

## MOBILE-VTON: High-Fidelity On-Device Virtual Try-On

Zhenchen Wan<sup>1,2\*</sup> Ce Chen<sup>3\*</sup> Runqi Lin<sup>1</sup> Jiaxin Huang<sup>2</sup> Tianxi Chen<sup>3</sup>  
 Yanwu Xu<sup>4</sup> Tongliang Liu<sup>1,2†</sup> Mingming Gong<sup>2,3†</sup>

<sup>1</sup>University of Sydney <sup>2</sup>MBZUAI <sup>3</sup>University of Melbourne <sup>4</sup>Google

{zwan0681, tongliang.liu}@sydney.edu.au,

ce.chen@student.unimelb.edu.au, mingming.gong@unimelb.edu.au

\* Equal contribution † Corresponding authors

	Ours	CatVTON	Ladi-VTON	StableVTON	IDM-VTON	BooW-VTON
Param	0.41B	0.86B	1.05B	1.84B	2.99B	2.99B
Mobile	✓	✗	✗	✗	✗	✗

Figure 1. Comparison of virtual try-on methods in terms of model size, mobile compatibility, and visual quality. Our model, with only 415M parameters, achieves competitive visual results while running entirely on mobile devices. The leftmost column shows the input person and garment images. All outputs are generated or super-resolved at  $1024 \times 768$  resolution. Zoom in for details.

### Abstract

Virtual try-on (VTON) has recently achieved impressive visual fidelity, but most existing systems require uploading personal photos to cloud-based GPUs, raising privacy concerns and limiting on-device deployment. To address this, we present MOBILE-VTON, a high-quality, privacy-preserving framework that enables fully offline virtual try-on on commodity mobile devices using only a single user image and a garment image. MOBILE-VTON introduces a modular TeacherNet–GarmentNet–TryonNet (TGT) architecture that integrates knowledge distillation, garment-

conditioned generation, and garment alignment into a unified pipeline optimized for on-device efficiency. Within this framework, we propose a Feature-Guided Adversarial (FGA) Distillation strategy that combines teacher supervision with adversarial learning to better match real-world image distributions. GarmentNet is trained with a trajectory-consistency loss to preserve garment semantics across diffusion steps, while TryonNet uses latent concatenation and lightweight cross-modal conditioning to enable robust garment-to-person alignment without large-scale pretraining. By combining these components, MOBILE-VTON achieves high-fidelity generation with low compu-

tational overhead. Experiments on VITON-HD and Dress-Code at 1024×768 show that it matches or outperforms strong server-based baselines while running entirely offline. These results demonstrate that high-quality VTON is not only feasible but also practical on-device, offering a secure solution for real-world applications. Code and project page are available at <https://zhenchenwan.github.io/Mobile-VTON/>.

## 1. Introduction

Virtual Try-On (VTON) technology has emerged as a transformative solution for the fashion and e-commerce industries, enabling users to visualize how garments would look on their bodies without physical trials [9, 15, 27, 36, 37, 44]. Recent advances in large-scale diffusion models have significantly improved visual fidelity and garment detail preservation [1, 3, 7, 14, 26, 39, 43, 47]. However, these models typically rely on cloud-based GPUs, requiring users to upload personal images for inference. This approach not only introduces latency and energy overhead but also raises serious privacy concerns, particularly under stringent data protection regulations. In response, our goal is to design a fully on-device VTON system that avoids any transmission of user data, while still achieving high-fidelity synthesis.

To bridge this gap, we identify three key challenges that limit efficient on-device deployment and guide the design of our system. (i) Diffusion-based VTON models typically have large parameter sizes, resulting in high memory consumption and latency that exceed the capabilities of mobile NPUs and GPUs. (ii) Garment representations can shift across diffusion timesteps, causing semantic drift and texture distortion. These inconsistencies reduce visual coherence, impair identity preservation, and complicate control over garment attributes such as structure, material, and pose. (iii) Existing approaches often rely on models pre-trained on large-scale image datasets before being fine-tuned for VTON tasks. Without such pretraining, these models struggle to acquire strong garment synthesis capabilities, particularly in low-data or minimal-input settings. This reliance on extensive upstream training pipelines limits the flexibility of lightweight architectures and poses a major obstacle to building standalone on-device systems that can be trained directly on the VTON task.

To address these challenges, we propose MOBILE-VTON, a unified, privacy-preserving framework optimized for on-device deployment, built around a TeacherNet–GarmentNet–TryonNet (TGT) architecture. (i) To enable efficient and realistic synthesis under mobile constraints, we introduce a Feature-Guided Adversarial (FGA) Distillation strategy that transfers fine-grained generative knowledge from a high-capacity teacher while leveraging adversarial supervision to align the student model’s outputs with

real-world image distributions. This allows compact student networks to synthesize realistic try-on results under tight mobile compute and memory constraints. (ii) To address semantic instability across diffusion steps, GarmentNet is trained with a trajectory-consistency objective that enforces garment feature invariance over timesteps. This improves robustness to pose and viewpoint changes and reduces semantic drift. (iii) To reduce reliance on large-scale pretraining, we design TryonNet as a garment-aware synthesis module that learns garment–body alignment directly within the virtual try-on task. It incorporates a cross-modal fusion mechanism that jointly encodes person and garment features in latent space, enabling effective training even in data-limited settings.

