

JanusMesh: Fast and Zero-Shot 3D Visual Illusion Generation via Cross-Space Denoising

Siang-Ling Zhang, Huai-Hsun Cheng, Tsung-Ju Yang, and Yu-Lun Liu*

National Yang Ming Chiao Tung University, Taiwan




Fig. 1: Zero-shot 3D Visual Illusion Generation. Given two different text prompts, our method creates a single 3D mesh that embodies a dual-semantic visual illusion. The generated shape appears unrecognizable from arbitrary viewpoints, but it perfectly reveals the two target semantics (*e.g.*, a peacock and a pineapple) when observed from specific camera angles (View θ_1 and View θ_2). Our framework achieves this intricate 3D illusion efficiently without requiring per-shape optimization.

Abstract. Creating 3D visual illusions, a single 3D mesh that reveals entirely different semantics from various viewing angles, is a fascinating but tough challenge. Existing optimization-based methods are slow and can produce oversaturated colors. In contrast, naive stitching approaches fail to produce geometrically coherent objects. This results in visible unnatural seams and semantic leaks. In this paper, we present a fast and training-free framework for generating text-driven 3D visual illusions. Our approach decouples the generation into two stages. First, we propose a cross-space dual-branch denoising process. This process dynamically decodes 3D latents into voxel space for CLIP-guided orientation alignment and Signed Distance Field (SDF) blending, which ensures seamless geometric fusion. Second, we introduce a view-conditioned texture synthesis module that projects and aggregates view-specific 2D diffusion priors onto the fused geometry. Extensive experiments demonstrate that our method generates highly realistic, dual-semantic 3D illusions in just 3–5 minutes. It significantly outperforms existing methods in geometric integrity, semantic recognizability, and efficiency. Project page: <https://siang1105.github.io/JanusMesh.github.io/>

* Corresponding author.




Fig. 2: Comparison of 3D visual illusion generation methods. (a) **SDS-Based Methods** suffer from severe over-saturation and slow optimization. (b) **Direct Concatenation** exposes unnatural geometric seams and semantic leakage at target views (red arrows). (c) **Our method** creates a seamless, dual-semantic coherent 3D mesh. Unlike previous approaches, our method does not require training. It generates high-quality 3D visual illusions in just 3–5 minutes while completely preventing geometric interference between the two semantics.

Keywords: 3D Visual Illusion · Dual-Branch Denoising · View-Conditioned Texturing

1 Introduction

Visual illusions have long fascinated human perception by challenging our understanding of physical reality. While diffusion models have enabled computational optical illusions in 2D, bringing this concept into the 3D world—creating a single object that presents entirely different semantics depending on the viewing angle—remains a formidable challenge. As illustrated in Fig. 1, our goal is zero-shot *3D visual illusion generation*: synthesizing a single 3D mesh that presents, for instance, a peacock from θ_1 and a pineapple from θ_2 , while appearing as an abstract geometry from arbitrary angles.

Earlier 3D optical illusion works focused on 2D projections or surface patterns such as shadow art, wireframe silhouettes, or height fields. With text-to-image diffusion models, zero-shot 2D illusion methods (e.g., Visual Anagrams [20]) synchronize noise predictions across views to create multi-interpretation images. Extending this to true 3D meshes, existing approaches rely on SDS to optimize a 3D representation so its renders match different prompts at different viewpoints [37].

However, existing methods fall short in producing high-quality 3D illusions efficiently. As shown in Fig. 2, optimization-based approaches (e.g., Shape From Semantics [37]) require ~40 minutes per shape and suffer from severe color over-saturation. A naive Direct Concatenation of two separately generated objects produces unnatural geometric seams and visible backside leakage, breaking the perceptual illusion.

We present a zero-shot two-stage framework that generates coherent 3D visual illusions in 3–5 minutes. Stage 1 employs a dual-branch denoising process using TRELLIS [71], decoding latents into voxel space at each step, aligning objects via CLIP-guided orientation search, and merging them via SDF blending before re-encoding. Stage 2 performs view-conditioned texturing by projecting Stable Diffusion predictions onto the fused mesh. Our method demonstrates superior geometry coherence, texture realism, and semantic recognizability over existing baselines.

In summary, our main contributions are threefold:

- A zero-shot framework extending generative multi-view illusions from 2D to fully textured 3D meshes.
- A training-free two-stage architecture featuring dual-branch denoising with SDF blending and CLIP-guided alignment for geometric integrity, coupled with view-conditioned texturing for dual-semantic coherency.
- A rigorous evaluation protocol incorporating CLIP, GPT-4.1-mini, FID/KID, and a novel Object Detection metric, with experiments validating our method over baselines and demonstrating scalability to three-object illusions.

2 Related Work

Computational Optical Illusions. Prior work has explored appearance-varying 3D objects, including shadow art [54, 62, 67], wire art [24, 58], view-dependent heightfields [56], SDF-based anamorphic packing [11], and spatially ambiguous geometry [14]; in 2D, spatial-frequency decomposition [19, 55] shifts perceived content with viewing distance, while progressive vector sketching [9] transforms perceived semantics through sequential stroke addition. These approaches produce 2D projections, surface patterns, or—as with adversarially injected viewpoint-dependent content in 3D Gaussian Splatting [31]—view-specific artifacts, whereas our work generates a fully textured 3D mesh with distinct semantics from different viewpoints.

Illusion Generation with Diffusion Models. Text-to-image diffusion models [21, 61, 64] have enabled new illusion synthesis. SDS-based methods [2, 57] optimize for multiple prompts but converge slowly. Visual Anagrams [20] introduced zero-shot illusions via per-view noise averaging, extended to frequency decompositions [19], phase-transfer [18], and audio-visual spectrograms [8]; in 3D, Illusion3D [17] and LookingGlass [4] lift these priors into NeRF and anamorphic images, respectively. Our work extends the zero-shot spirit of [20] to native 3D latent space, producing a fully textured mesh.

