


# Geometry-Guided Reinforcement Learning for Multi-view Consistent 3D Scene Editing

Jiyuan Wang<sup>1,2,3</sup>, Chunyu Lin<sup>1,✉</sup>, Lei Sun<sup>2,\*</sup>, Zhi Cao<sup>1</sup>, Yuyang Yin<sup>1</sup>, Lang Nie<sup>4</sup>, Zhenlong Yuan<sup>2</sup>, Xiangxiang Chu<sup>2</sup>, Yunchao Wei<sup>1</sup>, Kang Liao<sup>3</sup>, and Guosheng Lin<sup>3,✉</sup>

<sup>1</sup>BJTU, <sup>2</sup>AMap Alibaba Group, <sup>3</sup>NTU, <sup>4</sup>CQUPT  
{wangjiyuan,cylin}@bjtu.edu.cn



**Fig. 1:** We propose **RL3DEdit**, a novel RL-based model for single-pass 3D editing. Our method achieves high-quality results across diverse editing scenarios, including ①motion edits, ②subject replacement, ③style transfer, ④background changes, and ⑤challenging scene addition.

**Abstract.** Leveraging the priors of 2D diffusion models for 3D editing has emerged as a promising paradigm. However, maintaining multi-view consistency in edited results remains challenging, and the extreme scarcity of 3D-consistent editing paired data renders supervised fine-tuning (SFT), the most effective training strategy for editing tasks, infeasible. In this paper, we observe that, while generating multi-view consistent 3D content is highly challenging, verifying 3D consistency is tractable, naturally positioning reinforcement learning (RL) as a feasible solution. Motivated by this, we propose **RL3DEdit**, a single-pass framework driven by RL optimization with novel rewards derived from the 3D foundation model, VGGT. Specifically, we leverage VGGT’s robust priors learned from massive real-world data, feed the edited images, and utilize the output confidence maps and pose estimation errors as reward signals, effectively anchoring the 2D editing priors onto a 3D-consistent manifold via RL. Extensive experiments demonstrate that RL3DEdit achieves stable multi-view consistency and outperforms state-of-the-art methods in editing quality with high efficiency. To promote the development of 3D editing, we will release the code and model.

Work done during the internship at AMap Alibaba Group, and NTU. ✉ Corresponding author; \*Project Leader.

## 1 Introduction

3D scene editing plays a pivotal role in applications such as AR/VR and gaming, demanding both high-fidelity semantic manipulation and strict geometric coherence. To achieve this, generating multi-view consistent images via 2D editors, followed by fine-tuning 3D representations such as 3D Gaussian Splatting (3DGS), has emerged as a promising paradigm for 3D scene editing. Especially, the recent 2D editing models (*e.g.*, FLUX-Kontext [32]) achieved remarkable progress in editing fidelity and multi-image correlated editing [59]. Despite this promising foundation, current 3D editing methods not only underexploit these powerful models, but also remain three primary limitations: (i) Geometric-conditioned methods [60, 73] guide 3D consistency with depth maps of the source image, failing to handle edits involving geometric changes; (ii) Optimization-based approaches [23, 49] iteratively refine 3D representations with single-view edits, suffering from low efficiency and blurry artifacts due to 3D-inconsistent signals. (iii) Attention-based models [11, 56] reproject attention features across viewpoints, yet struggle to guarantee fine-grained geometric consistency.

Particularly, attention manipulation is similarly a suboptimal solution in the early stage of 2D editing [24], but is ultimately surpassed by data-driven supervision [32], suggesting that supervision remains the most effective pathway. However, the extreme scarcity of 3D-consistent editing paired data makes this approach infeasible. Recently, reinforcement learning (RL) [22] has demonstrated strong potential in advancing 2D editing models [39], offering a promising alternative to address the above challenges. Instead of relying on explicit supervision from carefully constructed datasets, RL algorithms optimize editing models using feedback from verifiable reward models (VRMs). In the 3D area, verifying multi-view consistency is significantly more tractable than generating consistent images, making RL a promising approach to acquire 3D consistency without massive paired data.

Motivated by this, the key remaining problem is to identify a robust 3D verifier that can effectively measure the consistency of editing results. In this paper, we leverage the 3D foundation model, VGGT [54], to serve as this verifier. Drawing an analogy to Score Distillation Sampling (SDS [45])—which leverages a frozen 2D diffusion model to assess image quality—we argue that a frozen VGGT, trained on massive real-world 3D data, can provide intrinsic feedback on multi-view consistency. Through empirical analysis, we observe that its confidence maps, originally designed for error tolerance, can serve as an effective proxy for evaluating consistency across views. Furthermore, the camera pose prediction can also offer explicit feedback on viewpoint arrangement. To preserve the 2D editor’s editing fidelity, we additionally design an anchor strategy that aligns edited outputs with high-quality single-view references. Unlike traditional 3D verifiers that are easily “reward-hacked” [58] by textureless or blurry images, the 3D foundation model acts as a robust, geometry-aware reward model by leveraging data-driven priors, thus providing stable guidance for the RL process.

Based on the above analysis, we propose **RL3DEdit**, a single-pass framework that augments a 2D editor’s 3D consistency prior with RL guided by geometry-



aware rewards. To satisfy the prerequisite of RL optimization, we explore base editors’ multi-image joint editing capabilities and select FLUX-Kontext, whose inherent cross-view attention lays the foundation for achieving consistency. During training, the RL process explores diverse editing candidates and evaluates their consistency via the VGGT reward model. The RL algorithm, GRPO [22], optimizes the editing model toward predicting 3D consistent results. At inference, the edited multi-view images are reconstructed into 3DGS to yield the final edited 3D scene.

Our method requires no per-scene/prompt fine-tuning and effectively preserves the powerful 2D foundation’s capabilities after RL optimization. Trained on limited samples, it successfully learns 3D consistency priors and shows promising generalization to unseen conditions. Moreover, RL3DEdit can handle geometry-changing instructions, and avoids iterative optimization, enabling single-pass inference that is over  $2\times$  faster than previous methods. Extensive experiments confirm that our approach achieves superior editing quality and multi-view consistency with high efficiency. Our contributions can be summarized as follows:

- We propose a novel 3D editing RL framework, which effectively empowers 2D editors with 3D capabilities through a tractable 3D consistency verifier, thereby bypassing the scarcity of paired training data.
- We identify that the 3D foundation model with data-driven priors, like VGGT, can serve as a superior verifier. Furthermore, we explore its usage and design tailored rewards to enforce geometric consistency and preserve editing quality.
- Our optimization-free RL3DEdit model achieves SoTA 3D editing quality with high efficiency.

