

000 001 002 003 004 005 DREAMCS: GEOMETRY-AWARE TEXT-TO-3D GENER- 006 ATION WITH UNPAIRED 3D REWARD SUPERVISION 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026

ABSTRACT

027 While text-to-3D generation has attracted growing interest, existing methods of-
028 ten struggle to produce 3D assets that align well with human preferences. Current
029 preference alignment techniques for 3D content typically rely on hardly-collected
030 preference-paired multi-view 2D images to train 2D reward models, when then
031 guide 3D generation — leading to geometric artifacts, such as the Janus face
032 problem and geometric incompleteness, due to their inherent 2D bias. To ad-
033 dress these limitations, we construct 3D-MeshPref, the first large-scale unpaired
034 3D preference dataset, featuring diverse 3D meshes annotated by a large language
035 model and refined by human evaluators. We then develop RewardCS, the first re-
036 ward model trained directly on unpaired 3D-MeshPref data using a novel Cauchy-
037 Schwarz divergence objective, enabling effective learning of human-aligned 3D
038 geometric preferences without requiring paired comparisons. Building on this,
039 we propose DreamCS, a unified framework that integrates RewardCS into text-to-
040 3D pipelines — enhancing both implicit and explicit 3D generation with human
041 preference feedback. Extensive experiments show DreamCS outperforms prior
042 methods, producing 3D assets that are both geometrically faithful and human-
043 preferred. Code and models will be released publicly.
044

1 INTRODUCTION

045 Text-to-3D generation has emerged as a key technique for automating 3D asset creation in do-
046 mains such as gaming, film, digital comics, and virtual reality. Recent advances include one-stage
047 optimization-based 2D lifting methods (Shi et al., 2023b; Wang et al., 2023b; Chen et al., 2023; Lin
048 et al., 2023; Wang et al., 2023a; Liu et al., 2023a; Shi et al., 2023a; Qian et al., 2023; Wu et al.,
049 2024c; Zhu et al., 2023b; Kwak et al., 2024), two-stage approaches that separately model geometry
050 and appearance (Lin et al., 2023; Chen et al., 2023), and end-to-end methods (Tsalicoglou et al.,
051 2024; Jun & Nichol, 2023). Despite these impressive advancements, generated 3D assets often mis-
052 align with human preferences, highlighting the need for preference-aware generation frameworks.
053

054 To address this gap, reinforcement learning from human feedback (RLHF) (Christiano et al., 2017;
055 Yang et al., 2024) has recently been incorporated into text-to-3D pipelines (Xie et al., 2024; Ye
056 et al., 2024; Zhou et al., 2025), showing improvements in fidelity and realism. However, cur-
057 rent approaches face two key challenges. First, they rely on costly paired preference-labeled 2D
058 data: for each prompt, multiple 3D assets must be rendered from diverse viewpoints, and pre-
059 ferred/dispreferred image pairs must be manually annotated. This process is computationally in-
060 tensive, time-consuming, and difficult to scale, especially since large preference gaps are rare in
061 practice. Second, these methods provide only 2D view-dependent supervision. Rewards derived
062 from rendered images (e.g., ImageReward (Xu et al., 2023)) may favor assets that appear plausi-
063 ble from some viewpoints while ignoring structural flaws from others, leading to artifacts such as
064 Janus faces or incomplete geometry (Fig. 1). The root cause lies in the absence of 3D-aware reward
065 signals: existing diffusion-based 2D lifting pipelines excel at semantic alignment but lack explicit
066 global geometric supervision, limiting their ability to ensure consistent and plausible 3D structure.
067

068 **Contributions.** In this work, we take a step forward by proposing a novel 3D-reward guided frame-
069 work for text-to-3D generation. To eliminate paired training data and provide geometry-level feed-
070 back during generation, we respectively resolve three associated key challenges by proposing the
071 first large-scale preference-unpaired 3D dataset, the first 3D reward model trained on our dataset
072 with theoretical guarantees, and the first 3D-reward guided text-to-3D generative framework.
073

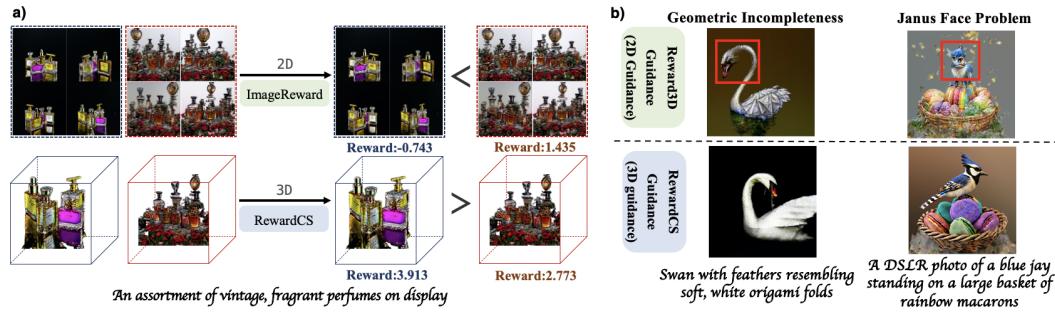


Figure 1: **Comparison of 2D- vs. 3D-based reward models.** a) 2D reward model ImageReward (Xu et al., 2023) assigns high scores to geometrically flawed 3D assets, while our 3D reward model RewardCS better aligns with human preference. b) while DreamFusion (Poole et al., 2022) guided by 2D Reward3D (Ye et al., 2024) produces 3D assets with geometric defects and the Janus problem, while DreamFusion with our RewardCS yields geometrically consistent 3D content.

First, collecting paired preference-labeled 3D data is prohibitively expensive due to complex generation and annotation. To overcome this, we build **3D-MeshPref**, the first large-scale preference-unpaired 3D mesh dataset with over 30,000 samples, each containing a text prompt, 3D asset, and preference reward score. We curate meshes from Cap3D (Luo et al., 2023), evaluate them with Llama-Mesh (Wang et al., 2024) on geometric fidelity, semantic alignment, and structural plausibility, and refine scores with human verification. Assets with higher rewards are designated as preferred, while lower ones serve as dispreferred examples.

Second, prior preference alignment frameworks (Ouyang et al., 2022; Yang et al., 2024; Ye et al., 2024; Zhou et al., 2025) rely on paired comparisons and thus cannot operate in our unpaired settings. Accordingly, we introduce **RewardCS**, the first 3D geometry-aware reward model trained on unpaired data. To enable this, we propose a distribution-level training objective based on Cauchy–Schwarz (CS) divergence (Jenssen et al., 2006), treating preferred and dispreferred assets as samples from two distributions. Optimizing their CS divergence encourages higher rewards for geometrically and semantically superior assets. We also prove the equivalence between CS-divergence learning on unpaired data and vanilla paired preference supervision, providing theoretical guarantee.

Finally, existing text-to-3D frameworks lack native geometry-level feedback and face challenges in mesh representation, differentiability, and reward compatibility. So we develop **DreamCS**, the first 3D-reward guided text-to-3D framework. DreamCS integrates RewardCS into both implicit and explicit pipelines through three innovations: (1) differentiable meshization for end-to-end gradient flow, (2) adaptive mesh fusion for reward compatibility without sacrificing detail, and (3) progressive reward guidance balancing coarse structure refinement with fine-grained optimization.

Extensive experiments on the GPTEval3D benchmark (Wu et al., 2024a) show that DreamCS consistently improves geometric alignment and mesh quality for one-stage and two-stage pipelines like MVDream (Shi et al., 2023b), DreamFusion (Poole et al., 2022), and Magic3D (Lin et al., 2023). Moreover, our 3D reward guidance complements 2D-based methods, and can yield further gains.

2 RELATED WORK

Text-to-3D Generation. Recent advances in text-to-3D generation have largely been driven by the use of pretrained 2D diffusion models, given the scarcity and limited diversity of high-quality 3D datasets. A common strategy in works like DreamFusion (Poole et al., 2022) and SJC (Wang et al., 2023a) involves using Score Distillation Sampling (SDS) to optimize 3D representations like NeRF (Mildenhall et al., 2021)—by using gradients from a 2D diffusion model’s denoising process as supervision to refine the 3D scene. Unfortunately, SDS suffers from issues like inconsistency, the Janus problem, and poor structural fidelity (Wang et al., 2023b). To mitigate these, new methods adopt either one-stage pipelines (Shi et al., 2023b; Zhu et al., 2023b; Wang et al., 2023a; Wu et al., 2024c) that jointly optimize geometry and appearance, or two-stage pipelines (Lin et al., 2023; Chen et al., 2023; Wang et al., 2023b) that decouple them. However, most of them still struggle with alignment to human preferences, particularly in terms of global 3D geometry.

Learning from Human Feedback. Human feedback is increasingly vital for aligning generative models with user intent. In language, RLHF (Ouyang et al., 2022) and DPO (Rafailov et al., 2023)

108 have successfully aligned LLMs to human values, and similar strategies have extended to images (Xu
 109 et al., 2023; Liang et al., 2024) and videos (Wu et al., 2024b; Liu et al., 2025a). In 3D generation,
 110 however, feedback has only recently been explored. DreamReward (Ye et al., 2024) trains a multi-
 111 view image-based reward model from annotated renderings, while DreamDPO (Zhou et al., 2025)
 112 applies DPO using human preferences from multi-view comparisons. Yet both approaches remain
 113 constrained to 2D supervision, aligning view-specific appearance rather than global 3D structure
 114 and often yielding inconsistent geometry. To address this gap, we introduce RewardCS, a 3D reward
 115 model for direct geometric supervision, and further design DreamCS to integrate this feedback into
 116 optimization, ensuring more consistent, plausible, and human-aligned 3D content.

3 METHODOLOGY

119 We first construct a large-scale preference-unpaired 3D mesh dataset, **3D-MeshPref**, in Section 3.1.
 120 Building on this, we develop the first 3D-geometry-aware reward model, **RewardCS**, trained on the
 121 3D-MeshPref dataset, detailed in Section 3.2. To effectively integrate the 3D reward model, Re-
 122 wardCS, into text-to-3D generation pipelines, we propose **DreamCS** in Section 3.3 — a framework
 123 that aligns generated 3D content with human preferences through geometry-aware 3D supervision.

3.1 3D-MESH_PREF: LARGE-SCALE PREFERENCE-UNPAIRED 3D MESH DATASET

124 A major challenge in 3D preference alignment is the lack of labeled datasets, as generating and an-
 125 notating high-quality 3D assets is resource-intensive. Moreover, the complexity of explicit 3D rep-
 126 resentations hinders the construction of consistent paired preference data, limiting the development
 127 of effective 3D reward models. To overcome this, we propose **3D-MeshPref**, the first large-scale
 128 preference-unpaired 3D mesh dataset. It includes 30,000+ samples, each with a text prompt, 3D
 129 asset, and a preference reward score. Our pipeline combines automated LLM-based scoring with
 130 human refinement, enabling scalable annotation while preserving alignment with human judgment.

131 We build our dataset from annotated point clouds in
 132 Cap3D (Luo et al., 2023), sourced from 3D datasets like
 133 Objaverse (Deitke et al., 2023) and ABO (Collins et al.,
 134 2022). ABO contains high-resolution 3D household ob-
 135 jects across 63 categories; Objaverse offers a large-scale
 136 and semantically rich 3D collection of over 21,000 cat-
 137 egories. For diversity, we respectively collect 8,000 and
 138 15,000 samples from ABO and Objaverse, which are then
 139 converted into high-quality mesh representations using Me-
 140 shAnythingV2 (Chen et al., 2024), which consistently out-
 141 performs other meshilization methods in both mesh quality
 142 and computational efficiency. To ensure mesh quality and
 143 compatibility with our 3D reward model, we apply QEM mesh
 144 simplification algorithm (Wu et al., 2004) to ensure the output with a maximum of 16,384 trian-
 145 gular faces. After filtering for mesh
 146 quality and completeness, we obtain 20,000+ 3D meshes. To adapt the 3D reward model to interme-
 147 diate assets generated during SDS optimization, we augment the dataset with over 10,000+ meshes
 148 optimized using SDS-based methods, such as DreamFusion (Poole et al., 2022) and MVDream (Shi
 149 et al., 2023b), under text prompts in Objaverse. These meshes are sampled at early stages of the
 150 optimization process. (see Appendix C) This data augmentation exposes the reward model to both
 151 early- and late-stage geometries, improving robustness and generalization across the optimization.

152 We use Llama-Mesh (Wang et al., 2024), a capable open-source LLM fine-tuned for 3D evaluation,
 153 to score each mesh on a 0-5 Likert scale (Joshi et al., 2015) based on prompt alignment, structural
 154 realism, and visual fidelity. We observe that LLM-based ratings tend to overestimate quality relative
 155 to human perception. Accordingly, human raters verify and re-evaluate them. **See annotation guide-
 156 lines and human refinement process in Appendix F**. We then partition the dataset using a reward
 157 threshold (3.5-4.0): meshes scoring ≥ 4.0 show strong geometric integrity and semantic fidelity
 158 and are labeled as preferred (47% of all 3D asserts), while those ≤ 3.5 typically indicate structural
 159 defects or poor alignment with the prompt and are dispreferred (53%). This clear margin excludes
 160 ambiguous cases, ensuring reliable partitioning—a distribution visualized in Fig. 2. This separation
 161 provides a strong supervisory signal for learning to distinguish relative quality in 3D meshes, and
 162 also provides the training of RewardCS balanced preference pairs—that is, each training instance
 163 contains one preferred and one dispreferred mesh sample, regardless of their raw score values, en-

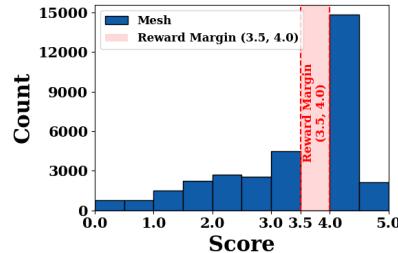


Figure 2: Annotated score distribution of meshes in **3D-MeshPref**.

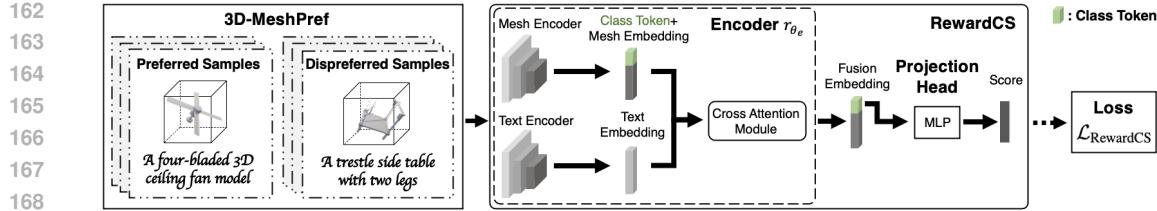


Figure 3: RewardCS is trained on 3D-MeshPref using Cauchy-Schwarz objective.

suring that the learning signal remains unbiased toward either class and avoids reward skewing due to class imbalance. We provide examples and detailed rationale for 3D-MeshPref in Appendix C.

3.2 REWARDCS: 3D GEOMETRY-AWARE REWARD MODEL

Learning reward models for 3D assets is crucial for guiding high-quality text-to-3D generation (Xie et al., 2024; Ye et al., 2024), but it is hindered by the scarcity of paired preference annotations, which are costly and hard to scale due to 3D complexity. Prior preference alignment methods — both general (e.g., DPO (Rafailov et al., 2023) and its variants (Meng et al., 2024; Lai et al., 2024)) and 3D-specific (Zhou et al., 2025; Ye et al., 2024) — rely on paired data of the form $\{\mathbf{m}_i^+, \mathbf{m}_i^-, \mathbf{c}_i\}_{i=1}^m$, where each tuple includes a prompt \mathbf{c}_i drawn from prompt distribution p_C , and a pair of 3D meshes: a preferred one \mathbf{m}_i^+ and a dispreferred one \mathbf{m}_i^- . These samples are drawn from prompt-dependent distributions: $\mathbf{m}_i^+ \sim p^+(\cdot|\mathbf{c}_i)$ and $\mathbf{m}_i^- \sim p^-(\cdot|\mathbf{c}_i)$. The goal is to train reward models to assign higher rewards to preferred assets and guide text-to-3D generative models. However, this paradigm fails when such paired data are unavailable, which is common in real-world 3D settings.

3.2.1 REWARDCS: A NOVEL FRAMEWORK FOR UNPAIRED 3D PREFERENCE LEARNING

To overcome this limitation, we introduce **RewardCS**, the first geometry-aware 3D reward model that learns from unpaired preference data. Rather than relying on explicit preferred-dispreferred pairs, RewardCS leverages a distributional training objective based on Cauchy-Schwarz (CS) divergence (Jensen et al., 2006), enabling effective preference modeling at scale. Our method is built upon the newly introduced 3D-MeshPref dataset, which contains only unpaired samples: high-quality (preferred) and low-quality (dispreferred) 3D meshes associated with diverse prompts.

Let $\{\mathbf{n}_i^+\}_{i=1}^m \sim p^+(\cdot|\mathbf{c}_i)$ and $\{\mathbf{n}_j^-\}_{j=1}^n \sim p^-(\cdot|\mathbf{c}_j)$ denote two sets of unpaired preferred and dispreferred 3D meshes, each conditioned on independently sampled prompts $\{\mathbf{c}_i\}_{i=1}^m$ and $\{\mathbf{c}_j\}_{j=1}^n$ from a distribution p_C . Standard preference alignment methods (Ouyang et al., 2022; Christiano et al., 2017; Ye et al., 2024) are inapplicable in this unpaired setting because they require joint comparisons across samples from the same prompt. Instead, we formulate preference learning as a distribution matching problem. We define a 3D reward model as a function $r_\theta : \mathcal{M} \times \mathcal{C} \rightarrow \mathbb{R}$, parameterized by θ , which maps a 3D mesh $\mathbf{m} \in \mathcal{M}$ and the text condition $\mathbf{c} \in \mathcal{C}$ to a scalar reward. Our reward model, RewardCS, consists of an encoder r_{θ_e} that transforms the input pair into a unified embedding $\mathbf{z} \in \mathcal{Z}$ and a projection head to generate a scalar reward. Crucially, the model is designed to learn the reward conditioned on the text prompt (see Section 3.2.2). To train the 3D reward model, we encode unpaired preferred and dispreferred samples into latent embeddings via the encoder r_{θ_e} :

$$\{\mathbf{x}_i\}_{i=1}^m = r_{\theta_e}(\mathbf{n}_i^+, \mathbf{c}_i) \sim p(\mathbf{x}), \quad \{\mathbf{y}_j\}_{j=1}^n = r_{\theta_e}(\mathbf{n}_j^-, \mathbf{c}_j) \sim p(\mathbf{y}). \quad (1)$$

We treat these embeddings as samples from two distinct distributions $p(\mathbf{x})$ and $p(\mathbf{y})$, representing the overall structure of high- and low-quality 3D meshes in the latent space. Then rather than comparing individual mesh pairs, we aim to separate the distributions $p(\mathbf{x})$ and $p(\mathbf{y})$ in the embedding space. It allows the reward model to generalize over the global statistical properties of high- and low-quality 3D assets, enabling robust learning even when direct pairwise supervision is unavailable. To this end, we introduce a distribution-level loss based on CS divergence (Yin et al., 2025), which quantifies dissimilarity between two distributions in a symmetric and differentiable manner:

$$D_{CS}(p(\mathbf{x}) \parallel p(\mathbf{y})) = -\log \left(\left(\int p(\mathbf{x})p(\mathbf{y}) d\mathbf{x}d\mathbf{y} \right)^2 / \left(\int p(\mathbf{x})^2 d\mathbf{x} \int p(\mathbf{y})^2 d\mathbf{y} \right) \right). \quad (2)$$

Compared to the Kullback–Leibler divergence (Van Erven & Harremos, 2014), the CS divergence offers a tighter generalization bound (Yin et al., 2024) and is more robust than Jensen-Shannon divergence, which lacks a closed-form for Gaussians (Fuglede & Topsoe, 2004; Nielsen, 2019).

