

# RealisHuman: A Two-Stage Approach for Refining Malformed Human Parts in Generated Images

Benzhi Wang<sup>1,2\*</sup>, Jingkai Zhou<sup>4\*</sup>, Jingqi Bai<sup>1,2\*</sup>,  
Yang Yang<sup>1,2</sup>, Weihua Chen<sup>4†</sup>, Fan Wang<sup>4</sup>, Zhen Lei<sup>1,2,3†</sup>

<sup>1</sup> State Key Laboratory of Multimodal Artificial Intelligence Systems,  
Institute of Automation, Chinese Academy of Sciences

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>3</sup> Centre for Artificial Intelligence and Robotics, Hong Kong Institute of Science & Innovation,  
Chinese Academy of Sciences

<sup>4</sup> Alibaba Group

wangbzzz@foxmail.com, {yang.yang, zlei}@nlpr.ia.ac.cn, baijingqi2022@ia.ac.cn  
{zhoujingkai.zjk, kugang.cwh, fan.w}@alibaba-inc.com

## Abstract

In recent years, diffusion models have revolutionized visual generation, outperforming traditional frameworks like Generative Adversarial Networks (GANs). However, generating images of humans with realistic semantic parts, such as hands and faces, remains a significant challenge due to their intricate structural complexity. To address this issue, we propose a novel post-processing solution named RealisHuman. The RealisHuman framework operates in two stages. First, it generates realistic human parts, such as hands or faces, using the original malformed parts as references, ensuring consistent details with the original image. Second, it seamlessly integrates the rectified human parts back into their corresponding positions by repainting the surrounding areas to ensure smooth and realistic blending. The RealisHuman framework significantly enhances the realism of human generation, as demonstrated by notable improvements in both qualitative and quantitative metrics.

**Code** — <https://github.com/Wangbenzhi/RealisHuman>

## Introduction

Diffusion models have emerged as a powerful approach in the field of visual generation, significantly surpassing traditional frameworks such as Generative Adversarial Networks (GANs) (Goodfellow et al. 2014). These models function as parameterized Markov chains, showcasing an exceptional capability to convert random noise into complex images through a sequential refinement process. Starting with noise, diffusion models progressively enhance the visual quality, ultimately producing high-fidelity representations. With ongoing technological advancements, diffusion models have shown substantial promise in image generation and various related tasks (Podell et al. 2024; Rombach et al. 2022a; Guo et al. 2023).

Despite their remarkable performance in generating a diverse range of objects, diffusion-based models encounter

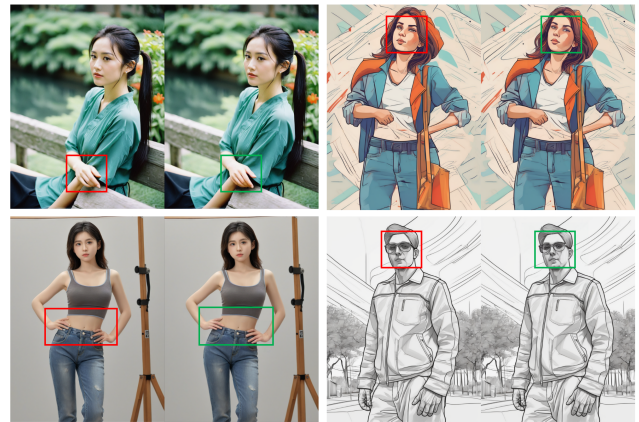


Figure 1: Illustration of our repair results. Each pair consists of two images: the left image is the original, and the right image is the repair result.

significant challenges when reconstructing realistic human features, particularly faces and hands. The intricate structural complexity of these parts, coupled with the limited information preserved after VAE encoder downsampling (Kingma and Welling 2013), often leads to incorrect hand structures or distorted faces. As depicted in Fig.1, these inaccuracies highlight the difficulties these models face in human image generation.

To address this issue, HandRefiner (Lu et al. 2023) proposed a lightweight post-processing solution that employs a conditional inpainting approach to correct malformed hands while preserving other image regions. Utilizing a hand mesh reconstruction model, HandRefiner ensures accurate finger counts and hand shapes, fitting the desired hand pose. By leveraging ControlNet modules, HandRefiner reintegrates correct hand information into the generated images, enhancing overall image quality. However, this method has two notable limitations. As illustrated in Fig.3, HandRefiner often fails to maintain consistency in skin tone and texture due to missing reference information. It also struggles with recon-

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

\*Equal contribution.

