

View-Consistent 3D Editing with Gaussian Splatting

Yuxuan Wang¹, Xuanyu Yi^{1,2}, Zike Wu¹, Na Zhao³, Long Chen⁵, and Hanwang Zhang^{1,4}

¹Nanyang Technological University ²Institute for Infocomm Research, A*STAR

³Singapore University of Technology and Design ⁴Skywork AI

⁵Hong Kong University of Science and Technology




Fig. 1: Capability highlight of our method: VcEDIT. Given a source 3D Gaussian Splatting and a user-specified text prompt, our VcEDIT enables versatile scene and object editing. By ensuring multi-view consistent image guidance, VcEDIT alleviates artifacts and excels in high-quality editing.

Abstract. The advent of 3D Gaussian Splatting (3DGS) has revolutionized 3D editing, offering efficient, high-fidelity rendering and enabling precise local manipulations. Currently, diffusion-based 2D editing models are harnessed to modify multi-view rendered images, which then guide the editing of 3DGS models. However, this approach faces a critical issue of multi-view inconsistency, where the guidance images exhibit significant discrepancies across views, leading to mode collapse and visual artifacts of 3DGS. To this end, we introduce View-consistent Editing (VcEDIT), a novel framework that seamlessly incorporates 3DGS into image editing processes, ensuring multi-view consistency in edited guidance images and effectively mitigating mode collapse issues. VcEDIT

employs two innovative consistency modules: the Cross-attention Consistency Module and the Editing Consistency Module, both designed to reduce inconsistencies in edited images. By incorporating these consistency modules into an iterative pattern, VCEDIT proficiently resolves the issue of multi-view inconsistency, facilitating high-quality 3DGS editing across a diverse range of scenes. Further code and video results are released at <http://yuxuanw.me/vcredit/>.

Keywords: 3D Editing · 3D Gaussian Splatting · Multi-view Consistency · Diffusion Model

1 Introduction

We consider the problem of text-driven 3D model editing: given a source 3D model and user-specified text instructions, the task is to modify the 3D model according to the instructions, as depicted in Fig. 1, ensuring both editing fidelity and preservation of essential source content [4, 5, 7, 18, 34]. This problem holds paramount importance across a variety of industrial applications, such as real-time outfit changes for 3D digital humans and immersive AR/VR interactive environments [3, 39, 43]. Recently, groundbreaking 3D Gaussian Splatting (3DGS) [19] has emerged as a promising “silver bullet” for 3D editing, notable for its efficient, high-fidelity rendering and explicit representation (3D anisotropic balls known as *Gaussians*) suitable for local manipulation.

In the context of editing 3DGS models, thanks to recent progress in large-scale pre-trained 2D diffusion models, existing methods [6, 8–10, 15] leverage off-the-shelf 2D editing models [1, 11, 12, 14, 28, 38] to guide optimization of the 3DGS model. As shown in Fig. 2(a), this pattern renders source 3DGS into multi-view 2D images, manipulates them via 2D editing models using text prompts, and then employs these adjusted images to fine-tune the original 3DGS. Beyond achieving plausible editing outcomes, this image-based pattern also facilitates user-friendly interaction, enabling users to pre-select their preferred edited images and personalize the editing workflow.

However, such image-guided 3DGS editing has a notorious **multi-view inconsistency** that cannot be ignored. Fig. 2(a) vividly illustrates that images edited separately using a state-of-the-art 2D editing model [38] manifest pronounced inconsistencies across views — the views of a man are edited to different styles of clowns. Utilizing these significantly varied edited images as guidance, the 3DGS model will struggle with the issue that few training images are coherent, whereas the majority display conflicting information. Unfortunately, the explicitness and inherent densification process of 3DGS make it especially vulnerable to multi-view inconsistency, which complicates 3DGS in densifying under-reconstruction regions or pruning over-reconstruction regions [35]. Consequently, training with the multi-view inconsistent guidance can lead to the mode collapse of 3DGS, characterized by the ambiguity between the source and target, as well as the flickering artifacts revealed in Fig. 2(a).

Thus, the crux lies in addressing the multi-view inconsistency of the image guidance. We conjecture such a problem stems from the lack of 3D awareness in 2D editing models; that is, they inherently process each view in isolation. Therefore, we introduce the **View-consistent Editing (VCEDIT)**, a high-quality image-guided 3DGS editing framework. This framework seamlessly incorporates 3DGS into image editing processes to achieve multi-view consistent guidance, thus effectively addressing the issue of 3DGS mode collapse. As illustrated in Fig. 2(b), VCEDIT employs specially-designed multi-view consistency modules within an iterative pattern.

Primarily, we design two effective consistency modules using the explicit nature and fast rendering capability of 3DGS: (1) The **Cross-attention Consistency Module (CCM)** that consolidates the multi-view cross-attention maps in the diffusion-based image editing model, thus harmonizing the model’s attentive 3D region across all views. More concretely, this process inverse-renders the original cross-attention maps from all views onto each Gaussian within the source 3DGS, thereby creating an averaged 3D map. Subsequently, this 3D map is rendered back to 2D, serving as the consolidated cross-attention maps to replace the originals for more coherent edits. (2) The **Editing Consistency Module (ECM)** that directly calibrates the multi-view inconsistent editing outputs: We fine-tune a source-cloned 3DGS with the editing outputs and then render the 3DGS back to images. Taking advantage of the rapid rendering speed of the 3DGS, this mechanism efficiently decreases incoherent content in each edited image.

