

VCD-Texture: Variance Alignment based 3D-2D Co-Denoising for Text-Guided Texturing

Shang Liu¹, Chaohui Yu¹, Chenjie Cao¹, Wen Qian¹, and Fan Wang¹

DAMO Academy Alibaba Group, Beijing China

{liushang.ls, huakun.ych, caochenjie.ccj, qianwen.qian,
fan.w}@alibaba-inc.com

Abstract. Recent research on texture synthesis for 3D shapes benefits a lot from dramatically developed 2D text-to-image diffusion models, including inpainting-based and optimization-based approaches. However, these methods ignore the modal gap between the 2D diffusion model and 3D objects, which primarily render 3D objects into 2D images and texture each image separately. In this paper, we revisit the texture synthesis and propose a **V**ariance alignment based 3D-2D **C**ollaborative **D**enoising framework, dubbed **VCD-Texture**, to address these issues. Formally, we first unify both 2D and 3D latent feature learning in diffusion self-attention modules with re-projected 3D attention receptive fields. Subsequently, the denoised multi-view 2D latent features are aggregated into 3D space and then rasterized back to formulate more consistent 2D predictions. However, the rasterization process suffers from an intractable variance bias, which is theoretically addressed by the proposed variance alignment, achieving high-fidelity texture synthesis. Moreover, we present an inpainting refinement to further improve the details with conflicting regions. Notably, there is not a publicly available benchmark to evaluate texture synthesis, which hinders its development. Thus we construct a new evaluation set built upon three open-source 3D datasets and propose to use four metrics to thoroughly validate the texturing performance. Comprehensive experiments demonstrate that VCD-Texture achieves superior performance against other counterparts.

Keywords: 3D Texture Synthesis · Diffusion Model · 3D Self-Attention · Rasterization Variance Alignment

1 Introduction

Textured 3D objects are essential in enhancing the realism and immersive experience in various computer graphics applications, including games, animation, and AR/VR environments [7, 10, 19, 20, 22, 46]. Traditionally, texture synthesis is a time-intensive and challenging task, often demanding specialized knowledge and extensive manual labor. With prominent advancements in diffusion models [32, 38, 39] trained with large-scale datasets [42, 45], some recent researches [3, 5, 8, 37] have pivoted toward the adoption of 2D text-conditioned



Fig. 1: Results of text-guided 3D shape textures generated by VCD-Texture. Our method could achieve high-quality texture synthesis with simple captions.

generative models, such as CLIP [35] and Stable Diffusion (SD) [38], to direct the creation of 3D textures. By leveraging the robust generalization properties inherent to pre-trained diffusion models, these methodologies enable to yield of high-quality textures. For instance, prior optimization-based [27, 28] and inpainting-based approaches [5, 37, 44] primarily render 3D objects into 2D images and texture each image through pre-trained 2D vision-language models. Very recently, driven by the advancement of 2D diffusion models, some literatures [3, 5, 24, 44] further improve the texturing performance by designing texturing schemes in the image pixel domain or latent feature domain.

However, these text-to-image diffusion-enhanced texturing methods merely focus on synchronized multi-view diffusion denoising, neglecting cross-view correspondences in 3D space, as well as the fundamental disparity between the 2D diffusion model and 3D objects. Moreover, we clarify that the feature aggregation-and-rasterization process [3, 24] suffers from a serious variance bias, which degrades the diffusion generation a lot. Thus, in the shape texturing task, there remain challenges in seamlessly embedding 2D text-to-image diffusion priors to strengthen the 3D texture synthesis.

In this paper, we primarily address these issues through 1) *unifying both 2D and 3D feature learning* during the text-to-image denoising process. Furthermore, 2) we *theoretically align the variance bias* that occurred in the aggregation-and-rasterization process for the multi-view feature fusion. Finally, 3) we *refine inconsistent regions with image inpainting*. Therefore, our method enjoys superior 3D consistency, unbiased high-frequency details, and high-fidelity textures as shown in Fig. 1.

Formally, we propose a novel 3D texture synthesis framework, called **Variance alignment based 3D-2D Collaborative Denoising (VCD-Texture)**. Particularly, VCD-Texture could be divided into two main processes: 3D-2D collaborative denoising and inpainting refinement. During the denoising steps of the Stable Diffusion (SD) [38], we modify the self-attention block in U-Net to learn features through both 2D and 3D respective fields with Joint Noise Prediction (JNP). JNP utilizes the rendering-projection relationship to lift 2D foreground latent features into the 3D space, substantially improving the cross-view feature corre-

lations. Moreover, we propose to use Multi-View Aggregation-and-Rasterization (MV-AR) to fuse multi-view latent predictions from SD. Specifically, MV-AR is enhanced by a Variance Alignment (VA) technique to eliminate the domain gap of rasterization, which is utilized to re-render features from the aggregated 3D space to separate 2D views. To further improve the texture consistency, especially for the intrinsic discrepancy in latent feature and image pixel domains, we leverage the inpainting refinement to rectify inconsistent regions detected by aggregated pixels with high variance.

Benefiting from the integration of aforementioned components of VCD-Texture, our model could properly synthesize high-fidelity 3D textures. However, the community still lacks a unified texture synthesis benchmark to evaluate the performance. For quantitative experiments, we sample three sub-sets from publicly available 3D datasets (Objaverse [11], ShapeNetSem [41], and ShapeNet [4]), and leverage four metrics (FID [13], ClipFID [21], ClipScore [35], ClipVar [1]) to thoroughly evaluate texture quality concerning fidelity, the semantic matching score between text and image, and the consistency of views. Qualitative and quantitative comparisons signify that our VCD-Texture is capable of generating 3D textures with high fidelity, exhibiting both global and local coherence. In addition, by harnessing the capabilities of pre-trained diffusion models, our approach, as a training-free method for 3D texture synthesis, demonstrates good time efficiency and remarkable and robust generalization ability in handling diverse 3D objects and complex textual descriptions.

