

# SplatFields: Neural Gaussian Splats for Sparse 3D and 4D Reconstruction

Marko Mihajlovic<sup>1\*</sup>, Sergey Prokudin<sup>1,3</sup>, Siyu Tang<sup>1</sup>, Robert Maier<sup>2</sup>,  
Federica Bogo<sup>2</sup>, Tony Tung<sup>2</sup>, and Edmond Boyer<sup>2</sup>  
ETH Zürich<sup>1</sup> Meta Reality Labs<sup>2</sup> Balgrist University Hospital<sup>3</sup>

[markomih.github.io/SplatFields](https://markomih.github.io/SplatFields)

**Abstract.** Digitizing 3D static scenes and 4D dynamic events from multi-view images has long been a challenge in computer vision and graphics. Recently, 3D Gaussian Splatting (3DGS) has emerged as a practical and scalable reconstruction method, gaining popularity due to its impressive reconstruction quality, real-time rendering capabilities, and compatibility with widely used visualization tools. However, the method requires a substantial number of input views to achieve high-quality scene reconstruction, introducing a significant practical bottleneck. This challenge is especially severe in capturing dynamic scenes, where deploying an extensive camera array can be prohibitively costly. In this work, we identify the lack of spatial autocorrelation of splat features as one of the factors contributing to the suboptimal performance of the 3DGS technique in sparse reconstruction settings. To address the issue, we propose an optimization strategy that effectively regularizes splat features by modeling them as the outputs of a corresponding implicit neural field. This results in a consistent enhancement of reconstruction quality across various scenarios. Our approach effectively handles static and dynamic cases, as demonstrated by extensive testing across different setups and scene complexities.

**Keywords:** Novel view synthesis · Gaussian splatting · Implicit models

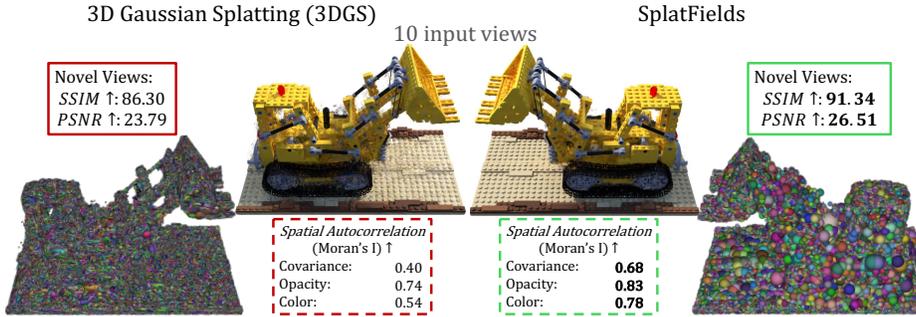
## 1 Introduction

Building a realistic replica of static and dynamic environments can revolutionize the world by transforming the way we interact, work, and engage online [67]. This ambitious vision has motivated a surge in recent research to develop new representations and rendering techniques that allow for comprehensive and photorealistic capture and reconstruction of scenes from multi-view imagery.

Recent advancements, notably the introduction of Neural Radiance Fields (NeRF) [46], have shown exceptional quality in photorealistic 3D reconstruction from casually captured images. This success comes from modeling a 3D scene as a neural field [77] and optimizing it through volume rendering techniques. The parameterization of the rendering volume using a continuous differentiable field presents several benefits. It enables a compact representation of the scene’s geometry and appearance through neural network weights, offering a more practical

---

\*Partly done during an internship at Meta.



**Fig. 1: SplatFields** regularizes 3D Gaussian Splatting (3DGS) [28] by predicting the splat features and locations via neural fields to improve the reconstruction under unconstrained sparse views. We measure spatial autocorrelation (Moran’s I [47]) of splat features in the local neighborhoods to assess their similarity and observe that better reconstruction quality achieved by our method corresponds to higher Moran’s I. The figure presents the results of a static reconstruction from ten calibrated images from Blender dataset [46]. Metrics are reported on the full test set; the rendered view is a novel view.

alternative to explicit volume modeling, which is often unfeasible. Crucially, for the focus of this work, the continuous nature and the spectral bias [59] of Multi-Layer Perceptrons (MLPa) introduce a *spatial bias*—nearby primitives are likely to exhibit similar features as predicted by the neural field MLP. This concept of implicitly modeling spatiotemporal signals has captured the research community’s attention in recent years [19], marking a significant shift in methods for 3D scene reconstruction and novel view synthesis. A substantial portion of research has also focused on adapting these methods for sparse view setups [49, 87] and enhancing training and rendering efficiency [8, 48, 53].

3D Gaussian Splatting (3DGS) [28] offers an alternative 3D reconstruction framework using point-based rasterization rather than computationally demanding volume rendering. The method quickly gained traction within both the computer vision and graphics communities due to its real-time rendering capabilities, potential compatibility with the standard rasterization pipelines, and the intuitive way of editing and combining the reconstructed scenes. This makes 3DGS a practical and scalable solution that is currently being rapidly adopted and supported by many 3D development platforms and visualization tools [9, 18, 70].

3D Gaussian Splatting represents the 3D scene as a set of unordered 3D Gaussian primitives, rendered from arbitrary views via rasterization, akin to traditional point splatting techniques [4, 92, 93]. Each rendering primitive comprises trainable parameters such as position, orientation, scale, color, and opacity, which are optimized by rendering the representation with respect to multi-view input images. The flexible parameterization, coupled with the efficient rasterization framework, is key to high-quality novel view synthesis results at scale. However, the flexibility of rendering primitives comes at the cost of requiring a

large number of input views to fully constrain the optimization process, making Gaussian splatting unsuitable for more practical captures from sparse views.

We analyze the performance of 3DGS and its 4D variants [82, 83] in sparse input view scenarios. We first show that splat-based techniques, with their independently modeled set of rendering primitives, are particularly vulnerable to the training view overfitting in such cases (see Fig. 1). In contrast, volumetric rendering techniques [69] which imply shared feature representations appear to be more robust in such scenarios as demonstrated in [44], at the expense of a significantly increased training time and suboptimal rendering efficiency. This key observation provides the basis for the method introduced in this work.

Our key idea is to regularize the behavior of independent Gaussian primitives by utilizing neural networks that regress splat features at different levels. First, inspired by [66], we aim to enforce the spatial bias through a hierarchical convolutional decoder [54] that outputs a tri-plane representation [7] of deep features associated with each splat. Please note that the tri-plane representation and the associated network are utilized only during the optimization phase to constrain the attributes of the Gaussian primitives; both are discarded thereafter for accelerated rendering and compatibility with established splat rasterization pipelines. The produced deep splat features are then utilized to condition neural fields [77] that model the geometric and appearance properties of Gaussian splats at various locations and time steps. This design is equipped with positional encoding [46] to represent high-frequency details while retaining the good spatial properties required to regularize Gaussian splatting.

We thoroughly analyze our representation (dubbed *SplatFields*) and demonstrate its superior reconstruction quality under sparse input views compared to alternative 3D Gaussian splatting techniques [16, 21, 28, 89]. We further present an effective extension of our optimization framework to model dynamic 4D scenes and propose a new forward-flow field formulation to model the dynamics of Gaussian splats, warping rendering primitives into the observation space. We observe that existing techniques that model 3D splat deformations either lack the modeling capacity due to simplified assumptions on scene motion [83] or have an insufficient spatial bias in the model [76], leading to suboptimal performance in sparse setups. Therefore, we introduce a forward-flow neural network for 3D Gaussians based on the recent ResFields MLP architecture [44]. Our method outperforms recent baselines [76, 82, 83] while retaining the key properties of Gaussian splatting, such as rendering efficiency and compatibility with existing frameworks. In summary, our key contributions are:

- We propose a novel optimization strategy, named *SplatFields*, which introduces spatial bias into the 3D Gaussian Splatting technique to stabilize the optimization process under sparse views.
- We extend our formulation to dynamic scenes, demonstrating superior reconstruction quality compared to recent state-of-the-art methods [76, 82, 83].
- We provide a detailed analysis of various modeling strategies, confirming the optimality of our framework for the tasks of sparse multi-view reconstruction.