## 2 Related Work

### 2.1 2D Image Editing Models

In the early stage, 2D editing faced a similar dilemma to current 3D editing: the scarcity of paired editing data. To address this, pioneering works manipulated the cross-attention maps of generative models [9, 24], achieving fine-grained editing. Subsequently, with the advent of paired editing datasets, several methods began directly training models to follow explicit editing instructions [6, 42, 69], significantly improving editing quality. Recently, with the support of massive data, high-quality unified editing models have emerged [32, 59], providing more powerful backbones for 3D editing.

### 2.2 3D Editing Models

3D editing can be broadly categorized into object editing and scene editing. Compared to objects, scenes are significantly more challenging due to the complex backgrounds and multiple entities.

**3D Object Editing.** Recent works [36, 61, 66] have achieved significant progress in 3D object editing by constructing dedicated datasets and fine-tuning on the voxel foundation model – Trellis [62]. However, scene editing is difficult to represent with voxels and poses greater challenges in data construction.

**3D Scene Editing.** Scene editing typically adopts 3DGS or NeRF [3] as the 3D representation. We categorize existing methods into four classes:

(i) *SDS-based methods* [13, 26, 29, 37, 44, 49, 76] leverage diffusion priors via Score Distillation Sampling (SDS) to optimize 3D representations, but typically suffer from blurry textures and over-smoothing.

(ii) *Iterative optimization-based methods*, pioneered by IN2N [23] and extended by [12, 26, 53, 57, 60, 63, 68], alternate between single-view editing and 3D optimization. However, the cross-view information is lacking in the editing process, which leads to inconsistent 3D representation optimization and prolonged editing time.

(iii) *Gaussian parameter manipulation methods* [47, 71] directly modify 3DGS parameters via semantic grounding or parameter-change predictors distilled from 2D editors. While efficient, these methods struggle with action-based edits (*e.g.*, “bowing head”), as such instructions are difficult to translate into explicit Gaussian parameter modifications.

(iv) *Multi-view consistent editing methods* have become the dominant paradigm. Similar to 2D editing, early works [11, 16, 19, 31, 56, 60] propagate, interact, or project attention features across views, yet accumulate alignment errors in geometrically discontinuous regions. The insufficient constraints also lead to residual fine-grained inconsistencies. Subsequent methods like EditSplat [10, 28, 34] employ multi-view fusion guidance, conditioning each newly edited view on adjacent/all previously edited ones. But they still produce visible artifacts under instructions that are difficult to quantify precisely (*e.g.*, “open mouth”—but how wide?).

Furthermore, most aforementioned methods rely on InstructPix2Pix, strictly limiting their 2D editing upper bound. While Tinker [73] recently adopted the powerful FLUX-Kontext, it relies on depth-map guidance (restricting it to geometry-preserving edits), introduces complex video-model pipelines, and demands massive paired data ( $\sim 25K$  samples). In contrast, our RL3DEdit achieves superior generalization with only 5% of the data via RL optimization.

### 2.3 Reinforcement Learning for 3D Tasks

To our knowledge, RL3DEdit is the first work to introduce RL into 3D editing. The RL paradigm, particularly GRPO [22], has achieved remarkable success in LLMs [21] and 2D editing [39], yet remains underexplored in the 3D domain and is mainly focused on 3D generation. Early works [55, 65] train human-preference reward models and employ RLHF to better align generated 3D assets with human intent. More recently, Nabla-R2D3 [40]/AR3D-R1 [52] adopt GRPO to enhance native 3D diffusion/autoregressive generation pipelines, improving texture detail and overall fidelity. Additionally, several works [8, 43, 70, 75] leverage RL




Fig. 2: Pipeline of RL3DEdit. Section 3.1 details the pipeline.

for 3D understanding and spatial layout, applying ensemble critics and physical feedback to enforce spatial and geometric constraints.

### 3 Methods

#### 3.1 3D Editing Pipeline with Reinforcement Learning

As shown in Fig. 2, we construct a novel 3D editing method based on RL. Specifically, given a 3D asset to be edited, we first render it from  $M$  viewpoints to obtain  $\{I_m\}_{m=1}^M$ , and feed them simultaneously into a 2D editor (denoted as  $\pi$ ) for joint multi-view editing. Ideally, at inference, the fine-tuned editor produces multi-view consistent images in a single forward pass, which are then fed into 3DGS reconstruction to obtain the final edited 3D scene, yielding an efficient editing workflow.

Therefore, the core challenge is to equip the 2D editor with a 3D-consistency prior. We address this via RL optimization. As shown in the upper part of Fig. 2 (FLUX-Kontext branch), during training, the GRPO algorithm explores a group of  $G$  edited results (through  $G$  independent inference passes [39]), each containing  $M$  edited views  $\{I'_m\}_{m=1}^M$ . To explicitly enforce both editing faithfulness and multi-view coherence, we first select an anchor view  $I'_a$  and substitute it with our pre-edited anchor image  $\tilde{I}_a$  (detailed in Sec. 3.4). Together with the other views, they are fed into a dedicated 3D-aware reward model (implemented via VGGT [54]), which is designed to jointly assess three critical aspects of multi-view consistency  $r^D, r^P, r^T$  and one term for editing quality  $r^a$ . These complementary rewards are combined to form the final composite reward  $R^i$ , which effectively guides our optimization toward consistent and high-quality



**Fig. 3:** Comparison of 2D editing capabilities before and after RL fine-tuning. Left: Visual editing results. Right: Quantitative evaluation using VIEScore [30] on GEdit-Bench-EN [41] ( $\uparrow$ , detailed in Sec. 4.2). Both demonstrate that RL3DEdit successfully preserves the original 2D editing fidelity of FLUX-Kontext.

3D-aware editing. The rewards  $\{R^i\}_{i=1}^G$  are used to compute the relative advantage via:  $A^i = (R^i - \text{mean}(\{R^j\}_{j=1}^G)) / \text{std}(\{R^j\}_{j=1}^G)$ , and the model is optimized by maximizing the following objective, which encourages the model to produce high-reward outputs:

$$J(\theta) = J_{\text{clip}}(\theta) - \beta D_{\text{KL}}(\pi_{\theta} \parallel \pi_{\text{ref}}),$$


where  $\pi_{\theta}/\pi_{\text{ref}}$  denote the fine-tuned/original 2D editors,  $J_{\text{clip}}(\theta)$  is computed from  $A^i$  (see Appendix for details and the formal definition of GRPO). In this way, the model can learn 3D-consistency priors without curated paired supervision.

