

NI-Tex: Non-isometric Image-based Garment Texture Generation

Hui Shan^{1,2,3}, Ming Li^{1,2,3}, Haitao Yang⁴, Kai Zheng³, Sizhe Zheng^{2,3},
 Yanwei Fu^{2,5}, Xiangru Huang^{3†}

¹Zhejiang University ²Shanghai Innovation Institute ³Westlake University
⁴University of Texas at Austin ⁵Fudan University




Figure 1. NI-Tex takes an image prompt and a target garment geometry as input, and generates high-quality PBR textures that faithfully transfer the textures and patterns from the input image to the target garment. Unlike other existing methods, the generation quality of NI-Tex does not degrade for challenging image-garment pairs with strong topological and geometric differences, enabling superior flexibility.

Abstract

Existing industrial 3D garment meshes already cover most real-world clothing geometries, yet their texture diversity remains limited. To acquire more realistic textures, generative methods are often used to extract Physically-based Rendering (PBR) textures and materials from large collections of wild images and project them back onto garment meshes. However, most image-conditioned texture generation approaches require strict topological consistency between the input image and the input 3D mesh, or rely on accurate mesh deformation to match to the image poses, which significantly constrains the texture generation quality and flexibility.

To address the challenging problem of non-isometric image-based garment texture generation, we construct 3D Garment Videos, a physically simulated, garment-centric dataset that provides consistent geometry and material

supervision across diverse deformations, enabling robust cross-pose texture learning. We further employ Nano Banana for high-quality non-isometric image editing, achieving reliable cross-topology texture generation between non-isometric image-geometry pairs. Finally, we propose an iterative baking method via uncertainty-guided view selection and reweighting that fuses multi-view predictions into seamless, production-ready PBR textures. Through extensive experiments, we demonstrate that our feedforward dual-branch architecture generates versatile and spatially aligned PBR materials suitable for industry-level 3D garment design. The code and dataset have been released on <https://github.com/SII-Hui/NI-Tex>.

1. Introduction

Acquiring high-quality 3D garment assets is important for a number of applications such as virtual reality, human avatars and physical simulation. While existing garment

[†]Corresponding author (huangxiangru@westlake.edu.cn).

datasets already cover a rich distribution of geometries, a key question is how to augment existing garment geometries with more diverse and photorealistic Physically-based Rendering (PBR) textures. A common strategy to acquire more realistic textures is to employ texture generation methods that synthesize desired PBR textures conditioning on a target geometry and an user-provided image prompt, which offers precise control over the generated textures.

Current texture generation models are typically trained on large-scale 3D datasets such as Objaverse [5] and TexVerse [31], which provide massive collections of 3D assets with PBR textures. These models usually use multi-view attention mechanisms to learn implicit image-geometry correspondences to faithfully transfer the visual appearance and material properties from the image to the target geometry. However, such correspondence learning relies heavily on geometric and topological similarity between the input image and the target geometry. When there exists a significant topological or geometric discrepancy, especially in cases of non-isometric deformation or garment topology variation, the generation quality deteriorates drastically, leading to distorted or inconsistent textures like Figure 2.

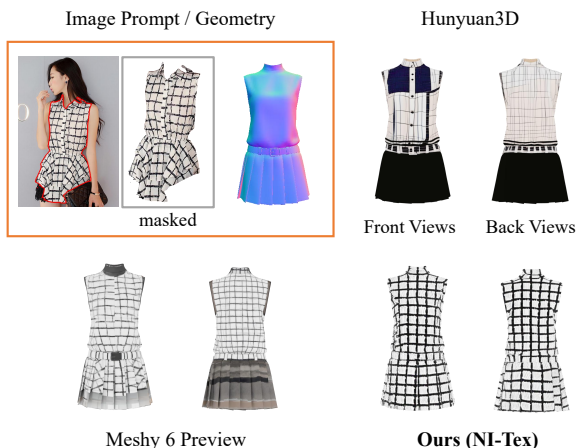


Figure 2. Texture generation becomes unreliable when the image prompt and the given mesh exhibit topology inconsistency in Hunyuan3D [12] and Meshy 6 Preview.

In this work, we address this critical challenge by proposing **NI-TeX**, a training framework that promotes non-isometric image-based garment texture generation, which focuses on robust and controllable texture generation under topological and geometric image-garment inconsistencies. NI-TeX builds upon a texture generation backbone pre-trained on Objaverse and TexVerse, and we further fine-tune it using datasets with rich garment deformation data such as BEDLAM [1] to improve the generalization for garment deformations.

To accommodate for topological image-garment differ-

ences, we introduce a novel data augmentation strategy that leverages image editing tools (e.g., Nano Banana) to modify the topology and geometry of the image prompt. We propose a set of carefully designed strategies to regulate this process for multiple purposes: 1) ensure diverse topology modifications; 2) reduce the sim-to-real gap when applying image editing models; and 3) preserve the texture consistency before and after modification. This process results in a new garment texture dataset with a special enhancement for non-isometric image-garment pairs. Extensive qualitative and quantitative experiments suggest that NI-TeX significantly outperforms existing baselines with superior robustness and generation quality under challenging topological and geometric differences.

Beyond dataset and model design, we further propose an iterative baking algorithm to improve the stability of multi-view texture integration. Specifically, we train an Uncertainty Quantification (UQ) model to predict the per-pixel uncertainty on the generated texture maps. This uncertainty representation is capable of capturing a wide range of baking artifacts, such as holes and blurriness. To acquire the training data for the UQ model, we employ the trained texture generation models and simulate the baking error on another set of edited image prompts. The uncertainty annotation can be easily collected by comparing the generated textures with the ground-truth texture maps.

