

TranSplat: Generalizable 3D Gaussian Splatting from Sparse Multi-View Images with Transformers

Chuanrui Zhang^{1*}, Yingshuang Zou^{1*}, Zhuoling Li², Minmin Yi³, Haoqian Wang^{1†}

¹Tsinghua University,

²The University of Hong Kong,

³E-surfing Vision Technology Co., Ltd

{zhang-cr22, zouys22}@mails.tsinghua.edu.cn, lizhuoling@connect.hku.hk,
yimm18@chinatelecom.cn, wanghaoqian@tsinghua.edu.cn

Abstract

Compared with previous 3D reconstruction methods like Nerf, recent Generalizable 3D Gaussian Splatting (G-3DGS) methods demonstrate impressive efficiency even in the sparse-view setting. However, the promising reconstruction performance of existing G-3DGS methods relies heavily on accurate multi-view feature matching, which is quite challenging. Especially for the scenes that have many non-overlapping areas between various views and contain numerous similar regions, the matching performance of existing methods is poor and the reconstruction precision is limited. To address this problem, we develop a strategy that utilizes a predicted depth confidence map to guide accurate local feature matching. In addition, we propose to utilize the knowledge of existing monocular depth estimation models as prior to boost the depth estimation precision in non-overlapping areas between views. Combining the proposed strategies, we present a novel G-3DGS method named TranSplat, which obtains the best performance on both the RealEstate10K and ACID benchmarks while maintaining competitive speed and presenting strong cross-dataset generalization ability.

Project Page: — <https://xingyoujun.github.io/transplat/>

Introduction

The task of generalizable sparse-view 3D reconstruction, which aims to recover 3D structure from a sparse set of images without scene-specific optimization, is crucial in many applications like virtual reality (Zou et al. 2024; Li et al. 2023). Significant progress has been achieved using various neural scene representations, such as NeRF (Mildenhall et al. 2021; Wang et al. 2022a) and 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023; Lu et al. 2024). Unlike NeRF, which relies on computationally expensive ray-marching (Xu et al. 2022b; Gao et al. 2024), 3DGS represents the scene with a set of 3D Gaussians benefiting from rasterization-based rendering. These characteristics make 3DGS more efficient because it bypasses the dense volumetric ray-marching process and enables both efficient and high-quality 3D reconstruction and novel view synthesis.

However, traditional 3D reconstruction and novel view synthesis methods mostly require scene-specific optimization (Barron et al. 2021; Pons-Moll et al. 2021), resulting in increased computational demands and longer processing times. Recently, there has been growing interest in Generalizable 3D Gaussian Splatting (G-3DGS) for sparse-view reconstruction. Some methods (Charatan et al. 2024; Chen et al. 2024) have emerged that are capable of reconstructing 3D scenes in a single forward pass without requiring extra optimization. These advancements achieve high-quality reconstruction results while reducing computational overhead, thereby enhancing the efficiency and applicability of 3D reconstruction across various contexts. For example, PixelSplat (Charatan et al. 2024) and MVSplat (Chen et al. 2024) unproject depth maps from corresponding views to serve as the centers of 3D Gaussians. However, the performances of these methods rely on accurate pixel-level matching based on depth, which is challenging to realize and thus limits the reconstruction precision of these methods. We observe that especially in the scenes with occlusions, insufficient texture, or repetitive patterns, the matching result is quite unsatisfactory (Yao et al. 2018; Sun et al. 2021; Pautrat et al. 2023; Wang et al. 2022b). Additionally, the reconstruction of non-overlapping areas between views is also challenging for existing methods due to the lack of matching pairs.

To address the aforementioned limitations, we propose **TranSplat**, a novel framework for generalizable sparse-view 3D reconstruction. TranSplat recovers 3D structures from sparse views by projecting each pixel into 3D Gaussian primitives leveraging the predicted depths and features. To precisely estimate 3D Gaussian centers, we generate depth estimates using a transformer-based module called the Depth-aware Deformable Matching Transformer (DDMT). This module prioritizes depth candidates with high confidence based on an initial depth distribution, which is calculated by the Coarse Matching module through assessing cross-view feature similarities. Additionally, TranSplat utilizes monocular depth priors and employs the Depth Refine U-Net to further refine the depth distribution. With the refined depth and image features, our method predicts all Gaussian parameters in parallel for each pixel. TranSplat is trained end-to-end using only ground-truth images for supervision.

*These authors contributed equally.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

