

GVGEN: Text-to-3D Generation with Volumetric Representation

Xianglong He^{1,2*}, Junyi Chen^{1,3*}, Sida Peng⁴, Di Huang¹, Yangguang Li⁵,
Xiaoshui Huang¹, Chun Yuan^{2†}, Wanli Ouyang^{1,6}, and Tong He^{1†}

¹Shanghai AI Laboratory ²Tsinghua Shenzhen International Graduate School

³Shanghai Jiao Tong University ⁴Zhejiang University

⁵VAST ⁶The Chinese University of Hong Kong

Abstract. In recent years, 3D Gaussian splatting has emerged as a powerful technique for 3D reconstruction and generation, known for its fast and high-quality rendering capabilities. Nevertheless, these methods often come with limitations, either lacking the ability to produce diverse samples or requiring prolonged inference times. To address these shortcomings, this paper introduces a novel diffusion-based framework, GVGEN, designed to efficiently generate 3D Gaussian representations from text input. We propose two innovative techniques: (1) *Structured Volumetric Representation*. We first arrange disorganized 3D Gaussian points as a structured form GaussianVolume. This transformation allows the capture of intricate texture details within a volume composed of a fixed number of Gaussians. To better optimize the representation of these details, we propose a unique pruning and densifying method named the Candidate Pool Strategy, enhancing detail fidelity through selective optimization. (2) *Coarse-to-fine Generation Pipeline*. To simplify the generation of GaussianVolume and empower the model to generate instances with detailed 3D geometry, we propose a coarse-to-fine pipeline. It initially constructs a basic geometric structure, followed by the prediction of complete Gaussian attributes. Our framework, GVGEN, demonstrates superior performance in qualitative and quantitative assessments compared to existing 3D generation methods. Simultaneously, it maintains a fast generation speed (~ 7 seconds), effectively striking a balance between quality and efficiency. Our project page is <https://gvgen.github.io/>.

Keywords: Text-to-3D Generation · Feed-forward Generation · 3D Gaussians

1 Introduction

The development of 3D models is a pivotal task in computer graphics, garnering increased attention across various industries, including video game design, film production, and AR/VR technologies. Among the different aspects of 3D

* Equal Contribution.

† Corresponding Authors.

modeling, generating 3D models from text descriptions has emerged as a particularly intriguing area of research due to its accessibility and ease of use. Various methods [17, 33, 35] have been proposed to handle the task. Still, it continues to present difficulties owing to the ambiguity of texts and intrinsic domain gap between text description and corresponding 3D assets.

Previous text-to-3D approaches can be broadly classified into two categories: optimization-based generation [9, 35, 43, 50] and feed-forward generation [20, 25, 33, 34, 42]. Optimization-based methods have become rather popular recently, due to the rapid development of text-to-image diffusion models [13, 38]. These methods usually optimize 3D objects conditioned on texts or images through Score Distillation Sampling (SDS) [35], distilling rich knowledge from 2D image generation models. Despite yielding impressive results, optimization-based methods face the Janus problem [35], manifesting as multiple faces or over-saturation problems. Additionally, the optimization of a single object can be prohibitively time-consuming, requiring extensive computational effort. Contrarily, feed-forward approaches strive to generate 3D assets directly from text descriptions, thus sidestepping the Janus problem and significantly hastening the generation process. Our work is closely related to feed-forward-based methods. However, feed-forward methods that utilize multi-view generation models often create lower-resolution 3D assets than their multi-view image counterparts. Moreover, models directly generating 3D objects from texts often encounter difficulties with semantics when complex prompts are used.

Different from previous feed-forward-based methods like [20] that follow a text-2D-3D framework, our method proposes to generate 3D assets via directly learning 3D representation. In this study, we introduce an innovative and streamlined coarse-to-fine generation pipeline, GVGEN, for generating 3D Gaussians directly from text descriptions. Leveraging the highly expressive and fast-rendering capabilities of 3D Gaussians, our method achieves not only promising results but also maintains rapid text-to-3D generation and rendering. As shown in Fig. 1, our method consists of two stages: GaussianVolume fitting and text-to-3D generation. In the first stage, we introduce GaussianVolume, a structured volumetric form composed of 3D Gaussians. Achieving this is challenging due to the sparse and unstructured nature of optimizing original Gaussians. To address this, we introduce a novel Candidate Pool Strategy for pruning and densification. The approach allows for fitting high-quality volumetric representation of Gaussians, rather than unordered points, making the generation process more conducive for a diffusion-based framework, as is utilized in the following step.