To further mitigate the multi-view inconsistency issue, we extend our VCEDIT to an **iterative pattern**: editing rendered images \rightarrow updating the 3DGS \rightarrow repeating. In situations where the image editing model yields overly inconsistent initial edits, it allows for the correction of initially inconsistent views in later iterations [15], which continuously refines the 3DGS and fosters a reciprocal cycle. Fig. 2(b) illustrates that the 3DGS of a “man” is iteratively guided to a consistent style that aligns with the desired “clown” target. To meet the demand for rapid iteration, VCEDIT integrates InfEdit [38], a high-quality, fast image editing model that bypasses the lengthy DDIM-inversion phase. Depending on the complexity of the editing scene and user instructions, the processing time in our VCEDIT for each sample ranged from 10 to 20 minutes.

As illustrated in Fig. 1, by incorporating consistency modules and the iterative pattern, VCEDIT significantly enhances the multi-view consistency in guidance images, leading to superior editing quality. We conducted comprehensive evaluations of VCEDIT in various real-world scenes. Both qualitative and quantitative experiments clearly indicate that VCEDIT can address the mode collapse of 3DGS, continuously outperforming other state-of-the-art methods. Our contributions can be summarized in three aspects:

- We propose effective multi-view **consistency modules** that harmonize the inconsistent multi-view guidance images.
- Based on our designed consistency module, we construct an effective **iterative pattern** that offers high-quality 3DGS editing.

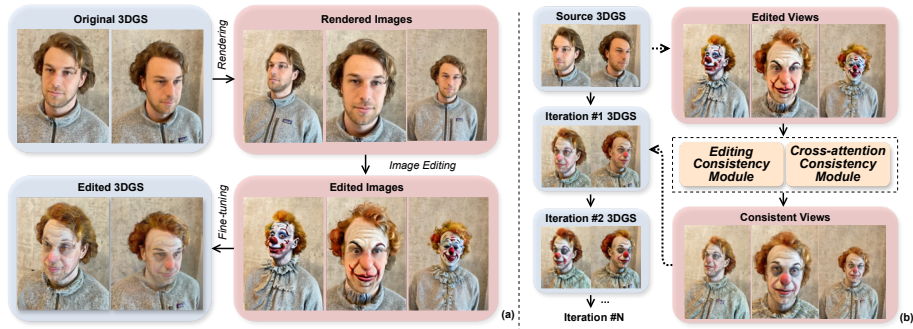


Fig. 2: (a): Current image-guided 3DGS editing pipeline and its **multi-view inconsistency issue**: The rendered views of a man are separately edited to different styles of clowns, leading to the mode collapse issue of learned 3DGS. (b): The iterative pattern and our consistency modules deployed in each iteration. The 3DGS is progressively guided to a coherent style that aligns with the “clown” through the iterative pattern.

- We pioneer a paradigm that integrates 3DGS characteristics with multi-view consistency design, inspiring further efforts and explorations.

2 Related Works

Text-guided Image Editing. Large-scale pretrained Text-to-image (T2I) diffusion models [32, 37] have been widely adopted in generative image editing, guided by specific text prompts. Textual Inversion [12] optimizes special prompt token(s) in text embedding space to represent the specified concept, whereas Prompt-to-Prompt leverages cross-attention information to maintain integrity in unmodified regions. Pix2pix-zero [28] utilizes embedding vector mechanisms to establish controllable editing. InstructPix2Pix [14] merges the capacities of GPT and Stable Diffusion [32] to generate a multi-modal dataset and fine-tunes Stable Diffusion to achieve instructed diverse editing. Our VCEDIT integrates InfEdit [38], which implies a virtual inversion strategy and attention control mechanisms for a controllable and faithful editing. Nevertheless, all the above works mainly focus on editing a single image, and obtaining high-quality paired data with multi-view consistency from the 2D diffusion model is extremely challenging, resulting in the application of 3D-aware editing impractical.

3D Field Editing. In the last few years, NeRF-based approaches [2, 17, 18, 20, 26, 27, 31, 41, 42, 47] have overtaken traditional point cloud and triangle mesh approaches [2, 27, 36, 44, 46] for 3D field editing. DreamEditor [48] updates textual attention areas with score distillation sampling (SDS) guidance [29]. DDS [16] and PDS [22] introduce an advance SDS loss with cleaner optimization direction for better instruction alignment. FocalDreamer [24] generates new geometries in specified empty spaces from the text input. Instruct-NeRF2NeRF [15] iteratively refines dataset images using the InstructPix2Pix diffusion model and

NeRF, a pattern widely adopted in subsequent research [9, 23]. However, the inherent computational demands and the implicit nature of NeRF usually require time-consuming fine-tuning for edit operations and lack of precise controllability, which cannot support practical interactive editing use cases. More recent works [6, 10] integrate Gaussian splatting into 3D field editing and surpasses previous 3D editing methods in terms of effectiveness, speed, and controllability. Our VCEdit builds upon the above image-based 3DGS editing approaches and further harmonizes the inconsistent multi-view image guidance they overlook, which enables more high-fidelity editing quality without mode collapse.