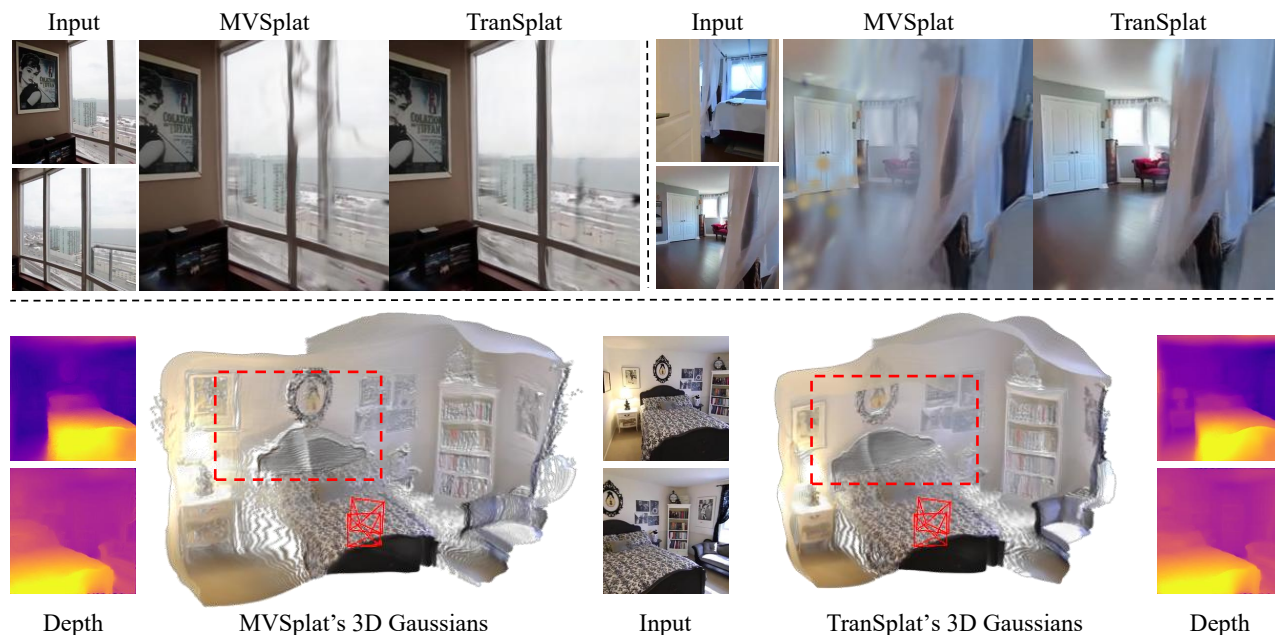


Figure 1: Given two-view images, TranSpLat achieves higher quality in both novel view synthesis and 3D Gaussian construction, particularly in challenging areas such as low-texture, repetitive patterns, and non-overlapping regions, compared to MVSpLat (Chen et al. 2024). Our novel view synthesis results contain fewer artifacts and maintain better geometric consistency.

We evaluate TranSpLat on two large-scale benchmarks: RealEstate10K (Zhou et al. 2018) and ACID (Liu et al. 2021a). Extensive experiments are conducted and demonstrate that TranSpLat achieves the best results in G-3DGS. Notably, compared to existing counterparts, TranSpLat presents strong cross-dataset generalization ability.

Comprehensively, our main contributions are as follows:

- We propose to utilize the depth confidence map to enhance matching between various views and correspondingly significantly improve the reconstruction precision in regions with insufficient texture or repetitive patterns.
- We propose a strategy that encodes the priors of monocular depth estimators into the prediction of Gaussian parameters, ensuring precise 3D Gaussian centers are estimated even in non-overlapping areas.
- The derived method TranSpLat achieves the best results on two large-scale benchmarks and presents strong cross-dataset generalization ability.

Related Work

Sparse-view Scene Reconstruction

With advancements in 3D scene reconstruction technologies such as NeRF and 3D Gaussian Splatting (3DGS), there is growing interest in reconstructing scenes from sparse inputs (Fan et al. 2024; Guo et al. 2024; Ni et al. 2024; Yu et al. 2021). The limited overlap in sparse inputs poses a significant challenge to scene reconstruction (Shi et al. 2024; Irshad et al. 2023). Existing sparse-view methods are primarily categorized into two main approaches: per-scene optimization methods and feed-forward inference methods. The former typically leverages multi-view geometry constraints to jointly optimize rendering results and camera

poses (Niemeyer et al. 2022; Truong et al. 2023; Deng et al. 2022). These methods may also improve novel-view results by incorporating additional prior knowledge, such as depths (Chung, Oh, and Lee 2024; Li et al. 2024). However, per-scene optimization methods require iterative processes to achieve the final 3D representation. In contrast, feed-forward inference methods reconstruct the entire scene in a single feed-forward pass, eliminating the need for subsequent optimization. These methods benefit from powerful priors from large datasets, which enables them to generalize effectively across different datasets.

3DGS for Generalizable Reconstruction

Currently, 3DGS-based generalizable reconstruction methods are attracting significant interest due to their rapid rendering speed and superior cross-dataset generalization capabilities. For example, PixelSpLat (Charatan et al. 2024) employs a multi-view epipolar transformer to accurately infer the real scale of scenes, effectively resolving the inherent scale ambiguity in real-world datasets. It predicts Gaussian parameters through dense probability distributions, which facilitates accurate and efficient 3D structure estimation. MVSpLat (Chen et al. 2024) reconstructs 3D scenes by constructing a cost volume that enables joint prediction of depth and Gaussian parameters. The centers of the Gaussians are determined by unprojecting the depth map into 3D space. LatentSpLat (Wewer et al. 2024) addresses reconstruction challenges in object-centric scenes with 360° views by introducing variational Gaussians. It uses a lightweight VAE-GAN decoder to generate RGB images of novel views. Despite these advancements, 3DGS-based methods still face challenges, such as achieving robust reconstructions in areas with occlusions, lack of texture, or repetitive patterns.

