Stablediff-Fabric: Controllable Generation of Ultra-Detailed Textile Images via Diffusion Models and Microstructure-Annotated Dataset

Jing Ma
Donghua University
jingma@dhu.edu.cn

Siyu Chen
Donghua University
230613202@mail.dhu.edu.cn

Yunhan Shang
Donghua University
230613210@mail.dhu.edu.cn

Yun Su*
Donghua University
suyun150@dhu.edu.cn

Abstract

The rise of the metaverse and sustainable fashion has accelerated demand for virtual garments that are both photorealistic and easily editable. Existing approaches, whether based on physics-based modeling or diffusion models, struggle to achieve realism and real-time controllability within a unified framework. In this work, we present Stablediff-Fabric, an interactive text-to-image system designed for high-quality fabric generation. The method synthesizes fabric appearances in seconds and enables intuitive editing of parameters such as weave structure, yarn density, and color. Specifically, we first construct a dedicated dataset comprising high-resolution fabric images annotated with structural, chromatic, and material properties and then use this dataset to finetune a pre-trained SDXL model. Quantitative and qualitative evaluations demonstrate that our model produces high-fidelity textures, captures fine structural details such as float count, and outperforms state-of-the-art baselines in realism and controllability. These results highlight the potential of our approach to enhance digital garment modeling pipelines and lower the technical barriers for virtual fashion design.

Keywords: ***Digital Clothing, Diffusion Models, Textile Image Generation, Microstructure Annotation, Ultra-Detailed Fabric***

1. Introduction

In recent years, the rise of the metaverse has increased the demand for advanced digital modeling technologies. As a key medium for visual expression, the simulation of clothing textures plays a critical role in digital environments. With market developments pushing for more realistic and diverse simulations of digital clothing, the demand for high-quality digital garments has grown substantially. On the market side, the growth of cloud-based fashion events and virtual fashion shows has intensified the need for companies to deliver clothing visuals that closely resemble real-world scenarios, strengthening consumer engagement. On the supply side, designers increasingly seek modeling assets that provide realistic fabric textures at lower cost and with easier accessibility. In addition, the integration of digital clothing significantly reduces the waste of fabric, energy and labor associated with traditional pattern-making and physical garment revisions. This digital approach has been shown to reduce fabric waste by more than 30% and energy consumption by an average of 15%, significantly contributing to the sustainable development of the fashion industry [2].

Current mainstream 3D garment modeling software, such as Vstitcher, CLO3D, and Style3D, offers comprehensive simulation capabilities for overall garment appearance. Their fabric libraries cover a wide range of conventional textiles and support combined adjustments of color and material properties. However, the quality of the rendering still does not meet expectations. The primary limitation is the lack of high-fidelity representation of fabric microstructure and surface properties.

The complex structure of fabrics makes their appearance difficult to simulate. As flexible materials composed of interlaced, knotted, or bonded yarns, fabrics exhibit intricate physical structures and visual properties. Factors such as fiber types, warp-weft interlacing methods, and yarn densities in both directions significantly affect the appearance of the fabric [18]. This microscopic textural complexity is the primary cause of limitations in fabric simulation. To better simulate this textural realism and improve user experience, prior research has focused on fabric appearance through texture tiling and warp-weft structure simulations. From early physical synthesis algorithms like Efros-Leung [6] and image quilting [5, 58], to deep learning-based approaches like

*Corresponding author.

TextureGAN [52, 54], which offers more controllable generation, texture synthesis algorithms have continually optimized the balance between efficiency, quality, and controllability. However, the results still tend to fall short of achieving true fabric realism. Shuang Zhao *et al*. [57] pioneered a scanning-based approach to establish the initial microscopic appearance of the fabric, providing a physical foundation for texture realism. Over time, fabric appearance modeling methods have advanced toward higher resolution and enhanced realism. Wang [49] laid the groundwork for further improvements in fabric texture generation in his classical work. Building on this foundation, we aim to develop a fabric image dataset based entirely on real-world photography. This dataset will be supplemented with comprehensive structural and appearance data obtained through multidimensional measurements, which is designed to enhance fabric appearance simulation, particularly in terms of detailed structure.

Users require a wide variety of fabric materials during clothing modeling, making editability and accessibility of fabric necessary. Previously, images synthesized through physics-based modeling or generated by neural networks such as GANs [37, 8, 20] proved difficult to edit. Users were often forced to repeatedly adjust parameters, spending hours experimenting to get the desired material, which hindered instant access to fabric resources. In recent years, diffusion models [15, 41] significantly advanced the controllability and fidelity of image generation, pushing these aspects to a next level. Extensive research on the generation of garment patterns and textures [12, 55, 23] has made considerable progress with the support of the diffusion model. Current 3D clothing models can already simulate basic textures of various materials, exhibit flexible physical properties, and adjust color with ease. However, due to the lack of authentic fabric images and comprehensive physical parameter inputs, fabric models still fall short in providing editable interactions for detailed structures such as weft density. These factors have a substantial impact on the fabric representation of garments [8].

To address these challenges, we introduce Stable-Diff Fabric, a textile image generation method based on the Stable diffusion model. Stablediff-Fabric allows users to interactively generate high-definition fabric assets through natural language, offering professional editing options, including specific fabric structures, types, colors, and warp/weft densities. First, we scanned and collected more than 300 high-resolution fabric sample images in our lab as the primary data set, and then employed manual annotation to create precise textual descriptions of these fabric samples. To enhance the model's capability, we designed data augmentation techniques specifically tailored for fabric images, expanding the original sample size to about 3,000. By fine-tuning the U-Net architecture of the pre-trained SDXL model, we obtained fabric material images suitable for clothing modeling and texturing applications. Finally, we propose a comprehensive evaluation methodology to assess the results from multiple dimensions, including structure, color, and warp/weft density. In summary, our key contributions are

- We propose a text-driven approach for generating textile images, allowing professional artists to create accurate fabric images on demand.
- To the best of our knowledge, we introduce the first ultra-detailed textile dataset comprising around 3k images with microstructural and appearance annotations obtained using professional instrumentation and processed using a customized augmentation pipeline.
- We develop a comprehensive set of metrics for assessing fabric appearance quality.

2. Related Work

2.1. Fabric Synthesis based on Physical Models

Early research on fabric texture synthesis mainly relied on physically based simulations. Zhao [58] introduced a structure-driven approach, while Shinohara [40] reconstructed yarns in 3D from X-ray scans and Zhao [57] generated voxel-based models from micro-CT data. Although computationally expensive, these works laid the foundation for digital fabric reconstruction. In parallel, patch-based methods emerged. Efros and Freeman's Image Quilting [5] demonstrated that visually coherent textures can be synthesized by stitching together small patches. Kwatra *et al*. [21] improved boundary quality via GraphCut, and Cohen *et al*. [3] applied Wang tiles [46] to efficiently generate non-periodic textures, opening new directions for fabric structure synthesis. Montazeri *et al*. [28, 27] further proposed a ply-level appearance model that represents both woven and knitted fabrics under varying lighting conditions. While this significantly improved realism, it offered limited flexibility for editing. Overall, these early approaches advanced virtual fabric modeling but remained inefficient, lacked interactivity, and depended heavily on physical samples.

Building upon this foundation, Kai Schroeder *et al*. [37] proposed "Visual Prototyping of Woven Cloth " that generates physically plausible initial fabric designs from parameters obtained through image inversion, while also supporting precise parameter editing. This approach partially overcomes the non-interactivity limitations of previous methods, though it remains somewhat complex in operation. Similarly, Deng [4] achieved control over both appearance and physical parameters. However, none of these approaches fully addressed the problem of reliance on physical fabric samples.