We evaluate MOBILE-VTON on three benchmark datasets: VITON-HD [3], DressCode [27], and VITON-HD In-the-Wild, which was introduced by BooW-VITON [49] which includes more realistic usage scenarios. All experiments are conducted at a resolution of 1024×768, using only a single person image and a single garment image as input. Despite operating in this minimal-input setting and running fully offline on commodity mobile devices, MOBILE-VTON achieves performance that is competitive with or superior to state-of-the-art server-based baselines, as measured by perceptual and structural metrics including LPIPS, SSIM, and CLIP-I. These results highlight the practical viability of our approach, demonstrating that high-quality virtual try-on can be achieved entirely on-device, with improved responsiveness and enhanced user privacy. In summary, our contributions are as follows:

- We present MOBILE-VTON, the first diffusion-based VTON system, to the best of our knowledge, capable of running entirely on commodity mobile devices without requiring any additional user-provided information beyond the garment and person images.
- We propose a novel TGT framework designed specifically for mobile deployment. We incorporate the proposed FGA Distillation for efficient and realistic generation, a trajectory-consistent GarmentNet to maintain semantic stability across diffusion steps, and a garment-aware TryonNet with perceptual priors, a lightweight adapter, and cross-modal fusion for accurate and texture-preserving garment-to-body alignment. The system is trained from scratch on task-specific data without external pretraining.
- We demonstrate that MOBILE-VTON achieves high visual quality under mobile constraints, matching or outperforming server-based baselines, while ensuring full offline operation and strong privacy guarantees.

## 2. Related Works

**GAN-based VTON.** Classical GAN pipelines [11] perform (i) body structure estimation, (ii) garment warping, and

(iii) image fusion. VITON-GAN and FW-GAN are early two-stage designs coupling geometric alignment with adversarial refinement [7, 19]. Later work improved locality and resolution: InsetGAN’s patch edits [10], HR-VITON’s high-resolution scaling [24], and GP-VTON’s layered parsing/warping for detail preservation [42]. ADGAN and CA-GAN target pose and style controllability and attribute-level edits [23, 25], while GarmentGAN and FitGAN broaden category coverage and fit [30, 32].

**Mask-based Diffusion-based VTON.** Leveraging diffusion’s fidelity [18], recent methods frame try-on as exemplar inpainting or latent diffusion with garment conditioning. DCI-VTON uses a warped-garment prior for detail retention [12]; StableVITON learns clothing–body correspondences in latent space to reduce explicit warping [21]. LaDI-VTON standardizes latent pipelines and pseudo-labeling, and OOTDiffusion injects garment features via self-attention [28, 50]. IDM-VTON adds garment semantics in cross/self-attention for in-the-wild fidelity, while TED-VITON adapts DiT with a garment adapter and text-preservation constraints for crisp logos and prints [4, 36]. StreetTryOn contributes an in-the-wild benchmark and a diffusion+DensePose framework for unpaired training [6].

**Mask-free Diffusion-based VTON.** To avoid errors from clothing masks, mask-free diffusion localizes try-on regions implicitly. BooW-VTON learns from large-scale unpaired data without explicit inpainting masks [49]. MF-VITON achieves high fidelity from only a person image and a garment [38]. Any2AnyTryon expands controllability and category coverage via adaptive position embeddings and a large open-source garment set [13]. Concurrent efforts explore training-free universality (OmniVTON) and direct garment–body correspondence within diffusion, targeting more general and robust try-on [35, 44].

### 3. Methodology

#### 3.1. Model Overview

As illustrated in Fig. 2, our proposed MOBILE-VTON framework introduces a novel modular TeacherNet–GarmentNet–TryonNet (TGT) architecture tailored for high-fidelity, on-device virtual try-on. Each component in the TGT design plays a distinct and complementary role. TeacherNet serves as a high-capacity reference model that guides student training. GarmentNet generates garment-conditioned intermediate features while preserving semantic and temporal consistency. TryonNet synthesizes the final try-on image by aligning person and garment representations through deep fusion.

At the core of this architecture is the Feature-Guided Adversarial (FGA) Distillation strategy, which enables

lightweight student networks to approximate the performance of the teacher model under mobile hardware constraints. FGA integrates two objectives: (i) a feature-level score matching loss that distills generative knowledge from TeacherNet into student U-Nets, enhancing texture fidelity and generation stability; and (ii) an adversarial supervision term that improves photorealism by aligning the student’s output distribution with that of real images.

**TeacherNet** is built upon Stable Diffusion 3.5 Large [8] and remains frozen during training. It serves as a high-capacity generative model that provides supervisory signals for guiding lightweight student networks. Instead of performing explicit pixel-level regression, TeacherNet is treated as a score function oracle that provides gradient-based supervision. During training, given an input pair consisting of a person image  $\mathbf{X}_p$  and a garment image  $\mathbf{X}_g$ , we first encode it into a latent space using a shared VAE encoder  $\mathcal{E}$ . Here,  $\mathbf{X}_p$  is used to supervise the student TryonNet, while  $\mathbf{X}_g$  provides supervision for the student GarmentNet. Gaussian noise is then added at a random timestep  $t$ , and the resulting noisy latent  $\tilde{\mathbf{z}}^{(t)}$  is passed to the frozen teacher denoiser  $D_t$ . The TeacherNet outputs denoising predictions and provides score estimates  $s_{\text{true}}(\tilde{\mathbf{z}}^{(t)}, t)$ , representing the gradients of the log-probability with respect to the input noise. These scores serve as ground-truth supervision for guiding the student networks.

**Light-UNets** implement the student-side architecture for this distillation framework. They consist of two lightweight networks, GarmentNet and TryonNet, adapted from the SnapGen architecture [2] for efficient on-device inference. To match the conditioning capabilities of the teacher, we integrate mobile-optimized attention and dual text encoders (CLIP-G [31] and CLIP-L [31]). Both GarmentNet and TryonNet operate on noisy latent representations  $\tilde{\mathbf{z}}^{(t)}$  at various diffusion timesteps  $t$ , producing denoised estimates and corresponding score functions  $s_{\text{fake}}(\tilde{\mathbf{z}}^{(t)}, t)$  that represent the student’s learned distribution. These outputs define the “fake” distribution in the distillation setup and are directly compared to the teacher’s score estimates using a DMD2-style loss [46]. This approach eliminates the need for trajectory-level regression and supports lightweight generation without compromising visual fidelity.

