

DiffBody: Diffusion-based Pose and Shape Editing of Human Images

Yuta Okuyama, Yuki Endo, and Yoshihiro Kanamori
 University of Tsukuba

okuyama.yuta.sw@alumni.tsukuba.ac.jp, {endo, kanamori}@cs.tsukuba.ac.jp



Figure 1. Results of editing the pose and body shape of a fullbody human image using our method. For various pose and body shape inputs, our method can generate realistic human images, while preserving the clothing textures and facial identity of reference person images.

Abstract

Pose and body shape editing in a human image has received increasing attention. However, current methods often struggle with dataset biases and deteriorate realism and the person’s identity when users make large edits. We propose a one-shot approach that enables large edits with identity preservation. To enable large edits, we fit a 3D body model, project the input image onto the 3D model, and change the body’s pose and shape. Because this initial textured body model has artifacts due to occlusion and the inaccurate body shape, the rendered image undergoes a diffusion-based refinement, in which strong noise destroys body structure and identity whereas insufficient noise does not help. We thus propose an iterative refinement with weak noise, applied first for the whole body and then for the face. We further enhance the realism by fine-tuning text embeddings via self-supervised learning. Our quantitative and qualitative evaluations demonstrate that our method outperforms other existing methods across various datasets.

1. Introduction

Editing the pose and body shape of a human image is a task to change the orientation, position, and slenderness/fatness of the subject’s limbs and torso, and outputs a realistic image of the same person. This task has been actively studied due to its potential applications, such as visual simulation after dieting and efficient fashion photo shooting.

Recent techniques for editing pose and shape in a human image can be broadly categorized into two approaches, i.e., 1) the approach based on image warping and generative adversarial networks (GANs) and 2) the diffusion-based approach [4]. The former approach utilizes image warping and well preserves the person’s identity, but it often causes artifacts with large edits (see Figures 9 and 10). The latter approach yields high-quality output images with diverse poses and shapes thanks to the diffusion models. However, it often struggles with identity preservation; the output human figures tend to have different clothes or faces from those of the input image (see Figure 2).




Figure 2. Limitations of the state-of-the-art methods for human pose editing by Ren et al. [31] and Bhunia et al. [4]. These methods often struggle with identity preservation for input images not observed in a test set corresponding to a training set (see the red boxes).

In this paper, we propose a method to enable large edits for various poses and shapes in a single human image with identity preservation. Our large edit for poses and shapes is driven by a 3D parametric body model, followed by refinement using a diffusion model. Namely, we fit a 3D parametric body model [26] to the reference human image, project the reference image onto the 3D model, and change the pose and shape parameters to obtain a textured human model with a new pose and shape. This initial textured model often has artifacts (see Figure 3) because it exhibits occluded regions that were not visible in the reference image, and the target body shape of the texture projection is often inaccurate. We thus refine the rendered human image using an image-to-image translation technique with a diffusion model [22]. However, a critical pitfall lies in the difficulty of the noise strength control in the diffusion model; insufficient noise strength does not refine the visual artifacts, whereas excessively strong noise will significantly alter the person’s identity (see Figure 4).

To improve the visual artifacts while avoiding significant structural changes, we propose an iterative refinement with weak noise. Furthermore, to enhance the quality of refinements, we update text embeddings used for network conditioning during iterative refinements through self-supervised learning. Building upon these refinement processes, we adopt a staged pipeline in which we first refine the full-body image and then locally refine facial features to enhance realism. We quantitatively and qualitatively evaluate our method, demonstrating superior results compared to existing methods across various datasets.




Figure 3. Problem of person image projection onto 3D parametric body models. As shown in the red boxes, the initial textured model often has artifacts when its pose and shape are edited.




Figure 4. Refinement results of a textured model using a diffusion model with different noise strengths. Insufficient noise strength does not refine the visual artifacts, whereas excessively strong noise will significantly alter the person’s identity.