To further enhance realism, Wu [51] and Zhang [59] refined yarn-level modeling and rendering, achieving high structural fidelity. However, their methods lacked generalizability across diverse materials. As a result, most current CAD systems still depend on texture-based synthesis pipelines tied to real fabric samples, often producing rigid and mechanical visual effects. Our approach operates without physical samples, can generate a wide range of fabric types, and delivers superior realism in fabric texture rendering compared to traditional texture-synthesis methods.

2.2. Fabric Texture Generation based on Neural Networks

With the advent of deep learning, neural networks and generative models have been widely applied to texture synthesis. Venkatraman *et al*. [45] combined VGG19 with semantic segmentation to capture garment textures and boundaries for AI-driven clothing transfer. He *et al*. [13] realized single-image garment transfer but lost much of the fine-grained texture detail. Xue *et al*. [54] introduced an Inception-based convolutional network that improved material realism across multiple aspects, though it remained limited in free-form editing flexibility.

The introduction of Generative Adversarial Networks (GANs) [9] significantly advanced generative quality, particularly for garment silhouettes and photorealistic rendering. Wang *et al*. [47] translated fashion sketches into photorealistic clothing images, while Wang *et al*. [48] employed a physically grounded shading model to train GANs for color recovery in faded fabrics. Adeyokunnu *et al*. [1] generated richly detailed African textile designs, and several works [14, 32, 7, 10] explored GAN-based material texture synthesis, though without explicitly modeling microscopic fabric structures. Seamless texture generation has also progressed, with SeamlessGAN [33] and TipGAN [11] achieving infinite tiling, while Huang [20] directed GANs toward realistic suit textures. More specialized efforts, such as MSEmbGAN [19], improved embroidery realism and color fidelity. Despite these advances, most methods remain tailored to specific fabric categories or weave patterns, with limited validation across diverse clothing types.

Beyond 2D synthesis, numerous studies have investigated 3D garment generation with integrated texture modeling [39, 22, 42, 35]. While these approaches produce convincing macroscale patterns, they generally fail to capture the structural complexity of real textiles. Overall, convolutional and GAN-based methods improve realism at the image level but remain constrained by limited texture precision and insufficient support for user-driven editing.

2.3. High-Precision Texture Image Generation based on Diffusion Models

In recent years, advances in generative modeling have opened new directions for fabric rendering. Multimodal generation provides more flexible ways to obtain material appearances. Diffusion models, including DDPM [15] and DDIM [41], established the foundation for systems such as Stable Diffusion [34], MidJourney, and video generation architectures [16]. Among them, Stable Diffusion is widely adopted due to its open-source availability and high controllability.

Several works have adapted diffusion-based models for garment and fabric generation. DressCode [12] fine-tunes pre-trained latent diffusion models to produce clothing prints and PBR fabric textures, capturing some structural characteristics but constrained to repetitive patterns. GarmentDiffusion [23] generates vectorized 3D sewing patterns with centimeter-level accuracy from multimodal inputs. FabricDiffusion [56] applies diffusion to garment pattern generation and additionally corrects deformations in texture images, enabling texture transfer to arbitrary 3D garment shapes. Diffavatar [24] supports end-to-end generation from a single scan to a complete physical garment. With the exception of DressCode, however, these approaches do not directly address fine-grained fabric texture generation.

From the perspective of multimodal interactive generation, ControlMAT [43], MatFusion [36], and MatFuse [44] extend diffusion models to controllable material rendering, yet none explicitly target fabric textures. ReflectanceFusion [53] and Material Picker also demonstrate fabric-related assets, but the generated textures remain relatively coarse. In contrast, our method produces high-precision, fully editable fabric textures with realistic surface detail.

3. Methods

This approach targets breakthroughs in two aspects. First, improving the visual realism of generated fabric images. Second, lowering the user adoption barrier through text-based interaction. To achieve these goals, we designed the methodology, constructed a sample database, and selected foundational models.

3.1. Ultra-Detailed Textile Generation

Our method is designed to generate various fabric material images from textual descriptions while ensuring their authenticity and editability, we first selected pre-trained generation models and built a generation framework.

3.1.1 Stable diffusion

First, we note that diffusion models [15] [41] perform exceptionally well in image generation tasks. By iteratively

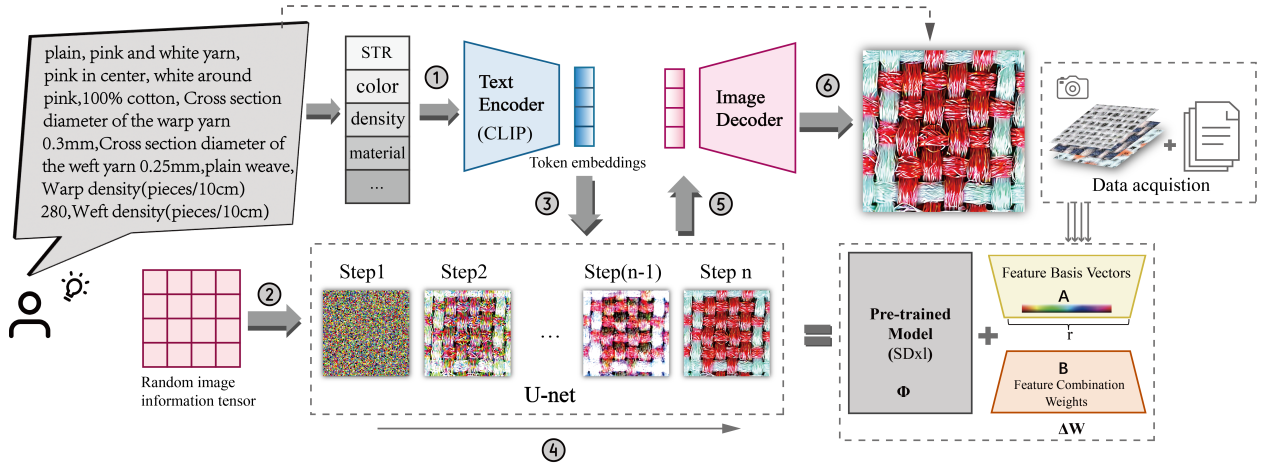

Figure 1. **Overview of Stablediff-Fabric pipeline.** (1) User Input and Text Encoding: Users submit textual descriptions containing fabric physical parameters. The CLIP text encoder converts these descriptions into token embeddings. (2) Random Initialization: Generate a noise tensor as initial image information for the diffusion model to progressively denoise and optimize. (3) Feed text embeddings into the U-Net architecture. (4) U-Net progressively fuses noisy images with text vectors across multiple time steps to generate image features. (5) Output image features. (6) Decode latent features into images.

denoising, they produce high-quality images with rich detail and realism. Unlike traditional Generative Adversarial Networks, which suffer from mode collapse, diffusion models offer more stable training and greater sample diversity. They also provide higher fidelity, avoiding monotonous outputs. Moreover, they offer remarkable flexibility and controllability, allowing for autonomous image generation and seamless integration with multimodal inputs like text.

