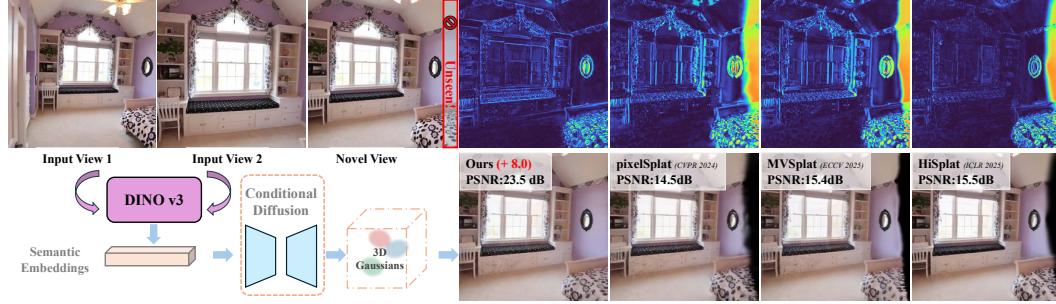


000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 CoDiffSPLAT: SPARSE-VIEW GENERALIZABLE 3D GAUSSIAN SPLATTING WITH SINGLE-STEP CONDI- TIONAL DIFFUSION

006
007 **Anonymous authors**
008 Paper under double-blind review



022 **Figure 1: Comparison between CoDiffSplat and baselines.** Previous methods are pixel-aligned
023 and depth-based, often failing to recover geometry in ambiguous or occluded regions (e.g., unsee
024 surfaces). In contrast, CoDiffSplat leverages semantic embeddings and conditional diffusion to
025 generate refined 3D Gaussians, enabling high-fidelity reconstruction of uncertain areas. Residual
026 maps between rendered results and ground truth are shown in the top-right corner.

ABSTRACT

030 Generalizable 3D Gaussian Splatting (G-3DGS) has emerged as a promising ap-
031 proach for novel view synthesis under sparse-view settings. However, existing
032 frameworks remain restricted by pixel-aligned Gaussian estimation, which strug-
033 gles in regions that are partially observed or occluded, often resulting in incom-
034 plete geometry and structural collapse. To overcome these challenges, we propose
035 CoDiffSplat, a new framework that couples semantic-conditioned latent diffusion
036 with 3D Gaussian splatting. Our design departs from conventional diffusion ap-
037 plied on image feature maps: instead, a lightweight single-step diffusion directly
038 refines Gaussian parameters, ensuring efficiency while preserving geometric con-
039 sistency. In addition, we introduce a Cross-View Entropy-Aware (CEA) module
040 that aggregates multi-view semantics and geometry into robust conditional em-
041 beddings, enabling diffusion to resolve ambiguities under occlusion and sparse
042 overlap. Comprehensive experiments on multiple benchmarks demonstrate that
043 CoDiffSplat consistently improves geometric quality and structural completeness,
044 especially under challenging extrapolation settings. Our study establishes con-
045 ditional diffusion as a scalable refinement mechanism for sparse-view 3D recon-
046 struction, advancing the reliability of generalizable Gaussian splatting.

1 INTRODUCTION

049 3D reconstruction is a cornerstone of computer vision, powering applications such as autonomous
050 driving, virtual reality, and augmented reality. Recently, 3D Gaussian Splatting (3DGS) Kerbl et al.
051 (2023) has emerged as a powerful paradigm for scene representation and novel view synthesis
052 (NVS) Buehler et al. (2001). By modeling a scene as a mixture of 3D Gaussians and leveraging
053 differentiable rasterization, 3DGS enables high-fidelity and real-time rendering from dense multi-
view images. However, conventional 3DGS methods typically rely on per-scene optimization with

054 dense input views, which limits scalability and requires costly data acquisition. To alleviate these
 055 constraints, generalizable 3DGS (G-3DGS) has been developed to reconstruct scenes from only a
 056 few input views. These approaches employ pre-trained feed-forward models Chen et al. (2025b);
 057 Charatan et al. (2024); Zhang et al. (2025); Wewer et al. (2025) that encode scene priors from large-
 058 scale datasets, enabling rapid inference without scene-specific optimization.

059 Despite this progress, existing G-3DGS frameworks remain fundamentally restricted by their strong
 060 reliance on pixel-aligned unprojection Charatan et al. (2024); Chen et al. (2025b). Specifically, each
 061 pixel is mapped to a fixed number of Gaussian primitives based on estimated depth, making recon-
 062 struction highly sensitive to depth errors. However, under sparse-view conditions, depth estimation
 063 often suffers from occlusions, weak textures, and limited viewpoint overlap. Moreover, rigid pixel-
 064 level alignment hinders the recovery of unobserved or partially visible regions, frequently producing
 065 ‘black holes’ or collapsed structures in the final rendering (Fig. 1).

066 To address these limitations, we propose to **shift the reconstruction paradigm from pixel-space**
 067 **alignment to refinement in the Gaussian domain.** While previous works attempted to address
 068 uncertainty through depth regularization or feature fusion, they remain inherently constrained by
 069 pixel-level priors. Motivated by the remarkable success of diffusion models in 3D content generation
 070 tasks (such as text-to-3D synthesis Lin et al. (2025); He et al. (2025); Cao et al. (2024)), we investi-
 071 giate whether diffusion can synthesize missing structures and compensate for uncertainty. Building
 072 on this intuition, we introduce **CoDiffSplat**, a novel G-3DGS framework that integrates conditional
 073 diffusion to refine initial pixel-aligned Gaussians into geometrically consistent and complete struc-
 074 tures. Unlike conventional diffusion pipelines, which require costly iterative denoising, we show
 075 that *a single-step refinement suffices* to correct geometric inconsistencies, turning diffusion into an
 076 efficient correction module rather than a full generative process. Consequently, our pipeline both
 077 *retains the speed advantages of feed-forward G-3DGS approaches and alleviates the dependency on*
 078 *perfect depth alignment.* However, a key challenge in NVS is the absence of explicit text prompts.
 079 While pseudo-captioning (e.g., BLIP Li et al. (2023)) offers weak supervision, it often overlooks
 080 fine-grained details Patni et al. (2024). To address this, we design a **Cross-View Entropy-Aware**
 081 **(CEA)** module that fuses multi-view semantic cues with geometry-uncertainty signals, yielding
 082 detail-preserving embeddings that guide the diffusion process toward challenging regions. Tech-
 083 nically, to mitigate the training difficulty and high computational cost associated with diffusion,
 084 CoDiffSplat performs denoising in latent Gaussian space using a lightweight DiT backbone.

