

HyperDiffusion: Generating Implicit Neural Fields with Weight-Space Diffusion

Ziya Erkoç¹ Fangchang Ma² Qi Shan² Matthias Nießner¹ Angela Dai¹

¹Technical University of Munich ²Apple

<https://ziyaerkoc.com/hyperdiffusion>

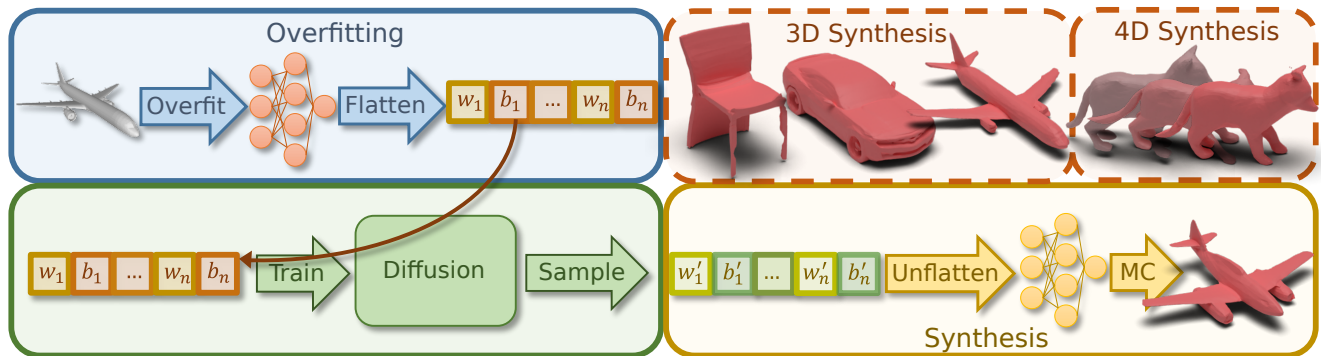


Figure 1: HyperDiffusion enables a new paradigm in directly generating neural implicit fields by predicting their weight parameters. We leverage implicit neural fields to optimize a set of MLPs that faithfully represent individual dataset instances (“Overfitting,” top-left). Our network, based on a transformer architecture, then models a diffusion process directly on the optimized MLP weights (“Diffusion,” bottom-left). This enables synthesis of new implicit fields (“Synthesis,” bottom-right).

Abstract

Implicit neural fields, typically encoded by a multilayer perceptron (MLP) that maps from coordinates (e.g., xyz) to signals (e.g., signed distances), have shown remarkable promise as a high-fidelity and compact representation. However, the lack of a regular and explicit grid structure also makes it challenging to apply generative modeling directly on implicit neural fields in order to synthesize new data. To this end, we propose HyperDiffusion, a novel approach for unconditional generative modeling of implicit neural fields. HyperDiffusion operates directly on MLP weights and generates new neural implicit fields encoded by synthesized MLP parameters. Specifically, a collection of MLPs is first optimized to faithfully represent individual data samples. Subsequently, a diffusion process is trained in this MLP weight space to model the underlying distribution of neural implicit fields. HyperDiffusion enables diffusion modeling over an implicit, compact, and yet high-fidelity representation of complex signals across 3D shapes and 4D mesh animations within one single unified framework.

1. Introduction

Recent years have seen profound development in implicit neural field models, demonstrating powerful representations, particularly 3D shape geometry [34, 6], neural radiance fields (NeRF) [31], and complex signals with higher-order derivative constraints [46]. Typically, an implicit neural field¹ maps an input coordinate location in n -dimensional space to the target signal domain. For example, an implicit surface representation

$$\{\mathbf{x} \in \mathbb{R}^n | f(\mathbf{x}, \theta) = 0\},$$

where $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is typically characterized by a multilayer perceptron (MLP). Notably, such neural fields efficiently represent sparse high-dimensional data in a relatively low-dimensional MLP weight space. This continuous mapping enables sampling at arbitrary-resolution for surface representations, eliminating explicit resolution constraints inherent to classical point, mesh, or voxel representations. One can then easily reconstruct the mesh underlying

¹Implicit neural fields are also referred to as coordinate fields or coordinate-based networks. We use these terms interchangeably.

ing this compact representation through methods such as Marching Cubes [28]. It has enabled faithful 3D reconstruction where an MLP is optimized to fit a set of point observations [34], as well as preliminary results in higher-dimensional 4D reconstruction and tracking [27].



(a) Generated 3D shapes.



(b) Generated 4D animation sequence.