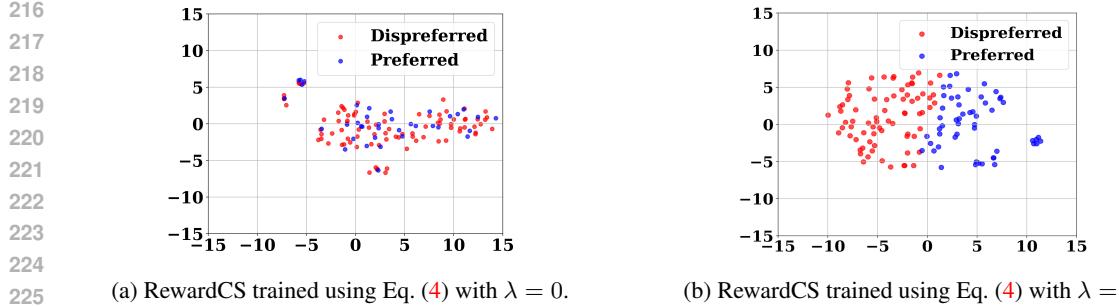


Figure 4: **t-SNE visualizations of class token embeddings based on different RewardCS in 3D-MeshPref data.** The addition of \mathcal{L}_{div} enables RewardCS to separate the underlying data distributions of preferred and dispreferred mesh samples ($p(\mathbf{x})$ and $p(\mathbf{y})$) in the learned embedding space.

This makes CS divergence well-suited for unpaired 3D reward learning, as maximizing it helps the model capture semantic and geometric cues that differentiate preferred from dispreferred assets without explicit supervision for each prompt.

Estimation of CS-Divergence. Since true data distribution densities $p(\mathbf{x})$ and $p(\mathbf{y})$ are unknown, we estimate them via kernel density estimation (KDE) (Weglarczyk, 2018). Let $\{\mathbf{x}_i\}_{i=1}^m$ and $\{\mathbf{y}_j\}_{j=1}^n$ denote the embeddings of preferred and dispreferred samples. The empirical CS divergence is:

$$\hat{D}_{CS}(p(\mathbf{x}) \parallel p(\mathbf{y})) = \log \left(\frac{1}{m^2} \sum_{i,j=1}^m \kappa(\mathbf{x}_i, \mathbf{x}_j) \right) + \log \left(\frac{1}{n^2} \sum_{i,j=1}^n \kappa(\mathbf{y}_i, \mathbf{y}_j) \right) - 2 \log \left(\frac{1}{mn} \sum_{i=1}^m \sum_{j=1}^n \kappa(\mathbf{x}_i, \mathbf{y}_j) \right), \quad (3)$$

where $\kappa(\cdot, \cdot)$ is a kernel function. This estimator possesses desirable properties: it is symmetric and differentiable. Importantly, it supports unequal numbers of preferred and dispreferred unpaired samples ($m \neq n$), enabling distributional alignment without requiring pairwise annotations.

Overall Training Objective of RewardCS. To train our RewardCS model r_{θ} , we combine a regression loss with the CS divergence loss. The final training objective is given by:

$$\mathcal{L}_{\text{RewardCS}}(\theta) = \mathcal{L}_{\text{MSE}}(\theta) + \lambda \mathcal{L}_{\text{div}}(\theta), \quad (4)$$

where \mathcal{L}_{MSE} is the mean squared error for reward prediction, $\mathcal{L}_{\text{div}} = -\hat{D}_{CS}(p(\mathbf{x}); p(\mathbf{y}))$ encourages separation between preferred and dispreferred mesh embeddings, and λ is a hyperparameter.

Fig. 4 shows 2D t-SNE visualizations of mesh embeddings from RewardCS trained with $\lambda = 0$ and $\lambda = 1$. Adding \mathcal{L}_{div} helps the reward model distinguishing preferred and dispreferred meshes in the embedding space, and enables RewardCS to align with human preferences from unpaired data. To further validate the CS divergence objective, we conduct an ablation study to examine its impact on clustering embedding quality and downstream performance (see Appendix D.2).

Theoretical Justification. We analyze CS divergence for training the 3D reward model on unpaired preference data by demonstrating its asymptotic equivalence to paired preference supervision. Optimizing the CS divergence is equivalent to optimizing a quantity in a kernel feature space (Jenssen et al., 2006; Yin et al., 2025):

$$\hat{D}_{CS}(p(\mathbf{x}) \parallel p(\mathbf{y})) = -2 \log \frac{\langle \boldsymbol{\mu}_x, \boldsymbol{\mu}_y \rangle_{\mathcal{H}}}{\| \boldsymbol{\mu}_x \|_{\mathcal{H}} \| \boldsymbol{\mu}_y \|_{\mathcal{H}}}, \quad \text{with} \quad \boldsymbol{\mu}_x = \frac{1}{m} \sum_{i=1}^m \varphi(\mathbf{x}_i), \quad \boldsymbol{\mu}_y = \frac{1}{n} \sum_{i=1}^n \varphi(\mathbf{y}_i), \quad (5)$$

where a characteristic kernel $\kappa(\mathbf{x}, \mathbf{y})$ is defined as $\kappa(\mathbf{x}, \mathbf{y}) = \langle \varphi(\mathbf{x}), \varphi(\mathbf{y}) \rangle_{\mathcal{H}}$, and φ maps samples to a Reproducing Kernel Hilbert Space (RKHS) \mathcal{H} (Smola et al., 2007). In our case, we need to compute the means as

$$\begin{aligned} \boldsymbol{\mu}_x^{\text{paired}} &= \rho(\{(\mathbf{m}_i^+, \mathbf{c}_i)\}_{i=1}^m), \quad \boldsymbol{\mu}_y^{\text{paired}} = \rho(\{(\mathbf{m}_i^-, \mathbf{c}_i)\}_{i=1}^m), \\ \boldsymbol{\mu}_x^{\text{unpaired}} &= \rho(\{(\mathbf{n}_i^+, \mathbf{c}_i)\}_{i=1}^m), \quad \boldsymbol{\mu}_y^{\text{unpaired}} = \rho(\{(\mathbf{n}_j^-, \mathbf{c}_i)\}_{i=1}^n), \end{aligned} \quad (6)$$

where $\rho(\{(\mathbf{m}_i, \mathbf{c}_i)\}_{i=1}^k) = \frac{1}{k} \sum_{i=1}^k \varphi(r_{\theta_e}(\mathbf{m}_i, \mathbf{c}_i))$. Then we can define the CS divergence loss on unpaired data and paired data:

$$\hat{D}_{CS}^{\text{paired}} = -2 \log \frac{\langle \boldsymbol{\mu}_x^{\text{paired}}, \boldsymbol{\mu}_y^{\text{paired}} \rangle_{\mathcal{H}}}{\| \boldsymbol{\mu}_x^{\text{paired}} \|_{\mathcal{H}} \| \boldsymbol{\mu}_y^{\text{paired}} \|_{\mathcal{H}}}, \quad \hat{D}_{CS}^{\text{unpaired}} = -2 \log \frac{\langle \boldsymbol{\mu}_x^{\text{unpaired}}, \boldsymbol{\mu}_y^{\text{unpaired}} \rangle_{\mathcal{H}}}{\| \boldsymbol{\mu}_x^{\text{unpaired}} \|_{\mathcal{H}} \| \boldsymbol{\mu}_y^{\text{unpaired}} \|_{\mathcal{H}}}, \quad (7)$$

Then we pose necessary assumptions widely used in RLHF and divergence analysis (Zhu et al., 2023a; Zou et al., 2025; Kim et al., 2021).

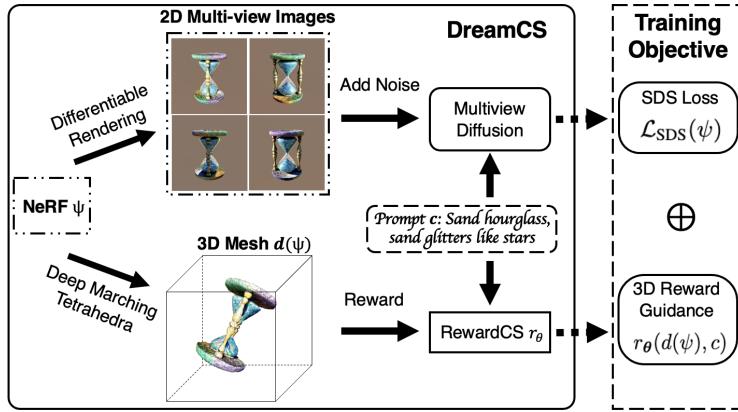


Figure 5: Framework of DreamCS: integrate RewardCS into the SDS model for NeRF optimization.

Assumption 1. *a)* The kernel $\kappa(\cdot, \cdot)$ is bounded, i.e. $\sup_{\mathbf{x}, \mathbf{y} \in \text{dom}(\kappa)} \kappa(\mathbf{x}, \mathbf{y}) \leq K$, and characteristic, meaning the kernel mean embedding is injective. *b)* The population mean embeddings of the preferred and dispreferred samples satisfy: $\|\mu_x^{\text{paired}}\|_{\mathcal{H}}, \|\mu_x^{\text{unpaired}}\|_{\mathcal{H}}, \|\mu_y^{\text{paired}}\|_{\mathcal{H}}, \|\mu_y^{\text{unpaired}}\|_{\mathcal{H}} > 0$.

Assumption 1 (a) and (b) are standard assumptions in kernel methods (Hofmann et al., 2008; Gretton et al., 2006; Muandet et al., 2017), and hold in our case. We use the Gaussian kernel $\kappa_{\sigma}(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|_2^2/(2\sigma^2))$, which is bounded and characteristic, and our data distributions are well-separated, satisfying the non-zero mean embedding condition.

Theorem 1 (Asymptotic Equivalence). *Suppose Assumption 1 holds. With a constant $C > 0$, the empirical CS divergences $\hat{D}_{\text{CS}}^{\text{paired}}$ and $\hat{D}_{\text{CS}}^{\text{unpaired}}$ computed from paired and unpaired data satisfy:*

$$|\hat{D}_{\text{CS}}^{\text{paired}} - \hat{D}_{\text{CS}}^{\text{unpaired}}| \leq C \cdot \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right) \xrightarrow{p} 0 \quad \text{as } m, n \rightarrow \infty, \quad (8)$$

See its proof in Appendix B. Thm.1 shows that the CS divergence computed from paired and unpaired preference data converges asymptotically at a convergence rate of $\mathcal{O}(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}})$. Thus, our RewardCS can learn a latent embedding \mathbf{z} such that maximizing CS divergence between embeddings of preferred and dispreferred samples induces reward separation even without paired data.

3.2.2 NETWORK ARCHITECTURE OF REWARDCS

As shown in Fig. 3, RewardCS contains two key components: an encoder r_{θ_e} to produce a unified representations of 3D geometry and text, and a prediction head to generate a scalar reward score. The encoder r_{θ_e} integrates three modules: a 3D mesh encoder for geometric feature extraction, a text encoder for prompt representation, and a cross-modal fusion module to integrate the two. 3D mesh encoder uses MeshMAE (Liang et al., 2022), a masked Vision Transformer (ViT) autoencoder for 3D meshes, which shows strong performance in capturing global geometric structure and its robustness to partial inputs and structural noise. Each input mesh is divided into 256 non-overlapping patches, with each patch consisting of 64 triangular faces. Each face is represented by a 10-dimensional feature vector capturing geometric attributes like area, angles, surface normals. This strategy enables 3D mesh encoder to capture both localized surface details and global geometric context. For text encoder, it adopts MeshCLIP (Song et al., 2023) to map input text prompt into a sequence of semantic embeddings. To fuse mesh and text information, we apply a cross-attention mechanism: mesh tokens with a learnable 128-dimensional class token act as queries, while text tokens serve as keys and values. During training, the pretrained text encoder is frozen. This design ensures that semantic information from the text condition is injected directly into the mesh representation, so that the resulting embeddings are contextually modulated by the prompt. The class token, which captures the global fused context, is then passed to an MLP-based projection head to produce the final reward score. More rationale of the network architecture is in Appendix D.1.

3.3 DREAMCS: 3D REWARD GUIDANCE FRAMEWORK

Integrating a 3D reward model into text-to-3D frameworks presents two key challenges. Firstly, existing reward methods are designed for 2D paired preferences and cannot directly handle preference-unpaired explicit 3D mesh representations, requiring a differentiable mechanism to convert implicit

324 fields into mesh form. Secondly, meshes generated during training may not meet the reward model’s
 325 structural requirements due to face count constraints, necessitating adaptive topology alignment. To
 326 address these, we propose **DreamCS**, the first geometry-aware 3D reward guidance framework, in-
 327 tegrating our RewardCS model into existing text-to-3D pipelines. As shown in Fig. 5, DreamCS
 328 enables end-to-end optimization via: (1) a differentiable meshization module for reward super-
 329 vision, (2) an adaptive mesh fusion algorithm for topology alignment, and (3) a progressive reward
 330 guidance scheme balancing exploration and reward optimization.

331 **Differentiable Meshization.** Text-to-3D models often operate on implicit representations such as
 332 NeRF or SDFs, which must be converted into explicit mesh geometry to be compatible with our
 333 RewardCS model. We introduce a differentiable meshization module that enables smooth con-
 334 version from these fields to triangle meshes without breaking the optimization flow. For SDS
 335 pipelines (Poole et al., 2022), we use DMTet (Shen et al., 2021) to extract high-quality isosur-
 336 faces differentially. This preserves the gradient flow from the reward signal back to the implicit 3D
 337 parameters. Compared to Marching Cubes, DMTet enables high-fidelity geometry extraction. For
 338 pipelines that optimize explicit meshes, this conversion step is bypassed.

339 **Adaptive Mesh Fusion.** Meshes produced above are often structurally incompatible with the in-
 340 put requirements of our reward model, RewardCS which expects meshes partitioned into 256 non-
 341 overlapping patches, each with 64 faces. To address this, we propose an adaptive mesh fusion
 342 algorithm that simplifies and reorganizes mesh topology while preserving geometric fidelity.

343 The algorithm iteratively merges adjacent faces using two criteria: (1) similarity in face normals and
 344 (2) topological adjacency (shared vertex). When two faces share a vertex and exhibit high normal
 345 similarity, we construct a new face using the shared vertex and one additional vertex from each
 346 original face. When two faces share an edge and have highly aligned normals, a new face is formed
 347 using the edge and a third vertex from either face. This fusion process reduces mesh complexity
 348 while maintaining fine structural details. Importantly, the entire fusion algorithm is differentiable
 349 and integrated into the training loop, maintaining gradient flow and supporting end-to-end learning.
 350 **We validate our proposed adaptive mesh fusion method on mesh geometric fidelity in Appendix E.2.**

351 **Progressive Reward Guidance.** To ensure faithful 3D generation, DreamCS integrates RewardCS
 352 through progressive guidance. At each optimization step, the current differentiable mesh – produced
 353 via our differentiable meshization module – is evaluated by the RewardCS model r_θ , which provides
 354 a gradient-based signal reflecting its geometric plausibility and alignment with input prompt.

355 Let the 3D generation be parameterized by an implicit representation ψ_t . At each optimization
 356 step t , the mesh $d(\psi_t)$ derived from ψ_t via DMTet is evaluated by the reward model r_θ to assign a
 357 scalar score that reflects mesh’s quality w.r.t. the input prompt c , which provides a gradient guidance
 358 signal based on the geometry of mesh vertices and flows back to the implicit field via mesh structure.
 359 Because RewardCS guidance derives from the high-level geometric properties of the mesh, it acts
 360 as a structural prior that promotes global geometric coherence and semantic alignment.

361 The optimization objective at optimization step t is defined as:

$$\mathcal{L}(\psi_t) = \mathcal{L}_{\text{SDS}}(\psi_t) - \alpha(t) \cdot r_\theta(d(\psi_t)|c), \quad (9)$$

362 where \mathcal{L}_{SDS} is vanilla score distillation loss (Poole et al., 2022), and $\alpha(t)$ is a weighting function.

363 To balance the SDS loss and the reward guidance, we adopt a progressive reward guidance sched-
 364 ule (Ye et al., 2024; Liu et al., 2025b). At the beginning of optimization, α is kept low so that
 365 the process is driven primarily by the SDS loss, enabling broad exploration of shape space. As
 366 optimization progresses, α increases linearly from α_{\min} to α_{\max} :

$$\alpha(t) = \alpha_{\min} + (\alpha_{\max} - \alpha_{\min}) \cdot \frac{t}{T}. \quad (10)$$

367 This schedule delays strong reward influence until the geometry stabilizes, balancing the SDS loss
 368 and the reward guidance and promoting coarse-to-fine refinement.

372 4 EXPERIMENT

373 **Setup.** We test our DreamCS by integrating our reward model RewardCS into one-stage SDS-based
 374 text-to-3D generative models like MVDream (Shi et al., 2023b) and DreamFusion (Poole et al.,
 375 2022), and two-stage models like Magic3D (Lin et al., 2023) and Fantasia3D (Chen et al., 2023).
 376 Moreover, we compare with other 3D preference alignment approaches, DreamReward (Ye et al.,
 377 2024), and DreamDPO (Zhou et al., 2025). All baselines use their official implementations where

378 available, otherwise Threestudio’s implementation (Liu et al., 2023b). Our DreamCS is initialized
 379 by pretrained ShapeNet (Chang et al., 2015) and fine-tuned on 3D-MeshPref by minimizing loss
 380 in Eq. (4). All methods run for 20,000 steps, with an additional 10,000 steps for refinement in
 381 texture/geometry-aware baselines. Rendering resolution follows a coarse-to-fine schedule (64×64
 382 for the first 5,000 steps, then 256×256), $\alpha_{\min} = 10$ and $\alpha_{\max} = 20$. We show that the trained
 383 MeshMAE and RewardCS can achieve strong alignment with the performance of LLaMA-Mesh,
 384 showing superiority in generalization and robustness (see details in Appendix D.3).

385 We use all 110 prompts in GPTEval3D (Wu et al., 2024a) for evaluation. We assess texture and
 386 geometry quality using CLIP (CP) (Radford et al., 2021) and VisionReward (VR) (Xu et al., 2024) on
 387 rendered 2D images. These image-text similarity metrics are computed by averaging the similarities
 388 between each view and the given prompt. We proposed a 3D Geometry-Asset Alignment Reward
 389 (GA) metric based on RewardCS trained with a different training configuration. It serves as an
 390 independent evaluator for 3D geometry-text alignment. MiniCPM-o (Yao et al., 2024), a multimodal
 391 vision-language model, rates text-asset alignment, 3D plausibility, and geometry-texture consistency
 392 on a 0–5 Likert scale. In addition, a user study with 30 participants on 60 prompts assesses these
 393 criteria. Finally, we report the proportion of 3D assets suffering from Janus artifacts (Proportion) by
 394 conducting an user study with 60 participants. See more details in Appendix F.

395 Table 1: Comparison of baselines using 110 GPTEval 3D prompts under 12 multi-view images.
 396 Metrics: CP = CLIP, VR = VisionReward, GA = 3D Geometry-Asset Alignment Reward.

Method	CP \uparrow	VR \uparrow	GA \uparrow	Method	CP \uparrow	VR \uparrow	GA \uparrow
DreamFusion	0.22	-3.21	2.53	Magic3D Stage1	0.20	-3.26	2.48
+Reward3D	0.23	-3.11	2.77	+Magic3D Stage2	0.21	-3.03	2.65
+DreamDPO	0.23	-2.98	2.79	Magic3D Stage1+DreamDPO	0.22	-3.05	2.56
+RewardCS	0.25	-2.11	2.96	+Magic3D Stage2+DreamDPO	0.23	-2.94	2.74
+Reward3D+RewardCS	0.25	-2.77	3.22	Magic3D Stage1+Reward3D	0.23	-3.09	2.63
MVDream	0.24	-3.31	2.79	+Magic3D Stage2+Reward3D	0.23	-3.32	2.55
+Reward3D	0.27	-3.12	2.87	Magic3D Stage1+RewardCS	0.23	-1.90	3.00
+DreamDPO	0.23	-2.13	2.87	+Magic3D Stage2+RewardCS	0.24	-1.58	3.22
+RewardCS	0.29	-2.11	2.96	Fantasia3D	0.22	-3.03	2.95
+Reward3D+RewardCS	0.29	-1.92	3.08	+RewardCS	0.25	-1.01	3.33

407 **Quantitative Comparison.** Tab. 1 presents a comparative analysis of one-stage (DreamFusion,
 408 MVDream) and two-stage (Magic3D, Fantasia3D) pipelines, along with their respective 2D-guided
 409 (Reward3D, DreamDPO) and our proposed 3D-guided (RewardCS) variants. As Reward3D (Ye
 410 et al., 2024) is finetuned to optimize ImageReward score, methods with Reward3D guidance often
 411 overfit the ImageReward (IR) metric, but do not generate high-quality 3D assets as shown in Fig. 6.
 412

413 Three key findings emerge from these results. 1) RewardCS consistently improves upon all vanilla
 414 baselines and outperforms 2D-guided variants in CP, GA and VR scores. In Magic3D, RewardCS
 415 achieves the highest GA (3.22), significantly outperforming the baseline (2.65) and the variant with
 416 2D guidance (2.55), while achieving the highest CP score (0.24) and VR score (-1.58). Similarly,
 417 in DreamFusion, RewardCS leads with a GA score of 2.96, surpassing both baseline (2.53) and
 418 the variant with 2D guidance (2.79), and the best CP (0.25) and VR score (-2.11). 2) RewardCS
 419 exhibits strong compatibility with diverse text-to-3D backbones, including both implicit 3D
 420 representation optimization and explicit mesh-based rendering pipelines. In MVDream, it improves the
 421 vanilla baseline from -3.31 to -2.11 in VR score and from 2.79 to 2.96 in GA score, outperforming
 422 the Reward3D variant (-3.12 VR, 2.87 GA). Furthermore, in Fantasia3D’s mesh-rendering-based
 423 pipeline, RewardCS achieves the highest CP and GA score and significantly improves the VR from
 424 -3.03 to -1.01. 3) 3D and 2D guidance prove complementary: integrated guidance in DreamFusion
 425 achieves the peak GA (3.22) score while improving VR to -2.77, demonstrating that Reward3D en-
 426 hances multi-view consistency while RewardCS optimizes geometric alignment. A discussion on
 427 inference computational cost is provided in Appendix G.4. Extended results on comparison with 3D
 428 controllable generation methods (Li et al., 2023; Huang et al., 2024) and advanced 3D generative
 429 backbones (Sun et al., 2023; Xiang et al., 2025) are provided in Appendix G.1 and G.2.

