

CGGS: Consistency–Augmented Geometric Gaussian Splatting for Ego-centric 3D Scene Generation

Zhenyu Sun, Xiaohan Zhang, Qi Liu[†], *Senior Member, IEEE*, and Huan Wang[†], *Senior Member, IEEE*

Abstract— Challenges remain in ego-centric 3D scene generation due to limited view overlap and the dominant influence of individual perspectives on scene interpretation. These factors hinder the creation of viewpoint-consistent and semantically aligned visual content, as well as the construction of accurate geometric structures. In this paper, we propose CGGS, a text-to-3D framework aiming to enhance 3D-content-awareness and address geometric distortions in ego-centric scene generation. Firstly, the Ego-centric Generator is proposed by fine-tuning a Multi-View Latent Diffusion Model with consistency-augmented loss to generate consistent, high-fidelity 2D content aligned with textual descriptions. Then, Layout Decorator leverages optical flow and point-track correspondence to estimate depth, therefore producing dense point clouds as coarse layouts from the ego-centric 2D priors. Building on this initialization, Geometric Refiner is proposed to enhance 3D Gaussian reconstruction via an entropy-based Mutual Information Depth Loss (MID) combined with a hierarchical optimization scheme for improving visual quality and geometric structure. Comprehensive experiments demonstrate that CGGS outperforms previous methods in generating coherent and accurate text-driven 3D scenes. Project page: <https://cggs-26.github.io/cggs26/>.

Index Terms—3D gaussian splatting, ego-centric generation, semantic alignment, global coherence.

I. INTRODUCTION

3D scene generation has recently gained significant attention, fueled by advances in generative models and strong image priors. In particular, generating 3D scenes from textual descriptions holds great promise for a wide range of real-world applications in AR/VR, robotics, and autonomous driving. With the rapid development of text-to-image generation [2], [3], [4], progress has been made toward text-to-3D generation. Latent Diffusion Models (LDMs) [4], [5] have been leveraged to optimize Neural Radiance Fields (NeRF) [6] via CLIP

embeddings [7], [8] or score distillation sampling (SDS) [9], [10], [11], [12]. However, these approaches often suffer from low rendering fidelity, multi-view inconsistencies, and limited scalability to scene-level 3D generation with fine-grained detail preservation. In contrast, the progressive expansion from text-driven 2D priors to 3D content [13], [14], [15], [16], [17], [18] enables high-quality synthesis but accumulates errors across iterations that induce style inconsistencies and structural discrepancies. Advance in 3D Gaussian splatting (3DGS) [19] and feed-forward architectures has significantly enhanced generalizability and the fidelity of complex 3D content representations, catalyzing holistic and realistic text-to-3D scene generation under both forward-facing and center-convergent viewpoint settings [20], [1], [21], [22], [23].

Despite these efforts, ego-centric 3D scene generation faces a fundamental dilemma regarding the choice of 2D priors: panoramic versus multi-view representations. While panoramic generation naturally ensures global continuity with a unified 360° field of view, the requisite equirectangular projection introduces severe geometric distortions—particularly near the poles—which fundamentally violate the pinhole camera assumption inherent in 3DGS and SfM pipelines. Such distortions inevitably lead to structural degradation and texture artifacts during the 3D lifting process, as illustrated in Fig. 1.

Conversely, multi-view generation synthesizes perspective images that are geometrically distortion-free and rich in local high-frequency details, offering a mathematically robust foundation for high-fidelity reconstruction. However, this paradigm inherently struggles with inter-view consistency due to the lack of a unified canvas.

To tackle with the aforementioned issues, we introduce **CGGS**, which unleashes the potential of latent diffusion models in text-image and image-image alignment, and learns the ego-centric 3D representation from the 2D images through a hierarchical 3D Gaussian optimization, as shown in Fig. 2. We leverage the real-world datasets Matterport3D [24], RealEstate-10k [25] and CO3Dv2 [26] to achieve domain-free, realistic 3D generation from textual descriptions. Following the settings of Correspondence-Aware Attention (CAA) for multi-view generation [27], we use Matterport3D [24] to fine-tune our ego-centric multi-view generator from stable diffusion model [28]. To enhance the semantic alignment and cross-view consistency, we introduce a consistency-augmented loss term as regularization to the LDM loss during the training of the CAA module. Building upon the synthesized views, a Flow-Depth Estimator is used to generate a dense point cloud as layout initialization. This approach can reconstruct a robust 3D structural layout of the scene from ego-centric 2D

This paper is supported by Young Scientists Fund of the National Natural Science Foundation of China (NSFC) (No. 62506305), Zhejiang Leading Innovative and Entrepreneur Team Introduction Program (No. 2024R01007), Key Research and Development Program of Zhejiang Province (No. 2025C01026), Scientific Research Project of Westlake University (No. WU2025WF003), Chinese Association for Artificial Intelligence (CAAI) & Ant Group Research Fund - AGI Track (No. 2025SCAAI-ANT-13). It is also supported by the research funds of the National Talent Program and Hangzhou Municipal Talent Program. It is also supported in part by the GJYC program of Guangzhou under Grant 2024D01J0081, in part by the ZJ program of Guangdong under Grant 2023QN10X455, and in part by the Fundamental Research Funds for the Central Universities under Grant 2025ZYGXZR053.

Zhenyu Sun, Xiaohan Zhang and Qi Liu are with the School of Future Technology, South China University of Technology, Guangzhou 511442, China (email: ftsunzhenyu@mail.scut.edu.cn; ftxiaohanzhang@mail.scut.edu.cn; dr-liuqi@scut.edu.cn).

Huan Wang is with the School of Engineering, Westlake University, Hangzhou 310030, China (email: wanghuan@westlake.edu.cn).

[†]Corresponding author.



Fig. 1. **Visualization of geometric distortions in panoramic generation** (using DreamScene360 [1] as an example). **1) Insufficient Text-Content Alignment:** Significant textual details are omitted in the generation, such as the absence of "mismatched frames" on the gallery wall despite being explicitly specified in the prompt. **2) Polar Geometric Distortions:** Due to the inherent nature of equirectangular projection, severe radial stretching and bending occur near the top and bottom boundaries (e.g., the warped ceiling and distorted sand), which violates perspective consistency. **3) Unreasonable Structural Artifacts:** The model fails to maintain physical continuity, most notably evidenced by the severed and floating tree trunks in the beach scene, as well as incoherent horizon lines that hinder valid 3D reconstruction.

priors, whereas conventional Structure-from-Motion methods (SfM) [29] typically struggle with such tasks. Based on the initial 3D layout, we further leverage the Mutual Information Depth Loss (MID) to refine the scene during 3D Gaussian optimization, combined with a hierarchical optimization strategy, maintaining rendering robustness. Collectively, our key contributions are as follows:

- **Ego-centric Generator:** A Multi-View Latent Diffusion Model is fine-tuned with our novel Consistency-Augmented Loss to produce ego-centric 2D priors that faithfully reflect the semantic intent of the textual descriptions and enhance cross-view consistency.
- **Layout Decorator:** A Flow-Depth Estimator guided by optical flow and point-track correspondences, transforming ego-centric 2D priors into a dense, coarse 3D layout. This approach addresses the inefficiencies and failure modes of direct SfM on ego-centric views.
- **Geometric Refiner:** Building upon the initial layouts, the hierarchical 3D Gaussian optimization supervised with Mutual Information Depth Loss (MID) iteratively sharpens structural details and enforces cross-view consistency, yielding geometrically precise and high-fidelity generation content.

II. RELATED WORK

2D Contents Generation. Generative Adversarial Networks (GANs) [30] were initially the leading method for image generation. Despite their success in creating 2D contents [31], [32], [33], [34], GANs struggle with textual prompt interpretation and dataset-specific biases. Diffusion models [35], [36], [37], [38] have emerged as a promising alternative for image

generation. They have built a strong foundation for customizing LDMs [2], [5], [4] to produce domain-specific contents from textual descriptions. Classifier-free guidance [39] is one such technique that has been employed to further enhance the fidelity to textual prompts. Moreover, several works [40], [41] further explore the application of perceptual loss on diffusion objectives to enhance image quality. In recent times, several works [42], [43], [44], [45], [46] have achieved panorama generation from texts, but faced challenges in dealing with image distortions caused by projection methods. In this study, we employ text-based multi-view generation to create ego-centric 2D priors, serving as systematic guidance for 3D scene generation from textual description.