Compared with previous methods, our method offers **two** key advantages. On one hand, verifying multi-view consistency only requires identifying geometric contradictions like ghosting artifacts. This is significantly more tractable than generating consistent images, which requires complex cross-view interactions and relies on scarce multi-view data. This asymmetry perfectly aligns with the RL paradigm, where consistency can naturally serve as the reward. On the other hand, 3D editing, like its 2D counterpart, must handle arbitrary instructions on diverse scenes. Learning such general editing priors in 3D would require prohibitively large paired datasets that currently do not exist. Fortunately, modern 2D foundation editors naturally possess this capability. As shown in Fig. 3, the fine-tuned editor preserves its original 2D editing capability, suggesting that our method primarily augments the 2D model with 3D-consistency priors rather than reshaping it, further highlighting the effectiveness of our method.

In the following, we detail the three core components of our pipeline: the 2D editor (Sec. 3.2), the 3D-aware reward model (Sec. 3.3), and the rewards design (Sec. 3.4).

### 3.2 Multi-Image Joint Editing

Multi-view consistent editing requires effective cross-view interaction during the editing process. In the RL paradigm, optimization relies on exploring diverse outputs and reinforcing those with high rewards. If a 2D editor processes each




**Fig. 4:** Multi-image joint editing comparison. FLUX-Kontext and Qwen-Image-Edit successfully swap the fur colors, while InstructPix2Pix fails due to the lack of cross-view interaction. Moreover, InstructPix2Pix must resize images to low resolution, causing detail loss in multi-image scenarios.

view independently, the probability of randomly producing mutually consistent multi-view images is essentially zero. Consequently, RL would have no successful samples to reinforce, causing the optimization to fail. Therefore, multi-image joint editing capability is a necessary prerequisite for RL to be effective.

Previous 3D editing methods typically adopt InstructPix2Pix [7] as their 2D backbone. However, its low processing resolution and local convolutional operations limit cross-image interaction. As shown in Fig. 4, given the instruction to swap the fur colors of two cats, InstructPix2Pix fails to access information from the other images and results in random color changes. In contrast, recent DiT-based models (*e.g.*, FLUX-Kontext [32], Qwen-Image-Edit [67]) are naturally suited for multi-image editing, as their Transformer architecture enables global attention across all inputs. Specifically, after tokenizing  $K$  input images as  $\{x^k\}_{k=1}^K$ , they are concatenated along the sequence dimension as  $X = \text{Concat}(x^1, \dots, x^K)$  and processed with self-attention, enabling attention weights to be distributed across images. This inherent mechanism facilitates efficient cross-view interaction, enabling RL to leverage it to achieve 3D consistency. Based on this, we adopt FLUX-Kontext as our baseline and further validate the applicability to Qwen-Image-Edit in ablation studies (Sec. 4.3).

### 3.3 Multi-View Consistent Verification

Having the RL framework and promising foundation editor, the remaining question is how to robustly verify the 3D consistency of edited images. In this paper, we choose a 3D foundation model, VGGT [54]. This is inspired by Score Distillation Sampling (SDS [46]), which employs the 2D foundation model Stable Diffusion (SD [48]) to supervise image quality. SD is trained on large-scale high-quality images, so that feeding it low-quality images produces a large loss of feedback. Analogously, VGGT is trained on millions of real-world 3D data samples, and feeding multi-view inconsistent edited images into VGGT can also produce meaningful corresponding feedback. Below, we explore this feedback form and employ VGGT as our reward model. Experiments (Sec. 4.3) show that such data-prior-based verification is better than traditional methods (*e.g.*, reprojection warping, Structure-from-Motion) in avoiding “reward-hacking”.



**Fig. 5:** Empirical analysis of VGGT’s depth confidence under progressively degraded 3D consistency. ①–⑤ visualize VGGT confidence predictions for the **same** set of 9 views, where individual views are gradually replaced by edited versions. ⑥ reveals a near-linear correlation between consistency degradation and average confidence. This validates VGGT as the multi-view consistency verifier. Detailed analysis is in Sec. 3.3.

**3D Foundation Model.** Formally, given  $N$  views  $(I_i)_{i=1}^N$  of the same 3D scene, VGGT learns a Transformer  $f$  such that  $f((I_i)_{i=1}^N) = (g_i, D_i, P_i, \tau_i)_{i=1}^N$ , where  $g_i, D_i, P_i, \tau_i$  denote the camera parameters, depth map, point map, and tracking features of the  $i$ -th view, respectively. During training, to handle varying regional difficulty, VGGT jointly predicts confidence maps with error tolerance. The loss function is represented as follows:

$$\mathcal{L}_{\text{geo}} = \|\text{conf}_{\text{geo}} \odot (\text{gêo} - \text{geo})\|_1 + \|\text{conf}_{\text{geo}} \odot (\nabla \text{gêo} - \nabla \text{geo})\|_1 - \alpha \|\log \text{conf}_{\text{geo}}\|_1,$$

where  $\text{conf}_{\text{geo}}$  denotes geometric confidence, including the uncertainty of depth ( $\text{conf}_D$ ) and point ( $\text{conf}_P$ ) predictions.

**Confidence Map Analysis.** To investigate the feedback, as illustrated in Fig. 5, we render 9 views  $\{I_i\}_{i=1}^9$  from a 3D asset and independently edit them to obtain  $\{I'_i\}$ . We then gradually replace  $I_i$  with its edited counterpart  $I'_i$  to disrupt multi-view consistency, and observe the changes in VGGT’s output. Our analysis reveals a strong correlation between the predicted confidence and 3D consistency: ① When all views are originally 3D-consistent, the depth confidence is uniformly high. ② Replacing the background of a single view, its background confidence drops significantly while the foreground character remains unaffected. ③ Further adding a pair of glasses in another view reduces the confidence in the corresponding eye region. ④ Replacing the person with Iron Man in a third view leads to globally low confidence, as both the background and foreground now exhibit cross-view inconsistencies. ⑤ Finally, when 8 out of 9 views are replaced, although all inputs conceptually belong to the same scene, they lack mutual 3D consistency, resulting in very low confidence scores. To quantitatively validate this, we select 100 sets of 9-view data and compute the average confidence under varying replacement ratios. As shown in ⑥, we observe that the predicted confidence decays almost linearly as consistency decreases. This empirical result confirms that VGGT’s confidence maps are a reliable and interpretable indicator



of multi-view consistency. Therefore, VGGT’s confidence can effectively serve as an indicator for 3D consistency and act as a reward.