†Corresponding Authors

structuring detailed hands when the regions are small. Additionally, it can introduce distortions in other areas, like the face, compromising the overall image integrity.

In this paper, we propose a novel post-processing solution named RealisHuman to address the challenge of refining malformed human parts. To ensure high-quality refinements in small regions, our method locates and crops the malformed areas, allowing us to concentrate on detailed local refinements. Compared to HandRefiner, our method is capable of refining various human parts, not just hands, while preserving intricate details such as skin tone and texture. This capability ensures that the refined parts are both realistic and consistent with the surrounding image. Additionally, our approach demonstrates strong generalization capabilities, effectively handling different styles of images, including cartoons, sketches, and so on. As shown in Fig.2, our RealisHuman framework operates in two stages. In the first stage, our goal is to generate rectified human parts that preserve the consistent details of the original malformed parts. By using the malformed parts as references, we extract detailed information through the Part Detail Encoder and DINOv2, ensuring the preservation of fine-grained details and enhancing the overall realism of the generated parts. Additionally, we incorporate 3D pose estimation results extracted from the malformed parts to guide the generation of human part images, ensuring that the poses are both accurate and realistic. After obtaining the rectified human parts, the subsequent challenge is to seamlessly integrate them into the original local image. We address this as an inpainting problem. Initially, the rectified human parts are placed back into their original positions, and the surrounding areas are masked. We then train a model capable of seamlessly blending the human parts with the surrounding areas, ensuring a smooth transition and realistic integration. Finally, the refined human parts are pasted into the original image, completing the process of malformed human parts refinement. This approach not only corrects structural inaccuracies but also maintains visual coherence with the original image, providing a robust solution for human parts refinement in image generation tasks. The RealisHuman framework significantly enhances the realism of human generation, as validated by comprehensive experiments demonstrating improvements in both qualitative and quantitative measures.

Our contributions are summarized as follows:

- We propose a novel post-processing framework named RealisHuman to address the task of refining human parts in generated images. Our method maintains consistent details with the original image, effectively handles small part refinements, and demonstrates strong generalization across different image styles.
- We propose a novel two-stage local refinement paradigm, which can be extended to the refinement of other structurally fixed objects, such as distorted logos.
- The RealisHuman framework significantly enhances the realism of human generation, as evidenced by extensive experiments demonstrating enhancements in both qualitative and quantitative metrics.

## Related Work

**Diffusion Model for Image Generation.** Diffusion models have emerged as a powerful approach in computer vision due to their ability to generate high-fidelity and diverse images. Despite their success, image generation’s high dimensionality poses computational challenges. The Latent Diffusion Model (LDM) (Rombach et al. 2022b) addresses this by performing denoising in a lower-dimensional latent space using a pre-trained autoencoder. This approach effectively balances computational efficiency with generative performance, representing a pivotal advancement in the scalability of diffusion-based image generation. Despite these advancements, controlling the generative process of diffusion models remains a challenge, particularly when precise semantic adherence is required. Diffusion models have achieved great success in producing realistic images that adhere to the semantic content provided by encoding text inputs into latent vectors via pre-trained language models like CLIP (Radford et al. 2021). However, relying solely on text descriptions for controlling the model is insufficient, especially when it comes to describing postures and actions (Ye et al. 2023). To enhance controllability and precision in generated imagery, researchers have explored the incorporation of additional control signals. ControlNet (Zhang, Rao, and Agrawala 2023) employs a trainable duplicate of the Stable Diffusion (SD) encoder architecture to extract features from conditional inputs. Similarly, T2I-Adapter (Mou et al. 2024) utilizes lightweight, composable adapter blocks for feature extraction. These additional conditional layers have proven instrumental in improving the model’s controllability under various conditions, such as pose, mask, and edge, thereby significantly influencing the direction of its output.

**Realistic Human Image Generation.** Diffusion models have been widely applied in pose-conditioned human image synthesis. Animate Anyone (Hu 2024) employs ReferenceNet, a symmetrical UNet, for capturing spatial details from reference images. MagicAnimate (Xu et al. 2024) uses ControlNet with DensePose (Güler, Neverova, and Kokkinos 2018) for more precise pose guidance. Champ (Zhu et al. 2024) integrates multiple control signals like depth, normal, semantic, and skeleton extracted from SMPL (Loper et al. 2023) models. Despite notable progress in generating high-quality synthetic human images, synthesizing realistic hands remains challenging due to their complex anatomy and skeletal representation limits. Some approaches have begun to specifically focus on generating higher-quality hands. Diffusion-HPC (Weng, Bravo-Sánchez, and Yeung 2023) introduces a technique that employs depth maps of human bodies rendered from reconstructed human body meshes, utilizing conditional diffusion models to correct morphological abnormalities in generated human bodies. Similarly, HandRefiner (Lu et al. 2023) proposes a post-processing approach that utilizes a reconstructed hand mesh to provide essential information about hand shape and location. While these methods can be effective in addressing distortions in hand morphology, they often fall short in preserving fine details such as skin tone consistency and texture.

