
FreeSplat: Generalizable 3D Gaussian Splatting Towards Free-View Synthesis of Indoor Scenes

Yunsong Wang Tianxin Huang Hanlin Chen Gim Hee Lee
School of Computing, National University of Singapore
yunsong@comp.nus.edu.sg gimhee.lee@nus.edu.sg
<https://github.com/wangys16/FreeSplat>

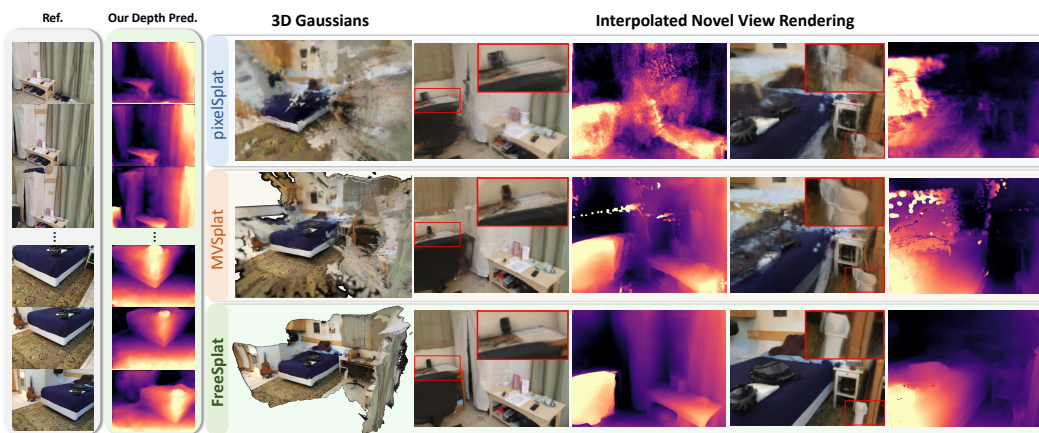


Figure 1: **Comparison between FreeSplat and previous methods.** pixelSplat [1] and MVSplat [2] fail to reconstruct geometrically consistent global 3D Gaussians, while our FreeSplat is proposed to accurately localize 3D Gaussians from long sequence input and support free view synthesis.

Abstract

Empowering 3D Gaussian Splatting with generalization ability is appealing. However, existing generalizable 3D Gaussian Splatting methods are largely confined to narrow-range interpolation between stereo images due to their heavy backbones, thus lacking the ability to accurately localize 3D Gaussian and support free-view synthesis across wide view range. In this paper, we present a novel framework FreeSplat that is capable of reconstructing geometrically consistent 3D scenes from long sequence input towards free-view synthesis. Specifically, we firstly introduce Low-cost Cross-View Aggregation achieved by constructing adaptive cost volumes among nearby views and aggregating features using a multi-scale structure. Subsequently, we present the Pixel-wise Triplet Fusion to eliminate redundancy of 3D Gaussians in overlapping view regions and to aggregate features observed across multiple views. Additionally, we propose a simple but effective free-view training strategy that ensures robust view synthesis across broader view range regardless of the number of views. Our empirical results demonstrate state-of-the-art novel view synthesis performances in both novel view rendered color maps quality and depth maps accuracy across different numbers of input views. We also show that FreeSplat performs inference more efficiently and can effectively reduce redundant Gaussians, offering the possibility of feed-forward large scene reconstruction without depth priors.

1 Introduction

Recent advancements has emerged [3, 4, 5, 6] in reconstructing 3D scenes from multiple viewpoints. Based on ray-marching-based volume rendering, Neural Radiance Fields [3, 7, 8, 9] is capable of learning the implicit 3D geometry and radiance fields without depth information. Nonetheless, computational cost remains to be the inherent bottleneck in ray-marching-based volume rendering, preventing it from real-time rendering. 3D Gaussian Splatting [10, 11, 12, 13] has recently been proposed as an efficient representation for photorealistic reconstruction of 3D scenes from multi-views. The explicit representation of 3D Gaussians are optimized to be densified in the textured regions, and the rasterization-based volume rendering avoids the costly ray marching scheme. Consequently, 3D Gaussian Splatting has achieved real-time rendering of high-quality images from novel views. Nonetheless, vanilla 3D Gaussian Splatting lacks generalizability and requires per-scene optimization.

Several attempts [1, 2, 14, 15, 16, 17] have been made to give 3D Gaussian Splatting generalization ability. Despite showing promising performance, these methods are limited to narrow-range scene-level view interpolation [1, 2, 15] and object-centric synthesis [14, 16]. The primary reason for the limitation is that these existing methods depend on dense view matching across multi-view images with transformers to predict Gaussian primitives, which consequently becomes computationally intractable with longer sequences and thus restricting the supervision of these methods to narrow-range interpolated views. As we show in Figure 4, supervision by narrow-range interpolated views often result in poorly localized 3D Gaussians that can become floaters when rendered from extrapolated views. Additionally, the problem is further aggravated by existing methods typically merging multi-view 3D Gaussians through simple concatenation and thus inevitably lead to noticeable redundancy in overlapping areas (*cf.* Table 2). In view of the above-mentioned problems, it is therefore imperative to design a method that is capable of long sequence reconstruction of global 3D Gaussians, which has the significant potential of supporting real-time rendering from arbitrary poses.

In this paper, we propose FreeSplat tailored for indoor long sequence free view synthesis. Unlike existing methods limited to view interpolation in narrow ranges, our method can effectively reconstruct explicit global 3DGS for novel view synthesis across wide view ranges. Our pipeline consists of Low-cost Cross-View Aggregation and Pixel-wise Triplet Fusion (PTF). In Low-cost Cross-View Aggregation, we introduce efficient CNN-based backbones and adaptive cost volumes formulation among nearby views for low-cost feature extraction and matching, then we leverage a Multi-Scale Feature Aggregation structure to broaden the receptive field of cost volume and predict Depths and Gaussian Triplets. Subsequently, we present Pixel-wise Alignment with progressive Gaussian fusion in PTF to adaptively fuse local Gaussian Triplets from multi-views and avoid Gaussian redundancy in the overlapping regions. Moreover, due to our efficient feature extraction and matching, we propose a Free-View Training (FVT) strategy to disentangle generalizable 3DGS with specific number of views and train the model on long sequences.

The **contributions** of our paper are summarized as follows:

1. We present Low-cost Cross-View Aggregation to predict initial Gaussian triplets, where the low computational cost makes it possible for feature matching between more nearby views and training on long sequence reconstruction;
2. We propose Pixel-wise Triplet Fusion to fuse Gaussian triplets, which can effectively reduce the Gaussian redundancy in the overlapping regions and aggregate multi-view 3D Gaussian latent features;
3. To the best of our knowledge, we are the first to explore generalizable 3DGS for long sequence reconstruction. Extensive experiments on indoor dataset ScanNet [18] and Replica [19] demonstrate our superiority on both image rendering quality and novel view depth rendering accuracy when given different lengths of input views.