Text-to-3D Generation. Optimization-based methods [6, 43, 57, 68] distill 2D diffusion priors via SDS, with improvements via interval score matching [42], rectified-flow distillation [73], and Gaussian acceleration [65, 75]. Feed-forward methods combine multi-view diffusion [46, 48, 50, 63] with reconstruction networks [23, 69, 72]. Application-driven variants extend generation to scene scale and alternative styles: training-free pipelines complete fully textured room-scale

meshes from sparse images [39], diffusion priors enable city-scale 3D scene creation with iterative geometry and texture refinement [35], and differentiable mesh optimization under 2D pixel-art supervision produces stylized 3D content [25]. Native 3D generative models [30, 34, 41, 70, 80] learn latent representations directly from 3D data; TRELLIS [71] encodes geometry and appearance in a sparse structured latent space via rectified flow [47], which we repurpose for dual-semantic mesh generation.

Synchronized Diffusion Denoising. Merging denoising trajectories enables compositional generation [45], seamless panoramas [1], and perceptual synchronization [36], with advanced samplers addressing compositional failures [15, 33]. Recent works average clean-image predictions rather than noise [32] and apply spatial guidance in 3D latents [16]. Closely related, multi-view diffusion outpainting coordinates denoising across views with geometry-aware strategies for sparse-view reconstruction [26], and 3D-aware 360° video diffusion decouples texture refinement from a 3D cache that enforces geometric consistency [7]. Unlike these consistency-enforcing frameworks, our dual-branch denoising uniquely enforces divergent semantics at target viewpoints via SDF fusion.

3D Texture Synthesis. Diffusion-based mesh texturing spans depth-conditioned inpainting [60], ControlNet-enhanced generation [5, 77], multi-view consistent UV-latents [3, 10, 49], and single-pass feed-forward models [76] with material [13] or appearance [74] extensions. Crucially, all these methods apply a single prompt uniformly. Our view-conditioned synthesis instead assigns distinct prompts to different angular sectors, back-projecting viewpoint-specific clean images via cosine-weighted blending.

CLIP-Guided 3D Understanding and Generation. CLIP’s [59] render-caption similarity enables zero-shot 3D generation [28], mesh stylization [52], NeRF manipulation [29, 66], and latent score distillation [51]. Leveraging this render-and-score paradigm, our CLIP-guided Orientation Search automatically selects the relative rotation that maximizes silhouette alignment between the two objects. This resolves geometric mismatches that would otherwise cause SDF fusion failures.

3 Method

3.1 Preliminaries

TRELLIS: Structured 3D Latent Representation. TRELLIS [71] is a two-stage Rectified Flow [44] 3D generator that first predicts a low-resolution sparse voxel structure, then refines it with high-dimensional appearance features. We perform geometry blending specifically during this first structural stage. Because direct spatial transformations distort latent distributions, we draw inspiration from LookingGlass [4]: at each denoising step, we decode the latent into voxel space, perform SDF blending, and re-encode the fused result back to the latent space. This cross-space denoising ensures the geometric validity of the blended mesh.




Fig. 3: Pipeline overview. (a) **Stage 1** employs dual-branch denoising. At each step, latents are decoded to voxel space, rotation-aligned, and fused via SDF blending (Fig. 4), then re-encoded to continue denoising, producing a single unified 3D mesh. (b) **Stage 2** applies view-conditioned texturing to the fused mesh. Estimated clean images $\hat{x}_{1|t}$ are predicted via Stable Diffusion, un-projected from viewpoints θ_1 and θ_2 , and iteratively aggregated via Mesh Texture Aggregation, producing a single object with distinct semantics at each target viewpoint.

Visual Anagrams. Visual Anagrams [20] generates 2D illusions by averaging noisy estimates across transformed views in a canonical space. Because this restricts transformations to be orthogonal, SyncTweedies [32] relaxes the constraint by averaging estimated clean latents instead: $\hat{z}_{1|t}^i = \pi_i(\frac{1}{N} \sum_j \pi_j^{-1}(z_{1|t}^j))$. We extend this synchronization principle from 2D pixels to 3D voxels, blending predicted clean geometries via SDF averaging to generate view-dependent dual-semantic meshes.

3.2 Overview

Given prompts y_1, y_2 and target viewpoints θ_1, θ_2 , our framework generates a single 3D mesh exhibiting y_1 at θ_1 and y_2 at θ_2 . A successful illusion requires: (1) **Semantic Recognizability** at target views; (2) **Geometric Integrity** from any viewpoint; and (3) **Illusion Effect** (concealing opposing semantics at non-target views), all evaluated in Sec. 4.1. As shown in Fig. 3, our pipeline achieves this via Stage 1 dual-branch geometric denoising (Sec. 3.3) and Stage 2 view-conditioned texturing (Sec. 3.4).

3.3 Dual-Branch Geometry Generation (Stage 1: Geometry)

Dual-Branch Denoising and Clean Latent Estimation. Stage 1 builds on TRELLIS’s Rectified Flow, starting from a shared initial noise z_t and performing two independent denoising branches conditioned on y_1 and y_2 , estimating the clean latent at each timestep t :

$$x_{1|t}^1 = z_t + u_\theta(z_t; t, y_1)(1 - t), \quad (1)$$




Fig. 4: SDF blending. (a) Given two rotation-aligned voxels, we compute their SDFs, take the element-wise average, and binarize with threshold τ to obtain the blended voxel. (b) 2D illustration: averaging the SDFs of a circle (A) and a triangle (B) produces a smooth intermediate contour lying geometrically between the two.

$$x_{1|t}^2 = z_t + u_\theta(z_t; t, y_2)(1 - t), \quad (2)$$

where u_θ is the Rectified Flow Network. We apply Classifier-Free Guidance (CFG) [22] with Interval CFG to improve generation quality while avoiding over-saturation at extreme noise levels.