430 **Visualization Comparison.** Fig. 6 compares visual generation results across methods. Vanilla base-
 431 lines such as frequently exhibit textural misalignment and the Janus face problem, as exemplified in
 432 Fantasia3D’s example, where the baseline produces a two-headed geometry. Baselines augmented
 433 with 2D guidance via Reward3D or DreamDPO improve aesthetics but still display artifacts like ge-

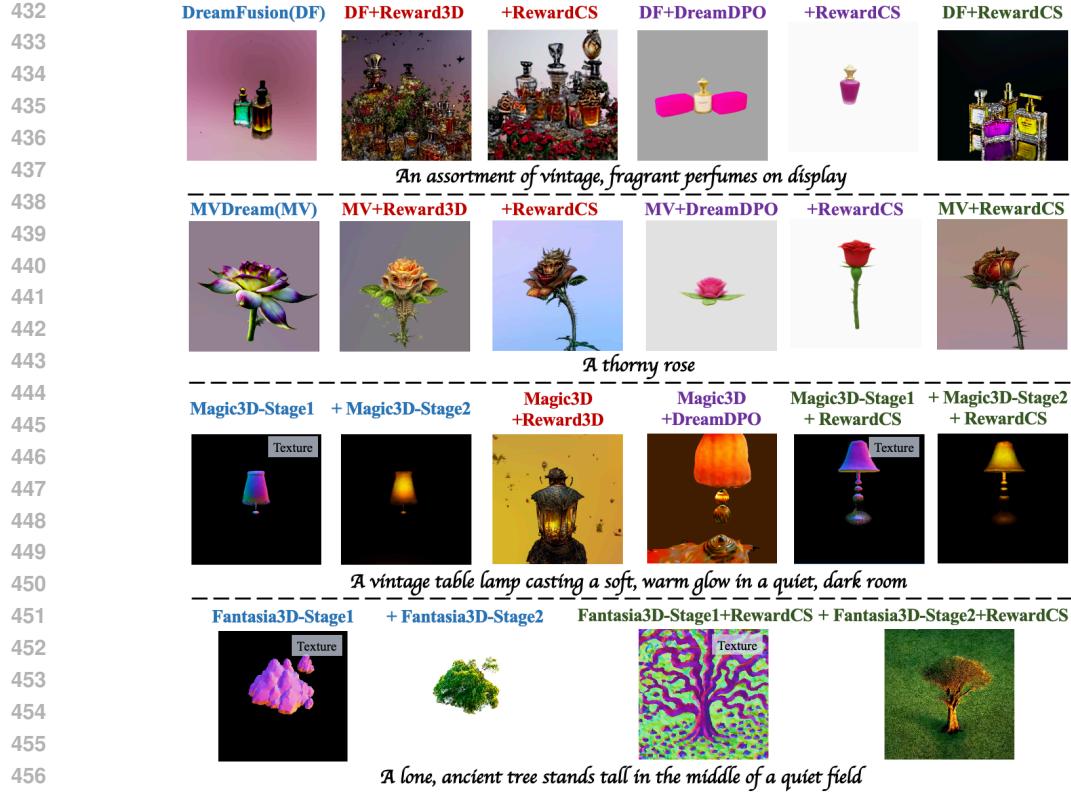


Figure 6: Comparisons with 1-stage generation pipelines (DreamFusion and MVDream) and 2-stage generation pipelines (Magic3D and Fantasia3D). More visualizations are provided in Appendix G.3.

Table 2: Comparison of methods on MiniCPM-o evaluation (left) and user study (right) using (a) 110 and (b) 30 representative GPTEval3D prompts, respectively. Metrics: T-A = Text-Asset Alignment, 3DP = 3D Plausibility, G-T = Geometry-Texture Alignment.

(a) MiniCPM-o evaluation on 110 prompts

Method	T-A ↑	3DP ↑	G-T ↑
MVdream	2.97	3.12	3.08
+Reward3D	3.38	3.35	3.22
+DreamDPO	3.42	3.30	3.11
+RewardCS	3.59	4.05	3.95
+Reward3D+RewardCS	3.61	3.98	3.81

(b) User study on 60 prompts

Method	T-A ↑	3DP ↑	G-T ↑
MVdream	2.90	2.87	2.91
+Reward3D	3.04	3.15	3.07
+DreamDPO	3.15	3.19	3.00
+RewardCS	3.21	3.72	3.59
+Reward3D+RewardCS	3.41	3.51	3.54

Table 3: Proportion of 3D assets suffering Janus artifacts using 60 GPTEval3D prompts.

(a) Evaluation on MVDream

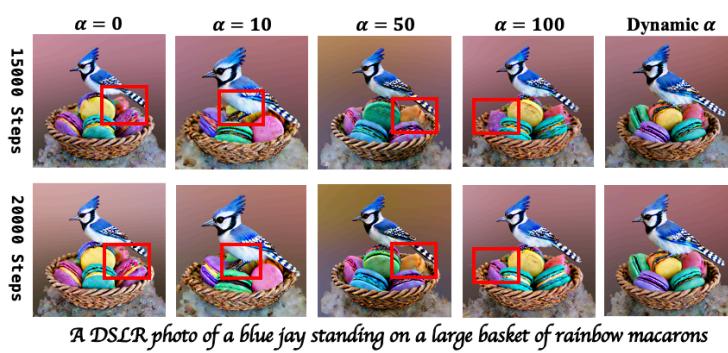
Method	Proportion ↓
MVDream	0.52
+Reward3D	0.44
+DreamDPO	0.43
+RewardCS	0.30
+Reward3D+RewardCS	0.28

(b) Evaluation on DreamFusion

Method	Proportion ↓
DreamFusion	0.61
+Reward3D	0.53
+DreamDPO	0.50
+RewardCS	0.41
+Reward3D+RewardCS	0.39

ometric incompleteness and floaters, as seen in DreamFusion, Magic3D and MVDream’s examples. In contrast, our DreamCS framework, which combines a baseline with RewardCS, demonstrates robust compatibility across all 3D generation backbones, consistently producing assets with superior text alignment, geometric completeness, and texture fidelity. Moreover, our RewardCS is also orthogonal to 2D guidance like Reward3D to further enhance text-to-3D generation.

Evaluation via MiniCPM-o and Human. Tab. 2 (a) reports evaluation results using MiniCPM-o, and Tab. 2 (b) reports findings from a user study. One can observe that RewardCS guidance outperforms 2D preference-based methods on geometry-related criteria while maintaining compet-

Figure 7: Ablation study on the weight of 3D guidance α in MVDream at 15000 and 20000 steps.

itive performance in text-asset alignment. In MiniCPM-o evaluations, MVDream with RewardCS achieves the highest scores in 3D Plausibility (4.05) and G-T Alignment (3.95), significantly exceeding both MVDream and MVDream with Reward3D. User studies confirm this trend, with RewardCS again leading in 3D Plausibility (3.72) and G-T Alignment (3.59).

Evaluation via Janus-Artifact Proportion. Regarding to the Janus problem, to the best of our knowledge, no established benchmark or quantitative metric currently exists for evaluating multi-view consistency or the Janus problem in text-to-3D generation. Consequently, we compare text-to-3D baselines and DreamCS using the proportion of generated 3D assets exhibiting Janus artifacts. As shown in Tab. 3, both MVDream and DreamFusion with RewardCS yields a lower rate of Janus artifacts (0.30 and 0.41) compared to both the vanilla baseline and 2D-guided variants.

Ablation Study. Fig. 7 shows the effect of varying guidance weight α in Eq. (9), balancing the 2D diffusion distillation score and DreamCS 3D guidance, on the quality of generated 3D assets using the MVDream backbone. Low values of $\alpha(0, 10)$ cause geometric ambiguity, while high $\alpha(50, 100)$ results in structural distortions. Our progressive strategy yields can improved results.

5 CONCLUSION

In this paper, we present DreamCS, the first 3D-guided text-to-3D generation framework explicitly designed for human preference alignment. Specifically, we construct a preference dataset on mesh data, 3D-MeshPref. Then we train a 3D reward model, RewardCS, on this dataset to encode human-aligned geometric priors. By incorporating RewardCS into the 3D generation pipeline, DreamCS produces geometrically plausible and texture-consistent 3D assets that are faithfully aligned with human preferences. Extensive experiments demonstrate that DreamCS consistently outperforms both vanilla and 2D-guided baselines across 3D geometry-oriented metrics and human evaluations.

Limitations. Due to limited budget, our evaluation of 3D generation relies on open-source LLMs like MiniCPM-o. Incorporating stronger models like GPT-4 (Wu et al., 2024a) may yield more reliable and fine-grained feedback, and will be considered when budget becomes available. In addition, the current absence of established benchmarks for assessing multi-view consistency and Janus artifacts limits the comprehensiveness of our evaluation. Developing standardized metrics related to the Janus problem would enable more robust comparisons in future work. Moreover, while DreamCS shows promising results, its performance is limited by text-to-3D generative backbone models like DreamFusion. In future work, we plan to explore advanced text-to-3D backbones.

540 ETHICS STATEMENT
541

542 The text-to-3D base models employed in this research may reflect biases or generate sensitive or
543 potentially offensive 3D content, intended solely for academic and scientific purposes. The opinions
544 expressed within generated outputs do not represent the views of the authors. We remain committed
545 to fostering the development of AI technologies which align with ethical standards and reflect
546 societal values.

548 REPRODUCIBILITY STATEMENT
549

550 We detail our work in the Methodology section (Section 3) and describe implementation details in
551 Section 4 and Appendix F.

553 REFERENCES
554

555 Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Cláudio Silva, and Gabriel Taubin. The
556 ball-pivoting algorithm for surface reconstruction. *IEEE transactions on visualization and com-*
557 *puter graphics*, 5(4):349–359, 2002.

558 Tadeusz Caliński and Jerzy Harabasz. A dendrite method for cluster analysis. *Communications in*
559 *Statistics-theory and Methods*, 3(1):1–27, 1974.

561 Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li,
562 Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d
563 model repository. *arXiv preprint arXiv:1512.03012*, 2015.

564 Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and
565 appearance for high-quality text-to-3d content creation. In *Proceedings of the IEEE/CVF interna-*
566 *tional conference on computer vision*, pp. 22246–22256, 2023.

568 Yiwen Chen, Yikai Wang, Yihao Luo, Zhengyi Wang, Zilong Chen, Jun Zhu, Chi Zhang, and Gu-
569 osheng Lin. Meshanything v2: Artist-created mesh generation with adjacent mesh tokenization.
570 *arXiv preprint arXiv:2408.02555*, 2024.

571 Paul F Christiano, Jan Leike, Tom Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep
572 reinforcement learning from human preferences. *Advances in neural information processing sys-*
573 *tems*, 30, 2017.

575 Jasmine Collins, Shubham Goel, Kenan Deng, Achleshwar Luthra, Leon Xu, Erhan Gundogdu,
576 Xi Zhang, Tomas F Yago Vicente, Thomas Dideriksen, Himanshu Arora, et al. Abo: Dataset and
577 benchmarks for real-world 3d object understanding. In *Proceedings of the IEEE/CVF conference*
578 *on computer vision and pattern recognition*, pp. 21126–21136, 2022.

579 David L Davies and Donald W Bouldin. A cluster separation measure. *IEEE transactions on pattern*
580 *analysis and machine intelligence*, (2):224–227, 1989.

581 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig
582 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-
583 tated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
584 *recognition*, pp. 13142–13153, 2023.

586 Herbert Edelsbrunner and Ernst P Mücke. Three-dimensional alpha shapes. *ACM Transactions On*
587 *Graphics (TOG)*, 13(1):43–72, 1994.

588 Bent Fuglede and Flemming Topsøe. Jensen-shannon divergence and hilbert space embedding. In
589 *International symposium onInformation theory, 2004. ISIT 2004. Proceedings.*, pp. 31. IEEE,
590 2004.

592 Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. A kernel
593 method for the two-sample-problem. *Advances in neural information processing systems*, 19,
2006.

594 Thomas Hofmann, Bernhard Schölkopf, and Alexander J Smola. Kernel methods in machine learning. 2008.
 595
 596

597 Tianyu Huang, Yihan Zeng, Zhilu Zhang, Wan Xu, Hang Xu, Songcen Xu, Rynson WH Lau, and
 598 Wangmeng Zuo. Dreamcontrol: Control-based text-to-3d generation with 3d self-prior. In *Pro-
 ceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5364–
 600 5373, 2024.

601 Robert Jenssen, Jose C Principe, Deniz Erdogmus, and Torbjørn Eltoft. The cauchy–schwarz diver-
 602 gence and parzen windowing: Connections to graph theory and mercer kernels. *Journal of the
 603 Franklin Institute*, 343(6):614–629, 2006.

604
 605 Ankur Joshi, Saket Kale, Satish Chandel, and D Kumar Pal. Likert scale: Explored and explained.
 606 *British journal of applied science & technology*, 7(4):396, 2015.

607 Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint
 608 arXiv:2305.02463*, 2023.

609
 610 Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Pro-
 ceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006.

611
 612 Taehyeon Kim, Jaehoon Oh, NakYil Kim, Sangwook Cho, and Se-Young Yun. Comparing
 613 kullback-leibler divergence and mean squared error loss in knowledge distillation. *arXiv preprint
 614 arXiv:2105.08919*, 2021.

615
 616 Min-Seop Kwak, Donghoon Ahn, Inès Hyeonsu Kim, Jin-Hwa Kim, and Seungryong Kim.
 617 Geometry-aware score distillation via 3d consistent noising and gradient consistency modeling.
 618 *arXiv preprint arXiv:2406.16695*, 2024.

619
 620 Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. Step-dpo: Step-
 621 wise preference optimization for long-chain reasoning of llms. *arXiv preprint arXiv:2406.18629*,
 622 2024.

623
 624 Zhiqi Li, Yiming Chen, Lingzhe Zhao, and Peidong Liu. Mvcontrol: Adding conditional control
 625 to multi-view diffusion for controllable text-to-3d generation. *arXiv preprint arXiv:2311.14494*,
 626 2023.

627
 628 Yaqian Liang, Shanshan Zhao, Baosheng Yu, Jing Zhang, and Fazhi He. Meshmae: Masked au-
 629 toencoders for 3d mesh data analysis. In *European conference on computer vision*, pp. 37–54.
 Springer, 2022.

630
 631 Youwei Liang, Junfeng He, Gang Li, Peizhao Li, Arseniy Klimovskiy, Nicholas Carolan, Jiao Sun,
 632 Jordi Pont-Tuset, Sarah Young, Feng Yang, et al. Rich human feedback for text-to-image genera-
 633 tion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 634 pp. 19401–19411, 2024.

635
 636 Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten
 637 Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content
 638 creation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 639 pp. 300–309, 2023.

640
 641 Jie Liu, Gongye Liu, Jiajun Liang, Ziyang Yuan, Xiaokun Liu, Mingwu Zheng, Xiele Wu, Qiulin
 642 Wang, Wenyu Qin, Menghan Xia, et al. Improving video generation with human feedback. *arXiv
 643 preprint arXiv:2501.13918*, 2025a.

644
 645 Jie Liu, Gongye Liu, Jiajun Liang, Ziyang Yuan, Xiaokun Liu, Mingwu Zheng, Xiele Wu, Qiulin
 646 Wang, Wenyu Qin, Menghan Xia, et al. Improving video generation with human feedback. *arXiv
 647 preprint arXiv:2501.13918*, 2025b.

648
 649 Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick.
 650 Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international
 651 conference on computer vision*, pp. 9298–9309, 2023a.

648 Ying-Tian Liu, Yuan-Chen Guo, Vikram Voleti, Ruizhi Shao, Chia-Hao Chen, Guan Luo, Zixin
 649 Zou, Chen Wang, Christian Laforte, Yan-Pei Cao, et al. threestudio: a modular framework for
 650 diffusion-guided 3d generation. *cg.cs.tsinghua.edu.cn*, 2023b.

651

652 Tiange Luo, Chris Rockwell, Honglak Lee, and Justin Johnson. Scalable 3d captioning with pre-
 653 trained models. *Advances in Neural Information Processing Systems*, 36:75307–75337, 2023.

654

655 Yu Meng, Mengzhou Xia, and Danqi Chen. Simpo: Simple preference optimization with a
 656 reference-free reward. *Advances in Neural Information Processing Systems*, 37:124198–124235,
 657 2024.

658

659 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and
 660 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications
 661 of the ACM*, 65(1):99–106, 2021.

662

663 Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, Bernhard Schölkopf, et al. Kernel
 664 mean embedding of distributions: A review and beyond. *Foundations and Trends® in Machine
 665 Learning*, 10(1-2):1–141, 2017.

666

667 Utkarsh Nath, Rajeev Goel, Eun Som Jeon, Changhoon Kim, Kyle Min, Yezhou Yang, Yingzhen
 668 Yang, and Pavan Turaga. Deep geometric moments promote shape consistency in text-to-3d
 669 generation. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*,
 670 pp. 4331–4341. IEEE, 2025.

671

672 Frank Nielsen. On the jensen–shannon symmetrization of distances relying on abstract means.
 673 *Entropy*, 21(5):485, 2019.

674

675 Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong
 676 Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to fol-
 677 low instructions with human feedback. *Advances in neural information processing systems*, 35:
 678 27730–27744, 2022.

679

680 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d
 681 diffusion. *arXiv preprint arXiv:2209.14988*, 2022.

682

683 Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-
 684 Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image
 685 to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint
 686 arXiv:2306.17843*, 2023.

687

688 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 689 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 690 models from natural language supervision. In *International conference on machine learning*, pp.
 691 8748–8763. PMLR, 2021.

692

693 Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea
 694 Finn. Direct preference optimization: Your language model is secretly a reward model. *Advances
 695 in Neural Information Processing Systems*, 36:53728–53741, 2023.

696

697 Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analy-
 698 sis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

699

700 Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra:
 701 a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information
 702 Processing Systems*, 34:6087–6101, 2021.

703

704 Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen,
 705 Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base
 706 model. *arXiv preprint arXiv:2310.15110*, 2023a.

707

708 Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view
 709 diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023b.

702 Alex Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A hilbert space embedding for
 703 distributions. In *International conference on algorithmic learning theory*, pp. 13–31. Springer,
 704 2007.

705

706 Yupeng Song, Naifu Liang, Qing Guo, Jicheng Dai, Junwei Bai, and Fazhi He. Meshclip: Efficient
 707 cross-modal information processing for 3d mesh data in zero/few-shot learning. *Information
 708 Processing & Management*, 60(6):103497, 2023.

709 Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu.
 710 Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. *arXiv preprint
 711 arXiv:2310.16818*, 2023.

712

713 Christina Tsalicoglou, Fabian Manhardt, Alessio Tonioni, Michael Niemeyer, and Federico Tombari.
 714 Textmesh: Generation of realistic 3d meshes from text prompts. In *2024 International Conference
 715 on 3D Vision (3DV)*, pp. 1554–1563. IEEE, 2024.