The contributions of this paper are summarized as follows:

1. We propose a self-attention-based JNP block, which incorporates both 2D and 3D features to promote consistency of predicted multi-view noise.
2. We design MV-AR to generate consistent textures, and apply VA to theoretically address the rasterization variance bias problem.
3. We design a texture conflict identification and an inpainting refinement pipeline to mitigate the discrepancy between the feature domain and the pixel domain.

2 Related Works

Text-to-Image Diffusion Models. In recent years, numerous advanced diffusion models [36, 38, 39, 50] have emerged, demonstrating their prowess in crafting high-fidelity images finely tuned by textual prompts. Among them, the immensely acclaimed SD [38] stands out, which is honed on a diverse text-image compendium and intricately interwoven with the fixed text encoder from CLIP [35]. The following works propose to add more conditions to text-to-image generation, including semantic segmentation [38], sketch [50], depth map [50], and other conditions [29, 50], which greatly promote the development and application of text-to-image generation. Driven by the success of text-to-image diffusion models, many works have explored text-conditional diffusion models in other modalities, *e.g.*, text-to-video [43], and text-to-3D [33], and text-guided texturing [3, 5, 44]. In this work, we focus on the field of 3D texture generation.

3D Shape and Texture Generation. Traditional Texture methods leverage rule-based or optimization-based approaches [2, 6, 23, 47] to tile exemplar patterns to 3D assets. However, due to the limited computation and capacity, these methods are hard to synthesize complicated 3D models. In the recent few years, deep learning methods [12, 16] have been widely applied to the synthesis of 3D textures. Gramgan [34] trains deep models to generate textures by non-linearly combining learned noise frequencies. SGAN [15] projects a single noise vector to the whole spatial space via spatial GAN [12] to synthesize texture. Kniaz [18] proposes a novel method for the generation of realistic 3D models with thermal textures using the SfM [17] pipeline and GAN. Texture Fields [30] utilizes a continuous 3D function parameterized with a neural network to achieve high-frequency textures. AUV-Net [9] learns to map 3D texture into 2D plane by embedding 3D surfaces into a 2D aligned UV space.

Recently, with the development of the vision language model, recent texture research has explored techniques to distill language models for 3D texture generation. Early approaches utilize CLIP [26, 27, 35, 40] to synthesize texture by improving the semantic similarity between rendering image and descriptive texture text. Subsequently, given the advancement of text-to-image models, such as Stable Diffusion [38], DALL-E [36], ControlNet [50], which enables the generation of high-fidelity images with given description text. To harness the capabilities of text-to-image models, TEXTure [37], Text2Tex [5] and Repaint3D [44] utilize Depth-aware Stable Diffusion (Depth-SD) to design an inpainting texturing scheme, which gradually paints the texture map of a 3D model from multiple viewpoints. Later, Texfusion [3] converts the progressive inpainting schema to the latent feature domain, which obtains consistent latent images via project-and-inpaint in each denoising time-step. In further, SyncMVD [24] drops the auto-regressive inpainting schema and treats each view equally, which designs an aggregate-and-rendering process to synchronized multi-view latent features in each denoising time-step, and this allows the diffusion processes from different views to reach a consensus of the generated content early in the process and hence ensures the texture consistency. Additionally, some studies, like Paint3D [49] and UV-Diffusion [48], involve training texture models from scratch. However, due to the limitations of 3D texture datasets, these models often exhibit suboptimal generalizability and produce textures that fall short of realism.

3 Methods

3.1 Preliminary

Diffusion Models. The diffusion model [14] comprises a forward process $q(\cdot)$ and a reverse denoising process $p_\theta(\cdot)$. The forward step incrementally corrupts data x_0 with noise $\epsilon \sim \mathcal{N}(0, 1)$ to a noisy sequence x_1, \dots, x_T , following a Markov chain as:

$$q(x_t|x_{t-1}, y) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I), \quad x_t = \sqrt{\beta_t}x_0 + \sqrt{1 - \beta_t} \cdot \epsilon, \quad (1)$$

where β_t is the variance schedule for $t = 1, \dots, T$; y is the condition. The reverse process denoises the pure data from x_T as:

$$p_\theta(x_{t-1}|x_t, y) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t, y), \Sigma_\theta(x_t, t, y)), \quad (2)$$

where μ_θ and Σ_θ indicate mean and variance predictions achieved by the trainable network parameterized with θ .

Stable Diffusion (SD). SD [38] projects all features into the latent space to save the diffusion computation. SD also employs an encoder-decoder-based U-Net for noise prediction. As shown in Fig. 2(b), the multi-scale transformer block comprised in U-Net consists of self-attention and cross-attention modules. Self-attention lets the network prioritize long-range relevant features within a 2D image, while cross-attention aligns the denoising process with textual descriptions for controllable generation.