Geometry Blending in Voxel Space. The clean latent estimates $x_{1|t}^1$ and $x_{1|t}^2$ are decoded via the Sparse Structure Decoder into voxel representations v_1 and v_2 . We rotate v_2 by θ_2 to align both voxels in a common reference frame. Since directly averaging binary occupancy grids lacks geometric continuity, we convert both into Signed Distance Fields (SDFs), average them element-wise, and binarize with threshold τ to obtain the blended geometry $\hat{x}_{1|t}$:

$$\text{SDF}_{\text{blend}} = \frac{\text{SDF}(v_1) + \text{SDF}(v_2)}{2}, \quad (3)$$

$$\hat{x}_{1|t} = [\text{SDF}_{\text{blend}} < \tau]. \quad (4)$$

The blending process is illustrated in Fig. 4. After blending, $\hat{x}_{1|t}$ is rotated by $-\theta_2$ to restore v_2 ’s original coordinate frame, then re-encoded into $\hat{z}_{1|t}^1$ and $\hat{z}_{1|t}^2$ via the Sparse Structure Encoder for the next denoising step.

3.4 View-Conditioned Texture Synthesis (Stage 2: Texture)

View-Aware Texture Prediction. Because the fused Stage 1 mesh contains unnatural geometry, direct TRELIS texturing fails. Therefore, we treat texturing as an independent stage using a depth-conditioned ControlNet [78] (Stable Diffusion [61]) to predict clean images $\hat{x}_{1|t}$. At each denoising step, we render the mesh from θ_1 and θ_2 , predict textures, and un-project them onto the 3D surface. Finally, Mesh Texture Aggregation merges these multi-view contributions using cosine-weighted blending (based on surface normals), ensuring high-quality, seamless textures.

View-Dependent Texture Selection. We determine which branch’s texture to use based on the current camera angle. Taking $\theta_1 = 0^\circ$ as an example, view-points within $270^\circ - 90^\circ$ adopt the texture estimate from y_1 , while the remaining




Fig. 5: Noise Guidance. Given two prompts y_1 and y_2 , two single-semantic voxels $\hat{v}_1$ and $\hat{v}_2$ are independently generated and concatenated at θ_1 and θ_2 to form v_{guide} , encoded into z_{guide} via the Sparse Structure Encoder. **Noise Blending Guidance** mixes z_{guide} with pure noise via $\alpha \cdot z_{\text{guide}} + (1 - \alpha) \cdot z_{\text{noise}}$, injecting a mild spatial prior. **Space Control Guidance** interpolates between z_{guide} and noise at timestep t_0 , providing stronger structural constraints for geometrically challenging pairs.

angles adopt that from y_2 . Although switched via a hard cutoff, no visible seam appears in practice, as cosine-weighted blending naturally smooths the transition at the boundary.

3.5 Noise Guidance

The starting point of Rectified Flow denoising—pure random noise—lacks spatial structural constraint for dual-semantic fusion, making it prone to geometric interference and difficult convergence. We propose two noise guidance strategies that inject a spatial prior before denoising begins, improving geometric fusion quality, as illustrated in Fig. 5.

Noise Blending Guidance. We pre-generate two single-semantic voxels $\hat{v}_1$ and $\hat{v}_2$, halve and concatenate them at target angles θ_1 and θ_2 to form a guidance voxel v_{guide} , which is encoded into a guidance latent via the Sparse Structure Encoder. The guidance latent is then combined with pure Gaussian noise via a weighted sum as the initial denoising latent:

$$z_{\text{init}} = \alpha \cdot \text{Encoder}(v_{\text{guide}}) + (1 - \alpha) \cdot z_{\text{noise}}, \quad (5)$$

where α controls the guidance strength, balancing structural prior and generation diversity.

Space Control Guidance. Inspired by SpaceControl [16], we adapt the principle of manipulating the starting denoising timestep to our dual-semantic fusion setting. Similarly to Noise Blending Guidance, we first generate v_{guide} and encode it into a guidance latent z_{guide} , and interpolate between the guidance latent and pure noise at timestep t_0 :

$$z\{t_0\} = t_0 \cdot \text{Encoder}(v_{\text{guide}}) + (1 - t_0) \cdot z_{\text{noise}}. \quad (6)$$

A larger t_0 imposes stronger structural constraints, while a smaller t_0 preserves more generation diversity. In our 25-step setting, $t_0 = 10$ means the first 10 steps are guided before normal denoising resumes, striking a balance between structural constraint and generation freedom.




Fig. 6: CLIP-guided Orientation Search. **Anchor View Selection** identifies the best representative view I_1 for $\hat{v}_1$ from 4 orthogonal renders via CLIP text-image similarity (e.g., “pineapple” at 90°). **Cross-object Matching** then evaluates 28 3D rotations of $\hat{v}_2$, selecting the angle θ_2^* that maximizes CLIP image-image similarity to I_1 . This optimally aligns their silhouettes prior to SDF blending.

3.6 CLIP-guided Orientation Search

Existing multi-view illusion methods typically assume fixed viewpoints (e.g., 0° and 180°); however, different objects may exhibit significant orientation differences in their canonical poses. Directly fusing two geometrically misaligned objects at default angles often results in a blended mesh that fails to present clear semantic silhouettes. To address this, we propose a CLIP-based adaptive orientation search that automatically selects the optimal fusion orientation before denoising begins, as illustrated in Fig. 6.

Anchor View Selection for Object 1. We independently generate single-semantic voxels $\hat{v}_1$ and $\hat{v}_2$ from prompts y_1 and y_2 respectively. For object 1, we render 4 candidate views of $\hat{v}_1$ at 90° intervals around the Z-axis and select the view with the highest CLIP text-image similarity to y_1 as the representative view I_1 .

Cross-Object Orientation Matching for Object 2. For object 2, we sample discrete rotation combinations along the (X, Y, Z) axes, generating 28 candidate renders, and select the rotation with the highest CLIP image-image similarity to I_1 as the fusion orientation θ_2^* . Candidate angles are sampled at 90° intervals, covering $\{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$ per axis, balancing search efficiency and angular coverage. Denser sampling (e.g., 45°) could further improve alignment precision, which we leave as future work.