Figure 2: HyperDiffusion is a dimension-agnostic generative model. The same approach can be trained on data of various dimensionalities to synthesize high-fidelity examples. For instance, (a) 3D shapes, and (b) 4D animation sequences (3D and time).

The aforementioned implicit neural field representation also poses a new opportunity for generative modeling - that is, rather than model the raw 3D or 4D surface information, we aim to directly model the space of optimized neural field MLPs without relying on any ties to alternative explicit representations (e.g., points, meshes, and voxels) or approximation to latent manifolds (which typically require a pair of encoder and decoder networks pretrained on large-scale datasets). In particular, we observe that optimized neural field MLPs typically maintain an extremely compact representation of a high-dimensional 3D or 4D surface.

We thus propose HyperDiffusion, a new paradigm for generative modeling of neural fields. HyperDiffusion leverages diffusion modeling directly on the weight space of optimized neural fields, enabling generative modeling of high-dimensional data characterized by neural fields. Our HyperDiffusion approach is dimension-agnostic, and we apply the same method towards both unconditional 3D and 4D (3D and time) surface generation to demonstrate the power of directly generating neural fields as MLP weights.

Specifically, to model neural fields characterizing n -dimensional surface data, we consider a set of MLPs that

have been optimized to represent individual instances from a dataset. We then employ a transformer-based network to model the diffusion process directly on the optimized MLP weights. This enables generative modeling on a low-dimensional space, and we demonstrate its high-fidelity and diverse generative modeling capabilities, achieving state-of-the-art 3D and 4D surface synthesis.

Our contributions can be summarized as follows:

- We present the first approach to model the space of neural field MLP weights by diffusion modeling, enabling a new paradigm for high-dimensional generative modeling.
- Our MLP optimization of surface occupancy provides a low-dimensional weight space for effective unconditional diffusion modeling with a transformer-based architecture for both 3D and 4D surfaces at high fidelity.

2. Related Work

Neural Implicit Fields Neural implicit fields have shown promising results in representing high-fidelity geometry and appearance in 3D. One seminal early work is DeepSDF [34], which encodes the shapes of a class of objects as signed distance functions using a multi-layer fully-connected neural network. Neural Radiance Fields (NeRF) [31] encode a coordinate-based radiance field with a MLP and is capable of producing photo-realistic 2D renderings at novel views through volumetric rendering. To address the bias towards learning low-frequency details in a standard MLP, Fourier features [48] and periodic activation functions [46] have been proposed to improve representation of complex signals. In this work, we train on the weights of these multi-layer fully-connected neural networks and generate new weights that represent valid signals.

Generative Adversarial Networks (GANs) Generative Adversarial Networks (GANs) [18] has been shown to be capable of generating high-resolution 2D images, most notably StyleGAN [24, 25, 23, 42]. More recently, GANs have been adopted for 3D generation with various underlying representations. For instance, pi-GAN [9] modulates the MLP network which encodes a radiance field, EG3D [7] proposes an efficient tri-plane representation, and GMPI [53] directly generates multiplane images for efficient training and inference. A similar line of work focuses on generating textures for meshes [44, 17]. A general limitation with GANs is volatile training stability due to the game theoretical nature of generator and discriminator networks competing against each other. This motivates researcher to look into alternative generative models, including Diffusion Process. In this work, we adopt a diffusion model for the MLP weight generation.

