to TensoRF. It further extends the features by a time dimension. *D-TensoRF* [20] reconstructs dynamic scenes using matrix-matrix factorization, similar to the VM factorization of TensoRF, but replacing $\gamma_1(\mathbf{x}) = \text{Orthogonal-1D}(\mathbf{x})$ and $\mathbf{f}_1(\mathbf{x}) = \text{Vectors}(\mathbf{x})$ with $\gamma_1(\mathbf{x}) = \text{Orthogonal-2D}(\mathbf{x})$ and $\mathbf{f}_1(\mathbf{x}) = \text{2D-Maps}(\mathbf{x})$. *Quantized Fourier Features (QFF)* [28] factorizes internal features into bins of Fourier features, corresponding to our Factor Fields framework when using $N = 1$, $\gamma_1(\mathbf{x}) = \text{sinusoidal}(\mathbf{x})$, $\mathbf{f}_1(\mathbf{x}) = \text{2D-Maps}(\mathbf{x})$, and $\mathcal{P}(\mathbf{x}) = \text{MLP}(\mathbf{x})$ for the 2D signal representation.

Dictionary Field (DiF): Besides existing representations, our Factor Fields framework enables the design of novel representations with desirable properties. As an example, we now discuss the DiF representation which we have already introduced in Eq. (3). DiF offers implicit regularization, compactness and fast optimization while also generalizing

across multiple signals. The central idea behind the DiF representation is to decompose the target signals into two fields: a global field (i.e., the basis) and a local field (i.e., the coefficient). The global field promotes structured signal signatures shared across spatial locations and scales as well as between signals, while the local field allows for spatially varying content. More formally, DiF factorizes the target signal into coefficient fields $\mathbf{f}_1(\mathbf{x}) = \mathbf{c}(\mathbf{x})$ and basis functions $\mathbf{f}_2(\mathbf{x}) = \mathbf{b}(\mathbf{x})$ which differ primarily by their respective coordinate transformation: We choose the identity mapping $\boldsymbol{\gamma}_1(\mathbf{x}) = \mathbf{x}$ for $\mathbf{c}(\mathbf{x})$ and a periodic coordinate transformation $\boldsymbol{\gamma}_2(\mathbf{x})$ for $\mathbf{b}(\mathbf{x})$, see Fig. 3 (top). As representation of the two factor fields $\mathbf{f}_1$ and $\mathbf{f}_2$ we may choose any of the ones illustrated in Fig. 1 (bottom-left). To facilitate comparison with previous representations, we use the sawtooth function as the basis coordinate transformation $\boldsymbol{\gamma}_2$ and uniform grids (i.e., 2D Maps for 2D signals and 3D Grids for 3D signals) as the representation of the coefficient fields $\mathbf{f}_1$ and basis functions $\mathbf{f}_2$ for most of our experiments. Besides, we also systematically ablate the number of factors, the number of levels, the coordinate transformation and the field representation in Section 4.4.

4. Experiments

We now present extensive evaluations of our Factor Fields framework and DiF representation. We first briefly discuss our implementation and hyperparameter configuration. We then compare the performance of DiF with previously proposed representations on both per-signal reconstruction (optimization) and across-signal generalization tasks. At the end of this section, we examine the properties of our Factor Fields framework by varying the number of factors N, level number L, different types of transformation functions $\boldsymbol{\gamma}_i$ field representation $\mathbf{f}_i$, and field connector $\circ$.

4.1. Implementation

We implement our Factor Fields framework using vanilla PyTorch without customized CUDA kernels. Performance is evaluated on a single RTX 6000 GPU using the Adam optimizer [23] with a learning rate of 0.02.

We instantiate DiF using $L = 6$ levels with frequencies (linearly increasing) $f_l \in [2., 3.2, 4.4, 5.6, 6.8, 8.]$, and feature channels $K = [4, 4, 4, 2, 2, 2]^\top \cdot 2^\eta$, where η controls the number of feature channels. We use $\eta = 3$ for our 2D experiments and $\eta = 0$ for our 3D experiments. The model parameters θ are distributed across 3 model components: coefficients $\theta_{\mathbf{c}}$, basis $\theta_{\mathbf{b}}$, and projection function $\theta_{\mathcal{P}}$. The size of each component can vary greatly depending on the chosen representation.

In the following experiments, we refer to the default model setting as "DiF-Grid", which implements the coefficients $\mathbf{c}$ and bases $\mathbf{b}$ with learnable tensor grids, $\mathcal{P}(\mathbf{x}) = \text{MLP}(\mathbf{x})$, and $\boldsymbol{\gamma}(\mathbf{x}) = \text{Sawtooth}(\mathbf{x})$, where Sawtooth$(\mathbf{x}) =$

$\mathbf{x}$ mod 1.0. In the DiF-Grid setting, the total number of optimizable parameters is mainly determined by the resolution of the coefficient and basis grids $M_{\mathbf{c}}^l$, $M_{\mathbf{b}}^l$:

$$|\theta| = |\theta_{\mathcal{P}}| + |\theta_{\mathbf{c}}| + |\theta_{\mathbf{b}}| = |\theta_{\mathcal{P}}| + \sum_{l=1}^{L} {M_{\mathbf{c}}^l}^D + K_l \cdot {M_{\mathbf{b}}^l}^D \tag{10}$$

We implement the basis grid using linearly increasing resolutions $M_{\mathbf{b}}^l \in [32, 128]^T \cdot \frac{min(\mathbf{v} - \mathbf{u})}{1024}$ with interval $[32, 128]$ and scene bounding box $[u, v]$. This leads to increased resolution for modeling higher-resolution signals in our experiments. We use the same coefficient grid resolution $M_{\mathbf{c}}^l$ across all L levels for query efficiency and to lower per-signal memory footprint.

In the following, the different model variants of our Dictionary factorization are labeled "DiF-xx", where "xx" indicates the differences from the default setting "DiF-Grid". For example, "-MLP-B" refers to using an MLP basis representation, and "SL" stands for single level.

4.2. Single Signals

We first evaluate the accuracy and efficiency of our DiF-Grid representation on various multi-dimensional signals, comparing it to several recent neural signal representations. Towards this goal, we consider three popular benchmark tasks for evaluating neural representations: 2D image regression, 3D Signed Distance Field (SDF) reconstruction and Radiance Field Reconstruction / Novel View Synthesis. We evaluate each method's ability to approximate high-frequency patterns, interpolation quality, compactness, and robustness to ambiguities and sparse observations.