### 3.4 Overview of Reward Model

**Geometric Rewards.** As discussed above, we use the average depth and point confidence as geometric rewards:

$$r^D = \frac{1}{M} \sum_{m=1}^M \text{mean}(\text{conf}_D^m), \quad r^P = \frac{1}{M} \sum_{m=1}^M \text{mean}(\text{conf}_P^m). \quad (1)$$

**Relative Pose Reward.** Beyond multi-view consistency, we also consider the arrangement of camera viewpoints. Considering that editing may alter the absolute camera viewpoint, but the inter-view relationship is still preserved, we use the relative poses between adjacent views to measure camera perspective alignment. Let the VGGT-predicted extrinsics of view  $m$  be  $T_m = [R_m | t_m]$ , and define the relative transform  $T_m^{\text{rel}} = T_{m+1} T_m^{-1} = [R_m^{\text{rel}} | t_m^{\text{rel}}]$ . We normalize the translation as  $\tilde{t}_m^{\text{rel}} = t_m^{\text{rel}} / \|t_m^{\text{rel}}\|_2$  and use the reference (GT) relative pose  $(T_m^{\text{rel}})^*$ , yielding the reward:

$$r^T = \exp \left( -\frac{1}{M-1} \sum_{m=1}^{M-1} \left( \|R_m^{\text{rel}} - (R_m^{\text{rel}})^*\|_F^2 + \|\tilde{t}_m^{\text{rel}} - (\tilde{t}_m^{\text{rel}})^*\|_2^2 \right) \right). \quad (2)$$

**Anchor Reward.** As discussed in Sec. 3.1, RL3DEdit preserves the original 2D editing capability while enhancing geometric priors. This not only benefits from the RL fine-tuning paradigm [39], but also from the proposed anchor reward. We pre-compute single-image editing results for all views offline using FLUX-Kontext, followed by light quality filtering (empirically, >98% are directly usable). Although these results  $\tilde{I}$  are mutually 3D-inconsistent, they all satisfy FLUX-Kontext’s high-quality distribution (*i.e.*, correct semantics, detail preservation, etc.). So we can leverage them independently to guide the RL optimization and preserve FLUX-Kontext’s editing fidelity.

During training, we randomly sample an anchor index  $a \in \{1, \dots, M\}$  and retrieve  $\tilde{I}_a$  from the offline results (one GRPO-group shares the same index to ensure fair group-reward comparison; randomness is at the sample level). To evaluate the editing quality of  $\{I'_m\}_{m=1}^M$ , we consider two cases: (1) For the anchor view  $I'_a$ , we directly measure the editing error:

$$r^a = \exp \left( -\lambda \mathcal{L}_{\text{LPIPS}}(I'_a, \tilde{I}_a) \right), \quad (3)$$

where  $\mathcal{L}_{\text{LPIPS}}(\cdot, \cdot)$  denotes the perceptual similarity [72]. (2) For other views, we replace  $I'_a$  with  $\tilde{I}_a$  in the multi-view input, feeding  $\{I'_1, \dots, I'_{a-1}, \tilde{I}_a, I'_{a+1}, \dots, I'_M\}$  into VGGT to evaluate consistency. In this way,  $I'_{i \neq a}$  are optimized for higher geometric consistency rewards while  $I'_a$  is optimized for higher anchor reward, ensuring all views converge toward correct semantics and visual details, thereby preserving FLUX-Kontext’s editing fidelity.

Finally, the  $i$ -th overall reward in the GRPO-group is defined as:

$$R^i = w_D r^D + w_P r^P + w_T r^T + w_a r^a, \quad (4)$$

which we plug into our framework to optimize FLUX-Kontext.

## 4 Experiments

### 4.1 Implementation Details

We adopt FLUX-Kontext-dev [32] as our baseline and fine-tune it using LoRA with rank 32 and alpha 32. During training, we set  $M = 9$  and  $w_D = w_P = w_T = w_a = 0.25$ . Following Flow-GRPO [39], we employ Stochastic Differential Equations (SDE) to enhance exploration randomness, setting the SDE noise level to 0.8 and the group size to 16. However, unlike the 6-step denoising exploration used in prior work [39], our experiments show that 3D consistency demands higher image fidelity, so we adopt a 12-step exploration. To accelerate convergence, we incorporate improvements from MixGRPO [35], setting the window size to 4.

For training data, we collect 8 scenes from the IN2N [23], BlendedMVS [64], and Mip-NeRF360 [5] datasets, and construct 7~9 editing prompts per scene using a VLM [14], yielding 70 prompts in total (Appendix for details). For each scene, we sample 18~23 sets of  $M$ -view images as training samples, where the  $M=9$  views are evenly sampled around the scene to ensure sufficient visual overlap for consistency evaluation, yielding a total of 1,319 samples. Training was conducted for one epoch on an NVIDIA RTX A6000 GPU and took 42 hours. In inference, we employ 3DGS to reconstruct the edited 3D scene from the multi-view edited images (Appendix for details).

### 4.2 Comparison Analysis

**Comparative Models.** We compare RL3DEdit with the *open-source* SoTA 3D editing methods, including DGE [11], EditSplat [34], and GaussCtrl [60]. We note that previous methods adopt InstructPix2Pix [6] as the 2D editing model; however, as discussed earlier, its limitations in multi-image joint editing make it unsuitable as our baseline. For a fair comparison, we re-implement the strongest baseline, EditSplat, under the same backbone FLUX-Kontext (see Appendix for implementation details) and compare it with our proposed method. For each scene, we use identical text prompts for editing, then render the edited 3D scene from the same camera poses to obtain results from different methods at consistent novel viewpoints.