Building on diffusion models, Stable Diffusion [34] improves computational efficiency by training a first-stage autoencoder to compress images into a latent space, while preserving perceptual detail. As shown in Figure 1, in Step 1, natural language inputs are tokenized and encoded by a text encoder into embeddings, which serve as text conditions. In practice, Stable Diffusion commonly adopts a frozen CLIP text encoder. The U-Net acts as the denoising backbone, operating in the latent space with cross-attention to the text embeddings and random noise. Diffusion consists of two processes: a forward step and a reverse step. In the forward process, Gaussian noise is progressively added to the latent variable z_0, producing a noisy latent z_t that approximates an isotropic Gaussian distribution. In Figure 1, Step 4 illustrates the reverse denoising process, in which the U-Net iteratively predicts and removes noise. This procedure can be formally expressed as Equation (1):

$$z_{t-1} = \frac{1}{\sqrt{1-\beta_t}}\left(z_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(z_t, t, y)\right) + \sigma_t\epsilon, \quad (1)$$

where β_t refers the variance schedule at step t. $\bar{\alpha}_t$ is the cumulative product of $(1-\beta_i)$ up to step t. $\epsilon_\theta(\mathbf{z_t}, t)$ refers a neural network that predicts the noise component ϵ in $\mathbf{z_t}$ conditioned on the time step t. σ_t and ϵ determines the random depth of the sampling process.

The loss function presented in Equation (2), denoted as $L(\theta)$, is the foundational training objective for the diffusion model. It operates by compelling the U-Net to predict the noise component ϵ that was added to the original image $\mathbf{x}_0$ at a random timestep t during the forward diffusion process. The core of this loss is the L2-norm (Mean Squared Error) between the true noise ϵ and the noise predicted by the model ϵ_θ, calculated as an expectation over random timesteps t, original images $\mathbf{x}_0$, and noise samples ϵ:

$$L(\theta) := \mathbb{E}_{t,\mathbf{x}_0,\boldsymbol{\epsilon}}\left[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}, t)\|^2\right] \quad (2)$$

By minimizing this loss, the U-Net progressively learns the complex reverse process of denoising.

Finally, in Steps 5 and 6, the decoder maps the final denoised latent representation back into the image space, producing the generated fabric image. Since the entire diffusion process is performed in the latent space, the method substantially reduces computational cost and accelerates generation compared to previous models.

3.1.2 Choice of Pretrained-Model

We adopt SDXL [31] as the pre-trained backbone. SDXL employs a two-stage pipeline, consisting of a base latent diffusion model and a refinement stage, which enables coherent composition and detailed synthesis even at 1024×1024 resolution. It integrates dual text encoders (OpenCLIP ViT-bigG and CLIP ViT-L), whose embeddings are concate-

nated for conditioning, thereby improving prompt alignment. The combination of high-capacity latent diffusion and refinement enhances the recovery of fine-scale structures and high-frequency details, which is beneficial for fabric texture generation.

3.1.3 Low-Rank Adaptation

Regarding fine-tuning, Hu *et al.* [17] proposed *Low-Rank Adaptation (LoRA)*, which freezes the original model weights and injects trainable low-rank parameter matrices into each layer. In the context of image generation, LoRA is commonly applied to the *cross-attention layers* of Stable Diffusion, where it introduces low-rank updates $\Delta W = A \cdot B$ to the existing weight matrices. In our ultra-high-definition textile image generation task, we adopt the same scheme. As shown in the right portion of Figure 1, paired image-text data serve as training samples; the low-rank matrices A, B within cross-attention modules are optimized to learn the mapping from textual prompts to textile visual features. By only updating A and B, the number of trainable parameters is drastically reduced, significantly lowering computational and memory demands compared to full model fine-tuning—making LoRA especially suitable for resource-limited environments.

3.2. Collection of Fabric Data

The quality of data samples significantly impacts the performance of generative models. To maximize the realism and diversity of the generated textures, we utilized both laboratory-collected samples and image data augmentation techniques.

3.2.1 Data Collection and Annotation

To ensure structural diversity in training data, we collected multiple fabrics from the textile market, covering the three main types: plain, twill, and satin. The twill samples include variations such as two-up-two-down and two-up-one-down twills, while the satin samples comprise three-up-one-down and four-up-one-down satins. The color of fabric exceed 50 varieties. Materials include common fibers such as cotton, linen, silk, polyester, and spandex, as well as less common ones such as acetate, polyurethane, and Lycra. Each fabric was first examined under a microscope for preliminary screening and classified into plain, twill, or satin categories. To ensure clear imaging and facilitate the learning of fabric structures, we selected fabrics with minimal surface fuzziness. After screening, 152 plain weaves, 85 twills, and 64 satins were chosen, yielding a total of 301 samples. Each fabric was cut into $4cm \times 4cm$ specimens for subsequent measurements of fundamental physical properties.

Based on the known influencing factors of woven fabric structures [8], we measured the micro-structure of each textile image: weave structure (plain, twill, satin), yarn density, yarn count, number of full weave yarns in satin, float count, yarn material, other than their basic properties such as color and areal weight. Fabric images were captured using the $PlanC\ N-4x/0.10-\infty/-/FN22$ lens of a fiber fineness meter shown in Figure 9. For fabrics with distinct textures and thick yarns, warp and weft densities were measured using a fabric density mirror $Y511B$ shown in Figure 10. For fabrics with high warp and weft densities, following the same density measurement principle, samples were placed on the fiber fineness meter and covered with a thin scaled ruler. After capturing an image screenshot, warp and weft yarn counts were recorded and converted to the unified unit of roots/10 cm. Additionally, the cross-sectional diameters of warp and weft yarns were measured using $ImageJ$ software shown in Figure 11. And the areal weight of fabric was tested by electronic balance $JA5003N$ shown in Figure 12. All the label names and measurement methods are summarized in Table 1.

3.2.2 Data Augmentation

To further improve the ability of our model, we applied data augmentation to the 301-sample dataset, as shown in shown in Figure 2.

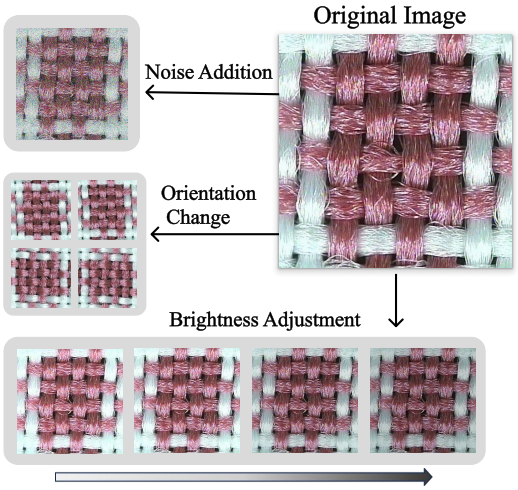

Figure 2. **Illustration of Image Data Augmentation**

First, images were manually reviewed for clarity, occlusion, and label consistency. The dataset was then expanded more than tenfold using augmentation techniques including rotation, mirroring, noise perturbation, and brightness adjustment. The final dataset contained 1,492 plain weave samples (50.0%), 849 twill weave samples (28.5%), and

Table 1. **Data collection of textile images and annotation of their labels**

Data/Label Name	Measurement Method	Numeric Type/Range
Fabric picture	Fiber fineness meter shooting	image (512×512)
Fabric structure	Manual judgment	Plain weave / Twill weave / Satin weave
Color	Manual judgment	–
Warp density (ends/10 cm)	Experimental Testing	110–1875
Weft density (ends/10 cm)	Experimental Testing	100–1300
Warp yarn cross-sectional diameter (mm)	ImageJ software measurement	0.0439–0.4001
Weft yarn cross-sectional diameter (mm)	ImageJ software measurement	0.0486–0.4303
Number of complete yarns in satin weave	Experimental Testing	–
Twist count	Experimental Testing	–
Yarn material	Provided by the supplier	–
Areal weight (g/m^2)	Provided by the supplier / Electronic balance measurement	17.875–419.375
Other descriptions (fabric name, process, visual effect, etc.)	Provided by the supplier / Manual judgment	–

640 satin weave samples (21.5%), increasing the total from 301 to 2,981 samples. All augmented samples were standardized to a resolution of 512×512 pixels, with detailed augmentation parameters recorded to enable reversible operations and subsequent ablation studies.