**Light-Adapter** is the unified image-to-text conditioning module. It adopts the decoupled cross-attention mechanism from IP-Adapter [45]. For improved mobile efficiency, we replace the large CLIP visual encoder with DINOv2-base [29], a compact yet semantically rich alternative. Given a garment image  $\mathbf{X}_g$ , visual features are extracted via DINOv2 and projected into key and value tensors using learnable linear layers:

$$(\mathbf{K}_d, \mathbf{V}_d) = \text{Proj}(\text{DINOv2}(\mathbf{X}_g)). \quad (1)$$




Figure 2. **Training overview of MOBILE-VTON.** The left side illustrates our Feature-Guided Adversarial (FGA) Distillation process, where a high-capacity TeacherNet supervises two lightweight student networks: GarmentNet and TryonNet. GarmentNet learns to reconstruct the garment using feature-level distillation loss  $\mathcal{L}_{\text{feature}}^G$ , while TryonNet is guided by both feature-level loss  $\mathcal{L}_{\text{feature}}^T$  and adversarial loss  $\mathcal{L}_{\text{GAN}}$ . The right side shows the main training pipeline of MOBILE-VTON, where TryonNet and GarmentNet are jointly trained with garment-aware supervision. The Light-Adapter injects garment semantics, and the model is optimized using latent concatenation, temporal consistency loss  $\mathcal{L}_{\text{cons}}$ , and reconstruction loss  $\mathcal{L}_{\text{Diff}}$  to generate high-fidelity try-on images.

These tensors encode the garment’s visual semantics and are used as conditioning inputs for TryonNet during cross-attention, providing detailed guidance for garment-to-person alignment and synthesis.

### 3.2. Feature-Guided Adversarial Distillation (FGA)

Inspired by DMD2 [46], we propose a supervision strategy that combines two complementary objectives: a feature-level distillation loss and an adversarial realism loss. Together, they enable the student networks to generate high-quality, realistic outputs while remaining lightweight enough for mobile deployment.

**Feature-Level Distillation.** To align the student’s generative behavior with that of the teacher, we adopt a score-based distillation objective. Rather than regressing pixel values or denoised outputs directly, we align the score functions predicted by the teacher and student networks at each diffusion timestep. At a given timestep  $t$ , the input image (either garment or person) is encoded into latent space via a shared VAE encoder  $\mathcal{E}$ , and Gaussian noise is added to produce the noisy latent  $\tilde{\mathbf{z}}^{(t)}$ . This latent is then passed through both the teacher denoiser  $D_t$  and the student denoiser  $D_s$  to obtain their respective score estimates:  $s_{\text{true}}(\tilde{\mathbf{z}}^{(t)}, t)$  from TeacherNet and  $s_{\text{fake}}(\tilde{\mathbf{z}}^{(t)}, t)$  from Light-UNets. The student is trained by minimizing the  $\ell_2$  distance between these two score functions, defined as:

$$\mathcal{L}_{\text{feature}} = \mathbb{E}_t \left\| s_{\text{true}}(\tilde{\mathbf{z}}^{(t)}, t) - s_{\text{fake}}(\tilde{\mathbf{z}}^{(t)}, t) \right\|_2^2. \quad (2)$$

This score-based supervision allows the lightweight student networks to approximate the teacher’s distributional behavior in latent space, thereby improving sample realism and generation stability under mobile constraints.

**Adversarial Realism Enhancement.** To further improve visual fidelity, we introduce an adversarial loss that encourages the student model to generate perceptually natural images. A lightweight discriminator  $D$  is trained to distinguish real person images  $X$  from generated images  $\hat{X}$ . TryonNet is optimized to produce outputs that can fool the discriminator, resulting in more photorealistic results. The adversarial loss is defined as:

$$\mathcal{L}_{\text{GAN}} = \mathbb{E}_{X \sim \mathcal{R}} [\log D(X)] + \mathbb{E}_{\hat{X} \sim \mathcal{G}} [\log(1 - D(\hat{X}))], \quad (3)$$

where  $\mathcal{R}$  and  $\mathcal{G}$  represent the distributions of real and generated images, respectively. This realism-oriented supervision complements the distillation loss by guiding the student model toward sharper and more coherent outputs.

### 3.3. Trajectory-Consistent GarmentNet (TCG)

Unlike conventional diffusion-based generators, GarmentNet in our framework is designed not for image synthesis but for extracting consistent garment features across diffusion timesteps. During training, we observe that naively perturbing garment inputs with noise often leads to semantic drift. Specifically, the garment representation may shift across timesteps, resulting in unstable textures, distorted

shapes, or inconsistent conditioning signals. These issues reduce reliability and coherence in the final try-on results.

To address this, we introduce a trajectory consistency constraint that enforces temporal stability in the garment encoding process. Rather than injecting noise during training, we apply the diffusion process deterministically at each timestep  $t \in [1, T]$  and require the model to reconstruct the original garment image consistently across all steps. This is formalized as the following objective:

$$\mathcal{L}_{\text{cons}} = \mathbb{E}_{t \sim [1, T]} \left[ \left\| \hat{X}_g^{(t)} - X_g \right\|_2^2 \right], \quad (4)$$

where  $\hat{X}_g^{(t)}$  is the predicted garment reconstruction at timestep  $t$ , and  $X_g$  is the original garment input. By minimizing this loss, the network is encouraged to produce garment features that remain stable along the diffusion trajectory. This temporal regularization improves the consistency of garment semantics, enabling more robust alignment and rendering in the downstream TryonNet module.

### 3.4. Garment-Aware TryonNet

The TryonNet module is responsible for synthesizing the final virtual try-on image by aligning garment and person representations through deep fusion. Building on ideas from IDM-VTON [4] and CatVTON [5], our design enhances garment-to-body interaction by integrating multi-level garment features and conditioning priors throughout the generation pipeline. Unlike previous methods, where TryonNet-style modules are typically initialized from large-scale pre-trained models (e.g., text-to-image diffusion), our TryonNet is trained from scratch on the downstream virtual try-on task. In this pretraining-free setting, the model lacks prior exposure to garment semantics, which may impair conditioning quality. To address this, we follow a strategy inspired by CatVTON [5], where garment and person images are spatially concatenated into a unified input. This design explicitly injects garment information into the synthesis pathway, enabling the model to learn alignment directly.