716 Tim Van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE
 717 Transactions on Information Theory*, 60(7):3797–3820, 2014.

718 Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jaco-
 719 brian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the
 720 IEEE/CVF conference on computer vision and pattern recognition*, pp. 12619–12629, 2023a.

721

722 Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Pro-
 723 lificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation.
 724 *Advances in Neural Information Processing Systems*, 36:8406–8441, 2023b.

725 Zhengyi Wang, Jonathan Lorraine, Yikai Wang, Hang Su, Jun Zhu, Sanja Fidler, and Xiaohui
 726 Zeng. Llama-mesh: Unifying 3d mesh generation with language models. *arXiv preprint
 727 arXiv:2411.09595*, 2024.

728 Stanislaw Weglarczyk. Kernel density estimation and its application. In *ITM web of conferences*,
 729 volume 23, pp. 00037. EDP Sciences, 2018.

730

731 Tong Wu, Guandao Yang, Zhibing Li, Kai Zhang, Ziwei Liu, Leonidas Guibas, Dahua Lin, and
 732 Gordon Wetzstein. Gpt-4v (ision) is a human-aligned evaluator for text-to-3d generation. In
 733 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22227–
 734 22238, 2024a.

735 Xun Wu, Shaohan Huang, Guolong Wang, Jing Xiong, and Furu Wei. Boosting text-to-video gener-
 736 ative model with mllms feedback. In *The Thirty-eighth Annual Conference on Neural Information
 737 Processing Systems*, 2024b.

738

739 Yong Wu, Yuanjun He, and Hongming Cai. Qem-based mesh simplification with global geometry
 740 features preserved. In *Proceedings of the 2nd international conference on Computer graphics and
 741 interactive techniques in Australasia and South East Asia*, pp. 50–57, 2004.

742

743 Zike Wu, Pan Zhou, Xuanyu Yi, Xiaoding Yuan, and Hanwang Zhang. Consistent3d: Towards
 744 consistent high-fidelity text-to-3d generation with deterministic sampling prior. In *Proceedings of
 745 the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9892–9902, 2024c.

746

747 Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen,
 748 Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation.
 749 In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 21469–21480,
 750 2025.

751

752 Desai Xie, Jiahao Li, Hao Tan, Xin Sun, Zhixin Shu, Yi Zhou, Sai Bi, Sören Pirk, and Arie E
 753 Kaufman. Carve3d: Improving multi-view reconstruction consistency for diffusion models with
 754 rl finetuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern
 755 Recognition*, pp. 6369–6379, 2024.

756

757 Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao
 758 Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation.
 759 *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023.

756 Jiazheng Xu, Yu Huang, Jiale Cheng, Yuanming Yang, Jiajun Xu, Yuan Wang, Wenbo Duan, Shen
 757 Yang, Qunlin Jin, Shurun Li, et al. Visionreward: Fine-grained multi-dimensional human prefer-
 758 ence learning for image and video generation. *arXiv preprint arXiv:2412.21059*, 2024.

759

760 Kai Yang, Jian Tao, Jiafei Lyu, Chunjiang Ge, Jiaxin Chen, Weihan Shen, Xiaolong Zhu, and Xiu Li.
 761 Using human feedback to fine-tune diffusion models without any reward model. In *Proceedings
 762 of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8941–8951, 2024.

763

764 Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li,
 765 Weilin Zhao, Zihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint
 766 arXiv:2408.01800*, 2024.

767

768 Junliang Ye, Fangfu Liu, Qixiu Li, Zhengyi Wang, Yikai Wang, Xinzhou Wang, Yueqi Duan, and
 769 Jun Zhu. Dreamreward: Text-to-3d generation with human preference. In *European Conference
 770 on Computer Vision*, pp. 259–276. Springer, 2024.

771

772 Wenzhe Yin, Shujian Yu, Yicong Lin, Jie Liu, Jan-Jakob Sonke, and Efstratios Gavves. Domain
 773 adaptation with cauchy-schwarz divergence. *arXiv preprint arXiv:2405.19978*, 2024.

774

775 Wenzhe Yin, Zehao Xiao, Pan Zhou, Shujian Yu, Jiayi Shen, Jan-Jakob Sonke, and Efstratios
 776 Gavves. Distributional vision-language alignment by cauchy-schwarz divergence. *arXiv preprint
 777 arXiv:2502.17028*, 2025.

778

779 Youcai Zhang, Xinyu Huang, Jinyu Ma, Zhaoyang Li, Zhaochuan Luo, Yanchun Xie, Yuzhuo Qin,
 780 Tong Luo, Yaqian Li, Shilong Liu, et al. Recognize anything: A strong image tagging model.
 781 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 782 1724–1732, 2024.

783

784 Ruowen Zhao, Junliang Ye, Zhengyi Wang, Guangce Liu, Yiwen Chen, Yikai Wang, and Jun Zhu.
 785 Deepmesh: Auto-regressive artist-mesh creation with reinforcement learning. *arXiv preprint
 786 arXiv:2503.15265*, 2025.

787

788 Zhenglin Zhou, Xiaobo Xia, Fan Ma, Hehe Fan, Yi Yang, and Tat-Seng Chua. Dreamdpo: Aligning
 789 text-to-3d generation with human preferences via direct preference optimization. *arXiv preprint
 790 arXiv:2502.04370*, 2025.

791

792 Banghua Zhu, Michael Jordan, and Jiantao Jiao. Principled reinforcement learning with human feed-
 793 back from pairwise or k-wise comparisons. In *International Conference on Machine Learning*,
 794 pp. 43037–43067. PMLR, 2023a.

795

796 Junzhe Zhu, Peiye Zhuang, and Sanmi Koyejo. Hifa: High-fidelity text-to-3d generation with ad-
 797 vanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023b.

798

799 Xiandong Zou, Wanyu Lin, Yuchen Li, and Pan Zhou. Hps: Hard preference sampling for human
 800 preference alignment. *arXiv preprint arXiv:2502.14400*, 2025.

801

802

803

804

805

806

807

808

809

810	CONTENTS OF APPENDIX	
811		
812	A Declaration of LLM usage	16
813		
814	B Theoretical Analysis of CS divergence	16
815		
816	C Details about 3D-MeshPref	19
817		
818	C.1 Dataset Sample Visualization	19
819		
820	C.2 Dataset Construction Rationale	21
821		
822	D Details about RewardCS	22
823		
824	D.1 Design Rationale for RewardCS Architecture	22
825		
826	D.2 Ablation Study on Latent-Space Separation Loss	22
827		
828	D.3 Ablation Study on RewardCS and MeshMAE	23
829		
830	E Details about RewardCS	24
831		
832	E.1 SDS Algorithm	24
833		
834	E.2 Discussion on Adaptive Mesh Fusion	24
835		
836	F Implementation Details	25
837		
838	F.1 Experimental Setup	25
839		
840	F.2 Annotation Pipelines	25
841		
842	F.3 Evaluation Pipelines	25
843		
844	G Extended Empirical Results	28
845		
846	G.1 Comparison with 3D Controllable Generation Methods	28
847		
848	G.2 Comparison with Advanced Text-to-3D Generation Pipeline	28
849		
850	G.3 Qualitative Visualization	29
851		
852	G.4 Additional Computational Cost from Meshization Guidance.	44
853		
854		
855		
856	A DECLARATION OF LLM USAGE	
857		
858		
859	We use Llama-Mesh to score each mesh in our 3D mesh dataset. Additionally, we employ	
860	MiniCPM-o to conduct the MLLM evaluation. All research ideas, methods, and experimental results	
861	presented are original contributions of the authors.	
862		
863		
864	B THEORETICAL ANALYSIS OF CS DIVERGENCE	
865		
866		
867	Theorem 2 (Asymptotic Equivalence). <i>Suppose Assumption 1 holds. With a constant $C > 0$, the empirical CS divergences $\hat{D}_{\text{CS}}^{\text{paired}}$ and $\hat{D}_{\text{CS}}^{\text{unpaired}}$ computed from paired and unpaired data satisfy:</i>	
868	$\left \hat{D}_{\text{CS}}^{\text{paired}} - \hat{D}_{\text{CS}}^{\text{unpaired}} \right \leq C \cdot \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right) \xrightarrow{p} 0 \quad \text{as } m, n \rightarrow \infty, \quad (11)$	
869		
870	Here we provide our proof for Theorem 1.	

864 *Proof.* By Assumption 1(a), we know that the kernel κ is bounded and we have $\sup_{\mathbf{x}} \kappa(\mathbf{x}, \mathbf{x}) \leq$
 865 $K < \infty$, so the canonical feature map φ satisfies $\|\varphi(\mathbf{x})\|_{\mathcal{H}}^2 \leq K$ for all \mathbf{x} .

866 Next, we define the empirical mean embeddings for the paired and unpaired cases:

$$868 \quad \boldsymbol{\mu}_x^{\text{paired}} = \frac{1}{m} \sum_{i=1}^m \varphi(r_{\theta_e}(\mathbf{m}_i^+, \mathbf{c}_i)), \quad \boldsymbol{\mu}_y^{\text{paired}} = \frac{1}{m} \sum_{i=1}^m \varphi(r_{\theta_e}(\mathbf{m}_i^-, \mathbf{c}_i)), \quad (12)$$

$$871 \quad \boldsymbol{\mu}_x^{\text{unpaired}} = \frac{1}{m} \sum_{i=1}^m \varphi(r_{\theta_e}(\mathbf{n}_i^+, \mathbf{c}_i)), \quad \boldsymbol{\mu}_y^{\text{unpaired}} = \frac{1}{n} \sum_{j=1}^n \varphi(r_{\theta_e}(\mathbf{n}_j^-, \mathbf{c}_j)). \quad (13)$$

874 By the law of large numbers in Hilbert spaces, for i.i.d. samples, we obtain the weak convergence:

$$876 \quad \frac{1}{m} \sum_{i=1}^m \varphi(r_{\theta_e}(\mathbf{m}_i^+, \mathbf{c}_i)), \frac{1}{m} \sum_{i=1}^m \varphi(r_{\theta_e}(\mathbf{n}_i^+, \mathbf{c}_i)) \xrightarrow{p} \mathbb{E}_{\mathbf{c} \sim p_{\mathbf{C}}, \mathbf{m}^+ \sim p^+(\cdot | \mathbf{c})} [\varphi(r_{\theta_e}(\mathbf{m}^+, \mathbf{c}))], \quad (14)$$

$$879 \quad \frac{1}{n} \sum_{j=1}^n \varphi(r_{\theta_e}(\mathbf{m}_j^-, \mathbf{c}_j)), \frac{1}{n} \sum_{j=1}^n \varphi(r_{\theta_e}(\mathbf{n}_j^-, \mathbf{c}_j)) \xrightarrow{p} \mathbb{E}_{\mathbf{c} \sim p_{\mathbf{C}}, \mathbf{m}^- \sim p^-(\cdot | \mathbf{c})} [\varphi(r_{\theta_e}(\mathbf{m}^-, \mathbf{c}))]. \quad (15)$$

882 Thus, we have the convergence of the empirical mean embeddings to their respective population
 883 mean embeddings in the RKHS:

$$884 \quad \boldsymbol{\mu}_x^{\text{paired}}, \boldsymbol{\mu}_x^{\text{unpaired}} \xrightarrow{p} \boldsymbol{\mu}_x^*, \quad \boldsymbol{\mu}_y^{\text{paired}}, \boldsymbol{\mu}_y^{\text{unpaired}} \xrightarrow{p} \boldsymbol{\mu}_y^*, \quad (16)$$

886 where the population mean embeddings are defined as:

$$887 \quad \boldsymbol{\mu}_x^* = \mathbb{E}_{\mathbf{c} \sim p_{\mathbf{C}}, \mathbf{m}^+ \sim p^+(\cdot | \mathbf{c})} [\varphi(r_e(\mathbf{m}^+, \mathbf{c}))], \quad \boldsymbol{\mu}_y^* = \mathbb{E}_{\mathbf{c} \sim p_{\mathbf{C}}, \mathbf{m}^- \sim p^-(\cdot | \mathbf{c})} [\varphi(r_e(\mathbf{m}^-, \mathbf{c}))] \quad (17)$$

888 Applying the triangle inequality for the distance between the paired and unpaired empirical mean
 889 embeddings, we obtain:

$$891 \quad \|\boldsymbol{\mu}_x^{\text{paired}} - \boldsymbol{\mu}_x^{\text{unpaired}}\|_{\mathcal{H}} \leq \|\boldsymbol{\mu}_x^{\text{paired}} - \boldsymbol{\mu}_x\|_{\mathcal{H}} + \|\boldsymbol{\mu}_x^{\text{unpaired}} - \boldsymbol{\mu}_x\|_{\mathcal{H}} \xrightarrow{p} 0 \quad (18)$$

892 and similarly for $\boldsymbol{\mu}_y$.

894 Therefore, the sample means from the paired and unpaired cases converge to the same population
 895 mean in RKHS norm. Explicitly, we have:

$$896 \quad \|\boldsymbol{\mu}_x^{\text{paired}} - \boldsymbol{\mu}_x^{\text{unpaired}}\|_{\mathcal{H}} \xrightarrow{p} 0, \quad \|\boldsymbol{\mu}_y^{\text{paired}} - \boldsymbol{\mu}_y^{\text{unpaired}}\|_{\mathcal{H}} \xrightarrow{p} 0. \quad (19)$$

898 To quantify the rate of convergence, we apply standard results from the theory of Hilbert space-
 899 valued random variables (Smola et al., 2007). For all $\varepsilon > 0$:

$$900 \quad \mathbb{P} (\|\boldsymbol{\mu}_x^{\text{paired}} - \boldsymbol{\mu}_x\|_{\mathcal{H}} \geq \varepsilon) \leq 2 \exp \left(-\frac{m\varepsilon^2}{2K} \right), \quad (20)$$

903 and similarly for $\boldsymbol{\mu}_y^{\text{unpaired}}$.

904 Therefore, the convergence rate of the difference between the paired and unpaired empirical mean
 905 embeddings is given by:

$$907 \quad \|\boldsymbol{\mu}_x^{\text{paired}} - \boldsymbol{\mu}_x^{\text{unpaired}}\|_{\mathcal{H}} = \mathcal{O}_p \left(\frac{1}{\sqrt{m}} \right). \quad (21)$$

$$909 \quad \|\boldsymbol{\mu}_y^{\text{paired}} - \boldsymbol{\mu}_y^{\text{unpaired}}\|_{\mathcal{H}} = \mathcal{O}_p \left(\frac{1}{\sqrt{n}} \right). \quad (22)$$

912 We now consider the empirical CS divergence functional defined as:

$$913 \quad f(\boldsymbol{\mu}_x, \boldsymbol{\mu}_y) = -2 \log \left(\frac{\langle \boldsymbol{\mu}_x, \boldsymbol{\mu}_y \rangle_{\mathcal{H}}}{\|\boldsymbol{\mu}_x\|_{\mathcal{H}} \|\boldsymbol{\mu}_y\|_{\mathcal{H}}} \right) \quad (23)$$

916 which is continuous on $\mathcal{H} \setminus \{0\} \times \mathcal{H} \setminus \{0\}$. This condition is satisfied under Assumption 1, which
 917 ensures that the population mean embeddings are non-degenerate due to the boundedness and uni-
 versality of the kernel.

918 Hence, based on the continuous mapping theorem, we yield:
 919

$$920 \quad |f(\boldsymbol{\mu}_x^{\text{paired}}, \boldsymbol{\mu}_y^{\text{paired}}) - f(\boldsymbol{\mu}_x^{\text{unpaired}}, \boldsymbol{\mu}_y^{\text{unpaired}})| \xrightarrow{p} 0 \quad (24)$$

921 That is,

$$922 \quad |\hat{D}_{\text{CS}}^{\text{paired}} - \hat{D}_{\text{CS}}^{\text{unpaired}}| \xrightarrow{p} 0 \quad (25)$$

923 as $m, n \rightarrow \infty$.

924 Furthermore, under Assumption 1(b), the CS divergence functional $f(\cdot, \cdot)$ is Lipschitz continuous,
 925 which allows us to quantify the rate of convergence. Using the previously established rates, we can
 926 conclude:
 927

$$928 \quad |f(\boldsymbol{\mu}_x^{\text{paired}}, \boldsymbol{\mu}_y^{\text{paired}}) - f(\boldsymbol{\mu}_x^{\text{unpaired}}, \boldsymbol{\mu}_y^{\text{unpaired}})| = \mathcal{O}_p \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right). \quad (26)$$

929 Thus, the convergence rate is bounded by:
 930

$$931 \quad |\hat{D}_{\text{CS}}^{\text{paired}} - \hat{D}_{\text{CS}}^{\text{unpaired}}| = \mathcal{O}_p \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right) \leq C \cdot \left(\frac{1}{\sqrt{m}} + \frac{1}{\sqrt{n}} \right) \quad (27)$$

932 where C is a constant that depends on the properties of the kernel function and the Lipschitz constant
 933 of the empirical CS divergence. \square
 934

935
 936
 937
 938
 939
 940
 941
 942
 943
 944
 945
 946
 947
 948
 949
 950
 951
 952
 953
 954
 955
 956
 957
 958
 959
 960
 961
 962
 963
 964
 965
 966
 967
 968
 969
 970
 971

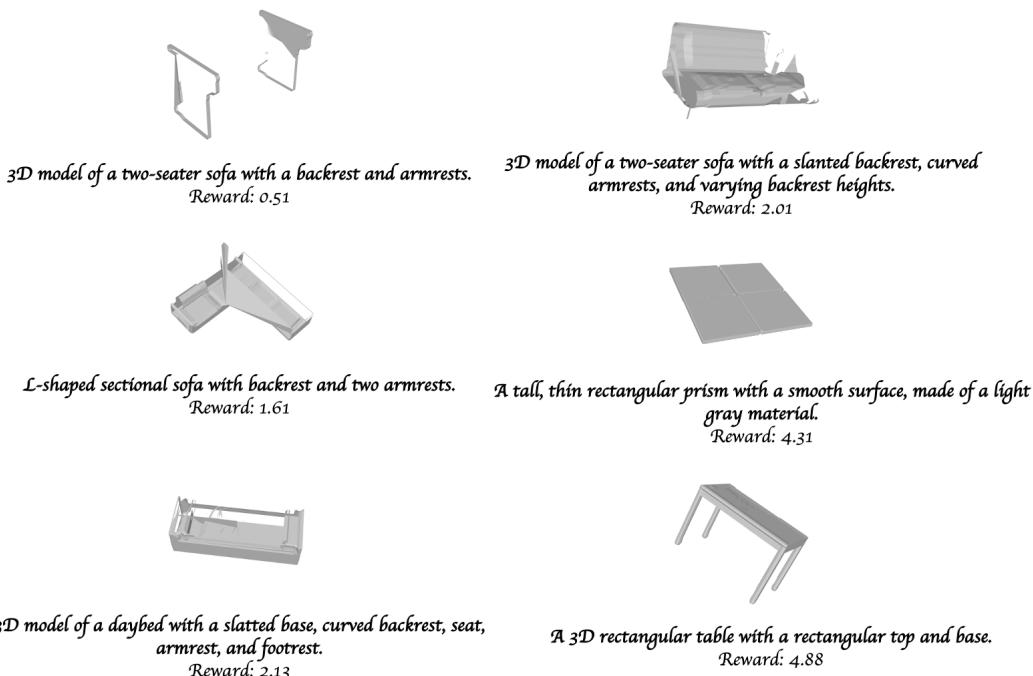
972 C DETAILS ABOUT 3D-MESH_PREF
973974 C.1 DATASET SAMPLE VISUALIZATION
975

Figure 8: Example instances from our large-scale preference-unpaired dataset 3D-MeshPref.