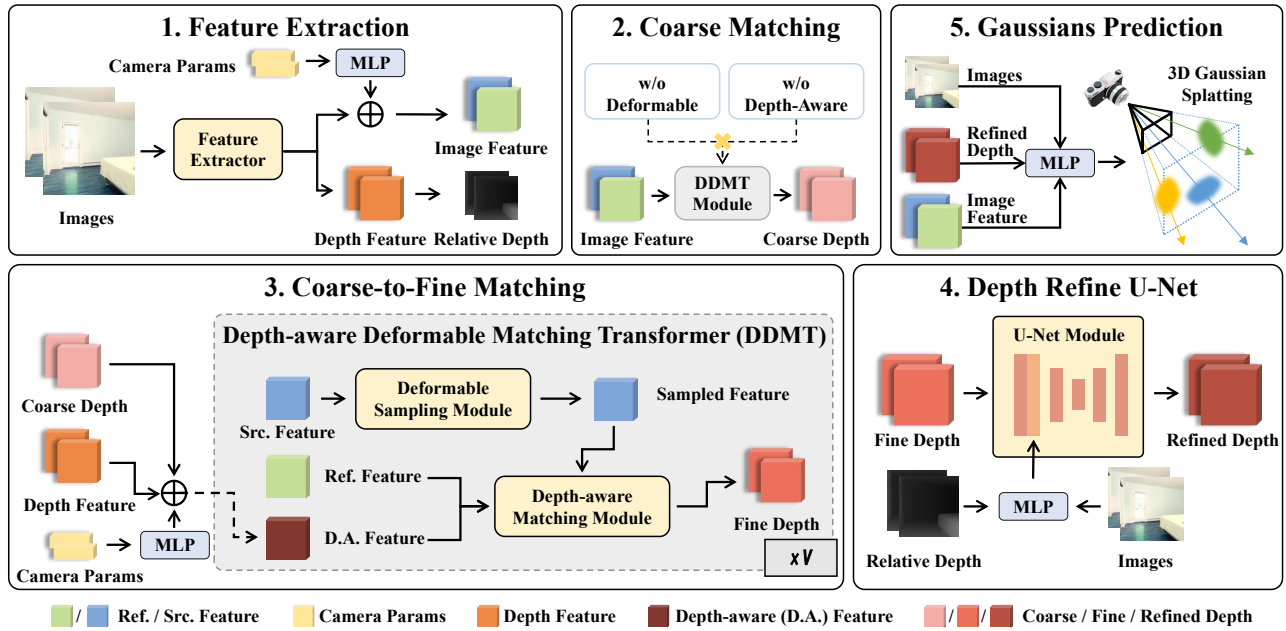


Figure 2: **Framework of TranSplat.** Our method takes multi-view images as input and first extracts image features and monocular depth priors. Specifically, we compute multi-view feature similarities using our proposed Depth-Aware Deformable Matching Transformer module to obtain a geometry-consistent depth distribution. The Depth Refine U-Net is then employed to further refine the depth prediction. Finally, we predict pixel-wise 3D Gaussian parameters to render novel views.

Multi-view Stereo

Multi-View Stereo (MVS) techniques enable the reconstruction of 3D structures from 2D images (Collins 1996). Recent learning-based MVS methods (Yao et al. 2018; Gu et al. 2020; Ding et al. 2022; Peng et al. 2022; Xu et al. 2023) utilize 3D cost volumes to predict depth for each view and then fuse these depth maps to obtain the final point cloud. These methods typically require dense depth maps for effective model training. In contrast, our proposed method directly extracts the 3D representation from sparse views, eliminating the need for separate depth estimation and point cloud fusion processes. Notably, our approach requires only RGB images for supervision, simplifying data acquisition and reducing reliance on costly depth sensors.

Method

We introduce **TranSplat**, a novel generalizable sparse-view scene reconstruction network that employs a transformer-based architecture. The overview of our proposed method is illustrated in Figure 2. Given K sparse-view images $\{\mathbf{I}^i\}_{i=1}^K$ (where $\mathbf{I}^i \in \mathbb{R}^{H \times W \times 3}$) and their corresponding camera parameters, our objective is to predict a 3D Gaussian representation of the scene through a single forward pass. The Gaussian parameters include position μ , opacity α , covariance Σ , and color c (represented as spherical harmonics). Using these parameters, we can render any novel views.

Feature Extraction

We employ a standard CNN and Transformer architecture (Xu et al. 2022a, 2023) to extract multi-view image features $\{\mathbf{F}^i\}_{i=1}^K$ (where $\mathbf{F}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$), with C denoting

the channel dimension. Notably, we incorporate camera parameters into each CNN feature using a squeeze-excitation (SE) layer (Hu, Shen, and Sun 2018) to provide global spatial information. In our self- and cross-attention layers, we adopt the local window attention mechanism of the Swin Transformer (Liu et al. 2021b) to enhance information exchange between different views. Furthermore, we utilize the large vision model DepthAnythingV2 (Yang et al. 2024) to obtain monocular depth priors. We extract the last layer of the DepthAnythingV2 module as depth features $\{\mathbf{D}^i\}_{i=1}^K$ (where $\mathbf{D}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$), as well as the relative depth outputs $\{\mathbf{d}_{rel}^i\}_{i=1}^K$ (where $\mathbf{d}_{rel}^i \in \mathbb{R}^{H \times W}$).

Coarse Matching

Following feature extraction, acquiring an initial depth distribution is essential for generating depth confidence maps. To simplify the code implementation, the DDMT module is utilized directly, bypassing the need for deformable sampling and depth-aware matching. This approach is taken because those models rely on depth information to produce attention maps for depth candidates.