To enhance the realism of human generation, we propose a two-stage post-processing method named RealisHuman.

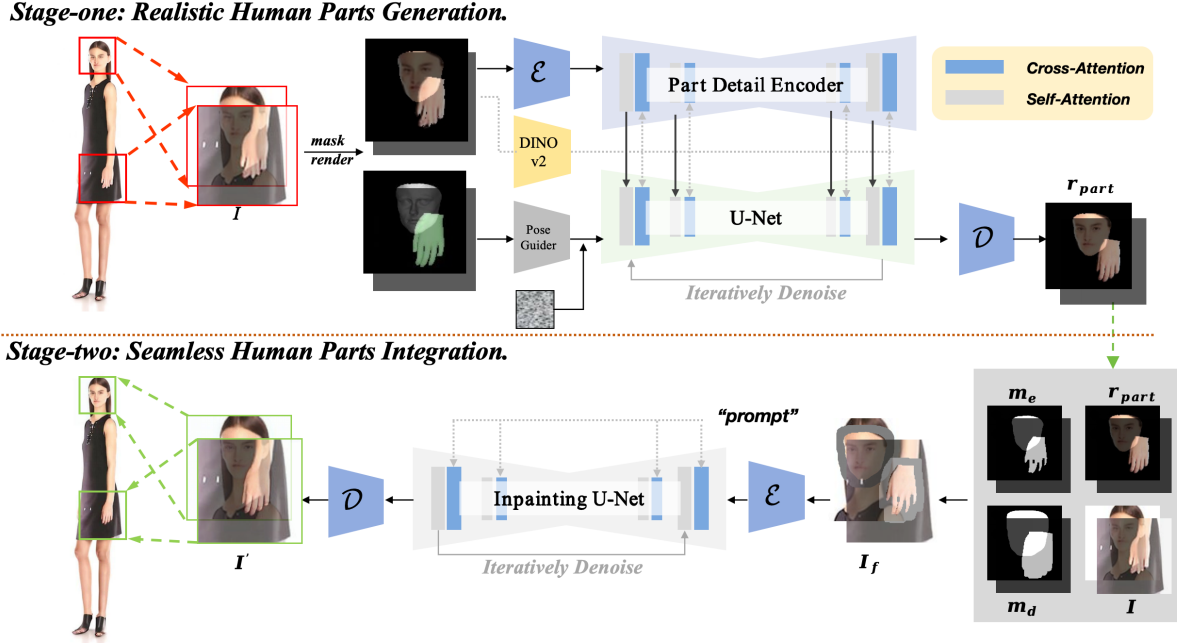


Figure 2: Details of our RealisHuman. Our method separates the task of refining malformed human parts into two distinct stages. In the first stage, we focus on generating realistic human parts using the Part Detail Encoder. Given an image containing malformed human parts, we begin by locating and cropping the target regions. Subsequently, we filter the background of the target regions, creating a reference image that provides essential part details, such as skin tone. We also estimate the 3D structure of the human parts to serve as pose guidance. Leveraging both the reference images and the part structures, we generate realistic human parts  $r_{part}$  with accurate structures and detailed information. In the second stage, our goal is to seamlessly integrate the refined human parts into the corresponding regions of the original image, resulting in the refined image  $I'$ . To avoid a cut-and-paste appearance, we repaint the area between the background and the rectified human parts, ensuring a seamless integration and a more natural overall appearance.

In the first stage, our method rectifies malformed parts by utilizing detailed information and 3D pose estimation results from the original malformed parts. In the second stage, we seamlessly integrate the rectified human parts back into the original image to complete the refinement process.

## Method

Our goal is to refine the malformed parts while preserving the consistent details of the original parts. The overall framework pipeline is depicted in Fig.2. To ensure the realism of the rectified human parts, the pipeline is divided into two distinct stages. In the first stage, the rectified human parts are generated under the guidance of the parts meshes and the malformed part images. In the second stage, the rectified human parts obtained from the first stage are integrated back into the local image, followed by repainting the surrounding region to achieve the final results.

