

Coarse-to-Fine Latent Diffusion for Pose-Guided Person Image Synthesis

Yanzuo Lu¹ Manlin Zhang¹ Andy J Ma^{1,2,3*} Xiaohua Xie^{1,2,3} Jianhuang Lai^{1,2,3,4}

¹School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou, China

²Guangdong Province Key Laboratory of Information Security Technology, China

³Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

⁴Pazhou Lab (HuangPu), Guangzhou, China

{luyz5, zhangmlin3}@mail2.sysu.edu.cn, {majh8, xiexiaoh6, stsljh}@mail.sysu.edu.cn

Abstract

Diffusion model is a promising approach to image generation and has been employed for Pose-Guided Person Image Synthesis (PGPIS) with competitive performance. While existing methods simply align the person appearance to the target pose, they are prone to overfitting due to the lack of a high-level semantic understanding on the source person image. In this paper, we propose a novel Coarse-to-Fine Latent Diffusion (CFLD) method for PGPIS. In the absence of image-caption pairs and textual prompts, we develop a novel training paradigm purely based on images to control the generation process of a pre-trained text-to-image diffusion model. A perception-refined decoder is designed to progressively refine a set of learnable queries and extract semantic understanding of person images as a coarse-grained prompt. This allows for the decoupling of fine-grained appearance and pose information controls at different stages, and thus circumventing the potential overfitting problem. To generate more realistic texture details, a hybrid-granularity attention module is proposed to encode multi-scale fine-grained appearance features as bias terms to augment the coarse-grained prompt. Both quantitative and qualitative experimental results on the DeepFashion benchmark demonstrate the superiority of our method over the state of the arts for PGPIS. Code is available at <https://github.com/YanzuoLu/CFLD>.

1. Introduction

Pose-Guided Person Image Synthesis (PGPIS) aims to translate the source person image into a specific target pose while preserving the appearance as much as possible. It has a wide range of applications, including film production, virtual reality, and fashion e-commerce. Most

*Corresponding author.



Figure 1. (a) The appearance of person images varies significantly given only a textual prompt for image generation by using Stable Diffusion [35] or ControlNet [55] with OpenPose guidance [3]. (b) Simply aligning the source appearance to the target pose without a semantic understanding of person image can easily lead to overfitting, such that the generated images become distorted and unnatural. (c) Our method learns the coarse-grained prompt for a comprehensive perception of the source image and injects fine-grained appearance features as bias terms, thus generating high-quality images with better generalization performance.

existing methods along this line are developed based on Generative Adversarial Networks (GANs) [6, 18, 20, 25–28, 33, 34, 38, 41, 44, 54, 56, 61, 62]. Nevertheless, the GAN-based approach may suffer from the instability of min-max training objective and difficulty in generating high-quality images in a single forward pass.

As a promising alternative to GANs for image generation, diffusion models synthesize more realistic images progressively from a series of denoising steps. The recently prevailing text-to-image latent diffusion model, such as Stable Diffusion (SD) [35] may now generate compelling person images conditioned on a given textual prompt. The appearance of the generated person can be determined by well-designed prompts [19, 31] or prompt learning [58, 59]. With more reliable structural guidance [29, 55], the synthe-

sized person images can be further constrained to specific poses. Though the text-to-image diffusion generates realistic images from textual prompts with high-level semantics, its training paradigm requires extensive image-caption pairs that are labor-expensive to collect for PGPIS. More importantly, due to the differing information densities between language and vision [11], even the most detailed textual descriptions inevitably introduce ambiguity and may not accurately preserve the appearance as illustrated in Fig. 1(a).

More recently, several diffusion-based approaches have emerged for PGPIS. A texture diffusion module is proposed by PIDM [1] to model the complex correspondence between the appearance of source image and the target pose. Since the denoising process at the high-resolution pixel level is computationally expensive, PoCoLD [9] reduces both the training and inference costs by mapping pixels to low-dimensional latent spaces with a pre-trained Variational Autoencoder (VAE) [7]. In PoCoLD, the correspondence is further exploited by a pose-constrained attention module based on additional 3D Densepose [8] annotations. While both the PIDM and PoCoLD generate more realistic texture details by aligning the source image to the target pose, they lack a **high-level semantic understanding of person images**. Therefore, they are prone to overfitting and poor generalization performance when synthesizing exaggerated poses that are vastly different from the source image or rare in the training set. As demonstrated in Fig. 1(b), the generated images become distorted and unnatural in these cases, which is in line with several GAN-based approaches.

In this work, we propose a novel Coarse-to-Fine Latent Diffusion (CFLD) method for PGPIS. Our approach breaks the conventional training paradigm which leverages textual prompts to control the generation process of a pre-trained SD model. Instead of conditioning on the human-generated signals, i.e. languages that are highly semantic and information-dense, we facilitate a coarse-to-fine appearance control method purely based on images. To obtain the aforementioned semantic understanding specific to person images, we endeavor to decouple the fine-grained appearance and pose information controls at different stages by introducing a *perception-refined decoder*. The perception of the source person image is achieved by randomly initializing a set of learnable queries and progressively refining them in the following decoder blocks via cross-attention. The decoder output serves as a coarse-grained prompt to describe the source image, focusing on the common semantics across different person images, e.g. human body parts and attributes such as age and gender. Moreover, we design a *hybrid-granularity attention* module to effectively encode multi-scale fine-grained appearance features as bias terms to augment the coarse-grained prompt. In this way, the source image is able to be aligned with the target pose by supplementing only the necessary fine-grained details under the

guidance of the coarse-grained prompt, thus achieving better generalization as illustrated in Fig. 1(c).