To fix the baking artifacts, we iteratively select new views that maximize the average per-pixel uncertainty, and apply simple reweighting strategy to robustly merge multi-view generated textures into integrated mesh textures.

To summarize, our contributions are:

- We are the first to adopt a feedforward architecture to effectively tackle the challenging task of non-isometric image-based texture generation. A novel image editing based framework for addressing non-isometric image-garment inconsistencies, effectively leveraging powerful image editing tools to enhance controllability and robustness of image-based garment texture generation.
- A new dataset designed to improve the generation quality for non-isometric image garment pairs, which will be publicly released to facilitate future research.
- An uncertainty-aware iterative baking algorithm that improves the stability and quality of multi-view texture integration.

2. Related work

2.1. 3D Texture Generation

3D Texture Generation methods study how to add textures to given 3D geometries (represented by mesh) from user-defined text or image prompts. TEXTure [20] and Text2Tex [2] generate images from different viewpoints using diffusion models. These images are then back-projected

onto the surface of the geometry. To further improve the alignment between the generated images and the given geometry, Easi-Tex [18] uses edge maps to guide the diffusion process. However, images from different viewpoints often cannot fully cover the entire geometry. Thus, some faces of the mesh remain untextured. This problem is more evident for objects with complex structures. Paint3D [28] additionally introduces an inpainting step on the UV map of the geometry, allowing all faces to be textured. By incorporating global texture information, it further improves the texture consistency across multiple viewpoints. Although the above methods achieve good results in 3D texture generation, they overlook the process of lighting and shadows in the generated images. As a result, the produced textures fail to deliver realistic visual effects under new lighting conditions.

2.2. PBR Texture Generation

PBR texture generation methods output the object’s base color (albedo) and its perceptual properties (metallic and roughness) separately. The Score-Distillation-Sampling (SDS) based optimization approach and the data-driven feed-forward approach are two common categories for PBR texture generation. Fantasia3D [3], Matlaber [26], and Paint-it [27] leverage the priors of pretrained 2D diffusion models and use SDS to distill the PBR texture attributes of the object. FlashTex [6] and DreamMat [30] train light ControlNet to achieve light-aware diffusion models by augmenting the training dataset with simulated reflection effect under various lighting conditions. However, SDS-based optimization is time-consuming and sometimes produces less realistic visual effects. Another line of research focuses on feed-forward training enabled by existing open-source 3D datasets [4, 5, 31]. MaterialAnything [11] and Hunyuan3D 2.1 [12] render PBR-annotated objects from the datasets (mainly from Objaverse) under different lighting conditions and viewpoints. They use albedo, metallic, and roughness maps for supervision in training, which makes the model material-aware. The feed-forward approach effectively reduces inference time, enabling efficient and large-scale PBR texture generation. Hunyuan3D 2.1 applies rigid operations such as scaling, translation, and rotation during rendering to augment training data for better generalization. However, the model still struggles to generate textures between image-geometry pairs with non-isometric differences such as topology changes.

2.3. Image-based Virtual Try-Off and Image-Prompt 3D Garment Texture Generation

TryOffDiff [22] is the first to introduce Virtual Try-Off (VTOFF), which extracts garment texture from given photos of clothed individuals and generates standardized 2D garment images. OMFA [16] and Voost [13] build a uni-

fied framework that combines Virtual Try-On (VTON) and VTOFF techniques to achieve better results. However, they are limited to generating 2D images and cannot bake the extracted garment textures onto a given 3D geometry. In the field of image-based 3D garment generation, methods such as Pix2Surf [17], Cloth2Tex [7], and Garment3DGen [21] typically require a 2D garment image with canonical pose and then align a template mesh to the garment image through mesh deformation, which suffers from compound errors from the mesh alignment procedure and has limited applicability. FabricDiffusion [29] successfully extracts sewing pattern textures and logos from garment images, but struggles with texture extraction beyond sewing patterns. Reweaver [14] focuses specifically on garment geometric reconstruction. Meanwhile, GarmentCrafter [24], and Dress-1-to-3 [15] adopt a pipeline that performs geometric reconstruction followed by texture generation. However, the reconstructed geometries rarely reach production-ready quality, while their texture synthesis also remains constrained by strict pose alignments or sewing patterns.

3. Preliminary

NI-Tex is designed for image-based garment texture generation with a special enhancement for non-isometric image-garment pairs. In this section, we first formally define the task of image-based garment texture generation (section 3.1), then briefly introduce the backbone network architecture that we adopted from Hunyuan3D (section 3.2).

3.1. Problem Setup

Image-based garment texture generation takes an input RGB image $I \in \mathbb{R}^{H \times W \times 3}$ and an input geometry represented by a mesh, and output multiple generated texture maps (albedo, roughness, metallic). Each texture map is defined as an image $a \in \mathbb{R}^{N \times N \times C}$ in the UV space of the input geometry. Albedo, roughness, metallic texture map contains $C = 3, 1, 1$ channels, respectively.

3.2. Network Architecture

To preserve the texture information from the input image prompt during training, we adopt a feedforward dual-branch architecture composed of a guidance branch and a generation branch like Hunyuan3D. The guidance branch extracts hierarchical features from the input image and feeds them into the corresponding layers of the training branch as reference information. The training branch takes multi-view normal and position maps as input, while leveraging the reference features to perform multi-channel inference of albedo and MR (metallic and roughness) attributes.