3 Preliminary

3.1 3D Gaussian Splatting

3D Gaussian splatting (3DGS) has emerged as a prominent efficient explicit 3D representation technique, employing anisotropic 3D Gaussians for intricate modeling [19]. Each Gaussian, represented as G , is characterized by its mean $\mu \in \mathbb{R}^3$, covariance matrix Σ , associated color $c \in \mathbb{R}^3$, and opacity $\alpha \in \mathbb{R}$. For better optimization, the covariance matrix Σ is decomposed into a scaling matrix $S \in \mathbb{R}^3$ and a rotation matrix $R \in \mathbb{R}^{3 \times 3}$, with Σ defined as $\Sigma = RSS^T R^T$. A Gaussian centered at μ is expressed as $G(x) = \exp(-\frac{1}{2}x^T \Sigma^{-1}x)$, where x denotes the displacement from μ to a point in space. In the splatting rendering process, we render the color c for a pixel by blending all sampled 3D points along the ray going through this pixel:

$$c = \sum_i c_i \alpha_i G(x_i) \prod_{j=1}^{i-1} (1 - \alpha_j G(x_j)). \quad (1)$$

An efficient tile-based rasterizer [19] allows for rapid forward and backward passes, supporting high-quality real-time rendering. The optimization of 3D Gaussian properties is conducted in tandem with an adaptive *density control mechanism*, involving densifying and pruning operations, where Gaussians are added and occasionally removed. By incorporating all the Gaussians G , the 3DGS model, denoted as \mathcal{G} , represents complete 3D scenes and objects.

3.2 Image-guided 3DGS Editing

Given a source 3D Gaussian Splatting (3DGS) model \mathcal{G}^{src} and a user-specified target prompt y , the goal of 3DGS editing is to transform \mathcal{G}^{src} into an edited version $\mathcal{G}^{\text{edit}}$ that aligns with y . In image-guided 3DGS editing, \mathcal{G}^{src} is first rendered from multiple views $\mathcal{V} = \{v\}$, generating a collection of source images \mathcal{I}^{src} . Subsequently, a 2D editing model transforms \mathcal{I}^{src} into edited images $\mathcal{I}^{\text{edit}}$ according to y . Finally, these edited images serve as training guidance to refine \mathcal{G}^{src} into the edited version $\mathcal{G}^{\text{edit}}$. Specifically, an editing loss L_{Edit} is calculated for each view between the real-time rendering and the edited images. The final

edited 3DGS model is obtained by minimizing the editing loss across all views \mathcal{V} :

$$\mathcal{G}^{\text{edit}} = \underset{\mathcal{G}}{\operatorname{argmin}} \sum_{v \in \mathcal{V}} \mathcal{L}_{\text{Edit}}(\mathcal{R}(\mathcal{G}, v), \mathcal{I}^{\text{edit}}), \quad (2)$$

where \mathcal{R} represents the rendering function that projects 3DGS to image given a specific view v . However, as introduced in Section 1, the edited images $\mathcal{I}^{\text{edit}}$ produced by 2D editing models exhibit severe multi-view inconsistency issue if no further control is conducted. This inconsistency in 2D guidance images will lead to mode collapse and low-quality editing in the edited 3DGS model.

4 Our Method

In this section, we introduce the View-Consistent Editing (VCEDIT), a novel approach for 3DGS editing. First, we outline the general pipeline of VCEDIT in Sec. 4.1. Subsequently, we introduce two innovative Consistency Modules designed to address the multi-view inconsistency issue. Particularly, in Sec. 4.2, we introduce the Cross-attention Consistency Module, which aims to reduce multi-view editing variance by improving the attention consistency. In Sec. 4.3, we present the Editing Consistency Module, which directly calibrates multi-view editing predictions to ensure consistency across all predictions. Finally, in Sec. 4.4, we extend our framework to an iterative pattern, which facilitates the iterative improvement of consistency between $\mathcal{I}^{\text{edit}}$ and $\mathcal{G}^{\text{edit}}$.

4.1 View-consistent Editing Pipeline

Fig. 3 illustrates the overview of our framework, VCEDIT. To perform image-guided 3DGS editing, we utilize the pre-trained 2D diffusion model [32, 37] to generate a series of edited images given a set of multi-view rendered images.

Inspired by [38], we adopt an annealed timestep schedule from the largest timestep T to 1, where we evenly sample N timesteps from the schedule, denoted as $t \in \{t_i\}_{i=1}^N$. Before the start of the first timestep t_1 , the rendered images, denoted as \mathcal{I}^{src} , are encoded into multi-view source latents z^{src} .

For each sampled timestep, the source latents z^{src} go through an editing process. Specifically, at timestep t , we first adopt the forward diffusion process to perturb the given rendered images with random Gaussian noise ε as follows:

$$z_t = \sqrt{\alpha_t} z^{\text{src}} + \sqrt{1 - \alpha_t} \varepsilon, \quad \varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (3)$$

where the α_t is a time-variant hyper-parameter, z^{src} is the encoded latent images. Then we feed the noisy latents z_t into a pre-trained diffusion model ε_θ to generate denoised images as edited latents based on the user’s target prompt y :

$$z^{\text{edit}} = (z_t - \sqrt{1 - \alpha_t} \cdot \varepsilon_\theta(z_t, y)) / \sqrt{\alpha_t}. \quad (4)$$

However, we observe that the edited images generated by simply adopting pre-trained 2D diffusion models lacks of multi-view consistency, which often leads

