

# IMAGGarment: Fine-Grained Garment Generation for Controllable Fashion Design

Fei Shen, Jian Yu, Cong Wang, Xin Jiang, Xiaoyu Du, and Jinhui Tang, *Senior Member, IEEE*

**Abstract**—This paper presents IMAGGarment, a fine-grained garment generation (FGG) framework that enables high-fidelity garment synthesis with precise control over silhouette, color, and logo placement. Unlike existing methods that are limited to single-condition inputs, IMAGGarment addresses the challenges of multi-conditional controllability in personalized fashion design and digital apparel applications. Specifically, IMAGGarment employs a two-stage training strategy to separately model global appearance and local details, while enabling unified and controllable generation through end-to-end inference. In the first stage, we propose a global appearance model that jointly encodes silhouette and color using a mixed attention module and a color adapter. In the second stage, we present a local enhancement model with an adaptive appearance-aware module to inject user-defined logos and spatial constraints, enabling accurate placement and visual consistency. To support this task, we release GarmentBench, a large-scale dataset comprising over 180K garment samples paired with multi-level design conditions, including sketches, color references, logo placements, and textual prompts. Extensive experiments demonstrate that our method outperforms existing baselines, achieving superior structural stability, color fidelity, and local controllability performance. Code, models, and datasets are publicly available at <https://github.com/muzishen/IMAGGarment>.

**Index Terms**—Fine-Grained Garment Generation, Multi-Conditional Generation, Fashion Design Applications, GarmentBench Dataset.

## I. INTRODUCTION

Fine-Grained garment generation (FGG) aims to synthesize high-quality garments with precise control over garment silhouette, color scheme, logo content, and spatial placement. As personalized fashion and the digital apparel market grow rapidly, fine-grained controllability [1]–[4] is increasingly crucial for applications in fashion design and e-commerce.

In traditional garment ideation [5], [6] and visualization [7], [8], designers analyze line drawings to establish silhouette and construction, then select color palettes and materials, and

finally arrange brand elements such as logos and trims. This manual workflow has two persistent drawbacks. First, it is time consuming: to match the specification, edits must be applied object by object and view by view; in a seasonal collection, even identical panels within the same board are recolored or relabeled one at a time, which does not scale. Second, it is error prone and inconsistent: small deviations in hue, shading, or logo placement arise across artists and rounds of revision, yielding mismatches across styles, sizes, and camera viewpoints. As project scope grows, these issues inflate turnaround time and complicate quality control and version management.

Recently, image synthesis [9]–[12] has made notable progress in tasks such as sketch-to-image generation [13]–[16] and logo insertion [17]–[19] (as illustrated in Fig. 1 (a)), demonstrating basic capabilities in structural and content-level control. However, these tasks [13], [17], [20] provide only coarse guidance and rely on single-condition inputs (*e.g.*, sketch or color), lacking the fine-grained controllability needed to model the nuanced interactions between global structure and local details in garment design. Although sequential or modular combinations may offer partial solutions, they [21]–[23] fail to explicitly disentangle and jointly model global attributes (*e.g.*, silhouette, color) and local appearance details (*e.g.*, logo content and spatial placement). Without unified control mechanisms, these approaches [21]–[23] often suffer from condition entanglement, conflicting objectives, and visual inconsistencies, ultimately falling short of the high standards required in real-world fashion design. In contrast, practical fashion design [5], [6] requires joint control over multiple interdependent factors: designers determine global attributes such as silhouette and color, followed by fine-tuning of local elements like logos and their placement. To support this process, a unified generation task that clearly separates and coordinates global and local attributes is essential for controllable and high-fidelity synthesis.

To address these limitations, we propose a new task: fine-grained garment generation (FGG), as illustrated in Fig. 1 (b). FGG is formulated as a unified multi-conditional garment synthesis task, taking a textual prompt, garment silhouette, color palette, and spatially constrained logos as joint inputs. It aims to generate garments that faithfully reflect high-level structural intent and fine-grained local styling cues. FGG is specifically designed to mirror real-world fashion workflows, where designers must coordinate diverse input modalities to express creative intent. Unlike conventional approaches that process each condition independently or sequentially, FGG emphasizes joint modeling and hierarchical reasoning across

Fei Shen is with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, 210094, China, and also with the NExT++ Research Centre, National University of Singapore, Singapore, e-mail: shenfei29@nus.edu.sg

Jian Yu, Xin Jiang, and Xiaoyu Du are with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, 210094, China. e-mail: jianyu@njust.edu.cn; xinjiang@njust.edu.cn; duxy@njust.edu.cn.

Cong Wang is with the State Key Laboratory for Novel Software Technology and the School of Computer Science, Nanjing University, Nanjing, 210023, China. e-mail: cw@smail.nju.edu.cn

Jinhui Tang is with the School of Computer Science and Engineering, Nanjing University of Science and Technology, Nanjing, 210094, China, and also with the College of Information Science and Technology and Artificial Intelligence, Nanjing Forestry University, Nanjing 210037, China, e-mail: jinhuitang@njust.edu.cn. (Corresponding author: Jinhui Tang.)