3.3. Evaluation Framework Design

We designed five metrics from three aspects to evaluate the quality of the fabric structure images we generate, covering both overall quality assessment and detailed perspective evaluation. These metrics will be applied to compare our model with the image generation model DressCode [12].

3.3.1 Overall Quality Assessment

We selected two common metrics, PSNR [29] and SSIM [50], to evaluate the overall quality of the generated fabric images. PSNR (Peak Signal-to-Noise Ratio) is a widely used objective measure of image quality. It quantifies distortion by computing pixel-level errors between the original and generated images. This metric primarily reflects signal fidelity and noise control, with a particular emphasis on pixel value accuracy. A higher PSNR value indicates smaller differences between the generated and original images, implying lower brightness loss and potentially clearer contours and details. In this study, PSNR is used to evaluate the accuracy of generated fabric images, as pixel-level precision is essential to ensure that fine structural details of fabrics are preserved.The equation is as (3) (4).

$$PSNR = 10 \cdot log_{10}(\frac{MAX_I^2}{MSE})(dB) \tag{3}$$

PSNR is measured in dB. MAX_I represents the maximum possible value of a pixel, and the final PSNR value is obtained through logarithmic transformation.

$$MSE = \frac{1}{MN}\sum_{i=1}^{M}\sum_{j=1}^{N}[I(i,j) - K(i,j)]^2 \tag{4}$$

MSE is the mean squared error at the pixel level between two images, serving as the basis for calculating PSNR.

SSIM (Structural Similarity Index) is a full-reference image quality metric used to assess the perceptual similarity between a test image and its reference. It combines three components—luminance similarity, contrast similarity, and structural similarity—to better reflect human visual perception than simple pixel-wise error metrics. In this work, we use SSIM to approximate subjective perceptual fidelity of generated fabric images relative to their references.

The design of SSIM is based on three independent comparisons: Comparison of mean luminance (μ_x, μ_y) between the two images, reflecting the impact of lighting variations.Comparison of standard deviation (σ_x, σ_y), measuring the intensity of light–dark differences within the images. Assessment of pixel-to-pixel dependency similarity through covariance (σ_{xy}), capturing structural features such as edges and textures.

$$\text{SSIM}(x,y) = \frac{(2\mu_x\mu_y + C_1)}{(\mu_x^2 + \mu_y^2 + C_1)} \cdot \frac{(2\sigma_{xy} + C_2)}{(\sigma_x^2 + \sigma_y^2 + C_2)} \tag{5}$$

3.3.2 Color Evaluation

For color similarity evaluation, we utilize the CIELAB color space [30]. This property ensures that the measured color difference aligns closely with human visual perception, making it particularly suitable for assessing textile materials. The primary color sets, $\{c_A\}$ and $\{c_B\}$, are first extracted from the images and represented in the CIELAB

space, where each color is defined by its L^* (lightness), a^* (green–red axis), and b^* (blue–yellow axis) values.

The overall color similarity between the two sets is then quantified by calculating the CIEDE2000 color difference metric for all pairs of primary colors [38]. The mean of the minimum pairwise ΔE_{00} values reflects the perceptual color difference, providing a more accurate and standardized measure of color similarity. The core calculation can be represented as (6):

$$\overline{\Delta E} = \frac{1}{N} \sum_{i=1}^{N} \min_{j} \Delta E_{00}(c_{A_i}, c_{B_j}) \tag{6}$$

3.3.3 Yarn Density and Fineness Evaluation

Automated evaluation of the warp and weft density and fineness of generated fabrics presents a significant challenge. To address this, we developed a method that detects warp/weft and void density based on wavelet transform [25], automating the measurement of these parameters. First, an algorithm batch-processes the void density across all images, screening out low-quality samples. For images whose void density falls within acceptable bounds, we apply warp/weft density detection to assess the accuracy of the generated fabric structure. For automated warp/weft density detection, we primarily employed projection analysis [26] and frequency domain computation. Decolorized fabric images were projected onto the warp and weft pixel points, with key frequency bands extracted and grayscale values summed to generate density waveform diagrams. The number of peaks in these diagrams corresponds to the yarn count in the image direction.

To mitigate errors from relying solely on automated evaluation, we also conduct manual image verification.

4. Experiment

We comprehensively evaluate our model. We first present results on multi-modal fabric generation, followed by editing capabilities such as fineness density, color brightness, and material control via text-to-image prompts. We then compare our method with DressCode both qualitatively and quantitatively. Finally, we conduct an ablation study to assess the contributions of each component.

4.1. Set up

All experiments were conducted on a single NVIDIA RTX 4090 GPU. We adopted Prodigy as the primary optimizer, which employs a dynamic learning rate strategy. This enabled us to set an initial learning rate of 1.0 without manual scheduling, ensuring stable convergence. For comparison, we also experimented with the Adam optimizer combined with a scheduler, which adjusted the learning rate according to the training loss curve. The dataset was divided into training, validation, and test sets in a 70%:15%:15% ratio, allowing the model to learn material patterns while preserving independent samples for evaluation. Training was performed on 512×512 pixel inputs for approximately 100,000 steps, requiring about 10 hours. During inference, SDXL served as the backbone model, while the official pre-trained VAE was employed to enhance high-frequency details. The DDPM2+ sampler produced complete and physically consistent material maps within 3–5 seconds using 20 sampling steps on the RTX 4090.

4.2. Generation Results

4.2.1 Qualitative Evaluation

As shown in Figure 3, we present six sets of real fabric images—plain weave, twill weave, and satin weave—captured using a scanner, alongside the textile images generated by our model. The left column shows the real images, while the right column displays the generated images. Our model effectively generates images that match the prompt, accurately capturing the fabric structure, including specialized parameters such as the number of floats and the density of warp and weft threads. With finer training data, more detailed labels, and the use of trigger words to control the overall fabric structure, our results align more closely with real fabric structures than previous methods.

4.2.2 Quantitative Evaluation

Table 2 presents a quantitative comparison between Stablediff-Fabric and DressCode, which we use [1]the code and checkpoint provided by the author. Our model outperforms DressCode in PSNR, SSIM, and ΔE_{00} values across all fabric types. Notably, it achieves relative PSNR gains of 11.9%,23.9%,and 12.2% on plain weave, twill weave, and satin weave, respectively. In terms of structural similarity, Stablediff-Fabric exhibits a 227.3% to 332.5% improvement in SSIM, highlighting a superior ability to capture intricate fabric details. It reflects that the model we built using high-definition datasets performs quite well at the level of human visual perception. Additionally, the ΔE_{00} values, which assesses color accuracy, shows a significant 13.9% to 28.8% reduction in color difference, further confirming the model's enhanced realism in fabric color generation.

4.3. Controllability Evaluation

4.3.1 Fineness and Density Control

In terms of controlling yarn fineness and density, as shown in Figure 4, Stablediff-fabric demonstrates excellent operability. Both qualitative prompts (e.g., 'finer warp yarns') and quantitative prompts (e.g., 'warp yarn diameter = 0.1695mm') effectively control fineness and density.