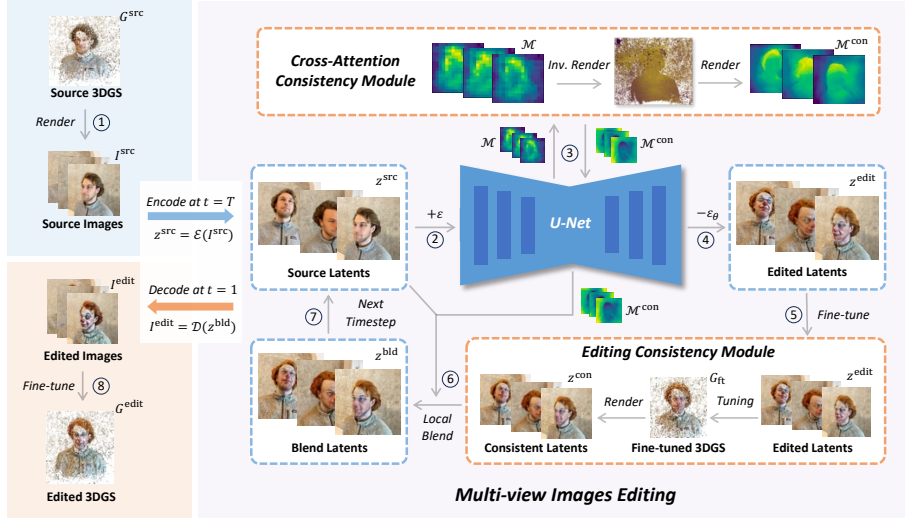


Fig. 3: The pipeline of our VCEDIT: VCEDIT employs an image-guided editing pipeline. In the image editing stage, the Cross-attention Consistency Module and Editing Consistency Module are employed to ensure the multi-view consistency of edited images. We provide a detailed overview in Sec. 4.1.

to mode collapse in the learned 3DGS [25]. Therefore, we propose the Cross-attention Consistency Module in Sec. 4.2, which consolidates multi-view cross-attention maps within each layer of the diffusion model’s U-Net architecture, improving the view-consistency in z^{edit} .

Upon obtaining the edited latents, we aim to further improve the multi-view consistency of the edited guidance images. To this end, we design the Editing Consistency Module, where we adopt a fine-tuning and rendering process using 3DGS to calibrate z^{edit} , producing a set of more consistent latents, denoted as z^{con} . Finally, a local blending process is conducted to preserve source information, generating the blended latents, denoted as z^{bld} , which serve as the output of the current timestep t . Details are demonstrated in Sec. 4.3.

Upon completion of one timestep, its output, z^{bld} , continually serves as the input of the succeeding timestep. This editing process is repeated as t progresses from t_1 to t_N . Finally, the last timestep’s output are decoded back to multi-view edited images, denoted as $\mathcal{I}^{\text{edit}}$, which serves as the view-consistent guidance images to fine-tune the G^{src} , where the training objective includes a MAE loss and a VGG-based LPIPS loss [33, 45]:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{MAE}}(\mathcal{R}(\mathcal{G}, v), \mathcal{I}^{\text{edit}}) + \lambda_2 \mathcal{L}_{\text{LPIPS}}(\mathcal{R}(\mathcal{G}, v), \mathcal{I}^{\text{edit}}) \quad (5)$$

In summary, VCEDIT is integrated with two effective consistency modules. In Sec. 6.1, our experiments thoroughly demonstrate that it can produce consistent image editing results which are directly used as guidance.

4.2 Cross-attention Consistency Module

As discussed in Sec. 1, existing image-guided 3DGS editing approaches [6, 10, 15] frequently encounter multi-view inconsistencies in the generated guidance images. We suggest that this inconsistency arises from the lack of cross-view information exchange during the editing process.

Therefore, we develop the **Cross-attention Consistency Module (CCM)** that consolidates the cross-attention maps from all views during the forward pass of the U-Net, where these maps indicate the pixel-to-text correlations between images and user prompts. In this way, we successfully enable cross-view information exchange within the backbone network in a seamless, plug-and-play manner, thereby facilitating multi-view consistency in the generation of edited images.

As illustrated in the *top* of Fig. 3, we first employ inverse-rendering [6, 19] to map the 2D cross-attention maps \mathbf{M} back onto the 3D Gaussians within \mathcal{G}^{src} , generating a 3D map \mathbb{M} . Given the j -th Gaussian in \mathcal{G}^{src} and the pixel \mathbf{p} in the v -th 2D map \mathbf{M}_v , the $o_j(\mathbf{p})$, $T_{j,v}(\mathbf{p})$, $\mathbf{M}_{v,\mathbf{p}}$ denote the opacity, transmittance matrix, and cross-attention weight, respectively. Specifically, for the j -th Gaussian in \mathcal{G}^{src} , its value in the 3D map \mathbb{M}_j is computed as:

$$\mathbb{M}_j = \frac{1}{C_j} \sum_{v \in V} \sum_{\mathbf{p} \in \mathbf{M}_v} o_j(\mathbf{p}) \cdot T_{j,v}(\mathbf{p}) \cdot \mathbf{M}_{v,\mathbf{p}} \quad (6)$$

The C_j is the count indicating the total number of values that assigned to j by $\mathbf{M}_{v,\mathbf{p}}$ across all \mathbf{p} and v . In the created 3D map \mathbb{M} , each Gaussian receives an attention weight towards each word in the prompt. Subsequently, we render the \mathbb{M} back to 2D space to obtain a set of consistent cross-attention maps, denoted as \mathbf{M}^{con} .