085 The main contributions are summarized as follows:

- 086 • We formulate sparse-view G-3DGS as a *latent Gaussian refinement* problem and present
 087 CoDiffSplat, which employs a lightweight DiT-based diffusion model with a single-step
 088 denoiser to relax rigid pixel alignment and restore missing geometry.
- 089 • We propose the CEA module, which combines semantic cues with depth-distribution en-
 090 tropy to emphasize uncertain regions and provide fine-grained conditional guidance for
 091 diffusion.
- 092 • We validate CoDiffSplat on standard benchmarks across interpolation and extrapolation
 093 settings, and CoDiffSplat consistently improves fidelity in both settings. In particular, rel-
 094 ative to the SOTA HiSplat Chen et al. (2025b), it achieves a **+2.32 dB** PSNR gain on
 095 RealEstate10K Zhou et al. (2018) in extrapolated views, while maintaining competitive
 096 computational efficiency.

097 2 RELATED WORK

100 2.1 NOVEL VIEW SYNTHESIS

101 Novel view synthesis (NVS) aims to render photo-realistic images from novel viewpoints using only
 102 a limited set of input images Buehler et al. (2001). Neural Radiance Fields (NeRF) Mildenhall et al.
 103 (2020); Yu et al. (2021); Pumarola et al. (2021); Barron et al. (2021; 2022) models scenes as con-
 104 tinuous volumetric radiance fields parameterized by neural networks. While NeRF-based methods
 105 have yielded impressive results in dense multi-view settings, they typically suffer from slow train-
 106 ing times, high memory usage, and suboptimal performance under sparse viewpoints due to their
 107 heavy reliance on per-ray MLP evaluations. In contrast, 3D Gaussian Splatting (3DGS) Kerbl et al.

(2023); Yu et al. (2024); Yang et al. (2024) introduces an explicit scene representation by modeling surfaces with anisotropic 3D Gaussian primitives. Each Gaussian is described by its position, covariance, color, and opacity, which can be differentiably rendered via a forward projection to the image plane. This explicit design significantly accelerates the rendering process compared to vanilla NeRF pipelines, yet many existing 3DGS approaches still assume relatively dense coverage of views for accurate geometry and appearance reconstruction. Consequently, their performance deteriorates for extremely sparse inputs, where geometric ambiguity and insufficient texture cues become major challenges.

116

117 2.2 SPARSE-VIEW GENERALIZABLE 3DGS

118

119 Sparse-View generalizable 3DGS methods focus on learning a feed-forward model capable of han-
 120 dling unseen scenes without per-scene re-optimization. PixelSplat Charatan et al. (2024) predict 3D
 121 Gaussian parameters from sparse multi-view inputs, leveraging an epipolar transformer for depth
 122 estimation. MVSplat Chen et al. (2025b) relies on cost-volume construction via plane sweeping to
 123 infer depth distributions. TranSplat Zhang et al. (2025) introduces a transformer-based architecture
 124 with depth-aware deformable matching for coarse-to-fine refinement. HiSplat Tang et al. (2025b)
 125 integrates hierarchical Gaussian features, leveraging iterative Gaussian alignment. eFreeSplat Min
 126 et al. (2024) eliminates epipolar priors by leveraging cross-view completion. DepthSplat Xu et al.
 127 (2025) bridges Gaussian splatting and depth estimation by leveraging pre-trained monocular depth
 128 features to enhance multi-view depth prediction. Despite these diverse approaches, most still rely
 129 heavily on estimated depth maps, which can become noisy or unreliable under sparse-view condi-
 130 tions. They also often utilize pixel-aligned Gaussian estimation, causing difficulties in recovering
 131 fine details or resolving ambiguities in unseen regions.

132

133 2.3 3D-AWARE DIFFUSION METHODS

134

135 Recent diffusion-based methods have expanded from 2D image generation to 3D-aware tasks such
 136 as image-to-3D and text-to-3D generation Lin et al. (2025); Cao et al. (2024); Nichol et al. (2022);
 137 Hong et al. (2024); Li et al. (2024); Tang et al. (2025a); Shi et al. (2024); Yang et al. (2025). Image-
 138 to-3D approaches attempt to recover 3D content from a single image, but they often suffer from
 139 ambiguity and incomplete geometry due to the lack of multi-view constraints. Text-to-3D pipelines
 140 leverage diffusion priors for novel asset creation, and can be adapted to image inputs through cap-
 141 tioning, thereby unifying the two paradigms. While diffusion provides strong generative flexibility,
 142 existing 3D-aware frameworks typically lack explicit multi-view consistency and assume access to
 143 dense supervision, which limits their applicability in sparse-view reconstruction. These limitations
 144 motivate our work, where we integrate diffusion-based priors with Gaussian splatting to enhance
 145 geometric completeness and robustness under sparse and extrapolated views.

146

147 3 METHODOLOGY

148

149 Following the G-3DGS framework Charatan et al. (2024); Chen et al. (2025b); Zhang et al. (2025);
 150 Min et al. (2024); Xu et al. (2025); Tang et al. (2025b), the input consists of V sparse-view images
 151 $\mathcal{I} = \{I_1, I_2, \dots, I_V\}$, where each image $I_i \in \mathbb{R}^{H \times W \times 3}$ is accompanied by its camera projection
 152 matrix derived from intrinsic and extrinsic parameters. The goal is to reconstruct the underlying 3D
 153 scene as a set of Gaussian primitives $\Theta = \{G_j\}_{j=1}^N$, where each primitive G_j is parameterized by
 154 its center μ_j , opacity α_j , covariance Σ_j , and color c_j . The number of Gaussians is typically set to
 155 $N = H \times W \times V$, corresponding to the input resolution and number of views. These primitives are
 156 subsequently rendered into novel views through differentiable Gaussian splatting Kerbl et al. (2023).