Multi-Channel Aligned Attention (MCAA) We adopt the MCAA module from MaterialMVP [10] to connect the

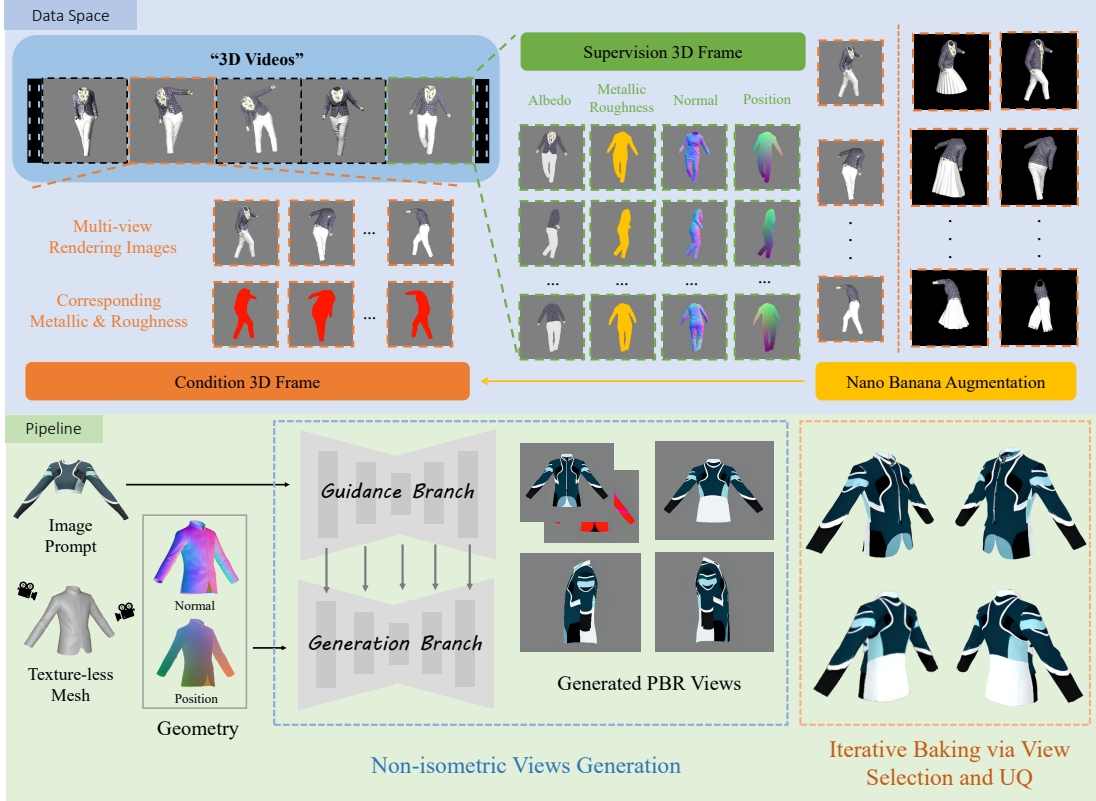


Figure 3. **Overview of NI-Tex.** In the data space (top), we construct our non-isometric training dataset from 3D Garment Videos by randomly selecting two frames, one as the condition 3D frame and the other as the supervision 3D frame, to enhance the model’s generalization across human poses, geometric deformations, and lighting variations. To further improve robustness to different garment topology, we apply Nano Banana for image editing on randomly rendered views of the condition 3D frame. In the pipeline (bottom), we render the input texture-less mesh to obtain normal and position maps as geometric constraints and employ a dual-branch architecture to achieve non-isometric PBR view generation. Finally, view selection and Uncertainty Quantification (UQ) are used to iteratively bake the results across multiple viewpoints.

two branches. As defined in Equation (1), we extract the reference texture features from the guidance branch and inject them into the albedo channel of the training branch:

$$\text{Attn}_{albedo} = \text{Softmax} \left(\frac{Q_{albedo} K_{ref}^T}{\sqrt{d}} \right) \cdot V_{ref} \quad (1)$$

To achieve spatial and geometric alignment between the MR and albedo attributes, we inject the albedo attention into the MR latent representation:

$$z_{MR}^{new} = z_{MR} + \text{Attn}_{albedo}, \quad (2)$$

4. Non-isometric Garment Texture Generation

NI-Tex proposes multiple novel strategies to enhance the generation quality for non-isometric image-garment pairs. First, to prepare training datasets for non-isometric image-garment pairs, we propose to use simulation data and image editing techniques to mitigate the generalization gap in

terms of human pose, garment topology, lighting and geometric deformation (section 4.1). To compensate for severe geometric and topological differences in test time, which can cause artifacts such as holes and view-inconsistent textures, we also propose an iterative baking procedure based on uncertainty quantification and view-selection (section 4.2). Finally, in section 4.3, we elaborate the training architecture of NI-Tex.

4.1. Dataset Preparation

In this section, we introduce how we construct our training dataset for enhancing non-isometric image-garment pairs.