2 Related Work

Novel View Synthesis. Traditional attempts in novel view synthesis mainly employed voxel grids [20, 21] or multiplane images [22]. Recently, Neural Radiance Fields (NeRF) [3, 5, 23, 24, 25] have drawn growing interest using ray-marching-based volume rendering to backpropagate image color error to the implicit geometry and radiance fields, such that the 3D geometry can be implicitly learned

to satisfy the multi-view color consistency. Nonetheless, one inherent bottleneck of NeRFs-based method is the computation intensity of ray marching, which requires the costly volume sampling in the implicit fields for each pixel during rendering. To this end, recently 3DGS [10, 26, 27, 11] have attracted increasing attention due to its high efficiency and photorealistic rendering. Instead of relying on MLPs to represent the coordinate-based implicit fields, 3DGS learns an explicit field using a set of 3D Gaussians. They optimize the 3D Gaussians parameters and perform adaptive density control to fit to the given set of images, such that the 3D Gaussians are encouraged to perform densification only in the textured regions and refrain from over-densification. During rendering, 3DGS performs tile-based rasterization to differentially accumulate color images from the explicit 3D Gaussian primitives, which is significantly faster than the ray-marching-based volume rendering and achieves real-time rendering speed.

Generalizable Novel View Synthesis. Another drawback of the traditional NeRF-based and 3DGS-based methods is the requirement of per-scene optimization instead of direct feeding-forward. To this end, there have been a line of work [28, 8, 7, 29] focusing on learning effective priors to predict 3D geometry from given images in a feed forward fashion, where the common practice is to project ray-marching sampled points onto given source views to aggregate multi-view features, conditioning the prediction of the implicit fields on source views instead of point coordinates. Recently, there have also been attempts towards generalizable 3DGS [1, 17, 2, 15, 14]. pixelSplat [1] and GPS-Gaussian [17] propose to predict pixel-aligned 3D Gaussian parameters in feed forward fashion. MVSPat [2] replaces the epipolar line transformer of pixelSplat with a lightweight cost volume to perform more efficient image encoding. GGRt [15] concatenates pixelSplat predicted 3D Gaussians in a sequence of images and simultaneously perform pose optimization. latentSplat [14] encodes 3D Variational Gaussians and leverages a discriminator to help produce more indistinguishable images. Nonetheless, existing methods do not reconstruct the global 3D Gaussians from arbitrary length of inputs, and are limited to view interpolation [1, 2, 17, 15] or object/human-centric scenes [17, 14]. In contrary, in this paper we focus on reconstructing large scenes from arbitrary length of inputs without depth priors, unleashing the potential of generalizable 3DGS for large scene explicit representation.

Indoor Scene Reconstruction. One line of efforts in feed-forward indoor scene reconstruction focuses on extracting 3D mesh using voxel volumes [30, 31, 32] and TSDF-fusion [33], while do not perform photorealistic novel view synthesis. On the other hand, the SLAM-based methods [34, 35, 36] require dense sequence of RGB-D input and per-scene tracking and mapping. Another paradigm of 3D reconstruction [37, 38, 39] learns implicit Signed Distance Fields from RGB input, while demanding intensive per-scene optimization. Another recent work SurfelNeRF [40] learns a feed-forward framework to map a sequence of images to 3D surfels which support photorealistic image rendering, while they rely on external depth estimator or ground truth depth maps. In contrary, we propose an end-to-end model without ground truth depth map input or supervision, enabling accurate 3D Gaussian localization using only photometric losses.

3 Preliminary

Vanilla 3DGS. 3D-GS [10] explicitly represents a 3D scene with a set of Gaussian primitives which are parameterized via a 3D covariance matrix Σ and mean μ :

$$G(\mathbf{p}) = \exp\left(-\frac{1}{2}(\mathbf{p} - \mu)^\top \Sigma^{-1}(\mathbf{p} - \mu)\right), \quad (1)$$

where Σ is decomposed into $\Sigma = \mathbf{R}\mathbf{S}\mathbf{S}^\top\mathbf{R}^\top$ using a scaling matrix \mathbf{S} and a rotation matrix \mathbf{R} to maintain positive semi-definiteness. During rendering, the 3D Gaussian is transformed into the image coordinates with world-to-camera transform matrix \mathbf{W} and projected onto image plane with projection matrix \mathbf{J} , and the 2D covariance matrix Σ' is computed as $\Sigma' = \mathbf{J}\mathbf{W}\Sigma\mathbf{W}^\top\mathbf{J}^\top$. We then obtain a 2D Gaussian G^{2D} with the covariance Σ' in 2D, and the color rendering is computed using point-based alpha-blending on each ray:

$$\mathbf{C}(\mathbf{x}) = \sum_{i \in N} \mathbf{c}_i \alpha_i G_i^{2D}(\mathbf{x}) \prod_{j=1}^{i-1} (1 - \alpha_j G_j^{2D}(\mathbf{x})), \quad (2)$$

where N is the number of Gaussian primitives, α_i is a learnable opacity, and \mathbf{c}_i is view-dependent color defined by spherical harmonics (SH) coefficients \mathbf{s} . The Gaussian parameters are optimized by a photometric loss to minimize the difference between renderings and image observations.

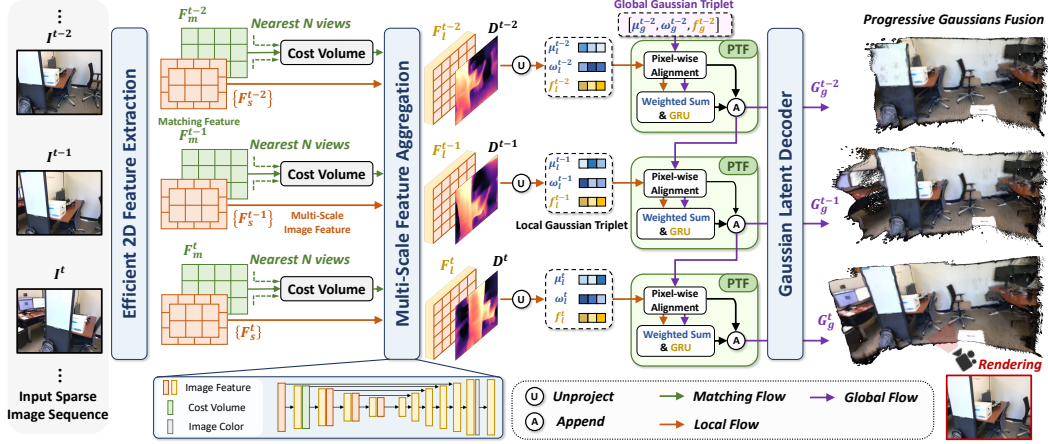


Figure 2: **Framework of FreeSplat.** Given input sparse sequence of images, we construct cost volumes between nearby views and predict depth maps and corresponding feature maps, followed by unprojection to Gaussian triplets with 3D positions. We then propose **Pixel-aligned Triplet Fusion (PTF)** module, where we progressively aggregate and update local/global Gaussian triplets based on pixel-wise alignment. The global Gaussian triplets can be later decoded into Gaussian parameters.

Generalizable 3DGS. Unlike vanilla 3DGS that optimizes per-scene Gaussian primitives, recent generalizable 3DGS [1, 17] predict pixel-aligned Gaussian primitives $\{\Sigma, \alpha, s\}$ and depths d , such that the pixel-aligned Gaussian primitives can be unprojected to 3D coordinates μ . The Gaussian parameters are predicted by 2D encoders, which are optimized by the photometric loss through rendering from novel views. However, existing methods are still limited to view interpolation within narrow view range, which leads to inaccurately localized 3D Gaussians that fail to support large scene reconstruction and view extrapolation (*cf.* Figure 1, 4). To this end, we propose FreeSplat towards global 3D Gaussians reconstruction with accurate localization that supports free-view synthesis.