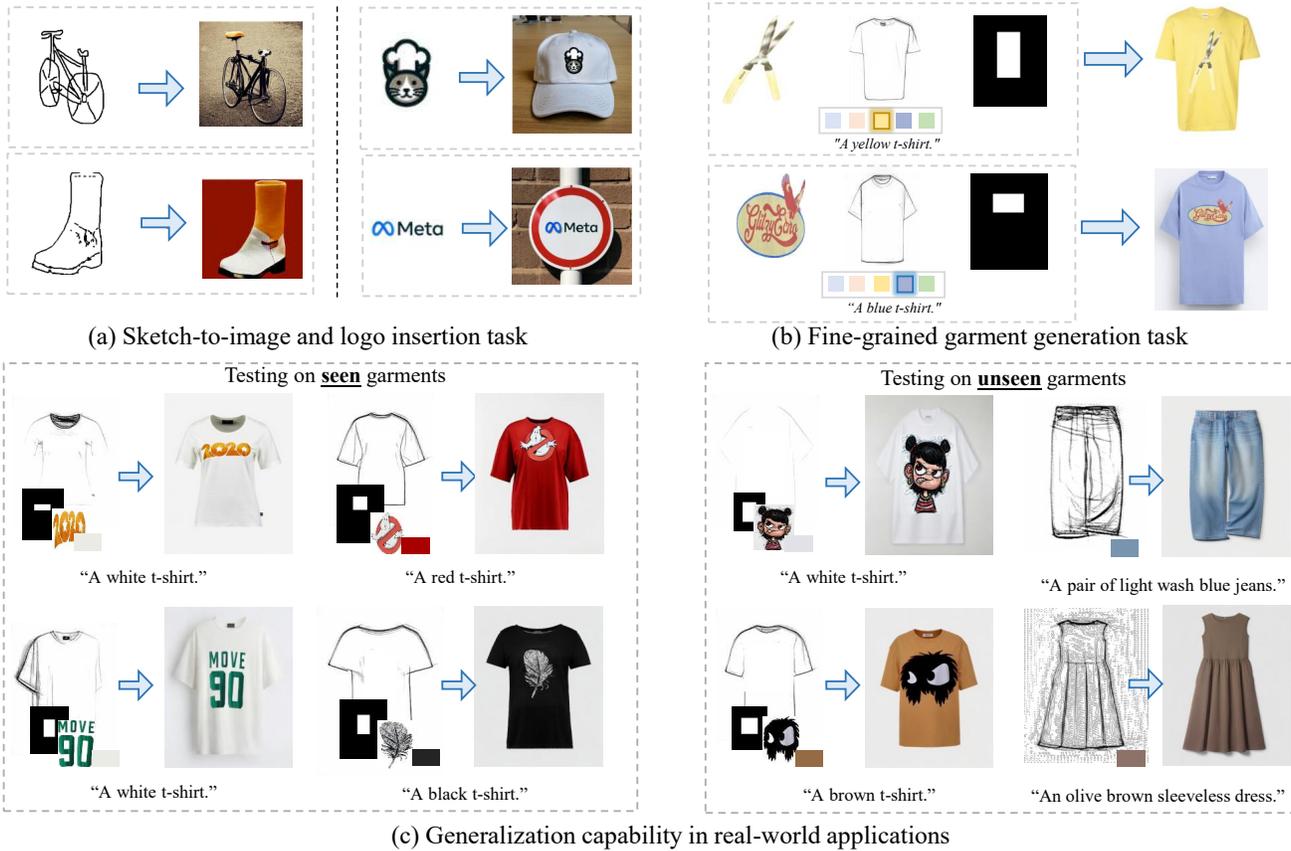


Fig. 1. Comparison of (a) existing sketch-to-image and logo insertion tasks with (b) our proposed fine-grained garment generation (FGG) task, which enables precise and controllable synthesis of garment structure, color, logo, and spatial placement. Unlike previous tasks that rely on a single input condition, FGG is tailored for real-world fashion design workflows by integrating multiple conditional controls.

input types. It goes beyond simple task combinations by enforcing consistent integration of global and local attributes within a unified generation framework, enabling nuanced control over the overall structure and detailed appearance. Specifically, **FGG task introduces three key challenges:** (1) maintaining visual and semantic consistency across heterogeneous input conditions, (2) resolving conflicts between global structures and localized visual elements, and (3) generalizing to unseen condition combinations without retraining (see Fig. 1(c)). FGG thus marks a fundamental shift from single-condition or loosely coupled pipelines toward a unified, design-intent-driven generation paradigm that better reflects the complexity of real-world garment design.

To this end, we propose IMAGGarment, a two-stage training and end-to-end inference framework tailored for fine-grained garment generation. Unlike prior methods that rely on single-condition inputs or simple feature fusion, our framework is explicitly designed to achieve fine-grained controllability under multiple, interdependent constraints. In the first stage, we propose a global appearance model with a mixed attention module and a color adapter to jointly encode garment silhouette and color palette, improving overall appearance fidelity and mitigating condition entanglement. In the second stage, we present a local enhancement model equipped with an adaptive appearance-aware module to inject user-defined logos and their spatial constraints, enabling precise logo placement while preserving global consistency. To further promote research in

this direction, we release GarmentBench, a large-scale dataset comprising over 180k garment samples annotated with rich multi-level design conditions, including silhouette sketches, color references, logo placements, and textual prompts. Extensive experiments demonstrate that IMAGGarment significantly outperforms existing baselines in terms of structural stability and local controllability. To summarize, the main contributions are listed as follows:

- We propose IMAGGarment, a controllable garment generation framework that enables precise control over garment structure, color, and logo placement, addressing the challenges of FGG.
- We design a mixed attention module, color adapter, and adaptive appearance-aware module to disentangle global structure from local attributes, achieving fine-grained visual control and accurate spatial control.
- We release GarmentBench, a large-scale dataset with diverse garments and rich multi-conditional annotations, serving as a valuable benchmark for controllable garment generation research.

The remainder of this paper is organized as follows. Section II surveys prior work on garment generation, encompassing GAN-based techniques and diffusion-based controllable generation. Section III describes the proposed IMAGGarment methodology, comprising a global appearance model with mixed attention and a color adapter, a local enhancement model with the A3 module, and the associated training and

inference strategies. Section IV presents the experimental protocol and results, including the GarmentBench dataset and evaluation metrics, implementation details, and results and analysis. Section V concludes the paper.

## II. RELATED WORK

### A. GAN-Based Methods

Early approaches [24]–[29] to garment generation predominantly build on generative adversarial networks (GANs) [30]–[32], with a major line devoted to sketch-to-image translation [33] that learns spatial mappings from structural cues. Representative systems such as DeepFaceDrawing [24] and DeepFaceEditing [25] decompose sketches into semantic components and progressively assemble photorealistic results, while DeepPortraitDrawing [26] extends this paradigm to full-body synthesis via local-to-global pipelines. Interactive frameworks [27] further introduce gating mechanisms for user-guided editing, and DALColor [34] combines WGAN-GP [35] with line-art colorization for refined appearance control. Beyond sketches, related GAN-based efforts explore pose- or part-guided generation [36], [37], leveraging learned warping or deformable alignment to better propagate structural constraints from sources to targets.

However, these methods [24]–[27] are largely restricted to single-condition settings (e.g., sketches or poses alone), making it difficult to support real-world fashion scenarios that require joint control over multiple factors such as silhouette, garment layers, color/pattern, and local embellishments. Moreover, adversarial training is prone to instability and visual artifacts [32], [36], [37], and the reliance on paired or carefully aligned supervision limits robustness to occlusion, diverse body shapes, and open-world catalogs. As a result, while GAN-based pipelines can produce plausible textures under constrained conditions, they struggle to achieve reliable, fine-grained, and multi-conditional controllability at scale.

### B. Diffusion-Based Methods

Diffusion models [38]–[40] have achieved strong progress in conditional image generation owing to their iterative denoising process and flexible conditioning interfaces. To improve controllability with minimal modification to large backbones, plugin-based approaches such as IP-Adapter [21], ControlNet [22], and BLIP-Diffusion [41] inject external conditions (e.g., reference images, structural maps, or language cues) through lightweight adapters. In parallel, reference-guided or dual-stream designs [42]–[45] propagate features from exemplars alongside text/image prompts, thereby strengthening identity preservation and fine control during sampling.

In fashion-related applications, DiffCloth [46] supports localized garment edits via part-specific textual prompts, enabling independent control over regions such as sleeves and collars. For logo-centric generation, AnyLogo [18] adopts a dual-state denoising strategy to retain subtle logo details; LogoSticker [19] performs token-based injection to flexibly place logo elements; and RefDiffuser [17] leverages expert-driven plugins to enhance texture fidelity and spatial alignment. Despite these advances, most methods emphasize either global

TABLE I  
DEFINITIONS OF MAIN SYMBOLS USED IN THIS PAPER.

Notation	Definition
$t$	Timestep
$Z_t$	Latent feature at $t$ step
$Z_m$	Output of mixed attention
$x_0$	Real image
$x_t$	Noisy data at $t$ step
$G$	Garment image
$L$	Logo image
$M$	Mask image
$C_g$	Feature of garment image
$C_l$	Feature of logo image
$C_m$	Feature of mask image
$C_s$	Feature of silhouette image
$C_c$	Feature of color image
$C_t$	Feature of text prompt
$\theta_g$	Global appearance model
$\theta_l$	Local enhancement model
$\epsilon$	Gaussian noise
$\bar{\alpha}_t$	Cumulative product of noise weights
$w$	Guidance scale
$\alpha$	Silhouette scale
$\beta$	Color scale

appearance control or localized editing in isolation. A unified framework that jointly models multiple design conditions, e.g., silhouette and layer topology together with color/pattern and local embellishments, while maintaining structural coherence across the denoising trajectory remains underexplored.

## III. METHODOLOGY

**Symbol Definition.** To introduce our IMAGGarment method more clearly, we define the main symbols used throughout the paper in TABLE I.

**Task Definition.** Given a garment silhouette, color palette, user-defined logo, location and an optional text description, fine-grained garment generation (FGG) aims to synthesize high-fidelity garment images with precise control over both global structure and local visual attributes. The key challenges lie in jointly modeling multi-conditional inputs, maintaining semantic and visual consistency across different design factors, and supporting controllable placement of fine-grained elements such as logos and color regions.