### Preliminary

**Latent Diffusion Models.** Our approach builds upon the foundation of Stable Diffusion (SD)(Rombach et al. 2022a), which originates from the Latent Diffusion Model (LDM).

LDMs are designed to operate within the latent space managed by an autoencoder, specifically  $\mathcal{D}(\mathcal{E}(\cdot))$ . A prime example of these models is Stable Diffusion (SD), which combines a Variational AutoEncoder (VAE)(Kingma and Welling 2013) and a time-conditioned U-Net(Ronneberger, Fischer, and Brox 2015) to estimate noise. For handling text inputs, SD uses a CLIP ViT-L/14(Radford et al. 2021) text encoder to transform textual queries into embeddings, denoted as  $c_{text}$ .

In the training phase, the model processes an image  $I$  and a corresponding text condition  $c_{text}$ . The image is encoded into a latent representation  $z_0 = \mathcal{E}(I)$ , which then undergoes a predefined sequence of  $T$  diffusion steps governed by a Gaussian process, resulting in a noisy latent representation  $z_T \sim \mathcal{N}(0, 1)$ . The objective of SD is to iteratively refine  $z_T$  back to  $z_0$ , using the following loss function:

$$L = \mathbb{E}_{E(I), c_{text}, \epsilon \sim \mathcal{N}(0,1), t} \left[ \|\epsilon - \epsilon_\theta(z_t, t, c_{text})\|_2^2 \right], \quad (1)$$

where  $t = 1, \dots, T$  denotes the timestep embedding.  $\epsilon_\theta$  denotes the trainable components within the denoising U-Net, which processes the noisy latents  $z_t$  and the text condition  $c_{text}$ . The architecture of the U-Net includes convolutional

layers (Residual Blocks) and both self-attention and cross-attention mechanisms (Transformer Blocks).

The training process involves encoding the image into a latent form  $z_0$  and subjecting it to a sequence of diffusion steps, producing  $z_T$ . The denoising U-Net is trained to predict and remove the noise added during these steps. Once trained, the model can generate  $z_0$  from  $z_T$  using a deterministic sampling method (such as DDIM(Song, Meng, and Ermon 2020)), and the final image is reconstructed through the decoder  $\mathcal{D}$ .

During inference, the initial latent  $z_T$  is sampled from a Gaussian distribution with the initial timestep  $T$  and gradually refined through iterative denoising steps to yield  $z_0$ . At each step, the U-Net predicts the noise present in the latent features corresponding to that specific timestep. The decoder  $\mathcal{D}$  then reconstructs the final image from  $z_0$ .

### Realistic Human Parts Generation

In the first stage, our objective is to generate realistic parts that maintain consistent detail and pose with the original images. This is achieved by using the guidance of meshes and reference information from the malformed parts. Leveraging these, we ensure the rectified human parts match the intended appearance and pose.

**Data preparation.** Suppose we have a series of original human images and corresponding generated images that contain malformed human parts, produced by algorithms such as (Rombach et al. 2022a; Zhu et al. 2024; Hu 2024; Wang et al. 2023a; Chang et al. 2024; Xu et al. 2024; Karras et al. 2023). We begin by locating and cropping the target part regions using the human skeleton estimation method (Yang et al. 2023). After isolating the target part regions, we employ the state-of-the-art (SOTA) mesh reconstruction method (Pavlakos et al. 2024; Wang et al. 2023b) to estimate the meshes for each part. Additionally, we render the meshes to produce depth maps and binary mask maps  $m$ . To reduce the influence of the background and focus on realistic human parts generation, we apply the mask  $m$  to filter out the background and obtain the foreground regions of the human parts as reference images  $I_{ref}$ .

**Part Detail Encoder.** Previous image-conditioned generation tasks (Wang et al. 2023a; Karras et al. 2023) have typically utilized the CLIP image encoder (Radford et al. 2021) to encode reference images. Specifically, these methods compress reference images from a spatial size of  $224 \times 224 \times 3$  into a one-dimensional vector of dimension 1024, and then employ cross-attention mechanisms to integrate the latent representation with this vector. However, these approaches face challenges in preserving appearance details, as encoding reference images into semantic-level features results in a loss of spatial representations. Previous works (Hu 2024; Cao et al. 2023; Chang et al. 2024) have demonstrated that the self-attention mechanism can significantly enhance the preservation of detail in reference images. Inspired by these findings, we introduce the Part Detail Encoder to improve the realism of rectified human parts by integrating detailed information from the reference images  $I_{ref}$ . The Part Detail Encoder shares the same architecture as the original Stable Diffusion (SD), comprising self-