2. Related work

Image-to-image translation. To synthesize human images according to specified poses, an early attempt based on image-to-image translation adopted a generative adversarial network (GAN) conditioned on keypoints [21]. Another attempt used a U-Net architecture combined with a variational autoencoder (VAE) [8]. However, these approaches cannot handle large deformation because they struggle to handle feature misalignment between input and output images with different poses. Some methods enable more flexible pose editing by using flow fields [11, 15, 20, 32], 2D UV maps of human bodies (i.e., DensePose) [1, 25], or human parsing maps [19, 43, 48]. In addition, the attention mechanism [38] further improved the quality of generated images by efficiently and extensively capturing image features [14, 31, 36, 45, 49]. Recently, PIDM [4] used a diffusion

model for pose editing and achieved state-of-the-art results thanks to its capability of image generation. These existing methods yield satisfactory results for a test set corresponding to a training set. However, for diverse human images outside the dataset, they often cause discrepancies between a generated person and a reference person in facial identity and clothing textures, resulting in limiting their applicability to in-the-wild data. In contrast, our method performs well for diverse inputs by refining textured 3D human models using general text-to-image diffusion models finetuned for a reference person image.

Image warping. Several pose editing methods perform image warping based on 3D human models to improve generalizability for diverse human images [16, 35, 41]. These methods first fit a 3D human model [26] to a reference person image. Then, the reference person image is projected onto the fitted 3D human model, whose pose is edited to obtain a final image with the desired pose. These methods inpaint invisible regions in reference images by using flow fields [16], GANs [41], or diffusion models [35]. However, these methods still struggle to handle invisible regions and sometimes yield highly distorted textures. Instead of directly inpainting invisible regions, our method extends an image-to-image translation technique using a diffusion model [22] to refine projected textures. Our iterative refinement approach enables us to generate more plausible textures in invisible regions.

Text-to-image translation. There have also been approaches utilizing large-scale language models and diffusion models for synthesizing person images from text and pose information [24, 44]. Combining these methods with DreamBooth [34], which finetunes diffusion models with several reference images, enables pose editing of a specific person in a reference image. The benefits of large-scale language models and diffusion models allow us to generate realistic human images with diverse poses. However, preserving the facial identity and clothing textures of the reference images remains challenging in this approach.

Body shape editing. Zhou et al. [47] proposed the first method for body shape editing using 3D human models. In this method, the user interactively fits a 3D parametric body model, SCAPE [2], to an input fullbody image. Then, after editing the body shape of the 3D model by specifying height and weight parameters, the user can obtain output images deformed via image warping based on the 3D model. MovieReshape [47] automated the fitting of SCAPE models by taking a video as input. FBBR [30] is a method that directly applies image warping based on 2D flow field to an input human image. The user can specify a positive or negative value indicating the amount of deformation, and

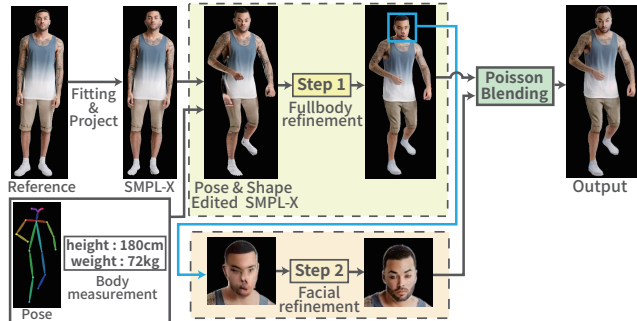


Figure 5. Overview of our method. Our method first computes a textured SMPL-X model from a reference person image. The SMPL-X model is then deformed according to given body pose, height, and weight parameters. To compensate for occluded and distorted textures resulting from texture projection, our method performs step-by-step refinement for the rendered image of the SMPL-X model using diffusion models.

the CNN estimates a flow field for the input reference person image. However, it yields significant distortion when extensively editing body shapes. In our method, we perform image deformation based on 3D human models, while locally refining the invisible and distorted regions of the deformed images using diffusion models. This allows large deformation with less distortion.

3. Method

Figure 5 shows an overview of our method. We first reconstruct a 3D human model from a reference image and then improve its appearance via diffusion-based refinement. The reconstruction procedure of 3D human models is as follows. We first fit SMPL-X [26], which is a 3D parametric body model, to a reference image. Next, we obtain a texture of the SMPL-X model by projecting the reference image. We then specify body shape parameters (height and weight) and keypoints (joint positions) to manipulate the pose and body shape of the SMPL-X model. Finally, we obtain a coarse human image by rendering the textured SMPL-X model.