Cross-pose Augmentation. We construct 3D Garment Videos based on the **BEDLAM**[1] dataset. BEDLAM contains nearly 1,700 distinct garment albedo textures, each associated with hundreds of motion sequences, and each sequence consists of hundreds of frames. Since the original dataset does not provide roughness or metallic attributes, we augment the selected motion sequences by assigning

PBR material properties to each frame. Specifically, we set roughness and metallic values as:

$$\text{roughness} \sim \mathcal{U}(0, 1), \quad \text{metallic} = 0 \quad (3)$$

For each physically-simulated motion sequence in BED-LAM, we extract per-frame garment geometry and form a sequence $V = \{M_1, M_2, M_3, \dots, M_n\}$. The geometric deformations of each garment M_i are driven by coherent temporal human motion, with a consistent albedo texture map shared across all frames (See Figure 3). We also denote a as the albedo texture map. We observe that fabric specular appearance is mainly determined by roughness, while metallic maps have little visual effect. Additionally, a single garment usually has uniform material properties, so its PBR attributes should remain consistent.

During training, we randomly sample two frames from each motion sequence (Figure 6, left). One frame is used as the **condition 3D frame** and the other as the **supervision 3D frame**. For the condition frame, we select one illuminated view as the input image prompt and use its normal and position maps from ten viewpoints as geometric constraints. For the supervision frame, we use its PBR texture attributes from ten viewpoints as supervision. This cross-frame supervision enables the model to learn texture consistency and variations across poses, resulting in more stable texture generation. Moreover, since any two frames in a 3D garment video can form a valid pair, the dataset expands combinatorially, scaling hundreds of thousands of frames into tens of billions of training samples.

Cross-topology Augmentation. To enable texture generation across diverse garment topologies (e.g., generating textures from skirts or shorts images onto long-pants meshes), we randomly sample rendered views from 3D Garment Videos and apply **Nano Banana** to edit garment topology while preserving the original textures. We use illuminated renderings instead of albedo images, as they yield more realistic results and reduce the domain gap during inference. The edited images replace the original condition images, while supervision remains from the original supervision 3D frames. This image–garment training effectively distills texture identity consistency from Nano Banana. To avoid incorrect distillation, we preserve the semantic integrity of garment textures following three key principles (Figure 4):

1. **Category-wise consistency.** When editing full-body garments, textures of upper and lower parts (e.g., tops and skirts or pants) should not drift or swap.
2. **Inner–outer consistency.** For layered outfits, outer garment textures must remain distinct from inner ones.
3. **Allowance for auxiliary human parts.** We allow Nano Banana to occasionally generate extra human regions, which encourages the model to focus on garment materials and ignore irrelevant areas in wild images.




Figure 4. We use Nano Banana to edit garment topology while preserving texture consistency (first row). We ensure category-wise and inner–outer texture consistency to avoid texture swaps or layering confusion (second and third rows). Additional human body generations (fourth row) are acceptable, as image–garment training improves NI-Tex’s understanding of garment textures.

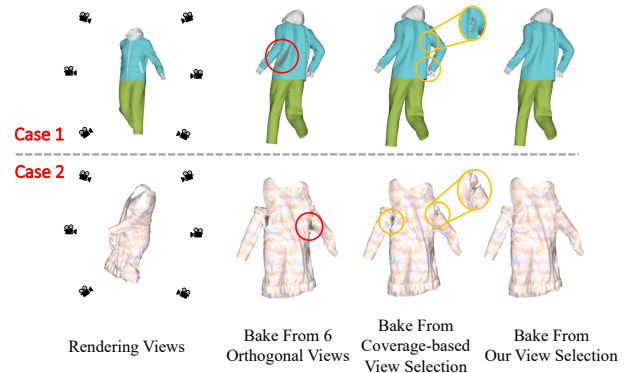


Figure 5. We render each garment mesh from six viewpoints and bake the results to check coverage. Self-occlusion under orthogonal views leaves many regions missing. Coverage-based view selection improves coverage but still misses small areas, while ours achieves full mesh coverage.

4.2. Iterative Baking via View Selection and Uncertainty Quantification

While the generative model trained from the augmented dataset (section 4.1) produces reasonable multi-view texture maps. The final step of 3D garment texture generation is to merge the multi-view texture maps into the view-independent texture maps on the input mesh. In this paper, we denote this procedure as *baking*. The baking procedure typically introduces unacceptable visual artifacts, such as