Mesh Rendering [31]. Given a 3D mesh M with vertices V , faces F , and attributes A , the rendering process converts the mesh into a 2D image by triangle rasterization. Each face $f \in F$ is consisted of three vertices $V_{f_1}, V_{f_2}, V_{f_3} \in V$ with attributes $A_{f_1}, A_{f_2}, A_{f_3}$, where f_1, f_2, f_3 indicate indexes of vertices stored in face. For each pixel value \mathbf{I}_i , the rendering process entails tracing a ray that intersects the mesh at point P_i within triangle face f_k . Attribute of pixel \mathbf{I}_i are interpolated using barycentric coordinates $(\alpha_{i1}, \alpha_{i2}, \alpha_{i3})$, which satisfy:

$$\alpha_{i1} + \alpha_{i2} + \alpha_{i3} = 1, \quad \alpha_{i1}, \alpha_{i2}, \alpha_{i3} \geq 0. \quad (3)$$

Then, the pixel value \mathbf{I}_i is computed from a linear combination:

$$\mathbf{I}_i = \alpha_{i1}A_{f1} + \alpha_{i2}A_{f2} + \alpha_{i3}A_{f3}. \quad (4)$$

Given camera view $C_{n \in N}$ with N views at all, we aim to establish a mutual render-projection relation between 2D image plane \mathbf{I} and 3D mesh M . To achieve this, we utilize the above rendering approach to render M at the same size as image \mathbf{I} . Thereby, we obtain rendering mapping function $R(\cdot)$, which rasterizes mesh vertices to the 2D image plane, and back projection function $R^{-1}(\cdot)$, which projects 2D plane values back to corresponding vertices.

3.2 3D-2D Collaborative Denoising

Given mesh M , text condition y , and camera $C_{n \in N}$ with N views, we first re-mesh the initial mesh into a coarse one as M_c , containing J_c vertices. Next, we render M_c into a latent space, which allows us to align the VCD-Texture learning process to the formulation of the SD U-Net.

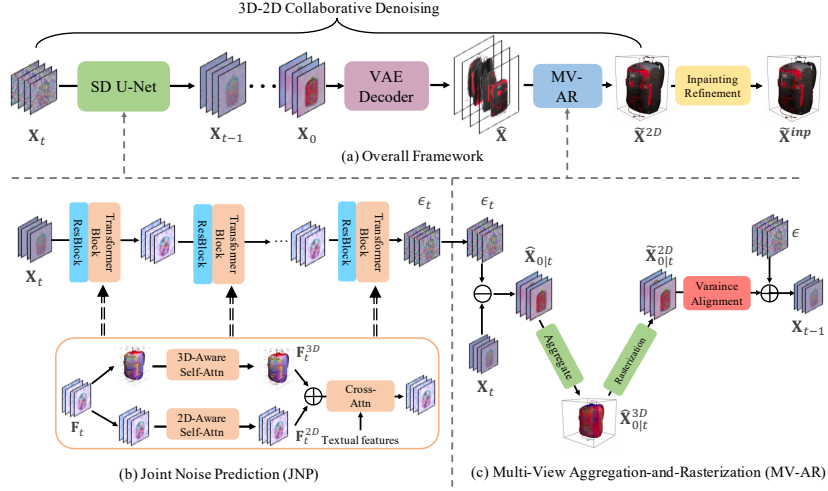


Fig. 2: The framework of VCD-Texture: (a) shows the overall process, including 3D-2D collaborative denoising and inpainting refinement; (b) shows the JNP in SD U-Net; (c) indicates the MV-AR with VA. Note that we only apply the aggregation sub-process of MV-AR to denoised multi-view images $\hat{\mathbf{I}}$ to achieve texture $\hat{\mathbf{I}}^{3D}$.

Joint Noise Prediction (JNP). During the denoising step of time step t in U-Net, we indicate N views' input feature tensor to each transformer block as $\mathbf{F}_t \in \mathbb{R}^{N \times \hat{h} \times \hat{w} \times c}$, where \hat{h}, \hat{w}, c are height, width, and channels of the feature. Specifically, VCD-Texture unifies both 2D and 3D self-attention learning for diffusion denoising to achieve consistent feature presentation indicated as \mathbf{F}_t^{2D} and \mathbf{F}_t^{3D} respectively. Since \mathbf{F}_t^{2D} can be separately learned by the inherent 2D self-attention of SD for each view, we incorporate another 3D self-attention branch as shown in Fig. 2(b) with tailored 3D receptive fields for \mathbf{F}_t^{3D} . Formally, we first apply the inverse projection $R^{-1}(\cdot)$ to obtain the spatial mapping of every foreground feature to the 3D space. Then we split the 3D space into a series of discrete volumetric grids with grid size G_t . In the 3D attention branch, multi-view features \mathbf{F}_t^{3D} can only attend to the ones within the same 3D grid, but it enjoys cross-view information which is ignored in \mathbf{F}_t^{2D} . Note that we empirically found that the whole process is training-free and the tensor shape can be also retained because we only adjust the attention receptive fields of \mathbf{F}_t^{3D} while all parameters are frozen. More implementation details about the 3D self-attention are discussed in the supplementary. After that, we average \mathbf{F}_t^{2D} and \mathbf{F}_t^{3D} as the final output of self-attention, which considers both globally long-range consistency for each separate 2D view and locally cross-view correlations.

Since the non-overlapping grid-based splitting with the same grid size suffers from limited feature interaction [25], we interactively two different grid sizes G_t across the whole denoising process with different time steps. Thus, our model could eliminate the isolation with overlapped 3D receptive fields.