### A. Overall Framework

To address the above challenges, we propose IMAGGarment, a conditional diffusion framework tailored for fine-grained garment generation. Our framework comprises two components: a global appearance model (stage I) and a local enhancement model (stage II), which explicitly disentangle and jointly control the global appearance and local details under multi-conditional guidance, enabling accurate synthesis of garment silhouette, color, and logo placement. As illustrated in Fig. 2, the global appearance model first generates a latent of coarse garment image conditioned on the textual prompt, garment silhouette, and color palette. Subsequently, the local enhancement model refines this latent representation by integrating user-defined logo and spatial constraint, producing the final high-fidelity garment image with fine-grained controllability. Specifically, the global appearance model (Section III-B)

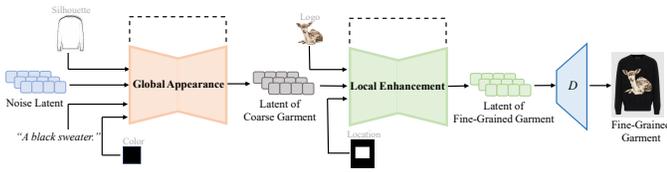


Fig. 2. **Visualization of the IMAGGarment inference pipeline.** The global appearance model generates coarse latent from textual prompts, silhouettes, and colors. The local enhancement model then injects user-defined logos and spatial location constraints to produce the fine-grained garment.

leverages our proposed mixed attention module and color adapter to effectively capture global appearance features from textual descriptions, silhouettes, and colors, while mitigating entanglement among these conditions. The local enhancement model (Section III-C) introduces an adaptive appearance-aware module ( $A^3$  Module) that injects logo content and spatial location constraint into the latent space, achieving precise logo placement. Finally, the training and inference strategies used in IMAGGarment are summarized in Section III-D.

### B. Stage I: Global Appearance Model

**Motivation.** Existing garment generation methods [21]–[23] typically rely on single-condition inputs (*e.g.*, sketch or text), causing entangled features and limited controllability. To resolve this, we propose a global appearance model that explicitly disentangles silhouette, color, and text, enabling precise multi-conditional control.

**Architecture.** As shown in the left of the Fig. 3, our global appearance model comprises two shared frozen VAE encoders, one frozen VAE decoder, a trainable silhouette UNet, a frozen text encoder, a trainable color adapter, and a denoising UNet with the proposed mixed attention. Specifically, we first utilize the frozen VAE encoder to project the input reference silhouette into the latent space. Subsequently, we employ a trainable silhouette UNet (structurally identical to the denoising UNet but without cross attention) to extract fine-grained silhouette features, which are then integrated into the frozen denoising UNet via our proposed mixed attention module. Meanwhile, textual features obtained from the frozen CLIP text encoder and color features extracted by the proposed color adapter are further fused into the denoising UNet through cross attention. After multiple denoising iterations, the model generates coarse garment images that precisely align with the reference silhouette and faithfully reflect user-specified color.

**Mixed Attention.** To effectively incorporate reference silhouette features into the denoising UNet without compromising the generative capability of the original UNet, we propose a mixed attention module. As shown in Fig. 3, we extend all self attention layers in the denoising UNet to the proposed mixed attention, which introduces two additional learnable projection layers to align the silhouette features  $C_s$  with the latent features  $Z_t$ . Formally, the mixed attention is defined as:

$$Z_m = \text{Softmax} \left( \frac{QK^T}{\sqrt{d}} \right) V + \alpha \cdot \text{Softmax} \left( \frac{Q(K')^T}{\sqrt{d}} \right) V', \quad (1)$$

where  $\alpha$  is a hyperparameter controlling the strength of silhouette conditioning. The projections are computed as follows:

$$Q = Z_t W_q, \quad K = Z_t W_k, \quad V = Z_t W_v, \quad K' = C_s W'_k, \quad V' = C_s W'_v \quad (2)$$

where  $W_q, W_k, W_v$  are frozen parameters of linear projection layers, whereas  $W'_k, W'_v$  are newly introduced learnable parameters of projection layers initialized from  $W_k$  and  $W_v$ , respectively. Our mixed attention facilitates the seamless integration of silhouette features into the denoising UNet, thus ensuring that generated garments maintain precise spatial alignment with the reference silhouette.

**Color Adapter.** Accurate color manipulation is essential for generating garments with fine-grained visual details, significantly enhancing visual quality and realism. However, as the base model’s textual prompts cannot reliably produce the intended colors, discrepancies often arise between the generated and expected colors. To address this issue, we propose a dedicated color adapter that explicitly treats color as an independent controllable factor. Specifically, given a reference color image, we extract color features  $C_c$  using a frozen CLIP image encoder combined with a trainable linear layer. Subsequently, these color features are integrated into the denoising UNet via a cross attention mechanism, jointly with textual features  $C_t$  obtained from the frozen CLIP text encoder:

$$Z_{new} = \text{Softmax} \left( \frac{QK_t^T}{\sqrt{d}} \right) V_t + \beta \cdot \text{Softmax} \left( \frac{QK_c^T}{\sqrt{d}} \right) V_c, \quad (3)$$

where  $Q = Z_t W_q, K_t = C_t W_k^t, V_t = C_t W_v^t$ , and  $K_c = C_c W_k^c, V_c = C_c W_v^c$ . Here,  $W_k^t, W_v^t$  denote frozen parameters of the original cross attention layers in the denoising UNet, while  $W_k^c, W_v^c$  are newly introduced trainable projection layers. The hyperparameter  $\beta$  modulates the adapter’s influence, ensuring precise alignment between generated colors and user specifications.

### C. Stage II: Local Enhancement Model

**Motivation.** Existing methods [18], [19] typically neglect detailed logo integration or treat it as a separate task, causing poor spatial alignment and visual inconsistency. To address this limitation, we propose a local enhancement model equipped with an adaptive appearance-aware ( $A^3$ ) module, explicitly injecting user-defined logos and spatial constraints into the latent space. This design enables precise, consistent control over localized garment details, significantly enhancing visual fidelity.

**Architecture.** As illustrated on the right of Fig. 3, the local enhancement model comprises a frozen VAE encoder and decoder, a denoising UNet, and an adaptive appearance-aware module ( $A^3$  module). The  $A^3$  module fuses local conditions, such as logos and spatial constraints, by concatenating them along spatial or channel dimensions, enabling precise control over fine-grained visual elements. Given a garment, logo, and placement mask, the model adaptively adjusts the logo’s size and position while preserving its visual fidelity. To reduce redundancy and focus on local detail refinement, we optimize only the self attention layers of the denoising UNet and discard

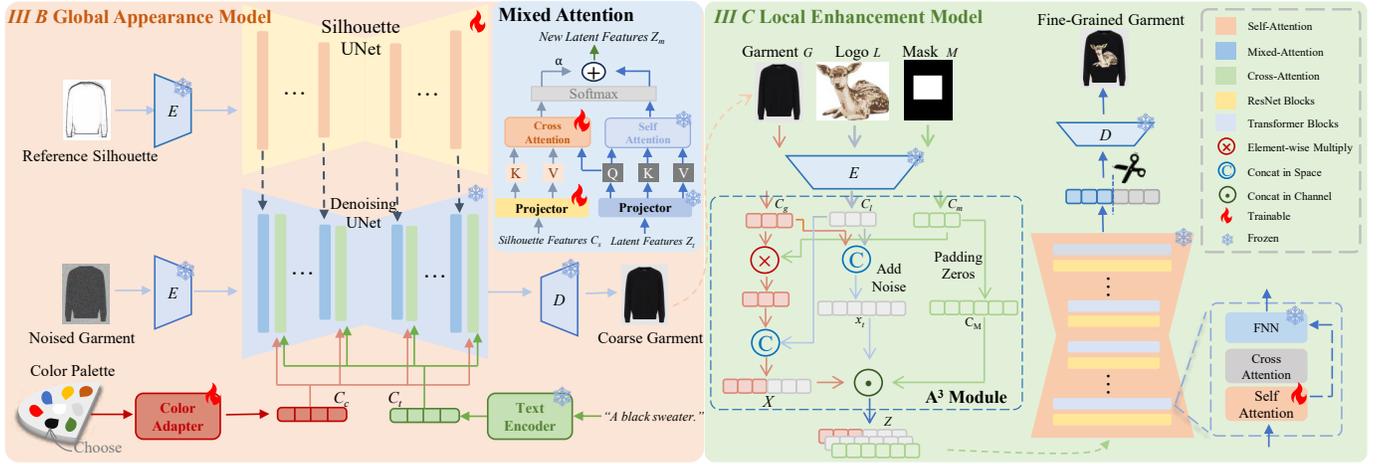


Fig. 3. **Overview of our IMAGGarment framework.** IMAGGarment is a two-stage conditional diffusion framework for fine-grained garment generation. The global appearance model first synthesizes a coarse latent representation from the input text prompt, silhouette, and color palette using a parallel UNet with mixed attention and a color adapter. The local enhancement model then refines this latent by injecting user-defined logos and location constraints through the proposed  $A^3$  module, enabling precise logo placement and high-fidelity garment generation.