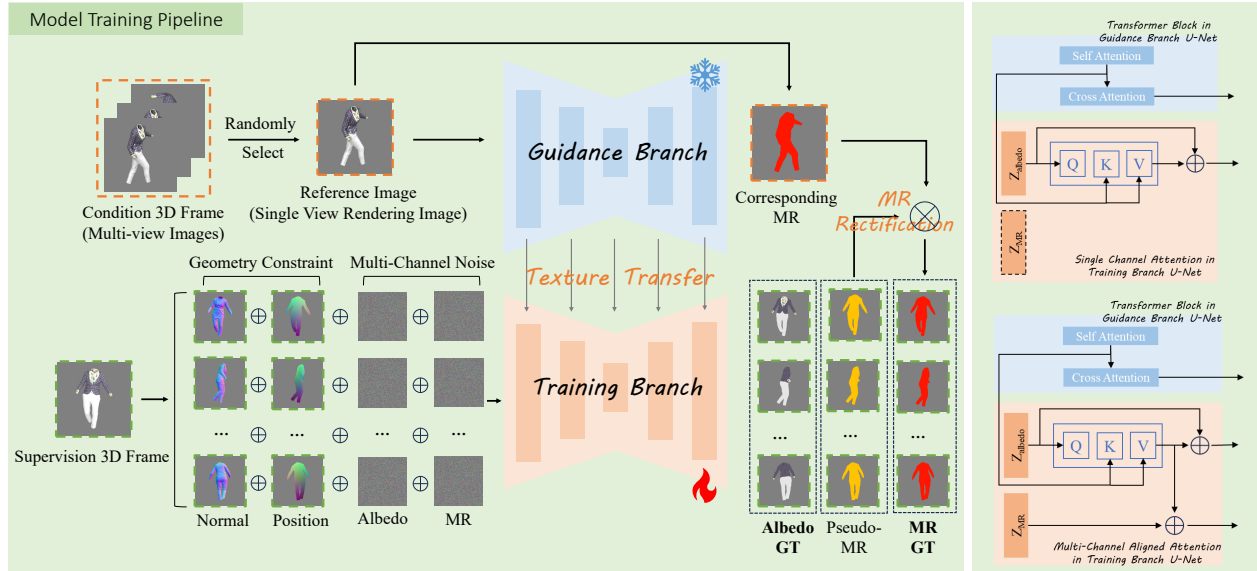


Figure 6. **(left)** In the model training pipeline, we use an illuminated image from the condition 3D frame as the reference image, which is encoded by the guidance branch to perform texture transfer to the training branch. The supervision 3D frame provides multi-view normal and position images as geometric constraints, and its PBR (albedo and MR) maps for supervision. During supervision, MR Rectification is applied to the MR maps of the supervision frame to obtain refined ground truth. **(right)** The architecture of our switchable multi-channel U-Net. Since images generated by Nano Banana cannot maintain consistent MR properties, we apply single-channel attention for albedo-only supervision (right top), and multi-channel aligned attention for joint albedo-MR supervision (right bottom).

blurriness or holes due to incompleteness or inconsistency in the predicted multi-view texture maps.

To fix the artifacts, we first train an Uncertainty Quantification (UQ) model that predicts unsatisfactory areas in the multi-view texture maps. We then propose an UQ-based view-selection algorithm to augment more views at test time. The multi-view texture map predictions in all (original and augmented) views are weighted by the predicted uncertainty to produce robust texture maps on the mesh. The process continues in an iterative way until no more progress can be made. In the following, we elaborate the details of each step.

Uncertainty Quantification. We formulate the Uncertainty Quantification task as an Image Quality Assessment (IQA) problem following Active View Selector (AVS) [25]. Given a single-view texture map $I \in \mathbb{R}^{H \times W \times 3}$, the UQ model adopts a residual net [9] architecture and learns to predict per-pixel uncertainty scores $S \in \mathbb{R}^{H \times W}$ on the multi-view texture maps via supervised learning.

The training data for UQ model is collected via an error simulation procedure: for a ground-truth mesh and its associated texture maps. We render the textured mesh in 10 views (including the front and back views) and use Nano Banana to edit a random view. The image editing procedure is identical to the one defined in section 4.1. We then apply our trained texture generation model using the edited image as the image prompt, and optimize the latent code in the dif-

fusion procedure to enforce an as-close-as-possible match between the predicted front and back view texture maps and the corresponding ground truth texture maps. This forms an optimization problem:

$$\min_z \|\mathcal{G}^{\text{front}}(z) - T_{\text{gt}}^{\text{front}}\|^2 + \|\mathcal{G}^{\text{back}}(z) - T_{\text{gt}}^{\text{back}}\|^2 \quad (4)$$

where z is the latent code, $\mathcal{G}^{(\cdot)}$ maps the latent code to the predicted front- or back-view texture map, $T_{\text{gt}}^{(\cdot)}$ represent the ground truth texture maps in front or back view. Once the problem is solved, we collect the predicted texture maps and the ground truth maps in all views, forming 10 pairs of texture maps. Each pair is automatically associated with a per-pixel uncertainty map computed by comparing each pair of texture maps in terms of SSIM.

View Selection. The view selection algorithm takes a set of predefined candidate views and compute a uncertainty score of each candidate view by averaging the per-pixel uncertainty. The candidate view with the maximal uncertainty score is selected by the algorithm, and we make another inference using the multi-view texture generation model for all views, including the one we just selected. Compared to coverage-based view-selection methods, our UQ models is capable of capturing intermediate baking errors and artifacts during the baking procedure, which leads to superior baking results (See section 5.2).

Implementation details. The iterative process stops when we reach a maximum number of views N_{view} or the newly

selected view has a uncertainty score that is lower than a threshold ϵ . More details are deferred to Appendix.

4.3. Training Architecture

Switchable Multi-channel U-Net. Images generated by Nano Banana often exhibit inconsistent surface reflections, making it unnecessary to optimize their MR attributes during training. To improve training efficiency, we introduce the switchable multi-channel U-Net (see Figure 6, right), which allows the MR channel to be deactivated when only albedo optimization is required.

Training Loss. The model ϵ_θ generates two diffusion noises ϵ_t^{MR} and ϵ_t^{Albedo} conditioned on the input image I at training timestep t . In the multi-channel optimization stage, loss can be designed as:

$$\mathcal{L}_1 = \mathbb{E}_{\epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_t^{MR}\|_2^2 + \|\epsilon - \epsilon_t^{Albedo}\|_2^2 \right] \quad (5)$$

When the switchable multi-channel U-Net deactivates the MR channel, the optimization is performed solely on the albedo channel:

$$\mathcal{L}_2 = \mathbb{E}_{\epsilon \sim \mathcal{N}(0,1), t} \left[\alpha * \|\epsilon - \epsilon_t^{Albedo}\|_2^2 \right] \quad (6)$$

To alternately optimize \mathcal{L}_1 and \mathcal{L}_2 , we set the balancing factor $\alpha = 2$ to smooth the loss curve during training.