2D Image Regression: In this task, we directly regress RGB pixel colors from pixel coordinates. We evaluate our DiF-Grid on fitting four complex high-resolution images, where the total number of pixels ranges from $4\,\text{M}$ to $213\,\text{M}$. In Fig. 4, we show the reconstructed images with the corresponding model size, optimization time, and image PSNRs, and compare them to Instant-NGP [37], a state-of-the-art neural representation that supports image regression and has shown superior quality over prior art including Fourier Feature Networks [53] and SIREN [49]. Compared to Instant-NGP, our model consistently achieves higher PSNR on all images when using the same model size, demonstrating the superior accuracy and efficiency of our model. On the other hand, while Instant-NGP achieves faster optimization owing to its highly optimized CUDA-based framework, our model, implemented in pure PyTorch, leads to comparably fast training while relying on a vanilla PyTorch implementation without custom CUDA kernels which simplifies future extensions.

Signed-Distance Field Reconstruction: Signed Distance

| Summer Day | Albert | Pluto | Girl with a Pearl Earring |

$6114 \times 3734 \times 3$ (resolution) / $35.42\,\mathrm{M}$ (params)
$0{:}46$ vs. $4{:}13$ (mm:ss) / 42.4 vs. 49.0 dB (PSNR)

$1024 \times 1024 \times 4$ / $1.36\,\mathrm{M}$
$0{:}30$ vs. $1{:}13$ / 51.0 vs. 62.7

$8000 \times 8000 \times 3$ / $58.60\,\mathrm{M}$
$0{:}50$ vs. $5{:}32$ / 44.30 vs. 46.19

$8000 \times 9302 \times 3$ / $68.52\,\mathrm{M}$
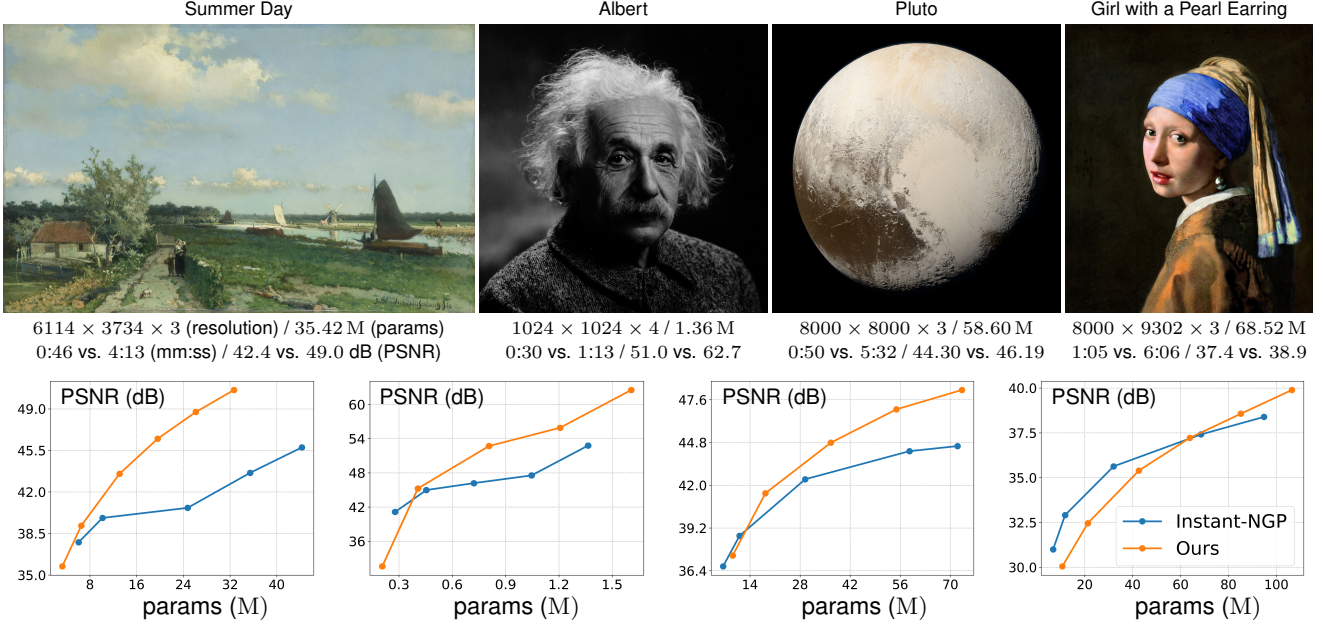$1{:}05$ vs. $6{:}06$ / 37.4 vs. 38.9

Figure 4: **2D Image Regression.** This figure shows images represented using our DiF-Grid model. The respective image resolutions and numbers of model parameters are shown below each image. Moreover, we also report a comparison to Instant-NGP (first number) in terms of optimization time and PSNR metrics (Instant-NGP vs Ours) at the bottom using the same number of model parameters. Note that our method achieves better reconstruction quality on all images when using the same model size. While optimization is slower than Instant-NGP, we use a vanilla PyTorch implementation without customized CUDA kernels. "Summer Day" credit goes to Johan Hendrik Weissenbruch and rijksmuseum. "Albert" credit goes to Orren Jack Turner. "Pluto" credit goes to NASA. "Girl With a Pearl Earring" renovation ©Koorosh Orooj (CC BY-SA 4.0).

Function (SDF), as a classic geometry representation, describes a set of continuous iso-surfaces, where a 3D surface is represented as the zero level-set of the function. We evaluate our DiF-Grid on modeling several challenging object SDFs that contain rich geometric details and compare with previous state-of-the-art neural representations, including Fourier Feature Networks [53], SIREN [49], and Instant-NGP [37]. To allow for fair comparisons in terms of the training set and convergence, we use the same training points for all methods by pre-sampling $8\,\mathrm{M}$ SDF points from the target meshes for training, with 80% points near the surface and the remaining 20% points uniformly distributed inside the unit volume. Following the evaluation setting of Instant-NGP, we randomly sample $16\,\mathrm{M}$ points for evaluation and calculate the geometric IOU metric based on the SDF sign

$$gIoU = \frac{\sum(s(\mathbf{X}) > 0) \cap (\hat{s}(\mathbf{X}) > 0)}{\sum(s(\mathbf{X}) > 0) \cup (\hat{s}(\mathbf{X}) > 0)} \qquad (11)$$

where $\mathbf{X}$ is the evaluation point set, $s(\mathbf{X})$ are the ground truth SDF values, and $\hat{s}(\mathbf{X})$ are the predicted SDF values.