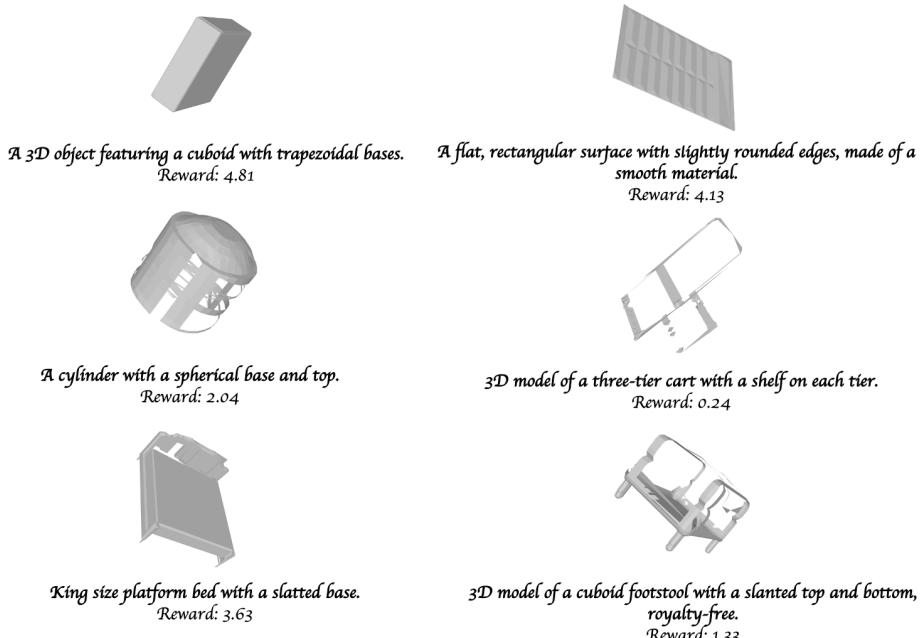


Figure 9: Example instances from our large-scale preference-unpaired dataset 3D-MeshPref.

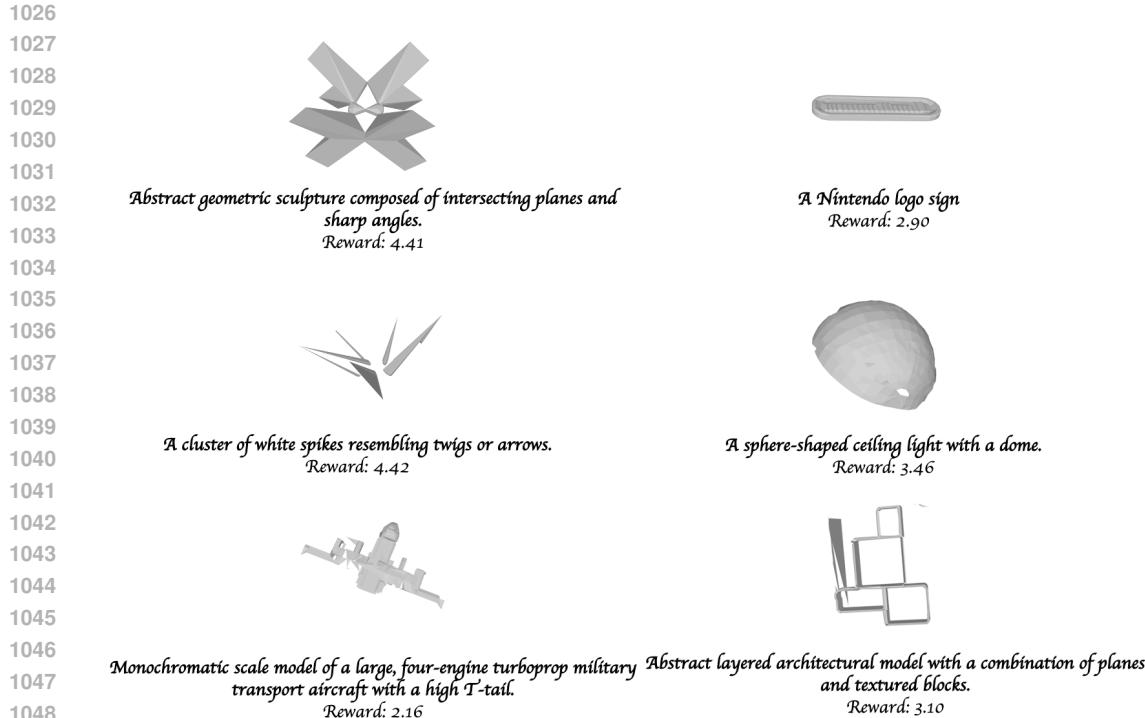


Figure 11: Example instances from our large-scale preference-unpaired dataset 3D-MeshPref.

1080 As presented in Figs. 8, 9, 10, 11, we provide diverse 3D mesh instances in our curated 3D-MeshPref.
 1081 3D-MeshPref 3D assets are semantically diverse, high-quality. Each example instance contains a
 1082 text prompt, 3D asset, and a preference reward score.
 1083

1084 C.2 DATASET CONSTRUCTION RATIONALE

1085 We detail the rationale behind our design choices for each component and describe the measures we
 1086 took to ensure data quality.
 1087

1088 **(a) Meshilization Method.** We explored both traditional point cloud-to-mesh methods, such as
 1089 Alpha Shapes (Edelsbrunner & Mücke, 1994), Ball Pivoting Algorithm (Bernardini et al., 2002),
 1090 Poisson Reconstruction (Kazhdan et al., 2006), and deep learning-based approaches, including Me-
 1091 shAnythingV2 (Chen et al., 2024) and DeepMesh (Zhao et al., 2025). Through systematic compar-
 1092 isons, we found that MeshAnythingV2, particularly with its adjacent mesh tokenization design, con-
 1093 sistently outperforms other alternatives in both mesh quality and computational efficiency. Based
 1094 on these observations, we selected MeshAnythingV2 as it offers the best trade-off among current
 1095 methods.

1096 **(b) Scoring.** Given the current scarcity of 3D reward models, we adopted Llama-Mesh (Wang et al.,
 1097 2024) — a large language model fine-tuned for 3D mesh evaluation — which, to our knowledge,
 1098 is the most capable open-source solution for mesh scoring. Its strong mesh understanding and
 1099 generalization ability derive from its large-scale supervised training data and its alignment with
 1100 natural language instruction, making it well-suited for scoring prompt-conditioned meshes in our
 1101 setting.

1102 To further validate the scoring capability of Llama-Mesh, we follow the DeepMesh (Zhao et al.,
 1103 2025) preference pair construction pipeline and construct a preference mesh dataset with 300 prefer-
 1104 ence paired 3D meshes. Each pair includes a human-verified positive and negative example. Llama-
 1105 Mesh evaluates all samples using the annotation criteria described in Appendix F.2 and produces
 1106 a ranking for each pair. Comparing these rankings with the ground-truth preferences, we find that
 1107 Llama-Mesh correctly identifies the preferred mesh in 93.6% of cases, demonstrating strong con-
 1108 sistency with human judgment and supporting its suitability as an initial scorer for constructing
 1109 3D-MeshPref.

1110 Due to the scarcity of large-scale human-annotated 3D data, prior work in 3D reward modeling (Ye
 1111 et al., 2024) and 3D generation (Zhao et al., 2025) often relies on automated 3D asset pipelines with
 1112 human verification. 3D-MeshPref adopts this practical approach, enabling scalable and effective
 1113 training of RewardCS while maintaining alignment with human preference.

1114 **(c) Human Refinement.** We recognize that LLM-based ratings may sometimes overestimate mesh
 1115 quality relative to human perception. To address this, we incorporate manual verification and cor-
 1116 rection into the data pipeline to ensure the final dataset reflects human preferences more faithfully.

1117 In addition, to adapt the 3D reward model to intermediate assets generated during SDS optimization,
 1118 we augment the dataset with over 10,000+ meshes optimized using SDS-based methods, such as
 1119 DreamFusion (Poole et al., 2022) and MVDream (Shi et al., 2023b), under text prompts in ABO and
 1120 Objaverse. These meshes are sampled at early stages of the optimization process. Thus, we sample
 1121 meshes not only at the final stages of SDS optimization but also at early stages of the process. By
 1122 including these intermediate 3D geometries in 3D-MeshPref, the reward model is exposed to a wider
 1123 spectrum of possible outputs, from coarse to refined structures.

1124 In summary, we carefully selected the most capable mesh generation and scoring models currently
 1125 available, and augmented them with human refinement to construct 3D-MeshPref. In addition, we
 1126 include the intermediate 3D meshes across the SDS optimization in 3D-MeshPref to further extend
 1127 the domain of our proposed RewardCS.

1128
 1129
 1130
 1131
 1132
 1133

1134 **D DETAILS ABOUT REWARDCS**
 1135

1136 **D.1 DESIGN RATIONALE FOR REWARDCS ARCHITECTURE**
 1137

1138 Existing mesh encoders typically fall into two categories: graph-based and transformer-based.
 1139 Transformer-based encoders like MeshMAE offer superior geometric representations given the mesh
 1140 input compared to the graph-based encoders, while there is currently no mesh encoder that jointly
 1141 models both high-resolution geometry and texture semantics due to lack of large-scale 3D mesh
 1142 data. Thus, we adopt MeshMAE, a transformer-based encoder, in RewardCS due to its strong per-
 1143 formance in capturing global geometric structure and its robustness to partial inputs and structural
 1144 noise. MeshMAE is pretrained using masked autoencoding on large-scale 3D mesh data, making it
 1145 well-suited for downstream tasks requiring shape understanding.

1146 **D.2 ABLATION STUDY ON LATENT-SPACE SEPARATION LOSS**
 1147

1148 We provide empirical evidence demonstrating the necessity of the separation loss \mathcal{L}_{div} in Equation. 4
 1149 in RewardCS. Our findings validate that this component improves embedding quality and text-to-3D
 1150 performance.

1151
 1152 **Table 4: Clustering performance under different λ values.**
 1153

λ	SCI \uparrow	CHI \uparrow	DBI \downarrow
0	-0.21	0.52	13.65
0.5	0.38	50.23	6.89
1	0.64	212.92	0.97
10	0.58	340.30	0.86

1160
 1161 We conduct ablations with four different λ values: 0, 0.5, 1, and 10. For each setting, we train a
 1162 separate RewardCS model and evaluate the quality of the learned embeddings using three standard
 1163 clustering metrics across 10 test batches from the 3D-MeshPref dataset: Silhouette Coefficient In-
 1164 dex (SCI) (Rousseeuw, 1987), Calinski–Harabasz Index (CHI) (Caliński & Harabasz, 1974), and
 1165 Davies–Bouldin Index (DBI) (Davies & Bouldin, 2009). Across all three metrics shown in Table 4,
 1166 we observe consistent improvements in cluster separability between preferred and dispreferred sam-
 1167 ples as λ increases, confirming that introducing \mathcal{L}_{div} effectively enhances the structure of the latent
 1168 space for reward learning.

1169
 1170 **Table 5: Performance on 3D Generation using RewardCS trained under different λ values.**
 1171

λ	IR	VR	GA	T-A(M)	3DP(M)	G-T(M)	T-A(U)	3DP(U)	G-T(U)
0	-0.61	-4.10	1.97	2.91	2.80	2.62	2.25	2.48	2.23
0.5	-0.56	-3.28	2.88	3.18	3.31	3.72	3.17	3.23	3.01
1	-0.43	-2.10	2.95	3.21	4.09	4.02	3.27	3.69	3.65
10	-1.48	-4.28	1.26	2.38	2.23	2.07	2.21	2.06	2.07

1172
 1173 Then We incorporate RewardCS trained with different λ values into MVDream and evaluate differ-
 1174 ent settings on GPTEval3D. We report results across the same metrics in our Tab. 1 and Tab. 2, where
 1175 MLLM evaluation is denoted as M and user study is denoted as U. As shown in Tab. 5, RewardCS
 1176 trained with $\lambda = 0.5$ and 1 consistently outperforms the $\lambda = 0$ variant across all metrics, with
 1177 $\lambda = 1$ yielding the strongest performance overall. These results demonstrate the positive impact of
 1178 on practical 3D asset quality.

1179
 1180 These empirical findings support our theoretical results (Theorem 1), which show that optimizing
 1181 CS divergence between unpaired preferred and dispreferred embeddings induces meaningful reward
 1182 separation. Our experiments confirm that this theoretical formulation can enhance embedding qual-
 1183 ity and improve generation fidelity in practice.

1188 D.3 ABLATION STUDY ON REWARDCS AND MESHMAE
11891190 **Generalization Ability of RewardCS.** Our evaluation (Table 2) is based on the GPTEval3D bench-
1191 mark, which consists of 110 test prompts that are out-of-distribution relative to the 3D-MeshPref
1192 dataset. These experiments demonstrate that RewardCS variants consistently outperform vanilla
1193 text-to-3D generation baselines, producing 3D assets with better 3DP and G-T, and thus it can gen-
1194 eralize well.1195 we curated a novel evaluation setup to measure generalization of RewardCS. We sampled 200 dis-
1196 tinct noun tags from the Recognize Anything (Zhang et al., 2024) tag set, and refined them into
1197 descriptive prompts using DeepSeek-R1. These prompts do not overlap with those in Cap3D, en-
1198 suring zero prompt-level data leakage. We used these prompts to generate 3D assets using two
1199 text-to-3D backbones: MVDream and DreamFusion, each generating 100 assets. These outputs are
1200 distinct from Cap3D assets and represent unseen real-world inference settings.1201 Then to assess how well RewardCS scores align with human-perceived quality, we used LLaMA-
1202 Mesh, a strong baseline scorer, to evaluate the generated assets based on three criteria F: (1) prompt
1203 alignment, (2) structural plausibility, and (3) geometric completeness. We then compared RewardCS
1204 scores against LLaMA-Mesh scores using: Mean Squared Error (MSE) to quantify absolute align-
1205 ment between RewardCS and LLaMA-Mesh outputs; and Classification Accuracy (Acc) by dis-
1206 cretizing the score interval [0,5] into 20 equal bins. A prediction is counted as correct if the Re-
1207 wardCS and LLaMA-Mesh scores fall in the same bin.1208 RewardCS achieves 0.06 and 83.10 in MSE and Acc, respectively, and thus show strong alignment
1209 with LLaMA-Mesh across both metrics, even on generated samples that lie far outside the training
1210 distribution. These results demonstrate that RewardCS effectively generalizes to high-quality 3D
1211 assets produced by text-to-3D pipelines like MVDream and DreamFusion, despite being trained on
1212 preference data from a different source.1213 **Robustness of MeshMAE.** Our evaluation is conducted on the GPTEval3D benchmark. The results
1214 show that RewardCS variants consistently outperform vanilla text-to-3D generation baselines, pro-
1215 ducing high-quality 3D assets. This performance indicates that MeshMAE—the mesh encoder used
1216 in RewardCS—is robust to geometric variation and is capable of effectively guiding the generation
1217 framework.1218 To further evaluate whether MeshMAE can effectively capture mesh quality, we constructed 100
1219 paired prompts from 100 unique noun tags in the Recognize Anything (Zhang et al., 2024) dataset.
1220 Each pair consists of a positive prompt with detailed object descriptions and a negative prompt,
1221 derived from the positive one, but modified to include geometry-related defects (e.g., holes, cracks,
1222 asymmetry, and surface degradation). Then, using MVDream, we generated 3D assets from these
1223 prompt pairs, producing meshes with topology differences.1224 To establish reference scores, we used LLaMA-Mesh, a strong mesh-level quality assessor, to score
1225 each asset on three criteria: (1) prompt alignment, (2) realism and structural plausibility, and (3)
1226 geometric completeness. This evaluation process reflects the same preference criteria used in con-
1227 structing the 3D-MeshPref dataset.1228 Finally, we used RewardCS (with MeshMAE as the encoder) to score the same 3D assets and com-
1229 pared the outputs to those from LLaMA-Mesh using: Mean Squared Error (MSE) between Re-
1230 wardCS and LLaMA-Mesh scores, and Classification Accuracy (Acc) by dividing the score range
1231 into 20 bins and checking score agreement within bins.1232 Reward CS achieves the MSE of 0.10 and Acc of 80.98. The results show strong alignment between
1233 RewardCS and LLaMA-Mesh across both metrics, indicating that MeshMAE captures meaningful
1234 distinctions in geometric quality and structural attributes across meshes with varying topology.
12351236
1237
1238
1239
1240
1241

1242 **E DETAILS ABOUT REWARDCS**
 1243

1244 **E.1 SDS ALGORITHM**
 1245

1246 We consider a differentiable rendering function $g(\psi, m)$ which maps a 3D scene representation
 1247 parameterized by $\psi \in \Psi$ and a camera parameter $m \in \mathcal{M}$ to a 2D image x .

1248 Let $\phi(x_t \mid c)$ denote a pretrained 2D conditional diffusion model, where x_t is the noisy image at
 1249 diffusion timestep t and c is a conditioning text prompt. The corresponding noise prediction network
 1250 is $\epsilon_\phi(x_t, t, c)$.

1251 Following the SDS framework (Poole et al., 2022), we can compute the gradient of the SDS loss
 1252 with respect to the 3D parameters ψ via the chain rule at step t :

$$1254 \quad \nabla_{\psi_t} \mathcal{L}_{\text{SDS}}(\psi_t; m, c) = \mathbb{E}_{t, \epsilon, m} \left[w(t) (\epsilon_\phi(x_t, t, c) - \epsilon) \frac{\partial x_t}{\partial \psi_t} \right], \quad (28)$$

1256 where the latent x_t is obtained by
 1257

$$1258 \quad x_t = \sqrt{\bar{\alpha}_t} g(\psi_t, m) + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad (29)$$

1259 with $\bar{\alpha}_t$ the cumulative product of the diffusion schedule, $\epsilon \sim \mathcal{N}(0, I)$, and $w(t)$ a timestep-
 1260 dependent weighting factor.
 1261

1262 **E.2 DISCUSSION ON ADAPTIVE MESH FUSION**
 1263

1264 To assess the impact of the Adaptive Mesh Fusion (AMF) step on geometric fidelity, we evaluate
 1265 two standard surface-level metrics that quantify the deviation between the mesh before (A) and after
 1266 (B) applying AMF. All meshes are normalized by the diagonal length of their bounding box for
 1267 scale-invariant evaluation.

1268 **Chamfer Distance (CD)** measures the average distance from points on one surface to the closest
 1269 points on the other surface, computed as:
 1270

$$1271 \quad \text{CD}(A, B) = \frac{1}{|A|} \sum_{x \in A} \min_{y \in B} \|x - y\|_1 + \frac{1}{|B|} \sum_{y \in B} \min_{x \in A} \|y - x\|_1, \quad (30)$$

1274 where $|A|$ and $|B|$ denote the cardinality of points on meshes A and B , respectively.

1275 **95%-Hausdorff Distance (HD)** captures the worst-case deviation between the two surfaces while
 1276 remaining robust to outliers:
 1277

$$1278 \quad \text{HD}(A, B) = \max \left(\text{Quantile}_{95} \left[\min_{y \in B} \|x - y\|_2 \right]_{x \in A}, \text{Quantile}_{95} \left[\min_{x \in A} \|y - x\|_2 \right]_{y \in B} \right). \quad (31)$$

1281 We evaluate 40 prompts from GPTEval3D for three baselines: MVDream, DreamFusion, and
 1282 DreamCraft3D. Measurements are reported for meshes before and after AMF at SDS steps 4000,
 1283 8000, 12000, 16000, and 20000. The averaged results are shown in Table 6.

1284 **Table 6: Geometric deviation of meshes before and after the Adaptive Mesh Fusion step in**
 1285 **DreamCS. CD and HD denote Chamfer Distance and Hausdorff Distance (95%-quantile), respec-**
 1286 **tively.**

Method	CD ($\times 10^{-3}$)	HD ($\times 10^{-2}$)
MVDream	3.1	2.9
DreamFusion	4.0	3.2
DreamCraft3D	2.7	2.4

1294 As shown Table 6, the geometric deviation introduced by AMF is negligible (less than 0.35% average
 1295 error), confirming that AMF preserves geometric details throughout optimization.

1296 **F IMPLEMENTATION DETAILS**
12971298 **F.1 EXPERIMENTAL SETUP**
1299

1300 The conversion of point clouds to meshes was performed using the MeshAnythingV2 frame-
1301 work (Chen et al., 2024) across 8 L40-S GPUs. For dataset annotation, we employed 4 L40-S
1302 GPUs. The training of the reward model, DreamCS, was conducted on 8 x L40-S GPUs using our
1303 proposed 3D-MeshPref dataset. The reward model was fine-tuned based on MeshMAE checkpoint
1304 for 100 epochs using the AdamW optimizer with a learning rate of 1e-3 in total around 90 GPU
1305 hours. **The CS divergence is computed at the batch level, and in our setup we fix the total batch**
1306 **size to $m + n = 256$.** The Geometry-Asset Alignment Reward (GA) metric, derived from Re-
1307 wardCS, which is trained on approximately 10,000 annotated meshes from the Objaverse-XL subset
1308 of Cap3D (Luo et al., 2023). Since GA provides a model-agnostic and architecture-agnostic mea-
1309 sure of geometric-semantic consistency, it offers a fair basis for comparing DreamCS with other
1310 text-to-3D generation approaches. This metric functions as an independent evaluator for assessing
1311 the alignment between 3D mesh geometry and textual descriptions.