5. Experiments

5.1. Experimental Setting

We select over **100K** PBR-textured meshes from Objaverse and **90K** from TexVerse as our general 3D training dataset. In addition, we use more than **150K** diffuse-textured meshes from Bedlam to build 3D Garment Videos for cross-pose augmentation. To enable cross-topology texture learning, we sample frames from the 3D Garment Videos and use the Nano Banana to generate about **50K** edited images.

Our diffusion backbone is based on Stable Diffusion 2.1, trained on 8 H200 GPUs for about 10 days with a batch size of 2 and an image resolution of 512x512. For evaluation, we use both industrial well-rendered images and wild images from DeepFashion2 [8] masked by SAM2 [19] as input prompts, and select industrial 3D meshes together with Hunyuan3D-generated meshes as target geometries.

Baselines. We compare our results under both geometries with several classic or state-of-the-art commercial models, including Paint3D, Hyper3D OmniCraft Texture Generator (Deemos Rodin), Hunyuan3D (Tencent), and Meshy 6 Preview (Meshy AI). Among them, Paint3D outputs only albedo maps, while the others generate both albedo and MR maps. Except for Paint3D, all model results are sourced from their official websites via paid access or granted credits.

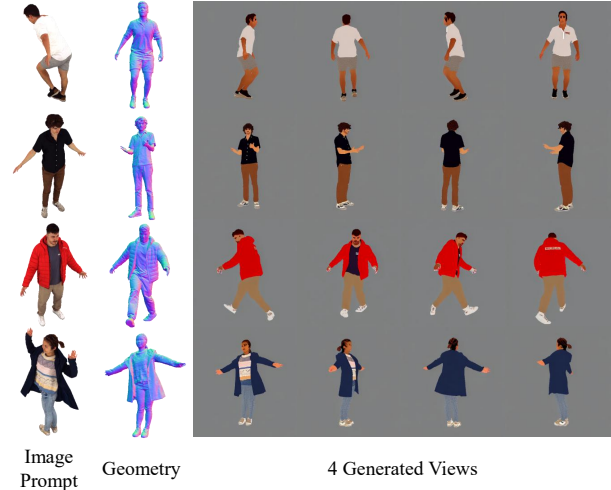


Figure 7. We use 4D-Dress as our test set to evaluate the model’s ability to generate textures across different human poses while maintaining consistent garment topology.

5.2. Qualitative Evaluation

Generation for Industrial Meshes. We perform experiments with two categories of input images in Figure 8 (left). The first group (top), named as well-rendered images, is rendered from existing 3D assets and are aligned with the target meshes in pose. This setting is specifically verify the effectiveness of the **cross-topology texture generation** in NI-TeX. The second group (bottom) of input image prompts is sourced from DeepFashion2, comprising real-world and web-collected images whose poses and geometries exhibit large discrepancies from the target 3D meshes. Despite this gap, NI-TeX successfully captures the correct texture information.

Generation for Generated Meshes. We generate meshes using Hunyuan3D text prompts. Compared to industrial meshes, these generated meshes exhibit more folds, and their poses and geometries present significant gaps relative to the target images from DeepFashion2. Despite this, NI-TeX produces textures that not only closely align with simple sewing patterns, but also faithfully preserve structural garment details (see Figure 8, right).

Effectiveness of Cross-pose Texture Generation. The 4D-Dress dataset [23] provides per-frame 3D captures of real human motion, though its texture maps often contain baked shadows. We use 4D-Dress to maintain consistent garment topology across poses, enabling a focused evaluation of NI-TeX’s cross-pose texture generation (See Figure 7). For each sequence, we randomly sample a mesh and viewpoint to render the input image prompt, and use another pose of the same subject as the target mesh.

Baking Strategy. We compare our baking strategy with the coverage-based baking strategy on multiple complex gar-