Fig. 5 shows a quantitative and qualitative comparison of all methods. Our method leads to visually better results, it recovers high-frequency geometric details and contains less noise on smooth surfaces (e.g., the elephant face). The high visual quality is also reflected by the highest gIoU value

of all methods. Meanwhile, our method also achieves the fastest reconstruction speed, while using less than half of the number of parameters used by CUDA-kernel enabled Instant-NGP, demonstrating the high accuracy, efficiency, and compactness of our factorized representation.

Radiance Field Reconstruction: Radiance field reconstruction aims to recover the 3D density and radiance of each volume point from as multi-view RGB images. The geometry and appearance properties are updated via inverse volume rendering, as proposed in NeRF [36]. Recently, many encoding functions and advanced representations have been proposed that significantly improve reconstruction speed and quality, such as sparse voxel grids [16], hash tables[37] and tensor decomposition[10].

In Table 1, we quantitatively compare DiF-Grid with several state-of-the-art fast radiance field reconstruction methods (Plenoxel [16], DVGO [51], Instant-NGP [37] and TensoRF-VM [10]) on both synthetic [36] as well as real scenes (Tanks and Temple objects) [24]. Our method achieves high reconstruction quality, significantly outperforming NeRF, Plenoxels, and DVGO on both datasets, while being significantly more compact than Plenoxels and DVGO. We also outperform Instant-NGP and are on par with TensoRF regarding reconstruction quality, while being highly compact with only $5.1\,\mathrm{M}$ parameters, less than one-third of

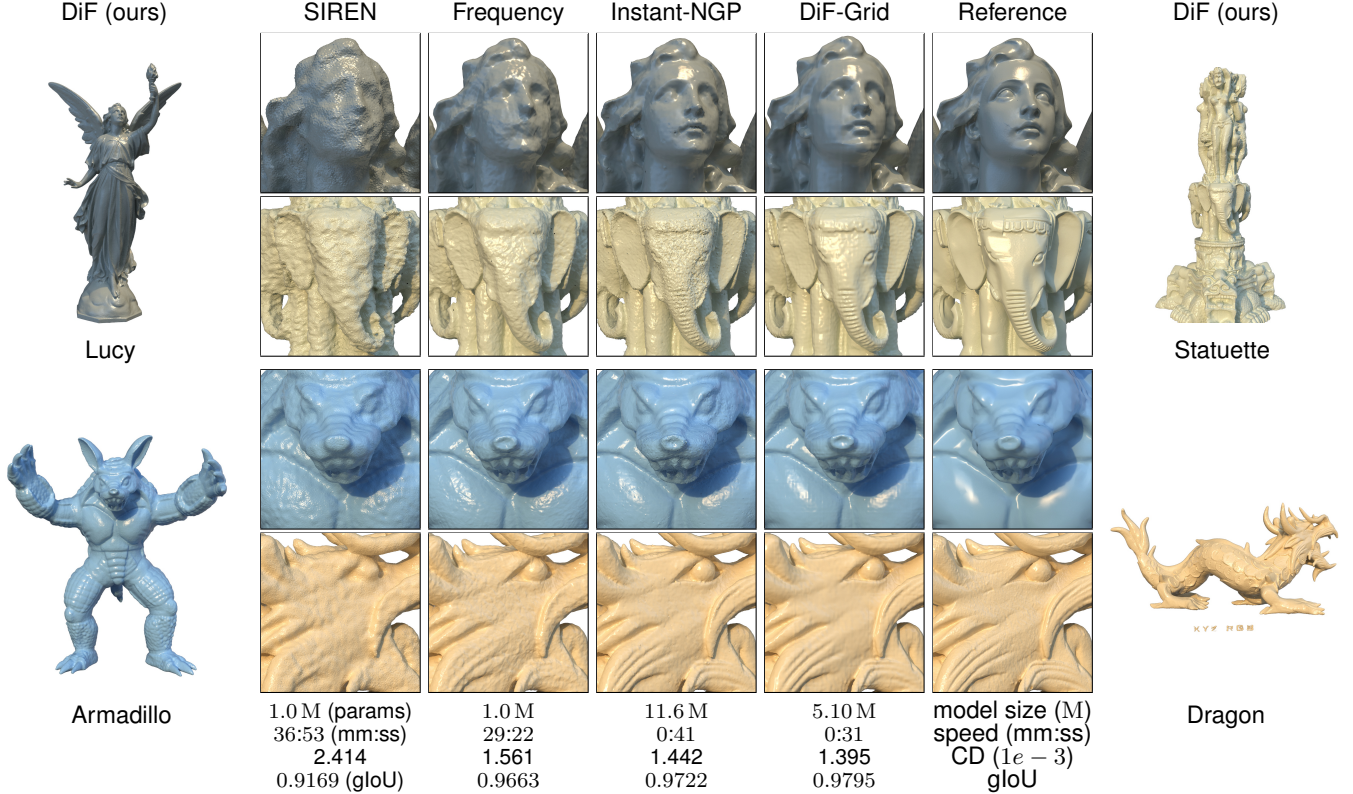| DiF (ours)	SIREN	Frequency	Instant-NGP	DiF-Grid	Reference	DiF (ours)
Lucy	1.0 M (params)	1.0 M	11.6 M	5.10 M	model size (M)	Statuette
	36:53 (mm:ss)	29:22	0:41	0:31	speed (mm:ss)	
Armadillo	2.414	1.561	1.442	1.395	CD ($1e-3$)	Dragon
	0.9169 (gIoU)	0.9663	0.9722	0.9795	gIoU	

Figure 5: **Signed-Distance Field Reconstruction.** We reconstruct SDFs from 8.0 M training points. We show qualitative visual comparisons on the top and quantitative comparisons on the bottom including the number of parameters, reconstruction time and gIoU. DiF-Grid and iNGP [37] are trained for $10k$ iterations, while SIREN [49] and NeRF with Frequency Encodings [53] are trained for $200k$ iterations.

TensoRF-VM and one-half of Instant-NGP. Our DiF-Grid also optimizes faster than TensoRF, at slightly over 10 minutes, in addition to our superior compactness. Additionally, unlike Plenoxels and Instant-NGP which rely on their own CUDA framework for fast reconstruction, our implementation uses the standard PyTorch framework, making it easily extendable to other tasks.

In general, our model leads to state-of-the-art results on all three challenging benchmark tasks with both high accuracy and efficiency. Note that the baselines are mostly single-factor, utilizing either a local field (such as DVGO and Plenoxels) or a global field (such as Instant-NGP). In contrast, our DiF model is a two-factor method, incorporating both local coefficient and global basis fields, hence resulting in better reconstruction quality and memory efficiency.