157

158 As illustrated in Figure 2, CoDiffSplat adopt a latent conditional diffusion paradigm tailored for
 159 sparse-view reconstruction. Specifically, our pipeline comprises two branches: one employ a multi-
 160 view Gaussian encoder to generate a coarse initialization of latent Gaussian parameters, and another
 161 for extracting a Cross-view Entropy-Aware (CEA) embedding as condition of subsequent diffusion
 162 model. Instead of operating over image-space noise or raw latent feature maps, our diffusion back-
 163 bone denoises in the latent space of Gaussian parameters, refining coarse geometry and appearance
 164 under semantic guidance. The refined latent Gaussians are decoded via an upsampling decoder to
 165 produce full-resolution Gaussians.

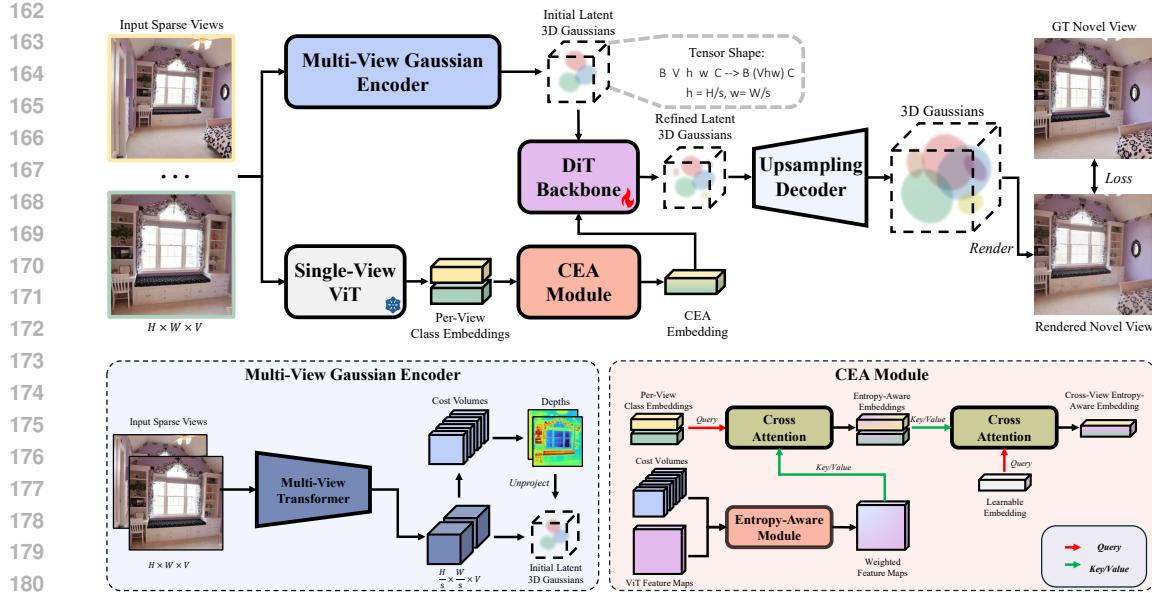


Figure 2: **Overview of CoDiffSplat.** Given sparse input views, our framework first constructs an initial set of latent 3D Gaussians through the multi-view Gaussian encoder. In parallel, a frozen single-view ViT extracts per-view class embeddings, which are fused by the proposed Cross-view Entropy-Aware (CEA) module to produce a unified multi-view semantic embedding. This embedding conditions the DiT backbone to refine the latent Gaussians, followed by an upsampling decoder that restores them to the original spatial resolution. The refined 3D Gaussians are rendered to synthesize novel views, and the entire model is optimized end-to-end using photometric loss.

3.1 GAUSSIAN INITIALIZATION

A reliable initialization is essential for stable diffusion-based refinement. We therefore introduce a multi-view Gaussian encoder that constructs latent 3D Gaussians directly from sparse input views. By aggregating multi-view features and inferring coarse scene geometry, this module produces a structured and pixel-aligned Gaussian representation. These initialized Gaussians provide strong geometric priors, serving as an effective starting point for the subsequent diffusion process.

3.1.1 LATENT FEATURE EXTRACTION

Diffusion models typically require many time steps to converge. Recently, Stable Diffusion Rombach et al. (2022) introduced an approach that shifts the denoising process into a learned latent representation, significantly easing optimization. Inspired by this paradigm, we employ a CNN and Transformer Vaswani et al. (2017) architecture to project multi-view images into latent space. Specifically, each input image I^i is first passed through a shallow ResNet He et al. (2016) to produce $s \times$ downsampled feature maps. To efficiently integrate cross-view information, we then leverage a multi-view Swin Transformer Liu et al. (2021b) which contains self- and cross-attention layers to obtain multi-view-aware features $\mathbf{F}^i \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s} \times C}$, where C is the feature dimension. This remains computationally tractable while preserving cross-view interactions. We then utilize the multi-view features to initialize latent Gaussian parameters.

3.1.2 COARSE MATCHING

To establish a reliable geometric prior for the subsequent diffusion process, we construct cost volumes with plane sweeping Xu et al. (2023); Yao et al. (2018) to model multi-view feature correspondences, which facilitates the initialization of coarse Gaussian parameters. Specifically, for each view i , we uniformly sample D depth candidates $\{d_m\}_{m=1}^D$ in the inverse depth domain between near and far planes. Features from other views ($\mathbf{F}^j, j \neq i$) are warped to view i using camera parameters and depth candidate d_m , producing D warped features $\{\mathbf{F}_{d_m}^{j \rightarrow i}\}_{m=1}^D$. The correlation $C_{d_m}^i$ between \mathbf{F}^i

216 and $\mathbf{F}_{d_m}^{j \rightarrow i}$ is computed with the dot-product operation:
 217

$$218 \quad 219 \quad 220 \quad C_{\text{dim}}^i = \frac{F_i \cdot F_{d_m}^{j \rightarrow i}}{\sqrt{C}}, \quad m = 1, 2, \dots, D. \quad (1)$$

221 We average $C_{d_m}^i$ across all other views to form the cost volume $\mathbf{C}^i \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s} \times D}$. Subsequently,
 222 we apply the softmax operation to compute per-view depth map:
 223

$$224 \quad \mathbf{Z}^i = \text{softmax}(\mathbf{C}^i) \cdot \mathbf{\Lambda}, \quad (2)$$