**Latent Concatenation (LC).** To promote accurate spatial alignment between the person and the garment, we first construct a height-wise concatenated image  $X_{\text{concat}} = \text{Concat}_{\text{height}}(X_p, X_g)$ , and obtain its clean latent representation via a shared VAE encoder as  $\mathcal{E}(X_{\text{concat}})$ , which is then perturbed with diffusion noise to form the noisy latent  $\mathbf{z}_{\text{concat}}$ . In parallel, we introduce an additional reference input  $X_{\text{condi}} = \text{Concat}_{\text{height}}(X_t, X_g)$ , which is similarly encoded to  $\mathcal{E}(X_{\text{condi}})$  to guide the model in preserving both person identity and garment appearance. We then concatenate the two latent codes along the batch dimension to form the final input to TryonNet:  $\mathbf{z}_{\text{input}} = \text{Concat}_{\text{batch}}(\mathbf{z}_{\text{concat}}, \mathcal{E}(X_{\text{condi}}))$ . This combined latent, along with fused conditioning tokens  $(\mathbf{K}_{\text{all}}, \mathbf{V}_{\text{all}})$ , is processed

through all denoising steps to produce a refined latent representation  $\mathcal{E}(\hat{X}_{\text{concat}})$ . Before decoding, we crop the top region of  $\mathcal{E}(\hat{X}_{\text{concat}})$  to isolate the person-relevant features, denoted as  $\mathcal{E}(\hat{X}_t) = \text{Crop}_{\text{top}}(\mathcal{E}(\hat{X}_{\text{concat}}))$ , and decode it via a lightweight VAE decoder  $\mathcal{D}(\cdot)$  [2] to reconstruct the final output image as  $\hat{X}_t = \mathcal{D}(\mathcal{E}(\hat{X}_t))$ .

**Garment-Aware Reconstruction.** To preserve garment details and spatial coherence, we define a reconstruction loss on the full concatenated image. The model is supervised to match the input pair  $(X_p, X_g)$  through a reconstruction of the combined latent representation:

$$\mathcal{L}_{\text{Diff}} = \left\| \hat{X}_{\text{concat}} - X_{\text{concat}} \right\|_2^2, \quad (5)$$

where  $X_{\text{concat}} = \text{Concat}_{\text{height}}(X_p, X_g)$  serves as the target image. This supervision enforces global structure preservation and ensures that garment textures and positions are faithfully retained in the final output.

**Feature Fusion in each layer.** In addition to latent concatenation, TryonNet integrates garment semantics from two complementary sources: (1) multi-scale garment features  $\mathbf{F}_g^{(i)}$  extracted from GarmentNet’s self-attention layers, and (2) visual key–value pairs  $(\mathbf{K}_d, \mathbf{V}_d)$  derived from the Light-Adapter. At each self-attention layer  $i$  of TryonNet’s U-Net backbone, the corresponding garment feature  $\mathbf{F}_g^{(i)}$  is concatenated with the intermediate hidden state  $\mathbf{H}_{\text{Tryon}}^{(i)}$  prior to the self-attention operation, forming an augmented representation  $\mathbf{H}_{\text{aug}}^{(i)}$ . This enables the model to jointly attend to both garment-specific structure and person-specific context across spatial scales, thereby improving alignment accuracy and enhancing the fidelity of garment details across layers.

In parallel, cross-attention is computed using two distinct branches of key–value pairs: one from the text prompt encoder  $(\mathbf{K}_{\text{text}}, \mathbf{V}_{\text{text}})$  and one from the Light-Adapter  $(\mathbf{K}_d, \mathbf{V}_d)$ . Following the fusion strategy of IP-Adapter [45], the output attention features are computed as:

$$\mathbf{Z} = \text{Attention}(\mathbf{Q}, \mathbf{K}_{\text{text}}, \mathbf{V}_{\text{text}}) + \text{Attention}(\mathbf{Q}, \mathbf{K}_d, \mathbf{V}_d), \quad (6)$$

where  $\mathbf{Q}$  is the query projected from the hidden state. This dual-branch conditioning allows TryonNet to integrate textual guidance and visual garment semantics in a balanced manner, enhancing control over garment appearance and ensuring consistent garment-to-body alignment.

### 3.5. Training Objective

All loss components described above are jointly optimized under a unified training objective tailored to the Feature-Guided Adversarial (FGA) framework. Each module is supervised by a set of objectives aligned with its functional role within the overall TGT architecture (Fig. 2).



Figure 3. Qualitative comparison on VITON-HD and DressCode between our approach and server-side baselines. Zoom in for finer details.

**GarmentNet Losses.** GarmentNet is optimized using two objectives: (i) a feature-level distillation loss  $\mathcal{L}_{\text{feature}}^G$  that aligns its generative behavior with TeacherNet via score matching (Eq. 2), and (ii) a trajectory consistency loss  $\mathcal{L}_{\text{cons}}$  that ensures semantic stability of garment features across diffusion steps (Eq. 4). The total GarmentNet objective is:

$$\mathcal{L}_{\text{GarmentNet}} = \lambda_1 \mathcal{L}_{\text{feature}}^G + \lambda_2 \mathcal{L}_{\text{cons}}. \quad (7)$$

**TryonNet Losses.** TryonNet is optimized using three complementary supervision signals: (i) a feature-level distillation loss  $\mathcal{L}_{\text{feature}}^T$  that aligns its generative score estimates with those of the teacher network (Eq. 2), (ii) an adversarial loss  $\mathcal{L}_{\text{GAN}}$  that encourages photorealistic synthesis through a lightweight discriminator (Eq. 3), and (iii) a garment-aware reconstruction loss  $\mathcal{L}_{\text{Diff}}$  that enforces spatial alignment and preserves fine-grained garment details (Eq. 5). The final training objective for TryonNet is defined as:

$$\mathcal{L}_{\text{TryonNet}} = \mathcal{L}_{\text{Diff}} + \lambda_1 \mathcal{L}_{\text{feature}}^T + \lambda_3 \mathcal{L}_{\text{GAN}}. \quad (8)$$

**Final Objective.** The total loss for training MOBILE-VTON integrates the objectives of both sub-networks:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{GarmentNet}} + \mathcal{L}_{\text{TryonNet}}, \quad (9)$$

where  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are hyperparameters balancing distillation, consistency, and realism. This unified objective enables both GarmentNet and TryonNet to learn complementary, garment-aware representations for high-quality virtual try-on, without requiring large-scale pretraining.