Text-to-3D Generation. DreamFields [47] pioneered the integration of vision-language models like CLIP [7] with NeRF [6] to synthesize 3D objects from textual descriptions, followed by [48] utilizing mesh as the 3D representation. Subsequent works employ 2D diffusion models to refine 3D representations through Score Distillation Sampling (SDS) [9], [10], [11], [49], [12] or Score Jacobian Chaining [50]. ProlificDreamer [51] further optimized this approach by introducing Variational Score Distillation (VSD) to deal with the over-saturation problem. Recent progress in text-to-image generation [2], [3], [28] with diffusion models has laid the groundwork for text-to-3D generation using LDMs [5], [4], [52], [18]. Nevertheless, most of these works mainly concentrate on object-level generation.

For scene-level 3D generation, various attempts [13], [53], [54], [14], [55], [15], [56] synthesize 3D scenes through progressive expansion by merging image inpainting models [5] and monocular depth estimation models [57], [58], [59].




Fig. 2. With text prompts as input, **CGGS** employs three core components: the Ego-centric Generator creates ego-centric 2D priors, the Layout Decorator proposes additional scene details, and the Geometric Refiner further enhances the geometric structure and visual quality.

Despite improvements for egocentric scenarios, these methods still lead to geometric and textural artifacts due to inherent inpainting limitations and depth-alignment errors. More recent works [16], [1], [20], [21] consider to develop 3D-scene generation from a panorama, but require extra multi-view constraints to reduce single-viewpoint limitations, thus producing suboptimal results of ego-centric scene representation. Several works [22], [23] propose end-to-end 3D generation frameworks, directly decoding 3D Gaussians from latent space. However, these approaches are restricted by relatively brief trajectory lengths and pose challenges in generating 3D models from purely outward-facing camera trajectories. Therefore, developing 3D scene generation from ego-centric perspectives with cross-modal generalization remains a challenging issue.

3D Scene Representation. The development of 3D Gaussian Splatting (3DGS) [19] has ushered in a new era of efficient, high-fidelity 3D reconstruction and synthesis [6], [60], [19], significantly reducing the rendering time compared to NeRF-based methods [6], [61], [62], [63], [64]. In particular, optimization over 3D Gaussians initialized from point clouds has emerged as a dominant paradigm. These point clouds are most often obtained via Structure-from-Motion (SfM) pipelines [29] or by back-projecting pixels through monocular depth estimator (MDE) [57], [58], [59]. Several recent works [65], [66] have advanced SfM by fusing global consistency constraints and dense geometric priors. Moreover, targeting sparse-view reconstruction, various methods [67], [68], [69], [70], [71], [72] map pixel-aligned features into 3D Gaussians and optimize them end-to-end, albeit at the cost of substantial GPU resources and large-scale training data. To further refine geometry, recent works [73], [74], [75] have introduced depth-priors to enforce consistency between 3DGS parameters and geometric structure.

Nevertheless, these works often focus on the object-centric situation, neglecting that ego-centric 3D reconstruction suffers from scarce cross-view overlap and pronounced viewpoint bias, hindering both semantic consistency and geometric accu-

racy. Distinctively, our CGGS first leverages an optical flow-guided depth estimator to construct a dense point cloud from ego-centric 2D priors, then refines this initialization with a Mutual Information Depth Loss (MID) and a hierarchical 3D Gaussian optimization to deliver view-consistent geometric representations.

III. PRELIMINARY

Latent Diffusion Models (LDMs) [4], [5], [2] operate generation in a learned latent space. They first train an auto-encoder [76] to compresses high-dimensional data x into a latent representation $z = \mathcal{E}(x)$, from which x can be approximately reconstructed via $\hat{x} = \mathcal{D}(z)$. Then a denoising network ϵ_θ is employed to reverse a gradual noising process applied to the latents. The training loss minimizes the difference between added noise ϵ and predicted noise ϵ_θ :

$$\mathcal{L} = \mathbb{E}_{x, \epsilon \sim \mathcal{N}(0, 1), t} \left[\|\epsilon - \epsilon_\theta(z_t, c, t)\|_2^2 \right], \quad (1)$$

where the noised latent at timestep t is represented as $z_t = \sqrt{\alpha_t} \mathcal{E}(x) + \sqrt{1 - \alpha_t} \epsilon$, and c denotes optional conditioning information. During generation, the model starts from random noise $z_T \sim \mathcal{N}(0, 1)$, iteratively denoises it using ϵ_θ , and finally decodes the resulting latent into the output image.

Correspondence-Aware Attention (CAA). MVDiffusion [27] introduces CAA blocks to ensure 3D-aware multi-view consistency. The CAA block processes N feature maps concurrently. For a source feature map F , cross-attention is performed with $(N - 1)$ target feature maps F^l . The message M is calculated as follows:

$$M = \sum_l \sum_{t^l \in \mathcal{N}} \text{SoftMax} (W_Q \bar{F}(s) \cdot W_K \bar{F}^l(t_*^l)) W_V \bar{F}^l(t_*^l), \quad (2)$$

$$\bar{F}(s) = F(s) + \gamma(0), \quad \bar{F}^l(t_*^l) = F^l(t_*^l) + \gamma(s_*^l - s), \quad (3)$$

where W_Q , W_K , and W_V are query, key, and value matrices. The position encoding $\gamma(\cdot)$ is added to target features based




Fig. 3. **Pipeline of CGGS.** It primarily comprises the Consistency-Augmented MV-LDM as Ego-centric Generator, the Flow-Depth Estimator as Layout Decorator, and the 3D Gaussian Optimization combined with MID Loss and the hierarchical optimization strategy, serving as Geometric Refiner.

on the 2D displacement between s_*^l and s , which provides relative location within local neighborhoods. CAA blocks are integrated into the pre-trained stable diffusion UNet [28], with other modules frozen during training to maintain original model functionality [27]. The following loss is used for training CAA blocks:

$$\mathcal{L} := \mathbb{E} \left[\left\{ \mathbf{Z}_t^i = \mathcal{E}(\mathbf{x}^i) \right\}_{i=1}^N, \left\{ \epsilon^i \sim \mathcal{N}(0, I) \right\}_{i=1}^N, \mathbf{y}, t \right] \left[\sum_{i=1}^N \left\| \epsilon^i - \epsilon_\theta^i(\{\mathbf{Z}_t^i\}, t, \tau_\theta(\mathbf{y})) \right\|_2^2 \right], \quad (4)$$

where ϵ_i and $\mathbf{Z}_t^i$ respectively denote the noise and the latent representation of the i^{th} image $\mathbf{x}^i$, and y represents the condition. However, simply optimizing the KL divergence between the multi-view image forward process and the denoising process to align their distributions is insufficient for the model to balance multi-view consistency and text-image alignment. Panfusion [46] uses a dual-branch structure and adds a panoramic perspective to enhance text-image consistency. Our approach, though, adopts a more elegant design without extra 2D architecture to leverage panoramic information.

IV. METHODOLOGY

We set up our problem formulation as follows. We consider the ego-centric multi-view images as $\mathcal{X} = \{x_i\}_{i=1}^N$, with $\mathcal{C} = \{c_i\}_{i=1}^N$ as the corresponding camera trajectory. The $\mathcal{C}$ can be described as known variable with different settings, and we consider each x_i obeys the distribution $p\{x_i | x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N, \mathcal{C}\}$. The images are interpolated to $\mathcal{X}'$ with N' views ($N' > N$), and we further model the ego-centric 2D priors as a motion sequence $\mathcal{M}$, aiming at deriving the dense point cloud as 3D layouts. Additionally, each image x_i is rendered from the 3D Gaussian representation $\mathcal{G}$ under camera pose c_i : $x_i = \mathcal{R}(\mathcal{G}, c_i)$, where $\mathcal{R}$ denotes the rendering operator. CGGS addresses this task through three synergistic stages, as shown in Fig. 3. Initially, a Multi-View Latent Diffusion Model (MV-LDM) is employed as the Ego-centric Generator to learn the conditional image distribution over specified camera trajectories and textual prompts, reinforced by a consistency augmentation loss as discussed in Sec. IV-A. Directly conducting conventional SfM, e.g. [29], on

ego-centric views yields inefficient and low-quality layouts. To address this, we introduce a Flow-Depth Estimator as Layout Decorator to generate a dense coarse point cloud, as elaborated in Sec. IV-B. Finally, to achieve geometrically precise reconstructions, the 3D Gaussian optimization is refined with Mutual Information Depth loss and a hierarchical strategy, as presented in Sec. IV-C.

A. Ego-centric Generator

A crucial requirement of ego-centric generation is the holistic coherence and the semantic alignment. The suboptimal text-content semantic alignment primarily arises from cross-view inconsistencies, including inconsistent representations of the same content across views and artifacts from fragmented structures or physically implausible structures across views. We address these issues as follows.