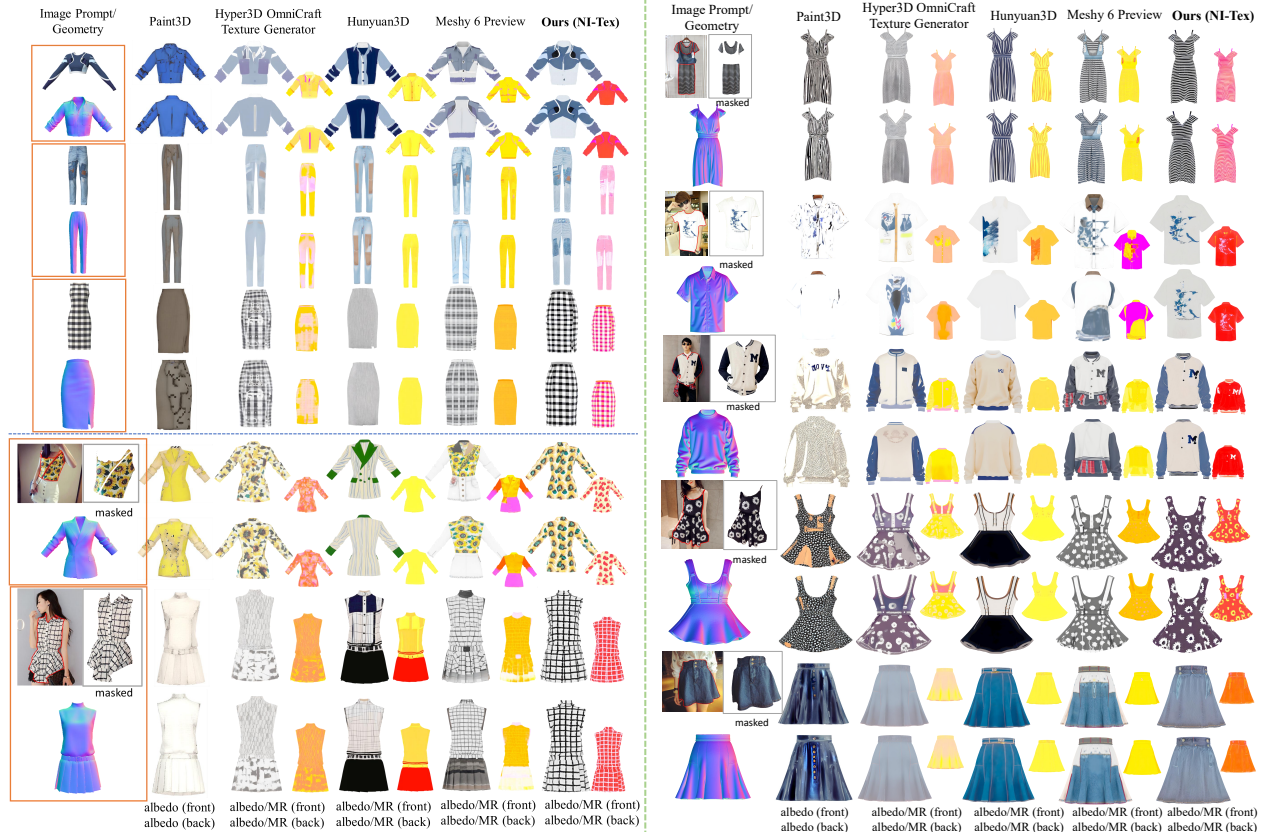


Figure 8. **(left)** NI-Tex results using industrial well-rendered images (top) and DeepFashion2 wild images masked by SAM2 (bottom) as prompts, with industrial meshes as targets. NI-Tex produces realistic industry-level PBR materials and aligns textures across complex garment surfaces. The top also shows robust cross-topology generation under consistent poses. **(right)** NI-Tex results using DeepFashion2 wild images as prompts with generated meshes as targets, demonstrating reliable PBR material generation.

ment meshes qualitatively in Figure 5. We quantitatively compare the two baking methods using metrics i^{UQ} and i^{cvg} .

$$i^{UQ} = \arg \max_i \sum_{p \in \mathcal{P}_i} U_p, \quad i^{cvg} = \arg \max_i \sum_{p \in \mathcal{P}_i} V_p \quad (7)$$

where \mathcal{P}_i is the set of pixels for view i , U_p is the predicted uncertainty of pixel p and $V_p = 1$ only if p is not covered by any views. Due to space limitations, quantitative results and more qualitative experiments are provided in the appendix.

5.3. Quantitative Evaluation

To quantitatively verify the effectiveness of NI-Tex in Table 1, we adopt the Kernel Inception Distance (KID) metric, which provides an unbiased estimate similar to FID, and also report FID for reference. We randomly select 10 textured objects from both industrial and generated meshes, render them from multiple viewpoints, and perform KID experiments using a fixed set of 42 random seeds.

Table 1. We compare Paint3D, Hyper3D OmniCraft Texture Generator, Hunyuan3D, Meshy, and NI-Tex using 10 image prompts and multiple rendered views of 10 textured objects in terms of KID and FID. NI-Tex achieves the best performance on both metrics.

Method	Paint3D	Hyper3D	Hunyuan3D	Meshy	Ours(NI-Tex)
KID ↓	0.0695	0.0471	0.0528	0.0383	0.0364
FID ↓	293.45	285.45	272.34	246.39	237.59

6. Conclusions

In summary, we propose a novel technique, NI-Tex, for non-isometric image-based garment texture generation. By constructing 3D Garment Videos and applying Nano Banana to edit images from videos, we create image-garment training pairs that enable NI-Tex to perform stable cross-pose and cross-topology texture generation. To produce complete textures on the mesh, we employ iterative baking with view selection and uncertainty quantification. Extensive experiments demonstrate that the generated materials are suitable for both industrial graphics design and commercial 3D generation models.