1312 **F.2 ANNOTATION PIPELINES**
1313

1314 **Llama-Mesh Annotation.** To annotate the quality of the converted 3D meshes in our dataset, we
1315 employ a structured three-criterion annotation rubric. Each criterion is rated on a Likert scale from
0.00 to 5.00, where 5.00 denotes the highest level of performance. The final score is calculated as
1316 the average of the three individual scores:

- 1317 • **Alignment with the provided description:** Measure the degree to which the 3D mesh accu-
1318 rately reflects the corresponding textual description. Deductions are applied for any semantic
1319 misalignment, omission of key elements, or inclusion of irrelevant features.
- 1320 • **Plausibility and realism of the 3D structure:** Assess the physical plausibility and visual
1321 realism of the 3D mesh. Artifacts that violate basic physical principles or render the object
1322 visually implausible result in lower scores.
- 1323 • **Completeness of geometry:** Evaluate the geometric and textural completeness of the 3D
1324 mesh. Penalties are applied for missing components, structural incompleteness, or ab-
1325 sent/erroneous textures that compromise model integrity.

1327 To control the potential biases in the automated assessments, rather than relying on a single anno-
1328 tation process, we score every mesh with three independent inference runs of Llama-Mesh and use
1329 the averaged score.

1331 **Human Refinement.** Meshes with a Llama-Mesh score greater than 2.0 are subsequently re-
1332 evaluated by human annotators using the interactive annotation panel shown in Figure 12. The
1333 same three-criterion rubric is applied, ensuring consistency with the automated Llama-Mesh scor-
1334 ing while allowing humans to correct potential errors or biases. Human raters assign scores on the
1335 same 0.0–5.0 Likert scale, and the final score is computed as the average across the three criteria,
1336 providing a refined assessment of each mesh’s quality.

- 1337 • **Alignment with the provided description:** Measure the degree to which the 3D mesh accu-
1338 rately reflects the corresponding textual description. Deductions are applied for any semantic
1339 misalignment, omission of key elements, or inclusion of irrelevant features.
- 1340 • **Plausibility and realism of the 3D structure:** Assess the physical plausibility and visual
1341 realism of the 3D mesh. Artifacts that violate basic physical principles or render the object
1342 visually implausible result in lower scores.
- 1343 • **Completeness of geometry:** Evaluate the geometric and textural completeness of the 3D
1344 mesh. Penalties are applied for missing components, structural incompleteness, or ab-
1345 sent/erroneous textures that compromise model integrity.

1346 **F.3 EVALUATION PIPELINES**
1347

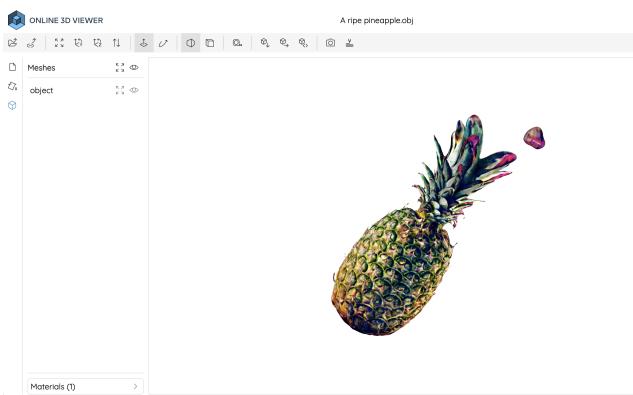
1348 **MiniCPM-o Evaluation.** To evaluate the quality of the generated 3D assets across MVDream-
1349 based pipelines, we use MiniCPM-o based on a three-part evaluation rubric. Each criterion is rated

1350 on a Likert scale from 0.00 to 5.00, where 5.00 denotes the highest performance. The final score is
 1351 calculated as the average of the three individual scores:
 1352

- 1353 • **Text Prompt and Asset Alignment:** Assess the semantic fidelity between the multi-view
 1354 renderings and the original textual prompt. The images should depict all referenced entities
 1355 and contextual elements with accurate visual attributes. Annotators are instructed to first
 1356 describe the output of each model and then evaluate how comprehensively and accurately
 1357 the visual content reflects the prompt.
- 1358 • **3D Plausibility:** Evaluate the inferred 3D structure by examining the set of RGB and normal
 1359 images. Reviewers consider the overall realism, solidity, and structural coherence implied by
 1360 the views. Deductions are made for implausible geometry, unnatural proportions, duplicated
 1361 or floating parts, and any visual artifacts that suggest physical inconsistency. High-quality
 1362 outputs should convey convincing, well-formed 3D shapes from all angles.
- 1363 • **Geometry-Texture Alignment:** Measure the local consistency between geometry and tex-
 1364 ture across views. Specific attention is paid to whether visual features, such as edges, pat-
 1365 terns, or object parts, are coherently represented in both RGB and normal maps. Discrepan-
 1366 cies between texture and underlying geometry reduce the score.

1367
 1368 **User Study.** To assess the perceptual quality of the generated 3D meshes, we conduct a user study
 1369 with 30 participants on 30 prompts based on three core criteria across MVDream-based pipelines,
 1370 as shown in Figures 12 and 13. To avoid bias, models were anonymized, and the response order
 1371 was randomized. Participants are asked to inspect the 3D meshes from multiple viewpoints and
 1372 assign a score on a Likert scale from 0.00 to 5.00 for each criterion, where 5.00 denotes the highest
 1373 performance. The final score is computed as the average of the three individual ratings:
 1374

- 1375 • **Text-Asset Alignment:** Measure how accurately the 3D mesh represents the content de-
 1376 scribed in the original text prompt. Participants assess whether all described objects, at-
 1377 tributes, and spatial arrangements are present and correctly rendered in the mesh. Partial or
 1378 incorrect representations lead to a lower score.
- 1379 • **3D Plausibility:** Evaluate the realism and structural integrity of the mesh geometry. Partici-
 1380 pants judge whether the mesh appears physically plausible and free of distortions, unrealis-
 1381 tic proportions, or non-functional geometry. 3D meshes with topological errors or unnatural
 1382 forms are penalized.
- 1383 • **Geometry-Texture Alignment:** Assess the coherence between the mesh geometry and its
 1384 texture. Annotators examine whether surface textures appropriately follow the underlying
 1385 shapes and contours of the geometry. Inconsistencies, such as textures that misalign with
 1386 surface features or obscure geometric details, result in a lower score.



1390
 1391 Figure 12: Interface built on Online 3D Viewer used in the user study for evaluating 3D meshes.
 1392 Participants could interact with the mesh (e.g., rotate, zoom).
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

1422

1423

1424

1425

1426

1427

1428

1429

1430

1431

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

User Study

You are presented with a set of 3D meshes generated in response to textual prompts. For each criterion, please assign a score between 0.00-5.00 based on the degree to which the mesh meets the stated objectives.

* Indicates required question

Text-to-Asset Alignment *

Evaluate the fidelity of the 3D mesh to the original text description.

Your answer

3D Plausibility *

Assess the physical realism and structural integrity of the mesh.

Your answer

Geometry-to-Texture Consistency *

Judge the alignment and coherence between the mesh geometry and its texture.

Your answer

Figure 13: User study questionnaire built on Google Form for evaluating 3D mesh quality. Participants can evaluate each mesh based on (1) Text-to-Asset Alignment, (2) 3D Plausibility, and (3) Geometry-to-Texture Consistency.

1458 **G EXTENDED EMPIRICAL RESULTS**

1460 **G.1 COMPARISON WITH 3D CONTROLLABLE GENERATION METHODS**

1462 3D controllable generation methods such as MVControl (Li et al., 2023), MT3D (Nath et al., 2025),
 1463 and DreamControl (Huang et al., 2024) can provide guidance for optimization-based 3D generation,
 1464 but they rely on strong 3D shape prior or external conditioning 2D inputs and are architecture-
 1465 dependent, making them difficult to apply in real-world text-to-3D settings. Among them, MVCon-
 1466 trol is the most compatible with MVDream framework, so we adopt it for a comparison.

1467 Under the MVDream backbone, we adopt the pretrained MVControl checkpoint. Given 30 evalua-
 1468 tion prompts from GPTEval3D, we first generate 2D image prior using stable diffusion 3 and then
 1469 use the MVcontrol method to generate the 3D assets. We report performance using the same met-
 1470 rrics in Tables 1, 2, and 3 of our main paper, where T-A, 3DP, and G-T are derived from MLLM
 1471 evaluation.

1472 **Table 7: Comparison with 3D controllable generation method MVControl.**

Model	CP (↑)	VR (↑)	GA (↑)	T-A (↑)	3DP (↑)	G-T (↑)	Proportion (↓)
MVDream	0.23	-3.30	2.77	2.95	3.10	3.07	0.55
MVControl	0.27	-3.09	2.86	3.42	3.47	3.51	0.40
MVDream+Reward3D	0.27	-3.13	2.85	3.38	3.34	3.20	0.46
MVDream+RewardCS	0.29	-2.15	2.96	3.58	4.04	3.93	0.31

1481 As shown in Table 7, RewardCS outperforms both MVControl and Reward3D across all metrics,
 1482 including CP, VR, GA, and especially in 3D plausibility and geometry–texture consistency. More-
 1483 over, RewardCS achieves the lowest Janus proportion (0.31), confirming that it effectively mitigates
 1484 the Janus issue—without requiring any auxiliary conditioning input.

1486 **G.2 COMPARISON WITH ADVANCED TEXT-TO-3D GENERATION PIPELINE**

1488 **Advanced mesh-based text-to-3D generation pipeline.** We extend our experiments to advanced
 1489 multi-stage SDS-based pipelines: ProlificDreamer (Wang et al., 2023b) and DreamCraft3D (Sun
 1490 et al., 2023) under 30 evaluation prompts from GPTEval3D. Following their default configurations,
 1491 for ProlificDreamer, we adopt 20000, 10000, and 10000 optimization steps for the object, geometry,
 1492 and texture stages, respectively. For DreamCraft3D, we use 5000 optimization steps for each of the
 1493 NeRF, geometry, and texture stages. In addition, we use Stable Diffusion 3 to generate the 2D image
 1494 prior. We adopt Zero123-XL as the backbone and use Omnidata for depth and normal estimation.
 1495 For both methods, the rendering resolution is set to 128x128 in the first stage and 512x512 in the
 1496 subsequent stages. We report performance using the same metrics in Tables 1, 2, and 3 of our main
 1497 paper, where T-A, 3DP, and G-T are derived from MLLM evaluation.

1498 **Table 8: Comparison with multi-stage SDS-based text-to-3D generation pipelines.**

Model	CP (↑)	VR (↑)	GA (↑)	T-A (↑)	3DP (↑)	G-T (↑)	Proportion (↓)
ProlificDreamer	0.26	-3.25	2.83	3.03	3.24	3.13	0.43
+Reward3D	0.28	-3.18	2.89	3.41	3.48	3.39	0.37
+RewardCS	0.29	-2.11	3.04	3.58	4.06	4.01	0.29
DreamCraft3D	0.28	-3.20	2.84	3.11	3.26	3.16	0.41
+Reward3D	0.29	-2.96	2.90	3.43	3.51	3.45	0.35
+RewardCS	0.30	-2.08	3.13	3.61	4.11	4.07	0.24

1508 As shown in Table 8, we observe that RewardCS consistently outperforms both the vanilla baselines
 1509 and the 2D reward variant. In multi-stage pipelines, 2D reward signals are often unstable and require
 1510 sensitive hyperparameter tuning; when optimization collapses in an early stage, the degradation
 1511 propagates to later stages. This underscores the robustness of 3D reward guidance.

1512 Additionally, RewardCS yields strong improvements on DreamCraft3D, which also adopts DMTet
 1513 as its underlying 3D representation. This structural compatibility allows RewardCS to provide more
 1514 accurate and geometry-aware feedback on-the-fly.

1515 **Advanced mesh-based text-to-3D generation pipeline.** To further demonstrate its generality, we
 1516 apply RewardCS to finetune the advanced end-to-end text-to-3D generation model TRELLIS-text-
 1517 base (Xiang et al., 2025). We construct a finetuning dataset by randomly selecting 100K prompts
 1518 from the TRELLIS-500K training set. For each prompt, a 3D asset is generated and differentially
 1519 rendered to DMTet or 2D multi-view images. We then use either Reward3D or RewardCS as the
 1520 reward guidance to finetune the TRELLIS backbone, following an objective similar to Equation 4.
 1521

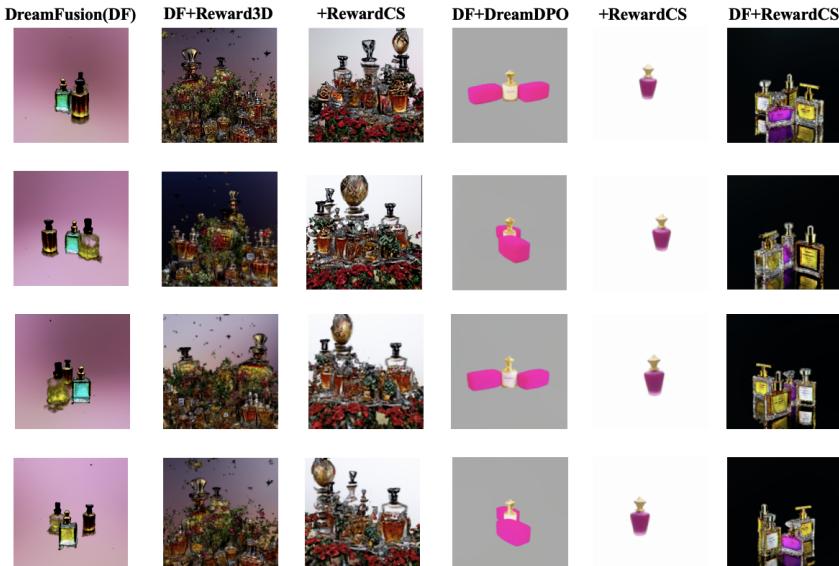
1522 Given 60 evaluation prompts from GPTEval3D, we report performance using the same metrics in
 1523 Tables 1, 2, and 3 of our main paper, where T-A, 3DP, and G-T are derived from MLLM evaluation.
 1524

1525 **Table 9: Comparison of TRELLIS with reward-based enhancements.**

Model	CP (\uparrow)	VR (\uparrow)	GA (\uparrow)	T-A (\uparrow)	3DP (\uparrow)	G-T (\uparrow)	Proportion (\downarrow)
TRELLIS	0.29	-1.53	2.86	3.13	3.30	3.51	0.43
+Reward3D	0.28	-1.47	2.90	3.20	3.35	3.58	0.39
+RewardCS	0.30	-1.39	3.02	3.25	3.43	3.60	0.35

1534 As shown in Table 9, these results demonstrate that RewardCS consistently improves text-asset
 1535 alignment, 3D plausibility, geometry-texture consistency, and reduces Janus artifacts, outperforming
 1536 both the base TRELLIS model and the 2D reward variant. This highlights RewardCS’s flexibility
 1537 and effectiveness as a model-agnostic reward signal for advanced end-to-end 3D generation frame-
 1538 works.

G.3 QUALITATIVE VISUALIZATION



An assortment of vintage, fragrant perfume on display

1563 **Figure 14: Qualitative comparisons with 1-stage generation pipelines: DreamFusion.**

1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583

DreamFusion(DF)	DF+Reward3D	+RewardCS	DF+DreamDPO	+RewardCS	DF+RewardCS
					
					
					
					

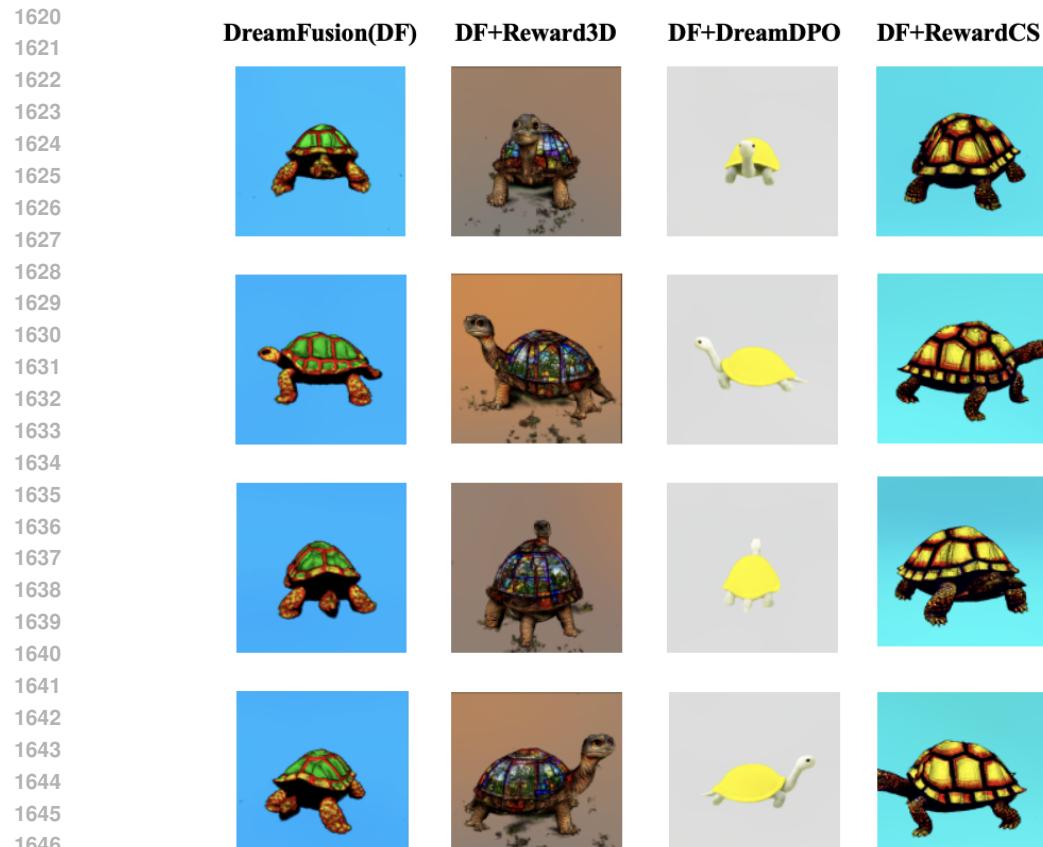
An old-fashioned rotary phone with a tangled cord

1586 Figure 15: Qualitative comparisons with 1-stage generation pipelines: DreamFusion.

1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599

Orange monarch butterfly resting on a dandelion

Figure 16: Qualitative comparisons with 1-stage generation pipelines: DreamFusion.



Tortoise with a shell that looks like pentagon glass

Figure 17: Qualitative comparisons with 1-stage generation pipelines: DreamFusion.

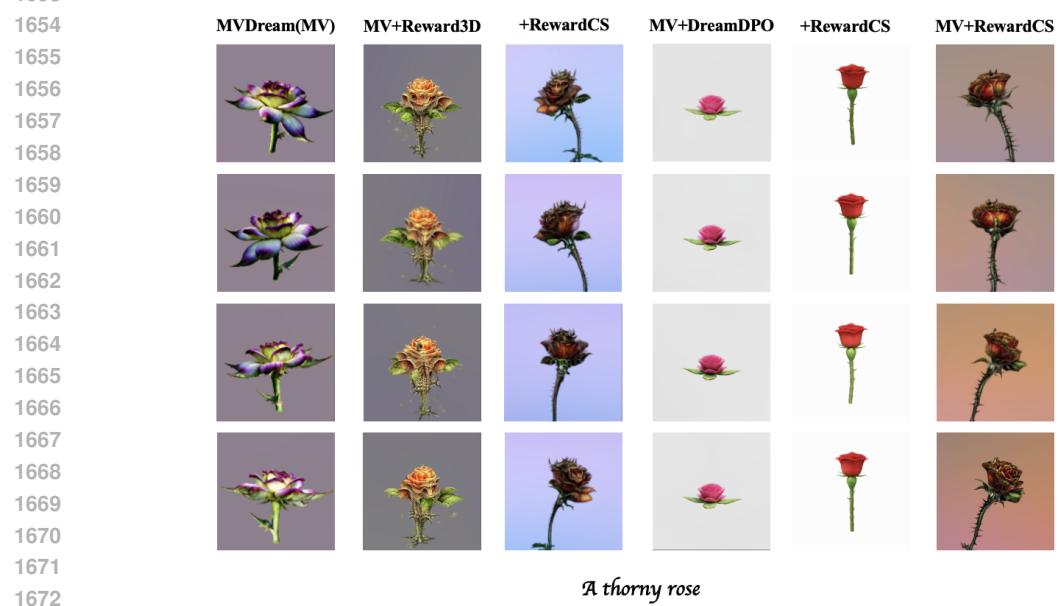


Figure 18: Qualitative comparisons with 1-stage generation pipelines: MVDream.