**Quantitative Comparison and Metrics.** We evaluate four dimensions: (i) *VIEScore* [30], a VLM-based metric (GPT-4.1 [1]) that jointly evaluates instruction-following and visual quality—we pioneer its usage in 3D editing to provide reproducible, objective assessment free from the subjective variance of user studies; (ii) *CLIP directional similarity* [17], following prior work [34, 49];

**Table 1:** Quantitative comparison with SoTA methods. Ph-Loss stands for photometric reprojection loss,  $\uparrow / \downarrow$  indicates higher/lower is better. Best results are highlighted in **bold**. Editing time is tested on an RTX A6000 GPU. Due to space constraints, detailed per-split metrics (in-distribution/zero-shot) and user study are provided in Appendix.

| Methods                   | VIEScore $\uparrow$<br>(Novel View Synthesis) | CLIP-dir $\uparrow$ | Ph-Loss $\downarrow$<br>(Edit-Results) | Avg. Editing<br>Time $\downarrow$ |
|---------------------------|-----------------------------------------------|---------------------|----------------------------------------|-----------------------------------|
| DGE [11]                  | 2.81                                          | 0.116               | 0.086                                  | 4min                              |
| GaussCtrl [60]            | 2.37                                          | 0.096               | 0.077                                  | 12min                             |
| EditSplat [34]            | 2.72                                          | 0.121               | 0.081                                  | 3.5min                            |
| EditSplat w/ FLUX-Kontext | 3.23                                          | 0.125               | 0.082                                  | 40min                             |
| RL3DEdit (Ours)           | <b>5.48</b>                                   | <b>0.147</b>        | <b>0.076</b>                           | <b>1.5min</b>                     |

(iii) *photometric reprojection loss* (Ph-Loss) [18] for multi-view consistency (definition in Appendix); and (iv) *average editing time*. The test data includes novel views (70 cases), unseen instructions (16 cases), and new scenes (14 cases), totaling 100 test cases (Appendix for detailed split). We note that all prior methods perform per-scene optimization, inherently using the same scene and instruction for both training and testing. Our 70 cases follow the same protocol for fair comparison. For the additional 30 cases, only RL3DEdit operates in a zero-shot setting, while others still optimize on corresponding data. As shown in Tab. 1, RL3DEdit achieves the best performance across all dimensions. It not only achieves a remarkable leap in editing fidelity and semantic alignment (VIEScore of **5.48** vs. 3.23), but also maintains the lowest Ph-Loss for rigorous 3D consistency. Crucially, this high-quality editing is achieved in just **1.5 minutes**—over  $2\times$  faster than traditional pipelines and over  $20\times$  faster than the other FLUX-based baseline.

**Qualitative Comparison.** Metrics alone cannot fully capture editing quality. We therefore present extensive visual results across Figs. 1, 2, 6, 7, and 8, covering diverse scenes and instructions; additional results are provided in the supplementary video. Specifically, Fig. 6 presents qualitative comparisons under some challenging instructions. (Row 1) EditSplat/GaussCtrl relies on depth warping/guiding, causing severe failures under geometric-changing prompts. DGE misinterprets the semantics. Only RL3DEdit successfully places the ball in front of the bear. (Row 2) Compared to the blurry results of other methods, RL3DEdit achieves a more realistic Minecraft-style appearance. (Row 3) Motion editing often causes artifacts due to the difficulty in quantifying action magnitude. DGE and EditSplat exhibit artifacts around the mouth, while GaussCtrl alters subject identity. Only RL3DEdit produces correct, high-quality results. (Row 4) For winter scene conversion, DGE merely whitens the scene; GaussCtrl generates snow but alters the dinosaur; EditSplat suffers ghost-artifacts, as massive multi-view inconsistency causes warping failure. Only RL3DEdit achieves semantically accurate editing. These improvements on challenging instructions vividly illustrate the performance advantages shown in Tab. 1, further demonstrating the effectiveness of our method.



**Fig. 6:** Qualitative comparison. RL3DEdit achieves much higher 3D editing quality than other methods, in addition (1st row)/transform (2nd row)/motion (3rd row)/style (4th row) editing.

### 4.3 Ablation Study


Due to computational constraints, the ablation models are all trained and tested on face scenes with over 200 samples to validate the effectiveness of each component. Experiments confirm that this data scale is sufficient to draw conclusions.

**Effect of Rewards.** Rewards  $r^D$  and  $r^P$  (depth and point confidence) both constrain 3D consistency, so we ablate them together. ① Removing  $r^D$  and  $r^P$  causes significant degradation in both novel-view editing quality and reprojection loss (Tab. 2). Fig. 7 shows severe ghosting artifacts due to inconsistency, demonstrating the importance of VGGT consistency rewards. ② Without  $r^T$ , subtle viewpoint shifts occur, which become evident in reprojection evaluation using GT extrinsics (Tab. 2). Fig. 7 shows displacement in wall details. ③ Without the anchor image to preserve editing priors, outputs degrade toward over-smoothed results (Fig. 7), primarily because 3D consistency is easier to achieve with low-frequency details, causing RL to optimize in this direction and fail to preserve FLUX-Kontext’s editing quality.

**Effect of Alternative Consistency Verifiers.** We compare VGGT against two widely used 3D-consistency measures to validate its necessity as a reward.

**Table 2:** Quantitative ablation study on reward components and alternative designs. RL3DEdit\* denotes the model trained with the same method but only on face data. The model with  $r^{\text{warp}}$  produces blurry images that achieve high consistency but not the desired editing quality. The extension to Qwen-Image-Edit demonstrates that stronger 2D editing models yield better results.

| Methods                         | VIEScore $\uparrow$ | Ph-Loss $\downarrow$ |
|---------------------------------|---------------------|----------------------|
| ① w/o ( $r^D, r^P$ )            | 2.11                | 0.193                |
| ② w/o $r^T$                     | 4.77                | 0.131                |
| ③ w/o $r^A$                     | 4.34                | 0.091                |
| ④ replaced w/ $r^{\text{SfM}}$  | 0.97                | 0.201                |
| ⑤ replaced w/ $r^{\text{warp}}$ | 1.41                | 0.065                |
| ⑥ Qwen-Image-Edit               | 5.43                | 0.079                |
| ⑦ RL3DEdit*                     | 5.26                | 0.077                |



**Fig. 7:** Qualitative ablation results. To better illustrate the causes of degradation, all variants use the same viewpoint for edited results except w/o ( $r^D, r^P$ ), which uses newly rendered views. In w/o  $r^T$ , red dashed lines highlight subtle viewpoint shifts compared to the original.

④ *SfM-based reward.* Structure-from-Motion requires no learned priors but relies on sparse feature matching (definition in Appendix). During RL training, the model quickly learns to produce textureless outputs that yield few matchable keypoints, trivially satisfying the SfM consistency check while destroying editing quality (Fig. 7, Tab. 2).

⑤ *Reprojection warping reward.* We compute photometric reprojection loss (Ph-Loss) [18] using VGGT-predicted depth and GT poses as the reward signal (definition in Appendix). Tab. 2 shows that the resulting model achieves the lowest Ph-Loss among all variants, yet produces severely blurred outputs (Fig. 7), confirming that Ph-Loss is easily “reward-hacked” by low-frequency images. It is worth noting that we still adopt Ph-Loss as an evaluation metric in Tab. 1, as it remains effective for measuring consistency among sharp images, while any blurriness is penalized by VIEScore, ensuring the fairness of our comparison.