Consider the source view $i \in K$ and the target view $j \in K$ (where $i \neq j$) as an example. We first construct a depth candidate $\mathbf{d}_{cand}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D}$ using the plane-sweep stereo approach (Yao et al. 2018). Then, we sample the j -th view feature $\mathbf{F}_{sample}^{ij} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D \times C}$ as follows:

$$\mathbf{F}_{sample}^{ij} = \text{Sample}(\mathbf{F}^j, \mathbf{P}^i, \mathbf{P}^j, \mathbf{d}_{cand}^i), \quad (1)$$

where Sample denotes the sampling operation (Xu et al. 2023), and \mathbf{P}^i and $\mathbf{P}^j \in \mathbb{R}^{4 \times 4}$ represent the camera projection matrices from the i -th and j -th views to the world coordinate system. D denotes the depth dimension.

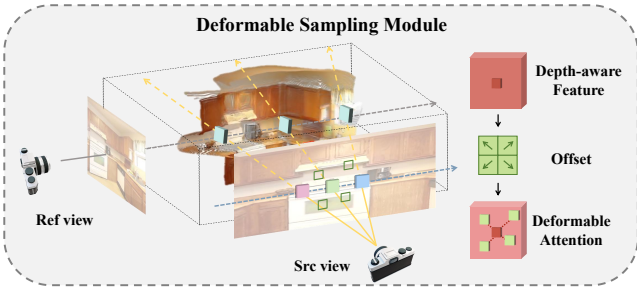


Figure 3: **The details of Deformable Sampling module.** We sample cross-view points using deformable attention, which enhances the aggregation of local spatial information.

Subsequently, we compute the dot product between \mathbf{F}^i and \mathbf{F}_{sample}^{ij} to obtain the coarse depth $\mathbf{d}_{coarse}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D}$:

$$\mathbf{d}_{coarse}^i = \frac{\mathbf{F}^i \otimes \mathbf{F}_{sample}^{ij}}{\sqrt{C}}, \quad (2)$$

where \otimes represents the element-wise product, and the results are summed along the channel dimension. When more than two views are provided as input, we apply the Coarse Matching module to the other views and compute the pixel-wise average across all views to obtain the final coarse depth.

Coarse-to-Fine Matching

Since the Coarse Matching module only performs pixel-aligned feature matching, it struggles with low-texture areas and repetitive patterns. To address this issue, we designed the Deformable Sampling (DS) module (in Figure 3), which aggregates the spatial information of local features. Furthermore, we observed that the Coarse Matching module equally prioritizes all depth candidates. To overcome this limitation, we propose the Depth-aware Matching Transformer (DMT) module (in Figure 4) to guide our network to focus on the correct depth candidates.

Similar to Section , we first consider the source view i and the target view j . To construct the depth-aware feature (d.a. feature) $\mathbf{F}_{d.a}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$, we fuse the coarse depth map \mathbf{d}_{coarse}^i , depth feature \mathbf{D}^i , and camera parameters \mathbf{P}^i through the depth-aware network:

$$\mathbf{F}_{d.a}^i = f_{\theta}(\mathbf{d}_{coarse}^i) + f_{\phi}(\mathbf{D}^i, \mathbf{P}^i), \quad (3)$$

where f_{θ} represents a Multi-Layer Perceptron (MLP) layer, and f_{ϕ} denotes an SE layer employed to inject position information into the depth feature.

Next, we apply the Deformable Sampling module to obtain the deformable sampled feature $\mathbf{F}_{d.s.}^{ij} \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D \times C}$ of the j -th view:

$$\mathbf{F}_{d.s.}^{ij} = \text{D.S.}(\mathbf{F}^j, \mathbf{P}^j, \mathbf{d}_{cand.}^i, \Delta \mathbf{p}^i), \quad (4)$$

where D.S. represents the Deformable Sampling module, and $\Delta \mathbf{p}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D \times P}$ represents the image-level sampling offsets, with P as the number of deformable points.

Our proposed Depth-Aware Matching module contains self- and cross-attention layers to enhance depth prediction. The coarse depth \mathbf{d}_{coarse}^i is first input into the self-attention

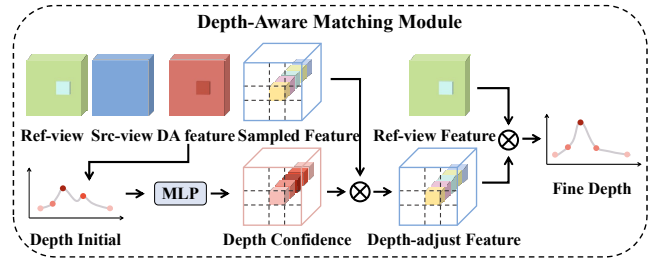


Figure 4: **The illustration of Depth-aware Matching Transformer.** We utilize the depth confidence map to integrate into our Depth-aware Matching Transformer, enabling the network to prioritize regions with higher depth confidence and improving the overall depth prediction accuracy. (Here D.A. feature represents depth-aware feature.)