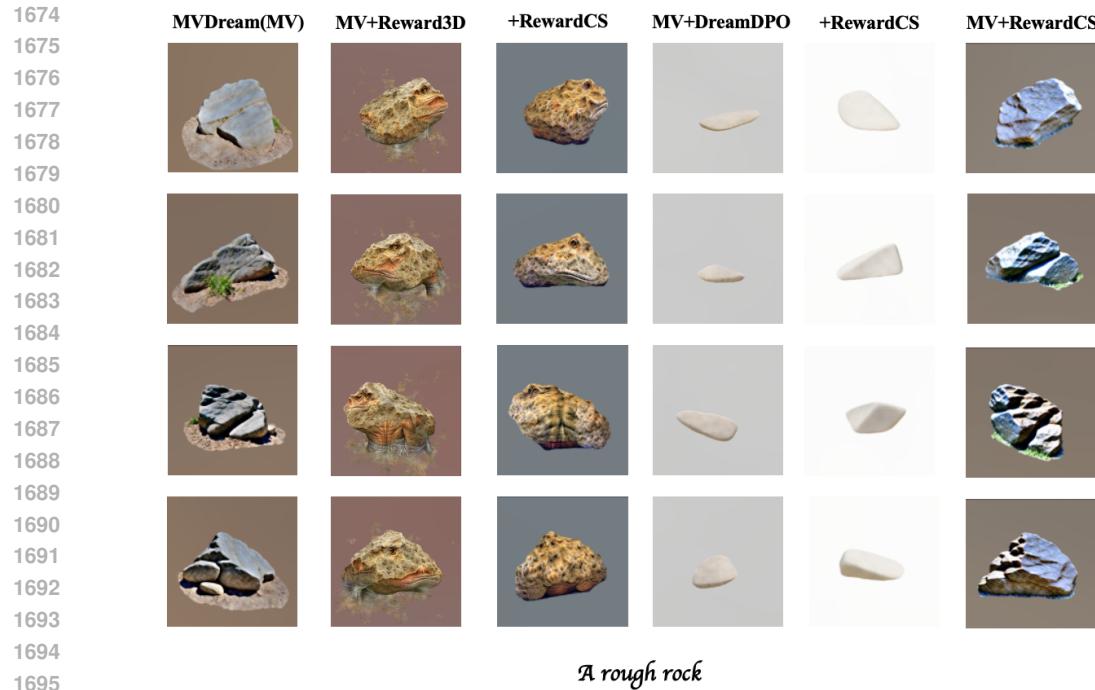
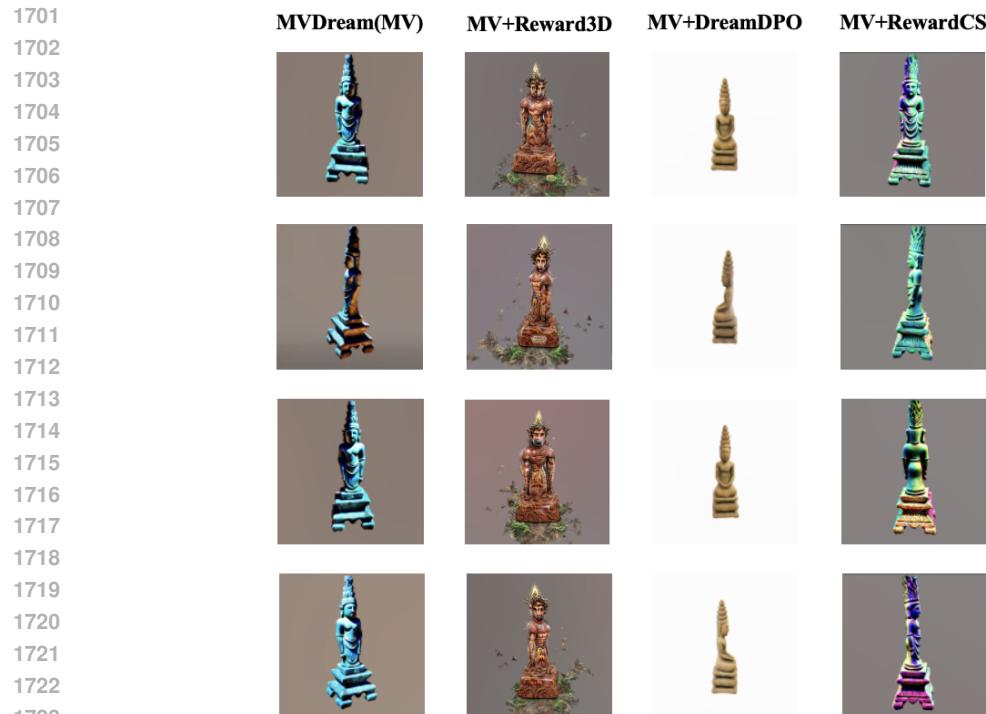
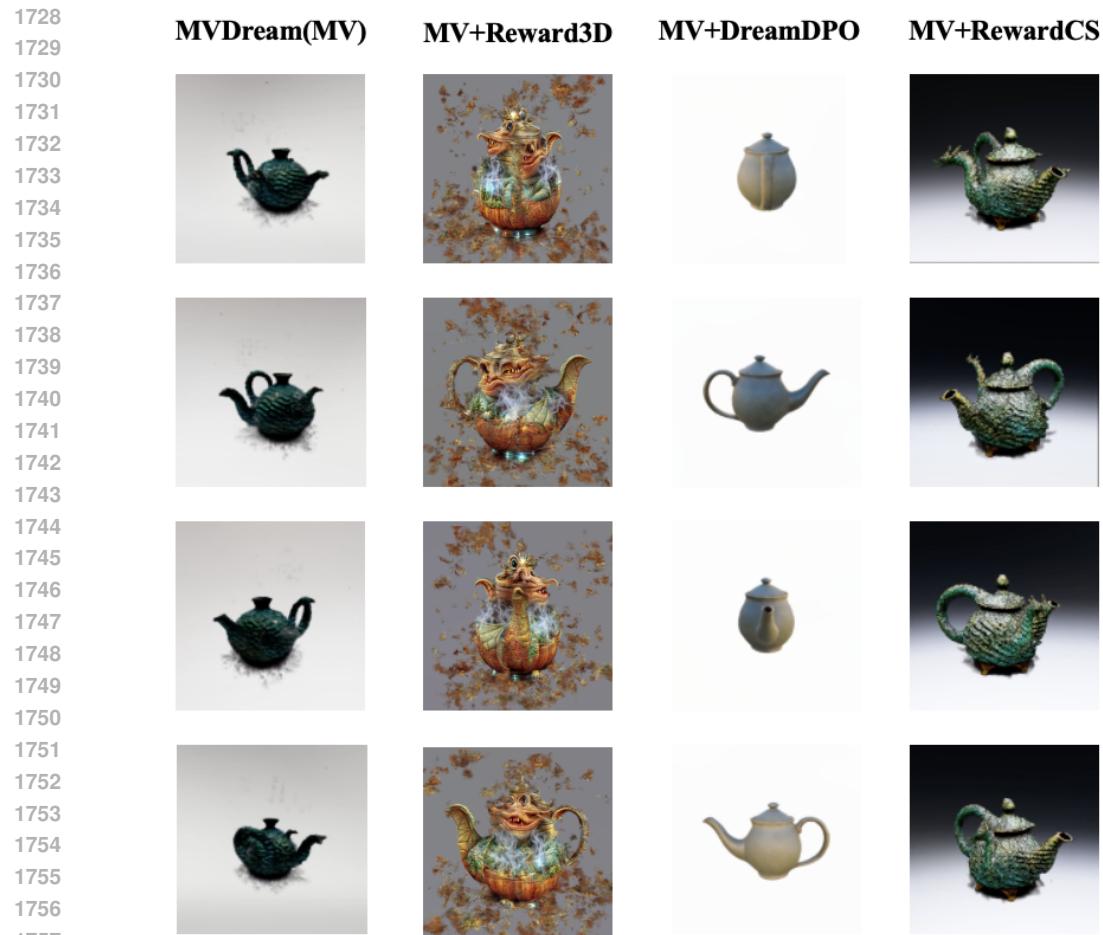


Figure 19: Qualitative comparisons with 1-stage generation pipelines: MVDream.



An intricately carved, solid, bronze figurine, with jagged contours depicting an ancient deity, the grain visible under a matte finish, on a stone altar

Figure 20: Qualitative comparisons with 1-stage generation pipelines: MVDream.



A medium-sized, hollow, asymmetrical teapot, crafted to look like a slumbering dragon, with a scaly, rough texture

Figure 21: Qualitative comparisons with 1-stage generation pipelines: MVDream.

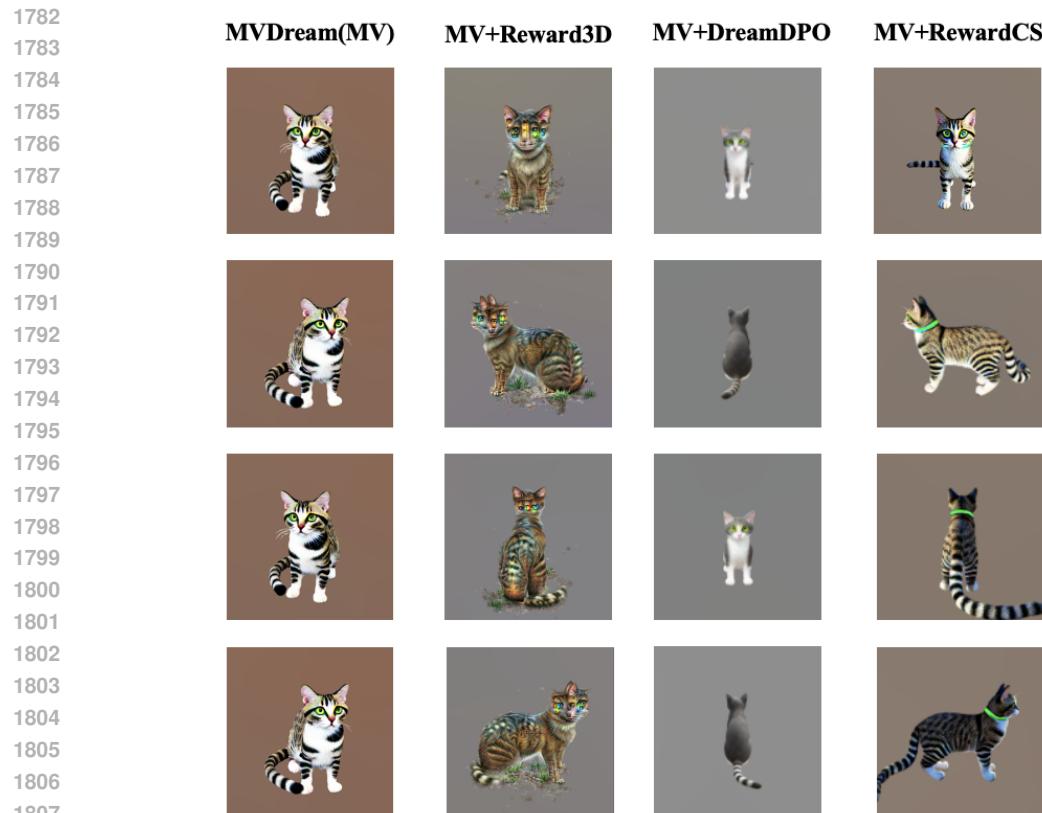
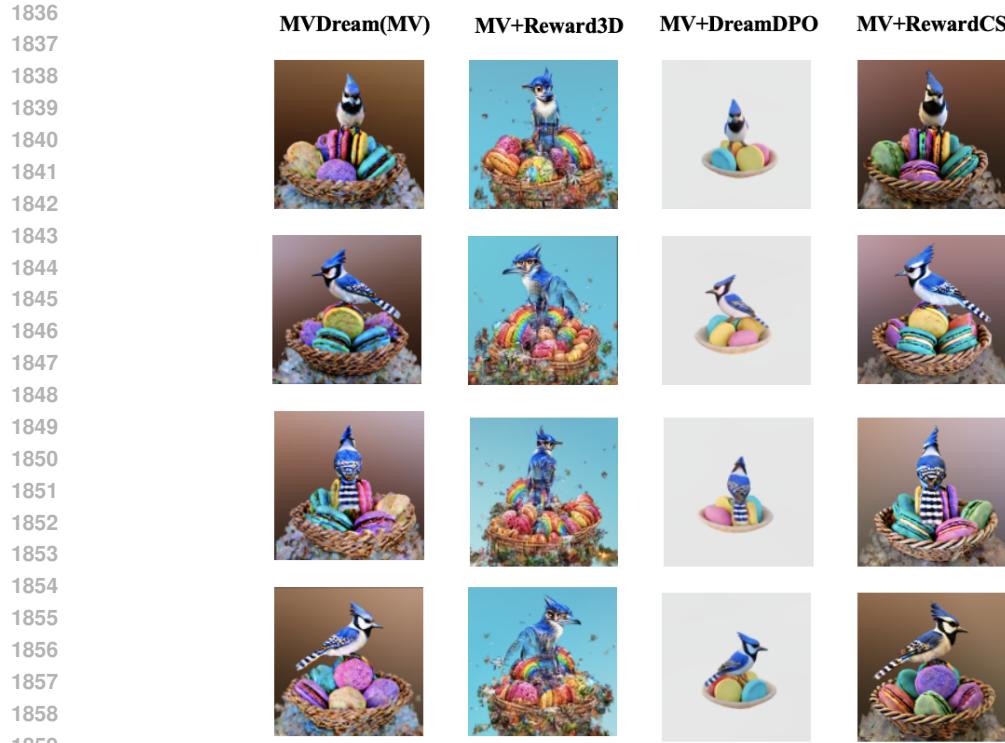
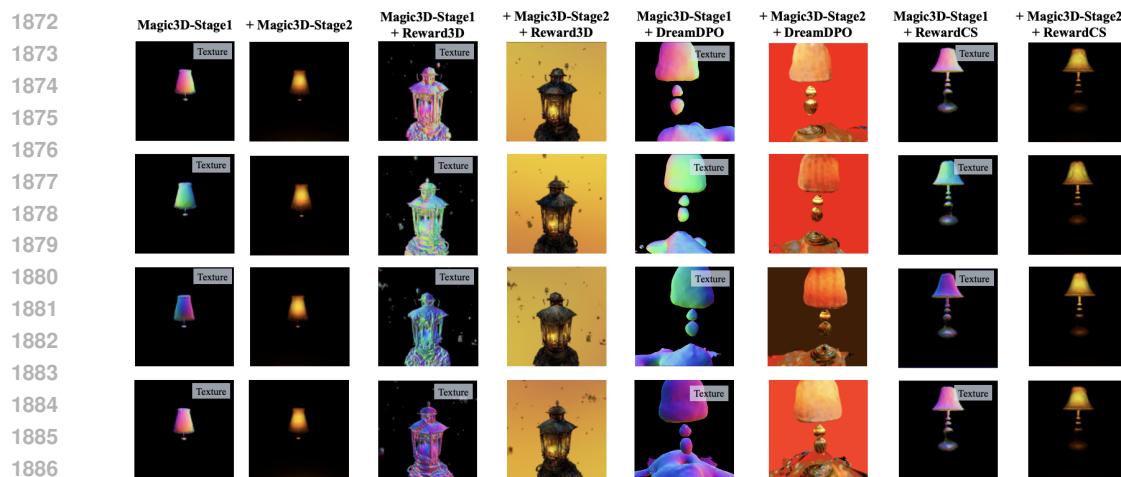
*A cat in a sitting pose with two different colored eyes*

Figure 22: Qualitative comparisons with 1-stage generation pipelines: MVDream.



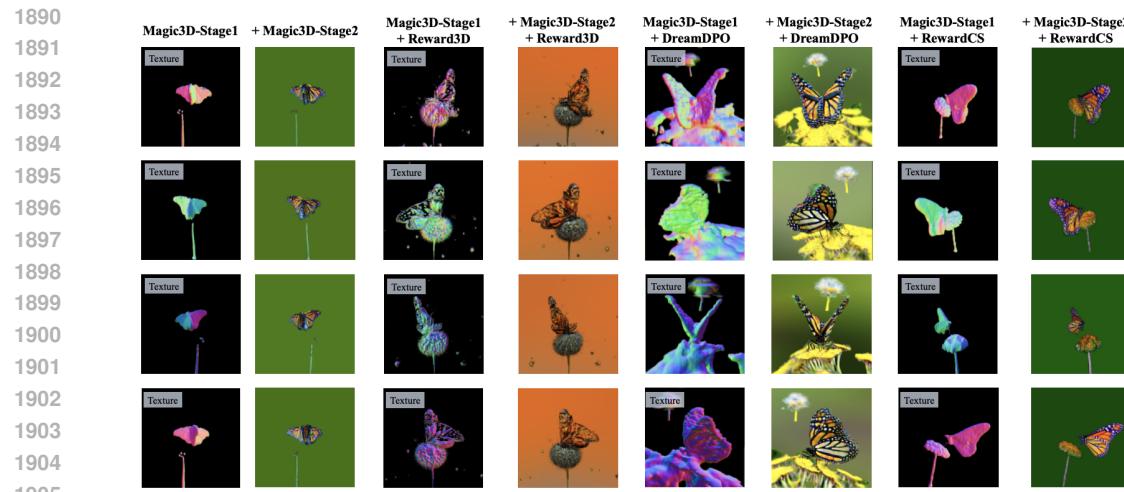
1860 *A DSLR photo of a blue jay standing on a large basket of rainbow macarons*
 1861

1862 Figure 23: Qualitative comparisons with 1-stage generation pipelines: MVDream.
 1863



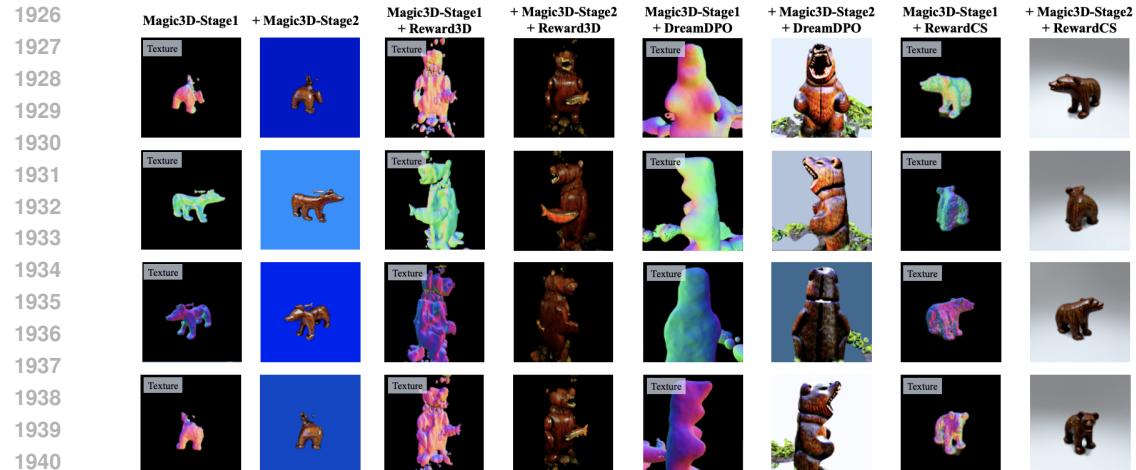
1880 *A vintage table lamp casting a soft, warm glow in a quiet, dark room*
 1881

1882 Figure 24: Qualitative comparisons with 2-stage generation pipelines: Magic3D.
 1883



Orange monarch butterfly resting on a dandelion

Figure 25: Qualitative comparisons with 2-stage generation pipelines: Magic3D.



A carved wooden bear in a preying pose

Figure 26: Qualitative comparisons with 2-stage generation pipelines: Magic3D.