Despite the GaussianVolume establishing a structured volumetric framework that integrates seamlessly with existing diffusion pipelines, the intrinsic complexity of rich features of 3D Gaussians presents significant challenges. Specifically, capturing the distribution of a vast amount of training data effectively becomes difficult, resulting in hard convergence for the diffusion model. Addressing these challenges, we partition the text-to-3D generation into two steps: coarse geometry generation and Gaussian attributes prediction. To be more specific, in the first step, we employ a diffusion model to generate the coarse geometry of ob-

jects, termed the Gaussian Distance Field (GDF) - an isotropic representation outlining the proximity of each grid point to the nearest Gaussian point’s center. Following this, the generated GDF, in conjunction with text inputs, is processed through a 3D U-Net-based model to predict the attributes of GaussianVolumes, ensuring enhanced control and model convergence.

To the best of our knowledge, this is the first study to directly feed-forward generate 3D Gaussians from texts, exploring new avenues for rapid 3D content creation and applications. Our main contributions are summarized as follows:

- We introduce GaussianVolume, a structured, volumetric form consisting of 3D Gaussians. Through the innovative Candidate Pool Strategy for pruning and cloning, we accommodate high-quality GaussianVolume fitting within a fixed volumetric resolution. This framework seamlessly integrates with existing generative networks, leveraging the inherent advantages of 3D Gaussians for explicit and efficient representation.
- We propose GVGEN, an efficient text-to-3D coarse-to-fine generation pipeline that first generates geometry volume and then predicts detailed Gaussian attributes, better controlling diverse geometry and appearances of generated assets. GVGEN achieves a fast speed (~ 7 seconds) compared with baseline methods, effectively striking a balance between quality and efficiency.
- Compared with existing baselines, GVGEN demonstrates competitive capabilities in both quantitative and qualitative aspects.

2 Related Works

2.1 Text-to-3D Generation

Generating 3D objects conditioned on texts has become a challenging yet prominent research area in recent years. Previous approaches [17, 31] utilize CLIP [36] as a prior for 3D asset optimization but lacked realism and fidelity. With the rise of text-to-image generative models [13, 38], Dreamfusion [35] leverages Score Distillation Sampling (SDS) to generate diverse 3D objects, drawing on the rich prior knowledge embedded in these models. Subsequently, some works [7, 22, 26, 27] focus on modeling and learning multi-modal attributes of objects (e.g. colors, albedo, normals, and depths), to enhance consistency. ProlificDreamer [43] and LucidDreamer [23] introduce novel losses to improve the quality. Some studies [16, 24, 39, 40, 51] explore predicting multi-views for objects, followed by using feed-forward methods or SDS-based optimization to generate 3D objects. Moreover, integrating the recently proposed 3D Gaussian Splatting [6, 19], works [9, 41, 48] lead to improvements in convergence speed for text-to-3D object generation. However, these methods still require significant generation time (\sim hours) per object or face 3D inconsistency of generated objects.

As the emergence of large-scale 3D datasets [11, 12], Feed-forward models are trained to reduce the generation time. Point-E [33] and Shap-E [18], trained on millions of 3D assets, generate point clouds and neural radiance fields respectively. 3D VADER [34] trains auto-decoders to fit latent volumes and uses

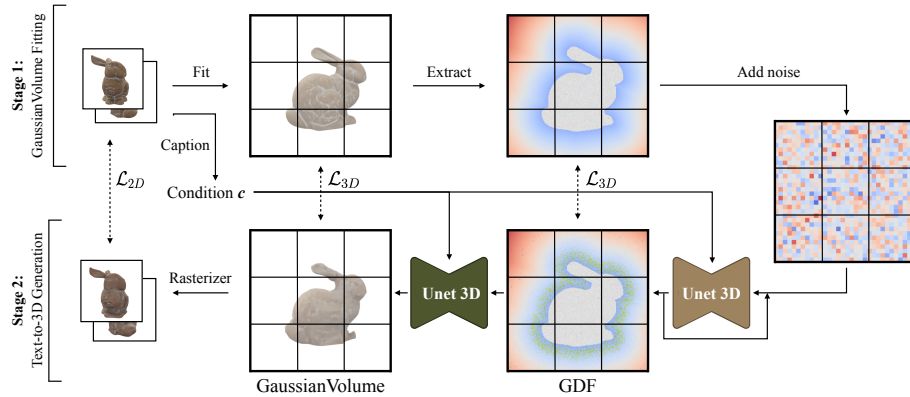


Fig. 1: Overview of GVGEN. Our framework comprises two stages. In the data pre-processing phase, we fit GaussianVolumes (Sec. 3.1) and extract coarse geometry Gaussian Distance Field (GDF) as training data. For the generation stage (Sec. 3.2), we first generate GDF via a diffusion model, and then send it into a 3D U-Net to predict attributes of GaussianVolumes.

them to train the diffusion model. Similarly, VolumeDiffusion [42] trains an efficient volumetric encoder to produce training data. Despite these advancements, text-to-3D Gaussian generation remains largely unexplored, primarily due to the complexity of organizing disorganized Gaussians. Our work introduces a coarse-to-fine pipeline to feed-forward generate 3D Gaussians from texts, by first generating the object’s coarse geometry and then predicting its explicit attributes, pioneering in the direct generation of 3D Gaussians from text descriptions.