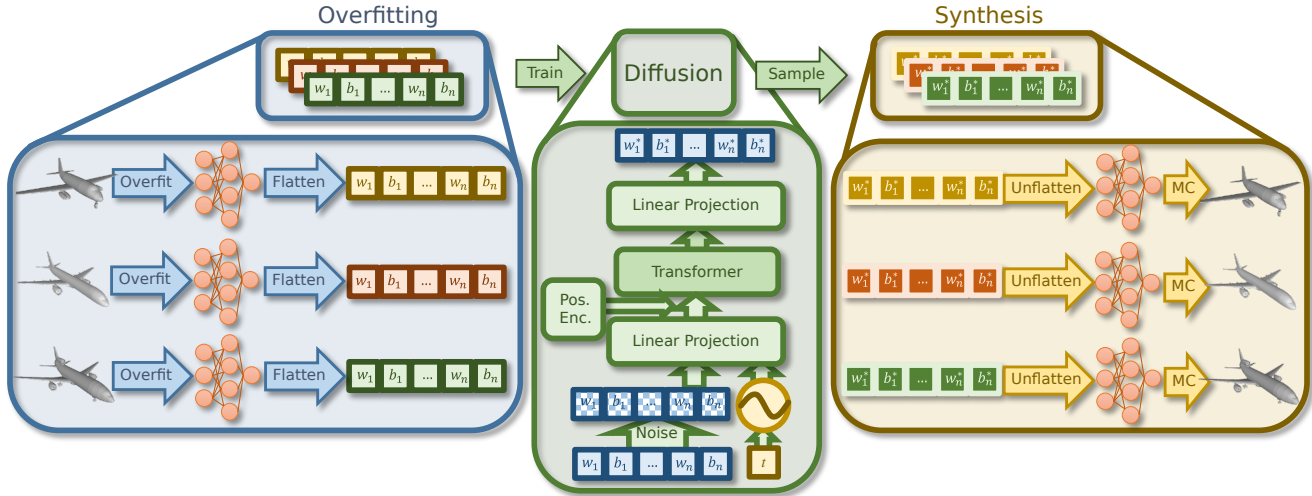


Figure 3: Overview of HyperDiffusion. We first fit a set of neural field MLPs to a dataset of 3D or 4D shapes in an overfitting process, producing high-fidelity shape representations that support arbitrary resolutions (left). To support the following diffusion process on the MLP weights, we use an arbitrary optimized MLP to initialize the rest of the MLP optimization. We then use a transformer-based architecture to model a diffusion process directly on these optimized MLP weights, predicting the denoised MLP weights and biases as flattened vectors (middle). This enables synthesis of new neural field representations as their MLP weights, from which meshes can be extracted through Marching Cubes. Our approach is agnostic to the resolution of the dataset instances, and we demonstrate effective 3D and 4D shape modeling.

2D Diffusion Models Diffusion probabilistic models [20, 47] have emerged as a powerful alternative to its precursors, such as GANs [18] and energy-based models (EBMs) [14], for generative modeling. They have been shown to outperform GANs on image synthesis tasks in terms of not only superior quality and fidelity [12, 40], but also the capability to enable text-conditioning [33, 41]. Specifically, Latent Diffusion [39, 4] enables high-resolution image synthesis by applying diffusion in the latent space of pretrained autoencoders. In contrast to Latent Diffusion, our method operates on the MLP weight space and generates new neural implicit fields.

3D/4D Diffusion Models Diffusion models have also been deployed in more challenging tasks beyond image synthesis, including 2D video generation [21], 3D shape and scene generation [8, 9, 11, 3, 37, 32], and 4D generation with animation of 3D shapes [45, 49]. In particular, MDM is a human motion diffusion model where the input is a sequence of human poses [49]. They use a transformer architecture as the denoising network. It takes in a noisy human pose sequence input and denoises it to generate a new motion animation. Unlike their approach, we can directly generate animated meshes instead of just human poses.

Diffusion modeling in a latent manifold for neural fields has also been concurrently proposed, including DiffusionSDF [11], 3DShape2VecSet [52], and LION [51]. Such

latent diffusion models require a high-quality latent manifold to be learned, which can be challenging with limited quantities of 3D and 4D data. In contrast, our approach to model the weights of optimized neural fields operates on inherently high quality shape representations that can be extremely well-fit per instance. In addition, Shue et al. [43] proposed a diffusion on triplane features. Compared to triplane diffusion, we observe that our approach on the MLP weight space enables leveraging the same diffusion framework for various input dimensionalities, enabling effective 4D generative modeling as well.

Cross-modality Diffusion Models A number of recent publications have been investigating unification of diffusion models across different data dimensions and modalities, including Functa [15], GEM [13], and GASP [16]. Most recently, Diffusion Probabilistic Fields [55] proposes an explicit field representation without a latent field parametrization, and formulates the generative model in a single-stage end-to-end training. Our method can be seen as an alternative solution to uniting methods designed across different data dimensions, since the implicit neural fields are essentially dimension-agnostic.

3. Method Overview

HyperDiffusion is an unconditional generative model for implicit neural fields encoded by MLPs. We operate di-

rectly on MLP weights, enabling generation of new neural implicit fields characterized by synthesized MLP parameters. Our training paradigm encompasses a two-phase approach, as shown in Figure 3.

In the first MLP overfitting step, detailed in Section 4, we optimize a collection of MLPs such that each MLP represents a faithful neural occupancy field of a data sample (e.g., a 3D shape) from the training set. This enables highly-accurate shape fitting due to the representation power of the neural fields. The optimized MLP weights are flattened into 1D vectors and passed to a downstream diffusion process as ground truth signals.