The code is publicly available: [markomih.github.io/SplatFields](https://markomih.github.io/SplatFields).

## 2 Related Work

**Implicit volumetric rendering** [14]. Novel View Synthesis (NVS) enables the generation of new images from arbitrary viewpoints using a given set of input images [46, 69]. Over the past few years, the predominant method of choice for NVS has been the Neural Radiance Field (NeRF) [46], which represents a 3D scene as a continuous neural field [77]. This field takes as input a location and view direction and predicts color and density. Then, the color of a pixel from an arbitrary viewpoint is rendered by casting a ray and employing volume rendering, which requires sampling multiple points along the ray and converting them into color and density values by querying the neural field. Numerous extensions have been proposed to handle various scenarios and setups such as sparse view reconstruction [10, 12, 13, 26, 43, 49, 87, 90], dynamic [5, 36, 51, 52, 56, 72] and large scale unbounded scenes [3, 68]. However, the implicit volume rendering process is inherently expensive due to the large number of sampled points whose predictions need to be integrated. Despite recent efforts to accelerate NeRFs [7, 8, 35, 48, 60, 64], achieving interactive rendering capabilities for regular scenes remains challenging without additional post-processing or compression [15, 61, 84, 86].

**Point-based rendering.** The limitations of volumetric rendering methods have led to a resurgence in point-based techniques [1]. The seminal work by [20] introduced the rasterization of fixed-size, unstructured point samples for NVS. However, this naive rendering approach often results in aliasing artifacts and images with holes. These issues have been partially addressed by employing splatting techniques, where points are rendered with extended sizes to cover multiple pixels, using shapes like circular ellipsoids or surfels [4, 92, 93]. The era of deep learning led to a new wave of point-based neural rendering methods which allowed differentiable point rendering [33, 62, 75, 85] and combined point-based rasterizers with 2D convolutional networks [2, 32, 45].

3D Gaussian Splatting (3DGS [28]) utilizes the volumetric composition of ordered splats to merge the advantages of volumetric representations with exceptional real-time rendering capabilities. Gaining rapid popularity due to its efficiency, 3DGS has been incorporated into a wide range of downstream tasks [9, 18]. Modifications to the original framework have enhanced its robustness to novel views [89], improved geometry reconstruction [21], and reduced model size [16]. Nevertheless, 3DGS’s dependence on numerous *independent* splatting primitives necessitates a large number of views for effective optimization, impacting its performance in sparse view reconstructions. Our work introduces neural networks to regularize splat behavior by regressing splat features based on their 3D location, introducing the spatial autocorrelation bias which substantially enhances reconstruction in sparse scenarios, as demonstrated in our experiments.

**Dynamic Gaussian splatting.** Several recent modifications of 3DGS have been proposed to extend its capabilities to dynamic sequences. The dynamic 3DGS [42] extends the basic pipeline by optimizing the motion of each splatting primitive and the change in its corresponding features. Although regularizations such as the enforcement of local rigidity and isometry losses [29, 55] help stabilize the learning process to a certain extent, the overall pipeline still requires a large

number of input views for each step and struggles to reconstruct scenes faithfully from sparse observations. The works closest to ours are [82, 83], which utilize MLPs to model time-dependent deformations of Gaussian splats. However, the single MLP used for modeling deformation in [83] lacks the capacity to represent non-trivial scene dynamics, while the parameterizations used in [76, 82] lack a substantial spatial and temporal autocorrelation bias, leading to suboptimal reconstructions in sparse view scenarios. We thoroughly analyze the behavior of the aforementioned methods and compare them with our approach on several dynamic scenes of varying complexity and view sparsity. We demonstrate that our model, which combines a triplane-based CNN generator [7] for splat features with ResFields-based [44] dynamics modeling, offers the optimal combination of expressivity and robustness in sparse capture scenarios.

A growing body of work also addresses dynamic scenes [11, 25, 37, 39, 41, 88] and template-based approaches [34] for modeling full-body [22, 23, 31, 50, 58, 79] and head avatars [57, 63]. In contrast, our model is capable of handling unbounded, topologically varying generic dynamic scenes.

### 3 Preliminaries: 3D Gaussian Splatting

In the following, we provide a brief overview of the Gaussian splatting rendering technique, which is a fundamental building block of our model.

**Scene representation.** 3DGS [28] parametrizes the 3D scene via static 3D Gaussian primitives  $\{\mathcal{G}_k\}_{k=1}^K$  that contain the geometric and appearance information. These rendering primitives are utilized for efficient differentiable rasterization-based volume splatting.

The geometry of each Gaussian splat  $\mathcal{G}_k$  is defined by the mean location  $\mathbf{p}_k \in \mathbb{R}^{3 \times 1}$ , the opacity value  $\alpha_k \in [0, 1]$ , and the covariance matrix  $\Sigma_k \in \mathbb{R}^{3 \times 3}$  defined in the world space. Each splatting primitive  $\mathcal{G}_k$  then induces the following Gaussian distribution in 3D space:

$$\mathcal{G}_k(\mathbf{x}) \propto \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_k)^T \Sigma_k^{-1} (\mathbf{x} - \mathbf{p}_k)\right), \quad (1)$$

where the covariance matrix is modeled by the scaling vector  $\mathbf{s}_k \in \mathbb{R}^3$  and the rotation matrix  $\mathbf{O}_k \in \mathbb{R}^{3 \times 3}$  (parameterized via quaternions) to ensure positive semi-definiteness:

$$\Sigma_k = \mathbf{O}_k \mathbf{s}_k \mathbf{s}_k^T \mathbf{O}_k^T. \quad (2)$$

The appearance of splats is view-dependent and described by  $C$  coefficients that are converted to color  $\mathbf{c}_k \in \mathbb{R}^3$  via spherical harmonics, similar to [86].

**Rendering.** Given an arbitrary camera viewpoint described by the rotation  $\mathbf{R} \in \mathbb{R}^{3 \times 3}$  and translation  $\mathbf{t} \in \mathbb{R}^{3 \times 1}$ , we can obtain the 2D coordinates of the splat center  $\mathbf{p}'_k \in \mathbb{R}^2$  on the image plane:

$$\mathbf{p}'_k = (\mathbf{R}\mathbf{p}_k + \mathbf{t})_{1:2} / (\mathbf{R}\mathbf{p}_k + \mathbf{t})_3. \quad (3)$$

Further, we can obtain the 2D projection  $\Sigma_k^{2D}$  of the covariance matrix on the image plane:

$$\Sigma_k^{2D} = (\mathbf{J}_k \Sigma'_k \mathbf{J}_k^T)_{1:2,1:2} \in \mathbb{R}^{2 \times 2}, \text{ where } \Sigma'_k = \mathbf{R} \Sigma_k \mathbf{R}^T, \quad (4)$$

and the Jacobian  $\mathbf{J}_k \in \mathbb{R}^{3 \times 3}$  is an affine approximation to the projective transformation (see [92] for details). The subscript  $1:2$  denotes row and column selection.

Using the image-space splat center and 2D covariance matrix, we obtain the 2D image-space Gaussian distribution induced by the corresponding splat:

$$\mathcal{G}_k^{2D}(\mathbf{x}') \propto \exp\left(-\frac{1}{2}(\mathbf{x}' - \mathbf{p}'_k)^T (\Sigma_k^{2D})^{-1} (\mathbf{x}' - \mathbf{p}'_k)\right), \quad (5)$$

Finally, we can predict the color  $\mathbf{c}(\mathbf{x}') \in \mathbb{R}^3$  at each pixel location  $\mathbf{x}' \in \mathbb{R}^2$  by blending the splats, sorted according to their projection depth:

$$\mathbf{c}(\mathbf{x}') = \sum_{k=1}^K \alpha_k \mathbf{c}_k \mathcal{G}_k^{2D}(\mathbf{x}') \prod_{j=1}^{k-1} (1 - \alpha_j \mathcal{G}_j^{2D}(\mathbf{x}')), \quad (6)$$

where  $\alpha_k$  is the learned opacity of the splat.