Our main contributions can be summarized as follows,

- We present a novel training paradigm in the absence of image-caption pairs to overcome the limitations when applying text-to-image diffusion to PGPIS. We propose a perception-refined decoder to extract semantic understanding of person images as a coarse-grained prompt.
- We formulate a new hybrid-granularity attention module to bias the coarse-grained prompt with fine-grained appearance features. Thus, the texture details of generated images are better controlled and become more realistic.
- We conduct extensive experiments on the DeepFashion [21] benchmark and achieve the state-of-the-art performance both quantitatively and qualitatively. User studies and ablations validate the effectiveness of our method.

2. Related Work

Pose-Guided Person Image Synthesis. Ma *et al.* [26] first presents the task of pose-guided person image synthesis and refines the generated images in an adversarial manner. To decouple the pose and appearance information, early approaches [6, 27] propose to learn pose-irrelevant features but fail to handle the complex texture details with vanilla convolutional neural networks. To alleviate this problem, auxiliary information is introduced to improve the generation quality, such as parsing [28] and UV maps [38]. Recent approaches [18, 20, 33, 34, 44, 62] focus on modeling the spatial correspondence between pose and appearance, with the more frequent use of parsing maps [25, 54, 61]. PIDM [1] and PoCoLD [9] are developed based on diffusion models to prevent from the drawbacks in the generative adversarial networks, including the instability of min-max training objective and difficulty in synthesizing high-resolution images. Both of these two diffusion-based methods extend the idea of spatial correspondence to model the relation between the appearance of source image and target pose via the cross-attention mechanism. We argue this leads to overfitting by simply aligning the source appearance to the target pose without a high-level semantic understanding of the person image. More concurrent works like MagicAnimate [50], Animate Anyone [16] and PCDMs [39] require multi-stage and progressive fully fine-tuning, while our pipeline is more efficient and end-to-end by freezing most parameters. And the training paradigm of IP-Adatper [52] is heavily relying on image-text pairs which are not available for PGPIS task.

Controllable Diffusion Models. Diffusion models have recently emerged and demonstrated their potential for high-resolution image synthesis. The core idea is to start with a simple noise vector and gradually transform it into a high-quality image through multiple denoising iterations. Beyond unconditional generation [15, 42, 43], various meth-

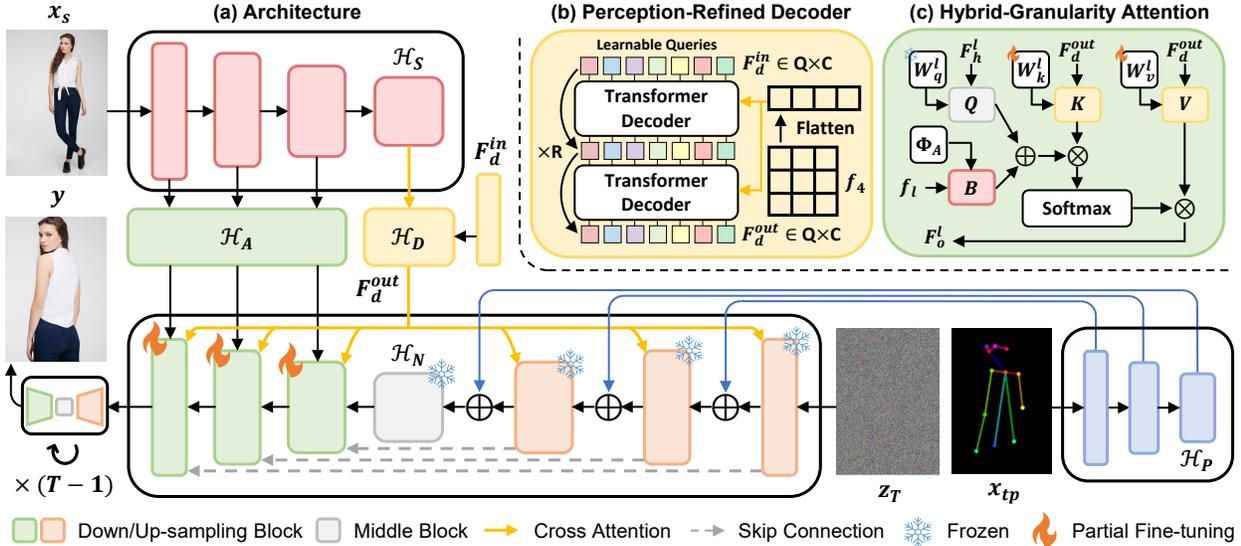


Figure 2. (a) Architecture of our proposed Coarse-to-Fine Latent Diffusion (CFLD) method. For pose-guided latent diffusion, we incorporate a lightweight pose adapter \mathcal{H}_P from [29] to add its output feature maps to the end of each down-sampling block of the pre-trained UNet \mathcal{H}_N for efficient structural guidance. To achieve a coarse-to-fine appearance control, we propose a perception-refined decoder \mathcal{H}_D and hybrid-granularity attention module \mathcal{H}_A , both of which take the multi-scale feature maps from a source image encoder \mathcal{H}_S as inputs. (b) The coarse-grained prompt is obtained by refining the learnable queries progressively in our proposed \mathcal{H}_D . (c) We encode the multi-scale fine-grained appearance features as bias terms in the up-sampling blocks for better texture details within \mathcal{H}_A .

ods have been introduced to incorporate user-supplied control signals into the generation process, enabling more controllable image generation. For instance, [5] introduces the usage of classifier gradients to condition on the generation, while [14] proposes a classifier-free control mechanism employing a weighted summation of conditional and unconditional outputs for controllable synthesis. Moreover, the Latent Diffusion Model (LDM) performs diffusion in the latent space and injects the conditioning signals via a specific encoder and cross-attention. Building upon the pre-trained LDM like Stable Diffusion (SD) [35], subsequent works have explored to bias the latent space by adding extra controls [29, 55], as well as further to provide users with control over the generated content [12, 45]. Rather than employing a high-level conditioning prompt throughout the generation, we design a coarse-to-fine conditioning process that adjusts the latent features at different stages within the UNet-based prediction network, providing better controllable pose-guided person image synthesis.