4.3. Generalization

Recent advanced neural representations such as NeRF, SIREN, ACORN, Plenoxels, Instant-NGP and TensoRF optimize each signal separately, lacking the ability to model multiple signals jointly or learning useful priors from mul-

Figure 6: **Radiance Field Reconstruction**. We evaluate our DiF using NeRF-Synthetic and Tanks and Temples datasets, our method is able to reconstruct high-quality surface details.

tiple signals. In contrast, our DiF representation not only enables accurate and efficient per-signal reconstruction (as demonstrated in Section 4.2) but it can also be applied to generalize across signals by simply sharing the basis field across signal instances. We evaluate the benefits of basis sharing by conducting experiments on image regression from partial pixel observations and few-shot radiance field reconstruction.

Method	BatchSize	Steps	Synthetic-NeRF				Tanks and Temples	
			Time ↓	Size (M)↓	PSNR↑	SSIM↑	PSNR↑	SSIM↑
NeRF [36]	4096	300k	∼35h	01.25 ●	31.01	0.947	25.78	0.864
Plenoxels [16]	5000	128k	11.4m ○	194.5	31.71	0.958	27.43	0.906
DVGO [51]	5000	30k	15.0m	153.0	31.95 ●	0.957	28.41 ●	0.911 ●
Instant-NGP [37]	10k-85k	30k	03.9m ●	11.64 ●	32.59 ○	0.960 ●	27.09	0.905
TensoRF-VM [10]	4096	30k	17.4m	17.95	33.14 ●	0.963 ●	28.56 ○	0.920 ○
DiF-Grid (Ours)	4096	30k	12.2m ●	05.10 ○	33.14 ●	0.961 ○	29.00 ●	0.938 ●

Table 1: **Novel View Synthesis with Radiance Fields.** We compare our method to previous radiance field reconstruction methods on the Synthetic-NeRF[36] and Tanks and Temples [24] datasets. We report the scores reported in the original papers whenever available. We also show average reconstruction time and model size for the Synthetic-NeRF dataset to compare the efficiency of the methods.

For these experiments, instead of DiF-Grid, we adopt DiF-MLP-B (i.e., (5) in the Table 3d) as our DiF representation, where we utilize a tensor grid to model the coefficient and 6 tiny MLPs (two layers with 32 neurons each) to model the basis. We find that DiF-MLP-B performs better than DiF-Grid in the generalization setting, owing to the strong inductive smoothness bias of MLPs.

Image Regression from Sparse Observations: Unlike the image regression experiments conducted in Sec. 4.2 which use all image pixels as observations during optimization, this experiment focuses on the scenario where only part of the pixels are used during optimization. Without additional priors, a single-signal optimization easily overfits in this setting due to the sparse observations and the limited inductive bias, hence failing to recover the unseen pixels.

We use our DiF-MLP-B model to learn data priors by pre-training it on 800 facial images from the FFHQ dataset[21] while sharing the MLP basis and projection function parameters. The final image reconstruction task is conducted by optimizing the coefficient grids for each new test image.

In Fig. 7, we show the image regression results on three different facial images with various masks and compare them to baseline methods that do not use any data priors, including Instant-NGP and our DiF-MLP-B without pre-training. As expected, Instant-NGP can accurately approximate the training pixels but results in random noise in the untrained mask regions. Interestingly, even without pre-training and priors from other images, our DiF-MLP-B is able to capture structural information to some extent within the same image being optimized; as shown in the eye region, the model can learn the pupil shape from the right eye and regress the left eye (masked during training) by reusing the learned structures in the shared basis functions. As shown on the right of Fig. 7, our DiF-MLP-B with pre-trained prior clearly achieves the best reconstruction quality with better structures and boundary smoothness compared to the baselines, demonstrating that our factorized DiF model allows for learning
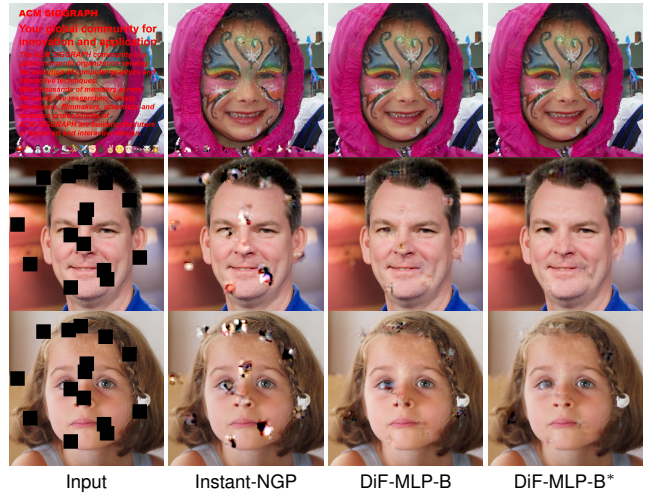

Figure 7: **Image Regression from Sparse Observations.** Results obtained by fitting each model to all unmasked pixels. We use randomly placed black squares as masks for the bottom two rows and an image of text and small icons as mask for the top row. The symbol * denotes pre-training of the basis factors using the FFHQ facial image set. Our pre-trained model (DiF-MLP-B*) learns robust basis fields which lead to better reconstruction compared to the per-scene baselines Instant-NGP and DiF-MLP-B.

and transferring useful prior information from the training set.

Few-Shot Radiance Field Reconstruction: Reconstructing radiance fields from few-shot input images with sparse viewpoints is highly challenging. Previous works address this by imposing sparsity assumptions [38, 22] in per-scene optimization or training feed-forward networks [61, 12, 26] from datasets. Here we consider 3 and 5 input views per scene and seek a novel solution that leverages data priors in pre-trained basis fields of our DiF model during the optimization task. It is worth-noting that the views are chosen in

a quarter sphere, thus the overlapping region between views is quite limited.

Specifically, we first train DiF models on 100 Google Scanned Object scenes [15], which contains 250 views per scene. During cross-scene training, we maintain 100 per-scene coefficients and share the basis c and projection function $\mathcal{P}$. After cross-scene training, we use the mean coefficient values of pre-trained coefficient fields as the initialization, while fixing the pre-trained functions (c and $\mathcal{P}$) and fine-tuning the coefficient field for new scenes with few-shot observations. In this experiment, we compare results from both DiF-MLP-B and DiF-Grid with and without the pre-training. We also compare with Instant-NGP and previous few-shot reconstruction methods, including PixelNeRF [61] and MVSNeRF [11], re-train with the same training set and test using the same 3 or 5 views. As shown in Table 2 and Fig. 8, our pre-trained DiF representation with MLP basis provides strong regularization for few-shot reconstruction, resulting in fewer artifacts and better reconstruction quality than the single-scene optimization methods without data priors and previous few-shot reconstruction methods that also use pre-trained networks. In particular, without any data priors, single-scene optimization methods (Instant-NGP and ours w/o prior) lead to a lot of outliers due to overfitting to the few-shot input images. Previous methods like MVSNeRF and PixelNeRF achieve plausible reconstructions due to their learned feed-forward prediction which avoids per-scene optimization. However, they suffer from blurry artifacts. Additionally, the strategy taken by PixelNeRF and MVSNeRF assumes a narrow baseline and learns correspondences across views for generalization via feature averaging or cost volume modeling which does not work as effectively in a wide baseline setup. On the other hand, by pre-training shared basis fields on multiple signals, our DiF model can learn useful data priors, enabling the reconstruction of novel signals from sparse observations via optimization.