**Training.** The collection of splats  $\{\mathcal{G}_k\}_{k=1}^K$  is optimized by minimizing the following rendering loss w.r.t the input images via the Adam [30] optimizer:

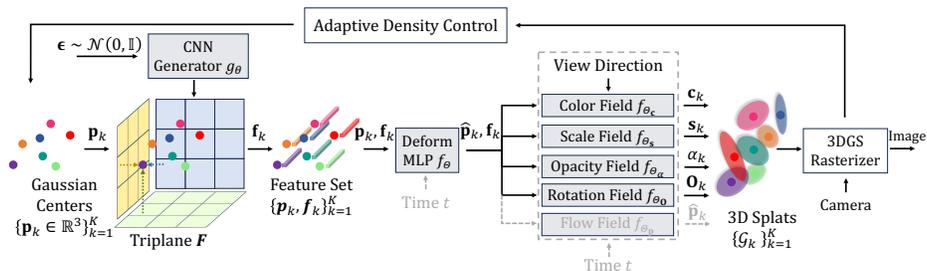
$$\mathcal{L} = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{D\text{-SSIM}}, \quad (7)$$

where the first term is a standard  $\mathcal{L}_1$  loss between the target and rendered images, and  $\mathcal{L}_{D\text{-SSIM}}$  is a differentiable version of structural similarity index [74]. As this optimization is highly sensitive to a local minima, 3DGS additionally employs periodic adaptive densification and pruning of splats through randomized sampling. We refer to [28] for further details.

## 4 SplatFields: Neural Gaussian Splats

**Limitations of 3DGS.** Modeling 3D scenes with irregularly spaced point primitives offers significant flexibility and facilitates rapid and efficient optimization when an extensive number of training views is provided. However, with limited views, these independent point primitives are prone to overfitting. Therefore, we advocate for integrating a spatial autocorrelation bias within the splats. This can be accomplished by deriving splat features through implicit neural models, presenting a viable method to constrain and regularize the otherwise ill-posed optimization in sparse-view environments.

**Key insight.** We analyze the optimization procedure of 3DGS under sparse view inputs and observe (Fig. 1) that the splats do not exhibit any local structure and display incoherent patterns. To quantify the local spatial autocorrelation of each splat, we select the five nearest neighbors and measure Moran’s I [47] of the splat’s attributes (color, opacity, covariance). We note that a low level of spatial autocorrelation is associated with overfitting to training views, which impedes



**Fig. 2: Overview.** SplatFields takes as input a point cloud (e.g., initialized from SfM [65]), for which it models the geometric (position  $\mathbf{p}_k$ , scale  $\mathbf{s}_k$ , rotation  $\mathbf{O}_k$ ) and appearance attributes (color  $\mathbf{c}_k$ , opacity  $\alpha_k$ ). These attributes represent the point set as 3D splats that are then rendered with the 3DGS rasterizer [28]. First, the point location set  $\{\mathbf{p}_k \in \mathbb{R}^3\}_{k=1}^K$  is encoded into features  $\{\mathbf{f}_k\}_{k=1}^K$  by sampling the tri-plane representation generated by a CNN generator  $g_\theta$  to provide a deep structural prior [71] on the feature values. These values are then propagated through a deformation MLP  $f_\theta$  to refine the point locations  $\hat{\mathbf{p}}_k$ . The new point set, along with the features, is then propagated through a series of compact neural fields to predict the properties of rendering primitives  $\{\mathcal{G}_k\}_{k=1}^K$  that are rendered with respect to arbitrary viewpoints. During the optimization, we adopt the adaptive density control [28] to periodically prune and densify the point set. SplatFields seamlessly adapts to 4D reconstruction by conditioning neural fields on the time step  $t$  and introducing an extra time-conditioned flow field. Gray blocks indicate learnable modules.

the learning of a structured 3D representation. This correlation is further showcased in our discussion and empirical evidence presented in Tab. 1.

The core idea of our method is to introduce a *spatial bias* during the optimization phase, which encourages nearby primitives to share similar features, thereby emulating the more continuous behavior characteristic of widely used implicit representations for volumetric rendering [44, 46, 69]. However, directly enforcing this constraint—ensuring local neighborhoods exhibit common patterns—yields sub-optimal performance (Tab 2) for volumetric point representations. To overcome this, we propose a novel neural framework, termed *SplatFields*, designed to adaptively regularize the optimization of 3DGS. Importantly, SplatFields straightforwardly extends to 4D, facilitating the reconstruction of dynamic scenes.

SplatFields (Fig. 2) builds on the core property of neural networks to discover local patterns and fit low frequencies of a signal first [59, 71]. To that end, we implement SplatFields as a neural generator that infers the attributes of Gaussian splats. The neural generator combines key properties of convolutional neural networks, which model local structured patterns, with multi-layer perceptrons that serve as global approximators. This approach straightforwardly extends to 4D reconstruction by conditioning the MLP networks on time  $t$ .

**Deep structural prior.** First, we follow the idea of a deep image prior [66, 71] and aim to utilize CNNs to model locally structured patterns of splat fea-

**Table 1: Impact of the spatial autocorrelation** on static scene reconstruction. Results on OwlII [78] dataset. See Section 5.1 for discussion

	Train			Test			Spatial Autocorrelation		
	View Synthesis			Novel View Synthesis			(Moran's I) $\uparrow$		
	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	SSIM $\uparrow$	PSNR $\uparrow$	LPIPS $\downarrow$	Color	Opacity	Covariance
3DGS [28]	99.85	44.07	0.493	91.68	27.50	8.881	0.547	0.670	0.232
SplatFields3D	98.87	37.58	3.274	<b>96.13</b>	<b>30.33</b>	<b>5.973</b>	<b>0.935</b>	<b>0.874</b>	<b>0.431</b>

tures. In the original work [71], the CNN takes as input low-dimension Gaussian noise  $\epsilon \sim \mathcal{N}(0, \mathbb{I})$  and gradually upsamples it into the desired image resolution; the weights of the network are then optimized to fit the observed noisy image. In our case, we aim to generate a 3D field of splat features; as 3D CNNs are computationally prohibitive, we use 2D CNNs that generate axis-aligned tri-plane representations [7, 54]. Overall, the step is a splat-based variation of the approach utilized in [66] for a fully volumetric NeRF-based sparse rendering.

Specifically, given a randomly initialized noise  $\epsilon$ , the convolutional network  $g_\theta$  regresses the three  $H \times W$ -resolution planes  $\mathbf{F}$ :

$$\mathbf{F} = g_\theta(\epsilon) \in \mathbb{R}^{3 \times H \times W \times l}, \quad (8)$$

where  $l$  denotes the feature dimension and  $\theta$  indicates learnable network weights. The overall CNN structure resembles the one originally proposed in [66].

**Neural splat fields.** Next, the splat center  $\mathbf{p}_k$  is projected onto each of the three feature planes to obtain feature vectors via bilinear interpolation. These features are then concatenated along the feature dimension and denoted as  $\mathbf{f}_k \in \mathbb{R}^{3l}$ . The feature and the initial point are propagated through a deformation MLP  $f_\theta$  which refines the position of the input point:

$$\hat{\mathbf{p}}_k = f_\theta(\mathbf{p}_k, \mathbf{f}_k, t), \quad (9)$$

where  $t$  indicates an optional time step input provided in the case of dynamic 4D reconstruction. Finally, the updated point location, along with the inferred feature vector, is provided as input to a set of compact (5-6 layers, 64-128 neurons) neural fields  $\{f_{\theta_c}, f_{\theta_s}, f_{\theta_\alpha}, f_{\theta_o}\}$  to obtain properties of Gaussian splats. Rather than using spherical harmonics to model color, we directly predict view-dependent color. The obtained splats are then rendered w.r.t. the input views to optimize the learnable modules by minimizing the photometric loss (Eq. 7).