225 where $\mathbf{\Lambda} = [d_1, d_2, \dots, d_D]$ are the depth candidates. The coarse depth map $\mathbf{Z}^i \in \mathbb{R}^{\frac{H}{s} \times \frac{W}{s}}$ is
 226 then unprojected to form preliminary Gaussian centers $\mathbf{\mu}$ using the camera parameters, and other
 227 Gaussian parameters are predicted by additional lightweight heads from feature maps. This ensures
 228 that the initial positions are geometrically consistent with the input views. The constructed cost
 229 volumes \mathbf{C}^i are further retained as conditional inputs to the CEA module, providing pixel-wise
 230 structural cues about the scene.
 231

232 3.2 CONDITIONAL DIFFUSION

233 While recent works have applied diffusion to 3D reconstruction Chen et al. (2025a); Wu et al. (2025;
 234 2024), sparse-view NVS suffers from a critical limitation: the generative prior lacks sufficient con-
 235 straints, resulting in unstable refinements and hallucinated geometry. We address this challenge
 236 with a conditional diffusion module tailored for Gaussian splatting. Guided by cross-view entropy-
 237 aware (CEA) embeddings, the diffusion directly refines Gaussian parameters in latent space through
 238 a single-step process, ensuring both efficiency and consistency.
 239

240 3.2.1 CROSS-VIEW ENTROPY-AWARE EMBEDDING EXTRACTION

241 Existing conditional embeddings, such as BLIP text features Li et al. (2023), mainly capture global
 242 semantics and tend to overemphasize large salient objects, suppressing fine-grained details and of-
 243 ten causes boundary artifacts Patni et al. (2024). Furthermore, since these embeddings are extracted
 244 independently for each view, they contain redundant or overlapping content, which further amplifies
 245 occlusion-induced information loss. To address these issues, we introduce a Cross-View Entropy-
 246 Aware (CEA) embedding that highlights uncertain or weakly constrained regions while consolidat-
 247 ing cross-view redundancy into a comprehensive representation.
 248

249 We first propose an entropy-aware module that leverages the cost volume \mathbf{C}^i to compute the match-
 250 ing entropy H^i , thereby identifying weakly constrained regions (e.g., unseen, occluded, or tex-
 251 tureless areas). A single-view ViT encoder pretrained with DINOv3Siméoni et al. (2025) provides
 252 per-view class embeddings \mathbf{E}_{CLS}^i as well as feature maps \mathbf{F}^i . At each pixel p in view i , the depth
 253 posterior is estimated by applying a softmax along the depth axis of the cost volume:
 254

$$255 \quad P^i(d | p) = \text{softmax}_d(\mathbf{C}^i(p, d)). \quad (3)$$

256 The matching entropy is then defined as:
 257

$$258 \quad H^i(p) = - \sum_{d \in \Lambda} P^i(d | p) \log P^i(d | p), \quad (4)$$

259 and further normalized into per-pixel weights $w^i(p)$. This entropy quantifies the degree of multi-
 260 view uncertainty, where larger values of $w^i(p)$ indicate more ambiguous or under-constrained re-
 261 gions. The feature map \mathbf{F}^i is reweighted accordingly to obtain an entropy-aware representation:
 262

$$263 \quad \tilde{\mathbf{F}}^i(p) = w^i(p) \mathbf{F}^i(p). \quad (5)$$

264 Subsequently, the class embedding \mathbf{E}_{CLS}^i is linearly projected into queries, while the weighted
 265 feature map $\tilde{\mathbf{F}}^i$ is projected into keys and values:
 266

$$267 \quad \mathbf{Q}^i = \mathbf{E}_{CLS}^i \mathbf{W}_Q, \quad \mathbf{K}^i = \tilde{\mathbf{F}}^i \mathbf{W}_K, \quad \mathbf{V}^i = \tilde{\mathbf{F}}^i \mathbf{W}_V. \quad (6)$$

268 A cross-attention mechanism then refines the class embedding into an entropy-aware embedding:
 269

$$270 \quad \tilde{\mathbf{E}}_{CLS}^i = \text{CrossAttn}(\mathbf{Q}^i, \mathbf{K}^i, \mathbf{V}^i) + \mathbf{E}^i, \quad (7)$$

270 which adaptively integrates fine-grained cues, especially from regions of high uncertainty.
271

272 To further eliminate inter-view redundancy and alleviate the masking effect of large salient objects,
273 we employ a cross-view Perceiver-style attention mechanism with a set of learnable latent
274 queries. Denote these learnable queries as \mathbf{Q}_ℓ . By concatenating the refined per-view embeddings
275 $\tilde{\mathbf{E}}^i$, we construct the keys and values \mathbf{K}_{mv} and \mathbf{V}_{mv} , respectively, and compute the CEA embeddings
276 through a cross-attention:

$$277 \quad \mathbf{E}_{\text{CEA}} = \text{CrossAttn}(\mathbf{Q}_\ell, \mathbf{K}_{\text{mv}}, \mathbf{V}_{\text{mv}}). \quad (8)$$

278 The resulting representation \mathbf{E}_{CEA} aggregates cross-view semantics while suppressing redundant
279 biases, thereby providing a stable and informative conditioning signal for the subsequent diffusion.
280

281 3.2.2 GAUSSIAN-STRUCTURED REPRESENTATION

282 Although Gaussian primitives are inherently unordered point sets, our initialization procedure
283 imposes a pixel-to-Gaussian mapping, enabling the latent parameter tensor to be structured as
284 $\Theta_l \in \mathbb{R}^{B \times V \times h \times w \times C}$, where B denotes the batch size, V the number of views, $h \times w = \frac{H}{s} \times \frac{W}{s}$
285 the latent spatial resolution, and C the dimensionality of Gaussian parameters. This structured rep-
286 resentation aligns with the pixel domain and facilitates efficient parameter organization. However,
287 directly applying convolutional architectures such as UNet Ronneberger et al. (2015) to this ten-
288 sor may introduce spurious grid-based inductive biases, which can potentially oversmooth Gaussian
289 distributions and compromise geometric fidelity. To better respect the unordered nature of Gaussian
290 primitives, we instead employ a transformer-based diffusion backbone. Specifically, we flatten the
291 structured tensor into a sequence of tokens:
292

$$\Theta_l = \text{rearrange}(\Theta_l, B \ V \ h \ w \ C \rightarrow B \ (Vhw) \ C), \quad (9)$$

293 and process it using a Diffusion Transformer (DiT) backbone Peebles & Xie (2023), which is well
294 suited for irregular and unordered data.
295