[1] https://ihe-kaii.github.io/DressCode/

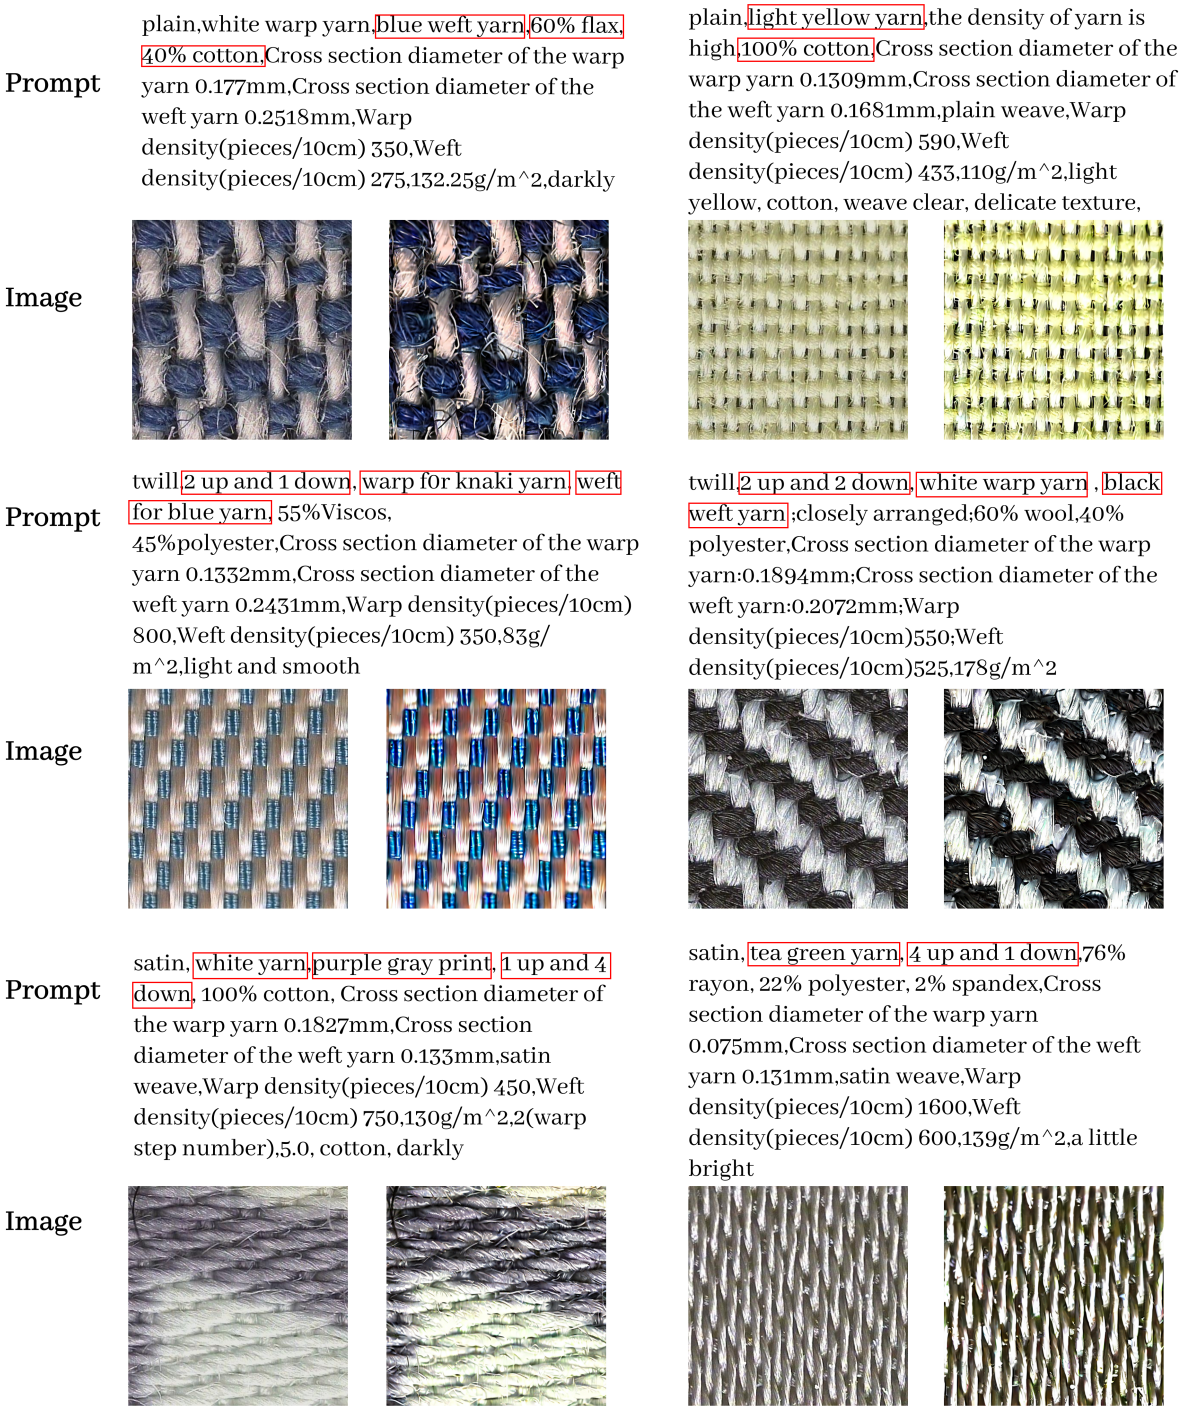

Figure 3. **Generation results of Stablediff-Fabric.** Generation results are divided into three groups: plain weave, twill weave, and satin weave. The first row shows plain weave. As seen in the first image, the model effectively distinguishes warp and weft directions while generating distinct colors. The second row showcases twill weave. Beyond accurate color reproduction, it also precisely renders the fly count within the twill structure. The third row shows satin weave. Despite its complex overall structure, the model successfully renders the fabric image alined with the prompt.

Table 2. **Quantitative comparison between Stablediff-Fabric and DressCode.** Pw: Plain weave, Tw: Twill weave, Sw: Satin weave.

	PSNR ↑			SSIM ↑			ΔE_{00} ↓		
Method	Pw	Tw	Sw	Pw	Tw	Sw	Pw	Tw	Sw
DressCode	10.68	9.94	10.88	0.0733	0.0549	0.0686	23.1585	26.9207	22.9440
Stablediff-Fabric	11.96	12.35	12.22	0.2395	0.2374	0.2623	19.9508	19.1652	18.6026

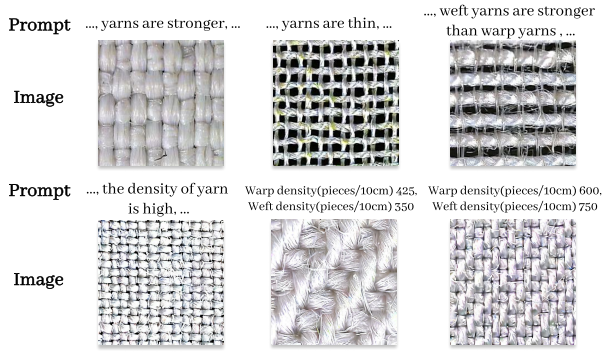

Figure 4. **Fineness and Density Control**. It demonstrates the model's ability to control fabric density and fineness.

4.3.2 Color and Brightness Control

In terms of color and brightness rendering, as shown in Figure 5, Stablediff-Fabric provides an intuitive mechanism for controllable generation. Prompt modifications directly control color and brightness, enabling faithful reproduction of visual variations.