3. Method

3.1. Preliminary

Our method builds on top of the text-to-image latent diffusion model, i.e., Stable Diffusion (SD) [35] with high-quality image generation ability. There are two main stages in the SD model: a Variational Autoencoder (VAE) [7] that maps between raw-pixel space and low-dimensional latent

space and an UNet-based prediction model [36] for denoising diffusion image generation. It follows the general idea of Denoising Diffusion Probabilistic Model (DDPM) [15], which formulates a forward diffusion process and a backward denoising process of $T = 1000$ steps. The diffusion process progressively adds random Gaussian noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ to the initial latent z_0 , mapping it into noisy latents z_t at different timesteps $t \in [1, T]$,

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad (1)$$

where $\bar{\alpha}_1, \bar{\alpha}_2, \dots, \bar{\alpha}_T$ are derived from a fixed variance schedule. The denoising process learns the UNet $\epsilon_\theta(z_t, t, c)$ to predict the noise and reverse this mapping, where c is the conditional embedding output by e.g. the CLIP [32] text encoder in [35]. The optimization can be formulated as,

$$\mathcal{L}_{mse} = \mathbb{E}_{z_0, c, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t, c)\|_2^2]. \quad (2)$$

3.2. Coarse-to-Fine Latent Diffusion

Architecture and Overview. Fig. 2(a) shows the architecture of our proposed method. For concise illustration, we omit the encoder \mathcal{E} and decoder \mathcal{D} of the VAE [7] model in this figure. In the training phase, we are given sets of the source image x_s , source pose x_{sp} , target pose x_{tp} , and ground-truth image x_g . The source image passes through an image encoder \mathcal{H}_S (e.g. swin transformer [22]), from which we extract a stack of multi-scale feature maps $F_s = [f_1, f_2, f_3, f_4]$ for a coarse-to-fine *appearance*

control. The coarse-grained prompts are learned by our Perception-Refined Decoder (PRD) \mathcal{H}_D and serve as conditional embeddings in both down-sampling and up-sampling blocks of the UNet \mathcal{H}_N . While the down-sampling block in \mathcal{H}_N remains intact in our method, we reformulate the up-sampling block with our Hybrid-Granularity Attention module (HGA) \mathcal{H}_A to bias the coarse-grained prompt with fine-grained appearance features for more realistic textures. More details about \mathcal{H}_D and \mathcal{H}_A will be presented later.

For efficient *pose control*, we adopt a lightweight pose adapter \mathcal{H}_P that consists of several ResNet blocks [10]. The output feature maps of \mathcal{H}_P are added directly to the end of each down-sampling block as in [29]. This requires no additional fine-tuning and explicitly decouples the fine-grained appearance and pose information controls. At different scales of down-sampling, the pose information is only aligned with the same coarse-grained prompts given by our PRD as conditional embeddings, rather than the different multi-scale fine-grained appearance features in the common practice [1, 9]. In this way, the HGA module learns all the pose-irrelevant texture details at the up-sampling stage and is not prone to overfitting. Denote the initial latent state for the ground-truth image as $z_0 = \mathcal{E}(x_g)$. The MSE loss in Eq. (2) is thus rewritten as,

$$\mathcal{L}_{mse} = \mathbb{E}_{z_0, x_s, x_{tp}, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t, x_s, x_{tp})\|_2^2]. \quad (3)$$

Perception-Refined Decoder. Instead of utilizing multi-scale appearance features as conditional embeddings as in the existing diffusion-based approaches [1, 9], we propose to decouple the controls from the fine-grained appearance and pose information at different stages. Thus we design a Perception-Refined Decoder (PRD) to extract semantic understanding of person images as a coarse-grained prompt, given the flattened last-scale output f_4 from \mathcal{H}_S as illustrated in Fig. 2(b). By revisiting how people perceive a person image, we find several common characteristics, i.e., human body parts, age, gender, hairstyle, clothing, and so on, as demonstrated in Fig. 1(a). This inspires us to maintain a set of learnable queries $F_d^{in} \in \mathbb{R}^{Q \times D}$ representing different semantics of person images. They are randomly initialized and progressively refined with the standard transformer decoders [47]. The source image conditioning f_4 interacts via the cross-attention module at each decoder block. After R blocks of refinement, we obtain the coarse-grained prompt F_d^{out} , which serves as the conditional embedding and inputs to both down-sampling and up-sampling in \mathcal{H}_N .

Hybrid-Granularity Attention. To precisely control the texture details of generated images, we introduce the Hybrid-Granularity Attention module (HGA) that is embedded in different scales ($l \in \{1, 2, 3\}$) of up-sampling blocks in \mathcal{H}_N , where we refer F_h^l, F_o^l to its input and output. Given the multi-scale feature maps f_l of the source image from \mathcal{H}_S , the HGA module aims to compensate for the missing

necessary details in the coarse-grained prompts. To achieve this, we formulate the HGA module that naturally follows a coarse-to-fine learning curriculum.