3.7 Extension to Three-Object 3D Illusions

Our framework naturally scales to three-object illusions by adding a third Stage 1 denoising branch. Given three prompts and a shared noise z_t , we fix target angles at 0° , 120° , and 240° . At each step, the three predicted voxels are SDF-fused, inverse-rotated, and re-encoded. Due to complex geometric conflicts, Noise Guidance is mandatory here, utilizing a merged prior of three pre-generated single-semantic objects. Specifically, since averaging three geometries imposes stronger geometric conflicts than the two-object case, we adopt Space Control




Fig. 7: Extension to three-object 3D illusion. Our framework scales to three semantics (e.g., “Grapes”, “A Pineapple”, “Bamboo”) by adding a third denoising branch, with rotation angles fixed at 0° , 120° , 240° to uniformly cover 360° . Noise Guidance steers all three branches toward their respective targets, ensuring each semantic is clearly presented at its target viewpoint.

Guidance with an increased $t_0 = 20$ out of 25 steps, providing stronger structural constraints to ensure each semantic is preserved at its target viewpoint. Stage 2 remains identical to the two-object case. See Fig. 7 for the pipeline and Sec. 4.4 for results.

4 Experiments

4.1 Experimental Setup

Baselines. We compare against **five** baselines: **Shape from Semantics** [38], which addresses the same task via Score Distillation Sampling (SDS) but requires ~ 40 minutes per object; **Shape from Semantics (VSD)** [38], a stronger variant that replaces the SDS objective with Variational Score Distillation (VSD) [68], incurring ~ 4 hr/sample runtime; and **Direct Concatenation**, which independently generates two objects via TRELLIS [71], halves each along the midplane, and stitches them into a single mesh. We also evaluate **TRELLIS** [71] and **DreamBeast** [40], both prompted with “a single 3D object, front side is $\{y_1\}$, back side is $\{y_2\}$ ” to elicit view-dependent semantics from a single generation. TRELLIS is a state-of-the-art feed-forward 3D generation model, while DreamBeast is an SDS-based method designed for generating fantastical creatures with part-level semantic control. These two baselines assess whether existing generation methods can produce coherent visual illusions without task-specific design.

Data. We collect object prompts from Shape from Semantics [38] and supplement with additional common objects, yielding 60 distinct objects across five categories: 16 birds, 19 mammals, 5 reptiles and aquatic animals, 9 plants, and 11 man-made artifacts. Two objects are randomly sampled per experiment to form a prompt pair.

Implementation Details. All experiments are conducted on a single NVIDIA RTX 4090 GPU. For SDF blending, we use Truncated SDFs with truncation

distance $\text{clip_s} = 12$ and binarization threshold $\tau = 0.8$. Stage 1 runs for 25 denoising steps with one geometry blending operation per step, with CFG applied at guidance scale $\omega = 7.5$ within $t \in [0.5, 0.95]$ (Interval CFG); Stage 2 runs for 30 denoising steps. For CLIP-guided Orientation Search, we use OpenCLIP [27] ViT-B/32 pretrained on LAION-2B (1aion2b_s34b_b79k). For Noise Blending Guidance, we set $\alpha = 0.3$; for Space Control Guidance, we set $t_0 = 10$ out of 25 steps.

We organize experiments into three rotational configurations, with object *A* fixed at its canonical orientation: **Case 1**, object *B* is unrotated, revealing the front of object *A* at 0° and the back of object *B* at 180° ; **Case 2**, object *B* is rotated by 180° , revealing the front views of both objects at the two target viewpoints; and **Case 3**, object *B* is automatically rotated via CLIP-guided Orientation Search, typically placing the two objects approximately 180° apart.

Noise Blending and Space Control Guidance can be freely combined with Cases 1 and 2, whereas Case 3 requires no additional guidance since the fusion angle is already adapted via CLIP. Cases 1 and 2 require approximately 3 minutes in total (~ 1 min for Stage 1 and ~ 2 min for Stage 2), while Case 3 requires approximately 5 minutes due to the additional CLIP-guided Orientation Search (~ 3 min for Stage 1 and ~ 2 min for Stage 2). All configurations complete within 3–5 minutes, offering a significant efficiency advantage over the SDS-based baseline (~ 40 minutes).

Metrics. We design six quantitative metrics and a user study to evaluate generation quality and illusion effect:

(1) *CLIP Similarity* computes CLIP [27, 59] text-image similarity across 1,000 renders per viewpoint ($\pm 20^\circ$ jitter). Note that Direct Concatenation artificially inflates this score by trivially preserving single-view appearances, making CLIP alone insufficient.

(2) *GPT Accuracy (%)* prompts GPT-4.1-mini to identify target semantics from side-by-side renders of both viewpoints, reporting the proportion of responses matching the ground-truth prompt assignment as a 2-way accuracy.

(3) *FID & KID* measure visual realism between 1,000 renders and 1,000 reference images (20 views of 50 objects) from Objaverse 1.0 [12]. (4) *Object Detection Score* renders the midpoint angle and counts objects via OWLv2 [53] (Fig. 8), reporting average object count (ideally 1) and multi-object rate (proportion with > 1 detection). (5) *View-Conditional CLIP Contrast* computes CLIP similarity between each viewpoint render and the *opposite* prompt, penalizing cross-view semantic leakage




Fig. 8: Object detection at the junction viewpoint for “Bamboo” & “Dwarf Cottage”. (a) Direct Concatenation is detected as two separate objects (red and green boxes); (b) Ours is detected as a single unified object (red box only).




Fig. 9: Qualitative comparison with baselines. Left-to-right: View 1, blended mesh, View 2. (a) Shape From Semantics [38] suffers from over-saturation and geometry leakage. (b) Direct Concatenation exposes unnatural junction seams and opposing geometry at target views. (c) Ours produces a coherent mesh with clear, view-dependent semantics and no leakage. (d) Shape From Semantics (VSD) [38] still exhibits over-saturation and leakage despite its stronger objective. (e) TRELLIS [71], prompted with “a single 3D object, front side is A, back side is B”, fails to produce view-dependent semantics and collapses to a generic shape that reflects neither target interpretation. (f) DreamBeast [40] blends both semantics uniformly across all viewpoints, failing to achieve view-dependent separation.