Multi-View Aggregation-and-Rasterization (MV-AR). To fuse the multi-view latent predictions $\mathbf{X}_{n \in N}$, we first compute the view score $S_{n \in N}$, which quantifies the cosine similarity between the mesh normal and the screen view direction. Note that we omit the time step t for simplicity. The distance score $D_{n \in N}$ is computed by: $1 - d_i/Z_{far}$, where d_i denotes distance between pixel i and mesh surface, Z_{far} represent the upper bound of the whole scene. Then we derive the barycentric coordinate map $B_{n \in N} \in \mathbb{R}^{h \times w \times 3}$ from $R(\cdot)$, where channels are denoted by $(\alpha_{i1}, \alpha_{i2}, \alpha_{i3})$ respectively. Subsequently, we formulate the vertex latent features $\hat{\mathbf{X}}_{n,j} \in \mathbb{R}^4$ and vertex weights $W_{n,j} \in \mathbb{R}^1$ for view $n \in N$, vertex index $j \in J_c$ as:

$$\hat{\mathbf{X}}_{n,j} = \sum_{i=1}^{hw} \mathbf{X}_{n,i} \cdot \psi(B_{n,i}, \tau_b) / \eta, \quad (5)$$

$$W_{n,j} = \sum_{i=1}^{hw} S_{n,i} \cdot \psi(D_{n,i}, \tau_d) \cdot \psi(B_{n,i}, \tau_b) / \eta, \quad (6)$$

where $i \in hw$ indicate the 2D pixel indices; $\eta = \sum_{i=1}^{hw} \psi(B_{n,i}, \tau_b)$; ψ represents power function; both τ_b and τ_d denote the exponent. Furthermore, we aggregate $\hat{\mathbf{X}}_{n,j}$ into $\hat{\mathbf{X}}_j^{3D}$ across all N views as:

$$\hat{\mathbf{X}}_j^{3D} = \sum_{n=1}^N \hat{\mathbf{X}}_{n,j} \cdot \psi(W_{n,j}, \tau_w) / \omega, \quad \omega = \sum_{n=1}^N \psi(W_{n,j}, \tau_w). \quad (7)$$

With the aggregated 3D vertex feature $\hat{\mathbf{X}}_j^{3D}$, we further rasterize them back to the 2D plane with initial camera views and replace the foreground rendering area in predicted latent features \mathbf{X}_n with the re-rendered ones, resulting in $\hat{\mathbf{X}}_n^{2D}$. As 2D latent features $\hat{\mathbf{X}}_n^{2D}$ are rendered from the 3D feature mesh, they naturally achieve superior view consistency. We subsequently add the noise to the features to produce the latent features for step $t - 1$, as the diffusion iteration [38] displayed in Fig. 2(c).

Variance Alignment (VA). In diffusion models, controlling the proper step-wise variance through a designated noise schedule is critical for maintaining stability and producing high-quality images. However, we find that the rasterization process in latent space would cause variance degradation during the denoising, resulting in over-smoothed generations.

To provide an in-depth analysis of this phenomenon, we begin by examining the rendering formula (Eq. 4) and the unit constraint condition (Eq. 3). We observe that the rasterization process can be represented as a convex combination, which follows the same linear combination formula and coefficient conditions. Furthermore, the variance $Var(\cdot)$ represents the expectation of a square function, which is naturally convex. Thus the variance satisfies the fundamental formula and conditions of Jensen’s inequality as the convex function and the convex combination. Formally, Jensen’s inequality states that if $\varphi(\cdot)$ is a convex

function, and $z_{i \in N_z}$ are points in interval Z , where N_z is the number of sampling points, then for any non-negative weights λ_i that satisfy the condition: $\sum_{i=1}^{N_z} \lambda_i = 1$, the following inequality holds:

$$\varphi \left(\sum_{i=1}^{N_z} \lambda_i z_i \right) \leq \sum_{i=1}^{N_z} \lambda_i \varphi(z_i). \quad (8)$$

Thereby, for random variable set \mathbf{x}_i , and any non-negative weights λ_i satisfying the condition $\sum_{i=1}^{N_z} \lambda_i = 1$, we can indicate:

$$Var \left(\sum_{i=1}^{N_z} \lambda_i \mathbf{x}_i \right) \leq \sum_{i=1}^{N_z} \lambda_i Var(\mathbf{x}_i). \quad (9)$$

This means that the variance of the convex value combination is no larger than the convex combination of variance, and this inequality is deduced in the supplementary materials in detail.

Revisiting the rasterization process, we claim that the rasterized 2D latent feature $\tilde{\mathbf{X}}$ contain smaller variance compared to the aggregated one $\hat{\mathbf{X}}^{3D}$ as:

$$Var(\tilde{\mathbf{X}}^{2D}) \leq \sum_{u=1}^3 Var(\hat{\mathbf{X}}_u^{3D}) \cdot B_u^{3D} = Var(\hat{\mathbf{X}}^{3D}), \quad (10)$$

where B_u^{3D} indicates three barycentric coordinate banks with $u \in [1, 2, 3]$ triangle indices of each face. To address this issue, we propose a variance correction formulation defined as follows:

$$\tilde{\mathbf{X}}^{2D'} = \frac{\tilde{\mathbf{X}}^{2D} - \mu(\tilde{\mathbf{X}}^{2D})}{\delta^{2D}} \cdot \delta^{3D} + \mu(\hat{\mathbf{X}}^{3D}), \quad (11)$$

where $\mu(\cdot)$ represents the mean operation; δ denotes standard deviation. And $\delta^{3D} = \sqrt{Var(\tilde{\mathbf{X}}^{3D})}$ is achieved by:

$$Var(\tilde{\mathbf{X}}^{3D}) = \sum_{u_1=1}^3 (B_{u_1}^{3D})^2 \cdot Var(\hat{\mathbf{X}}^{3D}) \quad (12)$$

$$+ \sum_{\substack{u_1, u_2=1 \\ u_1 \neq u_2}}^3 2 \cdot B_{u_1}^{3D} \cdot B_{u_2}^{3D} \cdot Cov(\hat{\mathbf{X}}_{u_1}^{3D}, \hat{\mathbf{X}}_{u_2}^{3D}), \quad (13)$$

where $Cov(\cdot, \cdot)$ represents covariance function. Eq. 11 meticulously aligns the feature variance after the rasterization to the mesh feature $\hat{\mathbf{X}}^{3D}$, which is capable of synthesizing realistic textures with more high-frequency details. We also empirically verified the effectiveness of VA in Fig. 5 and related ablation studies of Sec. 4.4.