all cross attention layers, as the global appearance model has already encoded the textual information.

**$A^3$  Module.** To precisely integrate fine-grained logo details into designated garment regions, we introduce the adaptive appearance-aware ( $A^3$ ) module. By fusing image-based conditions across specific dimensions, our  $A^3$  module enables precise and consistent logo integration. Specifically, given a coarse garment image  $G$ , a logo image  $L$ , and a binary placement mask  $M$ , we first encode them using a frozen VAE encoder to obtain their corresponding latent features:  $C_g \in \mathbb{R}^{4 \times \frac{H}{8} \times \frac{W}{8}}$  and  $C_l \in \mathbb{R}^{4 \times \frac{H}{8} \times \frac{W}{8}}$ . The mask  $M$  is resized via nearest-neighbor interpolation to match the latent resolution, resulting in  $C_m \in \mathbb{R}^{1 \times \frac{H}{8} \times \frac{W}{8}}$ . We then construct the spatially aligned conditional input as:

$$X = \text{Concat}(C_g \otimes C_m, C_l), \quad X \in \mathbb{R}^{4 \times \frac{H}{8} \times \frac{W}{4}}, \quad (4)$$

where  $\otimes$  denotes element-wise multiplication and  $\text{Concat}$  indicates spatial concatenation along the width dimension. To align with  $X$ , the resized mask  $C_m$  is zero-padded to obtain  $C_M \in \mathbb{R}^{1 \times \frac{H}{8} \times \frac{W}{4}}$ . Next, we concatenate the garment and logo features to form a clean latent representation:

$$x_0 = \text{Concat}(C_g, C_l), \quad (5)$$

and inject noise consistent with the diffusion process:

$$x_t = \sqrt{\bar{\alpha}_t} \cdot x_0 + \sqrt{1 - \bar{\alpha}_t} \cdot \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (6)$$

where  $x_0$  denotes the clean latent feature obtained by concatenating garment and logo features, and  $x_t \in \mathbb{R}^{4 \times \frac{H}{8} \times \frac{W}{4}}$  is the corresponding noisy latent at diffusion timestep  $t$ .  $\bar{\alpha}_t$  is the cumulative product of the noise schedule coefficients, and  $\epsilon$  is the Gaussian noise sampled from  $\mathcal{N}(0, \mathbf{I})$ . Finally, the full model input is obtained by concatenating the noisy latent  $x_t$ , the padded mask  $C_M$ , and the aligned conditional input  $X$  along the channel dimension:

$$Z = \text{Concat}(x_t, C_M, X), \quad Z \in \mathbb{R}^{9 \times \frac{H}{8} \times \frac{W}{4}}. \quad (7)$$

This channel-wise concatenation allows the model to jointly reason over appearance, spatial constraints, and guidance

signals, while maintaining compatibility with the UNet architecture for spatially aware logo synthesis.

#### D. Training and Inference

**Training.** The training process is divided into two stages, each targeting a specific set of objectives with separate optimization strategies. We first train the global appearance model independently to generate a semantically coherent garment representation conditioned on silhouette and color. After verifying its performance, we freeze it and train the local enhancement model to inject fine-grained logos guided by spatial masks. This sequential training avoids gradient interference between heterogeneous objectives and ensures each module converges toward its task-specific goal. Both stages adopt mean squared error (MSE) loss to supervise the denoising process.

**Stage I.** The global appearance model  $\theta_g$  is trained to synthesize garments that align with the target silhouette and color under textual guidance. To preserve the generative capacity of the pretrained denoising UNet, we freeze all parameters except those of the silhouette UNet and the cross-attention projections in the mixed attention module. Given silhouette features  $C_s$ , text embeddings  $C_t$ , and color features  $C_c$ , we adopt a decoupled training strategy with  $L_{\text{silhouette}}$  and  $L_{\text{color}}$  losses:

$$L_{\text{silhouette}} = \mathbb{E}_{x_0, \epsilon, C_t, C_s, t} \|\epsilon - \epsilon_{\theta_g}(x_t, C_t, C_s, t)\|^2, \quad (8)$$

$$L_{\text{color}} = \mathbb{E}_{x_0, \epsilon, C_t, C_c, t} \|\epsilon - \epsilon_{\theta_g}(x_t, C_t, C_c, t)\|^2,$$

where  $\epsilon$  is the added noise and  $\epsilon_{\theta_g}$  is the prediction from the global appearance model at timestep  $t$ .

**Stage II.** The local enhancement model  $\theta_l$  refines the coarse latent by injecting logos at user-defined locations. To reduce overhead, we fine-tune only the self-attention layers of the logo UNet. Given logo feature  $C_l$ , spatial mask  $C_m$ , and garment latent  $C_g$ , the training objective  $L_{\text{logo}}$  is:

$$L_{\text{logo}} = \mathbb{E}_{x_0, \epsilon, C_l, C_m, C_g, t} \|\epsilon - \epsilon_{\theta_l}(x_t, C_l, C_m, C_g, t)\|^2, \quad (9)$$

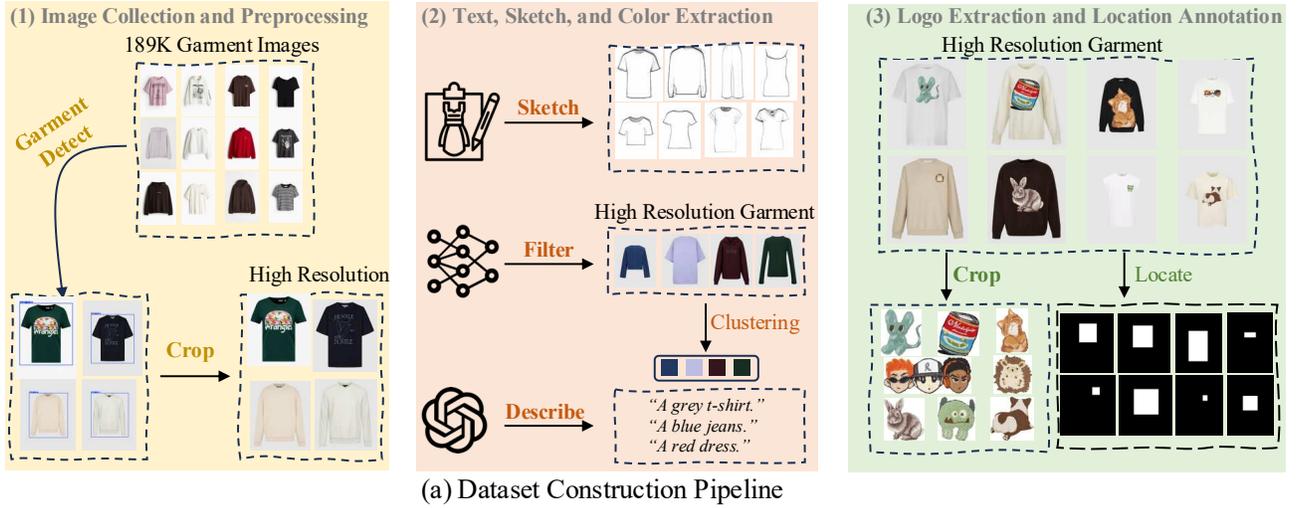


Fig. 4. Overview of GarmentBench dataset construction pipeline and samples. (a) Data construction pipeline for GarmentBench. (b) Example samples with multimodal annotations: silhouette, logo, text, logo location, and color.

where  $\epsilon_{\theta_l}$  denotes the prediction from the local enhancement model.