4 Our Methodology

4.1 Overview

The overview of our method is illustrated in Figure 2. Given a sparse sequence of RGB images, we build cost volumes adaptively between nearby views, and predict depth maps to unproject the 2D feature maps into 3D Gaussian triplets. We then propose the Pixel-aligned Triplet Fusion (PTF) module to progressively align the global with the local Gaussian triplets, such that we can fuse the redundant 3D Gaussians in the latent feature space and aggregate cross-view Gaussian features before decoding. Our method is capable of efficiently exchanging cross-view features through cost volumes, and progressively aggregating per-view 3D Gaussians with cross-view alignment and adaptive fusion.

4.2 Low-cost Cross-View Aggregation

Efficient 2D Feature Extraction. Given a sparse sequence of posed images $\{I^t\}_{t=1}^T$, we first feed them into a shared 2D backbone to extract multi-scale embeddings F_e^t and matching feature F_m^t . Unlike [1, 2] which rely on patch-wise transformer-based backbones [41, 42] that can lead to quadratically expensive computations, we leverage pure CNN-based backbones [43, 44] for 2D feature extraction for efficient performance on higher resolution inputs.

Adaptive Cost Volume Formulation. To explicitly integrate camera pose information given arbitrary length of input images, we propose to adaptively build cost volumes between nearby views. For current view I^t with pose P^t and matching feature $F_m^t \in \mathbb{R}^{C_m \times \frac{H}{4} \times \frac{W}{4}}$, we adaptively select its N nearby views $\{I^{t_n}\}_{n=1}^N$ with poses $\{P^{t_n}\}_{n=1}^N$ based on pose proximity, and construct cost volume via plane sweep stereo [45, 46]. Specifically, we define a set of K virtual depth planes $\{d_k\}_{k=1}^K$ that are uniformly spaced within $[d_{near}, d_{far}]$, and warp the nearby view features to each depth plane d_k

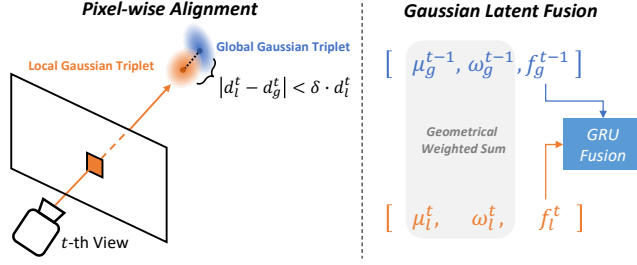


Figure 3: **Visual illustration of PTF.** The PTF incrementally projects current global Gaussians to input views and computes their pixel-wise distance with local Gaussians. Nearby local Gaussians are then fused using a lightweight Gate Recurrent Unit (GRU) network [49].

of current view:

$$\tilde{\mathbf{F}}_m^{t_n, k} = \text{Trans}(\mathbf{P}^{t_n}, \mathbf{P}^t) \mathbf{F}_m^{t_n}, \quad (3)$$

where $\text{Trans}(\mathbf{P}^{t_n}, \mathbf{P}^t)$ is the transformation matrix from view t_n to t . The cost volume $\mathbf{F}_{cv}^t \in \mathbb{R}^{K \times \frac{H}{4} \times \frac{W}{4}}$ is then defined as:

$$\mathbf{F}_{cv}^t(k) = f_\theta \left(\left(\frac{1}{N} \sum_{n=1}^N \cos(\mathbf{F}_m^t, \tilde{\mathbf{F}}_m^{t_n, k}) \right) \oplus \left(\frac{1}{N} \sum_{n=1}^N \tilde{\mathbf{F}}_m^{t_n, k} \right) \right), \quad (4)$$

where $\mathbf{F}_{cv}^t[k]$ is the k -th dimension of \mathbf{F}_{cv}^t , $\cos(\cdot)$ is the cosine similarity, \oplus is feature-wise concatenation, and $f_\theta(\cdot)$ is a 1×1 CNN mapping to dimension of 1.

Multi-Scale Feature Aggregation. The embedding of the cost volume plays a significant part to accurately localize the 3D Gaussians (*cf.* Table 5). To this end, inspired by previous depth estimation methods [47, 33], we design a multi-scale encoder-decoder structure, such that to fuse multi-scale image features with the cost volume and propagate the cost volume information to broader receptive fields. Specifically, the multi-scale encoder takes in \mathbf{F}_{cv}^t and the output is concatenated with $\{\mathbf{F}_s^t\}$ before sending into a UNet++ [48]-like decoder to upsample to full resolution and predict a depth candidates map $\mathbf{D}_c^t \in \mathbb{R}^{K \times H \times W}$, and Gaussian triplet map $\mathbf{F}_l^t \in \mathbb{R}^{C \times H \times W}$. We then predict the depth map through soft-argmax to bound the depth prediction between near and far:

$$\mathbf{D}^t = \sum_{k=1}^K \text{softmax}(\mathbf{D}_c^t)_k \cdot d_k. \quad (5)$$

Finally, the pixel-aligned Gaussian triplet map \mathbf{F}_l^t is unprojected to 3D Gaussian triplet $\{\boldsymbol{\mu}_l^t, \boldsymbol{\omega}_l^t, \mathbf{f}_l^t\}$, where $\boldsymbol{\mu}_l^t \in \mathbb{R}^{3 \times HW}$ are the Gaussian centers, $\boldsymbol{\omega}_l^t \in \mathbb{R}^{1 \times HW}$ are weights between (0, 1), and $\mathbf{f}_l^t \in \mathbb{R}^{(C-1) \times HW}$ are Gaussian triplet features.

4.3 Pixel-wise Triplet Fusion

One limitation of previous generalizable 3DGS methods is the redundancy of Gaussians. Since we need multi-view observations to predict accurately localized 3D Gaussians in indoor scenes, the pixel-aligned Gaussians become redundant in frequently observed regions. Furthermore, previous methods integrate multi-view Gaussians of the same region simply through their opacities, leading to suboptimal performance due to lack of post aggregation (*cf.* Table 5). Consequently, inspired by previous methods [31, 40], we propose the Pixel-wise Triplet Fusion (PTF) module which can significantly remove redundant Gaussians in the overlapping regions and explicitly aggregate multi-view observation features in the latent space. We align the per-view local Gaussians with global ones using Pixel-wise Alignment to select the redundant 3D Gaussian Triplets, and progressively fuse the local Gaussians into the global ones.