Fig. 3: The illustration of inpainting refinement. (a) shows the image view rendered from the initially inconsistent texture $\hat{\mathbf{I}}^{3D}$; (b) represents the dilated inpainting mask rendered from 3D mask \mathbf{M} ; (c) is the depth map rendered from the input mesh; (d) indicates the updated final texture through our inpainting refinement.

3.3 Inpainting Refinement

Due to the intrinsic discrepancy in resolution size and channel number between the latent and image domains, consistent latent cannot guarantee consistent images. In this section, we introduce inpainting refinement to further alleviate the inconsistency in pixel values across different views of the image domain. Similar to the latent aggregation process, we generate the initial texture via pixel domain aggregation. We first identify those inconsistent vertices by the variance of related aggregation pixels. Then, we incorporate pixel domain texture inpainting approach [44] to refine those inconsistent regions, and all intermediate results of the inpainting refinement are shown in Fig. 3. Specifically, given predicted multi-view images $\mathbf{I}_{n_f \in N_f}$, where $N_f = 4$ views are sampled from the total N views, we calculate the variance of vertex j as:

$$Var_j = \sum_{n_f=1}^{N_f} (\hat{\mathbf{I}}_{j,n_f}^{3D} - \mu(\hat{\mathbf{I}}_j^{3D}))^2 / (N_f - 1), \quad (14)$$

where $\hat{\mathbf{I}}_{n_f}^{3D} = R^{-1}(\mathbf{I}_{n_f}) \in \mathbb{R}^{J_f \times 3}$ indicates color repository re-projected from $n_f \in N_f$ sampled images; $\mu(\hat{\mathbf{I}}_j^{3D})$ indicates the mean operation across all N_f views; J_f denotes vertex number. Note that $J_f < J_c$, while meshes with J_f, J_c vertexes are used for the fine-grained pixel space and the coarse latent space, respectively. Variance quantifies the dispersion of data points in a dataset. Since the value of each vertex is accumulated from N_f pixel values, we can employ a predetermined threshold λ to discern vertices that exhibit inconsistency. We construct an indicator function as a 3D mask $\mathbf{M} \in \mathbb{R}^{J_f \times 1}$:

$$\mathbf{M}_j = \begin{cases} 1, & Var_j > \lambda \\ 0, & Var_j \leq \lambda \end{cases}. \quad (15)$$

After that, we perform pixel domain inpainting refinement through rendering both $\hat{\mathbf{I}}^{3D}$ and 3D mask \mathbf{M} with $R(\cdot)$ into four separate views and 2D masks. Then we follow Repaint3D [44] which first dilates inpainting mask with 8×8 kernel, and then utilize Depth-SD to inpaint them through the same prompts

and depth rendered by the input mesh. Subsequently, the inpainting procedure is carried out sequentially until all views have been inpainted and finally produces the refined texture $\tilde{\mathbf{I}}^{3D}$ as shown in the rightmost of Fig. 2(a).

4 Experiments

4.1 Experimental Settings

Implementation Details. We utilize the planar isotropic remeshing algorithm proposed in Repaint3D [44] to re-mesh input mesh, and use the PyTorch3D library¹ to render 3D objects and employ Depth-SD to conduct 3D-2D collaborative denoising and inpainting refinement. In the collaborative denoising stage, we engage $N = 9$ equidistant viewpoints with each view separated by the angular of 40° . The grid size G_t is altered within two values: 0.34 and 0.25, while the whole volume space is normalized to $[-1, 1]$. The distance upper bound Z_{far} is set to 5. Additionally, exponents τ_b , τ_w and τ_f are set to 2.0, 3.0 and 6.0 respectively, where τ_f is used for pixel aggregation. In the inpainting refinement stage, we commence by selecting $N_f = 4$ view images at angular positions of $0^\circ, 80^\circ, 160^\circ$, and 280° for the initial texture aggregation. Subsequently, a variance threshold λ of 0.005 is employed to identify inconsistent areas. For parameters related to Depth-SD, we set the denoising step as 50 in the denoising and inpainting stages, while the 3D-2D collaborative denoising is applied to the first 45 steps.

Datasets. Despite the extensive study on texture, there remains a lack of standardized datasets and evaluation metrics as a solid benchmark. To comprehensively evaluate the performance of the comparison methods, we construct the largest 3D shape texture evaluation dataset, which consists of three subsets (named *SubObj*, *SubShape*, and *SubTex*) from three open-source 3D datasets (Objaverse [11], ShapeNetSem [41], and ShapeNet [4]). *SubObj* is inherited from Text2Tex, containing 410 objects from Objaverse, we exclude some too-simple shapes, which remain 401 objects. Inspired by the sampling methodology of Text2Tex, we extract 445 3D shapes across categories from the ShapeNetSem dataset to form the *SubShape*, with three meshes sampled per category based on vertex count distribution. In addition, we sample publicly authorized meshes with prompts from the Texfusion dataset to assemble 43 mesh-prompt pairs, primarily sourced from ShapeNet, into the subset *SubTex*.

Metrics. After reviewing relevant metrics, we employ the Fréchet Inception Distance (FID) [13] and the CLIP-based extension of FID, denoted as CLIP-FID [21], to measure texture fidelity. Moreover, we utilize the Clip-Score [35] metric to quantify the correspondence score between the rendered image and the textual prompt. We also utilize CLIP-Var [1] to ascertain the consistency across multiple rendering views following [1]. In detail, following the evaluation setting of Repaint3D [44], we render the textured mesh with size of 512 from eight distinct camera perspectives, evenly spaced at 45° intervals, and overwrite the background regions with white color to emphasize the foreground object for more

¹ <https://github.com/facebookresearch/pytorch3d>

Table 1: Quantitative comparison of different texture methods on three datasets.