2.2 Differentiable 3D Representation

Since the proposal of Neural Radiance Field [30], various differentiable neural rendering methods [2, 3, 5, 32] have emerged, demonstrating remarkable capabilities in scene reconstruction, novel view synthesis, and 3D generation tasks. Instant-NGP [32] employs feature volumes for accelerations by querying features only at the corresponding spatial positions. Works [3, 5] decompose features into lower dimensions for faster training and less storage. However, they still have some shortcomings compared with point-based rendering methods [4, 47] in the aspects of rendering speed and explicit manipulability.

In recent research, 3D Gaussian Splatting [19] has received widespread attention. It adopts anisotropic Gaussians to represent scenes, achieving real-time rendering and facilitating downstream tasks like 3D generation [37, 46, 49], scene editing [8] and dynamic scene rendering [44]. We introduce a strategy to fit 3D Gaussians as a structured volumetric form called GaussianVolume, which allows for good integration with the existing generation pipeline, leveraging the advantages of 3D Gaussians and achieving fast generation and rendering speeds.




Fig. 2: Illustration of GaussianVolume Fitting. We organize a fixed number of 3D Gaussians in a volumetric form, termed GaussianVolume. By using position offsets to express slight movements from grid points to Gaussian centers, we can capture the details of objects. The proposed Candidate Pool Strategy (CPS) (Sec. 3.1) enables effective pruning and densification with a pool storing pruned points.

3 Methodology

Our text-to-3D generation framework (Fig. 1), GVGEN, is delineated into two pivotal stages: GaussianVolume fitting (Sec. 3.1) and text-to-3D generation (Sec. 3.2). Initially, in the GaussianVolume fitting stage, we propose a structured, volumetric form of 3D Gaussians, termed GaussianVolume. We fit the GaussianVolume as training data for the generation stage. This stage is crucial for arranging disorganized point-cloud-like Gaussian points as a format more amenable to neural network processing. To address this, We use a fixed number (i.e. fixed volume resolution for the volume) of 3D Gaussians to form our GaussianVolume (Sec. 3.1), thereby facilitating ease of processing. Furthermore, we introduce the Candidate Pool Strategy (CPS) for dynamic pruning, cloning, and splitting of Gaussian points to enhance the fidelity of fitted assets. Our GaussianVolume maintains high rendering quality with only a small number (32,768 points) of Gaussians.

In the phase of generation Sec. 3.2, we first use a diffusion model conditioned on input texts to generate the coarse geometry volume, termed Gaussian Distance Field (GDF), which represents the geometry of generated objects. Subsequently, a 3D U-Net-based reconstruction model utilizes the GDF and text inputs to predict the attributes of the final GaussianVolume, thus achieving the generation of detailed 3D objects from text descriptions.

3.1 Stage1: GaussianVolume Fitting

The original 3D Gaussian Splatting offers fast optimization and real-time rendering. However, its lack of structure makes it challenging for 3D neural networks to process effectively. Directly generating 3D Gaussians through networks involves integrating existing point cloud networks, treating Gaussians as carriers of semantically rich features. However, previous work [45] has pointed out the

Algorithm 1: GaussianVolume Fitting with Candidate Pool Strategy

Input: The number of Gaussian points N with assigned positions p , total iteration times T , multi-view images I with camera poses V , prune and densify conditions τ_p, τ_d

Output: GaussianVolume G , a volumetric form composed of a fixed number of 3D Gaussians

```

1   $t \leftarrow 0$ ; // Iteration Time
2   $G \leftarrow \text{Initialize}(N, p)$ ; // Initialize Gaussian Points
3   $P \leftarrow \emptyset$ ; // Initialize Candidate Pool
4  while  $t < T$  do
5     $\hat{I} \leftarrow \text{Rasterize}(G, V)$ ; // Rendering
6     $\mathcal{L} \leftarrow \text{Loss}(I, \hat{I})$ ; // Loss
7     $G \leftarrow \text{Optimize}(\nabla \mathcal{L})$ ; // Update Gaussian Attributes
8    if  $\text{IsRefinementIteration}(t)$  then
9       $G_p \leftarrow \text{PrunePoints}(G, \tau_p)$ ; // Determine Pruned Points
10      $G, P \leftarrow \text{AddPointsToPool}(G_p, G, P)$ ; // Make  $G_p$  "Deactivated"
11      $G_d \leftarrow \text{DensifyPoints}(G, \tau_d)$ ; // Determine Densified Points in  $G$ 
12      $G_{new} \leftarrow \text{FindPointsInPool}(G_d, P)$ ; // Find Added Points in  $P$ 
13      $G, P \leftarrow \text{RemovePointsFromPool}(G_{new}, G, P)$ ; // Make  $G_{new}$ 
        "Activated"
14   end
15   if  $\text{EndRefinementIteration}(t)$  then
16      $\text{ReleasePool}(G, P)$ ; // Make All Points "Activated"
17   end
18 end

```

difficulty of directly generating point clouds with only 3-dimensional color features. Extending such methods, especially for high-dimensional Gaussians, may make the learning process very challenging. We also attempted to extend previous work [29] using point cloud diffusion to generate 3D Gaussians, while the results were disappointing.