**Inference.** IMAGGarment supports end-to-end inference through a two-stage pipeline operating in a shared latent space. The global appearance model first generates a latent of coarse garment image conditioned on the input text prompt, silhouette, color, and mask. This process is guided by classifier-free guidance (CFG) [47]:

$$\check{\epsilon}_{\theta_g}(x_t, C_t, C_s, C_c, t) = w \cdot \epsilon_{\theta_g}(x_t, C_t, C_s, C_c, t) + (1 - w) \cdot \epsilon_{\theta_g}(x_t, t) \quad (10)$$

here,  $w$  is the CFG scale and  $x_t$  denotes the noisy latent at timestep  $t$ . The coarse latent is then refined by the local enhancement model, which incorporates user-defined logos and spatial constraints through the  $A^3$  module. We apply conditional CFG:

$$\check{\epsilon}_{\theta_l}(x_t, C_l, C_m, C_g, t) = w \cdot \epsilon_{\theta_l}(x_t, C_l, C_m, C_g, t) + (1 - w) \cdot \epsilon_{\theta_l}(x_t, C_m, C_g, t) \quad (11)$$

## IV. EXPERIMENTS

### A. Dataset and Metrics

**Dataset Construction.** As shown in Fig. 4 (a), we construct and release GarmentBench, a large-scale dataset for fine-grained garment generation, containing multi-modal design conditions such as text, sketches, colors, logos, and location masks. It serves as a controllable and extensible benchmark for advancing personalized fashion generation. The construction process is as follows:

(1) **Image Collection and Preprocessing.** We collect over 189K high-quality garment images from the internet, covering a wide range of categories such as tops, bottoms, and dresses. To eliminate background distractions and focus on the garment region, we apply YOLOv8 [48] for clothing detection and perform tight cropping to obtain clean garment-centric images for further processing.

(2) **Text, Sketch, and Color Extraction.** For each image, we automatically generate three auxiliary conditions to simulate real-world design guidance: textual descriptions generated by the multi-modal LLM Qwen-VL-Chat [49], covering key attributes such as color, silhouette, and style; structural sketches

obtained using Informative-Drawings [50], providing shape and layout priors; and color palettes extracted from single-color garments identified via ResNet50 [51] and clustered using K-means [52].

**(3) Logo Extraction and Location Annotation.** To support logo insertion and spatial control, we further extract local design elements such as logos and prints. We use YOLOv8 to detect visually distinct regions (e.g., anime characters, animal patterns), followed by manual verification to ensure label quality. We also annotate spatial locations and generate binary masks to serve as precise spatial constraints. In total, GarmentBench contains 189,966 garment-condition pairs with rich fine-grained annotations.

**Dataset Description.** As shown in Fig. 4 (b), we present representative samples from the GarmentBench dataset, which include fine-grained garment images paired with multi-modal conditions such as textual descriptions, structural silhouettes, color references, logos, and spatial location masks. Additionally, we randomly sample images from the Fashion-ControlNet-Dataset-V3<sup>1</sup> and apply the same preprocessing pipeline as GarmentBench to construct a test set with 1,267 image-condition pairs for evaluation and comparative analysis.

**Dataset Statement.** GarmentBench is curated from publicly available fashion imagery under a non-commercial research intent. All personal identifiers were removed; third-party logos and brand marks are included solely to evaluate controllability and remain the property of their respective owners. We release only derived annotations and source URLs (not raw images), together with license notices and a takedown procedure; exact split indices and random seeds are provided for reproducibility.

**Evaluation Metrics.** We adopt four metrics to comprehensively evaluate visual quality, conditional consistency, and fine-grained controllability. Fréchet inception distance (FID) [53] measures the distribution similarity between generated and real images, reflecting overall realism. Color structure similarity (CSS) [54] assesses the consistency of color distribution, measuring color controllability. Lastly, Logo location accuracy (LLA) [55] quantifies the spatial deviation between generated and target logo positions, reflecting spatial precision. Learned perceptual image patch similarity (LPIPS) [56] reflects human-perceived visual similarity, effectively capturing structural and textural consistency. These metrics comprehensively assess quality and controllability in fine-grained garment generation.

### B. Implementation Details

In our experiments, both the silhouette UNet and the denoising UNet are initialized with the pretrained Stable Diffusion v1.5 model<sup>2</sup>. The local enhancement model is based on the inpainting variant of Stable Diffusion v1.5<sup>3</sup>, with only the self-attention layers being fine-tuned to reduce computational cost. We adopt OpenCLIP ViT-H/14<sup>4</sup> as the CLIP image encoder. All input images are resized to  $512 \times 640$  resolution. We

TABLE II  
QUANTITATIVE COMPARISONS ON GARMENTBENCH. OURS ACHIEVES THE TOP RESULTS ACROSS ALL METRICS, WITH BEST IN BOLD.

Method	FID ↓	CSS ↓	LLA ↑	LPIPS ↓
BLIP-Diffusion* [41]	101.99	104.44	0.13	0.68
ControlNet-Garment* [22]	41.22	83.30	0.36	0.41
AnyDoor* [59]	38.08	68.24	0.65	0.17
IP-Adapter-Garment* [21]	37.95	92.95	0.36	0.43
<b>IMAGGarment (Ours)</b>	<b>17.63</b>	<b>36.16</b>	<b>0.72</b>	<b>0.10</b>

\* denotes re-implemented by us for a fair comparison.

use the AdamW optimizer [57] with a constant learning rate of  $1 \times 10^{-5}$ . The global appearance model and the local enhancement model are trained for 150K and 50K steps, respectively, using a batch size of 20. During inference, we adopt the DDIM sampler [58] with 50 sampling steps. Unless otherwise specified, the silhouette weight  $\alpha$  and color weight  $\beta$  in Eq.1 and Eq.3 are set to 0.6 and 1.0. The classifier-free guidance (CFG) scale  $w$  in Eq.10 and Eq.11 is set to a default value of 7.0.

### C. Baseline Comparisons

Due to the absence of prior work tailored to fine-grained garment generation with multi-condition control, we compare our method against four representative baselines: BLIP-Diffusion [41], AnyDoor [59], ControlNet [22], and IP-Adapter [21]. For subject-driven generation methods, BLIP-Diffusion [41] leverages a learnable Q-Former to align textual and visual embeddings in the latent space, initially designed for subject-preserving generation from text-image pairs. AnyDoor [59] combines identity and detail encoders to reconstruct personalized content, which we adapt to conditions of garment appearance and logo inputs. For plugin-based baselines, we extend ControlNet [22] and IP-Adapter [21] by duplicating and modifying their conditional branches to support multi-conditional inputs, such as silhouette, color, and logo. The adapted versions are referred to as ControlNet-Garment and IP-Adapter-Garment. Specifically, for ControlNet-Garment, we input silhouette, color, logo and mask maps into the ControlNet branch and inject them at each downsampling block, following standard practice. For IP-Adapter-Garment, we extend the official implementation to accept silhouette, color, logo and mask embeddings, which are concatenated and injected via cross-attention. To ensure task relevance, all methods are fine-tuned on our GarmentBench dataset with support for logo-specific conditioning. All methods are trained and evaluated under identical training protocols, input resolutions, and hardware setups. The corresponding quantitative and qualitative results are presented in Table II and Fig. 5, respectively, with detailed analysis provided below.