Here we leverage the diffusion process from MVDiffusion [27] to simultaneously synthesize N multi-view images that collectively capture a 360-degree scene coverage. However, during the training of multi-view generation, each view may introduce distinct gradient directions due to variations in perspective, leading to conflicting optimization signals. These discrepancies can hinder the effective minimization of the score matching loss, complicating the learning of a coherent representation across views. Consequently, the capacity to accurately capture the underlying distribution diminishes, adversely affecting both inter-view consistency and alignment with the textual input.

To address the aforementioned issue, we propose a consistency-augmented-loss on the latent diffusion objectives when training the CAA blocks, which boosts the alignment of gradients across perspectives. Specifically, based on the Equ. (4), the score matching loss for each view x_i is augmented with a regularization term defined as:

$$\mathcal{L}_{\text{aug}} = \mathbb{E} \left[\left\{ \mathbf{Z}_t^i = \mathcal{E}(\mathbf{x}^i) \right\}_{i=1}^N, \left\{ \epsilon^i \sim \mathcal{N}(0, I) \right\}_{i=1}^N, \mathbf{y}, t \right] \left[\sum_{i=1}^N \left\| \phi(\epsilon^i) - \phi(\epsilon_\theta^i(\{\mathbf{Z}_t^i\}, t, \tau_\theta(\mathbf{y}))) \right\|_2^2 \right], \quad (5)$$

where ϕ denotes the multi-layer CNN ϕ with L layers, whose hierarchical composition and parameter sharing impose

a structured inductive bias that aligns per-view gradient updates into a common subspace, thereby serving as a conflict harmonizer. In our framework, ϕ is instantiated as a VGG-16 network [77]. We utilize the full stack of convolutional layers, therefore the $\mathcal{L}_{aug}$ integrates structural signals across multiple spatial scales in a unified manner.

Crucially, we employ He Initialization [78] and keep the weights frozen without loading any pre-trained parameters. This configuration transforms ϕ into a *structured random projection*: the hierarchical convolutional layers impose a strong inductive bias for capturing spatial frequencies and geometry, while the random initialization ensures the feature space remains isotropic and free from the semantic prejudices (e.g., class-specific dominance) inherent in pre-training on ImageNet [79]. This yields a feature space that jointly exhibits (i) non-linear hierarchical structure, (ii) spatially localized multi-scale representation, and (iii) unbiased, stationary feature geometry—properties that are not simultaneously achieved by handcrafted filters or Fourier transform-based methods. Because ϕ is identical and frozen across all views, it defines a stationary metric space for alignment. Applying the chain rule to the augmented loss, the gradient with respect to the denoising parameters θ is derived as:

$$\nabla_{\theta} \mathcal{L}_{aug} = \mathbb{E} \left[\left\{ \mathbf{z}_t^i = \mathcal{E}(\mathbf{x}^i) \right\}_{i=1}^N, \left\{ \epsilon^i \sim \mathcal{N}(0, I) \right\}_{i=1}^N, \mathbf{y}, t \right] \left[\sum_{i=1}^N \underbrace{2 (J_{\epsilon^i})^\top}_{\text{Backprop}} \underbrace{(J_{\phi})^\top}_{\text{Projection}} \underbrace{(\phi(\epsilon^i) - \phi(\epsilon^i))}_{\text{Feature Error}} \right], \quad (6)$$

where J_{ϕ} denotes the Jacobian of the harmonizer. Mathematically, the term $(J_{\phi})^\top$ acts as a subspace projection operator. Unlike standard noise-prediction gradients, which are often high-frequency and mutually orthogonal across different ego-centric views (leading to optimization conflicts), Equ. (6) enforces the update vectors of all views to be projected onto the shared singular subspace of J_{ϕ} . This alignment maximizes the cosine similarity between view-specific updates, yielding a consensus gradient that effectively counteracts incoherent noise while harmonizing the global geometric structure.

Thus, according to Equ. (4) and Equ. (5), the full training objective becomes:

$$\mathcal{L}_{total} = \mathcal{L} + \lambda_{aug} \mathcal{L}_{aug}. \quad (7)$$

This modification guarantees that each optimization step benefits from both precise score matching and a unified, multi-scale semantic prior that reduces inter-view gradient conflicts.

B. Layout Decorator

While contemporary off-the-shelf monocular depth estimators [57], [58], [59] achieve impressive per-pixel accuracy, they inherently suffer from scale ambiguity and inconsistent geometric deformations across independent ego-centric views. Direct back-projection from such inconsistent priors results in stratified and misaligned point clouds. To resolve this, we leverage flow and point correspondences as relative geometric constraints to enforce cross-view alignment. To fully exploit the 2D priors, we first interpolate the image sequences and

then upsample them according to predefined camera groupings. This step aims to enhance the image resolution and provide more detailed information. Then we treat the input image sequence $\mathcal{X}' = \{x_i\}_{i=1}^{N'}$ as a short video stream and model the motion sequences as $\mathcal{M} = \{\mathcal{X}', \mathcal{F}\}$. Here $\mathcal{F}$ represents the sequence of optical flow from off-the-shelf flow estimator [80], which is defined as:

$$\mathcal{F} = \{f_i : x_i \leftrightarrow x_{i+1}\}_{i=1}^{N'-1} \cup \{f_{N'} : x_{N'} \leftrightarrow x_1\}. \quad (8)$$

This establishes dense pixel-level correspondences across adjacent ego-centric views. Theoretically, optical flow can also be used to directly predict pixel correspondences across views:

$$p_t = p_s + f_{st}(p_s), \quad (9)$$

where p_t denotes the target pixels in image j , and p_s represents the source pixels in image i . When $j = i + 1$ for $i = 1, \dots, N' - 1$ and $j = 1$ for $i = N'$, it corresponds to the situation between neighboring viewpoints.

Relying solely on short-term flow can lead to cumulative drift. To mitigate this, we integrate long-term Point Tracks [81] to rectify flow divergence over extended temporal windows explicitly. Crucially, the depth estimation network functions as a geometric regularizer during this supervised process; it filters out high-frequency noise from the flow estimators while learning the dominant, consistent structural trends mandated by the point tracks. Thus, rather than propagating error, this joint optimization effectively harmonizes the scale and geometry of the initial depth priors, yielding the consistent per-frame depth maps $\mathcal{D} = \{d_i\}_{i=1}^N$.

Using the estimated depth d_i and the known intrinsic calibration K , the pixel p_s in the i^{th} view is back-projected into 3D space as $z_s \in \mathbb{R}^3$. Given the relative extrinsic transformation between the i^{th} and j^{th} cameras, the corresponding pixel coordinate p_t of p_s in the j^{th} image is obtained by:

$$\tilde{p}_t = K T_j T_i^{-1} \begin{pmatrix} z_s \\ 1 \end{pmatrix}, \quad (10)$$

where T_i and T_j are the 4×4 extrinsic matrices for views i and j , respectively. After dividing by its third component, the final pixel location $\hat{p}_t$ is produced. According to the results from Equ. (9), a correspondence loss can be defined as:

$$\mathcal{L}_{corr} = \|\hat{p}_t - p_t\|. \quad (11)$$

Minimizing this loss yields a supervised signal that iteratively enhances the depth estimation network, improving 3D layouts. This formulation jointly models the 3D layout from ego-centric observations, effectively addressing the limitations of traditional SfM pipelines [29] and mitigating the spatial misalignment issues commonly encountered in monocular depth estimation. Building upon the estimated depth maps $\mathcal{D} = \{d_i\}_{i=1}^{N'}$, the pixels are back-projected and merged into a unified coarse point cloud, yielding a consistent initial 3D scaffold that faithfully reflects the underlying scene geometry despite narrow baselines and frequent occlusions.

“This bathroom is a refreshing blend of spa-like tranquility and modern functionality, featuring a freestanding soaking tub beneath a large window. To the left, a sleek double vanity with vessel sinks offers ample space, while on the right, a glass-enclosed shower adds a touch of luxury.”



“A peaceful countryside farm at dusk, with golden fields of wheat stretching to the horizon. A red barn sits near a small farmhouse, and a tractor is parked in front. Cows graze in a nearby field, and a dirt road leads towards a distant village.”



Fig. 4. Generation results of CGGS for ego-centric multi-view priors, gaussian point clouds, novel view synthesis, and depth maps. Our method generates harmonious, domain-free 3D scenes from ego-centric views, highly aligned with complex textual descriptions.