Recent works [46, 52] propose hybrid representations combining 3D Gaussians with triplane-NeRF, utilizing point clouds for geometry and triplane features for texture querying, to finally decode 3D Gaussians. While these methods demonstrate effectiveness, the reconstructed results are heavily constrained by the accuracy of predicted point clouds and lose the direct operability advantage of 3D Gaussians. In light of these challenges, we introduce GaussianVolume, which is a volume composed of a fixed number of 3D Gaussians. This innovation facilitates the processing of 3D Gaussians by existing generation and reconstruction models, retaining the efficiency benefits of Gaussians. The illustration of the fitting process can be found in Fig. 2.

GaussianVolume In 3D Gaussian Splatting [19], features of 3D Gaussians G include: a center position $\mu \in \mathbb{R}^3$, covariance matrix Σ , color information $c \in \mathbb{R}^3$ (when SH order=0) and opacity $\alpha \in \mathbb{R}$. To better optimize the covariance matrix,

Σ is analogous to describing the configuration of an ellipsoid via a scaling matrix $S \in \mathbb{R}^3$ and a rotation matrix $R \in \mathbb{R}^{3 \times 3}$, satisfying:

$$\Sigma = RSS^T R^T. \quad (1)$$

In the implementation, S, R could be stored as a 3D vector $s \in \mathbb{R}^3$ and a quaternion $q \in \mathbb{R}^4$, respectively. With these differentiable features, 3D Gaussians can be easily projected to 2D splats and employ neural point-based α -blending technique to render 2D-pixel colors. An efficient tile-based rasterizer is used for fast rendering and backpropagation.

In this work, we represent each object with a volume $V \in \mathbb{R}^{C \times N \times N \times N}$ composed of 3D Gaussians, which is trained from multi-view images of the object. Here, C, N represent the number of feature channels and volume resolution respectively, and the number of 3D Gaussians is N^3 . We follow the convention and rendering methods in original works, but use position offsets $\Delta\mu$ to express slight movements between the position p of each grid point in the volume and the center μ of the Gaussian point it represents:

$$\mu = p + \Delta\mu. \quad (2)$$

During the fitting phase, we only apply the backward operation on the offsets $\Delta\mu$, which allows the expression of more fine-grained 3D assets and also imposes a restriction to form better structures. Due to the fixed number of Gaussians, we cannot directly apply the original strategy for densification control. Instead, we propose the Candidate Pool Strategy for effectively pruning and cloning.

Candidate Pool Strategy We can not directly apply the original interleaved optimization/density strategy to move 3D Gaussians, due to the dynamical changes in the number of Gaussians not suitable for a fixed number of Gaussians. Densifying or pruning Gaussians freely is not allowed since Gaussians are bijective with assigned grid points. A naive way to avoid the problem is to only adjust the offsets $\Delta\mu$ relying on gradient back-propagation. Unfortunately, we experimentally found that the movement range of Gaussian centers becomes largely limited without an optimization/density strategy, which leads to lower quality and inaccurate geometry (see Fig. 7(a)). For this reason, we propose a novel strategy (Algorithm 1) to densify and prune the fixed number of Gaussians. The key point to our strategy is storing pruned points in a candidate pool P for later densification.

We initially align Gaussian centers, μ , with assigned positions, p and set offsets $\Delta\mu = 0$ (Line 2 in Algorithm 1). The candidate pool P is initialized as an empty set (Line 3). This pool stores "deactivated" points, referring to pruned points during optimization – they are not involved in the forward and backward process. We optimize each asset for T iterations. Following original 3DGS [19], thresholds τ_p, τ_d on view-space position gradients are used for pruning and densifying every fixed iterations (*IsRefinementIteration*(t) in Line 8).



Fig. 3: Visualization of GaussianVolume Fitting. The rendering results demonstrate excellent reconstruction performance of GaussianVolumes.

During the refining process, once Gaussian points G_p are pruned with pre-defined threshold τ_p , they are added into candidate pool P , making them "deactivated" for rendering and optimization (Lines 9-10). For densification, we first determine points G_d that should be densified in the "activated" set of Gaussians G . Then newly added points G_{new} are selected from candidate pool P via some criterion (e.g., we use the nearest point within a certain distance $\epsilon_{offsets}$ to the densified points G_d). The corresponding coordinate offsets for the added points are calculated. They become "activated" for both forward and backward processes and are subsequently removed from the candidate pool P (Lines 11-13). At the end of optimization/density, all points in the candidate pool are reintroduced into the optimization for further refinement (Lines 15-17).

This strategy ensures that Gaussian points can adaptively move during the optimization process, to represent more intricate object shapes. Simultaneously, the resulting structured volumetric form maintains its physical meaning, demonstrating a balance between adaptability and well-defined structure.