296 3.2.3 SINGLE-STEP REFINEMENT

297 Motivated by recent findings that single-step diffusion suffices for refinement tasks Wu et al. (2025);
298 Lin et al. (2024); Qu et al. (2025), we design the diffusion stage as a *residual correction module*
299 rather than a full generative process. Specifically, the initialized latent Gaussians as Θ_l can be
300 served as coarse but noisy approximations of the target distribution. The refinement is performed
301 through a single-step correction:
302

$$\hat{\Theta}_l = \Theta_l + f_\theta(\Theta_l, \mathbf{E}_{\text{CEA}}), \quad (10)$$

303 where f_θ is a DiT-based predictor conditioned on CEA embeddings \mathbf{E}_{CEA} . This one-step formula-
304 tion mitigates structured noise while preserving geometric consistency, avoiding the overcorrection
305 and instability often observed in multi-step denoising Lin et al. (2024); Qu et al. (2025).
306

307 3.3 RENDERING AND TRAINING LOSS

309 Datasets for NVS do not provide explicit ground-truth 3DGS supervision, which prevents the use of
310 a conventional forward-reverse diffusion process. In particular, there is no well-defined target do-
311 main for injecting noise, and thus no tractable formulation for noise prediction. We therefore adopt
312 a straightforward latent estimation scheme that directly predicts the denoised Gaussian parameters.
313 Consequently, the training objective reduces to a photometric reconstruction loss. The generated
314 Gaussian parameters Θ are used to render novel views via 3DGS’s differentiable rasterization Kerbl
315 et al. (2023). The model is trained end-to-end using a photometric loss between rendered images
316 $\mathcal{R}(\theta)$ and ground truth target views \mathcal{I}_{gt} , combining ℓ_2 and LPIPS Zhang et al. (2018b) terms:
317

$$\mathcal{L}_{\text{photo}} = \|\mathcal{R}(\Theta) - \mathcal{I}_{\text{gt}}\|_2^2 + 0.05 \cdot \text{LPIPS}(\mathcal{R}(\Theta), \mathcal{I}_{\text{gt}}). \quad (11)$$

319 4 EXPERIMENTS AND DISCUSSIONS

320 4.1 EXPERIMENTAL SETTINGS

323 We train and evaluate our approach on two large-scale datasets, RealEstate10K Zhou et al. (2018)
324 and ACID Liu et al. (2021a). RealEstate10K contains home walkthrough videos from YouTube,
325

324
 325 **Table 1: Comparison of interpolated NVS.** We evaluate performance on the RealEstate10K and
 326 ACID datasets by rendering three novel interpolation views from two reference viewpoints, aver-
 327 aging across all scenes. The dataset’s training and testing split follows the identical protocol estab-
 328 lished by pixelSplat. Note that 3DGs-based methods render extremely fast (~ 500 FPS).

329 Method	330 RealEstate10K			331 ACID		332 Inference Time	
	333 PSNR \uparrow	334 SSIM \uparrow	335 LPIPS \downarrow	336 PSNR \uparrow	337 SSIM \uparrow	338 LPIPS \downarrow	339 (s)
pixelSplat Charatan et al. (2024)	25.89	0.858	0.142	28.14	0.839	0.150	0.104
MVSplat Chen et al. (2025b)	26.39	0.869	0.128	28.25	0.843	0.144	0.044
eFreeSplat Min et al. (2024)	26.45	0.865	0.126	28.30	0.851	0.140	0.061
TranSplat Zhang et al. (2025)	26.69	0.875	0.125	28.35	0.845	0.143	0.087
HiSplat Tang et al. (2025b)	27.21	0.881	0.117	28.75	0.853	0.133	0.510
CoDiffSplat (Ours)	27.56	0.888	0.114	28.77	0.855	0.133	0.089

335
 336 **Table 2: Comparison of extrapolated NVS on RealEstate10K.** We evaluate model performance
 337 on RealEstate10K by rendering three novel extrapolated views from two reference views, averaging
 338 across all scenes under identical training settings.

339 Method	340 PSNR \uparrow	341 SSIM \uparrow	342 LPIPS \downarrow
pixelSplat Charatan et al. (2024)	21.76	0.779	0.217
MVSplat Chen et al. (2025b)	21.92	0.787	0.199
TranSplat Zhang et al. (2025)	21.89	0.791	0.201
HiSplat Tang et al. (2025b)	22.01	0.794	0.191
CoDiffSplat (Ours)	24.33 (+2.32)	0.846 (+0.052)	0.152 (-0.039)

343 comprising 67,477 training scenes and 7,289 testing scenes. ACID consists of aerial nature footage
 344 captured by drones, split into 11,075 training scenes and 1,972 testing scenes. Both datasets are cal-
 345 briated via Structure-from-Motion (SfM) Schönberger & Frahm (2016), which provides per-frame
 346 camera intrinsics and extrinsics. Following prior works Charatan et al. (2024); Chen et al. (2025b),
 347 we use two context images as input and render three novel target views for each test scene. To assess
 348 the model’s comprehensive understanding of 3D scenes, we evaluate not only conventional interpo-
 349 lated NVS but also *extrapolated* NVS, where target viewpoints lie beyond the reference range. We
 350 adopt the training curriculum from pixelSplat Charatan et al. (2024), increasing the sampling interval
 351 of target views up to 45 frames before and after the reference views to accommodate extrapolation.
 352 To evaluate visual fidelity, we compare the images rendered by each method with the corresponding
 353 ground truth frames by computing a peak signal-to-noise ratio (PSNR), structural similarity index
 354 (SSIM) Wang et al. (2004), and perceptual distance (LPIPS) Zhang et al. (2018a). Please refer to the
 355 Appendix A.1 for implementation details. We also conduct zero-shot cross-dataset generalization
 356 experiments following the MVSplat protocol, which are discussed in detail in Appendix A.2.1.