Although the rendered image of a textured SMPL-X can reflect a specified pose and body shape (see Figure 3), it contains visual artifacts due to invisible areas and texture distortion. Therefore, we aim to reduce such artifacts by leveraging diffusion models, which are equipped with highly-expressive generative capability and local editability. We use the pre-trained latent diffusion model (LDM) [33] as a backbone of our method. To faithfully preserve a person’s identity, we use DreamBooth [34]. Specifically, we fine-tune the LDM using a body image and a face image cropped from a reference image. As a textual input of DreamBooth, we use a prompt containing a special token “sks” associ-




Figure 6. Overview of our refinement module. This module is based on an image-to-image translation technique using a diffusion model [22]. It first adds noise to the latent map computed from an input image using the VAE. Then, the diffusion model denoises the noised maps using a refinement mask and conditioning vector extracted from a prompt and keypoints. The final output image is obtained from the denoised latent map using the VAE.

ated with the reference person (see Figure 5). We also condition the LDM on keypoints of joint positions using T2I-Adapter [24] to reflect pose information effectively.

Using the finetuned diffusion model conditioned on a text prompt and keypoints, we refine a textured SMPL-X image. To do so, inspired by SDEdit [22], we perform image-to-image translation that can modify fine details of an input image while preserving its coarse structure via forward (noising) and reverse (denoising) processes. However, simply using SDEdit unnecessarily changes textures in visible areas while not sufficiently modifying visual artifacts in invisible and distorted regions. This is because textures in each area should be modified with different noise levels, but SDEdit uniformly adds noise to an entire image and performs denoising. To address this issue, we propose two-stage refinement steps, which modify each part with different noise levels. The first step is for refining fullbody textures in invisible areas, whereas the second step is for facial distorted textures. We describe these steps in the following sections.

3.1. Step 1: Fullbody refinement

In Step 1, we introduce a refinement module (Figure 6) for correcting unnatural textures in invisible areas caused by the projection of reference images. The refinement module takes as input an image to be modified, a refinement mask for invisible areas, a conditioning vector extracted from a prompt and keypoints. The input image is converted into a latent feature map using the VAE encoder of the LDM. Inspired by Blended Diffusion [3], we perform denoising in the refinement mask to modify only invisible areas. Specifically, we first add noise to the latent feature map through the forward process and store the noise map at each time step. Then, during the reverse process, we perform denoising while replacing the values outside the refinement mask with the noise map stored at the corresponding time step.

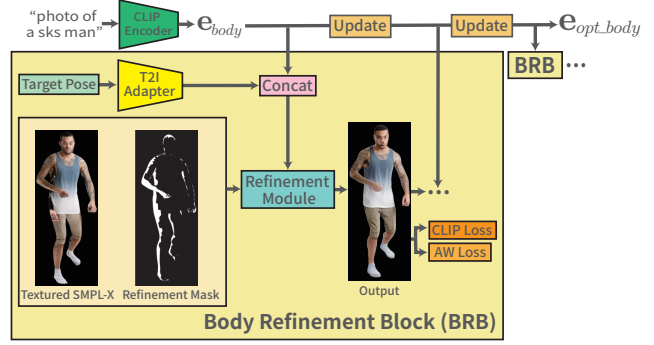


Figure 7. Overview of fullbody refinement in Step 1. The body refinement block (BRB) iteratively refines the invisible area of a textured SMPL-X model using our refinement module. At each refinement iteration, it also updates the text embedding e_{body} using the CLIP and AW losses. The process of BRB is also repeated while reinitializing the refined image to make the update of the text embedding stable.

However, in this process, excessively strong noise unnecessarily modifies human coarse structure, whereas weak noise cannot improve fine texture details sufficiently. In addition, the appropriate noise strength depends on input images, but adjusting individual noise strengths is cumbersome.

Iterative refinement. To adequately modify fine texture details while preserving coarse structure, our method iteratively refines images via multiple reverse processes with weak noise. As illustrated in Figure 7, we iterate the process that adds weak noise to an input image and performs denoising using the refinement module. During the iteration, we need to determine how many iterations are the best to obtain a high-quality output. This optimal iteration number is determined based on loss functions for evaluating the output image. We first use the Adaptive Wing (AW) loss [39] between joint heatmaps estimated by OpenPose [5] for the output image and the rendered SMPL-X image. In addition, we use CLIP similarity [29] between the output and reference images for each part [10] based on SMPL-X labeling. The final output is obtained when the sum of these loss functions become minimum during fixed iterations.

Text embedding optimization. In parallel with iterative refinement, we also aim to further improve the visual quality of the final output by optimizing the text embedding e_{body} conditioned on the network. At each iteration, we update the text embedding e_{body} via backpropagation of the above loss functions.