Training Loss The final loss for supervision is the original loss used in 3D Gaussian Splatting adding a regularization loss:

$$\mathcal{L}_{offsets} = Mean(\text{ReLU}(|\Delta\mu - \epsilon_{offsets}|)), \quad (3)$$

where $\epsilon_{offsets}$ is a hyper-parameter, to restrict the center of 3D Gaussians not too far from their corresponding grid points,

$$\mathcal{L}_{fitting} = \lambda_1 \mathcal{L}_1 + \lambda_2 \mathcal{L}_{SSIM} + \lambda_3 \mathcal{L}_{offsets} \quad (4)$$

After training, the points are sorted in spatial order according to volume coordinates as training data for the generation stage. Once the training is over, 2D images of target objects could be rendered at an ultra-fast speed since each GaussianVolume is rather lightweight. More implementation details can be found in the supplemental materials.

3.2 Stage2: Text-to-3D Generation

As discussed in work [42], 3D volume space is high-dimensional, which potentially poses challenges in training the diffusion model. Empirically, we found that even though we form 3D Gaussians as a volumetric structure, it still takes too long for model convergence when training with a large amount of data. Moreover, learning the data distribution of GaussianVolumes from diverse categories via a single diffusion model causes it to collapse to the single-step reconstruction, leading to the degradation of generation diversity. To overcome these challenges, we introduce a coarse-to-fine pipeline that first generates coarse geometry volumes (Gaussian Distance Field, GDF) and then predicts the attributes of the GaussianVolume. In the first step, we adopt a diffusion model conditioned on input texts to generate GDF. In the second step, we employ a 3D U-Net-based model to predict the attributes of Gaussians by inputting the GDF along with the text condition, leading to the final GaussianVolume.

Gaussian Distance Field Generation The Gaussian Distance Field (GDF) $F \in \mathbb{R}_0^{+1 \times N \times N \times N}$ stores an isotropic feature representing the fundamental geometry of a GaussianVolume. It measures the distance between each grid coordinate and its nearest Gaussian point, similar to the definition of Unsigned Distance Field. This attribute can be easily extracted from fitted GaussianVolumes via sorting algorithms. After obtaining the ground truth GDF, we train a diffusion model conditioned on texts to generate GDF, creating the coarse geometry of objects. The model is supervised by minimizing Mean Square Error (MSE) loss \mathcal{L}_{3D} between the ground truth GDF and the generated GDF, which is equivalent to predicting added noises in the diffusion process. Using diffusion models, the generation diversity with respect to object shape is introduced.

GaussianVolume Prediction After generating GDF, we send it into a U-Net-based model modified from SDFusion [10], along with text descriptions, to predict all attributes of the GaussianVolume. The reason why a reconstruction model is adopted is that we empirically found that the single-step model produces comparable predictions of Gaussian attributes with a diffusion model.

We use two kinds of losses in this phase: MSE loss \mathcal{L}_{3D} between ground truth GaussianVolume and the predicted one, and rendering loss \mathcal{L}_{2D} :

$$\mathcal{L}_{2D} = \lambda \mathcal{L}_1 + (1 - \lambda) \mathcal{L}_{SSIM} \quad (5)$$

for rendered images, which are composed of L1 and SSIM loss. The total loss is,

$$\mathcal{L} = \lambda_{3D} \mathcal{L}_{3D} + \lambda_{2D} \mathcal{L}_{2D} \quad (6)$$

Using a multi-modal loss balances global semantic and local details, and keeps the training process more stable. More implementation details about model architectures and data processing can be found in the supplemental materials.



Fig. 4: Comparisons with State-of-the-art Text-to-3D Methods. Our method achieves competitive visual results with better alignments with text conditions.

4 Experiments

4.1 Baseline Methods and Dataset

Baseline Methods We compare with both feed-forward-based methods and optimization-based methods. Across the former methods, we evaluate Shap-E [18] and VolumeDiffusion [42], both of which directly generate 3D assets. For latter ones, we consider DreamGaussian [41] as the baseline, where coarse 3D Gaussians undergo optimization with SDS loss [35] and are converted into meshes for further refinement. For VolumeDiffusion, we report results generated in the feed-forward process, without post-optimization using SDS via pretrained text-to-image models, to ensure fair comparisons with other methods.




Fig. 5: Text-to-3D Generation Results by GVGEN.

Dataset Our training dataset comprises the Objaverse-LVIS dataset [12], which contains $\sim 46,000$ 3D models in 1,156 categories. For text prompts for training, we use captions from Cap3D [28], which leverages BLIP-2 [21] to caption multi-view images of objects and consolidates them into single captions through GPT-4 [1]. For evaluation, we generate 100 assets and render 8 views for each asset. The camera poses for these views are uniformly sampled around the object.