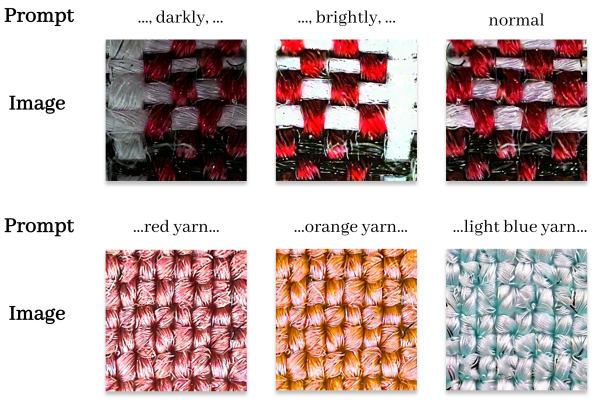

Figure 5. **Color and Brightness Control**. It demonstrates the model's ability to control fabric color and brightness.

4.3.3 Fabric Material Control

Prior works [12] [28] lacked the ability to generate fabrics with distinct material characteristics. As shown in Figure 6, Stablediff-fabric learns fabric texture characteristics from large-scale real-world image data by studying the descriptions of the physical and structural features of the surface of various types of fabric. Our model generates authentic textures across diverse materials (cotton, linen, silk, etc.), surpassing prior methods lacking material-specific fidelity.

4.4. Ablation Study

In Section 3, we introduced the primary methods for material generation and dataset preparation. We selected Low-Rank adaptation as our fine-tuning approach while applying multiple data augmentation techniques to our freshly collected dataset. We ablate data augmentation and fine-tuning strategies by comparing Stablediff-Fabric against a baseline without augmentation, evaluating across multiple dimensions. Through evaluation, we observe that our methods consistently achieve significant advantages.

4.4.1 General Quality

First, PSNR and SSIM metrics were used to compare the overall quality of the generated images, as shown in Table 3. While PSNR evaluates pixel-level similarity, SSIM focuses on perceptual similarity, which is more aligned with human visual perception. The results indicate that Stablediff-Fabric significantly outperforms the baseline method in both metrics.

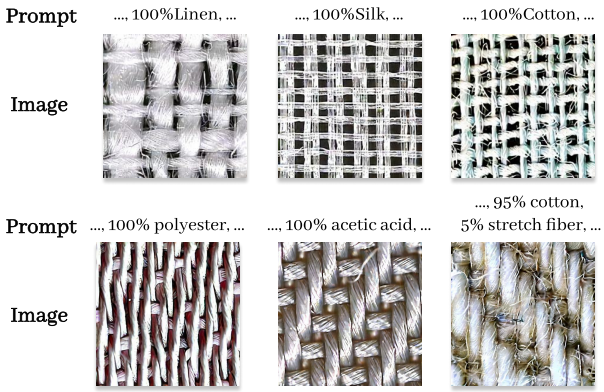

Figure 6. **Material Control**. It demonstrates the model's ability to control fabric material.

4.4.2 Fabric Structure

The structure of fabric plays a crucial role in generating detailed textures, as the warp and weft interlacing pattern directly affects physical properties such as smoothness and flatness. The accuracy of fabric structure simulation is primarily assessed by identifying the weave type (plain, twill, satin) and counting the number of floats in twill and satin weaves. A model is considered accurate if it correctly identifies both aspects. The accuracy is computed using the formula below, and the results are shown in Figure 7. Our model outperforms the baseline in simulating all three fabric structures, with data augmentation yielding a significant performance improvement.

$$\text{AccuracyRate} = \frac{\text{RightQuantity}}{\text{TestQuantity}} \tag{7}$$

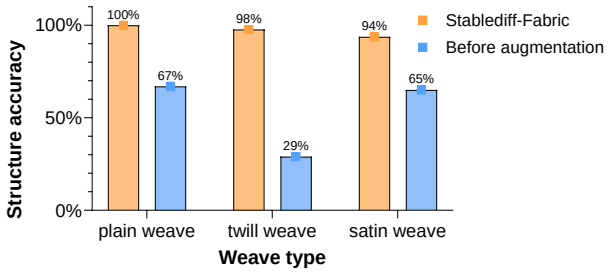

Figure 7. **Fabric structure accuracy comparison.** Both models achieved structure accuracy exceeding 88%, with data-augmented models achieving over 93% accuracy. Errors mainly arose in the generation of floating points for fabric structures.

4.4.3 Yarn density and fineness

For yarn density evaluation, we employed both automated calculations and manual assessments. As detailed in Section 3.3, fabric density was automatically computed using a wavelet-transform-based method, which provides high efficiency but may face limitations in complex structures. To enhance reliability, domain experts manually assessed the alignment accuracy of the generated images. The accuracy rates are shown in Figure 8. The results demonstrate that our method outperforms previous approaches, showing enhanced sensitivity to variations in input conditions for both yarn density and fineness.

4.4.4 Fabric Color

Color control remains a challenging aspect in generative models. As described in Section 3.4, we quantify color differences using the CIEDE2000 (ΔE_{00})color difference metric. As shown in Table 3, a smaller ΔE_{00} value indicates higher perceptual color similarity. Consistent with the visual results in Section 4.2, our method achieved higher color accuracy.

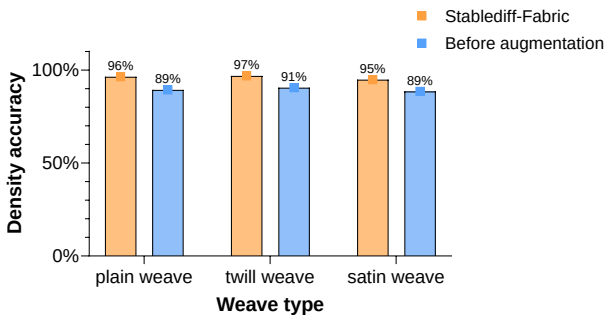

Figure 8. **Yarn density and fineness accuracy comparison.** Comparison of density accuracy before and after data augmentation. Both model sets achieved overall structural accuracy exceeding 88%, while models enhanced with data augmentation all reached structural accuracy above 94%. Errors primarily occurred in the resolution of the longitudinal and latitudinal directions.

5. Conclusion

5.1. Achievements

Stablediff-Fabric demonstrates that diffusion-based, text-driven generation can produce controllable, high-fidelity textile materials suitable for downstream garment modeling tasks. As the primary technical contribution, the method provides fine-grained control over structure, yarn density and fineness, color, and brightness while preserving material realism.

This claim is validated by both quantitative and qualitative evidence. Stablediff-Fabric achieves a PSNR of more than 20% higher than DressCode, with particularly strong gains in structurally complex weaves such as twill and satin. The SSIM results further confirm their advantage, reaching nearly four times the DressCode scores and indicating a clear improvement in perceptual realism. In terms of color reproduction, CIEDE2000 (ΔE_{00}) color difference metric shows that our outputs remain consistently closer to reference images. The proposes data-augmentation pipeline also plays a central role. It improves the accuracy of yarn fineness and density representation, enhances structural consistency, and delivers more faithful rendering of color and material properties, as confirmed by quantitative metrics and qualitative evaluation.

These findings highlight that Stablediff-Fabric goes beyond generating realistic images: it demonstrates that diffusion models can faithfully map textual descriptions into controllable fabric properties, from structure to color fidelity. This capability reduces reliance on physical sample collection, accelerates digital prototyping, and paves the way for scalable fabric generation and sustainable design

Table 3. Effect of data augmentation on model performance.