Input reinitialization. However, we often observed cases where the refined image becomes unnatural during itera-

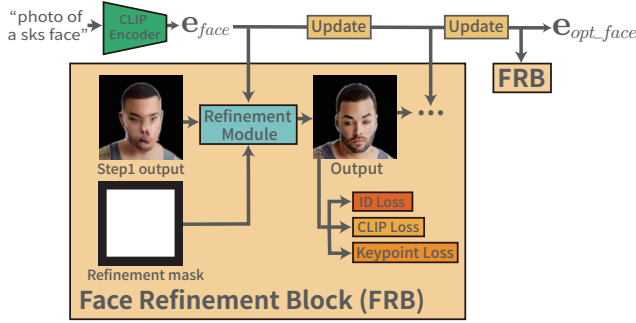


Figure 8. Overview of facial refinement in Step 2. The face refinement block (FRB) refines distorted facial textures in the output of Step 1. Similar to Step 1, we perform iterative refinement using our refinement module. The text embedding e_{face} is also updated according to the ID, CLIP, and Keypoint losses.

tive refinement before the text embedding optimization converges. As a result, the text embedding may also converge to an undesirable solution. To address this issue, we reinitialize the input image with the initial image every several iterations in iterative refinement. This approach makes text embedding optimization more stable and allows us to obtain higher-quality results, as shown in the ablation study in Section 4.3.

While this fullbody refinement process can improve the invisible areas of rendered SMPL-X images, unnatural distortions are still observed in face textures, which are further refined in Step 2.

3.2. Step 2: Facial refinement

Figure 8 shows an overview of Step 2. According to region labels obtained by SMPL-X [26], we automatically crop a facial region from the image refined in Step 1. After resizing the cropped image to 512×512 , our refinement module iteratively updates the image and text embedding e_{face} in a similar procedure to Step 1. To seamlessly merge face and body images after refinement, we keep the outer region of the face image unchanged by assigning 0 to the outer region of the refinement mask and 1 to the inner region. To optimize the text embedding e_{face} , we use the identity loss by MagFace [23], keypoint loss defined by MSE loss between keypoints of output and rendered SMPL-X images estimated using RetinaFace [7], and CLIP similarity [29]. We obtain a final output in Step 2 when the sum of these loss functions shows a minimum value during fixed iterations. The final result is generated by merging the output images in Steps 1 and 2 via Poisson blending [27].

4. Experiments

Datasets. We used various datasets such as DeepFashion [17], MonoPerfCap [40], Everybody Dance Now

(EDN) [6], YouTube 18 Dancers [13], and EHF [26]. From DeepFashion, we extracted 80 images containing fullbody humans and used 26 images as reference person images and 54 images as target poses. For the other datasets containing videos of different people, we used a single frame of each video as a reference and about 35 frames of each video as targets. In total, we used 51 reference images and 963 target images.

Evaluation metrics. We evaluated our method using the metrics including SSIM, PSNR, LPIPS [46], FID [9], and the AW and ID losses. For the AW loss, we employed the L2 loss between heatmaps estimated from generated and target images using OpenPose [5]. For the ID loss, we used the cosine similarity between facial features obtained using MagFace [23].

Implementation details. We implemented our method using Python and PyTorch and ran our program on NVIDIA RTX A6000. The image size used in the experiments was 512×512 . In our method, we finetuned a pre-trained Stable Diffusion (v1.4) [33] with a single reference image for each person using DreamBooth [34]. During the finetuning process, we used the AdamW optimizer with a learning rate of 1.0×10^{-6} . For inference, we used 30% noise for image refinement in Steps 1 and 2 and performed 100 iterations of refinement. For input reinitialization, which aims to make text embedding optimization stable, we reinitialize the input image every 5 iterations in iterative refinement. For text embedding optimization, we used the Adam optimizer [12] and the cosine annealing scheduler [18] with warmup, with minimum and maximum learning rates set to 4.0×10^{-4} and 5.0×10^{-4} , respectively. We used PyMaF [42] to estimate SMPL-X parameters. Our method took about 20 minutes for each finetuning and inference. The resulting images shown in our paper are trimmed to save space. See our supplementary material for more implementation details and additional results.