4.4. Influence of Design Choices in Factor Fields

Factor Fields is a general framework that unifies many previous representations with different settings. In this section, we aim to analyze the properties of these variations and offer a comprehensive understanding of the components of the proposed representation framework. We conduct extensive evaluations on the four main components of our Factor Fields framework: factor number N, level number L, coordinate transformation function γ_i, field representation f_i, and field connector $\circ$.

We present a comprehensive assessment of the representations' capabilities in terms of efficiency, compactness, reconstruction quality, as well as generalizability, with a range of tasks including 2D image regression (with all pixels), and per-scene and across-scene 3D radiance field reconstruction. Note that, the settings in per-scene and across-scene radiance

		3 views		5 views	
Method	Time↓	PSNR↑	SSIM↑	PSNR↑	SSIM↑
iNGP	03:38	14.74	0.776	20.79	0.860
DiF-Grid	13:39	18.13	0.805	20.83	0.847
DiF-MLP-B	18:24	16.31	0.804	22.26	0.900
PixelNeRF	00:00	21.37	0.878	22.73	0.896
MVSNeRF	00:00	20.50	0.868	22.76	0.891
PixelNeRF-ft	25:18	22.21	0.882	23.67	0.895
MVSNeRF-ft	13:06	18.51	0.864	20.49	0.887
DiF-Grid*	13:18	20.77	0.871	25.41	0.915
DiF-MLP-B*	18:44	21.96	0.891	26.91	0.927

Table 2: **Few-shot Radiance Field Reconstruction.** We show quantitative comparisons of few-shot radiance field reconstruction from 3 or 5 viewpoints regarding optimization time and novel view synthesis quality (PSNRs and SSIMs). Results are averaged across 8 test scenes. The results of Instant-NGP and our DiF models are generated based on per-scene optimization, while DiF models with * use pre-trained basis factors across scenes. We train the feed-forward networks of PixelNeRF and MVSNeRF using the same dataset we learn our shared basis factors, and the results of Pixel-NeRF and MVSNeRF are generated from the networks via direct feed-forward inference. Our DiF-MLP-B* with pre-trained MLP basis factors leads to the best reconstruction quality.

field reconstruction are the same as introduced in Section 4.2 and Section 4.3, while for the 2D image regression task, we use the same model setting as in Section 4.2 and test on 256 high fidelity images at a resolution of 1024×1024 from the DIV2K dataset [1]. To enable meaningful comparisons, we evaluate the variations within the same code base and report their performance using the same number of iterations number, batch size, training point sampling strategy and pyramid frequencies. Correspondingly, the results for Instant-NGP, EG3D, OccNet, NeRF, DVGO, TensoRF-VM/-CP in Tab. 3 are based on our reimplementation of the original methods in our unified Factor Fields framework with the corresponding design parameters shown in the tables.

Factor Number N: As illustrated in Eq. (4), the factor number N refers to the number of factors used to represent the target signal. We show comparisons between single-factor models (1,3,5) and two-factor models (2,4,6) in Table 3a. Compared to the models (1) iNGP, (3) EG3D and (5) DiF-no-C which only use a single factor, the models (2) DiF-Hash-B, (4) TensoRF-VM and (6) DiF-Grid use the same factor as (1), (3), (5) respectively, but are extended to two-factor models with an additional factor, leading to 3dB and 0.35 dB PSNR improvement in image regression and 3D radiance field reconstruction tasks, while also increasing
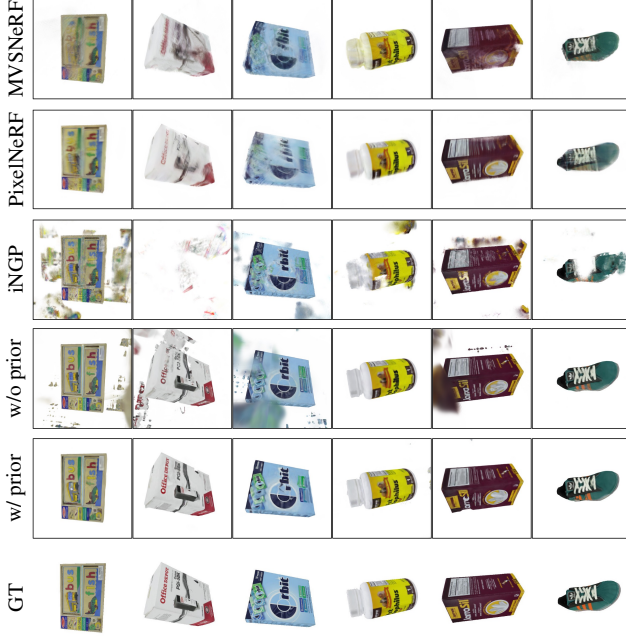
Figure 8: **Radiance Fields from 5 Views.** We visualize novel view synthesis results of six test scenes, corresponding to the quantitative results in Tab. 2. We show our DiF-MLP-B model w/ and w/o pre-trained data priors (bottom two rows) and compare it to Instant-NGP, PixelNeRF and MVSNeRF (top three rows). Our model with pre-trained basis factors can effectively utilize the learned data priors, resulting in superior qualitative results with fewer outliers compared to single-scene models (iNGP and ours w/o priors), as well as sharper details compared to feed-forward models (PixelNeRF and MVSNeRF).

training time ($\sim 10\%$) and model size ($\sim 5\%$) as expected. Despite the marginal computational overhead, introducing additional factors in the framework clearly leads to better reconstruction quality and represents the signal more effectively. The multiplication between factors allows the two factor fields to modulate each other's feature encoding and represent the entire signal more flexibly in a joint manner, alleviating the problem of feature collision in Instant-NGP and related problems of other single-factor models. Additionally, as shown in Table 1, multi-factor modeling (e.g., $N >= 2$) is able to provide more compact modeling while maintaining a similar level of reconstruction quality. Moreover, it allows for generalization by partially sharing the fields across instances, such as the across-scene radiance field modeling in the few-shot radiance reconstruction task.