Fig. 10: Qualitative comparison (CLIP-guided orientation). Left-to-right: View 1, blended mesh, View 2. (c) **Ours** uses CLIP-guided search to adaptively align compatible silhouettes, revealing clearer semantics at both viewpoints. (a)(b) Shape From Semantics [38] and Direct Concatenation suffer from silhouette misalignment and degraded fusion without adaptive rotation. (d)(e)(f) See Fig. 9 caption for descriptions of remaining baselines.

that naive stitching permits; a lower score reflects cleaner perceptual separation. (6) *Boundary Seam Score* counts triangles connected to fewer than three neighbors ($\text{deg}=1$, lower is better), averaged over 50 pairs; Direct Concatenation yields 1503.42 vs. **798.31** for ours (Fig. 11).

User Study. Beyond quantitative metrics, we conduct a user study to collect perceptual evaluations from human observers. A total of 50 participants took part in the study. Each participant was presented with rendered results from all three methods and asked to evaluate them across the following questions, each with four prompt-pair sub-questions:

- **Q1:** How recognizable are the intended semantics at each target viewpoint? (1: unrecognizable, 2: partially recognizable, 3: clearly recognizable)
- **Q2:** Which result better aligns with the intended semantics?
- **Q3:** Comparing CLIP-adaptive orientation (Case 3) versus fixed $0^\circ/180^\circ$ angles, which produces a more natural illusion effect?

4.2 Results and Analysis

Quantitative Comparison. Table 1 reports results over 50 randomly sampled prompt pairs. Although Direct Concatenation achieves the highest CLIP score

Table 1: Comparison with the baselines.

Method	CLIP $\uparrow$	CLIP (opp.) $\downarrow$	GPT Acc. (%) $\uparrow$	FID $\downarrow$	KID $\downarrow$	Object Detection			Boundary Seam (deg=1) $\downarrow$	Runtime
						Avg. Obj. Count	Multi-Obj. Rate (%)			
Shape From Semantic [38]	27.460	19.72	70	194.136	0.051	0.64	2		—	~40 min
Shape From Semantic (VSD) [38]	27.523	20.69	65	195.160	0.061	0.68	4		—	~3-4 hr
Direct Concat	29.030	20.38	76	187.886	0.067	2.1	56		1503.42	~3-5 min
Trellis	22.180	22.68	60	174.130	0.044	0.58	8		—	~2-3 min
DreamBeast	22.993	22.89	65	184.956	0.0935	0.76	18		—	~3-4 hr
Ours	28.170	19.26	84	185.555	0.051	0.86	18		798.31	~3-5 min




Fig. 12: User study. (Left) Semantic Recognizability: 78.5% of participants rated our results as clearly recognizable (score 3). **(Middle) Method Preference:** 71% of participants preferred our method over Shape From Semantics [38] (16%) and Direct Concatenation (13%). **(Right) Orientation Strategy:** 91% of participants found CLIP-guided orientation to produce a more natural illusion than fixed 0°/180° angles.

(29.030 vs. ours 28.170), this is a systematic artifact of naive per-viewpoint appearance preservation rather than genuine illusion quality (Sec. 4.1). Our method achieves the highest GPT Accuracy (84%), lowest FID (185.555), and best Object Detection scores (avg. count 0.86, multi-object rate 18% vs. 56% for Direct Concatenation), demonstrating superior semantic recognizability, visual realism, and geometric coherence. Table 1 further reports a lower Boundary Seam Score for ours (**798.31** vs. 1503.42 for Direct Concatenation), while Fig. 11 visually confirms smoother normals at the fusion boundary. All configurations complete within 3–5 minutes, matching Direct Concatenation and offering a significant efficiency advantage over Shape from Semantics [38] (~40 minutes).

Qualitative Comparison. Fig. 9 and Fig. 10 compare results under fixed-angle and CLIP-guided configurations respectively. Our method produces geometrically coherent results with clear per-viewpoint semantics, while Shape From Semantics [38] suffers from over-saturation and Direct Concatenation exposes visible junction seams.

User Studies. Our user studies strongly reinforce the quantitative findings. Participants overwhelmingly preferred our method over both baselines and rated our results as semantically recognizable, confirming the perceptual quality of our generated illusions. Additionally, the large majority found CLIP-guided orientation to produce more natural illusions than fixed angles, validating the effectiveness of our adaptive orientation search. Results are shown in Fig. 12.




Fig. 11: Normal map comparison at fusion boundary. Direct Concatenation exhibits sharp discontinuities, whereas ours produces smoother normals across the fusion region.




Fig. 13: Ablation on geometry blending (“carriage”/“sundae”). Left-to-right: View 1, blended mesh, View 2. (a) **Union**: yields conflicting junctions. (b) **Blur Avg**: loses fine details. (c) **Minkowski**: over-expands geometry. (d) **Polar Coord**: distorts asymmetric objects. (e) **SDF Avg (Ours)**: optimally balances geometric integrity and semantic preservation.

Input	(a) Latent Space Blend	(b) Voxel Space Blend (Ours)	Method	Latent	Ours
“Pumpkin Carriage” & “A Lion”			CLIP ↑	22.47	28.17
			GPT ↑	58	84
			FID ↓	195.6	185.6
			KID ↓	0.081	0.051
“A Hippopotamus” & “Blue Rose”			Obj. Cnt ↓	1.30	0.86
			Multi-Obj ↓	22	18

Fig. 14: Ablation on geometry blending in latent vs. voxel space. (a) **Latent Space Blend**: averaging $z_{1|t}^1$ and $z_{1|t}^2$ directly in latent space without the decode-SDF blend-re-encode procedure falls outside TRELLIS’s latent manifold, yielding incoherent and fragmented geometry. (b) **Voxel Space Blend (Ours)**: SDF blending in voxel space produces geometrically coherent meshes with clear view-dependent semantics. The quantitative comparison (right) confirms the necessity of our cross-space design.

4.3 Ablation Studies

Geometry Blending Strategy. We compare five voxel fusion strategies. **Union** takes the logical OR of two voxel grids, directly superimposing both structures without any shape coordination. **Blur Average** applies 3D Gaussian smoothing before element-wise averaging, introducing soft boundaries but losing fine geometric details. **Minkowski Blend** dilates each voxel with a spherical structuring element before merging, which tends to inflate the overall mesh volume. **Polar Coord Blend** fuses voxels slice-by-slice via 2D polar-coordinate boundary averaging, implicitly assuming star-shaped slices and thus failing for concave or non-symmetric objects. **SDF Average (Ours)** averages the Truncated SDFs of both occupancy grids and binarizes the result; the zero iso-surface naturally corresponds to the intermediate shape between the two objects, producing a stable and coherent blended surface. See Fig. 13 for visual comparisons.