By replacing the vanilla cross-attention maps \mathbf{M} with the consistent attention maps \mathbf{M}^{con} , our CCM conducts an 3D average pooling to unify model’s attentive region across different views. In this way, the U-Net model ensures uniform attention is given to the same 3D region across different views during the editing process, facilitating the production of more consistent edited latents. Our experiments, detailed in Sec. 6.1, validate that this cross-attention control mechanism significantly bolsters view consistency.

4.3 Editing Consistency Module

Incorporating the CCM described in Sec. 4.2 allows the diffusion models to produce z^{edit} with enhanced consistency. Yet, it is still inadequate for updating the 3DGS directly, since 3DGS is based on spatial physical modeling that requires coherent representation across views. Therefore, we introduce **Editing Consistency Module (ECM)** which distills the image guidance into an intermediate model incorporating physical modeling and consistency constraints. This intermediate model is then utilized to generate more coherent multi-view edited guidance images.

We illustrate the proposed module in the *bottom-right* of Fig. 3. Initially, we decode the multi-view edited latents z^{edit} into images, denoted as $\mathcal{D}(z^{\text{edit}})$,

which are subsequently used as guidance to fine-tune a copy of the original 3DGS model \mathcal{G}^{src} , which is denoted as \mathcal{G}_{ft} :

$$\mathcal{G}_{\text{ft}} = \underset{\mathcal{G}}{\operatorname{argmin}} \sum_{v \in \mathcal{V}} \mathcal{L}_{\text{edit}}(\mathcal{R}(\mathcal{G}, v), \mathcal{D}(z_v^{\text{edit}})) \quad (7)$$

Following the rapid fine-tuning process, we render the fine-tuned model \mathcal{G}_{ft} from each viewpoint $v \in \mathcal{V}$, yielding a collection of rendered images that exhibit significantly enhanced coherence. Finally, these images are encoded to the consistent latents, denoted as z^{con} .

Afterwards, we follow [38] and employ a Local Blend module to preserve the unedited details in the source image: Given the cross-attention maps \mathbf{M}^{con} , the source latents z^{src} , and the consistent latents z^{con} , the blending is conducted by:

$$z^{\text{bld}} = \mathbf{M}^{\text{con}} * z^{\text{con}} + (1 - \mathbf{M}^{\text{con}}) * z^{\text{src}} \quad (8)$$

Our empirical findings in *appendix* suggest that 3DGS is capable of rectifying inconsistencies when the discrepancies between views are minor. Our image editing process, unfolding across various diffusion timesteps following InfEdit [38], guarantees that z^{edit} undergoes only slight changes from z^{src} , maintaining a low level of inconsistencies at each step. Consequently, through continuously calibration at each timestep, the ECM effectively prevents the inconsistencies from accumulating. In our ablation study (Sec. 6.1), we demonstrate the significant capability of this mechanism to address multi-view inconsistencies.

4.4 Iterative Pattern Extension

To further alleviate the multi-view inconsistency issue, we evolve our VcEDIT into an iterative pattern [40]. Based on the image-guided editing pipeline detailed in Sec. 3.2, we define each sequence of transitions—from \mathcal{G}^{src} to \mathcal{I}^{src} , from \mathcal{I}^{src} to $\mathcal{I}^{\text{edit}}$, and from $\mathcal{I}^{\text{edit}}$ to $\mathcal{G}^{\text{edit}}$ —as an iteration. Upon completion of one iteration, the $\mathcal{G}^{\text{edit}}$ is forwarded to the subsequent iteration and serves as the new \mathcal{G}^{src} .

In this pattern, each iteration incrementally aligns the 3DGS more closely with the user’s prompt, facilitating smoother image editing. This fosters a mutual cycle, where $\mathcal{I}^{\text{edit}}$ and $\mathcal{G}^{\text{edit}}$ successively refine each other, resulting in ongoing enhancements of editing quality [15]. This leads to more consistent and higher-quality guidance images, which in turn refines the 3DGS for improved $\mathcal{G}^{\text{edit}}$. Taking advantage of our two effective Consistency Modules proposed in Sec. 4.2 and Sec. 4.3, only a few iterations are needed to achieve satisfying editing results. By integrating these strategies, our VcEDIT framework effectively ensures multi-view consistency in guidance images and excellence in 3DGS editing.

5 Experiments

5.1 Implementation Details

We implement our experiments using PyTorch based on the official codebase of 3D Gaussian splatting [19]. All of our experiments are conducted on a single



Fig. 4: Qualitative comparison with the DDS [16] and the GSEditor [6]: The *topmost* row demonstrate the original views, while the *bottom* rows show the rendering view of edited 3DGS. VcEDIT excels by effectively addressing the multi-view inconsistency, resulting in superior editing quality. In contrast, other methods encounter challenges with mode collapse and exhibit flickering artifacts.

NVIDIA RTX A6000 GPU. We used Adam optimizer [21] with a learning rate of 0.001 in both fine-tuning the source 3DGS and our Editing Consistency Module. For each 3DGS editing task, we conduct one to three editing iterations where the source 3DGS is optimized for 400 steps.

5.2 Baselines and Evaluation Protocol

We compare our proposed VcEDIT method with two baselines: (1) **Delta Distillation Sampling** (DDS) [16], which utilizes an implicit score function loss derived from a diffusion model as guidance, and (2) the state-of-the-art **GSEditor** [6], which employs image-guided editing and iteratively updates a training set to steer 3DGS modifications. To ensure a fair comparison, all methods incorporate the same local 3DGS update process pioneered by GSEditor. Our quantitative comparisons assess the similarity of appearance between the edited 3DGS and the user’s prompt with the CLIP-similarity metrics [13]. Furthermore, a user study was conducted to evaluate the effectiveness of editing methods from a human perspective. Please refers to *Appendix* for the detailed configuration of the user study.