Pixel-wise Alignment. Given the Gaussian triplets $\{\boldsymbol{\mu}_l^t, \mathbf{f}_l^t\}_{t=1}^T$, we start from $t = 1$ where the global Gaussians latent is empty. In the t -th step, we first project the global Gaussian triplet centers $\boldsymbol{\mu}_g^{t-1} \in \mathbb{R}^{3 \times M}$ onto the t -th view:

$$\mathbf{p}_g^t := \{\mathbf{x}_g^t, \mathbf{y}_g^t, \mathbf{d}_g^t\} = \mathbf{P}^t \boldsymbol{\mu}_g^{t-1}, \quad (6)$$

where $[\mathbf{x}_g^t, \mathbf{y}_g^t, \mathbf{d}_g^t] \in \mathbb{R}^{3 \times M}$ are the projected 2D coordinates and corresponding depths. We then correspond the local Gaussian triplets with the pixel-wise nearest projections within a threshold. Specifically, for the i -th local Gaussian with 2D coordinate $[\mathbf{x}_l^t(i), \mathbf{y}_l^t(i)]$ and depth $d_l^t(j)$, we first find its intra-pixel global projection set \mathcal{S}_i :

$$\mathcal{S}_i^t := \{j \mid [\mathbf{x}_g^t(j)] = \mathbf{x}_l^t(i), [\mathbf{y}_g^t(j)] = \mathbf{y}_l^t(i)\}, \quad (7)$$

where $[\cdot]$ is the rounding operator. Subsequently, we search for valid correspondence with minimum depth difference under a threshold:

$$m_i = \begin{cases} \arg \min_{j \in \mathcal{S}_i^t} \mathbf{d}_g^t(j) & \text{if } |\mathbf{d}_l^t(j) - \min_{j \in \mathcal{S}_i^t} \mathbf{d}_g^t(j)| < \delta \cdot \mathbf{d}_l^t(j) \\ \emptyset & \text{otherwise} \end{cases}, \quad (8)$$

where δ is a ratio threshold. We define the valid correspondence set as:

$$\mathcal{F}^t := \{(i, m_i) \mid i = 1, \dots, HW; m_i \neq \emptyset\}. \quad (9)$$

Gaussian Triplet Fusion. After the pixel-wise alignment, we remove the redundant 3D Gaussians through merging the validly aligned triplet pairs. Given a pair $(i, m_i) \in \mathcal{F}^t$, we compute the weighted sum of their center coordinates and sum their weights to restrict the 3D Gaussian centers to lie between the triplet pair:

$$\boldsymbol{\mu}_g^t(m_i) = \frac{\omega_l^t(i) \boldsymbol{\mu}_l^t(i) + \omega_g^{t-1}(m_i) \boldsymbol{\mu}_g^{t-1}(m_i)}{\omega_l^t(i) + \omega_g^t(m_i)}, \quad \text{where } \omega_g^t(m_i) = \omega_l^t(i) + \omega_g^{t-1}(m_i). \quad (10)$$

We then aggregate the aligned local and global Gaussian latent features through a lightweight GRU network:

$$\mathbf{f}_g^t(m_i) = \text{GRU}(\mathbf{f}_l^t(i), \mathbf{f}_g^{t-1}(m_i)), \quad (11)$$

and then append with the other unaligned local Gaussian triplets.

Gaussian primitives decoding. After the Pixel-wise Triplet Fusion, we can decode the global Gaussian triplets into Gaussian primitives:

$$\boldsymbol{\Sigma}, \boldsymbol{\alpha}, \mathbf{s} = \text{MLP}_d(\mathbf{f}_g^T) \quad (12)$$

and Gaussian centers $\boldsymbol{\mu} = \boldsymbol{\mu}_g^\top$. Our proposed fusion method can incrementally integrate the Gaussians with geometrical constraints and learnable GRU network for feature update. Consequently, our fusion method is capable of significantly removing redundant Gaussians and perform post feature aggregation across multiple views, and can be trained with the other framework components end-to-end with eligible computation overhead.

4.4 Training

Loss Functions. After predicting the 3D Gaussian primitives, we render from novel views following the rendering equations in Eq. (2). Similar to pixelSplat [1] and MVSplat [2], we train our framework using only photometric losses, *i.e.* a combination of MSE loss and LPIPS [50] loss, with weights of 1 and 0.05 following [1, 2].

Free-View Training. We propose a Free-View Training (FVT) strategy to add more geometrical constraints on the localization of 3D Gaussians, and to disentangle the performance of generalizable 3DGS with specific number of input views. To this end, we randomly sample T number of context views (in experiments we set T between 2 and 8), and supervise the image renderings in the broader view interpolations. The long sequence training is made feasible due to our efficient feature extraction and aggregation. We empirically find that FVT significantly contributes to depth estimation from novel views (*cf.* Table 3, 4).

5 Experiments

5.1 Experimental Settings

Datasets. We leverage the real-world indoor dataset ScanNet [18] for training. ScanNet is a large RGB-D dataset containing 1,513 indoor scenes with camera poses, and we follow [51, 40] to use

Table 1: **Generalizable Novel View Interpolation results on ScanNet [18]**. FreeSplat-*fv* is trained with our FVT strategy, and the other methods are all trained on specific number of views to form a complete comparison. Time(s) indicates the total time of encoding input images and rendering one image.

Method	2 views					3 views				
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time(s) \downarrow	#GS(k)	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Time(s) \downarrow	#GS(k)
NeuRay [9]	25.65	0.840	0.264	3.103	-	25.47	0.843	0.264	4.278	-
pixelSplat [1]	26.03	0.784	0.265	0.289	1180	25.76	0.782	0.270	0.272	1769
MVSplat [2]	27.27	0.822	0.221	0.117	393	26.68	0.814	0.235	0.192	590
FreeSplat-spec	28.08	0.837	0.211	0.103	278	27.45	0.829	0.222	0.121	382
FreeSplat-fv	27.67	0.830	0.215	0.104	279	27.34	0.826	0.226	0.122	390

Table 2: **Long Sequence (10 views) Explicit Reconstruction results on ScanNet**. The results of pixelSplat, MVSplat and FreeSplat-spec are given using their 3-views version.

Method	Time(s) \downarrow	#GS(k)	View Interpolation			View Extrapolation		
			PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
pixelSplat [1]	0.948	5898	21.26	0.714	0.396	20.70	0.687	0.429
MVSplat [2]	1.178	1966	22.78	0.754	0.335	21.60	0.729	0.365
FreeSplat-3views	0.599	882	25.15	0.800	0.278	23.78	0.774	0.309
FreeSplat-fv	0.596	899	25.90	0.808	0.252	24.64	0.786	0.277

100 scenes for training and 8 scenes for testing. To evaluate the generalization ability of our model, we further perform zero-shot evaluation on the synthetic indoor dataset Replica [19], for which we follow [52] to select 8 scenes for testing.

Implementation Details. Our FreeSplat is trained end-to-end using Adam [53] optimizer with an initial learning rate of $1e-4$ and cosine decay following [2]. Due to the large GPU requirements of [1, 2] given high-resolution images, all input images are resized to 384×512 and batch size is set to 1, to form a fair comparison between different methods. We mainly compare with previous generalizable 3DGS methods in 2, 3, 10 reference view settings, where the distance between input views is fixed, thus evaluating the models’ performance under different view ranges. For 10 views setting, we also choose target views that are beyond the given sequence of reference views to evaluate the view extrapolation results.

5.2 Results on ScanNet

View Interpolation Results. On ScanNet, we evaluate the generalizable novel view interpolation results given 2 and 3 reference views as shown in Table 1. Comparing to pixelSplat and MVSplat, our FreeSplat-spec consistently improves rendering quality and efficiency on 2-views setting and 3-views setting. Although slightly underperforming on SSIM comparing to NeuRay [9], we show significant improvements on PSNR and LPIPS over NeuRay and $300\times$ faster inference speed. Moreover, our FreeSplat-fv consistently offers competitive results given arbitrary number of views, and performs more similarly as FreeSplat-spec when number of input views increases.

Table 3: **Novel View Depth Rendering results on ScanNet**. \dagger : 10-views results of pixelSplat, MVSplat and FreeSplat-spec are given using their 3-views version.