Latent Space vs. Voxel Space Blending. While the above strategies all operate in voxel space, one may ask whether blending can be performed directly in latent space to avoid the decode-re-encode overhead. We compare against **Direct Latent Averaging**, which averages $z_{1|t}^1$ and $z_{1|t}^2$ directly without the decode-SDF blend-re-encode procedure. As shown in Fig. 14, direct latent blending falls outside TRELLIS’s [71] latent manifold, producing fragmented geometry and significantly worse quantitative performance than our voxel-space blending, confirming the necessity of the proposed cross-space design.

Noise Guidance. We compare three configurations: no guidance, Noise Blending Guidance, and Space Control Guidance. The optimal strategy depends on




Fig. 15: Ablation on noise guidance across varying geometric compatibilities: (a) no guidance, (b) Noise Blending, and (c) Space Control. (c) handles large geometric discrepancies (“Bamboo”/“Grapes”) best via strong spatial constraints. (b) is optimal for similar silhouettes but distinct semantics (“Succulent”/“Rooster”), avoiding semantic loss in (a) and residual artifacts in (c). For compatible pairs (“Frog”/“Turtle”), all settings perform comparably. Thus, optimal intervention strength depends on geometric alignment.




Fig. 16: Ablation on texture synthesis and orientation search. Each pair shows View 1, blended mesh, and View 2. (a) **w/o View-Conditioned Texture**: Standard texturing fails on fused geometries, blending semantics into indistinguishable appearances (e.g., uniform foliage or solid purple). (b) **w/ View-Conditioned Texture (Ours)**: Accurately assigns distinct, prompt-driven textures to each viewpoint. (c) **w/o CLIP-guided**: Fixed $0^\circ/180^\circ$ angles cause silhouette misalignments, failing to reveal clear semantics. (d) **w/ CLIP-guided (Ours)**: Adaptive rotation optimally aligns silhouettes, recovering distinct semantics at both views.

the geometric characteristics of the object pair: Space Control imposes stronger constraints via the guided latent state for the first t_0 steps, making it better suited for pairs with large silhouette discrepancies; Noise Blending provides a milder prior that benefits pairs with similar silhouettes but distinct semantics; for geometrically compatible pairs, no guidance suffices. We therefore treat noise guidance as an optional, pair-dependent auxiliary rather than a mandatory component. Qualitative comparisons are shown in Fig. 15.

View-Conditioned Texture Synthesis. We ablate the necessity of Stage 2 by comparing results with and without view-conditioned texturing (Fig. 16(a)(b)). Without Stage 2, TRELLIS fails to interpret the unnatural fused geometry, producing semantically incoherent textures at both viewpoints. Our view-conditioned synthesis assigns each viewpoint its own texture prediction, ensuring semantic consistency with the target prompt.




Fig. 17: Qualitative results of three-object 3D illusion generation. Each row shows the fused mesh and three target-viewpoint renders for a prompt triplet. Our method successfully generates a single coherent mesh that presents three distinct semantics at 0° , 120° , and 240° , respectively.

CLIP-Guided Orientation Search. We compare adaptive CLIP-guided rotation against fixed $0^\circ/180^\circ$ angles (Fig. 16(c),(d)). For canonically misaligned pairs (e.g., horizontal *“Rhinceros”*/upright *“Pineapple”*), fixed angles cause silhouette misalignment; CLIP rotation optimally aligns them to recover distinct semantics. It similarly prevents semantic collapse for *“Succulent”*/*“Cactus”*. However, for already compatible pairs, CLIP misclassifications can yield suboptimal angles. Thus, we recommend adaptive rotation specifically for geometrically challenging pairs, defaulting to fixed angles for compatible ones.

4.4 Applications

Our framework demonstrates strong performance across diverse semantic combinations, from structurally similar to geometrically distinct object pairs, as shown in Fig. 9 and Fig. 10. Beyond the two-object case, our method scales naturally to three-object illusion generation without architectural modifications (Sec. 3.7), as demonstrated in Fig. 17.

5 Conclusion

We present a zero-shot framework that extends visual illusions to true 3D geometry. Given text prompts, our method generates a single coherent mesh revealing distinct semantics from different viewpoints in under 5 minutes, without per-shape optimization. Furthermore, we introduce CLIP-guided Orientation Search for silhouette alignment and Noise Guidance to resolve geometric conflicts, demonstrating seamless scalability to three-object illusions without modifying the core fusion procedure.

Limitations. Our method has three limitations. First, it inherits failure cases from TRELLIS [71]; for example, TRELLIS may misinterpret *“bat”* as Batman, yielding incompatible geometries. Second, our framework requires a voxel-based latent representation and is therefore not directly applicable to token-based 3D generators such as Hunyuan3D [79]. Third, CLIP-guided Orientation Search scales poorly to three-object illusions due to ambiguous silhouette alignment, so we use fixed target angles (0° , 120° , 240°).