In the second step, detailed in Section 5, the aforementioned optimized MLP weights are passed into a diffusion network for training. This diffusion network is domain-agnostic without any assumptions or prior knowledge on the dimensionality of the underlying signal, since its input is a set of flattened MLP weight vectors. After training is completed, new MLP weights, which correspond to a valid neural implicit field, can be synthesized through the reverse diffusion process on a randomly sampled noise signal.

For 3D and 4D generation, the underlying meshes can be further extracted and visualized with Marching cubes [28].

4. Per-Sample MLP Overfitting

In this first step, we optimize MLPs for each data sample from the training set $\{S_i, i = 1, \dots, N\}$ and save their optimized network weights. Specifically, for a train sample S_i , the surface at iso-level τ is represented as

$$\{\mathbf{x} \in \mathbb{R}^n, f(\mathbf{x}, \theta_i) = \tau\}, \quad (1)$$

where f is an MLP parametrized by $\theta_i \in \mathbb{R}^h$ (i.e., weights and biases in Fig. 3) and \mathbf{x} represents a spatial location. The goal is to minimize the binary cross entropy loss function

$$L = \text{BCE}(f(\mathbf{x}, \theta_i), o_i^{\text{gt}}(\mathbf{x})), \quad (2)$$

where $o_i^{\text{gt}}(\mathbf{x})$ is the ground truth occupancy of \mathbf{x} with respect to S_i . Here, we exploit the representation power of neural fields, which can model high-dimensional surfaces with high accuracy. As these optimized MLPs serve as ground truth for the following diffusion processes, their ability to model high-fidelity shapes is crucial.

Unlike prior methods, we do not require any auto-encoding networks [1, 5] nor auto-decoding networks (e.g., DeepSDF [11]), which typically share the same network parameters across an entire dataset. In contrast, although we use the same MLP architecture for different samples in the training dataset, one set of MLP weights is optimized specifically for each data sample. In other words, there is no parameter sharing, and this per-sample optimization attains the best fidelity possible for the implicit neural field representation.

MLP Architecture and Training Our MLP architecture is a standard multi-layer fully-connected network with ReLU activation functions and an input positional encoding [31]. We use 3 hidden layers with 128 neurons each, finally outputting a scalar occupancy value. The same MLP architecture is shared across both 3D shape and 4D animation experiments, where each MLP encodes an occupancy field per-train sample. This enables dimension-agnostic paradigm for encoding of various data signals, in our case 3D and 4D shapes, where only the positional encoding is adapted for various-dimensional inputs.

To optimize a set of MLP weights and biases for an input 3D shape, we sample points randomly both inside and outside of the 3D surface. We normalize all train instances to $[-0.5, 0.5]^3$, and randomly sample $100k$ points within the space. To effectively characterize fine-scale surface detail, we further sample $100k$ points near the surface of the mesh. Both sets of points are combined and tested for inside/outside occupancy using generalized winding numbers [2]; these occupancies are used to supervise the overfitting process. We optimize each MLP with a mini-batch size of 2048 points, trained with a BCE loss for 800 epochs until convergence, which takes ≈ 6 minutes per shape.

MLP overfitting to 4D shapes is performed analogously to 3D. For each temporal frame, we sample $200k$ points and their occupancies, following the 3D shape sampling. The sampling process is repeated for each frame of an animation sequence. We optimize one set of MLP weights and biases for each animation sequence to represent each 4D shape.

Weight Initialization In order to encourage a smooth diffusion process over the set of optimized MLP weights, we guide the MLP optimization process with consistent weight initialization. That is, we initially optimize one set of MLP weights and biases θ_1 to represent the first train sample S_1 , and use the optimized weights of θ_1 to initialize optimization of the rest of the MLPs $\{\theta_2, \dots, \theta_N\}$.

5. MLP Weight-Space Diffusion

We then model the weight space of our optimized MLPs through a diffusion process. We consider each set of optimized MLP weights and biases $\{\theta_i\}$ as a flattened 1D vector. We use a transformer architecture \mathcal{T} , following [35], for our denoising network. As transformers have been shown to elegantly handle long vectors in the language domain, we find it to be a suitable choice for modeling the MLP weight space. \mathcal{T} predicts the denoised MLP weights directly, rather than the noise. Our h -dimensional vectors $\{\theta_i\}$ of MLP weights and biases are input to \mathcal{T} . Each θ_i is divided into 8 tokens by MLP layers, to be encoded by \mathcal{T} .

Modeling the MLP weights as a 1D vector for diffusion enables a general formulation for modeling neural fields, as the MLP weights are agnostic to varying-dimensional data.