Method	2 views				3 views			10 views \dagger		
	Abs Diff \downarrow	Abs Rel \downarrow	$\delta < 1.25 \uparrow$	Abs Diff \downarrow	Abs Rel \downarrow	$\delta < 1.25 \uparrow$	Abs Diff \downarrow	Abs Rel \downarrow	$\delta < 1.25 \uparrow$	
NeuRay [9]	0.358	0.200	0.755	0.231	0.117	0.873	0.202	0.108	0.875	
pixelSplat [1]	1.205	0.745	0.472	0.698	0.479	0.836	0.970	0.621	0.647	
MVSplat [2]	0.192	0.106	0.912	0.164	0.079	0.929	0.142	0.080	0.914	
FreeSplat-spec	0.157	0.086	0.919	0.161	0.077	0.930	0.120	0.070	0.945	
FreeSplat-fv	0.153	0.085	0.923	0.162	0.077	0.928	0.097	0.059	0.961	

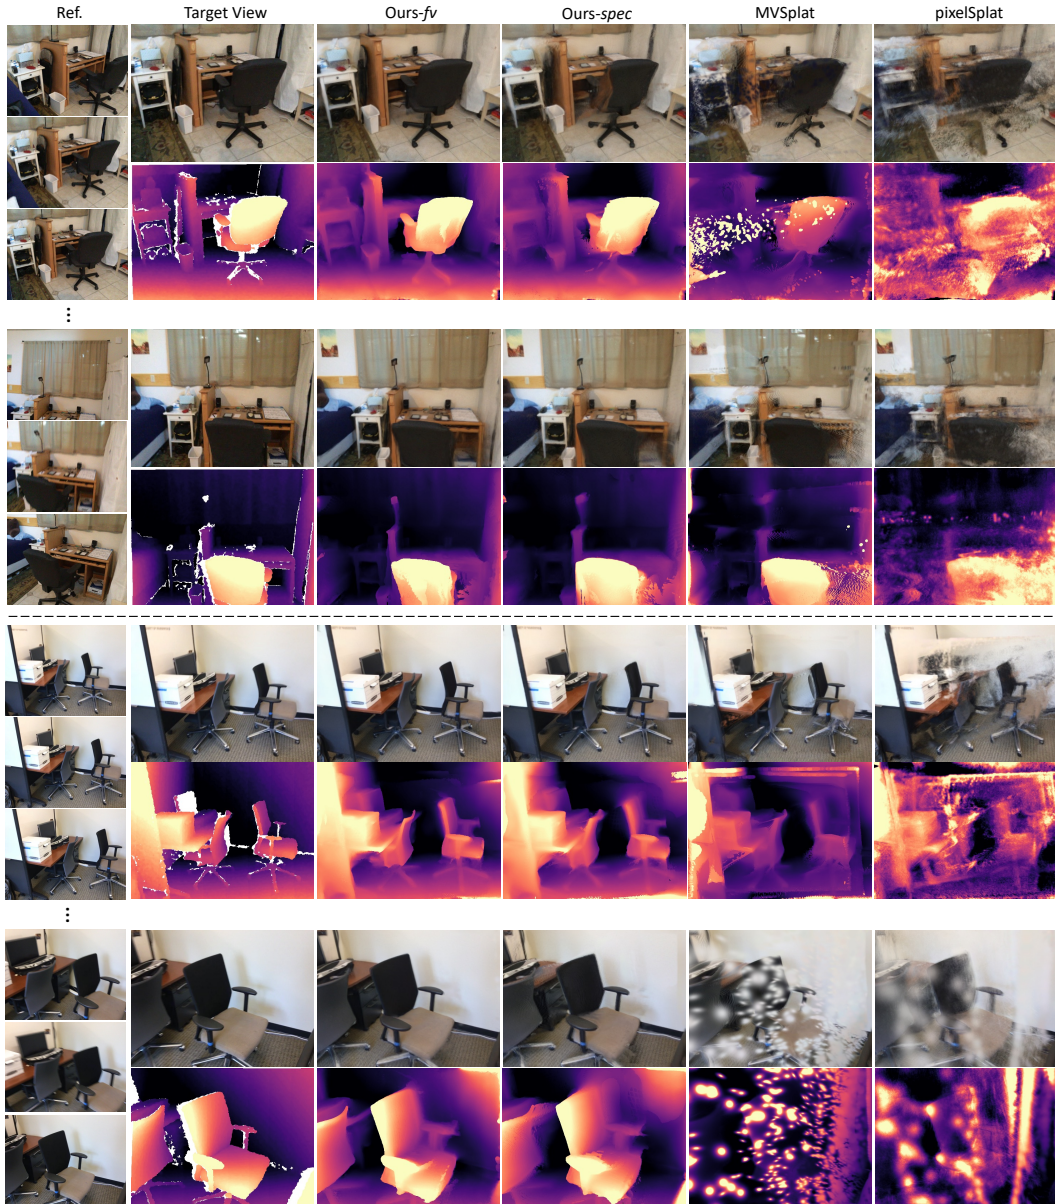


Figure 4: **Qualitative Results of Long Sequence Explicit Reconstruction.** For each sequence, the first two rows are view interpolation results, and the last two rows are view extrapolation results.

Long Sequence Results. As shown in Table 2, we further evaluate the long sequence results where we sample reference views with length of 10, and compare both view interpolation and extrapolation results. The results reveal that generalizable 3DGS methods underperform when given long sequence input images, which is due to the complicated camera trajectories in ScanNet, and the inaccuracy of 3D Gaussian localization that leads to errors when observed from wide view ranges. Our FreeSplat-3views significantly outperforms pixelSplat and MVSplat on view interpolation and view extrapolation results. Through our proposed FVT that can be easily plugged into our model due to our low requirement on GPU, our FreeSplat-fv consistently outperforms our 3-views version. Our PTF module can also reduce the number of Gaussians by around 55.0%, which becomes indispensable in long sequence reconstruction due to the pixel-wise unprojection nature of generalizable 3DGS. The qualitative results are shown in Figure 4, which clearly reveal that FreeSplat-spec outperforms MVSplat and pixelSplat in localizing 3D Gaussian and preserving fine-grained details, and FreeSplat-fv further improves on localizing and fusing multi-view Gaussians.

Table 4: **Zero-Shot Transfer Results on Replica [19].**

Method	3 Views					10 Views				
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$\delta < 1.25 \uparrow$	#GS(k)	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$\delta < 1.25 \uparrow$	#GS(k)
pixelSplat [1]	26.24	0.829	0.229	0.576	1769	19.23	0.719	0.414	0.375	5898
MVSplat [2]	26.16	0.840	0.173	0.670	590	18.66	0.717	0.360	0.565	1966
FreeSplat-spec	26.98	0.848	0.171	0.682	423	21.11	0.762	0.312	0.720	1342
FreeSplat-fv	26.64	0.843	0.184	0.682	421	21.95	0.777	0.290	0.742	1346

Table 5: **Ablation on ScanNet.** CV: Cost Volume, PTF: Pixel-wise Triplet Fusion, FVT: Free-View Training.