Datasets	Method	FID ↓	ClipFID ↓	ClipScore ↑	ClipVar ↑
SubTex	Texture [37]	150.21	26.92	26.90	82.37
	Text2Tex [5]	112.41	16.26	30.08	81.45
	SyncMVD [24]	65.30	16.76	28.78	81.93
	Repaint3D [44]	78.65	10.65	30.88	78.96
	VCD-Texture	56.29	6.84	31.65	83.97
SubShape	Texture [37]	64.78	18.08	27.20	82.32
	Text2Tex [5]	40.46	7.96	27.76	82.18
	SyncMVD [24]	32.44	6.18	28.76	82.76
	Repaint3D [44]	29.21	5.25	28.39	80.18
	VCD-Texture	19.46	2.37	28.98	82.93
SubObj	Texture [37]	65.30	16.76	28.78	81.93
	Text2Tex [5]	43.71	7.46	29.27	82.08
	SyncMVD [24]	34.00	5.60	30.08	84.52
	Repaint3D [44]	29.77	4.44	30.30	81.45
	VCD-Texture	21.19	2.33	30.42	83.64

focused comparisons. To compute the FID metric, following [3,44], ground-truth images are synthesized through Depth-SD, with each image being conditioned on the same prompt and synthesized from the designated evaluation viewpoints. The computation of Clip-Score is achieved upon the minimum value derived from eight cosine similarity scores, each one representing the similarity between the normalized features of the CLIP image and the corresponding text features. Similarly, CLIP-Var is determined by the minimal mutual cosine similarity scores among the eight normalized CLIP image features.

4.2 Quantitative Evaluation

In this section, we quantitatively compare our VCD-Texture method against state-of-the-art texturing methods, including TEXTure [37], Tex2Tex [5], Repaint3D [44], and SyncMVD [24], which have released their source codes. As for the comparison methods, we attain the results by running their official code with their default settings. As reported in Tab. 1, compared with the inpainting-based methods (Tex2Tex, TEXTure, and Repaint3D) and the optimization-based methods (SyncMVD), our method achieves lower FID and ClipFID and higher ClipScore and ClipVar results. The results demonstrate that our VCD-Texture method possesses superior performance in terms of both texture fidelity and multi-view consistency.

4.3 Qualitative Evaluation

Fig. 4 depicts textures generated by the same text prompt and 3D mesh. We can conclude that: 1) Inpainting-based methods, Texture [37], Text2Tex [5], and

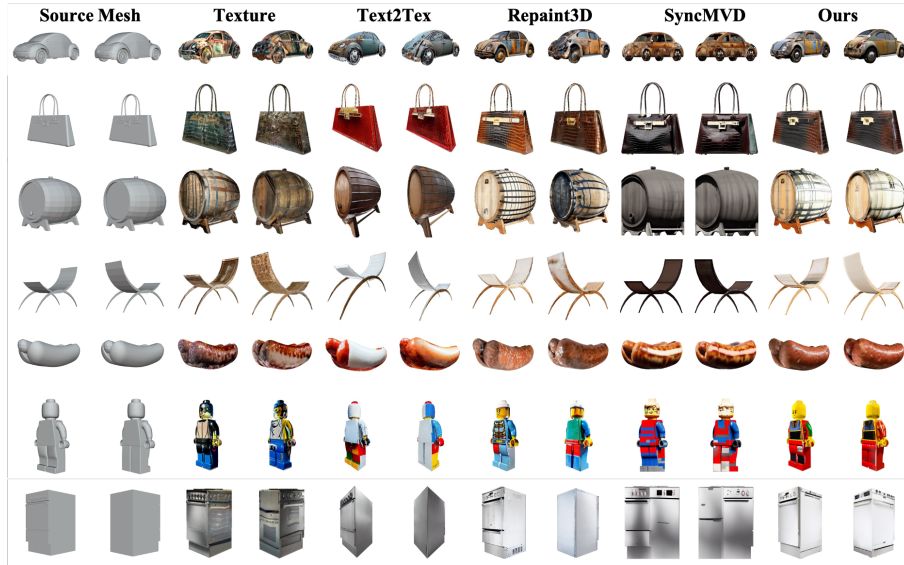


Fig. 4: Qualitative comparisons of text-guided texture synthesis. Prompts from top to down are: “old and rusty volkswagon beetle”, “crocodile skin handbag”, “barrel”, “half moon chaise”, “sausage”, “lego”, “electric oven”.

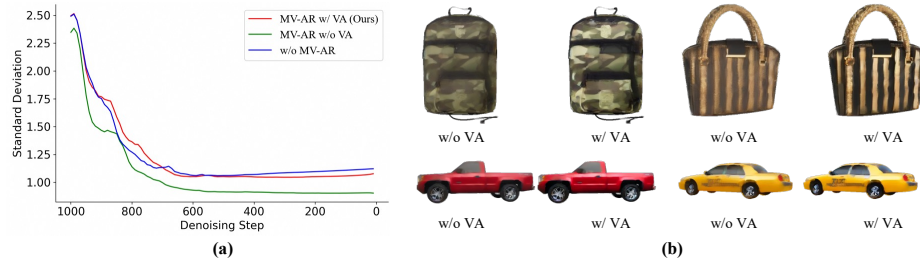
Repaint3D [44], tend to generate inconsistent textures on the two opposing views. 2) The optimization-based method, SyncMVD [24], can generate more consistent textures compared to the inpainting-based ones. However, SyncMVD tends to produce textures that are excessively smooth, resulting in a loss of intricate details and fine texture quality. 2) The optimization-based method, SyncMVD [24], can generate more consistent textures compared to the inpainting-based ones. However, SyncMVD tends to produce textures that are excessively smooth, resulting in a loss of intricate details and fine texture quality. This deficiency could be attributed to the frequent application of synchronized multi-view diffusion coupled with the absence of a VA. Such a process generates latent features with low variance, which in turn, leads to the creation of textures that are notably smooth yet exhibit high consistency. 3) Our VCD-Texture method is capable of generating more multi-view consistent and fidelity textures compared to both the inpainting-based and optimization-based methods.