4.1. Evaluation of pose editing

Quantitative comparison. Regarding pose editing, we compared our method with the existing methods including LWG [16], PGHA [41], NTED [31], PIDM [4], T2IA [24], and DINAR [35]. For T2IA, we finetuned the pre-trained diffusion model equipped with the T2I adapter using DreamBooth for each reference person image. For LWG, PGHA, and DINAR, we used official models trained on the iPER [16], 3D People [28], or Texel [37] dataset. For the remaining methods, we used official models trained on the DeepFashion dataset [17]. Table 1 shows the quantitative comparison on various datasets. The bold font indicates the best score for each metric on each dataset. As shown in the results, our method outperforms the existing methods in almost all metrics on the four datasets (i.e., Ev-



Figure 9. Qualitative comparison of editing the reference images to the target poses with our method and the existing methods (LWG [16], PGHA [41], DINAR [35], NTED [31], PIDM [4], and T2IA [24]) on each dataset.

Everybody Dance Now, EHF, MonoPerfCap, and YouTube 18 Dancers). In the results on DeepFashion, although NTED

performs the best, our method also outperforms the methods not trained on the DeepFashion dataset. Furthermore, our

Table 1. Quantitative comparison on the Everybody Dance Now (EDN) [6], EHF [26], MonoPerfCap [40], YouTube 18 Dancers (Y18D) [13], DeepFashion [17] datasets, and average scores across all datasets.

		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	ID \downarrow	AW \downarrow
EDN [6]	LWG [16]	17.000	0.751	0.251	48.322	0.352	7.70
	PGHA [41]	16.219	0.745	0.304	90.952	0.408	6.22
	DINAR [35]	15.398	0.725	0.276	52.515	0.308	6.78
	NTED [31]	16.153	0.742	0.292	90.843	0.480	5.92
	PIDM [4]	14.928	0.680	0.320	108.288	0.487	8.76
	T2IA [24]	5.027	0.400	0.627	229.689	0.454	3.90
	Ours	18.574	0.793	0.212	44.801	0.276	3.30
EHF [26]	LWG [16]	18.994	0.788	0.232	68.922	0.210	3.20
	PGHA [41]	17.763	0.764	0.289	124.662	0.389	2.75
	DINAR [35]	17.332	0.744	0.286	87.227	0.333	2.89
	NTED [31]	17.564	0.751	0.323	285.534	0.473	2.31
	PIDM [4]	15.125	0.681	0.379	279.378	0.490	3.28
	T2IA [24]	6.062	0.546	0.567	238.12	0.382	2.93
	Ours	20.356	0.815	0.216	58.887	0.180	1.62
MonoPerfCap [40]	LWG [16]	19.133	0.696	0.269	41.500	0.268	4.17
	PGHA [41]	17.849	0.670	0.339	88.296	0.400	4.13
	DINAR [35]	17.104	0.651	0.313	53.149	0.278	6.68
	NTED [31]	18.064	0.662	0.318	80.670	0.476	2.23
	PIDM [4]	16.547	0.649	0.350	92.661	0.476	4.78
	T2IA [24]	5.133	0.382	0.621	219.755	0.424	3.10
	Ours	19.720	0.716	0.242	46.432	0.175	1.82
Y18D [13]	LWG [16]	17.000	0.701	0.285	52.927	0.328	4.30
	PGHA [41]	16.583	0.689	0.328	107.667	0.393	4.65
	DINAR [35]	15.796	0.672	0.311	65.441	0.326	6.52
	NTED [31]	16.651	0.683	0.324	91.852	0.485	2.70
	PIDM [4]	15.676	0.659	0.341	102.664	0.459	4.57
	T2IA [24]	5.057	0.391	0.644	213.486	0.422	4.07
	Ours	18.048	0.731	0.254	52.170	0.283	2.54
DeepFashion [17]	LWG [16]	15.652	0.633	0.405	123.687	0.403	14.22
	PGHA [41]	14.939	0.621	0.450	180.599	0.533	11.29
	DINAR [35]	14.975	0.640	0.419	115.243	0.345	11.04
	NTED [31]	16.846	0.665	0.353	76.561	0.322	3.39
	PIDM [4]	15.854	0.665	0.383	78.285	0.407	5.58
	T2IA [24]	5.049	0.380	0.634	225.615	0.517	8.22
	Ours	15.738	0.642	0.376	98.572	0.296	7.13
Average	LWG [16]	17.556	0.714	0.289	67.072	0.312	6.72
	PGHA [41]	16.671	0.698	0.342	118.435	0.425	5.81
	DINAR [35]	16.121	0.686	0.321	74.715	0.318	6.78
	NTED [31]	17.064	0.701	0.322	125.362	0.447	3.31
	PIDM [4]	15.626	0.667	0.351	132.255	0.464	5.39
	T2IA [24]	5.266	0.420	0.619	225.333	0.440	4.44
	Ours	18.487	0.739	0.260	60.171	0.242	3.28

method shows the best average scores across all datasets. These findings suggest that our method works across multiple datasets.