References

1. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation (2023)
2. Burgert, R., Li, X., Leite, A., Ranasinghe, K., Ryoo, M.: Diffusion illusions: Hiding images in plain sight. In: ACM SIGGRAPH 2024 Conference Papers. pp. 1–11 (2024)
3. Cao, T., Kreis, K., Fidler, S., Sharp, N., Yin, K.: Textfusion: Synthesizing 3d textures with text-guided image diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4169–4181 (2023)
4. Chang, P., Sancho, S., Tang, J., Gross, M., Azevedo, V.: Lookingglass: Generative anamorphoses via laplacian pyramid warping. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 24–33 (2025)
5. Chen, D.Z., Siddiqui, Y., Lee, H.Y., Tulyakov, S., Nießner, M.: Text2tex: Text-driven texture synthesis via diffusion models. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 18558–18568 (2023)
6. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 22246–22256 (2023)
7. Chen, T.H., Chen, Y.H., Tu, T., Lee, J.Y., Wu, C.Y., Lin, F., Zhang, H., Paz, D., Huang, X., Guo, Y., et al.: Pantheon360: Taming digital twin generation via 3d-aware 360deg video diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11138–11149 (2026)
8. Chen, Z., Geng, D., Owens, A.: Images that sound: Composing images and sounds on a single canvas. *Advances in Neural Information Processing Systems* **37**, 85045–85073 (2024)
9. Cheng, H.H., Zhang, S.L., Liu, Y.L.: Stroke of surprise: Progressive semantic illusions in vector sketching. arXiv preprint arXiv:2602.12280 (2026)
10. Cheng, W., Mu, J., Zeng, X., Chen, X., Pang, A., Zhang, C., Wang, Z., Fu, B., Yu, G., Liu, Z., et al.: Mvpaint: Synchronized multi-view diffusion for painting anything 3d. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 585–594 (2025)
11. Debnath, S., Tiwari, A., Sadekar, K., Raman, S.: Rasp: revisiting 3d anamorphic art for shadow-guided packing of irregular objects. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5849–5858 (2025)
12. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13142–13153 (2023)
13. Deng, K., Omernick, T., Weiss, A., Ramanan, D., Zhu, J.Y., Zhou, T., Agrawala, M.: Flashtex: Fast relightable mesh texturing with lightcontrolnet. In: European conference on computer vision. pp. 90–107. Springer (2024)
14. Dodik, A., Yu, I., Chandra, K., Ragan-Kelley, J., Tenenbaum, J., Sitzmann, V., Solomon, J.: Meschers: Geometry processing of impossible objects. *ACM Transactions on Graphics (TOG)* **44**(4), 1–10 (2025)
15. Du, Y., Durkan, C., Strudel, R., Tenenbaum, J.B., Dieleman, S., Fergus, R., Sohl-Dickstein, J., Doucet, A., Grathwohl, W.S.: Reduce, reuse, recycle: Compositional generation with energy-based diffusion models and mcmc. In: International conference on machine learning. pp. 8489–8510. PMLR (2023)

16. Fedele, E., Engelmann, F., Huang, I., Litany, O., Pollefeys, M., Guibas, L.: Space-control: Introducing test-time spatial control to 3d generative modeling. arXiv preprint arXiv:2512.05343 (2025)
17. Feng, Y., Sanjay, V., Lutz, S., AlBahar, B., Ge, S., Huang, J.B.: Illusion3d: 3d multiview illusion with 2d diffusion priors. arXiv preprint arXiv:2412.09625 (2024)
18. Gao, X., Yang, S., Liu, J.: Ptdiffusion: Free lunch for generating optical illusion hidden pictures with phase-transferred diffusion model. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 18240–18249 (2025)
19. Geng, D., Park, I., Owens, A.: Factorized diffusion: Perceptual illusions by noise decomposition. In: European Conference on Computer Vision. pp. 366–384. Springer (2024)
20. Geng, D., Park, I., Owens, A.: Visual anagrams: Generating multi-view optical illusions with diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24154–24163 (2024)
21. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
22. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
23. Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: Lrm: Large reconstruction model for single image to 3d. arXiv preprint arXiv:2311.04400 (2023)
24. Hsiao, K.W., Huang, J.B., Chu, H.K.: Multi-view wire art. *ACM Trans. Graph.* **37**(6), 242 (2018)
25. Huang, Y.C., Chan, J., Chien, H.J., Liu, Y.L.: Voxify3d: Pixel art meets volumetric rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15398–15410 (2026)
26. Huang, Y.C., Chien, H.J., Lin, C.Y., Chen, Y.H., Liu, Y.L.: Gamo: Geometry-aware multi-view diffusion outpainting for sparse-view 3d reconstruction. arXiv preprint arXiv:2512.25073 (2025)
27. Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., Hajishirzi, H., Farhadi, A., Schmidt, L.: Openclip (Jul 2021). <https://doi.org/10.5281/zenodo.5143773>, <https://doi.org/10.5281/zenodo.5143773>, if you use this software, please cite it as below.
28. Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 867–876 (2022)
29. Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5885–5894 (2021)
30. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463 (2023)
31. Ke, B.H., Xie, Y.Z., Liu, Y.L., Chiu, W.C.: Stealthattack: Robust 3d gaussian splatting poisoning via density-guided illusions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 27400–27411 (2025)
32. Kim, J., Koo, J., Yeo, K., Sung, M.: Synctweedies: A general generative framework based on synchronized diffusions. arXiv preprint arXiv:2403.14370 (2024)
33. Kim, S., Lee, K., Choi, J.S., Jeong, J., Sohn, K., Shin, J.: Collaborative score distillation for consistent visual synthesis. arXiv preprint arXiv:2307.04787 (2023)
34. Lan, Y., Zhou, S., Lyu, Z., Hong, F., Yang, S., Dai, B., Pan, X., Loy, C.C.: Gaussiananything: Interactive point cloud latent diffusion for 3d generation (2025)