4.4 Ablation Study

Factor-by-factor Analyzation. We conduct a comprehensive ablation study on the SubTex datasets to validate the effect of different components. The experimental results are summarized in Tab. 2, from which we can draw the following conclusions:

Table 2: The ablation study for MV-AR, JNP, VA, the incorporating of distance score (DS), and inpainting refinement (IR).

No.	MV-AR	DS	JNP	VA	IR	Fid ↓	ClipFid ↓	ClipScore ↑	ClipVar ↑
(1)	✓					58.87	7.39	31.32	82.87
(2)	✓	✓				58.73	7.25	31.47	83.12
(3)	✓	✓	✓			58.02	7.16	31.57	83.62
(4)	✓	✓	✓	✓		56.05	6.88	31.62	83.94
(5)	✓	✓	✓	✓	✓	56.29	6.84	31.65	83.97

**Fig. 5:** The effectiveness of VA. (a) shows the standard deviation curve of three denoising policies; (b) showcases the qualitative comparison with and without (w.o) VA.

- Comparing the models (1) and (2), the distance score (DS) improves ClipFID and multi-view consistency. This validates pixel distance is a valid metric that can aid the view score in selecting more consistent pixels in the aggregating process.
- Comparing the models (2) and (3), the integration of JNP facilitates the exchange of information within the 3D domain, which improves the semantic metric ClipScore and also enhances the consistency across views.
- Comparing models (3) and (4), incorporation of VA leads to a significant improvement in the FID metric, this proves that rectifying feature variance to align original diffusion distribution after rasterization helps to generate high-fidelity images.
- Comparing models (4) and (5), since the inconsistent regions are small, thus inpainting refinement doesn't contribute much more improvements in evaluation metrics. However, the subsequent visual comparative analysis clearly reveals that inpainting refinement is instrumental in mitigating blurred regions and rectifying visual artifacts.

Variance Alignment Analysis. To elucidate the process of variance reduction, we present the standard deviation trajectories of latent features during the denoising phase in Fig. 5. The blue curve (without MV-AR) depicts the denoising process without consistent multi-view projection. The green curve (MV-AR without VA) illustrates the feature variance resulting from a consistent denoising projection, *i.e.*, lifting the feature to 3D space and rasterizing it back to 2D




Fig. 6: Qualitative results of the inpainting refinement.

without VA. The red curve (MV-AR with VA) represents the feature variance when VA is incorporated into the MV-AR.

A comparative analysis of these curves reveals that the green curve consistently maintains a lower trajectory than the standard blue line, indicative of the variance diminishing process and congruent with the theoretical proof’s conclusions. The trajectory of the red curve (MV-AR with VA), which adopts VA, displays a similar pattern to the blue curve and ultimately converges just below the blue line, which validates the efficacy of the proposed VA method. The lower endpoint of the red line can be attributed to the features generated by consistent projection being more consistent than those in the general denoising process.

Effectiveness of Inpainting Refinement. As illustrated in Fig. 6, we visualize the texture improvements w.r.t. the inpainting refinement process on four examples of *SubTex* dataset. We can observe that the dilated inpainting mask can accurately identify those inconsistent pixels, and the subsequent pixel domain inpainting refinement is capable of refining the inconsistent regions and achieving more fidelity and high-quality texturing results.

5 Conclusion

We propose a novel collaborative denoising 3D texture synthesis approach, VCD-Texture, to mitigate the gap between the 2D diffusion generation and 3D objects. VCD-Texture injects 3D geometries into the 2D diffusion denoising process, designing Joint Noise Prediction (JNP) and Multi-View Aggregation-and-Rasterization (MV-AR) modules to incorporate features in 2D and 3D space. Moreover, we theoretically analyze the variance bias issues caused by the rasterization in MV-AR, which is eliminated by the proposed Variance Alignment (VA) technique. To further reduce the intrinsic discrepancy in latent feature and image pixel domains, we design an inpainting refinement to rectify identified inconsistent regions. Through a collaborative denoising process and inpainting refinement process, VCD-Texture enables the generation of consistent textures with diverse and high-quality details.

References

1. Anonymous: Learning pseudo 3d guidance for view-consistent 3d texturing with 2d diffusion (2023), <https://openreview.net/forum?id=2A199SAhW3>
2. Balas, B.J.: Texture synthesis and perception: Using computational models to study texture representations in the human visual system. *Vision research* **46**(3), 299–309 (2006)
3. Cao, T., Kreis, K., Fidler, S., Sharp, N., Yin, K.: Texfusion: Synthesizing 3d textures with text-guided image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4169–4181 (2023)
4. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012* (2015)
5. Chen, D.Z., Siddiqui, Y., Lee, H.Y., Tulyakov, S., Nießner, M.: Text2tex: Text-driven texture synthesis via diffusion models. *arXiv preprint arXiv:2303.11396* (2023)
6. Chen, Y., Ip, H.H.: Texture evolution: 3d texture synthesis from single 2d growable texture pattern. *The Visual Computer* **20**, 650–664 (2004)
7. Chen, Y.: 3d texture mapping for rapid manufacturing. *Computer-Aided Design and Applications* **4**(6), 761–771 (2007)
8. Chen, Y., Chen, R., Lei, J., Zhang, Y., Jia, K.: Tango: Text-driven photorealistic and robust 3d stylization via lighting decomposition. *Advances in Neural Information Processing Systems* **35**, 30923–30936 (2022)
9. Chen, Z., Yin, K., Fidler, S.: Auv-net: Learning aligned uv maps for texture transfer and synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1465–1474 (2022)
10. Cula, O.G., Dana, K.J.: 3d texture recognition using bidirectional feature histograms. *International Journal of Computer Vision* **59**, 33–60 (2004)
11. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13142–13153 (2023)
12. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
13. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
14. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
15. Jetchev, N., Bergmann, U., Vollgraf, R.: Texture synthesis with spatial generative adversarial networks. *arXiv preprint arXiv:1611.08207* (2016)
16. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4401–4410 (2019)
17. Klodt, M., Vedaldi, A.: Supervising the new with the old: learning sfm from sfm. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 698–713 (2018)
18. Kniaz, V., Mizginov, V.: Thermal texture generation and 3d model reconstruction using sfm and gan. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **42**, 519–524 (2018)