Compared to these traditional methods, VGGT is trained on massive real-world 3D images. As a reward function, we think it can optimize consistency while preserving the capability to output high-quality images (further discussed



**Fig. 8:** Zero-shot editing results, where ①,②,③ are new instructions; ④,⑤ are new scenes. More results are available in the supplementary video.

in Appendix). This validates the necessity of using data-driven priors like VGGT as reward models.

**Extensions on Qwen-Image-Edit.** Both results in Tab. 2 ⑥ and Fig. 7 ⑥ demonstrate that our proposed framework can be further enhanced with more powerful editing models. With the rapid advancement of 2D editing models, our framework has strong potential.

**Zero-Shot Generalization.** As shown in Fig. 8, by preserving the priors of FLUX-Kontext, our method can generalize to unseen instructions and scenes.

## 5 Limitations and Future Work

**Limitations of the 2D Backbone.** Since RL3DEdit is fine-tuned on a 2D editor, its performance is bounded by the backbone’s inherent constraints. The primary limitation stems from attention sequence length: multi-view images share the same token capacity, forcing a trade-off between the number of views and per-image resolution. However, this is not an inherent defect of our framework. A natural extension is to use the anchor image as guidance and generate edited images in batches to cover more viewpoints, which we leave as future work. Moreover, with the rapid advancement of efficient attention mechanisms—especially streaming and causal attention, which have been successfully applied to long-sequence 3D perception [33]—context length continues to expand. As demonstrated in Sec. 4.3, our method can transfer to other backbones, and this limitation is also expected to diminish as foundation models evolve.

**Large Structural Deformations.** RL3DEdit may struggle with extremely drastic non-rigid deformations (*e.g.*, “make the person bow his head”), where the ambiguity in action magnitude leads to subtle multi-view discrepancies that compromise 3D reconstruction. We note that all baselines also fail on such cases, and enhancing generative priors for severe non-rigid deformations remains an open direction.

**Training Scale.** Our training scale is limited, primarily due to the computational overhead of GRPO—each sample requires exploring 16 candidate groups, with each group demanding a 12-step inference pass, resulting in  $\sim 2$  days of

training under our current setup. Nevertheless, RL3DEdit already achieves superior performance, and we believe the method can further improve with increased training scale.

## 6 Conclusion

We present RL3DEdit, an efficient 3D scene editing framework based on RL. Our core insight is that while generating 3D-consistent images is challenging, verifying 3D consistency is tractable, making RL an ideal solution. Building on this, we leverage the 3D foundation model VGGT to construct geometry-aware rewards and employ GRPO to effectively anchor the 2D editor’s prior onto the 3D consistency manifold. Experiments demonstrate that RL3DEdit achieves superior editing quality and multi-view consistency with minimal training data, while delivering over  $2\times$  speedup compared to existing methods. Our framework generalizes well and seamlessly transfers to other 2D editing models, offering a new paradigm for future 3D editing research.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) under Grant (62573039, U2441242).

## 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. Bao, Y., Ding, T., Huo, J., Li, W., Gao, Y.: Fisn: Finding spatial neighborhoods for generalizable novel view synthesis. *IEEE Transactions on Visualization and Computer Graphics* (2026)
3. Bao, Y., Ding, T., Huo, J., Li, W., Li, Y., Gao, Y.: Insertnerf: Instilling generalizability into nerf with hypernet modules. In: *International Conference on Learning Representations* (2024)
4. Bao, Y., Liao, J., Huo, J., Gao, Y.: Distractor-free generalizable 3d gaussian splatting. In: *The Fourteenth International Conference on Learning Representations* (2026), <https://openreview.net/forum?id=G33Iemmj3Z>
5. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields (2022), <https://arxiv.org/abs/2111.12077>
6. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions (2023), <https://arxiv.org/abs/2211.09800>
7. Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18392–18402 (2023)
8. Bucher, M.J., Armeni, I.: Respace: Text-driven autoregressive 3d indoor scene synthesis and editing (2025), <https://arxiv.org/abs/2506.02459>



9. Chefer, H., Alaluf, Y., Vinker, Y., Wolf, L., Cohen-Or, D.: Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models (2023), <https://arxiv.org/abs/2301.13826>
10. Chen, J.K., Bulò, S.R., Müller, N., Porzi, L., Kotschieder, P., Wang, Y.X.: Consistdreamer: 3d-consistent 2d diffusion for high-fidelity scene editing (2024), <https://arxiv.org/abs/2406.09404>
11. Chen, M., Laina, I., Vedaldi, A.: Dge: Direct gaussian 3d editing by consistent multi-view editing. arXiv preprint arXiv:2404.18929 (2024)
12. 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), <https://arxiv.org/abs/2311.14521>
13. 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 (2024), <https://arxiv.org/abs/2310.11784>
14. Comanici, G., Bieber, E., Schaeckermann, M., et al., I.P.: Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities (2025), <https://arxiv.org/abs/2507.06261>
15. Dang, J., Wu, X., Yan, H., Zheng, H., Zheng, W.S., Lai, J., Hu, B., Chua, T.S.: Decoupled seg tokens make stronger reasoning video segmenter and grounder. IEEE Transactions on Pattern Analysis and Machine Intelligence (2026)
16. Dong, J., Wang, Y.X.: Vica-nerf: View-consistency-aware 3d editing of neural radiance fields (2024), <https://arxiv.org/abs/2402.00864>
17. Gal, R., Patashnik, O., Maron, H., Chechik, G., Cohen-Or, D.: Stylegan-nada: Clip-guided domain adaptation of image generators (2021), <https://arxiv.org/abs/2108.00946>
18. Godard, C., Aodha, O.M., Firman, M., Brostow, G.: Digging into self-supervised monocular depth estimation (2019), <https://arxiv.org/abs/1806.01260>
19. Gomel, E., Wolf, L.: Diffusion-based attention warping for consistent 3d scene editing (2024), <https://arxiv.org/abs/2412.07984>
20. Gong, N., Li, Z., Dong, S., Bai, H., Ying, W., Wang, X., Fu, Y.: Sculpting features from noise: Reward-guided hierarchical diffusion for task-optimal feature transformation. arXiv preprint arXiv:2505.15152 (2025)
21. Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, e.a.: Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. Nature **645**(8081), 633–638 (Sep 2025). <https://doi.org/10.1038/s41586-025-09422-z>, <http://dx.doi.org/10.1038/s41586-025-09422-z>
22. Guo, D., Yang, D., Zhang, H., Song, J., Wang, P., Zhu, Q.e.a.: Deepseek-r1 incentivizes reasoning in llms through reinforcement learning. Nature **645**(8081), 633–638 (Sep 2025). <https://doi.org/10.1038/s41586-025-09422-z>, <http://dx.doi.org/10.1038/s41586-025-09422-z>
23. Haque, A., Tancik, M., Efros, A.A., Holynski, A., Kanazawa, A.: Instruct-nerf2nerf: Editing 3d scenes with instructions. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 19740–19750 (2023)
24. Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control (2022), <https://arxiv.org/abs/2208.01626>
25. Jiang, L., Mao, Y., Xu, L., Lu, T., Ren, K., Jin, Y., Xu, X., Yu, M., Pang, J., Zhao, F., et al.: Anysplat: Feed-forward 3d gaussian splatting from unconstrained views. ACM Transactions on Graphics (TOG) **44**(6), 1–16 (2025)