35. Lee, J.Y., Liu, Y.R., Tsai, S.R., Chang, W.C., Wu, C.H., Chan, J., Zhao, Z., Lin, C.H., Liu, Y.L.: Skyfall-gs: Synthesizing immersive 3d urban scenes from satellite imagery. arXiv preprint arXiv:2510.15869 (2025)
36. Lee, Y., Kim, K., Kim, H., Sung, M.: Syncdiffusion: Coherent montage via synchronized joint diffusions. *Advances in Neural Information Processing Systems* **36**, 50648–50660 (2023)
37. Li, L., Wang, C., Zhou, Y., Deng, B., Zhang, J.: Shape from semantics: 3d shape generation from multi-view semantics. arXiv preprint arXiv:2502.00360 (2025)
38. Li, L., Wang, C., Zhou, Y., Deng, B., Zhang, J.: Shape from semantics: 3d shape generation from multi-view semantics (2025), <https://arxiv.org/abs/2502.00360>
39. Li, M.F., Ku, Y.F., Yen, H.X., Liu, C., Liu, Y.L., Chen, A.Y., Kuo, C.H., Sun, M.: Genrc: Generative 3d room completion from sparse image collections. In: *European Conference on Computer Vision*. pp. 146–163. Springer (2024)
40. Li, R., Han, J., Melas-Kyriazi, L., Sun, C., An, Z., Gui, Z., Sun, S., Torr, P., Jakab, T.: Dreambeast: Distilling 3d fantastical animals with part-aware knowledge transfer. In: *2025 International Conference on 3D Vision (3DV)*. pp. 1243–1252. IEEE (2025)
41. Li, Y., Zou, Z.X., Liu, Z., Wang, D., Liang, Y., Yu, Z., Liu, X., Guo, Y.C., Liang, D., Ouyang, W., et al.: Triposg: High-fidelity 3d shape synthesis using large-scale rectified flow models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025)
42. Liang, Y., Yang, X., Lin, J., Li, H., Xu, X., Chen, Y.: Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6517–6526 (2024)
43. Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 300–309 (2023)
44. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
45. Liu, N., Li, S., Du, Y., Torralba, A., Tenenbaum, J.B.: Compositional visual generation with composable diffusion models. In: *European conference on computer vision*. pp. 423–439. Springer (2022)
46. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 9298–9309 (2023)
47. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022)
48. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023)
49. Liu, Y., Xie, M., Liu, H., Wong, T.T.: Text-guided texturing by synchronized multi-view diffusion. In: *SIGGRAPH Asia 2024 Conference Papers*. pp. 1–11 (2024)
50. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9970–9980 (2024)

51. Metzger, G., Richardson, E., Patashnik, O., Giryes, R., Cohen-Or, D.: Latent-nerf for shape-guided generation of 3d shapes and textures. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12663–12673 (2023)
52. Michel, O., Bar-On, R., Liu, R., Benaim, S., Hanocka, R.: Text2mesh: Text-driven neural stylization for meshes. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13492–13502 (2022)
53. Minderer, M., Gritsenko, A., Hounsby, N.: Scaling open-vocabulary object detection. *Advances in Neural Information Processing Systems* **36**, 72983–73007 (2023)
54. Mitra, N.J., Pauly, M.: Shadow art. *ACM Transactions on Graphics* **28**(5), 156–1 (2009)
55. Oliva, A., Torralba, A., Schyns, P.G.: Hybrid images. *ACM Transactions on Graphics (TOG)* **25**(3), 527–532 (2006)
56. Perroni-Scharf, M., Rusinkiewicz, S.: Constructing printable surfaces with view-dependent appearance. In: ACM SIGGRAPH 2023 conference proceedings. pp. 1–10 (2023)
57. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022)
58. Qu, Z., Yang, L., Zhang, H., Xiang, T., Pang, K., Song, Y.Z.: Wired perspectives: Multi-view wire art embraces generative ai. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6149–6158 (2024)
59. 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)
60. Richardson, E., Metzger, G., Alaluf, Y., Giryes, R., Cohen-Or, D.: Texture: Text-guided texturing of 3d shapes. In: ACM SIGGRAPH 2023 conference proceedings. pp. 1–11 (2023)
61. 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)
62. Sadekar, K., Tiwari, A., Raman, S.: Shadow art revisited: a differentiable rendering based approach. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 29–37 (2022)
63. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512* (2023)
64. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
65. 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)
66. Wang, C., Chai, M., He, M., Chen, D., Liao, J.: Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3835–3844 (2022)
67. Wang, C., Deng, B., Zhang, J.: Neural shadow art. *arXiv preprint arXiv:2411.19161* (2024)
68. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Advances in neural information processing systems* **36**, 8406–8441 (2023)
69. Wang, Z., Wang, Y., Chen, Y., Xiang, C., Chen, S., Yu, D., Li, C., Su, H., Zhu, J.: Crm: Single image to 3d textured mesh with convolutional reconstruction model. In: European conference on computer vision. pp. 57–74. Springer (2024)

70. Wu, S., Lin, Y., Zhang, F., Zeng, Y., Xu, J., Torr, P., Cao, X., Yao, Y.: Direct3d: Scalable image-to-3d generation via 3d latent diffusion transformer. *Advances in Neural Information Processing Systems* **37**, 121859–121881 (2024)
71. 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. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 21469–21480 (2025)
72. Xu, J., Cheng, W., Gao, Y., Wang, X., Gao, S., Shan, Y.: Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191* (2024)
73. Yang, X., Chen, C., Yang, X., Liu, F., Lin, G.: Text-to-image rectified flow as plug-and-play priors. *arXiv preprint arXiv:2406.03293* (2024)
74. Yeh, Y.Y., Huang, J.B., Kim, C., Xiao, L., Nguyen-Phuoc, T., Khan, N., Zhang, C., Chandraker, M., Marshall, C.S., Dong, Z., et al.: Texturedreamer: Image-guided texture synthesis through geometry-aware diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4304–4314 (2024)
75. Yi, T., Fang, J., Wang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: Gaussiandreamer: Fast generation from text to 3d gaussians by bridging 2d and 3d diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 6796–6807 (2024)
76. Yu, X., Yuan, Z., Guo, Y.C., Liu, Y.T., Liu, J., Li, Y., Cao, Y.P., Liang, D., Qi, X.: Texgen: a generative diffusion model for mesh textures. *ACM Transactions on Graphics (TOG)* **43**(6), 1–14 (2024)
77. Zeng, X., Chen, X., Qi, Z., Liu, W., Zhao, Z., Wang, Z., Fu, B., Liu, Y., Yu, G.: Paint3d: Paint anything 3d with lighting-less texture diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4252–4262 (2024)
78. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 3836–3847 (2023)
79. Zhao, Z., Lai, Z., Lin, Q., Zhao, Y., Liu, H., Yang, S., Feng, Y., Yang, M., Zhang, S., Yang, X., et al.: Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation. *arXiv preprint arXiv:2501.12202* (2025)
80. Zhao, Z., Liu, W., Chen, X., Zeng, X., Wang, R., Cheng, P., Fu, B., Chen, T., Yu, G., Gao, S.: Michelangelo: Conditional 3d shape generation based on shape-image-text aligned latent representation. *Advances in neural information processing systems* **36**, 73969–73982 (2023)