Qualitative comparison. Figure 9 shows the qualitative comparison of pose editing on various datasets. In the results of the first and second rows, the warping-based methods such as LWG [16], PGHA [41], and DINAR [35] yield stretched and unnatural textures in the invisible regions of the reference images. Moreover, in the third row, we can see that these methods burn the hand texture into the torso region. The image-to-image translation methods such as NTED [31] and PIDM [4] show good results on the DeepFashion dataset (the fifth row). However, they struggle to preserve clothing textures and facial identity on other




Figure 10. Qualitative comparison of editing the reference images to the target weights with our method and the method by Ren et al. [30].

datasets. The text-to-image approach, T2IA [24], also suffers from this problem and even generates different backgrounds from the reference images. Our method, on the other hand, consistently produces satisfactory results on all of the datasets. Our method successfully achieves a wide range of pose editing for a variety of person images, which is difficult to achieve with the existing methods.

4.2. Evaluation of body shape editing

To the best of our knowledge, there is no dataset available for objective quantitative evaluation of body shape editing with large deformation. We therefore conducted a qualitative comparison between our method and the state-of-the-art method proposed by Ren et al. [30]. The existing method controls the body shape using a warping strength as input instead of a body height and weight. For fair comparison, we searched for an appropriate warping strength so that the silhouette of the output image aligns with that of the SMPL-X model with a body height and weight specified in our method.

Qualitative comparison. Figure 10 shows the qualitative results. In the results of the existing method, increasing body size often causes significant distortion in the torso regions. In contrast, our method can create plausible images. In addition, our method can handle changes in the facial appearance that occur with changes in body weight.

Table 2. Quantitative results of the ablation study for diffusion-based refinement (Steps 1 and 2).

	LPIPS ↓	FID ↓	ID ↓	AW ↓
w/o Steps 1 & 2	0.242	55.862	0.232	3.68
w/o Step 2	0.239	45.949	0.403	1.74
Ours	0.242	46.432	0.175	1.82




Figure 11. Qualitative results of the ablation study for diffusion-based refinement (Steps 1 and 2). As shown in the red boxes in the second column, without both Steps 1 and 2, we can observe not only artifacts on the hands and pants but also a distorted face. The third column shows that Step 1 removes these artifacts (green boxes), but the identity of the face is still not improved (red boxes). Step 2 can address this problem, as shown in the green boxes in the fourth column.

4.3. Ablation studies

We conducted an ablation study on the MonoPerfCap dataset [40] to evaluate the effectiveness of our diffusion-based image refinement (i.e., Steps 1 and 2). Table 2 and Figure 11 show the quantitative and qualitative results, respectively. These results demonstrate that our fullbody refinement (Step 1) improves realism (FID) of the output images. This realism improvement also benefits the accuracy of pose estimation (Pose). However, we can see that facial identity (ID) degrades because the VAE used in the LDM cannot reconstruct relatively small faces accurately from low-dimensional latent maps. Our facial refinement (Step 2) can improve facial identity while preserving the scores of the other metrics.

In addition, we conducted another ablation study to evaluate our individual refinement approaches described in Section 3.1. Table 3 and Figure 12 show the quantitative and qualitative results, respectively. Here, “opt” means text embedding optimization, “iterate” means iterative refinement, and “reset” means input reinitialization. As shown in the comparison between “w/o opt & iterate & reset” and “w/o & iterate & reset,” text embedding optimization improves realism and facial identity. From the results of “w/o reset”

Table 3. Quantitative results of the ablation study for our refinement approaches including text embedding optimization (opt), iterative refinement (iterate), and input reinitialization (reset).