## 4. Experiments

### 4.1. Experiment Setup

**Baselines.** We compare our mobile-side, mask-free method MOBILE-VTON against both server-side, mask-based and mask-free state-of-the-art approaches. Mask-based baselines include SD-VITON [33], LaDI-VTON [28], StableVITON [21], and IDM-VTON [4], which require segmentation masks during inference. Mask-free baselines such as CatVTON [5] and BooW-VTON [49] rely entirely on diffusion-based generation without mask supervision. We follow a standard unpaired evaluation protocol and generate try-on images at  $1024 \times 768$  resolution, where supported. For lower-resolution outputs (e.g.,  $512 \times 384$ ), we apply bilinear interpolation or Real-ESRGAN [40] to ensure consistent output size and quality across methods.

Dataset Method	VITON-HD					DressCode Upper-body					Memory Usage(G)	On Mobile
	LPIPS↓	SSIM↑	CLIP-I↑	FID↓ UN	KID↓ UN	LPIPS↓	SSIM↑	CLIP-I↑	FID↓ UN	KID↓ UN		
<b>Server-side Mask-based methods</b>												
SD-VITON [33]	0.104	<b>0.890</b>	0.831	9.857	1.450	0.201	0.858	0.555	50.755	38.331	-	✗
LaDI-VTON [28]	0.166	0.873	0.819	9.386	1.590	0.157	0.905	0.773	22.689	2.580	7.63	✗
StableVITON [21]	0.142	0.875	0.838	9.371	1.990	0.113	0.910	0.844	19.712	2.149	13.84	✗
IDM-VTON [4]	<b>0.102</b>	0.868	<b>0.875</b>	<b>9.156</b>	<b>1.242</b>	0.065	0.920	<b>0.870</b>	<b>11.852</b>	<b>1.181</b>	18.47	✗
<b>Server-side Mask-Free methods</b>												
CatVTON [5]	0.161	0.872	0.832	11.275	2.847	0.0919	0.906	0.856	14.507	2.511	<b>5.80</b>	✗
BooW-VTON [49]	0.107	0.864	<b>0.852</b>	<b>9.036</b>	<b>1.011</b>	<b>0.0513</b>	<b>0.928</b>	<b>0.869</b>	<b>11.361</b>	<b>1.043</b>	18.47	✗
<b>Mobile-side Mask-Free methods</b>												
<b>Ours</b>	<b>0.088</b>	<b>0.893</b>	0.833	10.211	2.023	<b>0.053</b>	<b>0.935</b>	0.845	12.775	1.917	<b>2.84</b>	✓

Table 1. Quantitative comparison of models trained and evaluated on VITON-HD and DressCode upper-body datasets. **Bold** and underline indicate the best and the second best results, respectively. “UN” denotes the unpaired setting. FID score is multiplied by 100.

Dataset Method	VITON-HD In-the-wild				
	LPIPS↓	SSIM↑	CLIP-I↑	FID↓ UN	KID↓ UN
<b>Server-side Mask-based method</b>					
SD-VITON	0.174	0.857	0.804	17.151	5.326
StableVITON	0.201	0.863	0.848	12.983	2.804
IDM-VTON	<b>0.137</b>	0.832	<b>0.931</b>	<b>11.373</b>	<b>1.133</b>
<b>Server-side Mask-free method</b>					
CatVTON	0.139	<b>0.881</b>	<b>0.921</b>	12.118	1.736
BooW-VTON	0.148	0.864	0.842	<b>9.256</b>	<b>1.111</b>
<b>Mobile-side Mask-free method</b>					
<b>Ours</b>	<b>0.133</b>	<b>0.867</b>	0.902	12.095	1.503

Table 2. Quantitative comparison of models on the VITON-HD In-the-Wild dataset. “UN” indicated the unpaired setting. FID score is multiplied by 100.

**Evaluation Datasets.** We evaluate the performance of MOBILE-VTON on three datasets: VITON-HD [3], DressCode [27], and an In-the-Wild variant of VITON-HD. The third dataset is constructed following the protocol introduced in BooW-VITON [49], enabling evaluation under real-world conditions. VITON-HD contains 13,679 image–garment pairs, consisting of frontal-view images of female subjects and their corresponding upper-body clothing. Following established practice [4, 12, 21, 28, 34], we use 11,647 pairs for training and 2,032 for testing in both the original and In-the-Wild versions. DressCode includes 15,366 image–garment pairs focused on upper-body try-on. According to the official protocol, we select 1,800 samples from the upper-body subset for evaluation. All experiments are conducted at a resolution of  $1024 \times 768$  to ensure comparability with existing methods.

**Evaluation Metrics.** We assess MOBILE-VTON in both unpaired and paired evaluation settings, following standard protocols in previous virtual try-on studies [3–5, 28, 36, 38, 49]. In the unpaired setting, where the input garment differs from the original, we evaluate image realism and distributional similarity using Fréchet Inception Distance (FID) [17] and Kernel Inception Distance (KID) [20]. This setup closely reflects real-world application scenarios. In the paired setting, where the target garment matches the original, we measure visual fidelity using three metrics: Structural Similarity Index (SSIM) [41] for struc-

tural consistency, Learned Perceptual Image Patch Similarity (LPIPS) [48] for perceptual similarity, and CLIP image similarity (CLIP-I) [16] for semantic alignment.