**Quantitative Results.** As shown in Table II, IMAGGarment achieves the best performance across all four metrics on the GarmentBench dataset, demonstrating its superiority in controllable fine-grained garment generation. Compared to subject-driven methods (BLIP-Diffusion [41], AnyDoor [59]), which rely on global features for personalized reconstruction, IMAGGarment shows substantial improvements in FID,

<sup>1</sup>[https://huggingface.co/datasets/Abrumu/Fashion\\_controlnet\\_dataset\\_V3](https://huggingface.co/datasets/Abrumu/Fashion_controlnet_dataset_V3)

<sup>2</sup><https://huggingface.co/stable-diffusion-v1-5/stable-diffusion-v1-5>

<sup>3</sup><https://huggingface.co/stable-diffusion-v1-5/stable-diffusion-inpainting>

<sup>4</sup>[https://github.com/mlfoundations/open\\_clip](https://github.com/mlfoundations/open_clip)

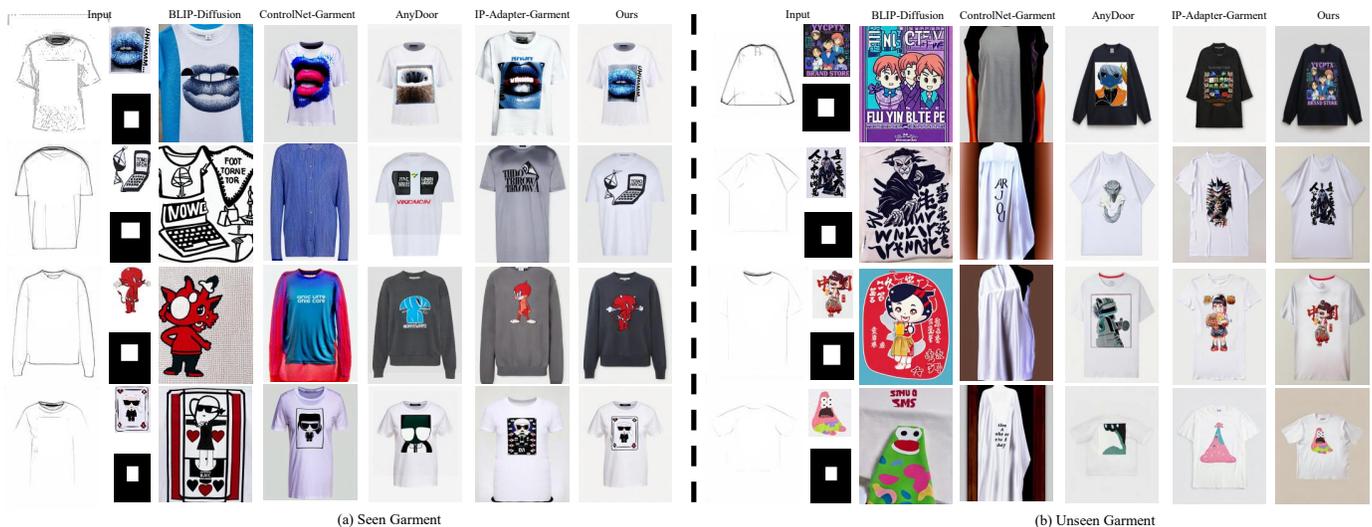


Fig. 5. **Qualitative results on seen and unseen GarmentBench samples.** The seen set uses original test pairs, while the unseen set involves randomly mixed conditions. IMAGGarment delivers the most consistent outputs, achieving accurate silhouette, color, and logo control across both settings.

TABLE III  
QUANTITATIVE ABLATION RESULTS ON GARMENTBENCH.

Method	FID ↓	CSS ↓	LLA ↑	LPIPS ↓
B0	139.33	104.54	0.15	0.64
B1	47.42	36.65	0.30	0.15
B2	30.19	97.05	0.56	0.33
B3	21.20	43.00	0.65	0.11
B4	46.16	108.25	0.52	0.38
<b>Full</b>	<b>17.63</b>	<b>36.16</b>	<b>0.72</b>	<b>0.10</b>

LPIPS, and CSS. These gains highlight the effectiveness of our mixed attention and color adapter modules in achieving coherent multi-condition fusion, resulting in more realistic, perceptually consistent, and color-faithful outputs. In contrast to plugin-based approaches (ControlNet-Garment [22], IP-Adapter-Garment [21]) that simply stack independent conditional branches, IMAGGarment yields significantly higher LLA, reflecting more precise logo placement. Our proposed  $A^3$  module drives these improvements, which adaptively injects spatial priors and logo features into the latent space for accurate local control. Overall, these results indicate that global-only conditioning or naive plugin stacking is insufficient for fine-grained control. By contrast, IMAGGarment provides an effective solution for multi-conditional garment synthesis, enabling precise coordination of global structure and local detail.

**Qualitative Results.** Fig. 5 presents qualitative comparisons on both seen and unseen garments. Notably, the seen test set refers to the designated test split of our GarmentBench dataset. In the absence of other suitable public datasets, we assess generalization using an unseen-composition test split constructed by randomly recombining input conditions (e.g., silhouette, color, logo) into combinations that never appear during training, thereby simulating real-world fashion-design scenarios. On seen garments, subject-driven methods (BLIP-Diffusion [41], AnyDoor [59]) reconstruct global appearance but lack spatial control. BLIP-Diffusion retains

logo identity yet fails at precise placement due to text-only conditioning, while AnyDoor introduces logo distortions and stylistic artifacts. Plugin-based baselines (ControlNet-Garment [22], IP-Adapter-Garment [21]) treat conditions independently, resulting in poor coordination. ControlNet-Garment suffers from cross-condition interference, and IP-Adapter-Garment often misplaces logos despite preserving structure. In contrast, IMAGGarment achieves accurate control over silhouette, color, and logo placement. On unseen garments, all baselines degrade notably. Subject-driven methods fail to generalize to novel layouts, AnyDoor distorts appearance, and BLIP-Diffusion struggles with logo positioning. Plugin-based methods also falter: ControlNet-Garment produces mismatched outputs, and IP-Adapter-Garment cannot interpret unseen spatial semantics. IMAGGarment remains robust, maintaining alignment across all conditions. This generalization stems from our  $A^3$  module, which effectively integrates spatial and visual cues in the latent space. These results validate the controllability and flexibility of our method in both seen and unseen settings.

#### D. Ablation Study

To validate the effectiveness of each component in our framework, we design a series of ablation variants within the IMAGGarment architecture: **B0** uses the vanilla Stable Diffusion v1.5 without any of our proposed modules, serving as the baseline. **B1** removes the local enhancement model (Stage II), evaluating the impact of omitting logo injection and spatial control. **B2** removes the global appearance model (Stage I), assessing the model’s performance without structured silhouette and color conditioning. **B3** removes the color adapter from the global appearance model, isolating the role of color guidance in generation. **B4** replaces our mixed attention with vanilla self-attention in the denoising UNet, testing the importance of spatial fusion with silhouette features. **Full** represents the complete IMAGGarment framework with all proposed modules integrated.



Fig. 6. **Qualitative ablation results on GarmentBench.** The “Full” configuration achieves the best results, highlighting the importance of each component.