attention and cross-attention layers, and is initialized with the original SD UNet. To achieve this, we use the reference images as input to the Part Detail Encoder and obtain intermediate outputs. To better integrate detailed information, we modify the input to the self-attention mechanism of the UNet. Specifically, we concatenate the intermediate outputs of the Part Detail Encoder with those of the original SD, and use this concatenated output as the input to the self-attention mechanism of the original SD. This approach ensures that fine-grained details are preserved, enhancing the overall realism of the generated human parts. The modified self-attention mechanism can be formulated as:

$$f_s = \text{softmax} \left( \frac{Q_o \cdot (K_o \oplus K_h)^T}{\sqrt{d}} \right) \cdot (V_o \oplus V_h), \quad (2)$$

where  $d$  is the feature dimension.  $Q_o$ ,  $K_o$ , and  $V_o$  denote the query, key, and value from the self-attention layers of the original SD, respectively. Meanwhile,  $K_h$  and  $V_h$  denote the key and value from the self-attention layers of the Part Detail Encoder.

Meanwhile, we employ DINOv2(Oquab et al. 2023) to get the image embedding  $c_r$  of the reference image, which is then passed into the Part Detail Encoder and the UNet through a cross-attention mechanism. This approach supplements the semantic-level features of the reference image. The depth map is processed through several convolution layers to obtain the pose condition  $c_p$ , which is then added to the noise latent before being input into UNet, as described in (Hu 2024).

**Training.** With the design of above, the loss term of this stage is computed as:

$$\mathcal{L}_1 = \mathbb{E}_{z_0, c_p, c_r, I_{ref}, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(z_t, c_p, c_r, I_{ref}, t)\|_2^2], \quad (3)$$

where  $\epsilon_\theta$  denotes the trainable parameters of the UNet and  $t$  is the timestep embedding.

### Seamless Human Parts Integration

Another issue is that directly pasting back the rectified human parts  $r_{part}$  introduces copy-and-paste artifacts in the edited region, making the generated image appear unnatural. To address this issue, we repaint the area between the background and the rectified human parts, seamlessly integrating them into the target region for a more natural appearance.

**Data Preparation.** Given an image containing human parts like the face or hands, we first locate and crop the target regions and obtain the binary masks using the same approach mentioned in the first stage. For each part, we dilate its binary mask  $m$  using the kernel  $k_d$  to obtain the dilated mask  $m_d = \text{dilate}(m, k_d)$ . Additionally, we erode the binary mask  $m$  with a small kernel  $k_e$  to obtain the eroded mask  $m_e = \text{erode}(m, k_e)$ . Using the eroded mask, we extract the eroded human part and paste it back into the corresponding region. The erosion process is crucial because the rectified human parts generated in the first stage often exhibit inharmonious edges, which significantly affect the repainting results. By eroding the human part regions, we aim to equip the model with the ability to complete human part edges during the repainting process. This approach helps mitigate

issues caused by inharmonious edges, resulting in a more natural and seamless integration of the rectified human parts into the target regions. Suppose the local human part image is denoted as  $I$ . The corresponding masked image and binary mask can be formulated with Eq. 4 and Eq. 5.

$$I_f = I \odot (1 - m_d) + I \odot m_e, \quad (4)$$

$$m_f = m_d - m_e. \quad (5)$$

Our goal is to predict the area where the binary mask  $m_f$  equals one while keeping the other areas unchanged, resulting in the final output  $I'$ . To achieve this, we first encode the masked image  $I_f$  to obtain the masked latent  $l_m = \mathcal{E}(I_f)$ . Next, we downsample the binary mask  $m_f$  to match the size of the masked latent  $l_m$ . Similar to SD-inpainting, we add five additional input channels for the UNet: four for the encoded masked image  $l_m$  and one for the mask  $m_f$ . Additionally, we initialize the model with SD-inpainting weights. With this design, the loss term for this stage is computed as follows:

$$\mathcal{L}_2 = \mathbb{E}_{z_0, l_m, m_f, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \epsilon_\theta(z_t, l_m, m_f, t)\|_2^2], \quad (6)$$

where  $t$  is the timestep embedding.