Level Number L**:** Our DiF model adopts multiple levels of transformations to achieve pyramid basis fields, similar to the usage of a set of sinusoidal positional encoding functions in NeRF [36]. We compare multi-level models (including DiF and NeRF) with their reduced single-level versions that only use a single transformation level in Table 3b. Note that Occupancy Networks (OccNet, row (1)) do not leverage positional encodings and can be seen as a single-level version of NeRF (row (2)) while the model with multi-level sinusoidal encoding functions (NeRF) leads to about 10dB PSNR performance boost for both 2D image and 3D reconstruction tasks. On the other hand, the single-level DiF models are also consistently worse than the corresponding multi-level models in terms of speed and reconstruction quality, despite the performance drops being not as severe as those in purely MLP-based representations.

Coordinate Transformation γ_i**:** In Table 3c, we evaluate four coordinate transformation functions using our DiF representation. These transformation functions include sinusoidal, triangular, hashing and sawtooth. Their transformation curves are shown in Fig. 3. In general, in contrast to the random hashing function, the periodic transformation functions (2, 3, 4) allow for spatially coherent information sharing through repeated patterns, where neighboring points can share spatially adjacent features in the basis fields, hence preserving local connectivity. We observe that the periodic basis achieves clearly better performance in modeling dense signals (e.g., 2D images). For sparse signals such as 3D radiance fields, all four transformation functions achieve high reconstruction quality on par with previous state-of-the-art fast radiance field reconstruction approaches [51, 37, 10].

Field Representation f_i**:** In Table 3d, we compare various functions for representing the factors in our framework (especially our DiF model) including MLPs, Vectors, 2D Maps and 3D Grids, encompassing most previous representations. Note that discrete feature grid functions (3D Grids, 2D Maps, and Vectors) generally lead to faster reconstruction than MLP functions (e.g. DiF-Grid is faster than DiF-MLP-Band DiF-MLP-C). While all variants can lead to reasonable reconstruction quality for single-signal optimization, our DiF-Grid representation that uses grids for both factors achieves the best performance on the image regression and single-scene radiance field reconstruction tasks. On the other hand, the task of few-shot radiance field reconstruction benefits from basis functions that impose stronger regularization. Therefore, representations with stronger inductive biases (e.g., the Vectors in TensoRF-VM and MLPs in DiF-MLP-B) lead to better reconstruction quality compared to other variants.

Field Connector $\circ$**:** Another key design choice of our Factor Fields framework and DiF model is to adopt the element-wise product to connect multiple factors. Directly concatenating features from different components is an alternative choice and exercised in several previous works [36, 8, 37]. In Table 3e, we compare the performance of

the element-wised product against the direct concatenation in three model variants. Note that the element-wise product consistently outperforms the concatenation operation in terms of reconstruction quality for all models on all applications, demonstrating the effectiveness of using the proposed product-based factorization framework.

5. Conclusion and Future Work

In this work, we present a novel unified framework for (neural) signal representations which factorizes a signal as a product of multiple factor fields. We demonstrate that Factor Fields generalizes many previous neural field representations (like NeRF, Instant-NGP, DVGO, TensoRF) and enables new representation designs. In particular, we propose a novel representation – DiF – with Dictionary factorization, as a new model in the Factor Fields family, which factorizes a signal into a localized coefficient field and a global basis field with periodic transformations. We extensively evaluate our DiF model on three signal reconstruction tasks including 2D image regression, 3D SDF reconstruction, and radiance field reconstruction. We demonstrate that our DiF model leads to state-of-the-art reconstruction quality, better or on par with previous methods on all three tasks, while achieving faster reconstruction and more compact model sizes than most methods. Our DiF model is able to generalize across scenes by learning shared basis field factors from multiple signals, allowing us to reconstruct new signals from sparse observations. We show that, using such pre-trained basis factors, our method enables high-quality few-shot radiance field reconstruction from only 3 or 5 views, outperforming previous methods like PixelNeRF and MVSNeRF in the sparse view / wide baseline setting. In general, our framework takes a step towards a generic neural representation with high accuracy and efficiency. We believe that the flexibility of our framework will help to inspire future research on efficient signal representations, exploring the potential of multi-factor representations or novel coordinate transformations.

References

[1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *CVPR Workshops*, 2017. 11

[2] Kara-Ali Aliev, Dmitry Ulyanov, and Victor S. Lempitsky. Neural point-based graphics. *arXiv.org*, 1906.08240, 2019. 5

[3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, 2022. 5

[4] Sai Bi, Zexiang Xu, Pratul Srinivasan, Ben Mildenhall, Kalyan Sunkavalli, Milo Haan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Neural reflectance fields for appearance acquisition, 2020. 5

[5] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, Miloš Hašan, Yannick Hold-Geoffroy, David Kriegman, and Ravi Ramamoorthi. Deep reflectance volumes: Relightable reconstructions from multi-view photometric images. 2020. 5

[6] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerd: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12684–12694, 2021. 5

[7] Mark Boss, Varun Jampani, Raphael Braun, Ce Liu, Jonathan Barron, and Hendrik Lensch. Neural-pil: Neural pre-integrated lighting for reflectance decomposition. *Advances in Neural Information Processing Systems*, 2021. 5

[8] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J. Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3d generative adversarial networks. In *CVPR*, 2022. 5, 6, 12

[9] Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. Pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *CVPR*, 2021. 5

[10] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV*, 2022. 1, 2, 4, 5, 6, 8, 10, 12

[11] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *ICCV*, 2021. 11

[12] Xu Chen, Yufeng Zheng, Michael Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 10

[13] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *CVPR*, 2021. 5

[14] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 5

[15] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. *arXiv preprint arXiv:2204.11918*, 2022. 11

[16] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022. 2, 6, 8, 10

[17] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. In *Advances In Neural Information Processing Systems*, 2022. 5

[18] Felix Heide, Wolfgang Heidrich, and Gordon Wetzstein. Fast and flexible convolutional sparse coding. In *CVPR*, 2015. 5

[19] Binbin Huang, Xinhao Yan, Anpei Chen, Shenghua Gao, and Jingyi Yu. PREF: phasorial embedding fields for compact neural representations. *arXiv.org*, 2022. 6

[20] Hankyu Jang and Daeyoung Kim. D-tensorf: Tensorial radiance fields for dynamic scenes. *arXiv.org*, 2022. 6