357 4.2 MAIN RESULTS

381 4.2.1 INTERPOLATED NOVEL VIEW SYNTHESIS

382 Table 1 shows that CoDiffSplat achieves better performance on both RealEstate10K and ACID.
 383 While prior G-3DGs approaches differ in architectural details, they all rely on pixel-aligned Gaus-
 384 sian parameter estimation in a purely feed-forward manner, which often leads to blurred details,
 385 geometric distortions, or failure in occluded regions (Figure 3). In contrast, our conditional dif-
 386 fusion leverages CEA embeddings that effectively fuse semantic and geometric cues across views,
 387 enabling accurate reconstruction of fine structures and hidden surfaces. Importantly, the single-step
 388 refinement introduces only marginal computational overhead (0.089s per frame), remaining compet-
 389 itive with lightweight baselines while being substantially faster than HiSplat (0.510s). These results
 390 demonstrate that CoDiffSplat achieves a favorable balance between efficiency and reconstruction
 391 fidelity, with enhanced generalization in both indoor and outdoor sparse-view settings. Additional
 392 qualitative comparisons on RealEstate10K are provided in Appendix A.2.5.

393 4.2.2 EXTRAPOLATED NOVEL VIEW SYNTHESIS

394 To further assess the generalization ability of our model, we evaluate on extrapolated viewpoints that
 395 fall outside the range of input views. This constitutes a more challenging setting for NVS, leading
 396 to increased uncertainty and unobserved regions. As summarized in Table 2, CoDiffSplat achieves

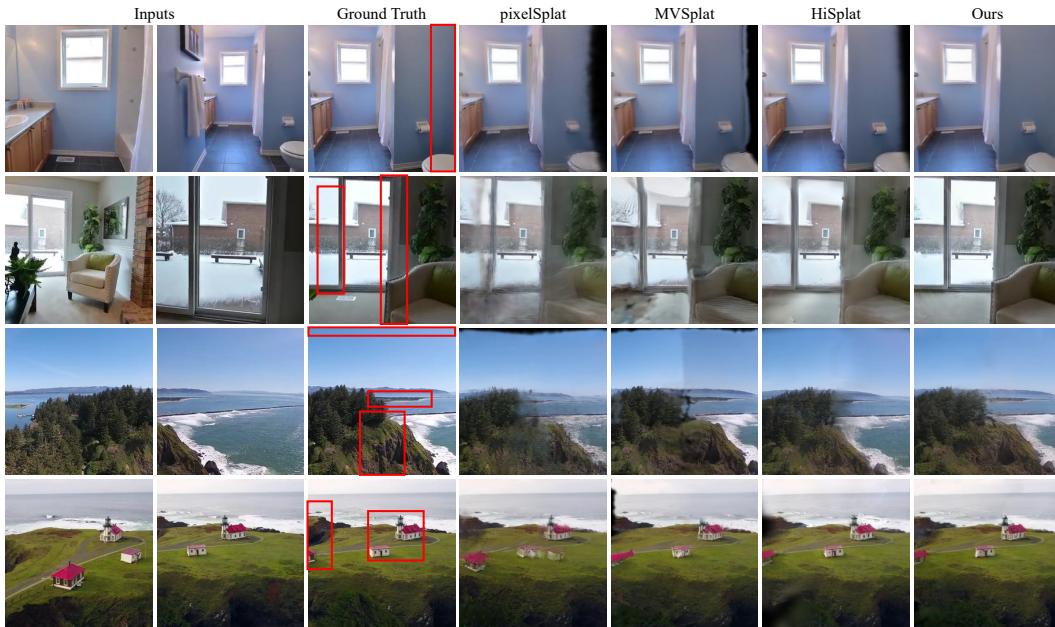


Figure 3: **Qualitative comparison of interpolated NVS.** The first two columns show sparse input views, while the third column presents the ground truth for the target interpolated view between them. CoDiffSplat better reconstructs fine details and occluded regions (highlighted in red) in both indoor (RealEstate10K, top two rows) and outdoor (ACID, bottom two rows) settings.



Figure 4: Qualitative results for extrapolated NVS on RealEstate10K. Baseline methods exhibit voids and distorted geometry in unseen regions, while CoDiffSplat reduces missing-geometry artifacts and better preserves boundary structures.

consistent improvements (+2.32 dB PSNR and 20% LPIPS reduction over the best baseline). As shown in Figure 4, pixel-aligned estimators (MVSplat and HiSplat) often produce “black holes” or geometric artifacts in unseen areas. In contrast, our method reduces missing-geometry artifacts and preserves fine structures. We attribute these gains to CEA conditioning and the single-step diffusion refinement, which regularizes Gaussian parameters with semantic cues beyond pixel-level correspondences.

4.3 ABLATION STUDIES

4.3.1 ABLATIONS ON DIFFUSION STRATEGY

We ablate our diffusion strategy from three aspects: (i) removing diffusion (*w/o Diffusion*), where the initial Gaussians are directly upsampled for rendering, (ii) replacing the DiT backbone with a

432
 433 **Table 3: Ablations of CoDiffSplat.** Models trained on RealEstate10K with two input views. *w/o*
 434 *Diffusion* removes the diffusion process and renders directly from the initial Gaussians. *w/ UNet*
 435 *Backbone* substitutes the DiT with a 2D UNet. *Full (10 steps)* introduces multi-step diffusion.

Setup	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o Diffusion	25.47 <i>(-2.09)</i>	0.849 <i>(-0.039)</i>	0.149 <i>(+0.035)</i>
w/ UNet Backbone	26.71 <i>(-0.85)</i>	0.873 <i>(-0.015)</i>	0.123 <i>(+0.009)</i>
Full (10 steps)	27.05 <i>(-0.51)</i>	0.876 <i>(-0.012)</i>	0.118 <i>(+0.004)</i>
Full (1 step)	27.56	0.888	0.114

440
 441 **Table 4: Ablations on conditional embeddings.** Models trained on RealEstate10K with two input
 442 views. CEA outperforms DINOv3, CLIP, and BLIP-2 under identical conditioning interfaces and
 443 training settings. Green indicates drops relative to the best model (CEA).

Setup	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Class Embedding	26.13 <i>(-1.43)</i>	0.859 <i>(-0.029)</i>	0.139 <i>(+0.025)</i>
CLIP Embedding	26.41 <i>(-1.15)</i>	0.867 <i>(-0.021)</i>	0.126 <i>(+0.012)</i>
BLIP Embedding	26.54 <i>(-1.02)</i>	0.870 <i>(-0.018)</i>	0.125 <i>(+0.011)</i>
CEA Embedding	27.56	0.888	0.114

444
 445 2D UNet (*w/ UNet Backbone*), and (iii) varying the number of diffusion steps (*Full (10 steps)*).
 446 As summarized in Table 3, excluding diffusion significantly degrades performance, indicating its
 447 necessity for effective refinement. Replacing DiT with a UNet also lowers accuracy, which we
 448 attribute to the inductive bias of convolutional backbones that enforce 2D grid structures, unsuitable
 449 for unordered 3D Gaussian primitives. Moreover, while classical diffusion models typically rely on
 450 multi-step denoising, we find that single-step diffusion yields better balance of fidelity and stability,
 451 outperforming the 10-step variant by +0.5dB PSNR. We hypothesize that additional steps introduce
 452 stochastic perturbations that accumulate into over-smoothing and hallucinated structures. Further
 453 qualitative comparisons are provided in Appendix A.2.2.