C. Geometric Refiner

Based on the point clouds and structural information generated by the Layout Decorator, Geometric Refiner further incorporates depth-aware geometric structure into the original 3DGS framework [19]. This guides the cloning and splitting of Gaussian ellipsoids, enabling finer detail representation while preserving the overall geometric consistency. Previous works have introduced depth-based regularization to improve geometric accuracy, such as applying epipolar constraints [82] or Pearson correlation-based depth supervision [1]. To enhance the ability to capture complex depth relationships and to mitigate its tendency toward oversmoothing, we further model the geometric supervision process using a mutual information depth loss, which provides a more expressive and robust signal for learning fine-grained depth structures. Specifically, for the perspective of the i^{th} camera, the depth map d_i from the Layout Decorator is considered as the ground truth depth, with its rendered depth map d_{render}^i calculated from the differential rasterization of 3DGS. We begin by flattening the rendered depth map d_{render}^i and the ground-truth depth map d_i into one-dimensional vectors r and g , respectively, which we treat

as samples of continuous random variables D_{render} and D_{gt} . The mutual information between these variables is defined as:

$$\mathcal{I}(D_{render}; D_{gt}) = \sum \sum p(r, g) \log \frac{p(r, g)}{p(r)p(g)}, \quad (12)$$

where $p(r, g)$ denotes the joint probability density of r and g , and $p(r)$, $p(g)$ are their marginal densities. Based on the operation, the MID loss is expressed as:

$$\mathcal{L}_{MID} = 1 - \mathcal{I}(D_{render}; D_{gt}). \quad (13)$$

Therefore, the total training loss is shown as follows:

$$\mathcal{L}_{total}^{rec} = (1 - \lambda_{MID})\mathcal{L}_{rgb} + \lambda_{MID}\mathcal{L}_{MID}, \quad (14)$$

where λ_{MID} is a scalar weight balancing this term against RGB objectives in the overall training loss.

Crucially, distinct from the common scale-invariant depth loss (e.g., Pearson correlation), which implicitly assumes a linear relationship between predicted and reference depths, the proposed MID loss operates on the statistical dependency between distributions. While Pearson-based objectives effectively handle global-scale ambiguity, they are prone to suppressing high-frequency geometric details and are sensitive




Fig. 5. Qualitative comparison between CGGS with other baselines. Our CGGS produces multi-view images with rich detail and superior semantic coherence, showcasing domain-agnosticity. Our results outperform other works with an accurately detailed description and unified 3D consistency. Specifically, DreamScene360 generates visual results with less major content in the horizon field; While Director3D is capable of depicting the content described in text prompts, it is constrained by a limited field of view; LucidDreamr causes undesirable style transfer, wrong stitches between concepts, and inconsistent content, as highlighted in the red box.

to outliers, often resulting in over-smoothed structures. In contrast, by maximizing the mutual information, $\mathcal{L}_{\text{MID}}$ enforces a stricter alignment of the underlying structural entropy. This enables the optimization to capture complex, non-linear geometric correspondences and preserve sharp depth discontinuities, thereby yielding higher fidelity in ego-centric 3D reconstruction.

During the Gaussian optimization stage, Geometric Refiner introduces supplementary cameras to mitigate the bias arising from a limited set of viewpoints. Inspired by the spatial supplementary camera arrangement in HoloDreamer [21] and the multi-stage virtual camera strategy of DreamScene360 [1], we progressively expand the base cameras with additional cameras sharing intrinsic parameters, providing diverse viewpoints from different hierarchies. As training proceeds, both the number of supplementary cameras and their relative pose offsets increase, depicted as:

$$\begin{aligned} C_k &= C_{k-1} \cup \{ (R_i \Delta R_k, t_i + \Delta t_k) \}, \quad \text{for } k = 1, 2, \dots, n \\ C_0 &= \{ (R_i, t_i) \}_{i=1}^{N'}, \quad (\text{base cameras}) \end{aligned} \quad (15)$$

where each ΔR_k and Δt_k denotes the rotational and translational offsets at stage k , and n is the total number of the hierarchical expansion stages. The loss at stage k is computed by rendering from all cameras in C_k and comparing against the depth and RGB supervision signals. This gradual expansion encourages reconstruction robustness and enhances the integrity of the 3D Gaussian optimization.

V. EXPERIMENTS

A. Implementation Details

CGGS leverages Matterport3D [24] datasets, RealEstate10k [25] and CO3Dv2 [26] as real-world multi-view datasets. The text prompts for scene description imitate the sentence structure used in MVDiffusion [27] and are expanded using GPT-4 [83] across various scenarios. The number of views N in the Ego-centric Generator is set to 8, with a horizontal field of view (FOV) $\Theta = 90^\circ$ and a rotation angle $\theta = 45^\circ$, and the λ_{aug} is assigned to 0.5. For Layout Decorator, the interpolation number $N' = 20$, with FOV $\Theta' = 60^\circ$. We set the λ_{MID} to 0.05 and the hierarchical expansion epochs n to 3. More details are reported below.

Datasets. For multi-view generation, we leverage Matterport3D [24] to enable text-driven ego-centric generation. For real-world multi-view datasets, we leverage RealEstate10k [25] and CO3Dv2 [26] for accurate flow depth estimation at both scene and object levels. Matterport3D is a large-scale indoor RGB-D dataset comprising 194,400 images and 10,800 panoramic views across 90 building-scale scenes. RealEstate10k consists of 10 million frames extracted from approximately 80,000 video clips sourced from about 10,000 YouTube videos. CO3Dv2 is a collection of 1.5 million frames from 18,619 videos, covering 50 MS-COCO object categories for 3D reconstruction tasks.

Ego-centric Generator. We fine-tune the CAA blocks with the improved loss function, as shown in Equ. (7). It takes about 35 ~ 40 hours on four NVIDIA RTX A6000 GPUs, with a

batch size equal to 4 on each GPU. Ego-centric Generator generates 8 ego-centric multi-view images with a resolution of 512×512 . For the sequence interpolation mentioned in Sec. IV-B, we first fuse the image sequence $\mathcal{X}$ into a pseudo-panorama with the resolution of 4096×968 , then project the view to acquire 20 multi-view images with the size as 512×512 .

Layout Decorator. The Flow-Depth Estimator is trained with RealEstate-10k and CO3Dv2 on one single NVIDIA RTX A6000 GPU for about 25 hours. During the inference stage, it takes about 10 minutes to lift a specific scene as a dense point cloud from the multi-view priors.

Geometric Decorator. The hierarchical optimization strategy is implemented as a staged process. During the warm-up phase, the 3DGS optimization proceeds with the default configuration. At a predefined initialization iteration, a set of m additional cameras is introduced, sharing the same intrinsic parameters $\mathbf{K}$ and extrinsic parameters $\mathbf{E}_{ij}$, and arranged around each base camera with extrinsics $\mathbf{E}_i$, where $j = 1, 2, \dots, m$. At each subsequent stage k , four additional cameras are further deployed around each base camera. As described in Sec. IV-C, the transformations ΔR_k and Δt_k represent the relative pose offsets between each auxiliary camera C_{ij} at stage k and its corresponding base camera C_i . These offsets are progressively increased across stages to introduce more significant viewpoint variation, thereby enhancing structural representation and optimizing scene reconstruction. The total number of hierarchical stages n is set to be 3, as mentioned in Sec. IV-C.

In Sec. IV-C, the theoretical training loss in Equ. (14) for 3D Gaussian optimization can be expressed as

$$\mathcal{L}_{\text{total}}^{\text{rec}} = (1 - \lambda_{\text{MID}})\mathcal{L}_{\text{rgb}} + \lambda_{\text{MID}}\mathcal{L}_{\text{MID}}, \quad (16)$$

and the λ_{MID} is set to be 0.05, as mentioned in Sec. V-A. Specifically, the detailed reconstruction training loss can be represented as:

$$\begin{aligned} \mathcal{L}_{\text{total}}^{\text{rec}} &= \lambda_{\text{SSIM}}(1 - \text{SSIM}) + \lambda_{\text{MID}}\mathcal{L}_{\text{MID}} \\ &\quad + (1 - \lambda_{\text{SSIM}} - \lambda_{\text{MID}})L_1, \end{aligned} \quad (17)$$

where L_1 denotes the per-pixel absolute difference between the reference image and the rendered image. Here λ_{SSIM} is set to be 0.2 in practice. We disable the opacity reset process to accelerate convergence and maintain high rendering quality. The remaining configurations are consistent with 3DGS [19]. It takes about 3 minutes per scene for optimization.

B. Generation Results

Our generation results are presented in Fig. 4. The ego-centric 2D priors exhibit strong cross-view consistency in both style and content and are well aligned with text prompts, reflecting the effectiveness of the Ego-centric Generator. Moreover, our method supports vivid 3D scene generation from complex prompts, resulting in geometrically coherent structures that highlight the effectiveness of the Layout Decorator and Geometric Refiner.