References

- [1] Michael J Black, Priyanka Patel, Joachim Tesch, and Jinlong Yang. Bedlam: A synthetic dataset of bodies exhibiting detailed lifelike animated motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8726–8737, 2023. 2, 4
- [2] Dave Zhenyu Chen, Yawar Siddiqui, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. Text2tex: Text-driven texture synthesis via diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 18558–18568, 2023. 2
- [3] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22246–22256, 2023. 3
- [4] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36:35799–35813, 2023. 3
- [5] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13142–13153, 2023. 2, 3
- [6] Kangle Deng, Timothy Omernick, Alexander Weiss, Deva Ramanan, Jun-Yan Zhu, Tinghui Zhou, and Maneesh Agrawala. Flashtex: Fast relightable mesh texturing with lightcontrolnet. In *European conference on computer vision*, pages 90–107. Springer, 2024. 3
- [7] Daiheng Gao, Xu Chen, Xindi Zhang, Qi Wang, Ke Sun, Bang Zhang, Liefeng Bo, and Qixing Huang. Cloth2tex: A customized cloth texture generation pipeline for 3d virtual try-on. In *2024 international conference on 3D vision (3DV)*, pages 602–611. IEEE, 2024. 3
- [8] Yuying Ge, Ruimao Zhang, Xiaogang Wang, Xiaoou Tang, and Ping Luo. Deepfashion2: A versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5337–5345, 2019. 7
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [10] Zebin He, Mingxin Yang, Shuhui Yang, Yixuan Tang, Tao Wang, Kaihao Zhang, Guanying Chen, Yuhong Liu, Jie Jiang, Chunchao Guo, et al. Materialmvp: Illumination-invariant material generation via multi-view pbr diffusion. *arXiv preprint arXiv:2503.10289*, 2025. 3
- [11] Xin Huang, Tengfei Wang, Ziwei Liu, and Qing Wang. Material anything: Generating materials for any 3d object via diffusion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 26556–26565, 2025. 3
- [12] Team Hunyuan3D, Shuhui Yang, Mingxin Yang, Yifei Feng, Xin Huang, Sheng Zhang, Zebin He, Di Luo, Haolin Liu, Yunfei Zhao, et al. Hunyuan3d 2.1: From images to high-fidelity 3d assets with production-ready pbr material. *arXiv preprint arXiv:2506.15442*, 2025. 2, 3
- [13] Seungyong Lee and Jeong-gi Kwak. Voost: A unified and scalable diffusion transformer for bidirectional virtual try-on and try-off. *arXiv preprint arXiv:2508.04825*, 2025. 3
- [14] Ming Li, Hui Shan, Kai Zheng, Chentao Shen, Siyu Liu, Yanwei Fu, Zhen Chen, and Xiangru Huang. Reweaver: Towards simulation-ready and topology-accurate garment reconstruction. *arXiv preprint arXiv:2601.16672*, 2026. 3
- [15] Xuan Li, Chang Yu, Wenxin Du, Ying Jiang, Tianyi Xie, Yunuo Chen, Yin Yang, and Chenfanfu Jiang. Dress-1-to-3: Single image to simulation-ready 3d outfit with diffusion prior and differentiable physics. *ACM Transactions on Graphics (TOG)*, 44(4):1–16, 2025. 3
- [16] Jinxi Liu, Zijian He, Guangrun Wang, Guanbin Li, and Liang Lin. One model for all: Partial diffusion for unified try-on and try-off in any pose. *arXiv preprint arXiv:2508.04559*, 2025. 3
- [17] Aymen Mir, Thiemo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7023–7034, 2020. 3
- [18] Sai Raj Kishore Perla. Easi-tex: edge-aware mesh texturing from single-image. 2024. 3
- [19] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 7
- [20] Elad Richardson, Gal Metzer, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. Texture: Text-guided texturing of 3d shapes. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–11, 2023. 2
- [21] Nikolaos Sarafianos, Tuur Stuyck, Xiaoyu Xiang, Yilei Li, Jovan Popovic, and Rakesh Ranjan. Garment3dgen: 3d garment stylization and texture generation. In *2025 International Conference on 3D Vision (3DV)*, pages 1382–1393. IEEE, 2025. 3
- [22] Riza Velioglu, Petra Bevandic, Robin Chan, and Barbara Hammer. Tryoffdiff: Virtual-try-off via high-fidelity garment reconstruction using diffusion models. *arXiv preprint arXiv:2411.18350*, 2024. 3
- [23] Wenbo Wang, Hsuan-I Ho, Chen Guo, Boxiang Rong, Artur Grigorev, Jie Song, Juan Jose Zarate, and Otmar Hilliges. 4d-dress: A 4d dataset of real-world human clothing with semantic annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 550–560, 2024. 7
- [24] Yuanhao Wang, Cheng Zhang, Gonçalo Frazão, Jinlong Yang, Alexandru-Eugen Ichim, Thabo Beeler, and Fernando De la Torre. Garmentcrafter: Progressive novel view synthesis for single-view 3d garment reconstruction and editing. *arXiv preprint arXiv:2503.08678*, 2025. 3
- [25] Zirui Wang, Yash Bhalgat, Ruining Li, and Victor Adrian Prisacariu. Active view selector: Fast and accurate active

- view selection with cross reference image quality assessment. *arXiv preprint arXiv:2506.19844*, 2025. 6
- [26] Xudong Xu, Zhaoyang Lyu, Xingang Pan, and Bo Dai. Mat-labrer: Material-aware text-to-3d via latent brdf auto-encoder. *arXiv preprint arXiv:2308.09278*, 2023. 3
- [27] Kim Youwang, Tae-Hyun Oh, and Gerard Pons-Moll. Paint-it: Text-to-texture synthesis via deep convolutional texture map optimization and physically-based rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4347–4356, 2024. 3
- [28] Xianfang Zeng, Xin Chen, Zhongqi Qi, Wen Liu, Zibo Zhao, Zhibin Wang, Bin Fu, Yong Liu, and Gang Yu. Paint3d: Paint anything 3d with lighting-less texture diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4252–4262, 2024. 3
- [29] Cheng Zhang, Yuanhao Wang, Francisco Vicente, Chenglei Wu, Jinlong Yang, Thabo Beeler, and Fernando De la Torre. Fabricdiffusion: High-fidelity texture transfer for 3d garments generation from in-the-wild images. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–12, 2024. 3
- [30] Yuqing Zhang, Yuan Liu, Zhiyu Xie, Lei Yang, Zhongyuan Liu, Mengzhou Yang, Runze Zhang, Qilong Kou, Cheng Lin, Wenping Wang, et al. Dreammat: High-quality pbr material generation with geometry-and light-aware diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4):1–18, 2024. 3
- [31] Yibo Zhang, Li Zhang, Rui Ma, and Nan Cao. Texverse: A universe of 3d objects with high-resolution textures. *arXiv preprint arXiv:2508.10868*, 2025. 2, 3