Figure 4: Denoising MLP parameters at various time steps, visualized with their corresponding shapes, from random noise (left) to fully denoised (right). The image shows the gradual change from 0 denoising steps, which are generated from random MLP weights, up to 500 steps corresponding to fully-denoised shape. More specifically, we leverage the DDIM [47] sampling strategy. Interestingly, noisy MLP weights do not necessarily correspond to a valid 3D shape; however, the iterative application of the denoising operator eventually converges to a quality output shape.

This makes HyperDiffusion flexible to a variety of neural field representations; in particular, we observe the neural field ability to compactly represent high-dimensional shape data and demonstrate generative modeling of MLPs representing 3D and 4D shapes, respectively.

During diffusion modeling, we apply standard gaussian noise t times to each vector θ . The noisy vector, along with the sinusoidal embedding of t , are then input to a linear projection. The projections are then summed up with a learnable positional encoding vector. Then, the transformer outputs denoised tokens that we pass through a final output projection to produce the predicted denoised MLP weights w^* . We train with a Mean Squared Error (MSE) loss between the denoised weights θ^* and the input weights θ .

We illustrate the denoising process for MLPs representing 3D shapes in Figure 4. Noisy MLPs correspond to invalid shapes, which are denoised to MLPs that represent valid 3D surfaces. We employ Denoising Diffusion Implicit Models (DDIM) [47] to sample new MLPs from the diffusion process.

Implementation Details Our 3-layer 128-dim MLPs contain $\approx 36k$ parameters, which are flattened and tokenized for diffusion. We use an AdamW [29] optimizer with batch size 32 and initial learning rate of $2e^{-4}$, which is reduced by 20% every 200 epochs. We train for ≈ 4000 epochs until convergence, which takes ≈ 4 days on a single A6000 GPU.

6. Results

6.1. Datasets

For 3D shape generation, we use the car, chair, and airplane categories of the ShapeNet [10] dataset. The car, chair and airplane categories have 3533, 6778, and 4045 shapes respectively. For the 4D shape generation task, we use 16-

frame animal animation sequences from the DeformingThings4D [27] dataset, comprising 1772 sequences. For both datasets, we split the data into non-overlapping partitions, including training (80%), validation (5%) and testing (15%) subsets.

6.2. Voxel-based Diffusion Baseline

While several existing methods tackle 3D shape generation, unconditional 4D shape generation remains underexplored. Thus, in addition to existing 3D baselines, we introduce a voxel-based diffusion model as a baseline for both 3D and 4D shape generation.

3D shapes are represented as dense occupancy grids, and a 3D UNet denoising network is applied on the 3D voxel grids. 4D shapes are represented similarly, with each frame of an animation sequence voxelized to a 3D occupancy grid, producing a 4D occupancy grid representing the full sequence. We use the same 3D UNet as a denoising network to synthesize 4D animation, as a 4D UNet became computationally intractable.

For our experiments, we use a voxel resolution of 24^3 for 3D shapes and 16×24^3 for 4D shapes (the maximum spatial resolution such that 4D grids could be tractably trained).

6.3. Evaluation Metrics

Evaluation of unconditional synthesis of 3D and 4D shapes can be challenging due to lack of direct correspondence to ground truth data. We thus follow prior works [51, 30, 54] in evaluating Minimum Matching Distance (MMD), Coverage (COV), and 1-Nearest-Neighbor Accuracy (1-NNA). For MMD, lower is better; for COV, higher

is better; for 1-NNA, 50% is the optimal.

$$\text{MMD}(S_g, S_r) = \frac{1}{|S_r|} \sum_{Y \in S_r} \min_{X \in S_g} D(X, Y),$$

$$\text{COV}(S_g, S_r) = \frac{|\{\arg \min_{Y \in S_r} D(X, Y) | X \in S_g\}|}{|S_r|},$$

$$1\text{-NNA}(S_g, S_r) = \frac{\sum_{X \in S_g} \mathbb{1}[N_X \in S_g] + \sum_{Y \in S_r} \mathbb{1}[N_Y \in S_r]}{|S_g| + |S_r|},$$

where in the 1-NNA metric N_X is a point cloud that is closest to X in both generated and reference dataset, i.e.,

$$N_X = \arg \min_{K \in S_r \cup S_g} D(X, K)$$

We use a Chamfer Distance (CD) distance measure $D(X, Y)$ for computing these metrics in 3D, and report CD values multiplied by a constant 10^2 . To evaluate 4D shapes, we extend to the temporal dimension for T frames:

$$D(X, Y) = \frac{1}{T} \sum_{t=0}^{T-1} CD(X[t], Y[t]).$$

We note that the MMD metric has been discussed to be unreliable as a measure of generation quality [51]. It suffers from not capturing diversity and coverage aspects of the generated set, and thus we also consider perceptual metric for 3D shape generation. In particular, we follow [52] and compute a Frchet Pointnet++ Distance [38] (FPD), analogous to the commonly used Frchet Inception Distance (FID score) [19] in the image domain. FPD instead uses a 3D Pointnet++ network (trained on the ModelNet [50]) for feature extraction from generated shapes. For FPD, lower scores are better.