**Implementation Details** All experiments were conducted using 8 NVIDIA A100 GPUs (80GB). As the official implementation of SnapGen [2] was not publicly available, we re-implemented the architecture based on the descriptions provided in the original paper. We adapted the model to the downstream MOBILE-VTON task using a two-stage training process. In the first stage, the model was trained for 140 epochs on a merged dataset combining the upper-body subsets of DressCode [27] and VITON-HD [3]. In the second stage, it was fine-tuned for 100 additional epochs on the higher-quality DressCode subset to further improve consistency and quality. All training stages used a batch size of 256 and the AdamW optimizer [22] with a learning rate of  $1e-4$  in the first stage and  $5e-5$  in the second. The optimizer parameters were set to  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ . The overall loss function was weighted using  $\lambda_1 = 1e-2$ ,  $\lambda_2 = 0.5$ , and  $\lambda_3 = 5e-3$  throughout both stages.

## 4.2. Qualitative Results

Figure 1 compares visual results on the VITON-HD In-the-Wild benchmark. Compared to server-based baselines such as CatVTON, LaDI-VTON, StableVITON, IDM-VTON, and BooW-VTON, our method produces competitive texture details and stable garment boundaries, with improved alignment in challenging regions such as logos and sleeves. Figure 3 shows results across varied poses and garment types on DressCode and VITON-HD. MOBILE-VTON consistently preserves garment structure and fabric appearance under pose and style variations. It accurately captures details such as lace, color transitions, and fabric deformation, where other methods often introduce smoothing or misalignment. Despite its compact size (0.41B parameters), MOBILE-VTON delivers high visual fidelity while being the only method that runs fully on mobile devices. Competing methods typically contain 2–7 times more parameters and require server-class GPUs for inference. These comparisons highlight both the visual quality and practical deploy-



Figure 4. Comparison of virtual try-on results with and without the TCG and LC. Zoom in for more details.

TCG	LC	LPIPS↓	SSIM↑	CLIP-I↑	FID↓ UN	KID↓ UN
✗	✗	0.119	0.874	0.798	11.231	3.136
✓	✗	0.111	0.879	0.805	10.814	2.744
✓	✓	<b>0.088</b>	<b>0.893</b>	<b>0.833</b>	<b>10.211</b>	<b>2.023</b>

Table 3. Quantitative comparison of models trained on the VITON-HD dataset with components. “UN” indicated the unpaired setting. FID score is multiplied by 100.

ability of our approach for real-world try-on applications.

### 4.3. Quantitative Results

As shown in Table 1, our method achieves the best overall SSIM and the second-lowest LPIPS on the DressCode Upper-body dataset, indicating strong structural preservation and perceptual quality. On VITON-HD, MOBILE-VTON ranks among the top performers across LPIPS, SSIM, and CLIP-I, while maintaining competitive FID and KID scores. Table 2 further reports results on the VITON-HD In-the-Wild benchmark. Our method performs on par with or better than server-based models in LPIPS and SSIM, while using significantly less memory. As reported in Table 1, MOBILE-VTON requires only 2.84 GB of memory, compared to 5 to 18 GB baselines. Importantly, most server-side baselines operate using segmentation masks to restrict generation to the garment region. In contrast, our model is completely mask free and must synthesize the entire image, including the body, clothing, and background, without any spatial priors. This makes realism-oriented metrics such as FID and KID considerably more difficult. Despite this, MOBILE-VTON still achieves competitive performance, demonstrating its robustness, generalization ability, and practical deployability in real-world settings.

### 4.4. Ablation Study

**Effect of Trajectory-Consistent GarmentNet (TCG).** As shown in Table 3, adding the TCG module improves all metrics on the VITON-HD dataset. LPIPS drops from 0.119

to 0.111, SSIM increases from 0.874 to 0.879, and CLIP-I improves from 0.798 to 0.805. FID and KID also decrease, indicating better image realism and distributional alignment. Figure 4 shows a visual comparison between models trained with and without TCG. Without TCG (column a), the model struggles to preserve spatial correspondence between garment colors and body structure. Critical elements such as stripes and logos appear faded, misaligned, or missing. For example, the red chevron in the top row is broken and misplaced, while the “FILA” logo in the second row is distorted and lacks contrast. In contrast, enabling TCG (column b) improves garment-to-body alignment. The model better preserves original garment colors and their spatial layout, producing outputs that more closely reflect the source design. Logos and stripe patterns are sharper and better positioned, indicating that TCG improves semantic stability and helps the network maintain color-to-location consistency across poses.

**Effect of Latent Concatenation (LC).** Building on TCG, adding LC further improves all metrics: LPIPS drops to 0.088, SSIM increases to 0.893, and CLIP-I rises to 0.833. FID and KID also decrease, indicating higher overall realism. This performance gain stems from LC’s ability to directly inject garment information into TryonNet, enabling it to synthesize garments from scratch rather than relying on pretraining. As shown in Figure 4 (column c vs. b), LC enhances garment generation quality and structural alignment. In the top row, stripe placement, sleeve length, and multi-color patterns better match the reference garment. In the second row, the pink-black-white color blocking and the “FILA” logo are rendered with improved clarity and placement. These results show that LC complements TCG by providing TryonNet with explicit garment geometry and appearance cues, overcoming the lack of large-scale pretraining and achieving final, high-fidelity synthesis.

## 5. Conclusion

We introduced MOBILE-VTON, a unified distillation-driven framework that enables high-fidelity virtual try-on directly on mobile devices. Built on the proposed TGT architecture, our approach transfers both generative behaviors and perceptual priors from a large diffusion teacher to lightweight student networks. GarmentNet extracts temporally consistent garment features, while TryonNet aligns garments with person representations through structured attention and latent concatenation. A lightweight Light-Adapter further facilitates efficient garment conditioning, and the training objective integrates distillation, adversarial supervision, and trajectory consistency. Overall, MOBILE-VTON demonstrates that diffusion-based try-on systems can achieve strong visual fidelity while remaining efficient enough for fully on-device deployment.