layer to aggregate local depth information. Next, we calculate the depth residual $\Delta \mathbf{d}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D}$ in the cross-attention layer:

$$\Delta \mathbf{d}^i = \frac{\mathbf{F}^i \otimes \sum_P (\mathbf{W}^i \otimes \mathbf{F}_{d.s.}^{ij})}{\sqrt{C}}, \quad (5)$$

where \otimes represents the element-wise product, and the results are summed along the channel dimension. $\mathbf{W}^i \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D \times P}$ represents the depth-aware attention weights. The sampling offsets and attention weights are predicted through the offset layer f_p and the attention layer f_w , respectively:

$$\Delta \mathbf{p}^i = f_p(\mathbf{F}_{d.a}^i), \quad \mathbf{W}^i = f_w(\mathbf{F}_{d.a}^i). \quad (6)$$

The fine depth is then computed as follows:

$$\mathbf{d}_{fine}^i = \mathbf{d}_{coarse}^i + \Delta \mathbf{d}^i. \quad (7)$$

Similar to Section , we get the fine depth of K views separately when there are more than two views as input.

Depth Refine U-Net

Recent methods (Charatan et al. 2024; Chen et al. 2024) primarily focus on cross-image information to infer depth. However, as the number of input views decreases, there are large areas without cross-view matches, resulting in unreliable depth outputs for those regions. To optimize the depth distribution of these non-overlapping areas, we propose providing strong monocular depth priors.

We designed the Depth Refine U-Net module to combine the accurate geometric consistency from the matching results with the reliable relative depth from monocular depth priors. Inspired by diffusion models (Rombach et al. 2022; Duan, Guo, and Zhu 2023), our Depth Refine U-Net module receives input images and relative depths as conditions to refine the input depths.

Taking input images $\{\mathbf{I}^i\}_{i=1}^K$, relative depths $\{\mathbf{d}_{rel.}^i\}_{i=1}^K$, and fine depths $\{\mathbf{d}_{fine}^i\}_{i=1}^K$ as input, our module outputs per-view depth residuals:

$$\mathbf{d}_{refine}^i = \mathbf{d}_{fine}^i + f_r(\mathbf{d}_{fine}^i, \mathbf{I}^i, \mathbf{d}_{rel.}^i), \quad (8)$$

where f_r denotes our Depth Refine U-Net module and $\{\mathbf{d}_{refine}^i\}_{i=1}^K$ ($\mathbf{d}_{refine}^i \in \mathbb{R}^{H \times W}$) represents the final refined depths.

Gaussian Prediction

In this section, we leverage our depth-aware features to predict the parameters of a set of 3D Gaussian primitives $\{(\mu_i, \alpha_i, \Sigma_i, c_i)\}_{i=1}^M$ that parameterize the scene. We set $M = H \times W \times K$ following (Charatan et al. 2024) as we predict a Gaussian primitive per pixel from K views.

Gaussian center μ . We utilize the final refined depth d_{refine}^i predictions to directly unproject every pixel to global 3D points using the camera parameters. These points are then directly selected as the centers of the 3D Gaussians.

Opacity α . For each pixel, the opacity should be set to 1 if the depth prediction is accurate, as the points must lie on the surface. The opacity can be regarded as the confidence of the depth prediction. Therefore, we can predict opacity using a simple MLP layer with the refined depth prediction as input.

Covariance Σ and Color c . Similar to other 3D Gaussian approaches (Kerbl et al. 2023; Chen et al. 2024), we calculate color from the predicted spherical harmonic coefficients and use a scaling matrix s and a rotation matrix $R(\theta)$ to represent the covariance matrix Σ .

$$\Sigma = R(\theta)^T \text{diag}(s) R(\theta). \quad (9)$$

These parameters are predicted using MLP layers.

Experiments

Experimental Setup

Datasets. We train and evaluate our model using the large-scale RealEstate10K (Zhou et al. 2018) and ACID (Liu et al. 2021a) datasets. RealEstate10K comprises home walk-through videos from YouTube, with 67,477 scenes for training and 7,289 scenes for testing. The ACID dataset, featuring aerial landscape videos, includes 11,075 training scenes and 1,972 testing scenes. Both datasets provide camera poses computed by SFM software. In line with MVSplat (Chen et al. 2024), we train our model using two context views and evaluate all methods from three novel target viewpoints for each test scene. Additionally, to assess cross-dataset generalization, we evaluate all methods on the multi-view DTU dataset (Jensen et al. 2014), selecting 16 validation scenes with four novel views each.

Metrics. To obtain quantitative results, we compare the novel RGB images produced by each method with ground-truth frames using the Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM)(Wang et al. 2004), and Perceptual Distance (LPIPS)(Zhang et al. 2018).

Implementation details. Input images are resized to 256×256 , following the method outlined in (Chen et al. 2024). In all experiments, the number of depth candidates is set to 128. We sample $P = 4$ deformable points in the Depth-Aware Deformable Matching Transformer for the main results. For the DepthAnythingV2 (Yang et al. 2024) module, we use the base size to balance training cost and result quality. All models are trained with a batch size of 14 on 7 RTX 3090 GPUs for 300,000 iterations using the Adam (Kingma 2014) optimizer. During inference, we measure speed and memory cost with one RTX 3090 GPU.