TABLE I

QUANTITATIVE COMPARISON BETWEEN CGGS AND OTHER BASELINES. WE BENCHMARK OUR METHODS WITH OTHER BRILLIANT PRIOR WORKS ACROSS 24 SCENES, COVERING INDOOR AND OUTDOOR ENVIRONMENTS. WE REPORT THE EVALUATION IN METRICS THAT REFLECT BOTH GENERATION QUALITY AND RECONSTRUCTION QUALITY. THE RESULTS INDICATE THAT CGGS ACHIEVES THE BEST OVERALL PERFORMANCE, GENERATING A SEMANTICALLY CONSISTENT 3D SCENE WITH HIGH VISUAL QUALITY AND PROPER GEOMETRIC STRUCTURE.

Method	3D Representations	Generation Quality					Reconstruction Quality		
		CLIP Score $\uparrow$	Sharp $\uparrow$	Color $\uparrow$	Resolution $\uparrow$	Q-Align $\uparrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Text2Room [13]	Mesh	24.732	0.215	0.210	0.231	0.697	20.915	0.844	0.169
LucidDreamer [14]	3DGS	<u>25.736</u>	0.216	<u>0.211</u>	0.224	0.764	25.667	0.824	0.163
Director3D [22]	3DGS	24.996	0.221	0.225	<u>0.232</u>	0.754	-	-	-
DreamScene360 [1]	3DGS	25.022	<u>0.219</u>	0.204	0.239	<u>0.828</u>	32.587	0.969	<u>0.0477</u>
CGGS	3DGS	26.253	0.218	<u>0.211</u>	0.231	0.839	37.345	0.977	0.0193

C. Qualitative Comparison

We conduct qualitative comparisons between our method and several baselines using 3DGS [19] as scene representation: LucidDreamer [14], Director3D [22], and DreamScene360 [1]. As shown in Fig. 5, CGGS generates 3D scenes with abundant geometric details and produces high-fidelity novel views. LucidDreamer exhibits abrupt shifts in overall style and content elements, even under modest viewpoint variation. Director3D tends to produce simple layouts when faced with intricately detailed text, resulting in visuals of suboptimal quality. DreamScene360 demonstrates strong panoramic consistency at the global level but struggles to reproduce fine structural details and fully reflect the textual descriptions within its horizontal views. These comparisons indicate that our CGGS outperforms other methods in generating realistic 3D scenes with proper geometric details.

D. Quantitative Comparison

We tabulate relative metrics of generation quality and reconstruction quality to assess our CGGS with other baselines [13], [14], [22], [1]. For evaluation purposes, we employ GPT-4 [83] to generate a diverse set of text descriptions from multiple types of scenes from public datasets [84], [27], [85], varying in complexity, style, and semantic content. (1) Generation Quality: we adopt CLIP-Score [86] to assess the semantic alignment. Q-Align [87] and CLIP-IQA [88], including Sharp, Colorful, and Resolution, are used to measure image perceptual quality. (2) Reconstruction Quality: for assessment of rendering quality, we adopt PSNR, SSIM, and LPIPS as metrics commonly used in 3D reconstruction tasks. CLIP-Score (CS) is computed between the textual descriptions of each scene and the corresponding generated multi-view images, and the average is reported. Q-Align and CLIP-IQA, which include Sharp (SH), Colorful (CL), and Resolution (RS), are evaluated on the multi-view images generated for each scene, with scores averaged per scene. PSNR, SSIM, and LPIPS are used to assess view rendering quality in methods that generate scenes with reference images.

For Text2Room, since its rendered views contain large black artifacts and missing geometry as shown in Fig. 6, we only calculated its performance metrics on indoor scenes in Tab. I. For LucidDreamer, which requires both a textual description and a single initial image as input, we use the first image



Fig. 6. **Poor quality of the rendered views from Text2Room [13] on outdoor scenes.** There are large black artifacts and missing geometry.

from the multi-view sequence generated by CGGS as its initial input. As for Director3D, since it does not utilize any intermediate reference images, we do not evaluate its image reconstruction quality.

According to Tab. I, our CGGS method demonstrates a comprehensive advantage over existing methods across multiple dimensions. In terms of generation quality, CGGS achieves the highest CLIP Score (26.253), indicating superior semantic alignment between the generated 3D scenes and textual descriptions. While competitive methods like Director3D and DreamScene360 exhibit strengths in local attributes such as sharpness and resolution, CGGS delivers the best overall perceptual performance, as evidenced by its leading Q-Align score (0.839). Simultaneously, CGGS exhibits exceptional performance in reconstruction fidelity, delivering a PSNR of 37.345 and an LPIPS of 0.0193, which underscores its high structural accuracy. Text2Room and LucidDreamer suffer from incongruent stitching artifacts, where unrelated semantic concepts are incoherently merged. This leads to degraded rendering quality, particularly evident in suboptimal LPIPS scores. These results demonstrate the robustness and effectiveness of CGGS in generating semantically consistent and visually high-fidelity 3D scenes.

E. Ablation Study

Ego-centric Generator. In the architecture of CGGS, the Ego-centric Generator is crucial for enhancing semantic alignment and cross-view coherence. With $\mathcal{L}_{aug}$ removed, as illustrated in Fig. 7, the content layout becomes chaotic, cross-view consistency of texture details degrades, and implausible

“This bedroom is a serene fusion of minimalist and bohemian styles, showcasing a large platform bed with an oversized knit blanket, and soft, layered pillows. To the left of the bed, a stylish rattan chair invites relaxation, while to the right, a wooden dresser complements the warm, neutral palette.”

“A tropical beach at sunset, with golden sand stretching along a calm, turquoise ocean. Palm trees sway gently in the breeze, and small waves lap at the shore. A few distant islands are visible on the horizon, and a wooden dock extends out into the water.”



Fig. 7. **Ablation study on consistency-augmented loss $\mathcal{L}_{aug}$.** Without $\mathcal{L}_{aug}$, cross-view texture discrepancies become pronounced, with abrupt background artifacts (e.g., exposed ceilings in bedroom scenes) and physically implausible anomalies (e.g., floating, distorted trees on beaches) emerging.

TABLE II

ABLATION STUDIES OF EGO-CENTRIC GENERATOR ON CONSISTENCY-AUGMENTED LOSS $\mathcal{L}_{aug}$. WE REPORT THE TRAINING TIME OF THE EGO-CENTRIC GENERATOR AND THE VANILLA VERSION WITHOUT $\mathcal{L}_{aug}$. THE QUANTITATIVE RESULTS ILLUSTRATE THAT OUR PROPOSED $\mathcal{L}_{aug}$ ENHANCES THE TRAINING OF THE DIFFUSION PROCESS, THUS IMPROVING THE SEMANTIC ALIGNMENT AND PERCEPTUAL QUALITY.

Settings	Epoch	Training Time	Multi-View	Panorama	CLIP-Score $\uparrow$	Sharp $\uparrow$	Colorful $\uparrow$	Resolution $\uparrow$	Q-Align $\uparrow$
w/o $\mathcal{L}_{aug}$	8	$\sim 40h$	$\checkmark$		25.869	0.218	0.212	0.228	0.911
ours	8	$\sim 37h$	$\checkmark$		25.949	0.219	0.213	0.228	0.914
w/o $\mathcal{L}_{aug}$	8	$\sim 40h$		$\checkmark$	25.686	0.218	0.216	0.230	0.809
ours	8	$\sim 37h$		$\checkmark$	26.251	0.217	0.215	0.229	0.812

image artifacts emerge. We further provide quantitative results for the ablation study of Ego-centric Decorator in Tab. II. For each configuration, we compute metrics over both the multi-view and holistic scopes to assess semantic alignment and perceptual quality in cross-view and global settings. The results illustrate that the introduction of $\mathcal{L}_{aug}$ can enhance the ability to produce semantically aligned and perceptually coherent outputs in both cross-view and global settings.

Furthermore, it is noteworthy that when $\mathcal{L}_{aug}$ is removed, output quality at the inter-view level exceeds that at the global level. However, upon introducing consistency augmentation, the semantic and visual coherence at the global level surpasses that at the inter-view level. This finding further demonstrates that our design is instrumental in alleviating the gradient conflicts inherent to multi-view generation.

Layout Decorator. We compare Layout Decorator with conventional SfM methods, COLMAP [29], by directly substituting the block in CGGS while keeping other modules same. To ensure a fair comparison, we evaluate COLMAP both with and without the identical camera trajectories used in our method. Results in Tab. III demonstrate that our methods provide more reliable 3D structure for subsequent 3D Gaussian optimization. Under the relatively sparse-view settings with little overlap across views, conventional SfM methods tend to diverge during optimization or result in suboptimal spatial structure.

Geometric Refiner. The results of ablation studies on


TABLE III

ABLATION STUDY OF DIFFERENT SfM METHODS. OUR CONFIGURATION PROVIDES MORE ROBUST 3D LAYOUTS FOR SUBSEQUENT 3D GAUSSIAN OPTIMIZATION.