Specifically, we propose to inject multi-scale texture details by biasing the queries of cross-attention in the up-sampling blocks as shown in Fig. 2(c), i.e.,

$$\begin{aligned} Q &= W_q^l F_h^l, & K &= W_k^l F_d^{out}, & V &= W_v^l F_d^{out}, \\ B &= \phi_A(f_l), & F_o^l &= \text{softmax}\left(\frac{(Q+B)K^T}{\sqrt{d}}\right)V, \end{aligned} \quad (4)$$

where W_q^l, W_k^l, W_v^l are specific projection layers for the l -th scale up-sampling block of dimension d . ϕ_A is a fine-grained appearance encoder that mainly consists of K transformer layers with a zero convolution [55] added in the beginning and the end. The zero convolution is a standard 1×1 convolution layer with both weight and bias initialized as zeros. It keeps the gradient of \mathcal{H}_A back to \mathcal{H}_S small enough in the early stage of training, so that the image encoder \mathcal{H}_S and the more easily converged perception-refined decoder \mathcal{H}_D can focus on learning to provide a high-level semantic understanding compatible with the pre-trained SD model. Since we have decoupled the controls of the fine-grained appearance and pose information at different stages, the target pose can be well controlled without overfitting during the down-sampling process. Therefore, such a design encourages the HGA module to slowly fill in more fine-grained textures to better align the generation with the source image during training. Note that W_k^l, W_v^l in the up-sampling blocks are trainable parameters. They are the only trainable parameters of the entire \mathcal{H}_N , which accounts for only **1.2%** of all the parameters in the pre-trained SD model.

3.3. Optimization

To assist the source-to-target pose translation, we follow the insights in [56] to conduct source-to-source self-reconstruction for training. The reconstruction loss is,

$$\mathcal{L}_{rec} = \mathbb{E}_{z_0, x_s, x_{sp}, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t, x_s, x_{sp})\|_2^2], \quad (5)$$

where $z_0 = \mathcal{E}(x_s)$ and z_t is the noisy latent mapped from z_0 at timestep t . The overall objective is written as,

$$\mathcal{L}_{overall} = \mathcal{L}_{mse} + \mathcal{L}_{rec}. \quad (6)$$

Moreover, we adopt the cubic function $t = (1 - (\frac{t}{T})^3) \times T$, $t \in \text{Uniform}(1, T)$ for the distribution of timestep t . It increases the probability of t falling in the early sampling stage and strengthens the guidance, which helps to converge faster and shorten the training time.

Sampling. Once the conditional latent diffusion model is learned, the inference can be performed and starts by sampling a random Gaussian noise $z_T \sim \mathcal{N}(0, I)$. The predicted latent \tilde{z}_0 is obtained by reversing the schedule in

Component	Default	Trainable Params.	
\mathcal{H}_S	Swin-B [22]	87.0M	
\mathcal{H}_A	$K = 4$	22.5M	
\mathcal{H}_D	$R = 8, Q = 16, C = 768$	97.7M	
\mathcal{H}_P	Adapter [29]	30.6M	
\mathcal{H}_N	up-sampling $\mathbf{W}_k^l, \mathbf{W}_v^l$	10.3M	

Method	Pose Info. & Annotation	Training Epochs	Trainable Params.
PIDM [1]	2D OpenPose [3]	300	688.0M
PoCoLD [9]	3D DensePose [8]	100	395.9M
CFLD (Ours)	2D OpenPose [3]	100	248.2M

Table 1. The default settings and the number of trainable parameters in each component of our method and comparison with other diffusion-based methods.

Eq. (1) using the denoising network ϵ_t at each timestep $t \in [1, T]$. We adopt the cumulative classifier-free guidance [2, 9, 14] to strengthen both the source appearance and target pose guidance, i.e.,

$$\begin{aligned} \epsilon_t = & \epsilon_\theta(\mathbf{z}_t, t, \emptyset, \emptyset) \\ & + w_{\text{pose}}(\epsilon_\theta(\mathbf{z}_t, t, \emptyset, \mathbf{x}_{tp}) - \epsilon_\theta(\mathbf{z}_t, t, \emptyset, \emptyset)) \\ & + w_{\text{app}}(\epsilon_\theta(\mathbf{z}_t, t, \mathbf{x}_s, \mathbf{x}_{tp}) - \epsilon_\theta(\mathbf{z}_t, t, \emptyset, \mathbf{x}_{tp})). \end{aligned} \quad (7)$$

When the source image \mathbf{x}_s is missing, we use learnable vectors as the conditional embeddings. The learnable vectors are trained with a probability of $\eta\%$ to drop both \mathbf{x}_s and \mathbf{x}_p during training. The outputs of the pose adapter \mathcal{H}_P will be set to all zeros if the target pose \mathbf{x}_{tp} is missing. We use the DDPM scheduler [15] with 50 steps as the same as in [1, 9]. Finally, the generated image is obtained by the VAE decoder $\mathbf{y} = \mathcal{D}(\tilde{\mathbf{z}}_0)$.

4. Experiments

4.1. Setup

Dataset. We follow [9, 34] to conduct experiments on the In-Shop Clothes Retrieval benchmark of DeepFashion [21] and evaluate on both the 256×176 and 512×352 resolutions. This dataset consists of 52,712 high-resolution person images in the fashion domain. The dataset split is the same as in PATN [62], where 101,966 and 8,570 non-overlapping pairs are selected for training and testing, respectively.