During inference, we paste the rectified human part  $r_{part}$  back into the corresponding region and predict the unknown area to ensure harmonious integration of the rectified human part. The formulation of  $I_f$  during the inference process is given by Eq. 7:

$$I_f = I \odot (1 - m_d) + r_{part} \odot m_e. \quad (7)$$

## Experiments

In this section, we begin by detailing the implementation aspects of our approach, followed by a description of the datasets and evaluation protocols used. Additionally, We present comparative experiments to benchmark our method against previous work, and conduct ablation studies to assess the efficacy of each component in our framework.

Our RealisHuman is trained in two stages: realistic human parts generation and seamlessly integrated the human parts. All experiments are conducted on 8 NVIDIA A800 GPUs. In the first stage, both the main UNet and the Part Detail Encoder are initialized from Real Vision v5.1, and all components are optimizable except for DINOv2(Oquab et al. 2023) and VAE encoder/decoder(Kingma and Welling 2013). Training is conducted for 50,000 steps with a batch size of 5. In the second stage, only the Inpainting U-Net is optimizable, which is initialized from SD-inpainting(Rombach et al. 2022a). We train the Inpainting U-Net for 20,000 steps with a batch size of 16. For both two stages, the learning rate is set to  $5e-5$ . The image is resize to a resolution of  $512 \times 512$ . The zero-SNR(Lin et al. 2024) and classifier-free guidance(CFG) (Ho and Salimans 2022) are enabled. The unconditional drop rate is set to  $1e-2$ . We employ HaMeR(Pavlakos et al. 2024) and 3DDFAv3(Wang et al. 2023b) to estimate the meshes for each human part. During inference, we adopt a DDIM sampler for 20 denoising steps. We set the hyper-parameter  $g_d$  to 5 and  $g_e$  to 0.05 times the perimeter of the mask. The images demonstrated

in our paper are generated by SDXL(Podell et al. 2024) and SDXL-LEOSAM\*.

## Datasets and Evaluation Protocol

We have collected a dataset comprising approximately 58,000 high-quality local hand images and 38,000 high-quality local face images for training our model. To demonstrate the effectiveness of our approach for refining malformed parts, we evaluate its performance on UBC Fashion(Zablotskaia et al. 2019) dataset. The human subjects in UBC Fashion exhibit clearly visible hands and faces. UBC Fashion consists of 500 training and 100 testing videos, each containing roughly 350 frames. We follow the official train/test split for both UBC Fashion. Specifically, we use Fréchet Inception Distance (FID)(Heusel et al. 2017) and the keypoint detection confidence scores of a hand detector or face detector(Lugaresi et al. 2019; Zhang et al. 2020) to evaluate the plausibility of the generated human parts.

## Results and Comparisons

We generate human images with pose guidance on the UBC Fashion dataset using the most advanced human synthesis methods(Hu 2024; Xu et al. 2024; Zhu et al. 2024). After generating the human images, we located and cropped the regions containing human parts and applied our RealisHuman framework to refine the malformed parts. To mitigate the influence of the relatively small size of human parts in the original images and to better evaluate the metrics, we focused the evaluation specifically on the regions containing human parts.

Method	Hand		Face	
	FID ↓	Det. Conf. ↑	FID ↓	Det. Conf. ↑
AnimateAnyone(Hu 2024)	14.26	0.86	20.55	0.82
AnimateAnyone+Ours	13.02	0.91	15.44	0.90
Champ(Zhu et al. 2024)	27.28	0.87	20.11	0.85
Champ+Ours	25.58	0.92	16.74	0.92
MagicAnimate(Xu et al. 2024)	57.73	0.90	43.12	0.87
MagicAnimate+Ours	55.18	0.94	38.81	0.92

Table 1: Comparison of FID and Det. Conf. scores before and after using our method.

In Tab.1, we report the FID and Det. Conf. scores before and after using our RealisHuman for both face and hand regions. The results demonstrate the effectiveness of our method. Specifically, we observe significant improvements in both metrics after applying our refinement process. The reduction in FID scores indicates that the refined images are perceptually closer to real images, showcasing enhanced realism. Similarly, the increase in Det. Conf. scores reflects improved detection confidence by the detectors, highlighting the structural accuracy and plausibility of the refined face and hand regions.