Method	Learning Type	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
COLMAP	progressive optimization	30.133	0.929	0.0860
COLMAP (w/ pose)	progressive optimization	30.362	0.928	0.0856
CGGS	feed-forward Network	37.345	0.997	0.0193

Geometric Refiner are reported in Tab. IV. Here, HO, PD, and MID briefly represent the hierarchical optimization, Pearson Depth loss, and MID loss, respectively. Configuration (e) is the full setting. The quantitative comparison indicates that relying solely on the depth supervision slightly degrades the rendering visual quality, due to the stricter structure constraints. In contrast, the combined application of hierarchical optimization and the MID loss yields substantial improvements in the geometric coherence of the 3D Gaussian primitives as well as in overall rendering performance. Qualitative results of ablation studies on Geometric Refiner are demonstrated in Fig. 8. The qualitative comparison demonstrates that incorporating MID and HO produces higher-fidelity visual content that most closely approximates the ground truth. It is worth noting that this hierarchical pipeline exhibits inherent robustness to upstream inconsistencies. The physical correspondences in the Layout Decorator and the statistical alignment driven by the MID loss in the Geometric Refiner effectively act as filtering

“A medieval castle perched on a hill, surrounded by a small village with stone houses and thatched roofs. A dirt road winds up to the castle gates, and villagers can be seen going about their daily tasks. The castle has tall towers, and banners flap in the wind from its walls.”



“This bedroom is a serene fusion of minimalist and bohemian styles, showcasing a large platform bed with an oversized knit blanket, and soft, layered pillows. To the left of the bed, a stylish rattan chair invites relaxation, while to the right, a wooden dresser complements the warm, neutral palette.”



GT

Ours (MID+HO)

MID

PD+HO

PD

w/o (MID+HO)

Fig. 8. Ablation studies of Geometric-Refiner on MID loss and hierarchical optimization. Here we demonstrate the qualitative comparison between the ground truth with different settings, including MID+HO, MID, PD+HO, PD, and w/o (MID+HO). The comparison covers both indoor and outdoor scenes. Our design of Geometric-Refiner provides the most accurate texture recovery, with fewer blurred blocks than other settings.

mechanisms, capable of repairing minor semantic or textural discrepancies present in the rendered ego-centric 2D priors.

TABLE IV
ABLATION STUDIES OF GEOMETRIC REFINER ON MID LOSS AND HIERARCHICAL OPTIMIZATION. WE EXPLORE THE EFFECTIVENESS OF OUR PROPOSED MID LOSS AND HIERARCHICAL OPTIMIZATION, AND COMPARE OUR DEPTH LOSS WITH THE CONVENTIONAL PEARSON DEPTH LOSS. THE RESULTS BELOW SUPPORT OUR DESIGN CHOICES.

	HO	PD	MID	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
(a)				36.087	0.972	0.0231
(b)		✓		35.883	0.971	0.0235
(c)			✓	35.982	0.972	0.0225
(d)	✓	✓		36.251	0.971	0.0238
(e)	✓		✓	37.345	0.997	0.0193

F. Generalization Analysis

Additional Qualitative Results. Although we simply fine-tune the CAA blocks in Ego-centric Generator with an exclusively indoor scene dataset, Matterport3D, the preservation of the other parts from stable diffusion enables our model to keep the capability of cross-domain generation. And the data from RealEstate-10k and CO3Dv2 provide powerful prior knowledge of real-world structure, therefore promoting the derivation of initial layout from Layout Decorator, further improving the 3D structure representation and visual details in the Geometric Refiner process.

We provide more generation results in Fig. 9 and Fig. 10. The content covers various indoor and outdoor scenarios. These visual results demonstrate the impressive ability of CGGS on novel view synthesis from complex textual descrip-

“This craft room is a creative mix of whimsical and functional design, featuring a large table covered in colorful supplies and projects in progress. To the left, open shelving displays organized bins of materials, while on the right, a cozy nook with a comfy chair is perfect for inspiration.”



“This game room is an exciting blend of retro and modern features, showcasing a classic pool table in the center surrounded by plush seating. To the left, a vintage arcade machine stands proudly, while on the right, a bar cart is stocked with snacks and drinks for entertaining.”



“A sprawling futuristic metropolis, with sleek silver buildings, flying cars zipping through the air, and massive digital billboards projecting advertisements. The city is illuminated by a soft blue glow, and a hovering train moves along elevated tracks high above the streets.”



“A vast, empty desert under the scorching midday sun. Sand dunes stretch endlessly in every direction, with only a few sparse cacti and dry shrubs dotting the landscape. In the distance, heat waves blur the horizon, and a solitary camel slowly walks across the sand.”



“This entryway is a welcoming fusion of rustic and contemporary elements, featuring a large wooden bench with colorful cushions. To the left, a set of hooks hangs for coats and bags, while on the right, a stylish console table displays a vase of fresh flowers and a mirror that reflects light.”



“This kitchen is a charming blend of rustic and modern, featuring a large reclaimed wood island with marble countertop, a sink surrounded by cabinets. To the left of the island, a stainless-steel refrigerator stands tall. To the right of the sink, built-in wooden cabinets painted in a muted.”



“A dense tropical rainforest with towering trees, thick underbrush, and a small waterfall cascading into a pool. The air is humid, and sunlight barely penetrates the canopy. Exotic birds and wildlife can be heard in the distance, and the vibrant green foliage creates a sense of overwhelming life.”



“A vast arctic landscape dominated by towering icebergs and glaciers. The icy surface stretches out in every direction, with cracks revealing deep blue crevices. A lone polar bear can be seen in the distance, walking across the frozen terrain under a clear, cold sky.”



Fig. 9. Additional generation results from CCGS. Our work can generate richly detailed, high-fidelity scenes with considerable diversity while ensuring cohesive semantic content and a harmonized visual style that faithfully reflects even the most intricate textual descriptions.

“An underwater coral reef teeming with colorful fish, sea turtles, and vibrant corals. Sunlight filters down through the clear blue water, casting dancing shadows on the ocean floor. Schools of fish swim among the coral, and a group of dolphins can be seen in the distance.”



“A bustling downtown area on a sunny afternoon, with people walking along busy streets, modern glass skyscrapers towering overhead, and a park with trees and benches in the background. Cars and buses pass by, and a fountain can be seen in the center of a square.”



A snowy landscape in the midst of winter, with thick blankets of snow covering the ground and tall pine trees. A cozy log cabin with smoke rising from the chimney sits in the middle of the scene, and a frozen lake sparkles under a clear blue sky. Footprints lead from the cabin into the forest.”



“A vibrant city at night, with neon signs glowing in bright colors, reflecting on wet streets after a recent rain. Skyscrapers with lit-up windows reach high into the sky, while streetlights cast soft glows on sidewalks with few people. A subway entrance is visible in the distance.”



Fig. 10. Additional generation results from CGGS. Our work can generate richly detailed, high-fidelity scenes with considerable diversity while ensuring cohesive semantic content and a harmonized visual style that faithfully reflects even the most intricate textual descriptions.

TABLE V
QUANTITATIVE COMPARISON ON OUT-OF-DOMAIN SCENES. OUR CGGS ACHIEVES COMPETITIVE IMAGE QUALITY AND STRONG SEMANTIC ALIGNMENT, AND IT SURPASSES PREVIOUS METHODS ON OVERALL GENERATIVE PERFORMANCE, INDICATED BY THE HIGHEST Q-ALIGN SCORE.

Method	CLIP Score $\uparrow$	CLIP-IQA			Q-Align $\uparrow$
		Sharp $\uparrow$	Color $\uparrow$	Resolution $\uparrow$	
Text2Room	24.14	0.210	0.220	0.224	0.669
LucidDreamer	25.84	0.207	<u>0.222</u>	0.221	0.676
Director3D	22.95	<u>0.218</u>	0.225	0.223	0.475
DreamScene360	25.36	0.217	0.216	0.232	0.791
CGGS	<u>25.80</u>	0.220	0.214	<u>0.227</u>	0.820

tions and provide compelling evidence for the effectiveness of our full pipeline. Specifically, they demonstrate that the Ego-centric Generator produces semantically faithful 2D priors, while the Layout Decorator successfully reconstructs coarse 3D layouts even in challenging ego-centric conditions. The Geometric Refiner further enhances geometric fidelity by enforcing structural consistency through hierarchical optimization, ultimately resulting in high-quality, cross-view-consistent reconstructions.