Objective metrics. We use four different metrics to evaluate the generated images quantitatively, including FID [13], LPIPS [57], SSIM [49] and PSNR. Both FID and LPIPS are based on deep features. The Fréchet Inception Distance (FID) calculates the Wasserstein-2 distance [46] between the distributions of generated and real images using Inception-v3 [37] features, and the Learned Perceptual Image Patch Similarity (LPIPS) leverages a network trained on human judgments to measure reconstruction accuracy in the perceptual domain. As for the Structural Similarity Index

Method	Venue	FID↓	LPIPS↓	SSIM↑	PSNR↑
<i>Evaluate on 256×176 resolution</i>					
PATN [62]	CVPR 19'	20.728	0.2533	0.6714	-
ADGAN [28]	CVPR 20'	14.540	0.2255	0.6735	-
GFLA [33]	CVPR 20'	9.827	0.1878	0.7082	-
PISE [54]	CVPR 21'	11.518	0.2244	0.6537	-
SPGNet [†] [25]	CVPR 21'	16.184	0.2256	0.6965	17.222
DPTN [†] [56]	CVPR 22'	17.419	0.2093	0.6975	17.811
NTED [†] [34]	CVPR 22'	8.517	0.1770	0.7156	17.740
CASD [†] [61]	ECCV 22'	13.137	0.1781	0.7224	17.880
PIDM [†] [1]	CVPR 23'	<u>6.812</u>	0.2006	0.6621	15.630
PIDM [‡] [1]	CVPR 23'	6.440	0.1686	0.7109	17.399
PoCoLD [9]	ICCV 23'	8.067	<u>0.1642</u>	<u>0.7310</u>	-
CFLD (Ours)		6.804	0.1519	0.7378	18.235
VAE Reconstructed		7.967	0.0104	0.9660	33.515
Ground Truth		7.847	0.0000	1.0000	$+\infty$
<i>Evaluate on 512×352 resolution</i>					
CoCosNet2 [60]	CVPR 21'	13.325	0.2265	0.7236	-
NTED [†] [34]	CVPR 22'	<u>7.645</u>	0.1999	0.7359	<u>17.385</u>
PoCoLD [9]	ICCV 23'	8.416	<u>0.1920</u>	<u>0.7430</u>	-
CFLD (Ours)		7.149	0.1819	0.7478	17.645
VAE Reconstructed		8.187	0.0217	0.9231	30.214
Ground Truth		8.010	0.0000	1.0000	$+\infty$

Table 2. Quantitative comparisons with the state of the arts in terms of image quality. [†] We strictly follow the evaluation implementation in NTED [34] and reproduce these results. [‡] Results are obtained using the generated images released by the authors.

Measure (SSIM) and Peak Signal to Noise Ratio (PSNR), they quantify the similarity between generated images and ground truths at the pixel level.

Subjective metrics. In addition to the objective metrics, we follow [9] to use the Jab [1, 25, 61] metric in our user study to calculate the percentage of generated images that were considered the best among all methods [1, 25, 34, 56, 61]. Moreover, in order to measure the similarity between the generated images and real data, we quantify the R2G and G2R metrics as many early approaches did [26, 41, 62]. R2G represents the percentage of real images considered as generated and G2R represents the percentage of generated images considered as real by humans.

Implementation details. Our method is implemented with PyTorch [30] and HuggingFace Diffusers [48] on top of the Stable Diffusion [35] with the version of 1.5. The source image is resized to 256×256 and the source image encoder \mathcal{H}_S is a standard Swin-B [22] pretrained on ImageNet [4]. The default settings and the number of trainable parameters in each component are summarized in Tab. 1. We train for 100 epochs using the Adam [17] optimizer with a base learning rate of $5e-7$ scaled by the total batch size. The learning rate undergoes a linear warmup during the first 1,000 steps and is multiplied by 0.1 at 50 epochs. For classifier-free guidance, we set w_{pose} and w_{id} to 2.0 during sampling, and drop the condition \mathbf{x}_s and \mathbf{x}_p with a probability of $\eta = 20(\%)$ during training.

	Source Image	Target Pose	Ground Truth	SPGNet	DPTN	NTED	CASD	PIDM	Ours	Source Image	Target Pose	Ground Truth	SPGNet	DPTN	NTED	CASD	PIDM	Ours
(1)																		
				0.684	0.562	0.492	0.629	0.591	0.495				0.346	0.319	0.302	0.306	0.334	0.191
(2)																		
				0.371	0.326	0.224	0.327	0.311	0.169				0.251	0.251	0.247	0.172	0.290	0.189
(3)																		
				0.242	0.178	0.162	0.170	0.196	0.158				0.374	0.343	0.185	0.199	0.221	0.148
(4)																		
				0.140	0.125	0.106	0.093	0.117	0.087				0.344	0.284	0.242	0.306	0.340	0.245
(5)																		
				0.243	0.226	0.172	0.180	0.179	0.152				0.358	0.279	0.312	0.308	0.314	0.220
(6)																		
				0.293	0.261	0.250	0.246	0.234	0.209				0.345	0.277	0.248	0.262	0.304	0.190
(7)																		
				0.397	0.380	0.315	0.328	0.302	0.258				0.242	0.205	0.242	0.199	0.226	0.191
(8)																		
				0.323	0.178	0.263	0.315	0.178	0.139				0.282	0.249	0.235	0.249	0.269	0.184
(9)																		
				0.306	0.264	0.252	0.254	0.226	0.177				0.317	0.263	0.190	0.266	0.339	0.132
(10)																		
				0.179	0.180	0.088	0.098	0.164	0.077				0.221	0.236	0.187	0.196	0.180	0.157

Figure 3. Qualitative comparisons with state-of-the-arts. To clarify the relation between objective and subjective metrics, we demonstrate the LPIPS measures and label the images with the first and second highest votes from user opinions as red and blue respectively.

4.2. Quantitative Comparison

We quantitatively compare our method with both GAN-based and diffusion-based state-of-the-art approaches in terms of objective metrics. The evaluation is performed on both 256×176 and 512×352 resolutions as the same as in [9, 34]. As shown in Tab. 2, our method significantly outperforms the state-of-the-art across all metrics on both resolutions. In particular, compared to the other two diffusion-based methods [1, 9] in Tab. 1, we achieve better reconstruction with simpler 2D-only pose annotations and fewer trainable parameters. The metrics for VAE [7] reconstructions and the ground truths are also provided for reference.