Main Results

Novel view synthesis quality. We report quantitative results in Table 1. Our evaluation compares TranSplat with state-of-the-art (SOTA) methods on scene-level novel view synthesis from sparse views. TranSplat outperforms the baselines across all metrics, with competitive inference speed (0.06s vs. 0.08s). Under identical training conditions, TranSplat not only exceeds all previous SOTA methods in visual quality, as indicated by the PSNR and SSIM metrics but also shows improvements in the LPIPS metric, which better aligns with human perception. Notably, TranSplat achieves superior performance compared to other SOTA methods with only 200K iterations (66% of 300K), demonstrating that our transformer-based architecture enhances convergence efficiency by focusing on accurate depth candidates.

The qualitative comparisons of the three methods are visualized in Figure 5. TranSplat achieves superior quality on novel view images across various scenes. The first row showcases a scene with a low-texture area, such as the "fridge surface," where our method excels. This improvement is due to our Depth-Aware Deformable Matching Transformer, which aids the network in aggregating long-term spatial information and concentrating on the correct areas through attention mechanisms. Our method also demonstrates superiority in mismatched scenes, such as the "in-room area" in the second row and the "lamp" in the third row. This enhancement is attributable to our Depth Refine U-Net, which utilizes monocular depth priors to refine Gaussian predictions in non-overlapping areas. Similar results are evident in the ACID datasets, as shown in the fourth row.

Reconstruction quality. TranSplat generates impressive 3D Gaussian primitives compared to the recent state-of-the-art method MVSplat (Chen et al. 2024), as shown in Figure 6. Our TranSplat method can produce high-quality depth maps in challenging scenes, such as those with repetitive patterns (e.g., "the walls"). Notably, our method significantly outperforms in generating 3D Gaussians in non-overlapping areas (e.g., "the roof on the right"). These results collectively demonstrate the superiority of our proposed method.

Cross-dataset Generation

Our proposed TranSplat demonstrates significant superiority in generalizing to *out-of-distribution* novel scenes, as illustrated in Table 2. Following the settings of MVSplat (Chen et al. 2024), we directly use the model trained on RealEstate10K (an indoor dataset) to test on ACID (an outdoor dataset) and DTU (an object-centric dataset). The results on the ACID dataset are nearly identical, as ACID consists of aerial landscape scenes that may be reduced to an image interpolation task. TranSplat outperforms other state-of-the-art methods on the DTU dataset by a substantial margin across all metrics. Even with only 100K (33% of 300K) iterations of training, our method exhibits excellent generalization in cross-dataset settings. This superiority is attributed to our Depth-Aware Deformable Matching Transformer, which efficiently adjusts the attention area as the dataset's scale changes. Additionally, the monocular depth prior contributes to more generalized Gaussian parameter predictions across different multi-view camera distributions.




Figure 5: **The qualitative comparisons with SOTA methods.** Our method outperforms other state-of-the-art methods across various scenes, with the first three rows from RealEstate10K and the last one from ACID. TranSplat excels in challenging regions thanks to the effectiveness of our transformer-based Depth-Aware Deformable Matching Transformer (DDMT).

Method	RealEstate10K			ACID			Consumption	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Memory \downarrow	Time \downarrow
pixelSplat(Charatan et al.)	25.89	0.858	0.142	28.14	0.839	0.150	5319 MB	0.1553 s
MVSplat(Chen et al.)	26.39	0.869	0.128	28.25	0.843	0.144	2906 MB	0.0667 s
Ours 200k	26.44	0.871	0.128	28.20	0.843	0.145	2228 MB	0.0851 s
Ours 300k	26.69	0.875	0.125	28.35	0.845	0.143	2228 MB	0.0872 s

Table 1: **Quantitative comparisons SOTA methods.** Our TranSplat outperforms others on large-scale benchmarks, offering competitive inference speed and memory efficiency. Remarkably, it achieves the best performance with only 200K iterations (66% of 300K) in the RealEstate10K dataset, a testament to the effectiveness of our transformer-based architecture.

We present additional results in Figure 7. TranSplat inherently improves rendering quality in novel views, even when scenes differ significantly in appearance. In contrast, the results of MVSplat contain artifacts, indicating poor geometric consistency. Our TranSplat method provides reliable depth information for mismatched areas, whereas MVSplat tends to force-fit the depth distribution of the training dataset.

Ablation Study

To evaluate the effect of our different modules, we conduct ablation experiments on the RealEstate10K (Zhou et al. 2018) dataset with a batch size of 9, trained for 100K iterations. All experiments show that our proposed modules contribute to performance improvements. Additional hyper-

parameter ablation details are provided in the Supplements.

Importance of Depth-aware Deformable Matching Transformer. The Depth-Aware Deformable Matching Transformer (DDMT) is responsible for maintaining multi-view geometric consistency. The DDMT module requires an initial depth estimate to generate depth-adjusted features for efficient matching. We present the results in Table 3. To assess the importance of the DDMT and the influence of different initial depths, we conduct ablation experiments with random depth estimates (1st row), only coarse matching (2nd row), and coarse-to-fine strategies (3rd row). We find that using random depth estimates as the initial input causes the DDMT module to fail in generating a meaningful depth con-




Figure 6: The qualitative comparisons on 3D Gaussians. TranSplat generates significantly higher-quality 3D Gaussians in scenes with repetitive patterns (e.g., walls) and non-overlapping areas (e.g., the right roof), showcasing the effectiveness of our proposed modules.