4.2 Qualitative and Quantitative Results

GaussianVolume Fitting We present the reconstruction results of GaussianVolume Fitting in Fig. 3. These results demonstrate that the proposed GaussianVolume can express high-quality 3D assets with a small number of Gaussian points under the volume resolution $N = 32$. Higher resolutions yield better rendering effects but require more computational resources for the generation stage. Additional ablation studies about fitting GaussianVolume could be found in Sec. 4.3 and the supplemental materials.




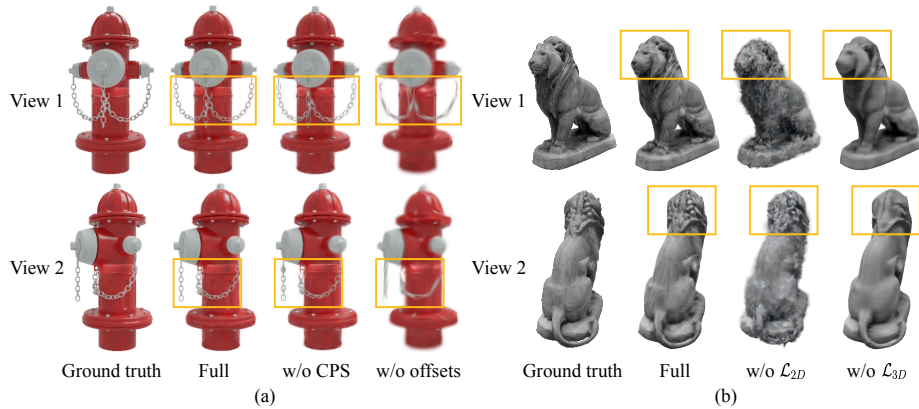
Fig. 6: Display of Generation Diversity. (a) displays diverse results with the same text prompt via GVGEN. (b) shows the comparisons among our image-conditioned GVGEN, OpenLRM [14] and TGS [52]. The single-view reconstruction models suffer from the problem of average patterns, leading to implausible shapes and textures in unseen regions, while GVGEN produces reasonable appearances and geometries.

Text-to-3D Generation We provide visual comparisons and quantitative analyses of our method and existing baselines in Fig. 4 and Tab. 1, respectively. As depicted in the figure, our model generates reasonable geometry and plausible textures. Shap-E [18] can produce rough shapes but sometimes result in misaligned appearances with respect to the input texts. VolumeDiffusion [42] tends to generate unrealistic textures. And DreamGaussian [41], which adopts an optimization-based approach, always produces over-saturated results. For quantitative results, we compare the CLIP score between rendered images and the corresponding texts, and also the generation time of these methods. We randomly select 300 text prompts for evaluation from the 40 LVIS categories containing the highest number of objects. For each object, we render 8 views in uniformly sampled camera poses. To measure generation time, we set the sampling steps to 100 for diffusion-based methods. The configuration of DreamGaussian [41] follows the default settings. Refer to Fig. 5 for more generation results.

Generation Diversity As depicted in Fig. 6(a), GVGEN can generate diverse assets conditioned on the same prompt. The generative diversity of our method not only differentiates it from reconstruction approaches but also heightens the imaginative capabilities of users. Furthermore, we develop an image-to-3D model conditioned on CLIP image embeddings and compare the results with the recently popular single-image reconstruction models TGS [52] and LRM [15] (See Fig. 6(b)). We use OpenLRM [14] as an alternative to the close-sourced LRM. The single-view reconstruction models suffer from the problem of average patterns, leading to implausible shapes and textures in unseen regions, while GV-

Table 1: Quantitative Comparisons against Baseline Methods on CLIP Score and Inference Speed.

Metrics	CLIP score \uparrow Time \downarrow	
Ours	28.53	7 sec
Shap-E	28.48	11 sec
VolumeDiffusion	25.09	7 sec
DreamGaussian	23.60	~ 3 min

**Fig. 7: Qualitative Results for Ablation Studies.** (a) represents visual comparisons of different GaussianVolume fitting methods. (b) stands for results using different losses to train the GaussianVolume prediction model.

GEN produces reasonable appearances and geometries. Such comparisons accentuate the critical difference between GVGEN and reconstruction methods.

4.3 Ablation Studies

GaussianVolume Fitting We first study the effects of strategies fitting GaussianVolume via visual comparisons. In these ablation experiments, 3D assets are trained with 72 images, with camera poses uniformly distributed at a distance of 2.4 from the world center, and evaluated with 24 images posed uniformly at a distance of 1.6. We set volume resolution $N = 32$, i.e. number of Gaussian points $N^3 = 32,768$. The full method, as used in the data preparation stage, performs the best in terms of rendering results, even with only 32,768 Gaussians. "w/o CPS" refers to optimizing coordinate offsets $\Delta\mu$ only via backpropagation, while "w/o offsets" means fixing Gaussian coordinates $\mu = p$ without optimizing point positions. Fig. 7(a) shows that using the full method produces the best rendering

Table 2: Quantitative Metrics for Ablation Studies. The left table analyzes GaussianVolume fitting strategies, while the right table compares different losses training the prediction model. For qualitative comparisons, see Fig. 7.