## References

- [1] Shuai Bai, Huiling Zhou, Zhikang Li, Chang Zhou, and Hongxia Yang. Single Stage Virtual Try-on via Deformable Attention Flows, 2022. arXiv:2207.09161 [cs]. 2
- [2] Jierun Chen, Dongting Hu, Xijie Huang, Huseyin Coskun, Arpit Sahni, Aarush Gupta, Anujraaj Goyal, Dishani Lahiri, Rajesh Singh, Yerlan Idelbayev, Junli Cao, Yanyu Li, Kwang-Ting Cheng, S.-H. Gary Chan, Mingming Gong, Sergey Tulyakov, Anil Kag, Yanwu Xu, and Jian Ren. SnapGen: Taming High-Resolution Text-to-Image Models for Mobile Devices with Efficient Architectures and Training. pages 7997–8008, 2025. 3, 5, 7
- [3] Seunghwan Choi, Sunghyun Park, Minsoo Lee, and Jaegul Choo. VITON-HD: High-Resolution Virtual Try-On via Misalignment-Aware Normalization. pages 14131–14140, 2021. 2, 7
- [4] Yisol Choi, Sangkyung Kwak, Kyungmin Lee, Hyungwon Choi, and Jinwoo Shin. Improving Diffusion Models for Authentic Virtual Try-on in the Wild, 2024. arXiv:2403.05139 [cs]. 3, 5, 6, 7
- [5] Zheng Chong, Xiao Dong, Haoxiang Li, Shiyue Zhang, Wenqing Zhang, Xujie Zhang, Hanqing Zhao, and Xiaodan Liang. CatVTON: Concatenation Is All You Need for Virtual Try-On with Diffusion Models, 2024. arXiv:2407.15886 [cs]. 5, 6, 7
- [6] Aiyu Cui, Jay Mahajan, Viraj Shah, Preeti Gomathinayagam, Chang Liu, and Svetlana Lazebnik. Street TryOn: Learning In-the-Wild Virtual Try-On from Unpaired Person Images. pages 8235–8239, 2024. 3
- [7] Haoye Dong, Xiaodan Liang, Xiaohui Shen, Bowen Wu, Bing-Cheng Chen, and Jian Yin. FW-GAN: Flow-Navigated Warping GAN for Video Virtual Try-On. pages 1161–1170, 2019. 2, 3
- [8] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, Kyle Lacey, Alex Goodwin, Yan-nik Marek, and Robin Rombach. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis, 2024. arXiv:2403.03206 [cs]. 3
- [9] Benjamin Fele, Ajda Lampe, Peter Peer, and Vitomir Struc. C-VTON: Context-Driven Image-Based Virtual Try-On Network. pages 3144–3153, 2022. 2
- [10] Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J. Mitra, Peter Wonka, and Jingwan Lu. InsetGAN for Full-Body Image Generation. pages 7723–7732, 2022. 3
- [11] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Networks, 2014. arXiv:1406.2661 [stat]. 2
- [12] Junhong Gou, Siyu Sun, Jianfu Zhang, Jianlou Si, Chen Qian, and Liqing Zhang. Taming the Power of Diffusion Models for High-Quality Virtual Try-On with Appearance Flow. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 7599–7607, New York, NY, USA, 2023. Association for Computing Machinery. 3, 7
- [13] Hailong Guo, Bohan Zeng, Yiren Song, Wentao Zhang, Chuang Zhang, and Jiaming Liu. Any2AnyTryon: Leveraging Adaptive Position Embeddings for Versatile Virtual Clothing Tasks, 2025. arXiv:2501.15891 [cs]. 3
- [14] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S. Davis. VITON: An Image-based Virtual Try-on Network, 2018. arXiv:1711.08447 [cs]. 2
- [15] Zijian He, Yuwei Ning, Yipeng Qin, Wangrun Wang, Sibe Yang, Liang Lin, and Guanbin Li. VTON 360: High-Fidelity Virtual Try-On from Any Viewing Direction, 2025. arXiv:2503.12165 [cs]. 2
- [16] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528, Online and Punta Cana, Dominican Republic, 2021. Association for Computational Linguistics. 7
- [17] 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, 2018. arXiv:1706.08500 [cs]. 7
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models, 2020. arXiv:2006.11239 [cs]. 3
- [19] Shion Honda. VITON-GAN: Virtual Try-on Image Generator Trained with Adversarial Loss. *Eurographics 2019 - Posters*, page 2 pages, 2019. arXiv:1911.07926 [cs]. 3
- [20] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-GAT-IT: Unsupervised Generative Attentional Networks with Adaptive Layer-Instance Normalization for Image-to-Image Translation, 2020. arXiv:1907.10830 [cs]. 7
- [21] Jeongho Kim, Gyojung Gu, Minh Park, Sunghyun Park, and Jaegul Choo. StableVITON: Learning Semantic Correspondence with Latent Diffusion Model for Virtual Try-On, 2023. arXiv:2312.01725 [cs]. 3, 6, 7
- [22] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization, 2017. arXiv:1412.6980 [cs]. 7
- [23] Robin Kips, Pietro Gori, Matthieu Perrot, and Isabelle Bloch. CA-GAN: Weakly Supervised Color Aware GAN for Controllable Makeup Transfer. pages 280–296. 2020. arXiv:2008.10298 [cs]. 3
- [24] Sangyun Lee, Gyojung Gu, Sunghyun Park, Seunghwan Choi, and Jaegul Choo. High-Resolution Virtual Try-On with Misalignment and Occlusion-Handled Conditions, 2022. arXiv:2206.14180 [cs]. 3
- [25] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable Person Image Synthesis with Attribute-Decomposed GAN, 2020. arXiv:2003.12267 [cs]. 3
- [26] Matiu Rahman Minar, T. Tuan, Heejune Ahn, Paul L. Rosin, and Yu-Kun Lai. CP-VTON+: Clothing Shape and Texture Preserving Image-Based Virtual Try-On. 2020. 2
- [27] Davide Morelli, Matteo Fincato, Marcella Cornia, Federico Landi, Fabio Cesari, and Rita Cucchiara. Dress Code: High-Resolution Multi-Category Virtual Try-On. pages 2231–2235, 2022. 2, 7