Name	N	$\mathbf{f}_i(\mathbf{x})$	$\boldsymbol{\gamma}_i(\mathbf{x})$	Time	Size	2D Images	Radiance Field	Few-shot RF
(1) iNGP*	1	Vectors	Hashing(x)	00:45/**12:02**/ -	0.92/6.80/ -	34.73/0.906 ●	32.56/0.958 ●	- / -
(2) DiF-Hash-B	2	Vectors; 3D grids	Hashing(x)	00:55/13:10/**04:45**	1.09/**4.37/3.28**	37.53/0.949 ○	32.80/0.960 ○	26.62/0.919 ○
(3) EG3D*	1	2D Maps	Orthog$_{2D}$(x)	- /14:17/ -	- /4.54/ -	- / -	30.01/0.935	- / -
(4) TensoRF-VM*	2	2D Maps; Vectors	Orthog$_{1,2D}$(x)	- /16:20/13:06	- /4.55/**4.93**	- / -	30.47/0.940	26.79/0.908 ●
(5) DiF-no-C	1	3D Grids	Sawtooth(x)	**00:41**/12:55/ -	**0.82**/4.55/ -	22.28/0.479	31.10/0.947	- / -
(6) DiF-Grid	2	3D Grids; 3D Grids	Sawtooth(x)	01:13/12:10/11:35	0.87/5.10/7.32	39.51/0.963 ●	33.14/0.961 ●	25.41/0.915 ●

(a) **Design Study on Number of Factors N.**

Name	N	$\mathbf{f}_i(\mathbf{x})$	$\boldsymbol{\gamma}_i(\mathbf{x})$	Time	Size	2D Images	Radiance Field	Few-shot RF
(1) OccNet*	1	x	x	02:17/100:23/ -	**0.38/0.43**/ -	13.90/0.437	20.60/0.849	- / -
(2) NeRF*	1	x	Sinusoidal(x)	02:18/65:50/ -	0.39/0.44/ -	28.99/0.816	27.81/0.919	- / -
(3) DiF-Hash-B	2	Vectors; 3D grids	Hashing(x)	00:55/13:10/4:45	1.09/4.37/**3.28**	37.53/0.949	32.80/0.960 ●	26.62/0.919 ●
(4) DiF-Hash-B-SL	2	Vectors; 3D grids	Hashing(x)	**00:41**/13:21/**03:03**	0.89/4.54/4.59	30.97/0.891	31.11/0.941 ●	24.13/0.881 ●
(5) DiF-Grid	2	3D Grids; 3D Grids	Sawtooth(x)	01:13/**12:10**/11:35	0.99/5.10/7.32	39.51/0.963 ●	33.14/0.961 ○	25.41/0.915 ○
(6) DiF-Grid-SL	2	3D Grids; 3D Grids	Sawtooth(x)	00:49/22:35/13:12	0.98/5.11/7.25	38.73/0.973 ○	31.08/0.942	23.88/0.882

(b) **Design Study on Pyramid Levels L.**

Name	N	$\mathbf{f}_i(\mathbf{x})$	$\boldsymbol{\gamma}_i(\mathbf{x})$	Time	Size	2D Images	Radiance Field	Few-shot RF
(1) DiF-Hash-B	2	Vectors; 3D Grids	Hashing(x)	**00:55**/13:10/4:45	1.09/**4.37/3.28**	37.53/0.949	32.80/0.960	26.62/0.919 ●
(2) DiF-Sin-B	2	3D Grids; 3D Grids	Sinusoidal(x)	**00:55**/12:19/**08:28**	0.99/5.10/7.32	38.21/0.953 ●	32.85/0.961 ●	25.43/0.908 ○
(3) DiF-Tri-B	2	3D Grids; 3D Grids	Triangular(x)	**00:55**/12:49/08:44	0.99/5.10/7.32	39.38/0.962 ○	32.95/0.960 ○	24.78/0.904
(4) DiF-Grid	2	3D Grids; 3D Grids	Sawtooth(x)	01:13/**12:10**/11:35	0.99/5.10/7.32	39.51/0.963 ○	33.14/0.961 ○	25.41/0.915 ●

(c) **Design Study on Coordinate Transformations $\boldsymbol{\gamma}_i$.**

Name	N	$\mathbf{f}_i(\mathbf{x})$	$\boldsymbol{\gamma}_i(\mathbf{x})$	Time	Size	2D Images	Radiance Field	Few-shot RF
(1) TensoRF-VM*	2	2D Maps; Vectors	Orthog$_{1,2D}$(x)	- /16:20/13:06	- /4.55/4.93	- / -	30.47/0.940	26.79/0.908 ○
(2) DiF-Grid	2	3D Grids; 3D Grids	Sawtooth(x)	01:13/**12:10**/11:35	0.99/5.10/7.32	39.51/0.963 ○	33.14/0.961 ○	25.41/0.915
(3) DiF-DCT	2	3D Grids; 3D Grids	Sawtooth(x)	00:53/ - / -	**0.18**/ - / -	23.16/0.606	- / -	- / -
(4) DiF-Hash-B	2	Vectors; 3D grids	Hashing(x)	00:55/13:10/10.15	1.09/4.37/3.28	37.53/0.949	32.80/0.960	26.53/0.924 ●
(5) DiF-MLP-B	2	MLP; 3D grids	Sawtooth(x)	01:24/18:18/18:23	**0.18**/0.62/2.53	28.76/0.819	29.62/0.932	26.91/0.927 ○
(6) DiF-MLP-C	2	3D Grid; MLP	Sawtooth(x)	01:13/13:38/**08:23**	0.87/4.54/4.86	34.72/0.910 ●	32.57/0.956 ●	23.54/0.875
(7) TensoRF-CP*	3	Vectors×3	Orthog$_{1D}$(x)	**00:43**/28:05/12:42	0.39/**0.29/0.29**	33.79/0.899	31.14/0.944	22.62/0.867