	LPIPS ↓	FID ↓	ID ↓	AW ↓
w/o opt & iterate & reset	0.244	56.093	0.404	1.93
w/o iterate & reset	0.240	47.770	0.183	2.04
w/o reset	0.277	56.670	0.187	2.37
Ours	0.242	46.432	0.175	1.82

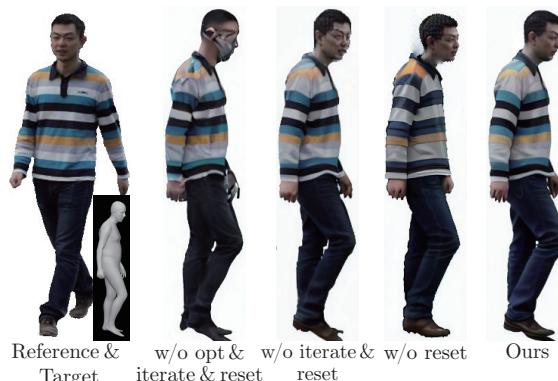


Figure 12. Qualitative results of the ablation study for our refinement approaches including text embedding optimization (opt), iterative refinement (iterate), and input reinitialization (reset).

and ours, we can also see that only iterative refinement does not perform well on its own but is effective when combined with input reinitialization.

5. Conclusion

In this paper, we proposed a method for editing the pose and body shape of a fullbody human image. Our method leverages a 3D parametric body model to control the pose and body shape using keypoints, height, and weight as input. Our two-step refinement pipeline is based on image-to-image translation with an LDM and refines the body and facial regions of the textured 3D human models. To improve the quality of the output images obtained via refinement, we introduced iterative refinement, text embedding optimization, and input reinitialization into the refinement pipeline. In the task of pose editing, our method exhibited quantitatively and qualitatively superior or comparable results compared to existing methods on multiple datasets. Our method also shows more plausible results in body shape editing. In the future, we will explore ways to speed up our method and enhance its capability to handle loose clothing like skirts.

References

- [1] Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. Pose with Style: Detail-preserving pose-guided image synthesis with conditional

- StyleGAN. *ACM Transactions on Graphics*, 40(6), 2021. 2
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. *ACM Transactions on Graphics*, 24(3):408–416, jul 2005. 3
- [3] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *CVPR*, pages 18208–18218, 2022. 4
- [4] 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. *CVPR*, pages 5968–5976, 2023. 1, 2, 5, 6, 7
- [5] Zhe Cao, Gines Hidalgo, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *TPAMI*, 43(1):172–186, 2021. 4, 5
- [6] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei Efros. Everybody dance now. In *ICCV*, pages 5932–5941, 2019. 5, 7
- [7] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. RetinaFace: Single-stage dense face localisation in the wild. In *CVPR*, pages 5202–5211, 2020. 5
- [8] Patrick Esser, Ekaterina Sutter, and Björn Ommer. A variational u-net for conditional appearance and shape generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8857–8866, 2018. 2
- [9] 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. *Advances in neural information processing systems*, 30, 2017. 5
- [10] Yangyi Huang, Hongwei Yi, Weiyang Liu, Haofan Wang, Boxi Wu, Wenxiao Wang, Binbin Lin, Debing Zhang, and Deng Cai. One-shot implicit animatable avatars with model-based priors. In *ICCV*, 2023. 4
- [11] Rishabh Jain, Krishna Kumar Singh, Mayur Hemani, Jingwan Lu, Mausoom Sarkar, Duygu Ceylan, and Balaji Krishnamurthy. VGFlow: Visibility guided flow network for human reposing. In *CVPR*, pages 21088–21097, 2023. 2
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [13] Jessica Lee, Deva Ramanan, and Rohit Girdhar. MetaPix: Few-Shot Video Retargeting. In *ICLR*, 2020. 5, 7
- [14] K. Li, J. Zhang, Y. Liu, Y. K. Lai, and Q. Dai. PoNA: Pose-guided non-local attention for human pose transfer. *IEEE Transactions on Image Processing*, 29:9584–9599, 2020. 2
- [15] Yining Li, Chen Huang, and Chen Change Loy. Dense intrinsic appearance flow for human pose transfer. In *CVPR*, 2019. 2
- [16] Wen Liu, Zhixin Piao, Zhi Tu, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping GAN with attention: A unified framework for human image synthesis. *TPAMI*, 44(9):5114–5132, 2022. 3, 5, 6, 7
- [17] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR 2016*, pages 1096–1104. IEEE Computer Society, 2016. 5, 7
- [18] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *ICLR*, 2017. 5
- [19] 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*, pages 10801–10810, 2021. 2
- [20] Liyuan Ma, Kejie Huang, Dongxu Wei, Zhao-Yan Ming, and Haibin Shen. FDA-GAN: Flow-based dual attention GAN for human pose transfer. *IEEE Transactions on Multimedia*, 25:930–941, 2021. 2
- [21] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. *Advances in neural information processing systems*, 30, 2017. 2
- [22] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jianjun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2022. 2, 3, 4
- [23] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. MagFace: A universal representation for face recognition and quality assessment. In *CVPR*, 2021. 5
- [24] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2I-Adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *CoRR*, abs/2302.08453, 2023. 3, 4, 5, 6, 7
- [25] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *ECCV*, pages 123–138, 2018. 2
- [26] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, pages 10975–10985, 2019. 2, 3, 5, 7
- [27] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, 2003. 5
- [28] Albert Pumarola, Jordi Sanchez, Gary Choi, Alberto Sanfeliu, and Francesc Moreno-Noguer. 3DPeople: Modeling the Geometry of Dressed Humans. In *ICCV*, pages 2242–2251, 2019. 5
- [29] 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*, volume 139, pages 8748–8763, 2021. 4, 5
- [30] Jianqiang Ren, Yuan Yao, Biwen Lei, Miaomiao Cui, and Xuansong Xie. Structure-aware flow generation for human body reshaping. In *CVPR*, pages 7744–7753, 2022. 3, 7
- [31] Yurui Ren, Xiaoqing Fan, Ge Li, Shan Liu, and Thomas H. Li. Neural texture extraction and distribution for controllable person image synthesis. *CoRR*, abs/2204.06160, 2022. 2, 5, 6, 7