Method	Re10k→DTU			Re10k→ACID		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
pixelSplat	12.89	0.382	0.560	27.64	0.830	0.160
MVSplat	13.94	0.473	0.385	<u>28.15</u>	<u>0.841</u>	<u>0.147</u>
Ours 100k	<u>14.72</u>	<u>0.517</u>	<u>0.344</u>	27.73	0.834	0.155
Ours 300k	14.93	0.531	0.326	28.17	0.842	0.146

Table 2: **Cross-dataset generalization results.** We conduct zero-shot tests on the DTU dataset (Jensen et al. 2014) using models trained on RealEstate10K (Zhou et al. 2018) and ACID (Liu et al. 2021a) without any fine-tuning. TranSplat demonstrates substantial generalization ability, outperforming other state-of-the-art methods on the DTU dataset, even with only 100K iterations of training.




Figure 7: **The qualitative comparisons of cross-dataset generalization.** Using a model trained on RealEstate10K, we directly render scenes from the ACID dataset (first row) and the DTU dataset (last three rows). TranSplat produces images with notable quality improvements, demonstrating its ability to adapt to diverse depth distributions and handle low-texture regions without requiring additional training.

confidence map, resulting in performance similar to that of simply using the coarse matching module. However, when providing coarse depth estimates, the DDMT module improves results, as the model can focus on depth-informative areas.

C.M.	DDMT	D.R.U.	C.P.E.	PSNR↑	SSIM↑	LPIPS↓
	✓			25.011	0.841	0.152
✓				25.020	0.842	0.150
✓	✓			25.231	0.846	0.148
✓	✓	✓		25.516	0.853	0.145
✓	✓	✓	✓	25.615	0.855	0.142

Table 3: Ablation study on the RealEstate10K (Zhou et al. 2018) dataset. C.M. and DDM refer to Coarse Matching and the Depth-Aware Deformable Matching module, respectively. D.R.U. and C.P.E. stand for Depth Refine U-Net and Camera Parameter Encoding. The ablation results demonstrate the effectiveness of our proposed methods, with the final line representing our complete model.

Importance of Depth Refine U-Net. We use the Depth Refine module to optimize the depth map by incorporating monocular depth information from the DepthAnythingV2 (Yang et al. 2024). To investigate the necessity of this module, we perform an ablation study without the Depth Refine U-Net (4th row in Table 3). We observe performance improvements with the Depth Refine module, as it provides reliable depth information for non-overlapping areas.

Importance of Camera Parameter Encoding. We investigate the significance of Camera Parameter Encoding (C.P.E.) for feature extraction by performing an ablation study on its use (5th row in Table 3). The results demonstrate that C.P.E. is essential for injecting global spatial information into feature maps, leading to improved performance.

Ablation of multi-view number. As discussed in Section , our proposed TranSplat is designed to process multi-view images. Our method benefits from additional input views, as more matchable areas are provided. We test our multi-view capability directly on the DTU dataset using a model trained on the RealEstate10K dataset with 2-view inputs. Compared to the 2-view results presented in Table 2, TranSplat achieves a PSNR of 15.14, an SSIM of 0.561, and an LPIPS of 0.330 with 3-view inputs, and a PSNR of 15.22, an SSIM of 0.570, and an LPIPS of 0.340 with 4-view inputs. We observe that TranSplat’s performance improves as the number of input views increases, even without additional training.

Conclusion

In this work, we introduce TranSplat, a novel generalizable sparse-view scene reconstruction network with a transformer-based architecture. Equipped with multi-view inputs, our method predicts a set of 3D Gaussian primitives to represent the scene. To address incorrect matches caused by similar areas, we propose the Depth-Aware Deformable Matching Transformer, which enhances attention to accurate depth distribution. To improve 3D Gaussians in non-overlapping areas, we designed the Depth Refine module, which optimizes the depth map using monocular depth priors. Our method achieves the SOTA performance on two large-scale scene-level reconstruction benchmarks and demonstrates superior cross-dataset generalization.

Acknowledgments

This research was funded through National Key Research and Development Program of China (Project No. 2022YFB36066)