26. Kamata, H., Sakuma, Y., Hayakawa, A., Ishii, M., Narihira, T.: Instruct 3d-to-3d: Text instruction guided 3d-to-3d conversion (2023), <https://arxiv.org/abs/2303.15780>
27. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
28. Khalid, U., Iqbal, H., Farooq, A., Hua, J., Chen, C.: 3dedit: 3d editing on the go! (2024), <https://arxiv.org/abs/2407.10102>
29. Kim, J., Lee, S., Shin, J., Choi, J., Shim, H.: Dreamcatalyst: Fast and high-quality 3d editing via controlling editability and identity preservation (2025), <https://arxiv.org/abs/2407.11394>
30. Ku, M., Jiang, D., Wei, C., Yue, X., Chen, W.: Viescore: Towards explainable metrics for conditional image synthesis evaluation (2024), <https://arxiv.org/abs/2312.14867>
31. Kwon, G., Park, J., Ye, J.C.: Unified editing of panorama, 3d scenes, and videos through disentangled self-attention injection (2024), <https://arxiv.org/abs/2405.16823>
32. Labs, B.F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., Kulal, S., Lacey, K., Levi, Y., Li, C., Lorenz, D., Müller, J., Podell, D., Rombach, R., Saini, H., Sauer, A., Smith, L.: Flux.1 kontext: Flow matching for in-context image generation and editing in latent space (2025), <https://arxiv.org/abs/2506.15742>
33. Lan, Y., Luo, Y., Hong, F., Zhou, S., Chen, H., Lyu, Z., Yang, S., Dai, B., Loy, C.C., Pan, X.: SStream3R: Scalable sequential 3d reconstruction with causal transformer (2025), <https://arxiv.org/abs/2508.10893>
34. Lee, D.I., Park, H., Seo, J., Park, E., Park, H., Baek, H.D., Shin, S., Kim, S., Kim, S.: Editsplat: Multi-view fusion and attention-guided optimization for view-consistent 3d scene editing with 3d gaussian splatting (2025), <https://arxiv.org/abs/2412.11520>
35. Li, J., Cui, Y., Huang, T., Ma, Y., Fan, C., Yang, M., Zhong, Z., Bo, L.: Mixgrpo: Unlocking flow-based grpo efficiency with mixed ode-sde (2025), <https://arxiv.org/abs/2507.21802>
36. Li, L., Huang, Z., Feng, H., Zhuang, G., Chen, R., Guo, C., Sheng, L.: Voxhammer: Training-free precise and coherent 3d editing in native 3d space (2025), <https://arxiv.org/abs/2508.19247>
37. 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), <https://arxiv.org/abs/2308.10608>
38. Li, Z., Yan, H., Li, S., Luo, K., Lu, L., Yang, X., Lin, W.: Diffpcn: Latent diffusion model based on multi-view depth images for point cloud completion. *arXiv preprint arXiv:2509.23723* (2025)
39. Liu, J., Liu, G., Liang, J., Li, Y., Liu, J., Wang, X., Wan, P., Zhang, D., Ouyang, W.: Flow-grpo: Training flow matching models via online rl (2025), <https://arxiv.org/abs/2505.05470>
40. Liu, Q., Liu, Z., Zhang, D., Jia, K.: Nabla-r2d3: Effective and efficient 3d diffusion alignment with 2d rewards (2025), <https://arxiv.org/abs/2506.15684>
41. Liu, S., Han, Y., Xing, P., Yin, F., Wang, R., Cheng, W., Liao, J., Wang, Y., Fu, H., Han, C., Li, G., Peng, Y., Sun, Q., Wu, J., Cai, Y., Ge, Z., Ming, R., Xia, L., Zeng, X., Zhu, Y., Jiao, B., Zhang, X., Yu, G., Jiang, D.: Step1x-edit: A practical framework for general image editing (2025), <https://arxiv.org/abs/2504.17761>