CV	PTF	FVT	3 views					10 views				
			PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$\delta < 1.25 \uparrow$	$\delta < 1.10 \uparrow$	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	$\delta < 1.25 \uparrow$	$\delta < 1.10 \uparrow$
✓			27.12	0.825	0.224	0.925	0.762	24.23	0.792	0.277	0.942	0.804
	✓		22.10	0.696	0.359	0.639	0.311	17.94	0.607	0.487	0.543	0.216
		✓	27.45	0.829	0.222	0.930	0.773	25.15	0.800	0.278	0.945	0.823
✓	✓		26.41	0.806	0.232	0.919	0.746	25.40	0.799	0.252	0.950	0.831
✓	✓	✓	27.34	0.826	0.226	0.928	0.764	25.90	0.808	0.252	0.961	0.858




Figure 5: **Qualitative Ablation Study.** The first and second row use input view lengths of 3 and 10.

Novel View Depth Estimation Results. We also investigate the correctness of 3D Gaussian localization of different methods through comparing their depth rendering results. We report the Absolute Difference (Abs. Diff), Relative Difference (Rel. Diff), and threshold tolerance $\delta < 1.25$ results from novel views in Table 3. We find that FreeSplat consistently outperforms pixelSplat and MVSplat in predicting accurately localized 3D Gaussians, where FreeSplat-fv reaches 94.9% of $\delta < 1.25$, enabling accurate unsupervised depth estimation on novel views. The improved depth estimation accuracy of FreeSplat-fv highlights the importance of depth estimation in supporting free-view synthesis across broader view range.

5.3 Zero-Shot Transfer Results on Replica

We further evaluate the zero-shot transfer results through testing on Replica dataset, with results in Table 4. Our view interpolation and novel view depth estimation results still outperforms existing methods. The long sequence results degrade due to inaccurate depth estimation and domain gap, indicating potential future work in further improving the depth estimation in zero-shot transferring.

5.4 Ablation Study

We conduct a detailed ablation study as shown in Table 5 and Figure 5. The results indicate that: 1) cost volume is essential in accurately localizing 3D Gaussians; 2) our proposed PTF module can consistently contribute to rendering quality and depth estimation results. The PTF module learns to incrementally fuse multi-view 3D Gaussians and contributes significantly when varying number of input views, and serves as a multi-view localization regularization that helps unsupervised depth estimation; 3) Our FVT module excels in long sequence reconstruction quality as well as novel view

depth rendering results, which provides stricter constraints on 3D Gaussian localization and can be seamlessly combined with the PTF module to fit to varying length of input views.

6 Conclusion

In this study, we introduced FreeSplat, a generalizable 3DGS model that is tailored to accommodate an arbitrary number of input views and perform free-view synthesis using the global 3D Gaussians. We developed a Low-cost Cross-View Aggregation pipeline that enhances the model’s ability to efficiently process long input sequences, thus incorporating stricter geometry constraints. Additionally, we have devised a Pixel-wise Triplet Fusion module that effectively reduces redundant pixel-aligned 3D Gaussians in overlapping regions and merges multi-view Gaussian latent features. FreeSplat consistently improves the fidelity of novel view renderings in terms of both color image quality and depth map accuracy, facilitating feed-forward global Gaussians reconstruction without depth priors.

7 Acknowledgement

This work is supported by the Agency for Science, Technology and Research (A*STAR) under its MTC Programmatic Funds (Grant No. M23L7b0021).

References

- [1] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. *arXiv preprint arXiv:2312.12337*, 2023.
- [2] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. *arXiv preprint arXiv:2403.14627*, 2024.
- [3] Ben Mildenhall, Pratul Srinivasan, Matthew Tancik, Jonathan Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020.
- [4] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021.
- [5] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022.
- [6] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [7] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021.
- [8] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021.
- [9] Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou, and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7824–7833, 2022.
- [10] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023.

- [11] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. *arXiv preprint arXiv:2311.16493*, 2023.
- [12] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898*, 2024.
- [13] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023.
- [14] Christopher Wewer, Kevin Raj, Eddy Ilg, Bernt Schiele, and Jan Eric Lenssen. latentsplat: Autoencoding variational gaussians for fast generalizable 3d reconstruction. *arXiv preprint arXiv:2403.16292*, 2024.
- [15] Hao Li, Yuanyuan Gao, Dingwen Zhang, Chenming Wu, Yalun Dai, Chen Zhao, Haocheng Feng, Errui Ding, Jingdong Wang, and Junwei Han. Ggrt: Towards generalizable 3d gaussians without pose priors in real-time. *arXiv preprint arXiv:2403.10147*, 2024.
- [16] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. *arXiv preprint arXiv:2312.09147*, 2023.
- [17] Shunyuan Zheng, Boyao Zhou, Ruizhi Shao, Boning Liu, Shengping Zhang, Liqiang Nie, and Yebin Liu. Gps-gaussian: Generalizable pixel-wise 3d gaussian splatting for real-time human novel view synthesis. *arXiv preprint arXiv:2312.02155*, 2023.
- [18] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.
- [19] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.
- [20] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019.
- [21] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhofer. Deepvoxels: Learning persistent 3d feature embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2437–2446, 2019.
- [22] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*, 2018.
- [23] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021.
- [24] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022.
- [25] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Zip-nerf: Anti-aliased grid-based neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19697–19705, 2023.
- [26] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023.
- [27] GuanJun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. *arXiv preprint arXiv:2310.08528*, 2023.
- [28] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021.

- [29] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022.
- [30] 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.
- [31] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. Neuralrecon: Real-time coherent 3d reconstruction from monocular video. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15598–15607, 2021.
- [32] Noah Stier, Alexander Rich, Pradeep Sen, and Tobias Höllerer. Vortex: Volumetric 3d reconstruction with transformers for voxelwise view selection and fusion. In *2021 International Conference on 3D Vision (3DV)*, pages 320–330. IEEE, 2021.
- [33] Mohamed Sayed, John Gibson, Jamie Watson, Victor Prisacariu, Michael Firman, and Clément Godard. Simplexcon: 3d reconstruction without 3d convolutions. In *European Conference on Computer Vision*, pages 1–19. Springer, 2022.
- [34] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022.
- [35] Chi Yan, Delin Qu, Dong Wang, Dan Xu, Zhigang Wang, Bin Zhao, and Xuelong Li. Gs-slam: Dense visual slam with 3d gaussian splatting. *arXiv preprint arXiv:2311.11700*, 2023.
- [36] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. *arXiv preprint arXiv:2312.02126*, 2023.
- [37] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021.
- [38] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35:25018–25032, 2022.
- [39] Zhaoshuo Li, Thomas Müller, Alex Evans, Russell H Taylor, Mathias Unberath, Ming-Yu Liu, and Chen-Hsuan Lin. Neuralangelo: High-fidelity neural surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8456–8465, 2023.
- [40] Yiming Gao, Yan-Pei Cao, and Ying Shan. Surfelfnerf: Neural surfel radiance fields for online photorealistic reconstruction of indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 108–118, 2023.
- [41] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arxiv 2020. arXiv preprint arXiv:2010.11929*, 2010.
- [42] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [43] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [44] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [45] Robert T Collins. A space-sweep approach to true multi-image matching. In *Proceedings CVPR IEEE computer society conference on computer vision and pattern recognition*, pages 358–363. Ieee, 1996.
- [46] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. Dpsnet: End-to-end deep plane sweep stereo. *arXiv preprint arXiv:1905.00538*, 2019.