Fig. 5: Extensive Results of VCEDIT: Our method is capable of various editing tasks, including face, object, and large-scale scene editing. The leftmost column demonstrates the original view, while the right four columns show the rendered view of edited 3DGS.

5.3 Qualitative Comparisons

We meticulously select three challenging scenarios ranging from human face editing to outdoor scene editing. In the first column of Fig. 4, the precision-demanding task of human face editing highlights the critical importance of accurate manipulation due to the sensitivity of facial features to imperfections. Subsequent columns feature objects placed within outdoor scenes with panoramic rendering views, posing a challenge for methods to maintain view-consistent guidance during editing.

Fig. 4 shows the editing results achieved by our VCEDIT along with those of the baseline methods. (1) We observe a notable editing failure in all the DDS results, where the 3DGS presents severe noisy coloration and shape distortion. (2) When compared to DDS, the output from GSEditor shows better alignment with the target prompt. However, the resulting 3DGS exhibits noticeable artifacts, including noisy coloration and blurred edges. (3) Unlike DDS and GSEditor, our VCEDIT consistently delivers clean and intricately detailed edits in all scenes. Our results confirm that VCEDIT exhibits superior editing performance compared to other baselines.

5.4 Quantitative Comparisons

Despite the subjective nature of 3D scene generation and editing, we adhere to established practices by utilizing the CLIP [30] text-to-image directional similarity metric [13] for quantitative analysis. This metric quantifies the semantic differences between the original and edited 3DGS scenes in relation to their corresponding prompts. Furthermore, we carried out a user study that collected

Table 1: Quantitative Comparison: Our VCEDIT performs in both user study evaluations and CLIP T2I Directional Similarity [12] metrics.

Metrics\Methods	DDS [16]	GSEditor [6]	VCEDIT
User Study	1.57%	34.49%	63.93%
CLIP similarity [13]	0.1470	0.1917	0.2108

data from 25 participants in 11 samples, each participant tasked with selecting the most superior edit from the edited outcomes of all comparison methods considering alignment to the target prompt and visual quality. Table 1 presents the quantitative results of the evaluation of our approach comparing with DDS and GSEditor in various test scenes. VCEDIT not only achieves superior results in the user study but also outperforms in CLIP text-to-image directional similarity. This performance indicates that VCEDIT provides high-quality edits that are closely aligned with the intended prompts, significantly outpacing the baseline methods in terms of editing quality and effectiveness.

5.5 Editing with Diverse Source

Fig. 5 highlights our VCEDIT’s adaptability in handling a wide range of challenging scenes and prompts, spanning from intricate face and object transformations to extensive adjustments in large-scale scenes. In the first row, our VCEDIT produces an editing result that vividly presents the facial details of the target, “grandma”. The second example highlights VCEDIT’s ability to transform a small “doll” into a “robot” while carefully preserving the doll’s key characteristics. In the last row, VCEDIT demonstrates its capacity for large-scale scene modifications, convincingly altering a woodland scene to depict it as “on fire”. These instances affirm VCEDIT’s competency across a diverse array of tasks.

6 Ablation and Discussion

6.1 Ablation on Consistency Modules

Sec. 4.2 and 4.3 respectively introduced the Cross-attention Consistency Module (CCM) and the Editing Consistency Module (ECM). This section presents an ablation study to assess the impact of these modules. Given the critical role of multi-view consistent guidance in image-guided 3D editing as discussed in Sec. 1, our analysis focuses on the view-consistency of the edited 2D images. To understand why VCEDIT outperforms the baseline methods, we revisit the cases of “clown” and “brown bear” from Sec. 5.3 for a detailed analysis.

We conducted a comparative analysis of the edited images produced by three variant versions of VCEDIT: the vanilla version *without any consistency module*; the version *with CCM but without ECM*; and the version that *employs both CCM and ECM*. The top 3 rows in Fig. 6 displays the comparison results: (1)



Fig. 6: *Top:* 2D edited guidance images using three variant versions of VcEDIT: vanilla version without any consistency module, version without ECM, and the original VcEDIT employs both CCM and ECM. *Bottom:* 2D edited guidance images using GSEditor

The edited images of the vanilla version contain significant view inconsistencies in the two samples. These incoherent edited contents among views can lead to the anomalies in the 3DGS. (2) Applying the CCM markedly enhances the view consistency. Compared to the vanilla version, the edited images show a reduction of the incoherent content. (3) When both consistency modules are employed, the results are in perfect coherence among the edited views. In conclusion, both the Cross-attention and Editing Consistency Modules play vital roles in addressing multi-view inconsistency, enabling VcEDIT to provide more accurate and detailed 3D editing through consistent image guidance.

In addition, we further compare the edited images produced by our VcEDIT and GSEditor in the *bottom 2 rows* of Fig. 6. The edited guidance images of our VcEDIT employing both CCM and ECM achieves significantly better multi-view consistency than that of using GSEditor, which explains the performance gap.