**Ablation of Architecture Design.** Table III presents the quantitative impact of each component in our proposed IMAGGarment. In B1, which removes the local enhancement stage, the model struggles to place logos precisely, leading to degraded LLA. Although the overall garment structure is preserved, the lack of spatial control prevents accurate logo integration. In B2, without the global appearance stage, the model fails to maintain silhouette and color consistency, resulting in significantly worse FID, LPIPS, and CSS. This demonstrates that local injection alone is insufficient to handle global garment layouts. B3 disables the color adapter, causing notable drops in CSS, highlighting its role in faithful color transfer and control. B4 replaces our mixed attention with standard self-attention, which weakens the fusion of silhouette guidance and causes drops in both LPIPS and FID, indicating reduced realism and structural coherence. The full IMAGGarment achieves the best performance across all metrics, validating the complementary design of each module’s effectiveness in handling multi-condition garment generation. Further, Fig. 6 shows qualitative comparisons. B1 fails to align logos spatially, while B2 produces distorted garments lacking color and silhouette guidance. Despite maintaining logo placement, B3 leads to color mismatch, and B4 generates less coherent garment layouts. In contrast, the full model successfully synthesizes garments with accurate silhouettes, precise logo placement, and faithful color reproduction, demonstrating the benefits

of our dual-stage design, color adapter, and mixed attention fusion. Overall, The “Full” configuration achieves the best results, highlighting the importance of each component.

#### E. More Results and Analysis

**Controllability Analysis.** We assess controllability by varying a single condition at a time (silhouette, color palette, or logo position) while keeping the others fixed. As shown in Fig. 7, each three column block visualizes the model’s response to one condition. Changing the silhouette (left block) yields garments that match the target shapes, indicating that the mixed attention module preserves structural alignment. Varying the color palette (middle block) produces the intended color distributions, validating the color adapter for color faithful generation. Adjusting the logo position (right block) achieves precise spatial relocation, showing that the  $A^3$  module effectively injects spatial priors for local control. Overall, IMAGGarment provides fine-grained and decoupled control of garment attributes suitable for practical design workflows. Non-varied attributes remain stable across manipulations, reflecting minimal cross-condition interference and consistent editing behavior. Sequential composition of edits across attributes produces similar outcomes regardless of edit order, which suggests low inter-attribute coupling. Control fidelity also holds under moderate changes of viewpoint and background, supporting robustness in real design scenarios.



Fig. 7. **Controllability visualization.** Each block varies one input condition while keeping others fixed. Left: Silhouette changes lead to consistent structural adaptation. Middle: Color palette variation results in accurate color transfer. Right: Logo mask adjustment yields precise spatial placement.

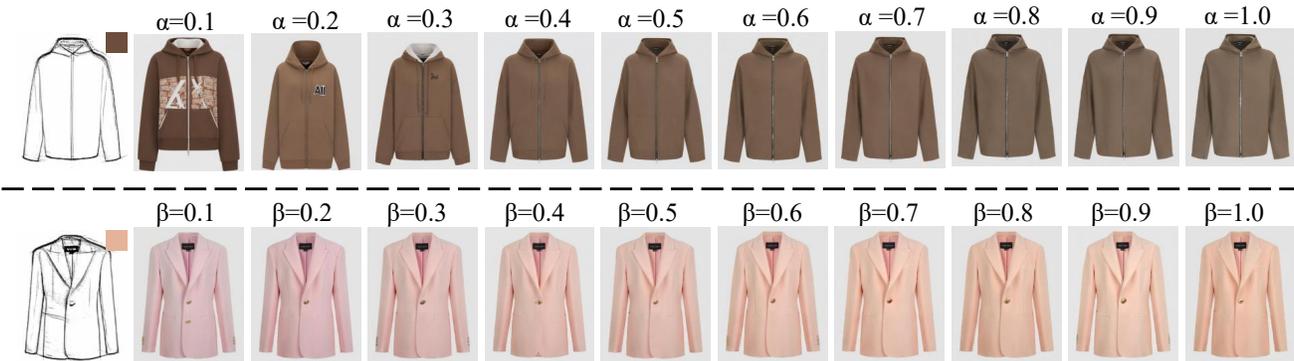


Fig. 8. **Hyperparameter analysis of silhouette weight  $\alpha$  and color weight  $\beta$ .**

**Hyperparameter Analysis.** We study the effect of two key hyperparameters in Eq.1 and Eq.3: the silhouette guidance weight  $\alpha$  and the color conditioning weight  $\beta$ . From Fig. 8, varying  $\alpha$  directly impacts the model’s ability to follow the reference silhouette. When  $\alpha$  is too low, the generated structure becomes blurry or deviates from the target shape; when too high, it may suppress color and text guidance. We empirically set  $\alpha = 0.6$  for balanced structural alignment. Similarly, the color weight  $\beta$  controls the influence of the color palette. As  $\beta$  increases, color consistency improves steadily, with  $\beta = 1.0$  yielding the best visual fidelity. Joint sweeps over  $(\alpha, \beta)$  indicate a broad stability region around  $\alpha \in [0.5, 0.7]$  and  $\beta \in [0.8, 1.1]$ , showing robustness to moderate mistuning. Interaction effects are mild: very large  $\alpha$  slightly narrows the effective range of  $\beta$ , while very large  $\beta$  can oversaturate colors and reduce shading nuance. We therefore adopt  $\alpha = 0.6$  and  $\beta = 1.0$  throughout all experiments.

## V. CONCLUSION

We propose IMAGGarment, a unified conditional diffusion framework for fine-grained garment generation with precise control over silhouette, color, and logo placement. By introducing mixed attention, color adapter, and the  $A^3$  module, our framework explicitly disentangles global structure (silhouette and color) from local attributes (logo content and spatial placement), enabling accurate spatial control and high-quality synthesis. To support this task, we construct GarmentBench, a large-scale benchmark with over 180K samples annotated with multi-level design conditions. Comprehensive experiments on both seen and unseen garments demonstrate that IMAGGarment achieves state-of-the-art results in structure fidelity, color consistency, and logo controllability. Code, models, and datasets are publicly available at <https://github.com/muzishen/IMAGGarment>.