- [47] Arda Duzceker, Silvano Galliani, Christoph Vogel, Pablo Speciale, Mihai Dusmanu, and Marc Pollefeys. Deepvideomvs: Multi-view stereo on video with recurrent spatio-temporal fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15324–15333, 2021.
- [48] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018.
- [49] Rahul Dey and Fathi M Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th international midwest symposium on circuits and systems (MWSCAS)*, pages 1597–1600. IEEE, 2017.
- [50] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [51] Xiaoshuai Zhang, Sai Bi, Kalyan Sunkavalli, Hao Su, and Zexiang Xu. Nerfusion: Fusing radiance fields for large-scale scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5449–5458, 2022.
- [52] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J. Davison. In-place scene labelling and understanding with implicit scene representation. 2021.
- [53] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

A Appendix / supplemental material

A.1 Experimental Environment

We conduct all the experiments on single NVIDIA RTX A6000 GPU. The experimental environment is PyTorch 2.1.2 and CUDA 12.2.

A.2 Additional Implementation Details

We set the number of virtual depth planes $K = 128$, matching feature dimension $C_m = 64$, and $d_{near} = 0.5, d_{far} = 15.0$ for cost volume formulation, and set $\delta = 0.05$ in Eq.(8) for pixel-wise alignment. To train the 3-views version of pixelSplat [1] on a single NVIDIA RTX A6000 GPU, we change their ViT patch size from 8×8 to 16×16 . During inference on the 10 views setting, the epipolar line sampling in pixelSplat and the cross-view attention in MVSpLat [2] are performed between nearby views similarly as ours to save GPU requirements and form a fair comparison. For the free-view version of FreeSplat we set the number of nearby views as $N = 4$ for training. For the testing of long sequence explicit reconstruction, cost volumes are formed between nearby 8 views.

A.3 Additional Experiments

Table 6: **Comparison on computational cost and whole scene reconstruction (30 input views).** We report the required GPU for Train / Test, the Encoding Time, the rendering FPS, and PSNR of novel views. - denotes that we are not able to run pixelSplat inference using 30 input views due to its increasing GPU requirement.

Method	GPU (GB)	Time (s)↓	FPS↑	PSNR↑
pixelSplat-3views	44.9 / -	-	-	-
MVSpLat-3views	34.8 / 44.0	3.004	39	17.57
FreeSplat-3views	16.9 / 21.0	1.191	57	21.33
FreeSplat-fv w/o PTF	42.2 / 21.0	1.006	39	21.82
FreeSplat-fv	42.2 / 21.0	1.205	72	22.32

Computational Cost. As shown in Table 6, we compare the required GPU memory for training and testing, the encoding time, rendering FPS, and PSNR for whole scene reconstruction. pixelSplat-3views and MVSpLat-3views already consume 30 50 GB GPU memory for training due to their quadratically increasing GPU memory requirement with respect to the image resolution / sequence length. Therefore, it becomes infeasible to extend their methods to higher resolution inputs or longer sequence training. In comparison, our low-cost framework design enable us to effectively train on long sequence inputs while requiring lesser GPU memory compared to the 3 views version of existing methods. Furthermore, our proposed PTF module can effectively reduce redundant 3D Gaussians, improving rendering speed from 39 to 72 FPS . This becomes increasingly important when reconstructing larger scenes since generalizable 3DGS normally perform pixel-wise unprojection, which can easily result in redundancy in the overlapping regions.

Table 7: **Results on RE10K and ACID with 2 input views.** We train our model on RE10K, and report its results on RE10K and ACID. * denotes that our model is trained on our previously downloaded 9,266 scenes instead of 11,075 scenes used by baselines.

Method	RE10K - 2 Views			ACID - 2 Views		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
pixelSplat	25.89	0.858	0.142	27.64	0.830	0.160
MVSpLat	26.39	0.869	0.128	28.15	0.841	0.147
Ours*	26.41	0.871	0.132	27.94	0.838	0.157

Experiments on RE10K and ACID. To further evaluate our model’s generalization ability across diverse domains, we train our model on RE10K using 2-View setting and 5-View setting respectively. The results are shown in Table 7, 8 and Figure 6. Note that for the 5-View setting inference, we sample input views with random intervals between 25 and 45 due to the limited sequence lengths in RE10K and ACID. In the 2-View setting, we perform better than pixelSplat and on par as MVSpLat on both datasets. In the 5-View setting, we outperform both baselines by a clear margin. We analyze the main causes of the above results as follows:

Table 8: **Results on RE10K and ACID with 5 input views.** The models are trained on RE10K with 5 input views.

Method	RE10K - 5 Views			ACID - 5 Views		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
pixelSplat	24.78	0.850	0.150	26.84	0.833	0.173
MVSplat	25.38	0.866	0.132	27.81	0.863	0.134
Ours	25.95	0.873	0.128	28.35	0.870	0.130



Figure 6: **Qualitative Results on RE10K and ACID.** We visualize the 2-Views results on RE10K and 5-Views results on ACID.

In the 2-view comparison experiments with the baselines, the image interval between the given stereo images were set to be large. On average, the interval between image stereo is 66 in RE10K and 74 in ACID, which is much larger than our indoor datasets setting (20 for ScanNet and 10 for Replica). Such large interval can result in minimum view overlap between the image stereo, which means that our cost volume can be much sparser and multi-view information aggregation is weakened. In contrast, MVSplat uses a cross-view attention that aggregates multi-view features through a sliding window which does not leverage camera poses. pixelSplat uses a heavy 2D backbone that can potentially become stronger monocular depth estimator. In our 5-view setting, we outperform both baselines by clear margins. This is partially due to the smaller image interval and larger view overlap between nearby views. As a result, our cost volume can effectively aggregate multi-view information, and our PTF module can perform point-level fusion and remove those redundant 3D Gaussians.

Therefore, our model is not specifically designed for highly sparse view inputs, but it is designed as a low-cost model that can easily take in much longer sequences of higher-resolution inputs, that is suitable for indoor scene reconstruction. Comparing to RE10K and ACID, real-world indoor scene sequences usually contain more complicated camera rotations and translations, which results in the requirement of more dense observations to reconstruct the 3D scenes with high completeness and accurate geometry. Consequently, our model is targeting the fast indoor scene reconstruction with keyframe inputs, which contain long sequences of high-resolution images, while existing works struggle to extend to such setting as evaluated in our main paper.

Comparison with SurfelNeRF. We further compare with SurfelNeRF as shown in Table 9 and Figure 7. We evaluate on the same novel views as theirs, sampling input views along their input sequences with an interval of 20 between nearby views. Note that the number of input views changes when the input length changes, while our FreeSplat-fv can seamlessly conduct inference with arbitrary numbers of inputs. Our method performs significantly better than SurfelNeRF in both rendering quality and efficiency. Our end-to-end framework jointly learns depths and 3DGS using an MVS-based backbone, while SurfelNeRF relies on depths and does not aggregate multi-view features to assist their surfel feature prediction.

Table 9: **Comparison with SurfNeRF.** We compare with SurfNeRF on the same sequences as their test set. * denotes the rendering speed is reported from SurfNeRF.

Method	Time (s)	FPS \uparrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
SurfNeRF	3.242	5*	24.20	0.694	0.477
FreeSplat-fv	0.302	224	27.06	0.818	0.223




Figure 7: **Qualitative Comparison with SurfNeRF.**

A.4 Additional Qualitative Results

2 and 3-View Interpolation Results. The qualitative results are shown in Figure 8, where FreeSplat more precisely localizes 3D Gaussians and captures more fine-grained details comparing to previous methods. FreeSplat can also localize 3D Gaussians more accurately and renders precise depth maps, supporting high-quality rendering from broader view range (*cf.* FreeSplat-*spec* results in Figure 4).