6.2 Ablation on Iterative Pattern

In Sec. 4.4, we expand our framework to include an iterative pattern to further reduce the adverse effects of multi-view inconsistency. This section analyzes the impact of this iterative approach on the 3DGS editing performance. We specifically re-examine the “man to grandma” and “man to clown” examples by comparing the 3DGS editing outcomes of the first and second iterations. Our findings, presented in Fig. 7, show that: (1) In the “man to grandma” case, the second iteration significantly refines facial details and clothing. (2) In the “man to clown” case, the second iteration further accentuates the “clown” features by



Fig. 7: Edited 3DGS after each iteration: Following the iterative pattern, the edited 3DGS is progressively refined by each iteration.

intensifying the makeup and smoothing facial details. (3) Even in the first iteration, our VCEDIT achieves relatively satisfying multi-view consistency, which is attributed to our effective Consistency Modules. These enhancements confirm that through reciprocal and continuous refinement of 3DGS and image guidance, the iterative pattern progressively steers the samples towards a consistent style that aligns with the desired target, leading to improved editing quality.

6.3 Limitations

In line with earlier efforts in 3D editing [6, 15] that utilize the image-guided editing pipeline, our VCEDIT depends on the 2D editing models to generate image guidance, which leads to two specific limitations: (1) Current diffusion-based image editing models occasionally fail to deliver high-quality image editing for intricate prompts, thus impacting the efficacy of 3D editing. (2) In non-rigid editing scenarios that request drastically changing of object shape, current diffusion models produce highly variant editing results among views, which makes it hard for our consistency modules to rectify the tremendous inconsistency.

7 Conclusion

In this paper, we propose the View-consistent Editing (VCEDIT) framework for 3D Editing using text instruction, aiming to address the multi-view inconsistency issue in image-guided 3DGS editing. By seamlessly incorporating 3DGS into the editing processes of the rendered images, VCEDIT ensures the multi-view consistency in image guidance, effectively mitigating mode collapse issues and achieving high-quality 3DGS editing. Several techniques are proposed to achieve multi-view consistency in the guidance images, including two innovative Consistency Modules integrated in an iterative pattern. Our extensive evaluations across various real-world scenes have demonstrated VCEDIT’s ability to outperform existing state-of-the-art methods, thereby setting a new standard for text-driven 3D model editing. The introduction of VCEDIT solves a critical problem in 3DGS editing, inspiring the exploration of future efforts.

Acknowledgements. This research is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-GC-2022-005). Xuanyu Yi is supported by the Agency for Science, Technology AND Research. Long Chen was supported by HKUST Special Support for Young Faculty (F0927) and HKUST Sports Science and Technology Research Grant (SSTRG24EG04). We also appreciate Ms. Sijing Lin and Dr. Zhongqi Yue for their contributions to the paper writing.

References

1. Bai, J., Dong, Z., Feng, A., Zhang, X., Ye, T., Zhou, K., Shou, M.Z.: Integrating view conditions for image synthesis (2023) [2](#)
2. Bao, C., Zhang, Y., Yang, B., Fan, T., Yang, Z., Bao, H., Zhang, G., Cui, Z.: Sine: Semantic-driven image-based nerf editing with prior-guided editing field. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20919–20929 (2023) [4](#)
3. Cao, Y., Cao, Y.P., Han, K., Shan, Y., Wong, K.Y.K.: Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. arXiv preprint arXiv:2304.00916 (2023) [2](#)
4. Chen, M., Xie, J., Laina, I., Vedaldi, A.: Shap-editor: Instruction-guided latent 3d editing in seconds. arXiv preprint arXiv:2312.09246 (2023) [2](#)
5. Chen, Y., Chen, A., Chen, S., Yi, R.: Plasticine3d: Non-rigid 3d editing with text guidance. arXiv preprint arXiv:2312.10111 (2023) [2](#)
6. 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 (2023) [2](#), [5](#), [8](#), [10](#), [12](#), [14](#)
7. Cheng, X., Yang, T., Wang, J., Li, Y., Zhang, L., Zhang, J., Yuan, L.: Progressive3d: Progressively local editing for text-to-3d content creation with complex semantic prompts. arXiv preprint arXiv:2310.11784 (2023) [2](#)
8. Dihlmann, J.N., Engelhardt, A., Lensch, H.: Signerf: Scene integrated generation for neural radiance fields. arXiv preprint arXiv:2401.01647 (2024) [2](#)
9. Dong, J., Wang, Y.X.: Vica-nerf: View-consistency-aware 3d editing of neural radiance fields (2024) [2](#), [5](#)
10. Fang, J., Wang, J., Zhang, X., Xie, L., Tian, Q.: Gaussianeditor: Editing 3d gaussians delicately with text instructions (2023) [2](#), [5](#), [8](#)
11. Feng, A., Qiu, W., Bai, J., Zhou, K., Dong, Z., Zhang, X., Ying, R., Tassiulas, L.: An item is worth a prompt: Versatile image editing with disentangled control (2024) [2](#)
12. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-Or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion. arXiv preprint arXiv:2208.01618 (2022) [2](#), [4](#), [12](#)
13. Gal, R., Patashnik, O., Maron, H., Bermano, A.H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators. ACM Transactions on Graphics (TOG) **41**(4), 1–13 (2022) [10](#), [11](#), [12](#)
14. Gan, Y., Park, S., Schubert, A., Philippakis, A., Alaa, A.M.: Instructcv: Instruction-tuned text-to-image diffusion models as vision generalists (2023) [2](#), [4](#)
15. Haque, A., Tancik, M., Efros, A.A., Holynski, A., Kanazawa, A.: Instruct-nerf2nerf: Editing 3d scenes with instructions (2023) [2](#), [3](#), [4](#), [8](#), [9](#), [14](#)