To evaluate the effectiveness of our method in refining hand images, we compare our method with the popular malformed hands refining method HandRefiner in Fig.3. Additionally, we conduct a detailed analysis to illustrate the advantages of our approach. As shown in Fig.3, each com-

\*<https://civitai.com/models/43977/leosams-helloworld-xl>



Figure 3: Comparison of hand refinement results. Each set of images displays, from left to right, the original image, our method’s repair result, and the HandRefiner method’s repair result.

parison figure consists of three horizontally aligned images: from left to right, they display the original image, our method’s repair result, and the HandRefiner method’s repair result. This figure presents a comprehensive comparison between our method and the HandRefiner method across several critical aspects: **(a) Preservation of Hand Details:** Our method excels at maintaining and matching the original details, such as the skin tone of the hands. It demonstrates superior consistency in preserving intricate details, accurately restoring textures and fine features of the hands. As a result, the repaired hands have a more natural and realistic appearance. **(b) Effectiveness in Small Hand Repair:** Compare to HandRefiner, our method is particularly effective in repairing smaller hands, meticulously restoring their details and shapes. **(c) Preservation of Other Regions:** Unlike HandRefiner, which can cause distortions in other areas such as the face while repairing hands, our method preserves the overall integrity and appearance of the image, ensuring that other regions remain unaffected. Effectively showcases these advantages, highlighting the superior performance of our method in hand repair tasks compared to the HandRefiner method. This comparison underscores the efficacy and reliability of our approach in producing high-quality hand restorations.

Additionally, we demonstrate the capability of RealisHuman in facial refinement. As shown in Fig.4, our method effectively addresses issues such as distorted facial features

and unfocused eyes in the original images, highlighting the efficacy of our approach. The results illustrate that RealisHuman can significantly enhance the realism and accuracy of facial features, further validating the robustness, versatility, and strong generalization capability of our method across various styles of human image restoration. For additional examples and details, please refer to the supplementary materials.

### Abalation Study

**Effect of the second stage.** As discussed above, we address the issue of copy-and-paste artifacts by repainting the transition area between the background and the rectified human parts, ensuring seamless integration into the target region for a more natural appearance. Fig.6 compares the results of directly pasting the rectified human part  $r_{part}$  with our method. It can be observed that our approach effectively integrates the rectified human parts into the surrounding area without introducing copy-and-paste artifacts.

**Effect of the eroded mask  $m_e$ .** As discussed above, the eroded mask  $m_e$  is used to mitigate the effects of inharmonious edges in the second stage. Without the eroded mask, these inharmonious edges can hinder the seamless integration of the rectified human parts with their surroundings, leading to the generation of discordant elements such as hair, watches, and other artifacts. We illustrate the impact of the eroded mask in Fig.6, comparing images processed with and



Figure 4: Comparison of face refinement results. The first row shows the original images, and the second row shows the images after face refinement.



Figure 5: Comparison of directly pasting the rectified human parts versus our method. The first row shows the results of direct pasting, which exhibit visible artifacts. The second row demonstrates the effectiveness of our method in achieving seamless integration.

without it. The first row shows the results without the eroded mask, where noticeable artifacts are present. The second row demonstrates the results when using the eroded mask, which effectively reduces edge artifacts and achieves a smoother integration.



Figure 6: Comparison of results with and without the eroded mask  $m_e$ . The first row shows visible edge artifacts without the eroded mask. The second row demonstrates improved integration using the eroded mask, reducing edge artifacts.

### Limitations and Discussion

While our method has demonstrated notable improvements in refining and reconstructing human hands, it still faces several challenges, as illustrated in Fig.7. Firstly, the method may struggle to accurately reconstruct interaction between

hands and objects. Secondly, it may fail to maintain consistency in the presence of objects. Thirdly, when the original hand is severely distorted, the method may be unable to estimate the correct hand pose, leading to unsuccessful hand reconstruction. Addressing these issues will be the focus of our future work, potentially incorporating more sophisticated modeling techniques or leveraging additional contextual information to improve performance in these areas.




Figure 7: Illustration of the limitations of our method.

### Conclusion

In this paper, we introduced RealisHuman, a novel post-processing solution for refining malformed human parts in generated images. Our method operates in two stages: first, generating realistic human parts using the original malformed human parts as the reference to maintain consistent details; second, seamlessly integrating the rectified human parts by repainting the surrounding areas. This framework effectively addresses the challenges of human parts generation and can be extended to other local refinement tasks, such as logo refinement. Comprehensive experiments demonstrate significant improvements in both qualitative and quantitative measures, validating the effectiveness and robustness of our approach.

### Acknowledgements

This work was supported in part by the Chinese National Natural Science Foundation Project 62276254, 62206276, U23B2054, and the InnoHK program.