For evaluation, each synthesized and ground truth shape is normalized by its mean and standard deviation on a per-shape basis. To evaluate point-based measures, we sample 2048 points randomly from all baseline outputs; for our approach and the voxel baseline, points are sampled from the extracted mesh surface, and for point cloud baselines, points are sampled directly from the synthesized outputs.

6.4. Unconditional 3D Generation

We demonstrate HyperDiffusion on unconditional neural field generation of 3D shapes. For our method, we show output meshes extracted with Marching Cubes [28].

Comparison to state of the art We compare with state-of-the-art 3D shape generation methods Point-Voxel Diffusion [54] (PVD), Diffusion Point Cloud [30] (DPC), ShapeGAN [26], as well as our voxel baseline in Table 1. HyperDiffusion achieves improved performance over these baselines in all shape categories and all evaluation metrics except MMD, which has been noted to be less reliable due




Figure 5: Synthesized 3D shapes with HyperDiffusion trained on ShapeNet [10]. Each shape is represented with a 3-layer 128-dim MLP. We extract the underlying isosurface from the MLPs with Marching Cubes [28].

to lack of sensitivity to low-quality results [51]. Notably, we achieve significantly improved FPD, which captures the perceptual quality of our synthesized shapes. Figure 6 further shows our capability to represent high-resolution detail as captured in our synthesized neural fields, in comparison with baselines.

What is the effect of MLP overfitting weight initialization? We analyze the effect of our weight initialization strategy for MLP optimization in Table 3, in comparison with random initialization per-MLP. Instead of initializing each MLP completely independently, our consistent weight initialization from a single optimized MLP produces improved results, particularly for the perceptual FPD metric. This could be attributed to the fact that with a consistent initialization the MLP weights stay relatively close to each other.

What is the impact of MLP positional encoding? We ablate the effect of positional encoding in our per-instance MLP representations in Table 4, where No PE denotes the same MLP architecture without any positional encoding. We see that applying positional encoding to the input coordinates significantly improves generation quality.

Novel shape synthesis. We explore the degree of novelty in our generated neural fields in 3D shape generation. Figure 8 our synthesized shapes (red) in comparison to the top-3 nearest neighbours by Chamfer distance from the training set (green). We show that our synthesized neural fields represent shapes not already present in the training set, but rather novel shape data. This is critical to the core idea of our approach since it supports the claim that running denoising diffusion over neural network parameter weights facilitates generalizability within the underlying data distribution of 3D meshes.

Category	Method	MMD ↓	COV (%) ↑	1-NNA (%) ↓	FPD ↓
Airplane	Voxel Baseline	6.0	28	94.1	38.9
	PVD [54]	3.4	39	76.3	5.8
	DPC [30]	3.1	46	74.7	18.7
	ShapeGAN [26]	6.6	18	93.8	70.4
	Ours	3.4	49	69.3	3.5
Car	Voxel Baseline	4.9	13	98.7	10.4
	PVD [54]	3.5	27	76.8	4.5
	DPC [30]	3.3	33	82.4	7.6
	ShapeGAN [26]	4.3	17	97.6	10.6
	Ours	3.4	36	73.1	2.6
Chair	Voxel Baseline	11.8	28	80.6	30.6
	PVD [54]	6.8	42	58.3	3.5
	DPC [30]	6.3	44	61.4	26.0
	ShapeGAN [26]	8.4	42	70.1	5.8
	Ours	7.1	53	54.1	1.7

Table 1: Quantitative comparison on unconditional 3D shape generation for the airplane, car and chair categories from ShapeNet [10]. Our synthesized neural fields outperform the baseline volumetric method and prior state of the art [54, 30], particularly on the perceptual FPD metric most representative of visual quality.

Methods	MMD ↓	COV (%) ↑	1-NNA (%) ↓
Voxel Baseline	21.9	35	85
Ours	15.5	45	62

Table 2: Quantitative evaluation of 4D unconditional generation of animation sequences from temporally deforming 3D shapes. HyperDiffusion enables a compact, high-fidelity representation space for synthesis, producing much more detailed, high-quality results than voxel-based diffusion.