**Splat norm regularization.** For static reconstruction, we add additional norm regularization  $\|\hat{\mathbf{p}}_k\|_2$  to the loss function to bias the resulting splats to not deviate significantly from the origin, similar to the floor loss considered in [42].

**4D reconstruction.** SplatFields is a flexible representation that straightforwardly extends to 4D reconstruction of dynamic scenes. It models temporal variations in the splat features by conditioning the corresponding neural field MLPs  $\{f_\theta, f_{\theta_c}, f_{\theta_s}, f_{\theta_\alpha}, f_{\theta_o}\}$  on the time step  $t$ . Additionally, we add a time-conditioned flow field  $f_p$  to warp the center of Gaussians  $\mathbf{p}_k$  to the desired time step  $t$ . To enhance the expressivity of the neural fields and allow for complex geometry changes, we utilize the recently proposed ResField MLP architecture [44].

**Table 3: Ablation study of SplatFields.** Blender dataset [46], setup from Tab. 2

	12 Views		10 Views		8 Views		6 Views		4 Views	
	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑
basic (MLP-only)	89.63	24.82	88.69	23.93	87.79	23.32	85.70	21.58	81.98	19.07
+ $\mathcal{L}_2$ -norm reg.	89.66	24.98	88.81	24.21	87.98	23.64	85.93	21.79	82.44	19.46
+tri-CNN	90.83	25.23	90.04	24.66	88.01	23.19	85.91	21.46	81.28	18.72
full model	91.18	25.80	90.32	24.94	88.94	23.98	86.62	22.26	82.27	19.16

## 5 Experiments

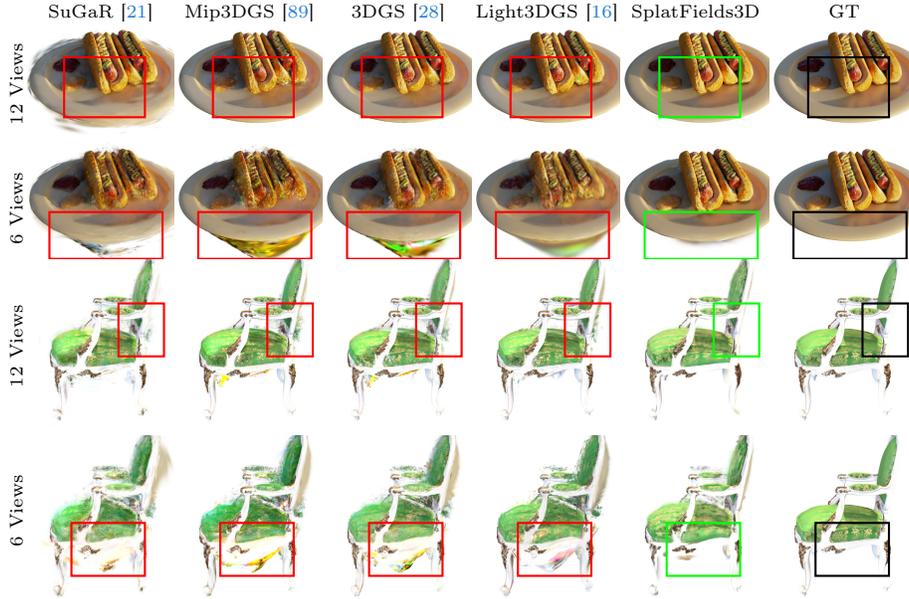
### 5.1 Static Scene Reconstruction

**Impact of the spatial autocorrelation.** First, we conduct a toy experiment to verify our intuition that the absence of the *spatial bias* hampers the reconstruction quality from sparse views. We utilize four sequences from the OwlII dynamic dataset [78]<sup>1</sup> and select the first frame from each. Each scene comprises nine training views and one validation view, on which we report training and test metrics. We compare 3DGS [28] and SplatFields (both initialized from visual hulls) and observe (Tab. 1) that 3DGS demonstrates extremely high fitting quality on the training views while poorly generalizing to the novel views. In contrast, SplatFields demonstrates a slightly lower training quality while achieving higher reconstruction quality on novel views. This observation is followed by computing the Moran’s I metric [47] which shows the amount of spatial correlation between nearby splat features; as hypothesized, the lower test-time error is in line with the increased level of spatial consistency of all groups of splat features. Fig. 1 presents both qualitative and quantitative results from the same experiment conducted on a scene from the Blender dataset [46].

**Static reconstruction from sparse views.** We benchmark SplatFields on Blender [46] under 6 and 12 input views (see Sup. Mat. for more extensive benchmarking). The main goal of this section is to showcase the efficiency of the utilized spatial regularization for 3DGS methods. We, therefore, focus on comparison against the recent splat-based techniques [16, 21, 28, 89] and SparseNeRF [73], leaving the comprehensive comparison against a broader range of NeRF-based methods for more challenging dynamic scenarios considered in Sec. 5.2.

Extensive quantitative results presented in Tab. 2 demonstrate that SplatFields consistently outperforms the respective baselines across varying numbers of input views. The achieved improvement is also verified by visually sharper reconstructions outlined in Fig. 3. More importantly, we observe that the relative gap in performance between our method and the baselines is increasing as the input views become more scarce, which confirms our intuition of the spatial bias as a powerful regularizer in such scenarios. We further validate our improvement on the real-world DTU dataset [27] for the challenging task of 3-view recon-

<sup>1</sup>All experiments presented in this publication were performed by ETH Zürich. ETH Zürich obtained the licenses for the data used in such experiments.

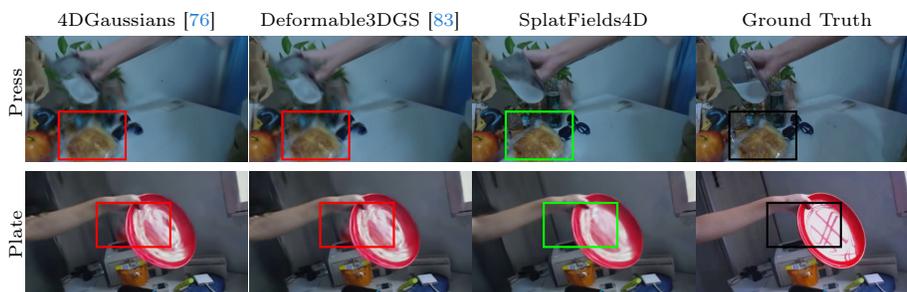
**Fig. 3: Static reconstruction of Blender [46] scenes for the setup from Tab. 2**