19. Kniss, J., Lefohn, A., Strzodka, R., Sengupta, S., Owens, J.D.: Octree textures on graphics hardware. In: ACM SIGGRAPH 2005 Sketches, pp. 16–es (2005)
20. Kry, P.G., James, D.L., Pai, D.K.: Eigenskin: real time large deformation character skinning in hardware. In: Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation. pp. 153–159 (2002)
21. Kynkäänniemi, T., Karras, T., Aittala, M., Aila, T., Lehtinen, J.: The role of imagenet classes in fr\`echet inception distance. arXiv preprint arXiv:2203.06026 (2022)
22. Lefebvre, S., Hoppe, H.: Appearance-space texture synthesis. *ACM Transactions on Graphics (TOG)* **25**(3), 541–548 (2006)
23. Lefebvre, S., Hoppe, H.: Appearance-space texture synthesis. *ACM Transactions on Graphics (TOG)* **25**(3), 541–548 (2006)
24. Liu, Y., Xie, M., Liu, H., Wong, T.T.: Text-guided texturing by synchronized multi-view diffusion. arXiv preprint arXiv:2311.12891 (2023)
25. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
26. Ma, Y., Zhang, X., Sun, X., Ji, J., Wang, H., Jiang, G., Zhuang, W., Ji, R.: X-mesh: Towards fast and accurate text-driven 3d stylization via dynamic textual guidance. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2749–2760 (2023)
27. Michel, O., Bar-On, R., Liu, R., Benaim, S., Hanocka, R.: Text2mesh: Text-driven neural stylization for meshes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13492–13502 (2022)
28. Mohammad Khalid, N., Xie, T., Belilovsky, E., Popa, T.: Clip-mesh: Generating textured meshes from text using pretrained image-text models. In: SIGGRAPH Asia 2022 conference papers. pp. 1–8 (2022)
29. Mou, C., Wang, X., Xie, L., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. arXiv preprint arXiv:2302.08453 (2023)
30. Oechsle, M., Mescheder, L., Niemeyer, M., Strauss, T., Geiger, A.: Texture fields: Learning texture representations in function space. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4531–4540 (2019)
31. Pineda, J.: A parallel algorithm for polygon rasterization. In: Proceedings of the 15th annual conference on Computer graphics and interactive techniques. pp. 17–20 (1988)
32. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
33. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022)
34. Portenier, T., Arjomand Bigdeli, S., Goksel, O.: Gramgan: Deep 3d texture synthesis from 2d exemplars. *Advances in Neural Information Processing Systems* **33**, 6994–7004 (2020)
35. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)

36. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* **1**(2), 3 (2022)
37. Richardson, E., Metzger, G., Alaluf, Y., Giryas, R., Cohen-Or, D.: Texture: Text-guided texturing of 3d shapes. *arXiv preprint arXiv:2302.01721* (2023)
38. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
39. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* **35**, 36479–36494 (2022)
40. Sanghi, A., Chu, H., Lambourne, J.G., Wang, Y., Cheng, C.Y., Fumero, M., Malekshah, K.R.: Clip-forge: Towards zero-shot text-to-shape generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18603–18613 (2022)
41. Savva, M., Chang, A.X., Hanrahan, P.: Semantically-Enriched 3D Models for Common-sense Knowledge. *CVPR 2015 Workshop on Functionality, Physics, Intentionality and Causality* (2015)
42. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems* **35**, 25278–25294 (2022)
43. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., et al.: Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792* (2022)
44. Wang, T., Kanakis, M., Schindler, K., Van Gool, L., Obukhov, A.: Breathing new life into 3d assets with generative repainting. In: *Proceedings of the British Machine Vision Conference (BMVC)*. BMVA Press (2023)
45. Wang, Z.J., Montoya, E., Munechika, D., Yang, H., Hoover, B., Chau, D.H.: Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. *arXiv preprint arXiv:2210.14896* (2022)
46. Wei, L.Y.: Tile-based texture mapping on graphics hardware. In: *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*. pp. 55–63 (2004)
47. Ying, L., Hertzmann, A., Biermann, H., Zorin, D.: Texture and shape synthesis on surfaces. In: *Rendering Techniques 2001: Proceedings of the Eurographics Workshop in London, United Kingdom, June 25–27, 2001* 12. pp. 301–312. Springer (2001)
48. Yu, X., Dai, P., Li, W., Ma, L., Liu, Z., Qi, X.: Texture generation on 3d meshes with point-uv diffusion. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4206–4216 (2023)
49. Zeng, X.: Paint3d: Paint anything 3d with lighting-less texture diffusion models. *arXiv preprint arXiv:2312.13913* (2023)
50. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3836–3847 (2023)