Init.	MMD ↓	COV (%) ↑	1-NNA (%) ↓	FPD ↓
Random	3.34	49	70.4	4.01
1st MLP	3.45	49	69.3	3.49

Table 3: Ablation on weight initialization for 3D shape generation (airplanes). Using a consistent initialization from a single optimized MLP leads to moderately improved results, most noticeably in the perceptual FPD metric.

	MMD ↓	COV (%) ↑	1-NNA (%) ↓	FPD ↓
No PE	3.75	44	76.5	6.40
With PE	3.45	49	69.3	3.49

Table 4: Ablation on positional encoding in the MLP for 3D shape generation (airplanes). Without positional encoding (No PE), performance noticeably degrades.

6.5. Unconditional 4D Generation

As HyperDiffusion can model neural fields that represent arbitrary dimension data, we exploit the compact rep-

resentation space of neural fields to model 4D sequences of deforming 3D shapes. The fourth dimension here corresponds to time. Table 2 and Figure 7 show that our approach can generate much higher quality animation sequences with more detailed representations than a voxel-based diffusion, which becomes quickly bound by its quartic growth in dimensionality. In Figure 7, we visualize 4 out of 16 generated frames for both animation sequences. The full animation sequences in our video and website show that we can generate temporally consistent 4D animations corresponding to meaningful actions (*e.g.*, jumping, rotating, and resting). Shape integrity is preserved throughout the generated animation sequence.

7. Limitations

While HyperDiffusion shows a promising approach towards directly generating neural fields, several limitations remain. For instance, the diffusion process currently operates only on optimized MLP parameters, without knowledge of any surface reconstruction. Our initial experiments with a naive secondary reconstruction loss did not improve performance, but we believe a more sophisticated formulation to encourage the diffusion process to be aware of the surface being represented would improve its generative modeling capabilities. Additionally, we operate on datasets of individual MLPs, while neural implicit scene representations for large-scale environments [22, 36] typically employ MLPs on a grid for greater spatial capacity. Extending to modeling multiple MLPs could enable larger-scale scene surface modeling.




Figure 6: Qualitative comparison for 3D shape generation. Voxel-based diffusion produces relatively low-resolution outputs, while state-of-the-art PVD [54] and DPC [30] synthesize discrete point clouds. In contrast, our neural field synthesis can generate a high-quality, continuous surface representation easily extracted as a mesh.

8. Conclusion

We proposed HyperDiffusion, a new generative modeling paradigm for neural implicit fields. We exploit the compact representation power of neural fields for modeling high-dimensional surface data, and model the weight space of the neural fields with a diffusion process. This enables high-fidelity surface representations by optimizing neural field MLPs to fit to individual train samples, and using the




Figure 7: Qualitative comparison of 4D animation synthesis. Our neural field synthesis generates not only more detailed animations but also achieves smoother temporal consistency. We also refer to our video for the animated shape results.




Figure 8: Novel shape generation vs nearest neighbor retrieval. For generated shapes (red) from our method, we look up the top-3 nearest neighbors (green) from the training set based on the Chamfer distance. As shown, our method does not simply memorize train samples and can generalize to novel shapes.

optimized MLP weights to train the downstream diffusion model. We can then synthesize new neural fields as their MLP weights from the diffusion model – a low-dimensional representation that decodes to high-fidelity shape surfaces for 3D shapes and 4D animation sequences of deforming shapes. Overall, we believe that our method is a first step to open up new possibilities for generative modeling of high-dimensional, complex data, and alternative representations in the context of diffusion models.

Acknowledgements. This work was supported by the Bavarian State Ministry of Science and the Arts coordinated by the Bavarian Research Institute for Digital Transformation (bidt), the ERC Starting Grant Scan2CAD (804724), and the German Research Foundation (DFG) Research Unit “Learning and Simulation in Visual Computing.” Apple was not involved in the evaluations and implementation of the code.