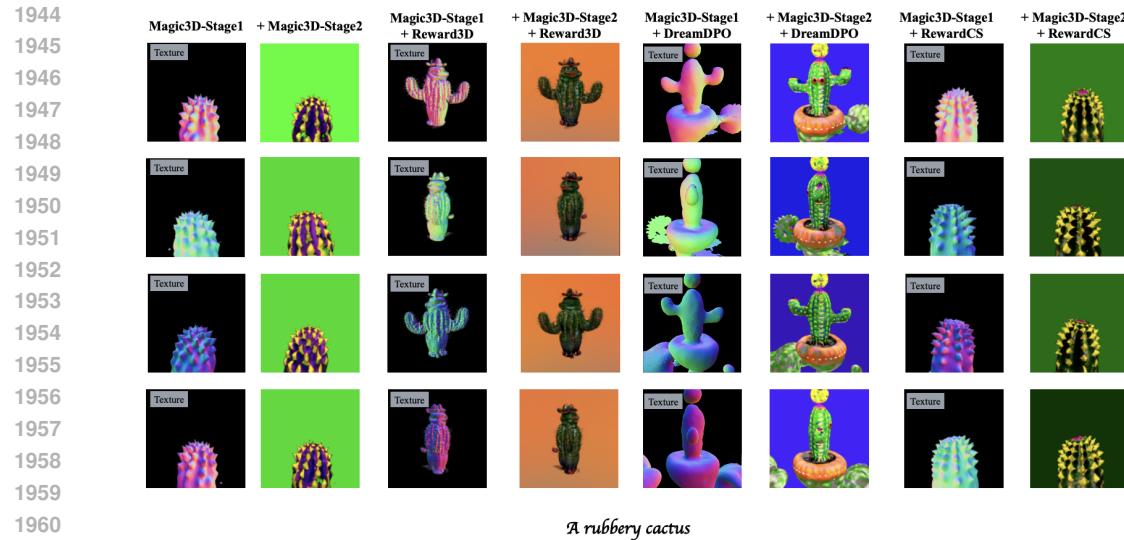


Figure 27: Qualitative comparisons with 2-stage generation pipelines: Magic3D.

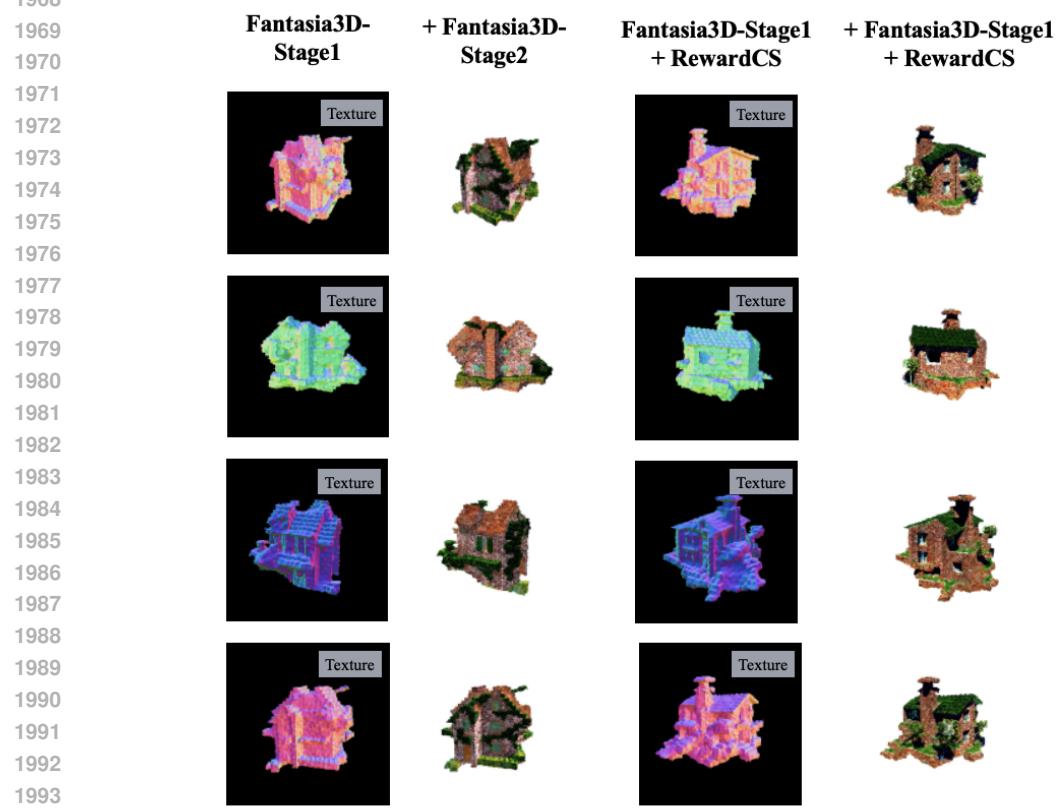


Figure 28: Qualitative comparisons with 2-stage generation pipelines: Fantasia3D.

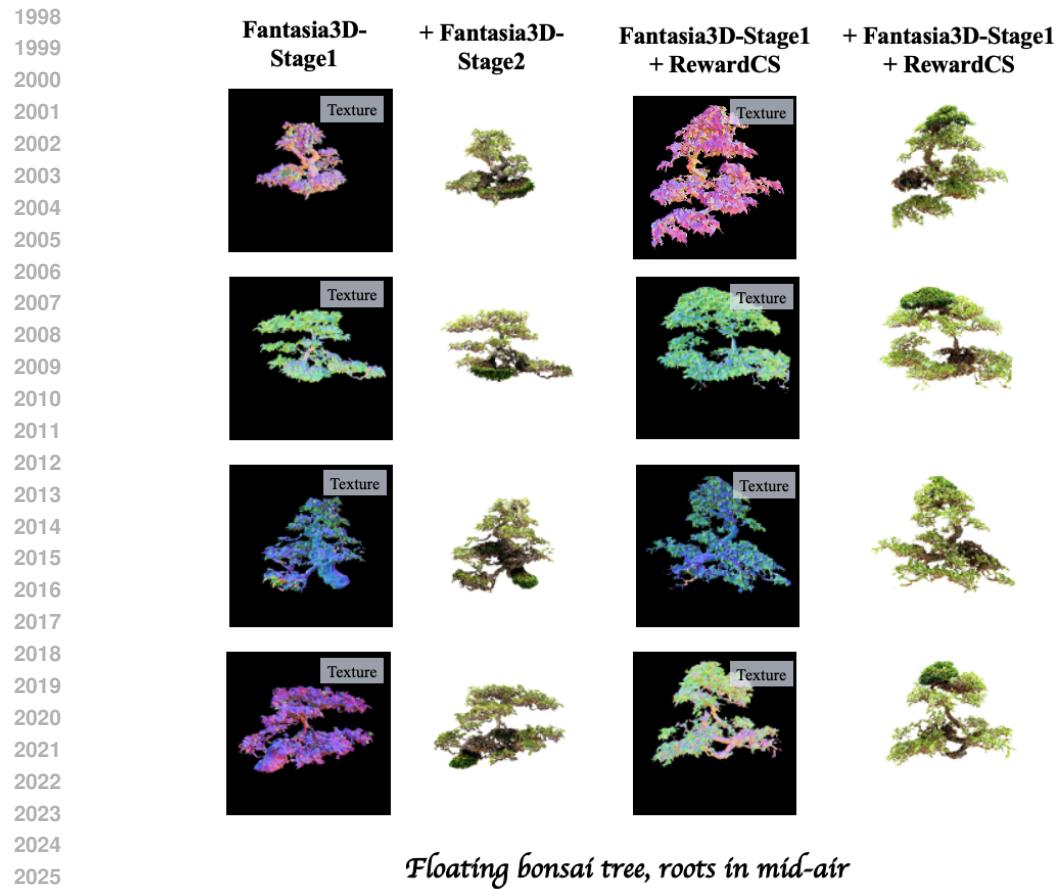
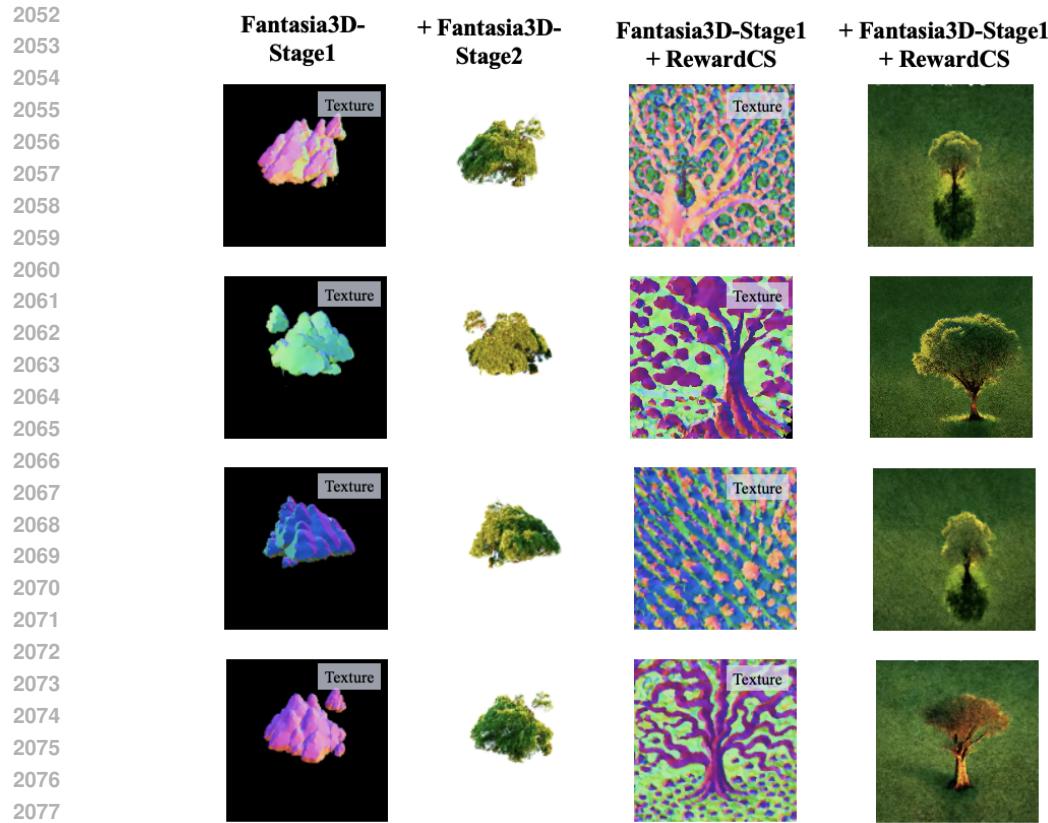
*Floating bonsai tree, roots in mid-air*

Figure 29: Qualitative comparisons with 2-stage generation pipelines: Fantasia3D.



A lone, ancient tree stands tall in the middle of a quiet field

Figure 30: Qualitative comparisons with 2-stage generation pipelines: Fantasia3D.

2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105

2160

2161

2162

2163

2164

2165

2166

2167

2168

2169

2170

2171

2172

2173

2174

2175

2176

2177

2178

2179

2180

2181

2182

2183

2184

2185

2186

2187

2188

2189

2190

2191

2192

2193

2194

2195

2196

2197

2198

2199

2200

2201

2202

2203

2204

2205

2206

2207

2208

2209

2210

2211

2212

2213

Trellis**Trellis+Reward3D****Trellis+RewardCS**

An ancient, twisted, and unsymmetric tower.

Figure 32: Qualitative comparisons with the end-to-end generation pipeline: Trellis.

2214

2215

2216

2217

2218

2219

2220

2221

2222

2223

2224

2225

2226

2227

2228

2229

2230

2231

2232

2233

2234

2235

2236

2237

2238

2239

2240

2241

2242

2243

2244

2245

2246

2247

2248

2249

2250

2251

2252

2253

2254

2255

2256

2257

2258

2259

2260

2261

2262

2263

2264

2265

2266

2267

Trellis



Trellis+Reward3D



Trellis+RewardCS



Floating bonsai tree roots in mid-air.

Figure 33: Qualitative comparisons with the end-to-end generation pipeline: Trellis.

2268

2269

2270

2271

2272

2273

2274

2275

2276

2277

2278

2279

2280

2281

2282

2283

2284

2285

2286

2287

2288

2289

2290

2291

2292

2293

2294

2295

2296

2297

2298

2299

2300

2301

2302

2303

2304

2305

2306

2307

2308

2309

2310

2311

2312

2313

2314

2315

2316

2317

2318

2319

2320

2321

Trellis**Trellis+Reward3D****Trellis+RewardCS***Four colorful and vibrant balloons tied together.*

Figure 34: Qualitative comparisons with the end-to-end generation pipeline: Trellis.

2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345

A collection of solid, irregularly shaped hand tools, with wooden handles and metal ends, well-used and slightly rusty, hanging on a pegboard in a workshop.

Figure 35: Qualitative comparisons with the end-to-end generation pipeline: Trellis.

These images present multi-view visualizations of text-to-3D results generated by one-stage and two-stage generation pipelines, respectively, across diverse text prompts. We find that text-to-3D vanilla baselines and variants enhanced with 2D guidance methods often produce 3D assets with geometric defects and the Janus problem, while 3D generation backbone with our RewardCS yields geometrically consistent 3D content. In conclusion, DreamCS outperforms text-to-3D vanilla baselines and variants enhanced with 2D guidance methods, **3D controllable generation methods, and advanced text-to-3D generative models** by achieving superior geometric consistency and texture fidelity in 3D generation.

G.4 ADDITIONAL COMPUTATIONAL COST FROM MESHIZATION GUIDANCE.

Regarding the inference cost, we would like to clarify that the computational overhead occurs only in the implicit NeRF-based SDS setting, since our 3D Reward Guidance framework employs DMTet for differentiable mesh extraction during optimization. For advanced mesh based text-to-3D generative pipeline and end-to-end 3D generative backbone, the computational cost incurred by RewardCS is negligible. We include these implicit NeRF-based SDS experiments to demonstrate that DreamCS is a model-agnostic guidance module compatible with general text-to-3D generative backbone.

Implicit NeRF-based SDS-based Generative Pipeline. While this introduces some overhead, we emphasize the following:

a) Significant Quality Gains Justify the Cost. Despite the added time, meshization-guided reward feedback significantly improves multi-view consistency, reduces geometric artifacts (e.g., floaters, Janus effects), and enhances 3D structural plausibility. These quality improvements are especially crucial in downstream applications like simulation, animation, and fabrication, where geometric plausibility is often more valuable than marginal runtime gains.

b) Quality-Efficiency Trade-off via Resolution Control. The meshization process is tunable via the isosurface resolution parameter. In our main experiments, we use a resolution of 128, which we find offers a good balance between efficiency and surface detail. Users can opt for lower resolutions

2376 during fast prototyping or higher resolutions for final asset generation, providing practical flexibility
 2377 in real-world workflows.

2378 **c) Future Modular Design.** Our framework is modular and fully compatible with advances in mesh
 2379 extraction techniques. Faster or more efficient meshization algorithms (e.g., learned meshing) can be
 2380 easily integrated into our pipeline to further reduce runtime without compromising guidance quality.
 2381

2382 We conducted the ablation study using MVDream and its RewardCS variants at three different iso-
 2383 surface resolution: 64, 128, and 256. Baseline refers to vanilla MVDream without meshization. The
 2384 evaluation was performed across 15 diverse prompts from the GPTEval3D benchmark, with each
 2385 3D asset generated through optimization over 20000 iterations (it) on a single L40S GPU.

2386 **Trade-off between Generation Speed and 3D Asset Quality.**

2388 Table 10: Performance metrics for different isosurface resolutions.

Resolution	Generation Speed (\uparrow)	IR (\uparrow)	VR (\uparrow)	GA (\uparrow)	T-A (\uparrow)	3DP (\uparrow)	G-T (\uparrow)
Baseline	0.64 it/s	-0.55	-3.46	2.72	2.97	3.03	2.99
64	0.59 it/s	-0.49	-2.96	2.81	3.04	3.37	3.32
128	0.50 it/s	-0.45	-2.11	3.02	3.23	4.03	4.01
256	0.38 it/s	-0.43	-1.84	3.18	3.34	4.10	4.09

2396 Table 11: Computational cost metrics for different isosurface resolutions.

Method	Wall-clock Time (\downarrow)	VRAM (\downarrow)
MVDream	8.5 hours	21872 MiB
+DreamDPO	8.9 hours	25787 MiB
+Reward3D	8.8 hours	26385 MiB
+RewardCS	11.0 hours	23784 MiB

2405 The results in Tab. 10 and 11 illustrate the trade-off between generation speed, computational cost
 2406 and generated 3D asset quality. We report performance using the same metrics in the main paper,
 2407 where T-A, 3DP, and G-T are derived from MLLM evaluation. We outline two key observations
 2408 from these results:

2409 i) Consistent improvement in mesh quality is observed across all metrics as the isosurface resolution
 2410 increases.
 2411
 2412 ii) A diminishing return in quality improvement appears at higher isosurface resolution. This indicates
 2413 that our 128-resolution setting achieves a balance between performance and efficiency.

2414 **Reducing Overhead via 3D Guidance Frequency.** To reduce computational overhead, we con-
 2415 ducted additional experiments under the 128-resolution setting, where we varied the RewardCS
 2416 guidance frequency during 3D asset generation. In particular, we reduced the guidance frequency
 2417 from the default of 1it^{-1} (i.e., applying 3D guidance at every step). We evaluated two configura-
 2418 tions with 3D guidance frequency of $\frac{1}{2}\text{it}^{-1}$ and $\frac{1}{3}\text{it}^{-1}$, corresponding to the guidance periods of 2
 2419 iterations and 3 iterations, respectively.

2420 Table 12: Performance metrics for different 3D guidance frequencies.

Frequency	Generation Speed (\uparrow)	IR (\uparrow)	VR (\uparrow)	GA (\uparrow)	T-A (\uparrow)	3DP (\uparrow)	G-T (\uparrow)
$\frac{1}{2}\text{it}^{-1}$	0.54 it/s	-0.48	-2.46	2.98	3.14	3.64	3.59
$\frac{1}{3}\text{it}^{-1}$	0.56 it/s	-0.51	-2.53	2.96	3.12	3.59	3.57

2427 As shown in Tab. 12, despite the reduced guidance frequency, both settings still outperform the
 2428 MVDream baseline across all quality metrics, while also achieving faster generation speed than the
 2429 standard RewardCS guidance configuration

2430
 2431 **Advanced Text-to-3D Generative Pipeline.** For advanced text-to-3D pipelines, where methods
 2432 optimize an explicit mesh rather than an implicit NeRF field, RewardCS becomes lightweight. When
 2433 integrated into mesh-based SDS systems such as Magic3D and DreamCraft3D, RewardCS operates
 2434 directly on explicit geometry and adds < 1% inference-time cost, making it negligible. In addition,
 2435 RewardCS can serve as a 3D prior to finetune current end-to-end text-to-3D models such as Trellis,
 2436 which produces a 3D asset in 3 minutes. As shown in Table 13, RewardCS improves 3D fidelity
 2437 and geometry metrics without introducing any extra inference cost in guiding DreamCraft3D and
 2438 TRELLIS under the setting detailed in Appendix G.2, which are better than MVdream.
 2439

Table 13: Comparison of different text-to-3D generative backbones.

Model	CP (↑)	VR (↑)	GA (↑)	T-A (↑)	3DP (↑)	G-T (↑)	Proportion (↓)	Latency (↓)
MVdream	0.23	-3.30	2.77	2.95	3.10	3.07	0.55	8.5 hours
DreamCraft3D	0.28	-3.20	2.84	3.11	3.26	3.16	0.41	5.4 hours
+Reward3D	0.29	-2.96	2.90	3.43	3.51	3.45	0.35	5.5 hours
+RewardCS	0.30	-2.08	3.13	3.61	4.11	4.07	0.24	5.4 hours
TRELLIS	0.29	-1.53	2.86	3.03	3.30	3.51	0.43	2.91 mins
+Reward3D	0.28	-1.47	2.90	3.10	3.35	3.58	0.40	2.91 mins
+RewardCS	0.30	-1.39	3.02	3.15	3.43	3.60	0.35	2.91 mins

2449 These results indicate that advanced mesh-based and end-to-end generative backbone not only out-
 2450 perform implicit SDS-based, such as MVdream, in quality but also achieve significantly lower
 2451 inference latency.
 2452

2453
 2454
 2455
 2456
 2457
 2458
 2459
 2460
 2461
 2462
 2463
 2464
 2465
 2466
 2467
 2468
 2469
 2470
 2471
 2472
 2473
 2474
 2475
 2476
 2477
 2478
 2479
 2480
 2481
 2482
 2483