It is worth noting that the results we obtain by running with the checkpoint provided by PIDM [1] suffer from severe overfitting, resulting in a large gap between the quantitative results of provided images and those from the checkpoint.

4.3. Qualitative Comparison

In Fig. 3, we present a comprehensive visual comparison with recent approaches that are publicly available and reproducible, including SPGNet [25], DPTN [56], NTED [34], CASD [61] and PIDM [1]. Our observations can be summarized as follows. (1) Both GAN-based and diffusion-based methods suffer from overfitting the human poses. When

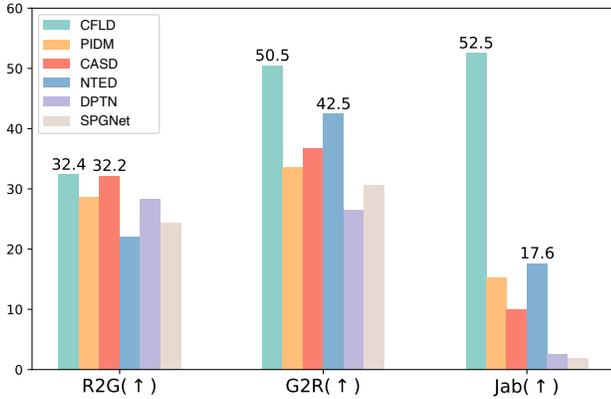


Figure 4. User study results in terms of R2G, G2R and Jab metrics.

generating some target poses that are extreme or not common in the training set, existing methods show severe distortions as demonstrated in rows 1–2. Since we decouple the controls of fine-grained appearance and pose information, our method circumvents the potential overfitting problem and always generates a reasonable pose with the conditioning coarse-grained prompt and fine-grained appearance bias. (2) For source images in rows 3–6 with more complex clothing, our generated images better preserve the textures details while aligning with the target pose thanks to the robust coarse-to-fine learning curriculum of hybrid-granularity attention module. For other methods, although they match in color, the clothes either exhibit blurring and distortion (SPGNet, DPTN, and CASD) or are spliced unnaturally in texture, creating a large gap from the source image (NTED and PIDM). (3) As for cases where the target pose requires visualization of areas invisible in the source image, our method exhibits strong understanding and generalization capabilities. With a semantic understanding of the source image provided by the perception-refined decoder, our method is aware of what should be predicted when the person turns around or sits down as illustrated in rows 7–10, such as different patterns on the front and back of clothes, the sitting chair, and lower body wear.

4.4. User Study

To verify the gap between generated and real images as well as our superiority over the state of the arts, we have recruited over 100 volunteers to perform the following two user studies following PIDM [1]. (1) For the R2G and G2R metrics, volunteers were asked to discriminate between 30 generated images and 30 real images from the testing set. Each volunteer could only see the generated images of a specific method, and the pairs of source image and target pose for generation were consistent across methods for a fair comparison. From the results in Fig. 4, chances of a real image being recognized as generated (R2G) are relatively low, and over half of the images we generated are recognized as real (G2R), demonstrating that our method gener-

Method	Biasing	Trainable	Prompt	LPIPS↓	SSIM↑
B1		K, V	M-S	0.2018	0.6959
B2		K, V	CLIP	0.2099	0.6944
B3		K, V	PRD	0.1615	0.7293
B4		Q, K, V	PRD	0.1742	0.7198
B5	Q	K, V	Swin	0.1912	0.7038
Ours	Q	K, V	PRD	0.1519	0.7378

Table 3. Quantitative results for ablation studies. M-S is short for multi-scale fine-grained appearance features similar to [1, 9].



Figure 5. Qualitative ablation results. Our approach has a high-level understanding of the source image rather than forced alignment. It is also less prone to overfitting through the complementary coarse-grained prompts and fine-grained appearance biasing.

ates more realistic images that are less likely to be judged as fake by humans. (2) For the Jab metric, each volunteer was asked to choose the best match to the ground truth from the generated images of different methods. Compared to other methods, our Jab score achieved 52.5 percent, significantly higher (+34.9) than the counterpart in second place, indicating that our method is more preferred and generates better texture details and pose alignment.

4.5. Ablation Study

We perform ablation studies at multiple baselines to compare with our method. The quantitative results are presented in Tab. 3. B1 is referenced from the other two diffusion-based approaches [1, 9] that incorporate multi-scale fine-grained appearance features as conditional prompts. we also experiment with CLIP image encoder [32] in B2 to produce descriptive coarse-grained prompts for source images, which is first explored by an image-editing approach [51] that are also conditioned purely on images. Together with the qualitative results in Fig. 5, we can see that even very simple textures sometimes fail to be preserved, suggesting that these prompts are not compatible with the pre-trained SD model. To provide coarse-grained features that are more specific to person images, we integrate the Perception-Refined Decoder (PRD) into B3. The reconstruction metrics (i.e., LPIPS and SSIM) in Tab. 3 re-



Figure 6. (a) Style transfer results of our method. The appearance in the reference image can be edited while maintaining the pose and appearance. This is achieved by masking out regions of interest in the reference image and requires no additional training. (b) The interpolation results show that texture details can be gradually shifted from one style to another in a smooth manner (from Style 1 to 2).