## REFERENCES

- [1] Aijia Zhang, Weiqiang Jia, Qiang Zou, Yixiong Feng, Xiaoxiang Wei, and Ye Zhang. Diffusion-cad: Controllable diffusion model for generating computer-aided design models. *IEEE Transactions on Visualization and Computer Graphics*, 2025.
- [2] Xiongzheng Li, Jing Huang, Jinsong Zhang, Xiaokun Sun, Haibiao Xuan, Yu-Kun Lai, Yingdi Xie, Jingyu Yang, and Kun Li. Learning to infer inner-body under clothing from monocular video. *IEEE Transactions on Visualization and Computer Graphics*, 29(12):5083–5096, 2022.
- [3] Nannan Zhang, Zhenyu Xie, Zhengwentai Sun, Hairui Zhu, Zirong Jin, Nan Xiang, Xiaoguang Han, and Song Wu. Viton-gun: Person-to-person virtual try-on via garment unwrapping. *IEEE Transactions on Visualization and Computer Graphics*, 2025.
- [4] Wen-Yang Zhou, Lu Yuan, Shu-Yu Chen, Lin Gao, and Shi-Min Hu. Lcnerf: Local controllable face generation in neural radiance field. *IEEE Transactions on Visualization and Computer Graphics*, 30(8):5437–5448, 2023.
- [5] Pinaki Nath Chowdhury, Tuanfeng Wang, Duygu Ceylan, Yi-Zhe Song, and Yulia Gryaditskaya. Garment ideation: Iterative view-aware sketch-based garment modeling. In *2022 International Conference on 3D Vision (3DV)*, pages 22–31. IEEE, 2022.
- [6] Yu Jin and Kyungho Lee. Human-ai co-creation in fashion design ideation and sketching: an empirical study. In *Proceedings of IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR), CVFAD Workshop, Seattle, USA, 2024*.
- [7] Funda Durupynar and Ugur Gudukbay. A virtual garment design and simulation system. In *2007 11th International Conference Information Visualization (IV'07)*, pages 862–870. IEEE, 2007.
- [8] Saikrupa PA et al. Smart stitch: A mobile app for personalized garment customization and stitching guidance. In *2025 International Conference on Data Science, Agents & Artificial Intelligence (ICDSAAI)*, pages 1–5. IEEE, 2025.
- [9] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. 2023. arXiv:2307.01952.
- [10] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. 2022. ArXiv:2210.02747.
- [11] Wanchao Su, Hui Ye, Shu-Yu Chen, Lin Gao, and Hongbo Fu. Drawinginstyles: Portrait image generation and editing with spatially conditioned stylegan. *IEEE transactions on visualization and computer graphics*, 29(10):4074–4088, 2022.
- [12] Changjian Chen, Fei Lv, Yalong Guan, Pengcheng Wang, Shengjie Yu, Yifan Zhang, and Zhuo Tang. Human-guided image generation for expanding small-scale training image datasets. *IEEE Transactions on Visualization and Computer Graphics*, 2025.
- [13] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. In *Proceedings of the ACM SIGGRAPH Conference*, pages 1–11, 2023.
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1125–1134, 2017.
- [15] Subhadeep Koley, Ayan Kumar Bhunia, Deeptanshu Sekhri, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. It’s all about your sketch: Democratising sketch control in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7204–7214, 2024.
- [16] Subhadeep Koley, Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Text-to-image diffusion models are great sketch-photo matchmakers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16826–16837, 2024.
- [17] Taewook Kim, Ze Wang, Zhengyuan Yang, Jiang Wang, Lijuan Wang, Zicheng Liu, and Qiang Qiu. Conditional text-to-image generation with reference guidance. 2024. ArXiv:2411.16713.
- [18] Jinghao Zhang, Wen Qian, Hao Luo, Fan Wang, and Feng Zhao. Anylogo: Symbiotic subject-driven diffusion system with gemini status. 2024. ArXiv:2409.17740.
- [19] Mingkang Zhu, Xi Chen, Zhongdao Wang, Hengshuang Zhao, and Jiaya Jia. Logosticker: Inserting logos into diffusion models for customized generation. In *Proceedings of European Conference on Computer Vision*, pages 363–378, 2024.
- [20] Mingzhe Yu, Yunshan Ma, Lei Wu, Changshuo Wang, Xue Li, and Lei Meng. FashionDPO: Fine-tune fashion outfit generation model using direct preference optimization. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 212–222, 2025.
- [21] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. 2023. ArXiv:2308.06721.
- [22] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.
- [23] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn 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.
- [24] Shu-Yu Chen, Wanchao Su, Lin Gao, Shihong Xia, and Hongbo Fu. Deepfacedrawing: deep generation of face images from sketches. *ACM Transactions on Graphics*, 39(4), August 2020.
- [25] Shu-Yu Chen, Feng-Lin Liu, Yu-Kun Lai, Paul L. Rosin, Chungpeng Li, Hongbo Fu, and Lin Gao. Deepfaceediting: deep face generation and editing with disentangled geometry and appearance control. *ACM Transactions on Graphics*, 40(4), July 2021.
- [26] Xian Wu, Chen Wang, Hongbo Fu, Ariel Shamir, Song-Hai Zhang, and Shi-Min Hu. Deepportraitdrawing: Generating human body images from freehand sketches, 2022. ArXiv:2205.02070.
- [27] Arnab Ghosh, Richard Zhang, Puneet K Dokania, Oliver Wang, Alexei A Efros, Philip HS Torr, and Eli Shechtman. Interactive sketch & fill: Multiclass sketch-to-image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1171–1180, 2019.
- [28] Wengling Chen and James Hays. Sketchygan: Towards diverse and realistic sketch to image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9416–9425, 2018.
- [29] Zeyu Li, Cheng Deng, Erkun Yang, and Dacheng Tao. Staged sketch-to-image synthesis via semi-supervised generative adversarial networks. *IEEE Transactions on Multimedia*, 23:2694–2705, 2020.
- [30] Antonia Creswell, Tom White, Vincent Dumoulin, Kai Arulkumaran, Biswa Sengupta, and Anil A Bharath. Generative adversarial networks: An overview. *IEEE Signal Processing Magazine*, 35(1):53–65, 2018.
- [31] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. 2014. ArXiv:1411.1784.
- [32] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 5084–5093, 2020.
- [33] Yifan Liu, Zengchang Qin, Zhenbo Luo, and Hua Wang. Auto-painter: Cartoon image generation from sketch by using conditional generative adversarial networks. 2017. ArXiv:1705.01908.
- [34] Yuanzheng Ci, Xinzhu Ma, Zhihui Wang, Haojie Li, and Zhongxuan Luo. User-guided deep anime line art colorization with conditional adversarial networks. In *Proceedings of the 26th ACM International Conference on Multimedia*, page 1536–1544, 2018.
- [35] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.
- [36] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Proceedings of the Conference on Neural Information Processing Systems*, page 405–415, 2017.
- [37] Aliaksandr Siarohin, Enver Sangineto, Stéphane Lathuilière, and Nicu Sebe. Deformable gans for pose-based human image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3408–3416, 2018.
- [38] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the Conference on Neural Information Processing Systems*, 2020.
- [39] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. 2020. ArXiv:2011.13456.
- [40] Junyao Gao, Yanan Sun, Fei Shen, Xin Jiang, Zhening Xing, Kai Chen, and Cairong Zhao. Faceshot: Bring any character into life. 2025. ArXiv:2503.00740.
- [41] Dongxu Li, Junnan Li, and Steven Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Proceedings of the Conference on Neural Information Processing Systems*, 36:30146–30166, 2023.

- [42] Fei Shen, Xin Jiang, Xin He, Hu Ye, Cong Wang, Xiaoyu Du, Zechao Li, and Jinhui Tang. Imagdressing-v1: Customizable virtual dressing. 2024. ArXiv:2407.12705.
- [43] Ente Lin, Xujie Zhang, Fuwei Zhao, Yuxuan Luo, Xin Dong, Long Zeng, and Xiaodan Liang. Dreamfit: Garment-centric human generation via a lightweight anything-dressing encoder. 2024. ArXiv:2412.17644.
- [44] Weifeng Chen, Tao Gu, Yuhao Xu, and Arlene Chen. Magic clothing: Controllable garment-driven image synthesis. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 6939–6948, 2024.
- [45] Yuhao Xu, Tao Gu, Weifeng Chen, and Chengcai Chen. Ootdiffusion: Outfitting fusion based latent diffusion for controllable virtual try-on. 2024. ArXiv:2403.01779.
- [46] Xujie Zhang, Binbin Yang, Michael C Kampffmeyer, Wenqing Zhang, Shiyue Zhang, Guansong Lu, Liang Lin, Hang Xu, and Xiaodan Liang. Diffcloth: Diffusion based garment synthesis and manipulation via structural cross-modal semantic alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23154–23163, 2023.
- [47] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. 2022. ArXiv:2207.12598.
- [48] Muhammad Hussain. Yolov5, yolov8 and yolov10: The go-to detectors for real-time vision, 2024. ArXiv:2407.02988.
- [49] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023. ArXiv:2308.12966.
- [50] Caroline Chan, Frédo Durand, and Phillip Isola. Learning to generate line drawings that convey geometry and semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7915–7925, 2022.
- [51] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015. ArXiv:1512.03385.
- [52] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967.
- [53] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Proceedings of the Conference on Neural Information Processing Systems*, 30, 2017.
- [54] Kai Zeng, Zhou Wang, Anmin Zhang, Zhaohui Wang, and Wenjun Zhang. A color structural similarity index for image quality assessment. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 660–664, 2014.
- [55] Masato Fujitake. Rl-logo: Deep reinforcement learning localization for logo recognition. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2830–2834. IEEE, 2024.
- [56] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.
- [57] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. ArXiv:1711.05101.
- [58] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022. ArXiv:2010.02502.
- [59] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6593–6602, 2024.