References

- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5855–5864.
- Charatan, D.; Li, S. L.; Tagliasacchi, A.; and Sitzmann, V. 2024. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19457–19467.
- Chen, Y.; Xu, H.; Zheng, C.; Zhuang, B.; Pollefeys, M.; Geiger, A.; Cham, T.-J.; and Cai, J. 2024. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images. *arXiv preprint arXiv:2403.14627*.
- Chung, J.; Oh, J.; and Lee, K. M. 2024. Depth-regularized optimization for 3d gaussian splatting in few-shot images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 811–820.
- Collins, R. T. 1996. A space-sweep approach to true multi-image matching. In *Proceedings CVPR IEEE computer society conference on computer vision and pattern recognition*, 358–363. Ieee.
- Deng, K.; Liu, A.; Zhu, J.-Y.; and Ramanan, D. 2022. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12882–12891.
- Ding, Y.; Yuan, W.; Zhu, Q.; Zhang, H.; Liu, X.; Wang, Y.; and Liu, X. 2022. Transmvsnet: Global context-aware multi-view stereo network with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8585–8594.
- Duan, Y.; Guo, X.; and Zhu, Z. 2023. DiffusionDepth: Diffusion Denoising Approach for Monocular Depth Estimation. *arXiv:2303.05021*.
- Fan, Z.; Cong, W.; Wen, K.; Wang, K.; Zhang, J.; Ding, X.; Xu, D.; Ivanovic, B.; Pavone, M.; Pavlakos, G.; et al. 2024. Instantsplat: Unbounded sparse-view pose-free gaussian splatting in 40 seconds. *arXiv preprint arXiv:2403.20309*.
- Gao, X.; Yang, Z.; Zhao, Y.; Sun, Y.; Jin, X.; and Zou, C. 2024. A General Implicit Framework for Fast NeRF Composition and Rendering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1833–1841.
- Gu, X.; Fan, Z.; Zhu, S.; Dai, Z.; Tan, F.; and Tan, P. 2020. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2495–2504.
- Guo, S.; Wang, Q.; Gao, Y.; Xie, R.; and Song, L. 2024. Depth-Guided Robust and Fast Point Cloud Fusion NeRF for Sparse Input Views. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1976–1984.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 7132–7141.
- Irshad, M. Z.; Zakharov, S.; Liu, K.; Guizilini, V.; Kollar, T.; Gaidon, A.; Kira, Z.; and Ambrus, R. 2023. Neo 360: Neural fields for sparse view synthesis of outdoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9187–9198.
- Jensen, R.; Dahl, A.; Vogiatzis, G.; Tola, E.; and Aanæs, H. 2014. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 406–413.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4).
- Kingma, D. 2014. Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Li, J.; Zhang, J.; Bai, X.; Zheng, J.; Ning, X.; Zhou, J.; and Gu, L. 2024. Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20775–20785.
- Li, Z.; Zhang, C.; Ma, W.-C.; Zhou, Y.; Huang, L.; Wang, H.; Lim, S.; and Zhao, H. 2023. Voxelformer: Bird’s-eye-view feature generation based on dual-view attention for multi-view 3d object detection. *arXiv preprint arXiv:2304.01054*.
- Liu, A.; Tucker, R.; Jampani, V.; Makadia, A.; Snavely, N.; and Kanazawa, A. 2021a. Infinite nature: Perpetual view generation of natural scenes from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14458–14467.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Lu, T.; Yu, M.; Xu, L.; Xiangli, Y.; Wang, L.; Lin, D.; and Dai, B. 2024. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20654–20664.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Ni, Z.; Yang, P.; Yang, W.; Wang, H.; Ma, L.; and Kwong, S. 2024. ColNeRF: Collaboration for Generalizable Sparse Input Neural Radiance Field. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4325–4333.
- Niemeyer, M.; Barron, J. T.; Mildenhall, B.; Sajjadi, M. S.; Geiger, A.; and Radwan, N. 2022. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5480–5490.
- Pautrat, R.; Suárez, I.; Yu, Y.; Pollefeys, M.; and Larsson, V. 2023. Gluestick: Robust image matching by sticking points and lines together. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9706–9716.

- Peng, R.; Wang, R.; Wang, Z.; Lai, Y.; and Wang, R. 2022. Rethinking depth estimation for multi-view stereo: A unified representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8645–8654.
- Pons-Moll, G.; Moreno-Noguer, F.; Corona, E.; and Pumarola, A. 2021. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Shi, R.; Wei, X.; Wang, C.; and Su, H. 2024. ZeroRF: Fast Sparse View 360deg Reconstruction with Zero Pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21114–21124.
- Sun, J.; Shen, Z.; Wang, Y.; Bao, H.; and Zhou, X. 2021. LoFTR: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8922–8931.
- Truong, P.; Rakotosaona, M.-J.; Manhardt, F.; and Tombari, F. 2023. Sparf: Neural radiance fields from sparse and noisy poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4190–4200.
- Wang, P.; Chen, X.; Chen, T.; Venugopalan, S.; Wang, Z.; et al. 2022a. Is Attention All That NeRF Needs? *arXiv preprint arXiv:2207.13298*.
- Wang, X.; Zhu, Z.; Huang, G.; Qin, F.; Ye, Y.; He, Y.; Chi, X.; and Wang, X. 2022b. Mvster: Epipolar transformer for efficient multi-view stereo. In *European Conference on Computer Vision*, 573–591. Springer.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4): 600–612.
- Weyer, C.; Raj, K.; Ilg, E.; Schiele, B.; and Lenssen, J. E. 2024. latentsplat: Autoencoding variational gaussians for fast generalizable 3d reconstruction. *arXiv preprint arXiv:2403.16292*.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofghi, H.; and Tao, D. 2022a. Gmflow: Learning optical flow via global matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8121–8130.
- Xu, H.; Zhang, J.; Cai, J.; Rezatofghi, H.; Yu, F.; Tao, D.; and Geiger, A. 2023. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Xu, Q.; Xu, Z.; Philip, J.; Bi, S.; Shu, Z.; Sunkavalli, K.; and Neumann, U. 2022b. Point-nerf: Point-based neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5438–5448.
- Yang, L.; Kang, B.; Huang, Z.; Zhao, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth Anything V2. *arXiv:2406.09414*.
- Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. Mvs-net: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, 767–783.
- Yu, A.; Ye, V.; Tancik, M.; and Kanazawa, A. 2021. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4578–4587.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhou, T.; Tucker, R.; Flynn, J.; Fyffe, G.; and Snavely, N. 2018. Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817*.
- Zou, Y.; Ding, Y.; Qiu, X.; Wang, H.; and Zhang, H. 2024. M2Depth: Self-supervised Two-Frame Multi-camera Metric Depth Estimation. *arXiv preprint arXiv:2405.02004*.