	12 Input Views									
	<i>mean</i>	Toy	Ficus	Hotdog	Chair	Mic	Ship	Drums	Materials	
SparseNeRF [73]	-	23.02	18.19	-	26.20	23.26	20.81	19.21	20.80	
SparseNeRF <i>wo. depth</i>	22.92	24.00	18.84	27.52	27.11	23.35	21.84	19.17	21.50	
SuGaR [21]	21.78	23.77	23.08	22.36	25.72	18.72	21.09	19.55	19.94	
ScaffoldGS [40]	23.82	23.65	22.78	26.34	25.80	28.28	21.17	20.47	22.06	
Mip3DGS [89]	24.86	24.65	25.62	26.53	26.25	28.40	22.52	21.98	22.94	
3DGS [28]	25.29	25.14	25.92	27.51	27.10	29.02	22.79	22.10	22.71	
Light3DGS [16]	25.39	25.08	27.53	27.10	27.40	28.04	23.02	22.07	22.90	
2DGS [24]	25.62	25.50	25.62	29.24	28.52	28.07	23.08	22.19	22.75	
3DGS <i>w. <math>\mathcal{L}_{Moran}</math></i>	25.44	25.26	26.55	28.96	27.91	27.87	22.33	21.98	22.65	
SplatFields3D	25.80	26.98	26.27	29.45	27.42	27.60	23.78	22.55	22.32	
	6 Input Views									
SparseNeRF [73]	-	20.86	18.03	-	22.75	22.40	19.33	16.24	19.54	
SparseNeRF <i>wo. depth</i>	20.86	22.62	17.63	25.84	22.65	20.72	19.85	17.25	20.30	
SuGaR [21]	19.07	19.89	20.61	20.80	21.92	18.26	17.72	16.86	16.53	
ScaffoldGS [40]	19.65	18.21	20.72	19.48	22.20	24.31	16.47	17.21	18.62	
Mip3DGS [89]	20.04	19.39	21.81	19.70	21.72	24.44	17.02	17.72	18.52	
3DGS [28]	20.62	19.80	22.25	21.16	22.75	25.21	17.58	17.77	18.48	
Light3DGS [16]	20.76	20.25	23.12	20.66	22.69	24.89	17.83	18.02	18.63	
2DGS [24]	20.74	19.38	21.93	23.85	23.26	24.48	16.92	17.91	18.17	
3DGS <i>w. <math>\mathcal{L}_{Moran}</math></i>	21.03	20.34	23.05	23.92	22.50	24.64	17.20	18.14	18.48	
SplatFields3D	22.26	22.41	22.26	26.19	25.03	24.84	19.33	18.97	19.05	

**Table 2: Sparse static scene reconstruction of Blender [46] scenes.** Reported numbers indicate PSNR metric on the novel views (“-” denotes failed runs). Colors denote the 1st, 2nd, and 3rd best-performing model. See Sec. 5.1 for discussion

**Table 4: Monocular reconstruction** of dynamic sequences from the NeRF-DS dataset [80] with recent state-of-the-art methods. The forward slash in FPS indicates the rendering speed without the neural network inference when the rendering primitives are extracted and stored for each frame *vs.* with the neural network inference

	Resources		$mean\uparrow$		$LPIPS\downarrow (\times 10^2)$							
	FPS $\uparrow$	t $\downarrow$	PSNR	SSIM	$mean$	Sieve	Plate	Bell	Press	Cup	As	Basin
3D-GS [28]	120+	15 <i>m.</i>	20.29	78.16	29.20	22.47	40.93	25.03	29.04	25.48	29.94	31.53
TiNeuVox [17]	< 1	30 <i>m.</i>	21.61	82.34	27.66	31.76	33.17	25.68	30.01	36.43	39.67	26.90
4DGaussians [76]	120+ / 50	30 <i>m.</i>	23.68	83.22	21.06	16.39	<b>23.80</b>	21.84	21.68	19.06	22.06	22.57
HyperNeRF [52]	< 1	1 <i>d.</i>	23.45	<b>84.88</b>	19.90	16.45	29.40	20.52	<b>19.59</b>	<b>16.50</b>	<b>17.77</b>	<b>19.11</b>
Deformable3DGS [83]	120+ / 30	1 <i>h.</i>	<b>23.54</b>	84.05	<b>19.79</b>	15.30	25.04	<b>15.93</b>	29.89	<b>15.38</b>	17.88	<b>19.10</b>
NeRF-DS [80]	< 1	1 <i>d.</i>	<b>23.60</b>	84.94	18.16	14.72	<b>19.96</b>	18.67	20.47	17.37	<b>17.41</b>	<b>18.55</b>
SplatFields4D	120+ / 30	1 <i>h.</i>	<b>23.84</b>	<b>85.17</b>	<b>17.86</b>	<b>14.72</b>	<b>22.43</b>	16.10	<b>19.26</b>	<b>15.67</b>	<b>17.71</b>	<b>19.11</b>



**Fig. 4: Monocular reconstruction** of sequences from [80] on the setup from Tab. 4.

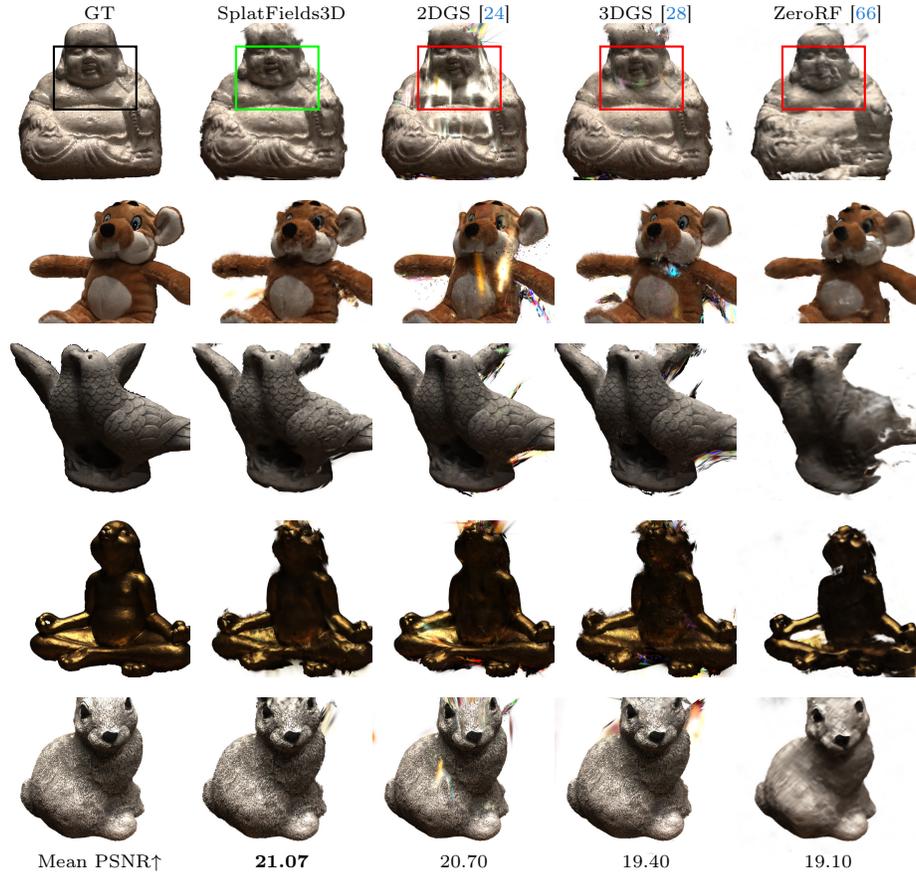
struction (Fig. 5) and demonstrate consistent improvement over NeRF- [66] and splatting-based [24, 28] baselines. See Sup. Mat. for more extensive evaluation.

Please note that all the methods, including ours, demonstrate real-time rendering performance with high interactive rates (120+FPS) during test time since the generator  $g_\theta$  is discarded after the training completion. We refer the reader to the supplementary for further details.

**Ablation of the triplane CNN generator.** We validate the impact of the proposed triplane CNN generator on the performance of the SplatFields model in Tab. 3. Here, the basic pipeline implies using only the set of MLPs to directly predict the splat rendering features (opacity, scale, etc.) and point displacements from the initial splat locations, without conditioning on the deep features produced by the triplane CNN. Results indicate that utilizing the deep features regressed by a CNN improves the quality, with the splat  $\mathcal{L}_2$ -norm regularization term further benefiting the reconstruction. Note that the regularization has a marginal improvement on the results of our pipeline that does not utilize the CNN feature generator, demonstrating the synergy of both modeling strategies.