16. Hertz, A., Aberman, K., Cohen-Or, D.: Delta denoising score (2023) [4](#), [10](#), [12](#)
17. Kamata, H., Sakuma, Y., Hayakawa, A., Ishii, M., Narihira, T.: Instruct 3d-to-3d: Text instruction guided 3d-to-3d conversion. arXiv preprint arXiv:2303.15780 (2023) [4](#)
18. Karim, N., Khalid, U., Iqbal, H., Hua, J., Chen, C.: Free-editor: Zero-shot text-driven 3d scene editing. arXiv preprint arXiv:2312.13663 (2023) [2](#), [4](#)
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) [2](#), [5](#), [8](#), [9](#)
20. Khalid, U., Iqbal, H., Karim, N., Hua, J., Chen, C.: Latenteditor: Text driven local editing of 3d scenes. arXiv preprint arXiv:2312.09313 (2023) [4](#)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
22. Koo, J., Park, C., Sung, M.: Posterior distillation sampling (2023) [4](#)
23. Li, J., Liu, S., Liu, Z., Wang, Y., Zheng, K., Xu, J., Li, J., Zhu, J.: Instructpix2nerf: Instructed 3d portrait editing from a single image (2024) [5](#)
24. Li, Y., Dou, Y., Shi, Y., Lei, Y., Chen, X., Zhang, Y., Zhou, P., Ni, B.: Focal-dreamer: Text-driven 3d editing via focal-fusion assembly (2023) [4](#)
25. Liu, M., Xu, C., Jin, H., Chen, L., Varma T, M., Xu, Z., Su, H.: One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. Advances in Neural Information Processing Systems **36** (2024) [7](#)
26. Mikaeili, A., Perel, O., Safaei, M., Cohen-Or, D., Mahdavi-Amiri, A.: Sked: Sketch-guided text-based 3d editing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14607–14619 (2023) [4](#)
27. Park, J., Kwon, G., Ye, J.C.: Ed-nerf: Efficient text-guided editing of 3d scene using latent space nerf. arXiv preprint arXiv:2310.02712 (2023) [4](#)
28. Parmar, G., Singh, K.K., Zhang, R., Li, Y., Lu, J., Zhu, J.Y.: Zero-shot image-to-image translation (2023) [2](#), [4](#)
29. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion (2022) [4](#)
30. 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) [11](#)
31. Richardson, E., Metzger, G., Alaluf, Y., Giryes, R., Cohen-Or, D.: Texture: Text-guided texturing of 3d shapes. arXiv preprint arXiv:2302.01721 (2023) [4](#)
32. 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) [4](#), [6](#)
33. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) [7](#)
34. Song, L., Cao, L., Gu, J., Jiang, Y., Yuan, J., Tang, H.: Efficient-nerf2nerf: Streamlining text-driven 3d editing with multiview correspondence-enhanced diffusion models. arXiv preprint arXiv:2312.08563 (2023) [2](#)
35. 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) [2](#)
36. Wang, C., He, M., Chai, M., Chen, D., Liao, J.: Mesh-guided neural implicit field editing (2023) [4](#)

37. Wu, Z., Zhou, P., Kawaguchi, K., Zhang, H.: Fast diffusion model (2023) [4](#), [6](#)
38. Xu, S., Huang, Y., Pan, J., Ma, Z., Chai, J.: Inversion-free image editing with natural language (2023) [2](#), [3](#), [4](#), [6](#), [9](#)
39. Yi, X., Deng, J., Sun, Q., Hua, X.S., Lim, J.H., Zhang, H.: Invariant training 2d-3d joint hard samples for few-shot point cloud recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14463–14474 (2023) [2](#)
40. Yi, X., Tang, K., Hua, X.S., Lim, J.H., Zhang, H.: Identifying hard noise in long-tailed sample distribution. In: European Conference on Computer Vision. pp. 739–756. Springer (2022) [9](#)
41. Yi, X., Wu, Z., Xu, Q., Zhou, P., Lim, J.H., Zhang, H.: Diffusion time-step curriculum for one image to 3d generation (2024) [4](#)
42. Yu, L., Xiang, W., Han, K.: Edit-diffnerf: Editing 3d neural radiance fields using 2d diffusion model. arXiv preprint arXiv:2306.09551 (2023) [4](#)
43. Yuen, S.C.Y., Yaoyuneyong, G., Johnson, E.: Augmented reality: An overview and five directions for ar in education. Journal of Educational Technology Development and Exchange (JETDE) **4**(1), 11 (2011) [2](#)
44. Zhang, K., Chen, Y., Yang, X., Zhang, W., Jin, C.: Point cloud part editing: Segmentation, generation, assembly, and selection (2023) [4](#)
45. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018) [7](#)
46. Zhang, Y., Xu, Q., Zhang, L.: Dragtex: Generative point-based texture editing on 3d mesh (2024) [4](#)
47. Zhou, X., He, Y., Yu, F.R., Li, J., Li, Y.: Repaint-nerf: Nerf editing via semantic masks and diffusion models. arXiv preprint arXiv:2306.05668 (2023) [4](#)
48. Zhuang, J., Wang, C., Liu, L., Lin, L., Li, G.: Dreameditor: Text-driven 3d scene editing with neural fields (2023) [4](#)