Method	PSNR ↑			SSIM ↑			ΔE_{00} ↓		
	Pw	Tw	Sw	Pw	Tw	Sw	Pw	Tw	Sw
Before-augmentation	9.6791			0.1889			23.2486	26.8508	24.0026
Stablediff-Fabric	12.1266			0.2438			19.9508	19.1652	18.6026

Pw: Plain weave, Tw: Twill weave, Sw: Satin weave. Stablediff-Fabric outperforms the baseline method in both PSNR and SSIM, demonstrating superior image quality and structural similarity. In terms of color reproduction, comparison before and after data augmentation shows that CIEDE2000 color difference (ΔE_{00}) decreased by over 15%, indicating significantly enhanced color fidelity. Overall, satin-weave samples exhibited lower metrics than other fabrics, likely due to the limited sample size.

workflows. By making high-quality, editable textile materials more accessible, our approach offers value for both academic research and industrial applications.

5.2. Limitations and Future Works

Firstly, our research is limited to the generation of woven fabric. It doesn't include other textile types, such as knit fabrics or leather, which represent possible areas for future exploration. The primary limitation of our research lies in capturing the diversity of fabric colors. While the dataset includes common fabric colors, it is not exhaustive. Due to the challenge of establishing one-to-one correspondences between nature's wide range of colors and general language descriptions, the model struggles to learn colors with precision, failing to detect subtle color differences. In experiments, we attempted data augmentation using multi-angle color transformations, but the results were suboptimal. Future work aims to incorporate more precise color descriptions into the model. Additionally, although the model achieves nearly 90% accuracy in generating warp and weft fineness and density on the test set, some images still show inaccuracies in determining warp and weft directions. Moreover, subtle differences in yarn fineness persist, making the model unsuitable for applications that demand higher precision in fabric structure. Moreover, mechanical and physical parameters are crucial for practical fabric rendering, yet the current method lacks textual feedback on these parameters and requires users to engage in iterative trial-and-error attempts to achieve desired results.

References

[1] O. Adeyokunnu, A. Abayomi-Alli, O. A. Alabi, O. Arogundade, D. W. Daniel, and B. O. Omoyiola. Classification and augmentation of african fabrics using convolutional neural network and generative adversarial network models. In *2024 IEEE 5th International Conference on Electro-Computing Technologies for Humanity (NIGERCON)*, pages 1–5. IEEE, 2024. 3

[2] D. Casciani and M. Bertolini. Towards a sustainable on-demand fashion industry: the impact of digital body measurement technologies. *Discover Sustainability*, 6(1):478, 2025. 1

[3] M. F. Cohen, J. Shade, S. Hiller, and O. Deussen. Wang tiles for image and texture generation. *ACM transactions on graphics (TOG)*, 22(3):287–294, 2003. 2

[4] W. Deng, W. Ke, Z. Deng, and X. Wang. Virtual design of woven fabrics based on parametric modeling and physically based rendering. *Computer-Aided Design*, 173:103717, 2024. 2

[5] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*, pages 341–346, New York, NY, USA, 2001. ACM. 1, 2

[6] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1033–1038. IEEE, 1999. 1

[7] R. Gal, Y. Alaluf, Y. Atzmon, O. Patashnik, A. H. Bermano, G. Chechik, and D. Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 3

[8] K. Gandhi. *Woven textiles: Principles, technologies and applications*. Woodhead Publishing, 2019. 2, 5

[9] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 3

[10] Y. Guo, C. Smith, M. Hašan, K. Sunkavalli, and S. Zhao. Materialgan: Reflectance capture using a generative svbrdf model. *arXiv preprint arXiv:2010.00114*, 2020. 3

[11] H. He, Z. Sun, J. Fan, and P. Mok. Tipgan: High-quality tileable textures synthesis with intrinsic priors for cloth digitization applications. *Computer-Aided Design*, 183:103866, 2025. 3

[12] K. He, K. Yao, Q. Zhang, J. Yu, L. Liu, and L. Xu. Dresscode: Autoregressively sewing and generating garments from text guidance. *ACM Transactions on Graphics (TOG)*, 43(4):1–13, 2024. 2, 3, 6, 9

[13] W. He, B. Song, N. Zhang, J. Xiang, and R. Pan. Modeling and realization of image-based garment texture transfer. *The Visual Computer*, 40(9):6063–6079, 2024. 3

[14] P. Henzler, V. Deschaintre, N. J. Mitra, and T. Ritschel. Generative modelling of brdf textures from flash images. *arXiv preprint arXiv:2102.11861*, 2021. 3

[15] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3

[16] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet. Video diffusion models. *Advances in neural information processing systems*, 35:8633–8646, 2022. 3

[17] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022. 5

[18] J. Hu. *Structure and mechanics of woven fabrics*, volume 35. CRC press, 2004. 1

[19] X. Hu, C. Yang, F. Fang, J. Huang, P. Li, B. Sheng, and T.-Y. Lee. Msembgan: Multi-stitch embroidery synthesis via region-aware texture generation. *IEEE Transactions on Visualization and Computer Graphics*, 31(9):5334–5347, 2024. 3

[20] Q. Huang, D. Ma, Y. Zhang, and G. Xu. Research on image synthesis of fabric replacement in suit customization. In *2022 IEEE International Conference on Imaging Systems and Techniques (IST)*, pages 1–6. IEEE, 2022. 2, 3

[21] V. Kwatra, I. Essa, A. Bobick, and N. Kwatra. Texture optimization for example-based synthesis. In *ACM Siggraph 2005 Papers*, pages 795–802. 2005. 2

[22] B. Li, X. Li, Y. Jiang, T. Xie, F. Gao, H. Wang, Y. Yang, and C. Jiang. Garmentdreamer: 3dgs guided garment synthesis with diverse geometry and texture details. In *2025 International Conference on 3D Vision (3DV)*, pages 1416–1426. IEEE, 2025. 3

[23] X. Li, Q. Yao, and Y. Wang. Garmentdiffusion: 3d garment sewing pattern generation with multimodal diffusion transformers. *arXiv preprint arXiv:2504.21476*, 2025. 2, 3

[24] Y. Li, H.-y. Chen, E. Larionov, N. Sarafianos, W. Matusik, and T. Stuyck. Diffavatar: Simulation-ready garment optimization with differentiable simulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4368–4378, 2024. 3

[25] J. Liu, H. Jiang, X. Liu, and Z. Chai. Automatic measurement for dimensional changes of woven fabrics based on texture. *Measurement Science and Technology*, 25(1):015602, 2013. 7

[26] S. McAuley, S. Hill, N. Hoffman, Y. Gotanda, B. Smits, B. Burley, and A. Martinez. Practical physically-based shading in film and game production. In *ACM SIGGRApH 2012 courses*, pages 1–7. 2012. 7

[27] Z. Montazeri, S. Gammelmark, H. W. Jensen, and S. Zhao. A practical ply-based appearance modeling for knitted fabrics. *arXiv preprint arXiv:2105.02475*, 2021. 2

[28] Z. Montazeri, S. B. Gammelmark, S. Zhao, and H. W. Jensen. A practical ply-based appearance model of woven fabrics. *ACM Transactions on Graphics (TOG)*, 39(6):1–13, 2020. 2, 9

[29] H. G. Musmann, P. Pirsch, and H.-J. Grallert. Advances in picture coding. *Proceedings of the IEEE*, 73(4):523–548, 2005. 6

[30] Y. Ohno. CIE fundamentals for color measurements. In *Proceedings of the International Conference on Digital Printing Technologies (NIP16)*, pages 540–545, Springfield, VA, USA, Jan 2000. Society of Imaging Science and Technology. 6

[31] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, and R. Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 4

[32] E. Richardson, G. Metzer, Y. Alaluf, R. Giryes, and D. Cohen-Or. Texture: Text-guided texturing of 3d shapes. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–11, 2023. 3