- [32] Yurui Ren, Xiaoming Yu, Junming Chen, Thomas H. Li, and Ge Li. Deep image spatial transformation for person image generation. In *CVPR*, pages 7690–7699, June 2020. 2
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10674–10685, 2022. 3, 5
- [34] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, pages 22500–22510, 2023. 3, 5
- [35] David Svitov, Dmitrii Gudkov, Renat Bashirov, and Victor Lempitsky. DINAR: Diffusion inpainting of neural textures for one-shot human avatars. In *ICCV*, pages 7062–7072, 2023. 3, 5, 6, 7
- [36] Hao Tang, Song Bai, Li Zhang, Philip HS Torr, and Nicu Sebe. Xingan for person image generation. In *ECCV*, 2020. 2
- [37] Texel. Texel 3D body model dataset. <https://texel.graphics/texel-3d-body-model-dataset/>. accessed 28 August 2023. 5
- [38] 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*, pages 5998–6008, 2017. 2
- [39] Xinyao Wang, Liefeng Bo, and Li Fuxin. Adaptive wing loss for robust face alignment via heatmap regression. In *ICCV*, pages 6971–6981, 2019. 4
- [40] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. MonoPerfCap: Human performance capture from monocular video. *ACM Trans. Graph.*, 37(2):27:1–27:15, May 2018. 5, 7, 8
- [41] Jae Shin Yoon, Lingjie Liu, Vladislav Golyanik, Kripasindhu Sarkar, Hyun Soo Park, and Christian Theobalt. Pose-guided human animation from a single image in the wild. In *CVPR*, pages 15039–15048, 2021. 3, 5, 6, 7
- [42] Hongwen Zhang, Yating Tian, Xinchu Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. PyMAF: 3D human pose and shape regression with pyramidal mesh alignment feedback loop. In *ICCV*, pages 11446–11456, 2021. 5
- [43] Jinsong Zhang, Kun Li, Yu-Kun Lai, and Jingyu Yang. PISE: Person image synthesis and editing with decoupled GAN. In *CVPR*, pages 7982–7990, 2021. 2
- [44] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, pages 3836–3847, 2023. 3
- [45] Pengze Zhang, Lingxiao Yang, Jian-Huang Lai, and Xiaohua Xie. Exploring dual-task correlation for pose guided person image generation. In *CVPR*, pages 7713–7722, 2022. 2
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5
- [47] Shizhe Zhou, Hongbo Fu, Ligang Liu, Daniel Cohen-Or, and Xiaoguang Han. Parametric reshaping of human bodies in images. *ACM transactions on graphics*, 29(4):1–10, 2010. 3
- [48] Xinyue Zhou, Mingyu Yin, Xinyuan Chen, Li Sun, Changxin Gao, and Qingli Li. Cross attention based style distribution for controllable person image synthesis. In *ECCV*, pages 161–178, 2022. 2
- [49] Zhen Zhu, Tengpeng Huang, Baoguang Shi, Miao Yu, Bofei Wang, and Xiang Bai. Progressive pose attention transfer for person image generation. In *CVPR*, pages 2347–2356, 2019. 2