Metrics	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Metrics	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Full	30.122	0.963	0.038	Full	35.03	0.9872	0.0236
w/o CPS	29.677	0.958	0.049	w/o \mathcal{L}_{3D}	35.21	0.9846	0.0268
w/o offsets	27.140	0.936	0.084	w/o \mathcal{L}_{2D}	29.55	0.9654	0.0444

results. Tab. 2 reports the average PSNR, SSIM, and LPIPS scores, supporting the effectiveness of the proposed strategy quantitatively.

Text-to-3D Generation Due to the computational resources required, it is impractical to evaluate each part of our design with models trained on the full Objaverse-LVIS dataset for our ablation study. Therefore, we train the models on a small subset of Objaverse-LVIS to validate the effectiveness of losses in the stage of predicting GaussianVolume attributes. Specifically, we pick 1,336 assets themed animals from 24 LVIS categories for training.

We randomly select $\sim 20\%$ of the data from the training data and evaluate quantitative metrics with rendered novel views. Fig. 7(b) demonstrates that the proposed multi-modal losses produce plausible results with fine-grained details, and quantitative results are listed in Tab. 2. The combination of two types of losses leads to more detailed textures while keeping the geometry smooth.

4.4 Limitations

GVGEN has shown encouraging results in generating 3D objects. However, its performance is constrained when dealing with input texts significantly divergent from the domain of training data. It is also time-consuming for fitting millions of training data to scale up for better diversity. Additionally, the volume resolution N is set as 32 to save computational resources, which limits the rendering effects of 3D assets with very complex textures. In the future, we will further explore how to generate higher-quality 3D assets in more challenging scenarios.

5 Conclusions

In conclusion, this paper explores the feed-forward generation of explicit 3D Gaussians conditioned on texts. We innovatively organize disorganized 3D Gaussian points into a structured volumetric form GaussianVolume with a novel pruning and densifying strategy, i.e. Candidate Pool Strategy, enabling feed-forward generation of 3D Gaussians via a coarse-to-fine generation pipeline. Our proposed framework, GVGEN, demonstrates remarkable efficiency in generating 3D Gaussians from texts. Experimental results demonstrate its competitive capabilities. This progress suggests potential extensions of our approach in tackling a broader spectrum of challenges within the field.

Acknowledgments. This work was done during Xianglong He and Junyi Chen’s internship at Shanghai Artificial Intelligence Laboratory. The work is supported by the National Key R&D Program of China (No.2022YFB4701400/4701402, No.2022ZD0160102), SSTIC Grant (KJZD20230923115106012, KJZD20230923114916032), and Beijing Key Lab of Networked Multimedia.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
2. Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5855–5864 (2021)
3. Cao, A., Johnson, J.: Hexplane: A fast representation for dynamic scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 130–141 (2023)
4. Chang, J.H.R., Chen, W.Y., Ranjan, A., Yi, K.M., Tuzel, O.: Pointersect: Neural rendering with cloud-ray intersection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8359–8369 (2023)
5. Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: European Conference on Computer Vision. pp. 333–350. Springer (2022)
6. Chen, G., Wang, W.: A survey on 3d gaussian splatting. arXiv preprint arXiv:2401.03890 (2024)
7. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. arXiv preprint arXiv:2303.13873 (2023)
8. Chen, Y., Chen, Z., Zhang, C., Wang, F., Yang, X., Wang, Y., Cai, Z., Yang, L., Liu, H., Lin, G.: Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. arXiv preprint arXiv:2311.14521 (2023)
9. Chen, Z., Wang, F., Liu, H.: Text-to-3d using gaussian splatting. arXiv preprint arXiv:2309.16585 (2023)
10. Cheng, Y.C., Lee, H.Y., Tulyakov, S., Schwing, A.G., Gui, L.Y.: Sdfusion: Multi-modal 3d shape completion, reconstruction, and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4456–4465 (2023)
11. Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., et al.: Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems* **36** (2024)
12. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13142–13153 (2023)
13. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
14. He, Z., Wang, T.: Openlrn: Open-source large reconstruction models. <https://github.com/3DTopia/OpenLRM> (2023)

15. Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: Lrm: Large reconstruction model for single image to 3d. arXiv preprint arXiv:2311.04400 (2023)
16. Huang, Z., Wen, H., Dong, J., Wang, Y., Li, Y., Chen, X., Cao, Y.P., Liang, D., Qiao, Y., Dai, B., et al.: Epidiff: Enhancing multi-view synthesis via localized epipolar-constrained diffusion. arXiv preprint arXiv:2312.06725 (2023)
17. Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 867–876 (2022)
18. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463 (2023)
19. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics **42**(4) (2023)
20. Li, J., Tan, H., Zhang, K., Xu, Z., Luan, F., Xu, Y., Hong, Y., Sunkavalli, K., Shakhnarovich, G., Bi, S.: Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. arXiv preprint arXiv:2311.06214 (2023)
21. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv preprint arXiv:2301.12597 (2023)
22. Li, W., Chen, R., Chen, X., Tan, P.: Sweetdreamer: Aligning geometric priors in 2d diffusion for consistent text-to-3d. arXiv preprint arXiv:2310.02596 (2023)
23. Liang, Y., Yang, X., Lin, J., Li, H., Xu, X., Chen, Y.: Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching (2023)
24. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023)
25. Liu, Y., Li, Q., Sun, Z.: One-shot face reenactment with dense correspondence estimation. Machine Intelligence Research pp. 1–13 (2023). <https://doi.org/10.1007/s11633-023-1433-9>, <https://www.mi-research.net/en/article/doi/10.1007/s11633-023-1433-9>
26. Liu, Z., Li, Y., Lin, Y., Yu, X., Peng, S., Cao, Y.P., Qi, X., Huang, X., Liang, D., Ouyang, W.: Unidream: Unifying diffusion priors for relightable text-to-3d generation. arXiv preprint arXiv:2312.08754 (2023)
27. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. arXiv preprint arXiv:2310.15008 (2023)
28. Luo, T., Rockwell, C., Lee, H., Johnson, J.: Scalable 3d captioning with pretrained models. Advances in Neural Information Processing Systems **36** (2024)
29. Melas-Kyriazi, L., Rupprecht, C., Vedaldi, A.: Pc2: Projection-conditioned point cloud diffusion for single-image 3d reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12923–12932 (2023)
30. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021)
31. Mohammad Khalid, N., Xie, T., Belilovsky, E., Popa, T.: Clip-mesh: Generating textured meshes from text using pretrained image-text models. In: SIGGRAPH Asia 2022 conference papers. pp. 1–8 (2022)
32. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) **41**(4), 1–15 (2022)

33. Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3d point clouds from complex prompts. arXiv preprint arXiv:2212.08751 (2022)
34. Ntavelis, E., Siarohin, A., Olszewski, K., Wang, C., Gool, L.V., Tulyakov, S.: Autodecoding latent 3d diffusion models. *Advances in Neural Information Processing Systems* **36** (2024)
35. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022)
36. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International conference on machine learning*. pp. 8748–8763. PMLR (2021)
37. Ren, J., Pan, L., Tang, J., Zhang, C., Cao, A., Zeng, G., Liu, Z.: Dreamgaussian4d: Generative 4d gaussian splatting. arXiv preprint arXiv:2312.17142 (2023)
38. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
39. Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model. arXiv preprint arXiv:2310.15110 (2023)
40. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512 (2023)
41. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653 (2023)
42. Tang, Z., Gu, S., Wang, C., Zhang, T., Bao, J., Chen, D., Guo, B.: Volumediffusion: Flexible text-to-3d generation with efficient volumetric encoder. arXiv preprint arXiv:2312.11459 (2023)
43. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in Neural Information Processing Systems* **36** (2024)
44. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. arXiv preprint arXiv:2310.08528 (2023)
45. Wu, Z., Wang, Y., Feng, M., Xie, H., Mian, A.: Sketch and text guided diffusion model for colored point cloud generation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 8929–8939 (2023)
46. Xu, D., Yuan, Y., Mardani, M., Liu, S., Song, J., Wang, Z., Vahdat, A.: Agg: Amortized generative 3d gaussians for single image to 3d. arXiv preprint arXiv:2401.04099 (2024)
47. Xu, Q., Xu, Z., Philip, J., Bi, S., Shu, Z., Sunkavalli, K., Neumann, U.: Point-nerf: Point-based neural radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5438–5448 (2022)
48. Yi, T., Fang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: Gaussiandreamer: Fast generation from text to 3d gaussian splatting with point cloud priors. arXiv preprint arXiv:2310.08529 (2023)
49. Yin, Y., Xu, D., Wang, Z., Zhao, Y., Wei, Y.: 4dgen: Grounded 4d content generation with spatial-temporal consistency. arXiv preprint arXiv:2312.17225 (2023)
50. Yu, X., Guo, Y.C., Li, Y., Liang, D., Zhang, S.H., Qi, X.: Text-to-3d with classifier score distillation. arXiv preprint arXiv:2310.19415 (2023)

51. Zhao, H., Zhang, J., Chen, Z., Yuan, B., Tao, D.: On robust cross-view consistency in self-supervised monocular depth estimation. *Machine Intelligence Research* **21**(3), 495–513 (2024). <https://doi.org/10.1007/s11633-023-1474-0>, <https://www.mi-research.net/en/article/doi/10.1007/s11633-023-1474-0>
52. Zou, Z.X., Yu, Z., Guo, Y.C., Li, Y., Liang, D., Cao, Y.P., Zhang, S.H.: Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. *arXiv preprint arXiv:2312.09147* (2023)