[33] C. Rodriguez-Pardo and E. Garces. Seamlessgan: Self-supervised synthesis of tileable texture maps. *IEEE transactions on visualization and computer graphics*, 29(6):2914–2925, 2022. 3

[34] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3, 4

[35] N. Sarafianos, T. Stuyck, X. Xiang, Y. Li, J. Popovic, and R. Ranjan. Garment3dgen: 3d garment stylization and texture generation. In *2025 International Conference on 3D Vision (3DV)*, pages 1382–1393. IEEE, 2025. 3

[36] S. Sartor and P. Peers. Matfusion: a generative diffusion model for svbrdf capture. In *SIGGRAPH Asia 2023 conference papers*, pages 1–10, 2023. 3

[37] K. Schröder, A. Zinke, and R. Klein. Image-based reverse engineering and visual prototyping of woven cloth. *IEEE transactions on visualization and computer graphics*, 21(2):188–200, 2014. 2

[38] G. Sharma, W. Wu, and E. N. Dalal. The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur*, 30(1):21–30, 2005. 7

[39] Y. Shen, J. Liang, and M. C. Lin. Gan-based garment generation using sewing pattern images. In *European Conference on Computer Vision*, pages 225–247. Springer, 2020. 3

[40] T. Shinohara. Expression of individual woven yarn of textile fabric based on yarn positional information extracted from its three dimensional ct image. In *2008 International Conference on Control, Automation and Systems*, pages 668–673. IEEE, 2008. 2

[41] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2, 3

[42] A. Srivastava, P. Manu, A. Raj, V. Jampani, and A. Sharma. Wordrobe: Text-guided generation of textured 3d garments. In *European Conference on Computer Vision*, pages 458–475. Springer, 2024. 3

[43] G. Vecchio, R. Martin, A. Roullier, A. Kaiser, R. Rouffet, V. Deschaintre, and T. Boubekeur. Controlmat: a controlled generative approach to material capture. *ACM Transactions on Graphics*, 43(5):1–17, 2024. 3

[44] G. Vecchio, R. Sortino, S. Palazzo, and C. Spampinato. Matfuse: controllable material generation with diffusion models.

In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4429–4438, 2024. 3

[45] K. Venkatraman, B. Logeshwar, and D. Aadhithyan. Garment-specific style transfer using vgg19 and semantic segmentation for high-fidelity clothing image generation. In *2024 First International Conference on Data, Computation and Communication (ICDCC)*, pages 509–514. IEEE, 2024. 3

[46] H. Wang. Proving theorems by pattern recognition—ii. *Bell system technical journal*, 40(1):1–41, 1961. 2

[47] J. Wang, R. Chen, Y. Shi, and J. Gao. Sketchfashion: Image translation from fashion sketch based on gan. In *2022 2nd International Conference on Electronic Information Technology and Smart Agriculture (ICEITSA)*, pages 1–6. IEEE, 2022. 3

[48] J. Wang, K. Shi, L. Wang, Z. Li, R. Pan, and W. Gao. Decoloration of multi-color fabric images for fabric appearance smoothness evaluation by supervised image-to-image translation. *IEEE Access*, 7:181284–181294, 2019. 3

[49] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 2

[50] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6

[51] K. Wu and C. Yuksel. Real-time cloth rendering with fiber-level detail. *IEEE transactions on visualization and computer graphics*, 25(2):1297–1308, 2017. 3

[52] W. Xian, P. Sangkloy, V. Agrawal, A. Raj, J. Lu, C. Fang, F. Yu, and J. Hays. Texturegan: Controlling deep image synthesis with texture patches. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8456–8465, 2018. 2

[53] B. Xue, C. Guarnera, S. Zhao, and Z. Montazeri. Reflectancefusion: Diffusion-based text to svbrdf generation. In *Eurographics Symposium on Rendering*. Eurographics Association, 2024. 3

[54] B. Xue, S. Zhao, H. W. Jensen, and Z. Montazeri. A hierarchical architecture for neural materials. In *Computer Graphics Forum*, volume 43, page e15116. Wiley Online Library, 2024. 2, 3

[55] Y.-Y. Yeh, J.-B. Huang, C. Kim, L. Xiao, T. Nguyen-Phuoc, N. Khan, C. Zhang, M. Chandraker, C. S. Marshall, Z. Dong, et al. Texturedreamer: Image-guided texture synthesis through geometry-aware diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4304–4314, 2024. 2

[56] C. Zhang, Y. Wang, F. Vicente, C. Wu, J. Yang, T. Beeler, and F. 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

[57] S. Zhao, W. Jakob, S. Marschner, and K. Bala. Building volumetric appearance models of fabric using micro ct imaging. *ACM Transactions on Graphics (TOG)*, 30(4):1–10, 2011. 2

[58] S. Zhao, W. Jakob, S. Marschner, and K. Bala. Structure-aware synthesis for predictive woven fabric appearance. *ACM Transactions on Graphics (TOG)*, 31(4):1–10, 2012. 1, 2

[59] Y. Zheng, Y. Chen, G. Fei, J. Dorsey, and E. Wu. Simulation of textile stains. *IEEE Transactions on Visualization and Computer Graphics*, 25(7):2471–2481, 2018. 3

Acknowledgments. This work is supported by the AI for Science Program, Shanghai Municipal Commission of Economy and Informatization (Grant Number 2025-GZLRGZN-BTBX-02014), the Fundamental Research Funds for the Central Universities (Grant Number 2232025D-15), and the Open Project Program of the State Key Laboratory of CAD&CG of Zhejiang University (Grant Number A2511).

Appendix

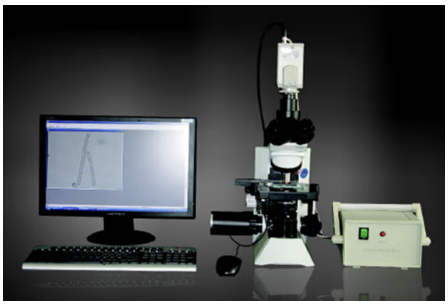

Figure 9. Device for fabric image capture and yarn fineness measurement

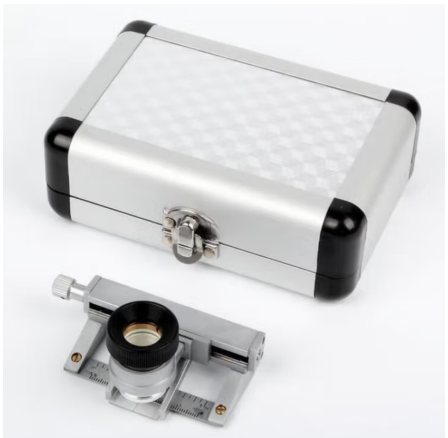

Figure 10. Device for fabric density measurement

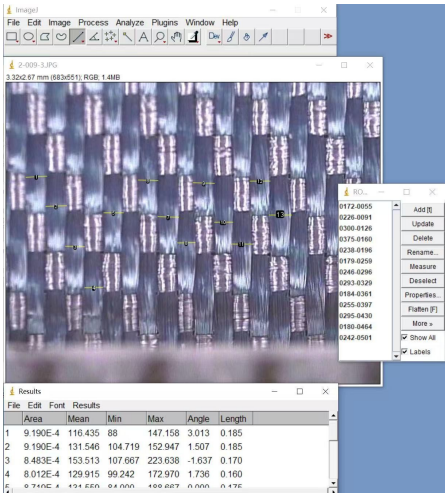

Figure 11. ImageJ software for measuring cross-sectional diameters of yarns

Figure 12. Device for measuring fabric areal weight