## 5.2 Dynamic Scene Reconstruction

**Monocular dynamic reconstruction.** We further evaluate our method on seven sequences of varying lengths (ranging from 424 to 881 frames) from the



**Fig. 5: Three-view reconstruction on DTU [27];** PSNR are averaged across all 15 scenes. See Tab. ?? for the individualized scores.

NeRF-DS dataset [80] and compare our method against both NeRF- [17, 52, 80] and 3DGS-based [76, 83] dynamic reconstruction baselines.

The results are reported in Tab. 4, and the qualitative comparison of the 3DGS-based methods is presented in Fig. 4. The results demonstrate that our method achieves competitive or superior reconstruction quality across all sequences, while also maintaining real-time rendering capabilities and facilitating accelerated training processes. However, we observed that the sequences in the dataset involve relatively small motion and large static parts. Therefore, we further analyze SplatFields on multi-view sequences with more challenging dynamics.

**Multi-view dynamic reconstruction.** Following [44], we use 4 sequences from OwlII [78]. We opt for the dataset as it has realistic and more complex motion compared to the commonly utilized synthetic sequences [56]. Each sequence is 100 frames long and comprises multi-view video streams, where we

**Table 5: Multi-view reconstruction** of dynamic sequences from the OwlII dataset [46] under varying number of input views. The reported metric is PSNR averaged across novel views. The forward slash in FPS indicates the rendering speed without the neural network inference when the rendering primitives are extracted and stored for each frame *vs.* with the neural network inference.

		Resources		10 Input Views				
		FPS $\uparrow$	t $\downarrow$	<i>mean</i>	Dancer	Exercise	Model	Basketball
4D NeRFs	DyNeRF [36, 44]		1 day	29.70	28.22	30.64	29.95	30.00
	TNeRF [36, 44]		1 day	30.39	29.12	31.00	30.71	30.71
	DNeRF [44, 56]	<1	1.5 day	30.25	29.39	30.63	30.63	30.35
	Nerfies [44, 51]		1.5 day	30.70	29.57	31.08	30.53	31.60
	HyperNeRF [44, 52]		2 days	30.36	30.09	30.39	30.88	30.08
Splatting	4D-GS [82]	120+	10h	28.05	28.11	29.09	29.06	25.94
	Deformable3DGS [83]	120+/30	8h	27.76	27.86	28.78	26.47	27.95
	4DGaussians [76]	120+/50	2h	29.80	28.46	30.21	30.69	29.82
	SplatFields4D (30k it)	120+/30	2h	30.88	30.46	30.78	31.14	31.15
	SplatFields4D (40k it)	120+/30	3h	30.96	30.57	30.85	31.20	31.24
	SplatFields4D (100k it)	120+/30	7h	31.12	30.79	30.99	31.33	31.39
SplatFields4D (200k it)	120+/30	14h	31.32	31.05	31.16	31.50	31.58	
Splatting	4D-GS [82]	120+	10h	26.20	26.99	26.34	27.41	24.07
	Deformable3DGS [83]	120+/30	8h	26.06	26.77	26.24	25.61	25.62
	4DGaussians [76]	120+/50	2h	28.16	27.34	28.10	29.60	27.62
	SplatFields4D (30k it)	120+/30	2h	29.46	29.38	28.84	29.80	29.83
	SplatFields4D (40k it)	120+/30	3h	29.53	29.46	28.90	29.85	29.90
	SplatFields4D (100k it)	120+/30	7h	29.66	29.66	29.01	29.96	30.02
SplatFields4D (200k it)	120+/30	14h	29.84	29.92	29.16	30.10	30.18	
Splatting	4D-GS [82]	120+	10h	21.42	22.89	20.80	21.60	20.40
	Deformable3DGS [83]	120+/30	8h	24.46	25.37	24.31	24.12	24.02
	4DGaussians [76]	120+/50	2h	26.52	26.13	26.27	27.34	26.36
	SplatFields4D (30k it)	120+/30	2h	28.04	28.36	27.31	28.44	28.07
	SplatFields4D (40k it)	120+/30	3h	28.10	28.43	27.35	28.48	28.13
	SplatFields4D (100k it)	120+/30	7h	28.22	28.61	27.44	28.56	28.25
SplatFields4D (200k it)	120+/30	14h	28.36	28.84	27.54	28.67	28.39	
Splatting	4D-GS [82]	120+	10h	17.40	17.70	16.86	18.35	16.71
	Deformable3DGS [83]	120+/30	8h	20.04	21.42	19.56	19.71	19.45
	4DGaussians [76]	120+/50	2h	21.31	21.49	21.05	21.90	20.80
	SplatFields4D (30k it)	120+/30	2h	21.88	22.60	20.73	21.83	22.34
	SplatFields4D (40k it)	120+/30	3h	21.89	22.63	20.74	21.83	22.35
	SplatFields4D (100k it)	120+/30	7h	21.92	22.73	20.75	21.82	22.36
SplatFields4D (200k it)	120+/30	14h	21.95	22.83	20.76	21.83	22.39	

**Table 6: Flow model ablation study.** Multi-view reconstruction task from Tab. 5. See Sec. 5.2 for discussion. Symbol “-” denotes failed runs

	10 Views		8 Views		6 Views		4 Views	
	SSIM $\uparrow$	PSNR $\uparrow$						
DCT [72]	95.10	28.72	94.53	27.34	94.00	26.42	90.94	21.68
DCT+ResFields [44]	96.56	30.99	96.05	29.66	95.46	28.22	91.47	22.12
offset [76]	95.20	28.91	94.74	27.83	94.15	26.73	90.85	21.63
offset+ResFields [44]	96.75	31.24	96.28	29.83	95.68	28.38	91.41	21.62
SE3 [83]	95.33	29.05	94.78	27.95	94.39	27.09	91.13	22.09
SE3+ResFields [44]	96.81	31.32	96.28	29.84	95.69	28.36	91.60	21.95
scaled SE3 [91]	95.05	28.78	92.42	23.88	-	-	-	-
scaled SE3+ResFields [44]	96.74	31.30	96.08	29.83	95.63	28.64	92.54	23.55

vary the number of input views from 4 to 10 to study the robustness of our and baseline models. For evaluation, we use 100 images of a rotating camera around the performer, where each image comes from a different time step.

We provide a comprehensive comparison of the SplatFields4D against the respective baselines in Tab. 5. Dynamic NeRF methods [36, 51, 52, 56] improved with ResFields [44] generally require significantly longer training times, while recent dynamic 3DGS-based methods [76, 82, 83] showcase suboptimal modeling capabilities and performance in the considered sparse scenario. Our method demonstrates a clear metric improvement over the baselines while retaining the key properties of 3DGS, such as interactive rendering speed and compatibility with the existing visualization pipelines. We specifically emphasize the difference in performance between our and the recent closely related method [82], where our combination of the triplane features and MLP-based dynamics modeling proves to be more robust compared to the HexPlane-based [6] approach in the case of rapid motion and sparse camera setups.

**Splat flow model ablation.** We compare the flow field  $f_{\mathbf{p}}$  modeling deformations via DCT basis [38], translation vectors [56], SE(3) transformation [51], and via scaled SE(3) transformation [91]. We also ablate the impact of implementing neural fields via ResFields [44] to further increase the modeling capacity of our pipeline without affecting its training speed. The results (Tab. 6) suggest that modeling flow as SE(3) achieves slightly better quality when the number of views is large. We further observe that implementing neural fields via ResFields further boosts the reconstruction quality across all setups.

## 6 Conclusion

In this work, we proposed an effective optimization strategy that introduces spatial and connectivity biases into the 3D Gaussian splats during optimization process by modeling them through a continuous neural field. We demonstrated that our optimization strategy considerably enhances reconstruction quality in the sparse setups, without the need for any external, data-driven priors. Furthermore, we introduced an effective extension of our method for reconstructing dynamic sequences and demonstrated state-of-the-art results under sparse views.