Domain-Free Generation & Generalization. An important observation is that the diverse examples presented in our paper include many out-of-domain scenarios not seen during training, such as urban scenes with pedestrians and vehicles, as well as scenes dominated by humans, animals, and other biological entities. We further conduct quantitative evaluation on generation performance between these methods, with 4 out-of-domain scenes—camels in the desert, people in the market, creatures underwater, pedestrians in urban daytime. As illustrated in Tab. V, CGGS competes with existing methods in image quality, and surpasses them in semantic alignment and the overall generation quality. This demonstrates that CGGS exhibits notable generalization ability, generating semantically faithful 3D scenes from textual descriptions with accurate details and consistent spatial structures. These results indicate that our proposed CGGS effectively tackles the challenges of ego-centric 3D generation and reconstruction, while preserving strong diversity and generalization performance.

Limitations and Future Work. Despite strong performance in text-driven ego-centric 3D synthesis, CGGS remains limited by per-scene optimization, which increases computation time and hampers generalization. Future work will explore dynamic scene synthesis under ego-centric settings and visual language navigation within the generated environments via LLMs.

Broader Impacts. This paper introduces a framework aiming at improving semantic alignment and overall perceptual coherence in ego-centric 3D scene generation. Due to the inherent viewpoint characteristics of the generated scenes, the system holds considerable promise for synthetic data generation in the autonomous driving domain. Moreover, the proposed reconstruction enhancements may substantially improve the fidelity of imagery reconstructed from vehicle-mounted cameras. Such synthetic reconstructions may exacerbate privacy infringements or be weaponized to create convincing digital forgeries, posing substantial risks to societal security.

VI. CONCLUSION

In this work, we propose **CGGS**, a novel text-to-3D framework designed to address the challenges of ego-centric 3D scene generation. Ego-centric Generator is introduced to synthesize high-fidelity 2D content aligned with textual prompts. Layout Decorator is proposed to produce reliable initialization from ego-centric 2D priors, then Geometric Refiner is leveraged for further optimization, using 3D Gaussians as scene representation. We validate the effectiveness of our methods via comprehensive experiments. We believe that CGGS not only enhances current ego-centric 3D scene generation approaches but also paves the way for more diverse and realistic 3D content creation and virtual environment exploration.

REFERENCES

- [1] S. Zhou, Z. Fan, D. Xu, H. Chang, P. Chari, T. Bharadwaj, S. You, Z. Wang, and A. Kadambi, "Dreamscene360: Unconstrained text-to-3d scene generation with panoramic gaussian splatting," in *ECCV*, 2024.
- [2] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen, "Glide: Towards photorealistic image generation and editing with text-guided diffusion models," *arXiv preprint arXiv:2112.10741*, 2021.
- [3] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with clip latents," *arXiv preprint arXiv:2204.06125*, 2022.
- [4] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. L. Denton, K. Ghasemipour, R. G. Lopes, B. K. Ayan, T. Salimans *et al.*, "Photorealistic text-to-image diffusion models with deep language understanding," in *NeurIPS*, 2022.
- [5] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High resolution image synthesis with latent diffusion models," in *CVPR*, 2022.
- [6] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," in *ECCV*, 2020.
- [7] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, and *et al.*, "Learning transferable visual models from natural language supervision," in *ICML*, 2021.
- [8] A. Jain, B. Mildenhall, J. T. Barron, P. Abbeel, and B. Poole, "Zero-shot text-guided object generation with dream fields," in *CVPR*, 2022.
- [9] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall, "Dreamfusion: Text-to-3d using 2d diffusion," in *ICLR*, 2023.
- [10] C.-H. Lin, J. Gao, L. Tang, T. Takikawa, X. Zeng, X. Huang, K. Kreis, S. Fidler, M.-Y. Liu, and T.-Y. Lin, "Magic3d: High-resolution text-to-3d content creation," in *CVPR*, 2023.
- [11] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich, "Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation," in *CVPR*, 2023.
- [12] Y. Shi, P. Wang, J. Ye, L. Mai, K. Li, and X. Yang, "Mvdream: Multi-view diffusion for 3d generation," in *ICLR*, 2024.
- [13] L. Höllein, A. Cao, A. Owens, J. Johnson, and M. Nießner, "Text2room: Extracting textured 3d meshes from 2d text-to-image models," in *ICCV*, 2023.
- [14] J. Chung, S. Lee, H. Nam, J. Lee, and K. M. Lee, "Luciddreamer: Domain-free generation of 3d gaussian splatting scenes," *arXiv preprint arXiv:2311.13384*, 2023.
- [15] J. Zhang, X. Li, Z. Wan, C. Wang, and J. Liao, "Text2nerf: Text-driven 3d scene generation with neural radiance fields," *IEEE Transactions on Visualization and Computer Graphics*, 2024.
- [16] H.-X. Yu, H. Duan, J. Hur, K. Sargent, M. Rubinstein, W. T. Freeman, F. Cole, D. Sun, N. Snavely, J. Wu *et al.*, "Wonderjourney: Going from anywhere to everywhere," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [17] J. Shriram, A. Trevithick, L. Liu, and R. Ramamoorthi, "Realdreamer: Text-driven 3d scene generation with inpainting and depth diffusion," in *3DV*, 2025.
- [18] Y. Zhou, D. Ye, H. Zhang, X. Xu, H. Sun, Y. Xu, X. Liu, and Y. Zhou, "Recurrent diffusion for 3d point cloud generation from a single image," *IEEE Transactions on Image Processing*, vol. 34, pp. 1753–1765, 2025.
- [19] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, 2023.
- [20] W. Li, F. Cai, Y. Mi, Z. Yang, W. Zuo, X. Wang, and X. Fan, "Scenedreamer360: Text-driven 3d-consistent scene generation with panoramic gaussian splatting," *arXiv preprint arXiv:2408.13711*, 2024.
- [21] H. Zhou, X. Cheng, W. Yu, Y. Tian, and L. Yuan, "Holodreamer: Holistic 3d panoramic world generation from text descriptions," *arXiv preprint arXiv:2407.15187*, 2024.
- [22] X. Li, Z. Lai, L. Xu, Y. Qu, L. Cao, S. Zhang, B. Dai, and R. Ji, "Director3d: Real-world camera trajectory and 3d scene generation from text," in *NeurIPS*, 2024.
- [23] Y. Yang, J. Shao, X. Li, Y. Shen, A. Geiger, and Y. Liao, "Prometheus: 3d-aware latent diffusion models for feed-forward text-to-3d scene generation," *arXiv preprint arXiv:2412.21117*, 2024.
- [24] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3D: Learning from RGB-D data in indoor environments," in *3DV*, 2017.
- [25] T. Zhou, R. Tucker, J. Flynn, G. Fyffe, and N. Snavely, "Stereo magnification: Learning view synthesis using multiplane images," *ACM Transactions on Graphics*, vol. 37, no. 4, pp. 65:1–65:12, 2018.