- [28] Davide Morelli, Alberto Baldrati, Giuseppe Cartella, Marcella Cornia, Marco Bertini, and Rita Cucchiara. LaDIVTON: Latent Diffusion Textual-Inversion Enhanced Virtual Try-On. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 8580–8589, New York, NY, USA, 2023. Association for Computing Machinery. [3](#), [6](#), [7](#)
- [29] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning Robust Visual Features without Supervision, 2024. arXiv:2304.07193 [cs]. [3](#)
- [30] Sonia Pecenkova, Nour Kaessli, and Reza Shirvany. FitGAN: Fit- and Shape-Realistic Generative Adversarial Networks for Fashion, 2022. arXiv:2206.11768 [cs]. [3](#)
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. ISSN: 2640-3498. [3](#)
- [32] Amir Hossein Raffiee and Michael Sollami. GarmentGAN: Photo-realistic Adversarial Fashion Transfer, 2020. arXiv:2003.01894 [cs]. [3](#)
- [33] Sang-Heon Shim, Jiwoo Chung, and Jae-Pil Heo. Towards Squeezing-Averse Virtual Try-On via Sequential Deformation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(5):4856–4863, 2024. Number: 5. [6](#), [7](#)
- [34] Siqi Wan, Yehao Li, Jingwen Chen, Yingwei Pan, Ting Yao, Yang Cao, and Tao Mei. Improving Virtual Try-On with Garment-focused Diffusion Models, 2024. [7](#)
- [35] Siqi Wan, Jingwen Chen, Yingwei Pan, Ting Yao, and Tao Mei. Incorporating Visual Correspondence into Diffusion Model for Virtual Try-On, 2025. arXiv:2505.16977 [cs]. [3](#)
- [36] Zhenchen Wan, Yanwu Xu, Zhaoqing Wang, Feng Liu, Tongliang Liu, and Mingming Gong. TED-VITON: Transformer-Empowered Diffusion Models for Virtual Try-On, 2024. arXiv:2411.17017 [cs]. [2](#), [3](#), [7](#)
- [37] Zhenchen Wan, Yanwu Xu, Dongting Hu, Weilun Cheng, Tianxi Chen, Zhaoqing Wang, Feng Liu, Tongliang Liu, and Mingming Gong. Mft-viton: High-fidelity virtual try-on with minimal input via a mask-free transformer-diffusion model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1985–1994, 2025. [2](#)
- [38] Zhenchen Wan, Yanwu xu, Dongting Hu, Weilun Cheng, Tianxi Chen, Zhaoqing Wang, Feng Liu, Tongliang Liu, and Mingming Gong. MF-VITON: High-Fidelity Mask-Free Virtual Try-On with Minimal Input, 2025. arXiv:2503.08650 [cs]. [3](#), [7](#)
- [39] Bochao Wang, Huabin Zheng, Xiaodan Liang, Yimin Chen, Liang Lin, and Meng Yang. Toward Characteristic-Preserving Image-based Virtual Try-On Network. pages 589–604, 2018. [2](#)
- [40] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-ESRGAN: Training Real-World Blind Super-Resolution With Pure Synthetic Data. pages 1905–1914, 2021. [6](#)
- [41] Zhou 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. [7](#)
- [42] Zhenyu Xie, Zaiyu Huang, Xin Dong, Fuwei Zhao, Haoye Dong, Xijin Zhang, Feida Zhu, and Xiaodan Liang. GP-VTON: Towards General Purpose Virtual Try-On via Collaborative Local-Flow Global-Parsing Learning. pages 23550–23559, 2023. [3](#)
- [43] Han Yang, Ruimao Zhang, Xiaobao Guo, Wei Liu, Wangmeng Zuo, and Ping Luo. Towards Photo-Realistic Virtual Try-On by Adaptively Generating  $\leftarrow$  Preserving Image Content, 2020. arXiv:2003.05863 [cs, eess]. [2](#)
- [44] Zhaotong Yang, Yuhui Li, Shengfeng He, Xinzhe Li, Yangyang Xu, Junyu Dong, and Yong Du. OmniVTON: Training-Free Universal Virtual Try-On, 2025. arXiv:2507.15037 [cs]. [2](#), [3](#)
- [45] 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 [cs]. [3](#), [5](#)
- [46] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T. Freeman. Improved Distribution Matching Distillation for Fast Image Synthesis, 2024. arXiv:2405.14867 [cs]. [3](#), [4](#)
- [47] Ruiyun Yu, Xiaoqi Wang, and Xiaohui Xie. VTNFP: An Image-Based Virtual Try-On Network With Body and Clothing Feature Preservation. pages 10511–10520, 2019. [2](#)
- [48] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric, 2018. arXiv:1801.03924 [cs]. [7](#)
- [49] Xuanpu Zhang, Dan Song, Pengxin Zhan, Tianyu Chang, Jianhao Zeng, Qingguo Chen, Weihua Luo, and Anan Liu. BooW-VTON: Boosting In-the-Wild Virtual Try-On via Mask-Free Pseudo Data Training, 2024. arXiv:2408.06047 [cs]. [2](#), [3](#), [6](#), [7](#)
- [50] Luyang Zhu, Dawei Yang, Tyler Zhu, Fitsum Reda, William Chan, Chitwan Saharia, Mohammad Norouzi, and Ira Kemelmacher-Shlizerman. TryOnDiffusion: A Tale of Two UNets. pages 4606–4615, 2023. [3](#)