Figure 7. Visualizing attention maps by different queries of the prompt decoder. The maps are averaged over all attention heads.

veal a significant improvement in the quality of generated images, which validates the effectiveness of our proposed PRD. While this can be confirmed qualitatively in Fig. 5, there is still a lack of textural details as indicated by the red box. To address this issue, we experiment with training more parameters in the UNet as B4 and instead observe a decrease in performance. This implies that the generalization ability of SD model is compromised, which is not our expectation. Thus we come up with the Hybrid-Granularity Attention (HGA) to bias the queries and achieve state-of-the-art results both quantitatively and qualitatively. In order to verify whether the source image encoder (i.e., Swin Transformer [22]) is able to learn sufficient information for HGA and give a useful prompt, we abandon the PRD in B5. The qualitative results in Fig. 5 demonstrate that both B4 and B5 are overfitting, only our method circumvents this problem by learning in a coarse to fine-grained manner.

Visualization. In Fig. 7, we visualize the effectiveness of different queries in \mathcal{H}_D . The attention maps reflect different human body parts of person images captured by learnable queries, which proves that we have a high-level understanding of the source images and thus less prone to overfitting.

4.6. Appearance Editing

Style Transfer. Our CFLD inherits the strong generation ability of SD model by freezing the vast majority of its pa-

rameters. Thus the style transfer can be achieved simply by masking without additional training. Specifically, we mark the regions of interest in the reference image \mathbf{y}^{ref} as a binary mask m . During sampling, the noise prediction is decomposed into $\epsilon'_t = m \cdot \epsilon_t + (1 - m) \cdot z_t^{ref}$, where the ϵ_t is based on the pose from \mathbf{y}^{ref} and the appearance from different styles of source images. Let z_t^{ref} be the noisy latent at timestep t mapped from $z_0^{ref} = \mathcal{E}(\mathbf{y}^{ref})$ as in Eq. (1). From the results in Fig. 6(a), our method generates realistic and coherent texture details in the regions of interest.

Style Interpolation. Additionally, our CFLD supports arbitrary linear interpolation of both coarse-grained prompts and fine-grained appearance biases. As shown in Fig. 6(b), our generated images are faithfully reproducing different styles with smooth transitions.

5. Conclusion

This paper presents a novel Coarse-to-Fine Latent Diffusion (CFLD) method for Pose-Guided Person Image Synthesis (PGPIS). We circumvent the potential overfitting problem by decoupling the fine-grained appearance and pose information controls. Our proposed Perception-Refined Decoder (PRD) and Hybrid-Granularity Attention module (HGA) enable a high-level semantic understanding of person images, while also preserving texture details through a coarse-to-fine learning curriculum. Extensive experiments demonstrate that CFLD outperforms the state of the arts in PGPIS both quantitatively and qualitatively. Our future work will investigate whether the CFLD can be extended to more downstream tasks that suffer from inferior data like person re-identification [23, 53] and domain adaptation [24, 40], since our training paradigm yields both a pre-trained feature network and a powerful generator for augmentation.

Acknowledgments. This work was supported in part by the National Natural Science Foundation of China (U22A2095, 62276281).