- [26] J. Reizenstein, R. Shapovalov, P. Henzler, L. Sbordone, P. Labatut, and D. Novotny, "Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction," in *ICCV*, 2021.
- [27] S. Tang, F. Zhang, J. Chen, P. Wang, and F. Yasutaka, "Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion," *arXiv preprint arXiv:2307.01097*, 2023.
- [28] StabilityAI, "Stable diffusion 2," <https://huggingface.co/stabilityai/stable-diffusion-2>, 2023.
- [29] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *CVPR*, 2016.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, 2014.
- [31] B. Li, X. Zhang, M. Cai, and Y. Hou, "Tree-gan: Generative adversarial network for 3d tree modeling," in *ECCV*, 2018.
- [32] H. Lin, Y. Yu, and Q. Chen, "Blockgan: A generative model for 3d block-based urban structures," in *CVPR*, 2018.
- [33] Y. Liu, L. Bao, Z. Yang, L. Yuan, and H. Li, "l-gan: Generative adversarial networks for 3d shape generation," in *CVPR*, 2017.
- [34] F. Wu, C. Li, J. Wu, J. Johnson, and L. Fei-Fei, "Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling," in *ECCV*, 2016.
- [35] Y. Song and S. Ermon, "Generative modeling by estimating gradients of the data distribution," in *NeurIPS*, 2019.
- [36] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *NeurIPS*, 2020.
- [37] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," in *ICLR*, 2021.
- [38] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, "Score-based generative modeling through stochastic differential equations," in *ICLR*, 2021.
- [39] J. Ho and T. Salimans, "Classifier-free diffusion guidance," *arXiv preprint arXiv:2207.12598*, 2022.
- [40] T. Berrada, P. Astolfi, M. Hall, M. Havasi, Y. Benchetrit, A. Romero-Soriano, K. Alahari, M. Drozdal, and J. Verbeek, "Boosting latent diffusion with perceptual objectives," *arXiv preprint arXiv:2411.04873*, 2025.
- [41] S. Lin and X. Yang, "Diffusion model with perceptual loss," *arXiv preprint arXiv:2401.00110*, 2025.
- [42] O. Bar-Tal, L. Yariv, Y. Lipman, and T. Dekel, "Multidiffusion: Fusing diffusion paths for controlled image generation," *arXiv preprint arXiv:2302.08113*, 2023.
- [43] J. Li and M. Bansal, "Panogen: Text-conditioned panoramic environment generation for vision-and-language navigation," *arXiv preprint arXiv:2305.19195*, 2023.
- [44] Q. Zhang, J. Song, X. Huang, Y. Chen, and M. Yu Liu, "Diffcollage: Parallel generation of large content with diffusion models," in *CVPR*, 2023.
- [45] H. Wang, X. Xiang, Y. Fan, and J.-H. Xue, "Customizing 360-degree panoramas through text-to-image diffusion models," in *WACV*, 2024.
- [46] C. Zhang, Q. Wu, C. Cruz Gambardella, X. Huang, D. Phung, W. Ouyang, and J. Cai, "Taming stable diffusion for text to 360° panorama image generation," in *CVPR*, 2024.
- [47] e. a. Zhang, "Zero-shot text-guided object generation with dream fields," *CVPR*, 2022.
- [48] X. Z. Wei, K. H. Lee, Y. H. Lee, J. Lee, and A. Y. K. Y. Cheung, "Text2mesh: Text-driven 3d mesh generation," in *CVPR*, 2022.
- [49] R. Chen, Y. Chen, N. Jiao, and K. Jia, "Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation," in *ICCV*, 2023.
- [50] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich, "Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation," in *CVPR*, 2023.
- [51] Z. Wang, C. Lu, Y. Wang, F. Bao, C. Li, H. Su, and J. Zhu, "Prolific-dreamer: High-fidelity and diverse text-to-3d generation with variational score distillation," in *NeurIPS*, 2024.
- [52] F. Hong, J. Tang, Z. Cao, M. Shi, T. Wu, Z. Chen, T. Wang, L. Pan, D. Lin, and Z. Liu, "3dtopia: Large text-to-3d generation model with hybrid diffusion priors," *arXiv preprint arXiv:2403.02234*, 2024.
- [53] J. Zhang, Y. Xu, W. Wang, J. Yang, Y. Shen, X. Li, L. Xie, and F. Yu, "Rgbd2: Generative scene synthesis via incremental view inpainting using rgbd diffusion models," in *CVPR*, 2023.
- [54] H. Ouyang, K. Heal, S. Lombardi, and T. Sun, "Text2immersion: Generative immersive scene with 3d gaussians," *arXiv preprint arXiv:2312.09242*, 2023.
- [55] R. Fridman, A. Abecasis, Y. Kasten, and T. Dekel, "Scenescape: Text-driven consistent scene generation," *arXiv preprint arXiv:2302.01133*, 2023.
- [56] H. Wang, Y. Liu, Z. Liu, W. Wang, Z. Dong, and B. Yang, "Vistadream: Sampling multiview consistent images for single-view scene reconstruction," *arXiv preprint arXiv:2410.16892*, 2024.
- [57] R. Ranftl, K. Lasinger, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 3, pp. 1624–1637, 2022.
- [58] S. F. Bhat, R. Birkel, D. Wofk, P. Wonka, and M. Müller, "Zoedepth: Zero-shot transfer by combining relative and metric depth," *arXiv preprint arXiv:2302.12288*, 2023.
- [59] L. Yang, B. Kang, Z. Huang, X. Xu, J. Feng, and H. Zhao, "Depth anything: Unleashing the power of large-scale unlabeled data," in *CVPR*, 2024.
- [60] M. You, M. Guo, X. Lyu, H. Liu, and J. Hou, "Learning a locally unified 3d point cloud for view synthesis," *IEEE Transactions on Image Processing*, vol. 32, pp. 5610–5622, 2023.
- [61] A. Yu, V. Ye, M. Tancik, and A. Kanazawa, "pixelnerf: Neural radiance fields from one or few images," in *CVPR*, 2021.
- [62] W. Bian, Z. Wang, K. Li, J. Bian, and V. A. Prisacariu, "Nope-nerf: Optimising neural radiance field with no pose prior," in *CVPR*, 2023.
- [63] G. Metzger, E. Richardson, O. Patashnik, R. Giryes, and D. Cohen-Or, "Latent-nerf for shape-guided generation of 3d shapes and textures," in *CVPR*, 2023.
- [64] J. Wynn and D. Turmukhambetov, "Diffusernerf: Regularizing neural radiance fields with denoising diffusion models," in *CVPR*, 2023.
- [65] L. Pan, D. Baráth, M. Pollefeys, and J. L. Schönberger, "Global structure-from-motion revisited," in *ECCV* 2024, 2024.
- [66] C. Smith, D. Charatan, A. Tewari, and V. Sitzmann, "Flowmap: High-quality camera poses, intrinsics, and depth via gradient descent," *arXiv preprint arXiv:2404.15259*, 2024.
- [67] J. Tang, Z. Chen, X. Chen, T. Wang, G. Zeng, and Z. Liu, "Lgm: Large multi-view gaussian model for high-resolution 3d content creation," in *ECCV*, 2024.
- [68] K. Zhang, S. Bi, H. Tan, Y. Xiangli, N. Zhao, K. Sunkavalli, and Z. Xu, "Gs-irm: Large reconstruction model for 3d gaussian splatting," in *ECCV*, 2024.
- [69] Y. Xu, Z. Shi, W. Yifan, H. Chen, C. Yang, S. Peng, Y. Shen, and G. Wetzstein, "Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation," in *ECCV*, 2024.
- [70] Y. Chen, H. Xu, C. Zheng, B. Zhuang, M. Pollefeys, A. Geiger, T.-J. Cham, and J. Cai, "Mvsplat: Efficient 3d gaussian splatting from sparse multi-view images," in *ECCV*, 2024.
- [71] D. Charatan, S. Li, A. Tagliasacchi, and V. Sitzmann, "Pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction," in *CVPR*, 2024.
- [72] S. Szymanowicz, C. Rupprecht, and A. Vedaldi, "Splatter image: Ultra-fast single-view 3d reconstruction," in *CVPR*, 2024.
- [73] J. Li, J. Zhang, X. Bai, J. Zheng, X. Ning, J. Zhou, and L. Gu, "Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization," in *CVPR*, 2024.
- [74] Z. Zhu, Z. Fan, Y. Jiang, and Z. Wang, "Fsgs: Real-time few-shot view synthesis using gaussian splatting," in *ECCV*, 2024.
- [75] H.-X. Yu, H. Duan, C. Herrmann, W. T. Freeman, and J. Wu, "Wonderworld: Interactive 3d scene generation from a single image," *arXiv:2406.09394*, 2024.
- [76] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *ICLR*, 2014.
- [77] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *ICLR*, 2015.
- [78] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *ICCV*, 2015.
- [79] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and F.-F. Li, "Imagenet: A large-scale hierarchical image database," in *CVPR*, 2009.
- [80] Z. Teed and J. Deng, "Raft: Recurrent all-pairs field transforms for optical flow," in *ECCV*, 2020.
- [81] N. Karaev, I. Rocco, B. Graham, N. Neverova, A. Vedaldi, and C. Rupprecht, "Cotrack: It is better to track together," in *ECCV*, 2024.
- [82] Y. Zheng, Z. Jiang, S. He, Y. Sun, J. Dong, H. Zhang, and Y. Du, "Nexusgs: Sparse view synthesis with epipolar depth priors in 3d gaussian splatting," *arXiv preprint arXiv:2503.18794*, 2025.
- [83] OpenAI, "Gpt-4 technical report," OpenAI, Tech. Rep., 2024. [Online]. Available: <https://platform.openai.com/docs/gpt-4>

- [84] A. Knapitsch, J. Park, Q.-Y. Zhou, and V. Koltun, “Tanks and temples: Benchmarking large-scale scene reconstruction,” in *ACM Transactions on Graphics (TOG)*, 2017.
- [85] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, “Scannet: Richly-annotated 3d reconstructions of indoor scenes,” in *CVPR*, 2017.
- [86] J. Hessel, A. Holtzman, M. Forbes, R. L. Bras, and Y. Choi, “CLIPScore: a reference-free evaluation metric for image captioning,” in *EMNLP*, 2021.
- [87] H. Wu, Z. Zhang, W. Zhang, C. Chen, L. Liao, C. Li, Y. Gao, A. Wang, E. Zhang, W. Sun *et al.*, “Q-align: Teaching lmms for visual scoring via discrete text-defined levels,” in *ICML*, 2024.
- [88] J. Wang, K. C. K. Chan, and C. C. Loy, “Exploring clip for assessing the look and feel of images,” in *AAAI*, 2023.