461 4.3.2 ABLATIONS ON CONDITIONAL EMBEDDINGS

462 To isolate the effect of conditioning, we compare the proposed Cross-view Entropy-Aware (CEA)
 463 embedding with three common alternatives: (i) a class embedding from a DINOv3-pretrained ViT,
 464 (ii) a CLIP image–text embedding, and (iii) a BLIP-2 pseudo-caption embedding. All embeddings
 465 are used under identical training and injection settings. As shown in Table 4, CEA achieves con-
 466 sistent gains over the alternatives; for example, it improves PSNR by +1.02 dB over the next best
 467 variant (BLIP) and reduces LPIPS by $\sim 8.8\%$. We attribute this gap to CEA’s cross-view fusion and
 468 entropy-aware weighting, which emphasize uncertain regions and thereby reduce missing-geometry
 469 artifacts and boundary erosion. For more qualitative discussions, please refer to the Appendix A.2.3.

473 5 CONCLUSION

474 In this work, we propose CoDiffSplat, a novel G-3DGS framework for novel view synthesis from
 475 sparse-view inputs. Unlike prior methods that rely on pixel-aligned Gaussian estimation, our
 476 approach leverages single-step conditional diffusion to refine the Gaussians, effectively compensat-
 477 ing and refining geometry in partially observed or uncertain scenes. To enhance both global se-
 478 mantic coherence and fine-grained detail awareness, we introduce a Cross-View Entropy-Aware
 479 module that produces semantically rich embeddings for conditioning the diffusion. Extensive ex-
 480 periments on multiple datasets demonstrate that CoDiffSplat significantly enhances reconstruction
 481 fidelity, particularly in occluded and unseen regions. These results highlight the promising poten-
 482 tial of diffusion-based paradigms in generalizable 3DGS. However, extreme extrapolation remains
 483 challenging, where geometric collapse can still occur. Furthermore, our current training is limited
 484 to RealEstate10K and ACID. Scaling to broader datasets or jointly training across diverse domains
 485 may unlock richer and more robust diffusion priors.

486 REFERENCES
487

488 Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and
489 Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields.
490 In *ICCV*, pp. 5835–5844, 2021. doi: 10.1109/ICCV48922.2021.00580.

491 Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-
492 nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, pp. 5460–5469, 2022. doi:
493 10.1109/CVPR52688.2022.00539.

494 Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstruc-
495 tured lumigraph rendering. In *SIGGRAPH*, pp. 425–432, 2001. doi: 10.1145/383259.383309.

496 Ziang Cao, Fangzhou Hong, Tong Wu, Liang Pan, and Ziwei Liu. Large-vocabulary 3d diffusion
497 model with transformer. In *ICLR*, 2024. URL <https://openreview.net/forum?id=q57JLSE2j5>.

498 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jegou, Julien Mairal, Piotr Bojanowski, and
499 Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pp. 9630–
500 9640, 2021. doi: 10.1109/ICCV48922.2021.00951.

501 David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. Pixelsplat: 3d gaus-
502 sian splats from image pairs for scalable generalizable 3d reconstruction. In *CVPR*, pp. 19457–
503 19467, 2024. doi: 10.1109/CVPR52733.2024.01840.

504 Minghao Chen, Roman Shapovalov, Iro Laina, Tom Monnier, Jianyuan Wang, David Novotny, and
505 Andrea Vedaldi. Partgen: Part-level 3d generation and reconstruction with multi-view diffusion
506 models. In *CVPR*, pp. 5881–5892, 2025a.

507 Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-
508 Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view
509 images. In *ECCV*, pp. 370–386, 2025b.

510 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
511 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit,
512 and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at
513 scale. In *ICLR*, 2021.

514 Gongfan Fang, Kunjun Li, Xinyin Ma, and Xinchao Wang. Tinyfusion: Diffusion transformers
515 learned shallow. In *CVPR*, pp. 18144–18154, 2025. doi: 10.1109/CVPR52734.2025.01691.

516 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-
517 nition. In *CVPR*, pp. 770–778, 2016.

518 Xianglong He, Junyi Chen, Sida Peng, Di Huang, Yangguang Li, Xiaoshui Huang, Chun Yuan,
519 Wanli Ouyang, and Tong He. Gvgen: Text-to-3d generation with volumetric representation. In
520 *ECCV*, pp. 463–479, 2025.

521 Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli,
522 Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. In *ICLR*, 2024.
523 URL <https://openreview.net/forum?id=s11U8vvsFF>.

524 Rasmus Jensen, Anders Dahl, George Vogiatzis, Engil Tola, and Henrik Aanæs. Large scale multi-
525 view stereopsis evaluation. In *CVPR*, pp. 406–413, 2014. doi: 10.1109/CVPR.2014.59.

526 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splat-
527 ting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. doi:
528 10.1145/3592433.

529 Jiahao Li, Hao Tan, Kai Zhang, Zexiang Xu, Fujun Luan, Yinghao Xu, Yicong Hong, Kalyan
530 Sunkavalli, Greg Shakhnarovich, and Sai Bi. Instant3d: Fast text-to-3d with sparse-view gen-
531 eration and large reconstruction model. In *ICLR*, 2024. URL [https://openreview.net/](https://openreview.net/forum?id=21DQLiH1W4)
532 [forum?id=21DQLiH1W4](https://openreview.net/forum?id=21DQLiH1W4).

540 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: bootstrapping language-image
 541 pre-training with frozen image encoders and large language models. In *ICML*, pp. 19730–19742,
 542 2023.

543

544 Chenguo Lin, Panwang Pan, Bangbang Yang, Zeming Li, and Yadong Mu. Diffsplat: Repurposing
 545 image diffusion models for scalable gaussian splat generation. In *ICLR*, 2025. URL <https://openreview.net/forum?id=ejZpoQkGK>.