References

- [1] Ankan Kumar Bhunia, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Jorma Laaksonen, Mubarak Shah, and Fahad Shahbaz Khan. Person image synthesis via denoising diffusion model. In *CVPR*, page 5968–5976, 2023. 2, 4, 5, 6, 7
- [2] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, pages 18392–18402, 2023. 5
- [3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, page 7291–7299, 2017. 1, 5
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, page 248–255, 2009. 5
- [5] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, pages 8780–8794, 2021. 3
- [6] Patrick Esser and Ekaterina Sutter. A variational u-net for conditional appearance and shape generation. In *CVPR*, page 8857–8866, 2018. 1, 2
- [7] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, page 12873–12883, 2021. 2, 3, 6
- [8] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *CVPR*, pages 7297–7306, 2018. 2, 5
- [9] Xiao Han, Xiatian Zhu, Jiankang Deng, Yi-Zhe Song, and Tao Xiang. Controllable person image synthesis with pose-constrained latent diffusion. In *ICCV*, page 22768–22777, 2023. 2, 4, 5, 6, 7
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, page 770–778, 2016. 4
- [11] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, page 16000–16009, 2022. 2
- [12] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv:2208.01626*, 2022. 3
- [13] 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. In *NeurIPS*, 2017. 5
- [14] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS Workshops*, 2021. 3, 5
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2, 3, 5
- [16] Liucheng Hu, Xin Gao, Peng Zhang, Ke Sun, Bang Zhang, and Liefeng Bo. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. *arXiv:2311.17117*, 2023. 2
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [18] Yining Li, Chen Huang, and Chen Change Loy. Dense intrinsic appearance flow for human pose transfer. In *CVPR*, page 3693–3702, 2019. 1, 2
- [19] Vivian Liu and Lydia B Chilton. Design guidelines for prompt engineering text-to-image generative models. In *CHI*, pages 1–23, 2022. 1
- [20] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *ICCV*, pages 5904–5913, 2019. 1, 2
- [21] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, page 1096–1104, 2016. 2, 5
- [22] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, page 10012–10022, 2021. 3, 5, 8
- [23] Yanzuo Lu, Manlin Zhang, Yiqi Lin, Andy J Ma, Xiaohua Xie, and Jianhuang Lai. Improving pre-trained masked autoencoder via locality enhancement for person re-identification. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 509–521. Springer, 2022. 8
- [24] Yanzuo Lu, Meng Shen, Andy J Ma, Xiaohua Xie, and Jianhuang Lai. Mlnet: Mutual learning network with neighborhood invariance for universal domain adaptation. In *AAAI*, 2024. 8
- [25] Zhengyao Lv, Xiaoming Li, Xin Li, Fu Li, Tianwei Lin, Dongliang He, and Wangmeng Zuo. Learning semantic person image generation by region-adaptive normalization. In *CVPR*, page 10806–10815, 2021. 1, 2, 5, 6
- [26] Liqian Ma, Xu Jia, Qianru Sun, B. Schiele, T. Tuytelaars, and L. Gool. Pose guided person image generation. In *NeurIPS*, 2017. 2, 5
- [27] Liqian Ma, Qianru Sun, Stamatios Georgoulis, Luc Van Gool, Bernt Schiele, and Mario Fritz. Disentangled person image generation. In *CVPR*, pages 99–108, 2018. 2
- [28] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *CVPR*, page 5084–5093, 2020. 1, 2, 5
- [29] Chong Mou, Xintao Wang, Liangbin Xie, Jing Zhang, Zhong-gang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv:2302.08453*, 2023. 1, 3, 4, 5
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 32, 2019. 5
- [31] Nikita Pavlichenko and Dmitry Ustalov. Best prompts for text-to-image models and how to find them. In *SIGIR*, pages 2067–2071, 2023. 1
- [32] 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 *ICML*, page 8748–8763, 2021. 3, 7
- [33] Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H. Li, and Ge Li. Deep image spatial transformation for person image generation. In *CVPR*, page 7690–7699, 2020. 1, 2, 5
- [34] Yurui Ren, Xiaoqing Fan, Ge Li, Shan Liu, and Thomas H. Li. Neural texture extraction and distribution for controllable person image synthesis. In *CVPR*, page 13535–13544, 2022. 1, 2, 5, 6
- [35] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, page 10684–10695, 2022. 1, 3, 5
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, page 234–241, 2015. 3
- [37] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *NeurIPS*, 29, 2016. 5
- [38] Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. Style and pose control for image synthesis of humans from a single monocular view. *arXiv:2102.11263*, 2021. 1, 2
- [39] Fei Shen, Hu Ye, Jun Zhang, Cong Wang, Xiao Han, and Wei Yang. Advancing pose-guided image synthesis with progressive conditional diffusion models. In *ICLR*, 2024. 2
- [40] Meng Shen, Yanzuo Lu, Yanxu Hu, and Andy J Ma. Collaborative learning of diverse experts for source-free universal domain adaptation. In *ACM MM*, pages 2054–2065, 2023. 8
- [41] Aliaksandr Siarohin, Enver Sangineto, Stephane Lathuiliere, and Nicu Sebe. Deformable gans for pose-based human image generation. In *CVPR*, page 3408–3416, 2018. 1, 5
- [42] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *ICLR*, 2021. 2
- [43] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 2
- [44] Hao Tang, Song Bai, Li Zhang, Philip H. S. Torr, and Nicu Sebe. Xinggan for person image generation. In *ECCV*, page 717–734, 2020. 1, 2
- [45] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *CVPR*, pages 1921–1930, 2023. 3
- [46] Leonid Nisonovich Vaserstein. Markov processes over denumerable products of spaces, describing large systems of automata. *Problemy Peredachi Informatsii*, 5(3):64–72, 1969. 5
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
- [48] Patrick von Platen, Suraj Patil, Anton Lozhkov, Pedro Cuenca, Nathan Lambert, Kashif Rasul, Mishig Davaadorj, and Thomas Wolf. Diffusers: State-of-the-art diffusion models. <https://github.com/huggingface/diffusers>, 2022. 5
- [49] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *TIP*, 13(4):600–612, 2004. 5
- [50] Zhongcong Xu, Jianfeng Zhang, Jianfeng Liew, Hanshu Yan, Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng Shou. Magicanimate: Temporally consistent human image animation using diffusion model. *arXiv:2311.16498*, 2023. 2
- [51] Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. Paint by example: Exemplar-based image editing with diffusion models. In *CVPR*, page 18381–18391, 2023. 7
- [52] Hu Ye, Jun Zhang, Siyi Liu, Xiao Han, and Wei Yang. Ip-adapt: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv:2308.06721*, 2023. 2
- [53] Junkun Yuan, Xinyu Zhang, Hao Zhou, Jian Wang, Zhongwei Qiu, Zhiyin Shao, Shaofeng Zhang, Sifan Long, Kun Kuang, Kun Yao, et al. Hap: Structure-aware masked image modeling for human-centric perception. *NeurIPS*, 36, 2024. 8
- [54] Jinsong Zhang, Kun Li, Yu-Kun Lai, and Jingyu Yang. Pise: Person image synthesis and editing with decoupled gan. In *CVPR*, page 7982–7990, 2021. 1, 2, 5
- [55] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, page 3836–3847, 2023. 1, 3, 4
- [56] Pengze Zhang, Lingxiao Yang, Jianhuang Lai, and Xiaohua Xie. Exploring dual-task correlation for pose guided person image generation. In *CVPR*, page 7713–7722, 2022. 1, 4, 5, 6
- [57] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, page 586–595, 2018. 5
- [58] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *CVPR*, page 16816–16825, 2022. 1
- [59] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *IJCV*, 130(9):2337–2348, 2022. 1
- [60] Xingran Zhou, Bo Zhang, Ting Zhang, Pan Zhang, Jianmin Bao, Dong Chen, Zhongfei Zhang, and Fang Wen. Cocosnet v2: Full-resolution correspondence learning for image translation. In *CVPR*, page 11465–11475, 2021. 5
- [61] Xinyue Zhou, M. Yin, Xinyuan Chen, Li Sun, Changxin Gao, and Qingli Li. Cross attention based style distribution for controllable person image synthesis. In *ECCV*, page 161–178, 2022. 1, 2, 5, 6
- [62] Zhen Zhu, Tengting Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *CVPR*, page 2347–2356, 2019. 1, 2, 5