References

- [1] Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, and Yaser Sheikh. Modeling facial geometry using compositional vaes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3877–3886, 2018. 4
- [2] Gavin Barill, Neil G Dickson, Ryan Schmidt, David IW Levin, and Alec Jacobson. Fast winding numbers for soups and clouds. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018. 4
- [3] Miguel Angel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, et al. Gaudi: A neural architect for immersive 3d scene generation. *arXiv preprint arXiv:2207.13751*, 2022. 3
- [4] Andreas Blattmann, Robin Rombach, Kaan Oktay, and Björn Ommer. Retrieval-augmented diffusion models, 2022. 3
- [5] Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J Davison. Codeslam—learning a compact, optimisable representation for dense visual slam. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2560–2568, 2018. 4
- [6] Rohan Chabra, Jan Eric Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *ECCV*, 2020. 1
- [7] 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, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022. 2
- [8] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. 3
- [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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5799–5809, 2021. 2, 3
- [10] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. 5, 6, 7
- [11] Gene Chou, Yuval Bahat, and Felix Heide. Diffusionsdf: Conditional generative modeling of signed distance functions. *arXiv preprint arXiv:2211.13757*, 2022. 3, 4
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 3
- [13] Yilun Du, Katie Collins, Josh Tenenbaum, and Vincent Sitzmann. Learning signal-agnostic manifolds of neural fields. *Advances in Neural Information Processing Systems*, 34:8320–8331, 2021. 3
- [14] Yilun Du and Igor Mordatch. Implicit generation and generalization in energy-based models. *arXiv preprint arXiv:1903.08689*, 2019. 3
- [15] Emilien Dupont, Hyunjik Kim, SM Eslami, Danilo Rezende, and Dan Rosenbaum. From data to functa: Your data point is a function and you should treat it like one. *arXiv preprint arXiv:2201.12204*, 2022. 3
- [16] Emilien Dupont, Yee Whye Teh, and Arnaud Doucet. Generative models as distributions of functions. *arXiv preprint arXiv:2102.04776*, 2021. 3
- [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. *arXiv preprint arXiv:2209.11163*, 2022. 2
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 2, 3
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 3
- [21] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022. 3
- [22] Chiyu Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Nießner, Thomas Funkhouser, et al. Local implicit grid representations for 3d scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6001–6010, 2020. 7
- [23] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021. 2

- [24] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. [2](#)
- [25] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020. [2](#)
- [26] Marian Kleineberg, Matthias Fey, and Frank Weichert. Adversarial generation of continuous implicit shape representations. *arXiv preprint arXiv:2002.00349*, 2020. [6](#), [7](#)
- [27] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12706–12716, 2021. [2](#), [5](#)
- [28] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987, Anaheim, California, USA, July 27-31, 1987*, pages 163–169, 1987. [2](#), [4](#), [6](#)
- [29] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [5](#)
- [30] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2837–2845, 2021. [5](#), [6](#), [7](#), [8](#)
- [31] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. [1](#), [2](#), [4](#)
- [32] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, and Matthias Nießner. Diffri: Rendering-guided 3d radiance field diffusion. *arXiv preprint arXiv:2212.01206*, 2022. [3](#)
- [33] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. [3](#)
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 165–174, 2019. [1](#), [2](#)
- [35] William Peebles, Ilija Radosavovic, Tim Brooks, Alexei A Efros, and Jitendra Malik. Learning to learn with generative models of neural network checkpoints. *arXiv preprint arXiv:2209.12892*, 2022. [4](#)
- [36] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 523–540. Springer, 2020. [7](#)
- [37] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. [3](#)
- [38] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. [6](#)
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. [3](#)
- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [3](#)
- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [3](#)
- [42] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. volume abs/2201.00273, 2022. [2](#)
- [43] J Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3d neural field generation using triplane diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20875–20886, 2023. [3](#)
- [44] Yawar Siddiqui, Justus Thies, Fangchang Ma, Qi Shan, Matthias Nießner, and Angela Dai. Texturify: Generating textures on 3d shape surfaces. *arXiv preprint arXiv:2204.02411*, 2022. [2](#)
- [45] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, and Yaniv Taigman. Text-to-4d dynamic scene generation. *arXiv:2301.11280*, 2023. [3](#)
- [46] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020. [1](#), [2](#)
- [47] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. [3](#), [5](#)
- [48] 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. [2](#)
- [49] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit H Bermano. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. [3](#)
- [50] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [6](#)

- [51] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. Lion: Latent point diffusion models for 3d shape generation. *arXiv preprint arXiv:2210.06978*, 2022. 3, 5, 6
- [52] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *arXiv preprint arXiv:2301.11445*, 2023. 3, 6
- [53] Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G Schwing, and Alex Colburn. Generative multiplane images: Making a 2d gan 3d-aware. *arXiv preprint arXiv:2207.10642*, 2022. 2
- [54] Linqi Zhou, Yilun Du, and Jiajun Wu. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5826–5835, 2021. 5, 6, 7, 8
- [55] Peiye Zhuang, Samira Abnar, Jiatao Gu, Alex Schwing, Joshua M Susskind, and Miguel Ángel Bautista. Diffusion probabilistic fields. In *International Conference on Learning Representations*. 3