Results on Replica. We show the qualitative results on Replica in Figure 10, where our superiority over MV-Splat and pixelSplat remains. The results indicate the generalization ability of FreeSplat across indoor datasets for the view interpolation task.

Results of Whole Scene Reconstruction. We also show qualitative results of our whole scene reconstruction in Figure 9. Despite the long input sequence (~ 40 images) covering the whole scene, FreeSplat can still perform efficient feed-forward in ~ 1 s on single NVIDIA RTX A6000, and can render high-quality images and accurate depth maps from novel views. On the other hand, it is still difficult to accurately predict depth of textureless (*e.g.* wall) and specular (*e.g.* light reflection on the floor) regions. However, we hope our work provides an initial step towards accurate geometry reconstruction without ground truth depth priors.

A.5 Limitations

Although our approach excels in novel view rendering depth estimation and support arbitrary number of input views, the GPU requirement becomes expensive (> 40 GB) when inputting extremely long image sequence (> 50). On the other hand, due to our unsupervised scheme of depth estimation, there is still a gap between our 3D reconstruction accuracy and the state-of-the-art methods with 3D supervision [33, 32] or RGB-D inputs [36, 35] (*e.g.* as shown in Figure 9, the textureless and specular regions). Our main focus is to explore the feed-forward indoor scene photorealistic reconstruction purely based on 2D supervision.

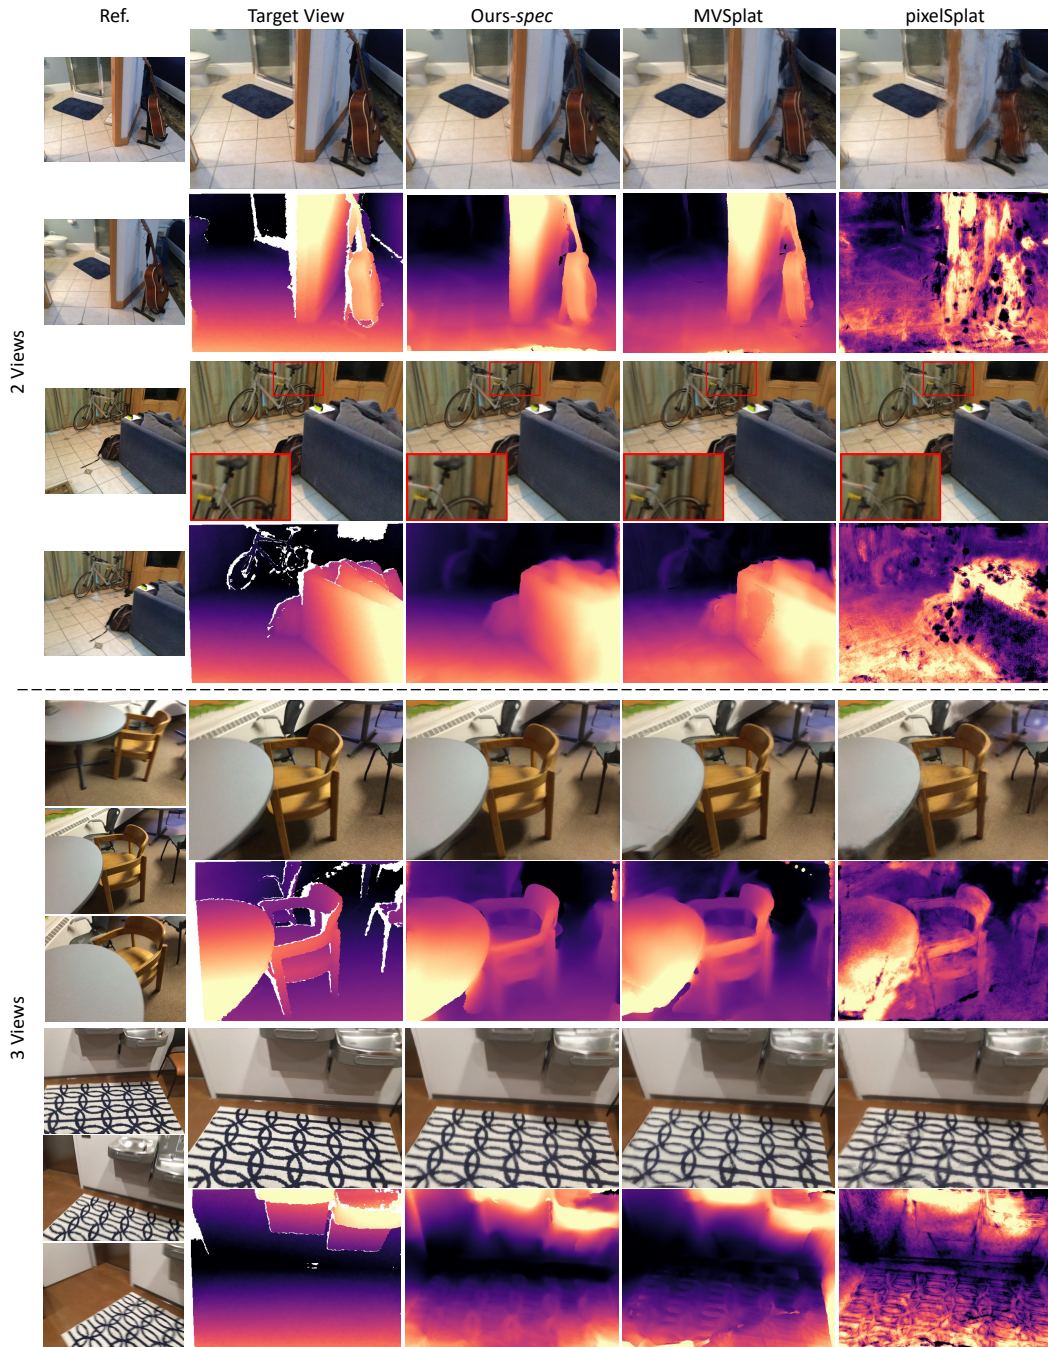


Figure 8: **Qualitative Results given 2 and 3 reference views.** We show the rendered color images (first row) and depth maps (second row) for each batch of reference views.




Figure 9: Qualitative Results on Replica.




Figure 10: Qualitative Results of whole scene reconstruction.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We mainly focus on generalizable 3DGS for indoor scene reconstruction, where we support arbitrary number of inputs through adaptive cost volume (Sec. 4.2) and gaussian fusion (Sec. 4.3).

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Discussed in Sec. A.5

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We do not propose new theory in this paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We include detailed description of our framework in Sec. 4 and implementation details in Sec. 5.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: Our code will be released upon paper acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The detailed descriptions of the experiment settings are illustrated in Sec. 5.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We follow our related works in the setting for error bars.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We report the experimental environment in Sec. A.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: To the best of our knowledge, our work is conducted with NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: To the best of our knowledge, we do not foresee societal impacts of our work.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our work does not pose such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We properly cite the used existing datasets and models.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide detailed descriptions about our method (Sec. 4) as well as its limitations (Sec. A.5). Our code will be released upon paper acceptance.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not involve crowdsourcing research.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not involve crowdsourcing research.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.