(d) **Design Study on Field Representations $\mathbf{f}_i$.**

Name	N	$\mathbf{f}_i(\mathbf{x})$	$\boldsymbol{\gamma}_i(\mathbf{x})$	Time	Size	2D Images	Radiance Field	Few-shot RF
(1) TensoRF-CP*	3	Vectors×3	Orthog$_{1D}$(x)	**00:43/28:05**/10:11	0.39/0.29/0.29	33.79/0.899 ●	31.14/0.944 ○	23.19/0.879 ○
(2) TensoRF-CP*-Cat	3	vectors×3	Orthog$_{1D}$(x)	00:47/39:05/**09:47**	0.40/0.31/0.31	25.67/0.683 ○	26.75/0.905 ○	21.43/0.856 ○
(3) TensoRF-VM*	2	2D Maps; vectors	Orthog$_{1,2D}$(x)	- /**16:20**/12:52	- /4.55/4.93	- / -	30.47/0.940 ○	26.99/0.911 ●
(4) TensoRF-VM*-Cat	2	2D Maps; vectors	Orthog$_{1,2D}$(x)	- /18:35/**07:02**	- /4.56/4.94	- / -	29.86/0.939 ○	24.67/0.885 ○
(5) DiF-Grid	2	3D Grids; 3D Grids	Sawtooth(x)	01:13/12:10/11:35	0.99/5.10/7.32	39.51/0.963 ●	33.14/0.961 ●	25.41/0.915 ●
(6) DiF-Grid-Cat	2	3D Grids; 3D Grids	Sawtooth(x)	**00:51/11:35/06:47**	1.00/5.10/7.32	37.76/0.946 ○	32.95/0.960 ○	24.71/0.894 ○

(e) **Performance comparison on element-wise product ∘ vs. concatenation.**

Table 3: **Influence of Design Choices in the Factor Fields Framework.** The comparison is done using the same code base and hyperparameter configuration including number of levels, frequency of each level, etc. Prior methods represented in our framework are labeled with * due to minor differences with respect to the original publications. Runtime and model size are reported separately for 2D Images / Radiance Fields / Few-shot Radiance Fields.

[21] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *arXiv.org*, 2018. 10

[22] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray

entropy minimization for few-shot neural volume rendering. In *CVPR*, 2022. 10

[23] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 7

[24] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Trans. on Graphics*, 36(4), 2017. 8, 10

[25] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. In *NIPS*, 2022. 5

[26] Jonas Kulhanek, Erik Derner, Torsten Sattler, and Robert Babuska. Viewformer: Nerf-free neural rendering from few images using transformers. In Shai Avidan, Gabriel J. Brostow, Moustapha Cisse, Giovanni Maria Farinella, and Tal Hassner, editors, *ECCV*, 2022. 10

[27] Alexandr Kuznetsov, Krishna Mullia, Zexiang Xu, Miloš Hašan, and Ravi Ramamoorthi. Neumip: multi-resolution neural materials. *ACM Trans. on Graphics*, 2021. 5

[28] Jae Yong Lee, Yuqun Wu, Chuhang Zou, Shenlong Wang, and Derek Hoiem. QFF: quantized fourier features for neural field representations. *arXiv.org*, 2022. 6

[29] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, and Zhaoyang Lv. Neural 3d video synthesis. *arXiv.org*, 2103.02597, 2021. 5

[30] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *arXiv.org*, 2011.13084, 2020. 5

[31] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. In *NeurIPS*, 2020. 5

[32] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. In *ACM Trans. on Graphics*, 2019. 5

[33] Julien Mairal, Francis R. Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In Andrea Pohoreckyj Danyluk, Léon Bottou, and Michael L. Littman, editors, *ICML*, 2009. 5

[34] Matthew, Aditya Kusupati, Alex Fang, Vivek Ramanujan, Aniruddha Kembhavi, Roozbeh Mottaghi, and Ali Farhadi. Neural radiance field codebooks. *arXiv.org*, 2023. 5

[35] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 3, 5

[36] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 3, 4, 5, 6, 8, 10, 12

[37] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. on Graphics*, 2022. 1, 2, 4, 6, 7, 8, 9, 10, 12

[38] Michael Niemeyer, Jonathan Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022. 10

[39] Michael Niemeyer, Lars M. Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 5

[40] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 3, 5

[41] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021. 5

[42] Songyou Peng, Kyle Genova, Chiyu "Max" Jiang, Andrea Tagliasacchi, Marc Pollefeys, and Thomas A. Funkhouser. Openscene: 3d scene understanding with open vocabularies. *arXiv.org*, 2022. 5

[43] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, 2020. 6

[44] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 5

[45] Gilles Rainer, Wenzel Jakob, Abhijeet Ghosh, and Tim Weyrich. Neural btf compression and interpolation. In *Computer Graphics Forum*, 2019. 5

[46] Ron Rubinstein, Michael Zibulevsky, and Michael Elad. Efficient implementation of the k-svd algorithm using batch orthogonal matching pursuit. Technical report, Computer Science Department, Technion, 2008. 5

[47] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *NeurIPS*, 2020. 5

[48] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d : Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. *arXiv.org*, 2022. 6

[49] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NIPS*, 2020. 1, 7, 8, 9

[50] Liangchen Song, Anpei Chen, Zhong Li, Zhang Chen, Lele Chen, Junsong Yuan, Yi Xu, and Andreas Geiger. Nerfplayer: A streamable dynamic scene representation with decomposed neural radiance fields. *arXiv.org*, 2022. 6

[51] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 1, 2, 4, 6, 8, 10, 12

[52] Towaki Takikawa, Alex Evans, Jonathan Tremblay, Thomas Müller, Morgan McGuire, Alec Jacobson, and Sanja Fidler. Variable bitrate neural fields. In *ACM Trans. on Graphics*, 2022. 6

[53] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *NeurIPS*, 2020. 7, 8, 9

[54] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: image synthesis using neural textures. *ACM Trans. on Graphics*, 2019. 5

[55] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. *arXiv.org*, 2021. 5

[56] 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. In *NeurIPS*, 2021. 5

[57] John Wright, Allen Y. Yang, Arvind Ganesh, S. Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. *PAMI*, 2009. 5

[58] Qiangeng Xu, Zexiang Xu, Julien Philip, Sai Bi, Zhixin Shu, Kalyan Sunkavalli, and Ulrich Neumann. Point-nerf: Point-based neural radiance fields. In *CVPR*, 2022. 5

[59] Jianchao Yang, John Wright, Thomas S. Huang, and Yi Ma. Image super-resolution via sparse representation. *TIP*, 2010. 5

[60] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. 5

[61] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, 2021. 10, 11

[62] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. In *NeurIPS*, 2022. 5

[63] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *CVPR*, 2021. 5

[64] Yuanqing Zhang, Jiaming Sun, Xingyi He, Huan Fu, Rongfei Jia, and Xiaowei Zhou. Modeling indirect illumination for inverse rendering. In *CVPR*, 2022. 5

[65] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: learning view synthesis using multiplane images. *ACM Trans. on Graphics*, 2018. 5

[66] Junqiu Zhu, Yaoyi Bai, Zilin Xu, Steve Bako, Edgar Velázquez-Armendáriz, Lu Wang, Pradeep Sen, Miloš Hašan, and Ling-Qi Yan. Neural complex luminaires: representation and rendering. *ACM Trans. on Graphics*, 2021. 5