546

547 Tianyu Lin, Zhiguang Chen, Zhonghao Yan, Weijiang Yu, and Fudan Zheng. Stable diffusion seg-
 548 mentation for biomedical images with single-step reverse process. In *MICCAI*, pp. 656–666,
 549 2024. doi: 10.1007/978-3-031-72111-3_62.

550

551 Andrew Liu, Ameesh Makadia, Richard Tucker, Noah Snavely, Varun Jampani, and Angjoo
 552 Kanazawa. Infinite nature: Perpetual view generation of natural scenes from a single image.
 553 In *ICCV*, pp. 14438–14447, 2021a. doi: 10.1109/ICCV48922.2021.01419.

554

555 Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo.
 556 Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pp. 10012–
 557 10022, 2021b.

558

559 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.

560

561 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and
 562 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, pp.
 563 405–421, 2020.

564

565 Zhiyuan Min, Yawei Luo, Jianwen Sun, and Yi Yang. Epipolar-free 3d gaussian splatting for
 566 generalizable novel view synthesis. In *NeurIPS*, 2024. URL <https://openreview.net/forum?id=i06tcLJEwA>.

567

568 Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system
 569 for generating 3d point clouds from complex prompts, 2022. URL <https://arxiv.org/abs/2212.08751>.

570

571 Suraj Patni, Aradhya Agarwal, and Chetan Arora. Ecodepth: Effective conditioning of diffusion
 572 models for monocular depth estimation. In *CVPR*, pp. 28285–28295, 2024.

573

574 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *CVPR*, pp.
 575 4195–4205, 2023.

576

577 Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural
 578 radiance fields for dynamic scenes. In *CVPR*, pp. 10313–10322, 2021. doi: 10.1109/CVPR46437.
 579 2021.01018.

580

581 Wentao Qu, Jing Wang, YongShun Gong, Xiaoshui Huang, and Liang Xiao. An end-to-end robust
 582 point cloud semantic segmentation network with single-step conditional diffusion models. In
 583 *CVPR*, pp. 27325–27335, 2025.

584

585 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 586 resolution image synthesis with latent diffusion models. In *CVPR*, pp. 10684–10695, 2022.

587

588 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomed-
 589 ical image segmentation. In *MICCAI*, pp. 234–241, 2015.

590

591 Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, pp.
 592 4104–4113, 2016. doi: 10.1109/CVPR.2016.445.

593

594 Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. MVDream: Multi-view
 595 diffusion for 3d generation. In *ICLR*, 2024. URL <https://openreview.net/forum?id=FUgrjq2pbB>.

594 Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose,
 595 Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel
 596 Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darct, Théo Moutakanni, Leonel Sentana,
 597 Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé
 598 Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3, 2025. URL <https://arxiv.org/abs/2508.10104>.

600 Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm:
 601 Large multi-view gaussian model for high-resolution 3d content creation. In *ECCV*, pp. 1–18,
 602 2025a.

603 Shengji Tang, Weicai Ye, Peng Ye, Weihao Lin, Yang Zhou, Tao Chen, and Wanli Ouyang. Hisplat:
 604 Hierarchical 3d gaussian splatting for generalizable sparse-view reconstruction. In *ICLR*, 2025b.
 605 URL <https://openreview.net/forum?id=SBzIbJobs8>.

606 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,
 607 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, volume 30, 2017.

608 Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error
 609 visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
 610 doi: 10.1109/TIP.2003.819861.

611 Christopher Wewer, Kevin Raj, Eddy Ilg, Bernt Schiele, and Jan Eric Lenssen. Latentsplat: Au-
 612 toencoding variational gaussians for fast generalizable 3d reconstruction. In *ECCV*, pp. 456–473,
 613 2025.

614 Jay Zhangjie Wu, Yuxuan Zhang, Haithem Turki, Xuanchi Ren, Jun Gao, Mike Zheng Shou, Sanja
 615 Fidler, Zan Gojcic, and Huan Ling. Difix3d+: Improving 3d reconstructions with single-step
 616 diffusion models. In *CVPR*, pp. 26024–26035, 2025.

617 Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P.
 618 Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. Reconfusion:
 619 3d reconstruction with diffusion priors. In *CVPR*, pp. 21551–21561, 2024.

620 Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, Fisher Yu, Dacheng Tao, and Andreas
 621 Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and
 622 Machine Intelligence*, 45(11):13941–13958, 2023.

623 Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and
 624 Marc Pollefeys. Depthsplat: Connecting gaussian splatting and depth. In *CVPR*, pp. 16453–
 625 16463, 2025. doi: 10.1109/CVPR52734.2025.01534.

626 Zeyu Yang, Zijie Pan, Chun Gu, and Li Zhang. Diffusion²: Dynamic 3d content generation via score
 627 composition of video and multi-view diffusion models. In *ICLR*, 2025.

628 Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d
 629 gaussians for high-fidelity monocular dynamic scene reconstruction. In *CVPR*, pp. 20331–20341,
 630 2024. doi: 10.1109/CVPR52733.2024.01922.

631 Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstruc-
 632 tured multi-view stereo. In *ECCV*, pp. 767–783, 2018.

633 Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from
 634 one or few images. In *CVPR*, pp. 4576–4585, 2021. doi: 10.1109/CVPR46437.2021.00455.

635 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-
 636 free 3d gaussian splatting. In *CVPR*, pp. 19447–19456, 2024. doi: 10.1109/CVPR52733.2024.
 637 01839.

638 Chuanrui Zhang, Yingshuang Zou, Zhuoling Li, Minmin Yi, and Haoqian Wang. Transplat: Gen-
 639 eralizable 3d gaussian splatting from sparse multi-view images with transformers. In *AAAI*, pp.
 640 9869–9877, 2025. doi: 10.1609/aaai.v39i9.33070.

648 Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable
649 effectiveness of deep features as a perceptual metric. In *CVPR*, pp. 586–595, 2018a. doi: 10.
650 1109/CVPR.2018.00068.

651 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
652 effectiveness of deep features as a perceptual metric. In *CVPR*, pp. 586–595, 2018b.

653 Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification:
654 learning view synthesis using multiplane images. *ACM Transactions on Graphics*, 37(4), 2018.
655 doi: 10.1145/3197517.3201323.

656

657

658

659

660

661

662

663

664

665

666

667

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701