**Limitations and future work.** The performance of our method noticeably diminishes in extremely sparse and highly dynamic scenarios, such as those involving rapid motion with as few as four views, exemplified by the OwlII dataset [78]. This performance is inferior when compared to the best-performing NeRF-based methods in similar sparse configurations [44]. Therefore, further exploration is required to narrow the performance gap between 3DGS- and NeRF-based methods in sparse settings. Future work should also consider incorporating learning-based priors [10, 26, 81, 87] as promising directions for advancement.

## References

1. Alexa, M., Gross, M., Pauly, M., Pfister, H., Stamminger, M., Zwicker, M.: Point-based computer graphics. In: SIGGRAPH notes (2004) [4](#)
2. Aliev, K.A., Sevastopolsky, A., Kolos, M., Ulyanov, D., Lempitsky, V.: Neural point-based graphics. In: ECCV (2020) [4](#)
3. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: CVPR (2022) [4](#)
4. Botsch, M., Hornung, A., Zwicker, M., Kobbelt, L.: High-quality surface splatting on today’s gpus. In: Proceedings Eurographics/IEEE VGTC Symposium Point-Based Graphics, 2005. pp. 17–141. IEEE (2005) [2](#), [4](#)
5. Cai, H., Feng, W., Feng, X., Wang, Y., Zhang, J.: Neural surface reconstruction of dynamic scenes with monocular rgb-d camera. In: NeurIPS (2022) [4](#)
6. Cao, A., Johnson, J.: Hexplane: A fast representation for dynamic scenes. CVPR (2023) [14](#)
7. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: CVPR (2022) [3](#), [4](#), [5](#), [8](#)
8. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: ECCV (2022) [2](#), [4](#)
9. Chen, G., Wang, W.: A survey on 3d gaussian splatting. arXiv preprint arXiv:2401.03890 (2024) [2](#), [4](#)
10. Chen, H., Gu, J., Chen, A., Tian, W., Tu, Z., Liu, L., Su, H.: Single-stage diffusion nerf: A unified approach to 3d generation and reconstruction. In: ICCV (2023) [4](#), [14](#)
11. Das, D., Wewer, C., Yunus, R., Ilg, E., Lenssen, J.E.: Neural parametric gaussians for monocular non-rigid object reconstruction. In: CVPR (2024) [5](#)
12. Deng, C., Jiang, C., Qi, C.R., Yan, X., Zhou, Y., Guibas, L., Anguelov, D., et al.: Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In: CVPR (2023) [4](#)
13. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. In: CVPR (2022) [4](#)
14. Drebin, R.A., Carpenter, L., Hanrahan, P.: Volume rendering. SIGGRAPH (1988) [4](#)
15. Duckworth, D., Hedman, P., Reiser, C., Zhizhin, P., Thibert, J.F., Lučić, M., Szeliski, R., Barron, J.T.: Smerf: Streamable memory efficient radiance fields for real-time large-scene exploration. arXiv preprint arXiv:2312.07541 (2023) [4](#)
16. Fan, Z., Wang, K., Wen, K., Zhu, Z., Xu, D., Wang, Z.: Lightgaussian: Unbounded 3d gaussian compression with 15x reduction and 200+ fps. arXiv preprint arXiv:2311.17245 (2023) [3](#), [4](#), [9](#), [10](#)
17. Fang, J., Yi, T., Wang, X., Xie, L., Zhang, X., Liu, W., Nießner, M., Tian, Q.: Fast dynamic radiance fields with time-aware neural voxels. In: SIGGRAPH Asia (2022) [11](#), [12](#)
18. Fei, B., Xu, J., Zhang, R., Zhou, Q., Yang, W., He, Y.: 3d gaussian as a new vision era: A survey. arXiv preprint arXiv:2402.07181 (2024) [2](#), [4](#)
19. Gao, K., Gao, Y., He, H., Lu, D., Xu, L., Li, J.: Nerf: Neural radiance field in 3d vision, a comprehensive review. arXiv preprint arXiv:2210.00379 (2022) [2](#)
20. Grossman, J.P., Dally, W.J.: Point sample rendering. In: Eurographics Workshop (1998) [4](#)

21. Guédon, A., Lepetit, V.: Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering. In: CVPR (2024) [3](#), [4](#), [9](#), [10](#)
22. Hu, L., Zhang, H., Zhang, Y., Zhou, B., Liu, B., Zhang, S., Nie, L.: Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. In: CVPR (2024) [5](#)
23. Hu, S., Hu, T., Liu, Z.: Gauhuman: Articulated gaussian splatting from monocular human videos. In: CVPR (2024) [5](#)
24. Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S.: 2d gaussian splatting for geometrically accurate radiance fields. In: SIGGRAPH (2024) [10](#), [11](#), [12](#)
25. Huang, Y.H., Sun, Y.T., Yang, Z., Lyu, X., Cao, Y.P., Qi, X.: Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In: CVPR (2024) [5](#)
26. Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: ICCV (2021) [4](#), [14](#)
27. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanæs, H.: Large scale multi-view stereopsis evaluation. In: CVPR (2014) [9](#), [12](#)
28. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ToG (2023) [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#), [10](#), [11](#), [12](#)
29. Kilian, M., Mitra, N.J., Pottmann, H.: Geometric modeling in shape space. In: SIGGRAPH (2007) [4](#)
30. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015) [6](#)
31. Kocabas, M., Chang, R., Gabriel, J., Tuzel, O., Ranjan, A.: Hugs: Human gaussian splats. In: CVPR (2024) [5](#)
32. Kopanas, G., Philip, J., Leimkühler, T., Drettakis, G.: Point-based neural rendering with per-view optimization. Computer Graphics Forum (2021) [4](#)
33. Lassner, C., Zollhofer, M.: Pulsar: Efficient sphere-based neural rendering. In: CVPR (2021) [4](#)
34. Lei, J., Wang, Y., Pavlakos, G., Liu, L., Daniilidis, K.: Gart: Gaussian articulated template models. In: CVPR (2024) [5](#)
35. Li, R., Gao, H., Tancik, M., Kanazawa, A.: Nerfacc: Efficient sampling accelerates nerfs. In: ICCV (2023) [4](#)
36. Li, T., Slavcheva, M., Zollhofer, M., Green, S., Lassner, C., Kim, C., Schmidt, T., Lovegrove, S., Goesele, M., Newcombe, R., et al.: Neural 3d video synthesis from multi-view video. In: CVPR (2022) [4](#), [13](#), [14](#)
37. Li, Z., Chen, Z., Li, Z., Xu, Y.: Spacetime gaussian feature splatting for real-time dynamic view synthesis. In: CVPR (2024) [5](#)
38. Li, Z., Wang, Q., Cole, F., Tucker, R., Snavely, N.: Dynibar: Neural dynamic image-based rendering. In: CVPR (2023) [14](#)
39. Lin, Y., Dai, Z., Zhu, S., Yao, Y.: Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In: CVPR (2024) [5](#)
40. Lu, T., Yu, M., Xu, L., Xiangli, Y., Wang, L., Lin, D., Dai, B.: Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In: CVPR (2024) [10](#)
41. Lu, Z., Guo, X., Hui, L., Chen, T., Yang, M., Tang, X., Zhu, F., Dai, Y.: 3d geometry-aware deformable gaussian splatting for dynamic view synthesis. In: CVPR (2024) [5](#)
42. Luiten, J., Kopanas, G., Leibe, B., Ramanan, D.: Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In: 3DV (2024) [4](#), [8](#)
43. Mihajlovic, M., Bansal, A., Zollhofer, M., Tang, S., Saito, S.: Keypointnerf: Generalizing image-based volumetric avatars using relative spatial encoding of keypoints. In: ECCV (2022) [4](#)