## References

- Cao, M.; Wang, X.; Qi, Z.; Shan, Y.; Qie, X.; and Zheng, Y. 2023. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22560–22570.
- Chang, D.; Shi, Y.; Gao, Q.; Xu, H.; Fu, J.; Song, G.; Yan, Q.; Zhu, Y.; Yang, X.; and Soleymani, M. 2024. MagicPose: Realistic Human Poses and Facial Expressions Retargeting with Identity-aware Diffusion. In *International Conference on Machine Learning*.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. 139–144.
- Güler, R. A.; Neverova, N.; and Kokkinos, I. 2018. Densepose: Dense human pose estimation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7297–7306.
- Guo, Y.; Yang, C.; Rao, A.; Wang, Y.; Qiao, Y.; Lin, D.; and Dai, B. 2023. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 6629–6640.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Hu, L. 2024. Animate Anyone: Consistent and Controllable Image-to-Video Synthesis for Character Animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8153–8163.
- Karras, J.; Holynski, A.; Wang, T.-C.; and Kemelmacher-Shlizerman, I. 2023. Dreampose: Fashion image-to-video synthesis via stable diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22623–22633.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Lin, S.; Liu, B.; Li, J.; and Yang, X. 2024. Common diffusion noise schedules and sample steps are flawed. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5404–5411.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2023. SMPL: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 851–866.
- Lu, W.; Xu, Y.; Zhang, J.; Wang, C.; and Tao, D. 2023. HandRefiner: Refining Malformed Hands in Generated Images by Diffusion-based Conditional Inpainting. *arXiv preprint arXiv:2311.17957*.
- Lugaresi, C.; Tang, J.; Nash, H.; McClanahan, C.; Uboweja, E.; Hays, M.; Zhang, F.; Chang, C.-L.; Yong, M. G.; Lee, J.; et al. 2019. Mediapipe: A framework for building perception pipelines. *arXiv preprint arXiv:1906.08172*.
- Mou, C.; Wang, X.; Xie, L.; Wu, Y.; Zhang, J.; Qi, Z.; and Shan, Y. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4296–4304.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Pavlakos, G.; Shan, D.; Radosavovic, I.; Kanazawa, A.; Fouhey, D.; and Malik, J. 2024. Reconstructing hands in 3d with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9826–9836.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2024. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *International Conference on Learning Representations*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022a. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022b. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Wang, T.; Li, L.; Lin, K.; Zhai, Y.; Lin, C.-C.; Yang, Z.; Zhang, H.; Liu, Z.; and Wang, L. 2023a. Disco: Disentangled control for realistic human dance generation. *arXiv preprint arXiv:2307.00040*.
- Wang, Z.; Zhu, X.; Zhang, T.; Wang, B.; and Lei, Z. 2023b. 3D Face Reconstruction with the Geometric Guidance of Facial Part Segmentation. *arXiv preprint arXiv:2312.00311*.
- Weng, Z.; Bravo-Sánchez, L.; and Yeung, S. 2023. Diffusion-hpc: Generating synthetic images with realistic humans. *arXiv preprint arXiv:2303.09541*.
- Xu, Z.; Zhang, J.; Liew, J. H.; Yan, H.; Liu, J.-W.; Zhang, C.; Feng, J.; and Shou, M. Z. 2024. Magicanimate: Temporally consistent human image animation using diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1481–1490.

- Yang, Z.; Zeng, A.; Yuan, C.; and Li, Y. 2023. Effective whole-body pose estimation with two-stages distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4210–4220.
- Ye, Y.; Li, X.; Gupta, A.; De Mello, S.; Birchfield, S.; Song, J.; Tulsiani, S.; and Liu, S. 2023. Affordance diffusion: Synthesizing hand-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22479–22489.
- Zablotskaia, P.; Siarohin, A.; Zhao, B.; and Sigal, L. 2019. Dwnet: Dense warp-based network for pose-guided human video generation. *arXiv preprint arXiv:1910.09139*.
- Zhang, F.; Bazarevsky, V.; Vakunov, A.; Tkachenka, A.; Sung, G.; Chang, C.-L.; and Grundmann, M. 2020. Mediapipe hands: On-device real-time hand tracking. *arXiv preprint arXiv:2006.10214*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhu, S.; Chen, J. L.; Dai, Z.; Xu, Y.; Cao, X.; Yao, Y.; Zhu, H.; and Zhu, S. 2024. Champ: Controllable and consistent human image animation with 3d parametric guidance. *arXiv preprint arXiv:2403.14781*.