42. Pan, X., Dong, L., Huang, S., Peng, Z., Chen, W., Wei, F.: Kosmos-g: Generating images in context with multimodal large language models (2024), <https://arxiv.org/abs/2310.02992>
43. Pan, Z., Liu, H.: Metaspatial: Reinforcing 3d spatial reasoning in vlms for the metaverse (2025), <https://arxiv.org/abs/2503.18470>
44. Park, J., Kwon, G., Ye, J.C.: Ed-nerf: Efficient text-guided editing of 3d scene with latent space nerf (2024), <https://arxiv.org/abs/2310.02712>
45. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion (2022), <https://arxiv.org/abs/2209.14988>
46. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion (2022), <https://arxiv.org/abs/2209.14988>
47. Qin, H., Sun, Y., Wang, M., Kong, M., Lu, M., Zhu, Q.: Variation-aware flexible 3d gaussian editing (2026), <https://arxiv.org/abs/2602.11638>
48. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2022), <https://arxiv.org/abs/2112.10752>
49. Shu, Z., Yu, J., Chao, K., Xin, S., Liu, L.: Gaussedit: Adaptive 3d scene editing with text and image prompts. *IEEE Transactions on Visualization and Computer Graphics* **31**(10), 7769–7780 (Oct 2025). <https://doi.org/10.1109/tvcg.2025.3556745>, <http://dx.doi.org/10.1109/TVCG.2025.3556745>
50. Sun, Y., Cheng, H., Lu, C., Li, Z., Wu, M., Lu, H., Zhu, J.: Hyperpoint: Multimodal 3d foundation model in hyperbolic space. *Pattern Recognition* (2025)
51. Sun, Y., Zhu, J., Cheng, H., Lu, C., Yang, Z., Chen, L., Wang, Y.: Align then adapt: Rethinking parameter-efficient transfer learning in 4d perception. *IEEE Transactions on Multimedia* (2026)
52. Tang, Y., Guo, Z., Zhu, K., Zhang, R., Chen, Q., Jiang, D., Liu, J., Zeng, B., Song, H., Qu, D., Bai, T., Xu, D., Zhang, W., Zhao, B.: Are we ready for rl in text-to-3d generation? a progressive investigation (2025), <https://arxiv.org/abs/2512.10949>
53. Wang, B., Dutt, N.S., Mitra, N.J.: Proteusnerf: Fast lightweight nerf editing using 3d-aware image context (2024), <https://arxiv.org/abs/2310.09965>
54. Wang, J., Chen, M., Karaev, N., Vedaldi, A., Rupprecht, C., Novotny, D.: VGGT: Visual geometry grounded transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2025)
55. Wang, W., Xu, H., Yang, Y., Liu, Z., Meng, J., Wang, H.: Mvreward: Better aligning and evaluating multi-view diffusion models with human preferences (2024), <https://arxiv.org/abs/2412.06614>
56. Wang, Y., Yi, X., Wu, Z., Zhao, N., Chen, L., Zhang, H.: View-consistent 3d editing with gaussian splatting (2025), <https://arxiv.org/abs/2403.11868>
57. Wang, Y., Yi, X., Wu, Z., Zhao, N., Chen, L., Zhang, H.: View-consistent 3d editing with gaussian splatting (2025), <https://arxiv.org/abs/2403.11868>
58. Weng, L.: Reward hacking in reinforcement learning. [lilianweng.github.io](https://lilianweng.github.io/posts/2024-11-28-reward-hacking/) (Nov 2024), <https://lilianweng.github.io/posts/2024-11-28-reward-hacking/>
59. Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., ming Yin, S., Bai, S., Xu, X., Chen, Y., Chen, Y., Tang, Z., Zhang, Z., Wang, Z., Yang, A., Yu, B., Cheng, C., Liu, D., Li, D., Zhang, H., Meng, H., Wei, H., Ni, J., Chen, K., Cao, K., Peng, L., Qu, L., Wu, M., Wang, P., Yu, S., Wen, T., Feng, W., Xu, X., Wang, Y., Zhang, Y., Zhu, Y., Wu, Y., Cai, Y., Liu, Z.: Qwen-image technical report (2025), <https://arxiv.org/abs/2508.02324>

60. Wu, J., Bian, J.W., Li, X., Wang, G., Reid, I., Torr, P., Prisacariu, V.A.: Gaussctrl: Multi-view consistent text-driven 3d gaussian splatting editing (2024), <https://arxiv.org/abs/2403.08733>
61. Xia, R., Tang, Y., Zhou, P.: Towards scalable and consistent 3d editing (2025), <https://arxiv.org/abs/2510.02994>
62. Xiang, J., Lv, Z., Xu, S., Deng, Y., Wang, R., Zhang, B., Chen, D., Tong, X., Yang, J.: Structured 3d latents for scalable and versatile 3d generation (2025), <https://arxiv.org/abs/2412.01506>
63. Xiao, H., Chen, Y., Huang, H., Xiong, H., Yang, J., Prasad, P., Zhao, Y.: Localized gaussian splatting editing with contextual awareness. In: 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). p. 5207–5217. IEEE (Feb 2025). <https://doi.org/10.1109/wacv61041.2025.00509>, <http://dx.doi.org/10.1109/WACV61041.2025.00509>
64. Yao, Y., Luo, Z., Li, S., Zhang, J., Ren, Y., Zhou, L., Fang, T., Quan, L.: Blendmvs: A large-scale dataset for generalized multi-view stereo networks (2020), <https://arxiv.org/abs/1911.10127>
65. Ye, J., Liu, F., Li, Q., Wang, Z., Wang, Y., Wang, X., Duan, Y., Zhu, J.: Dreamrewards: Text-to-3d generation with human preference (2024), <https://arxiv.org/abs/2403.14613>
66. Ye, J., Xie, S., Zhao, R., Wang, Z., Yan, H., Zu, W., Ma, L., Zhu, J.: Nano3d: A training-free approach for efficient 3d editing without masks (2025), <https://arxiv.org/abs/2510.15019>
67. Yin, F., Liu, S., Han, Y., Wang, Z., Xing, P., Wang, R., Cheng, W., Wang, Y., Li, A., Yin, Z., Chen, P., Zhang, X., Jiang, D., Zeng, X., Yu, G.: Reasonedit: Towards reasoning-enhanced image editing models (2025), <https://arxiv.org/abs/2511.22625>
68. Yu, L., Xiang, W., Han, K.: Edit-diffnerf: Editing 3d neural radiance fields using 2d diffusion model (2023), <https://arxiv.org/abs/2306.09551>
69. Yu, Q., Chow, W., Yue, Z., Pan, K., Wu, Y., Wan, X., Li, J., Tang, S., Zhang, H., Zhuang, Y.: Anyedit: Mastering unified high-quality image editing for any idea. arXiv preprint arXiv:2411.15738 (2024)
70. Yu, T., Li, X., Shen, Y., Liu, Y., Lourentzou, I.: Core3d: Collaborative reasoning as a foundation for 3d intelligence (2025), <https://arxiv.org/abs/2512.12768>
71. Zhang, Q., Xu, Y., Wang, C., Lee, H.Y., Wetzstein, G., Zhou, B., Yang, C.: 3ditscene: Editing any scene via language-guided disentangled gaussian splatting (2024), <https://arxiv.org/abs/2405.18424>
72. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric (2018), <https://arxiv.org/abs/1801.03924>
73. Zhao, C., Li, X., Feng, T., Zhao, Z., Chen, H., Shen, C.: Tinker: Diffusion’s gift to 3d—multi-view consistent editing from sparse inputs without per-scene optimization (2025), <https://arxiv.org/abs/2508.14811>
74. Zhao, Q., Li, Y., Sun, Q., Yan, Z.: Resilphase: Plug-and-play phase mapping and noise-resilient macro-trajectory extrapolation for diffusion acceleration (2026), <https://arxiv.org/abs/2606.26769>
75. Zhao, Y., Sun, S., Zhang, M., Shi, Y., Yang, X., Bian, J.: Scenerevis: A self-reflective vision-grounded framework for 3d indoor scene synthesis via multi-turn rl (2026), <https://arxiv.org/abs/2602.09432>
76. Zhuang, J., Wang, C., Liu, L., Lin, L., Li, G.: Dreameditor: Text-driven 3d scene editing with neural fields (2023), <https://arxiv.org/abs/2306.13455>