44. Mihajlovic, M., Prokudin, S., Pollefeys, M., Tang, S.: ResFields: Residual neural fields for spatiotemporal signals. In: ICLR (2024) [3](#), [5](#), [7](#), [8](#), [12](#), [13](#), [14](#)
45. Mihajlovic, M., Weder, S., Pollefeys, M., Oswald, M.R.: Deepsurfels: Learning on-line appearance fusion. In: CVPR (2021) [4](#)
46. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020) [1](#), [2](#), [3](#), [4](#), [7](#), [9](#), [10](#), [13](#)
47. Moran, P.A.: Notes on continuous stochastic phenomena. *Biometrika* (1950) [2](#), [6](#), [9](#)
48. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ToG* (2022) [2](#), [4](#)
49. Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In: CVPR (2022) [2](#), [4](#)
50. Pang, H., Zhu, H., Kortylewski, A., Theobalt, C., Habermann, M.: Ash: Animatable gaussian splats for efficient and photoreal human rendering. In: CVPR (2024) [5](#)
51. Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV (2021) [4](#), [13](#), [14](#)
52. Park, K., Sinha, U., Hedman, P., Barron, J.T., Bouaziz, S., Goldman, D.B., Martin-Brualla, R., Seitz, S.M.: Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ToG* (2021) [4](#), [11](#), [12](#), [13](#), [14](#)
53. Peng, S., Yan, Y., Shuai, Q., Bao, H., Zhou, X.: Representing volumetric videos as dynamic mlp maps. In: CVPR (2023) [2](#)
54. Peng, S., Niemeyer, M., Mescheder, L., Pollefeys, M., Geiger, A.: Convolutional occupancy networks. In: ECCV (2020) [3](#), [8](#)
55. Prokudin, S., Ma, Q., Raafat, M., Valentin, J., Tang, S.: Dynamic point fields. In: ICCV (2023) [4](#)
56. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR (2021) [4](#), [12](#), [13](#), [14](#)
57. Qian, S., Kirschstein, T., Schoneveld, L., Davoli, D., Giebenhain, S., Nießner, M.: Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In: CVPR (2024) [5](#)
58. Qian, Z., Wang, S., Mihajlovic, M., Geiger, A., Tang, S.: 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In: CVPR (2024) [5](#)
59. Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., Courville, A.: On the spectral bias of neural networks. In: ICML (2019) [2](#), [7](#)
60. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: ICCV (2021) [4](#)
61. Reiser, C., Szeliski, R., Verbin, D., Srinivasan, P., Mildenhall, B., Geiger, A., Barron, J., Hedman, P.: Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. *ToG* (2023) [4](#)
62. Rückert, D., Franke, L., Stamminger, M.: Adop: Approximate differentiable one-pixel point rendering. *ToG* (2022) [4](#)
63. Saito, S., Schwartz, G., Simon, T., Li, J., Nam, G.: Relightable gaussian codec avatars. In: CVPR (2024) [5](#)
64. Sara Fridovich-Keil and Alex Yu, Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR (2022) [4](#)
65. Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016) [7](#)
66. Shi, R., Wei, X., Wang, C., Su, H.: Zerorf: Fast sparse view 360deg reconstruction with zero pretraining. In: CVPR (2024) [3](#), [7](#), [8](#), [11](#), [12](#)

67. Singh, M., Fuenmayor, E., Hinchy, E.P., Qiao, Y., Murray, N., Devine, D.: Digital twin: Origin to future. *Applied System Innovation* (2021) [1](#)
68. Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretzschmar, H.: Block-nerf: Scalable large scene neural view synthesis. In: *CVPR* (2022) [4](#)
69. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Yifan, W., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., et al.: Advances in neural rendering. In: *Computer Graphics Forum* (2022) [3](#), [4](#), [7](#)
70. Tosi, F., Zhang, Y., Gong, Z., Sandström, E., Mattoccia, S., Oswald, M.R., Poggi, M.: How nerfs and 3d gaussian splatting are reshaping slam: a survey. *arXiv preprint arXiv:2402.13255* (2024) [2](#)
71. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: *CVPR* (2018) [7](#), [8](#)
72. Wang, C., Eckart, B., Lucey, S., Gallo, O.: Neural trajectory fields for dynamic novel view synthesis. *arXiv preprint arXiv:2105.05994* (2021) [4](#), [13](#)
73. Wang, G., Chen, Z., Loy, C.C., Liu, Z.: Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In: *ICCV* (2023) [9](#), [10](#)
74. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* (2004) [6](#)
75. Wiles, O., Gkioxari, G., Szeliski, R., Johnson, J.: Synsin: End-to-end view synthesis from a single image. In: *CVPR* (2020) [4](#)
76. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Xinggang, W.: 4d gaussian splatting for real-time dynamic scene rendering. In: *CVPR* (2024) [3](#), [5](#), [11](#), [12](#), [13](#), [14](#)
77. Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., Sridhar, S.: Neural fields in visual computing and beyond. In: *Computer Graphics Forum*. Wiley Online Library (2022) [1](#), [3](#), [4](#)
78. Xu, Y., Lu, Y., Wen, Z.: OwlII dynamic human mesh sequence dataset. In: *ISO/IEC JTC1/SC29/WG11 m41658, 120th MPEG Meeting* (2017) [8](#), [9](#), [12](#), [14](#)
79. Xu, Y., Chen, B., Li, Z., Zhang, H., Wang, L., Zheng, Z., Liu, Y.: Gaussian head avatar: Ultra high-fidelity head avatar via dynamic gaussians. In: *CVPR* (2024) [5](#)
80. Yan, Z., Li, C., Lee, G.H.: Nerf-ds: Neural radiance fields for dynamic specular objects. In: *CVPR* (2023) [11](#), [12](#)
81. Yang, C., Li, S., Fang, J., Liang, R., Xie, L., Zhang, X., Shen, W., Tian, Q.: Gaussianobject: Just taking four images to get a high-quality 3d object with gaussian splatting. *arXiv preprint arXiv:2402.10259* (2024) [14](#)
82. Yang, Z., Yang, H., Pan, Z., Zhu, X., Zhang, L.: Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In: *ICLR* (2024) [3](#), [5](#), [13](#), [14](#)
83. Yang, Z., Gao, X., Zhou, W., Jiao, S., Zhang, Y., Jin, X.: Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In: *CVPR* (2024) [3](#), [5](#), [11](#), [12](#), [13](#), [14](#)
84. Yariv, L., Hedman, P., Reiser, C., Verbin, D., Srinivasan, P.P., Szeliski, R., Barron, J.T., Mildenhall, B.: BakedSDF: Meshing neural SDFs for real-time view synthesis. In: *SIGGRAPH* (2023) [4](#)
85. Yifan, W., Serena, F., Wu, S., Öztireli, C., Sorkine-Hornung, O.: Differentiable surface splatting for point-based geometry processing. *ToG* (2019) [4](#)
86. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: PlenOctrees for real-time rendering of neural radiance fields. In: *ICCV* (2021) [4](#), [5](#)
87. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: *CVPR* (2021) [2](#), [4](#), [14](#)

88. Yu, H., Julin, J., Milacski, Z.A., Niinuma, K., Jeni, L.A.: Cogs: Controllable gaussian splatting. In: CVPR (2024) [5](#)
89. Yu, Z., Chen, A., Huang, B., Sattler, T., Geiger, A.: Mip-splatting: Alias-free 3d gaussian splatting. In: CVPR (2024) [3](#), [4](#), [9](#), [10](#)
90. Zhang, J., Yang, G., Tulsiani, S., Ramanan, D.: Ners: Neural reflectance surfaces for sparse-view 3d reconstruction in the wild. NeurIPS (2021) [4](#)
91. Zhang, Y., Prokudin, S., Mihajlovic, M., Ma, Q., Tang, S.: Degrees of freedom matter: Inferring dynamics from point trajectories. In: CVPR (2024) [13](#), [14](#)
92. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Ewa volume splatting. In: VIS (2001) [2](#), [4](#), [6](#)
93. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Surface splatting. In: PACMCGIT (2001) [2](#), [4](#)