Total Selfie: Generating Full-Body Selfies

BOWEI CHEN, University of Washington, USA BRIAN CURLESS, University of Washington, USA IRA KEMELMACHER-SHLIZERMAN, University of Washington, USA STEVE SEITZ, University of Washington, USA



Fig. 1. We generate full-body selfies (right), similar to a photo someone else would take of you at a given scene. At the start of the day, the users pre-capture a video of themselves with their shoes, pants, and outfit (left). They can also choose a target pose(s) for the day's photos. Then they can proceed to take any number of on-site image pairs (selfie + background; middle) to produce a full-body image at each location.

We present a method to generate full-body selfies – photos that you take of yourself, but capturing your whole body as if someone else took the photo of you from a few feet away. Our approach takes as input a precaptured video of your body, a target pose photo, and a selfie + background pair for each location. We introduce a novel diffusion-based approach to combine all of this information into high quality, well-composed photos of

Authors' addresses: Bowei Chen, University of Washington, 1410 NE Campus Pkwy, Seattle, WA, 98195, USA, boweiche@cs.washington.edu; Brian Curless, University of Washington, Seattle, USA, curless@cs.washington.edu; Ira Kemelmacher-Shlizerman, University of Washington, Seattle, USA, kemelmi@cs.washington.edu; Steve Seitz, University of Washington, Seattle, USA, seitz@cs.washington.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. © 2023 Association for Computing Machinery.

0730-0301/2023/8-ART \$15.00

https://doi.org/10.1145/nnnnnnnnnnn

you with the desired pose and background. **Project Webiste**: https://homes. cs.washington.edu/~boweiche/project_page/totalselfie/.

 $\label{eq:CCS} \text{Concepts:} \bullet \textbf{Computing methodologies} \to \textbf{Computer graphics}.$

Additional Key Words and Phrases: Full-Body Generation from Selfies, Stable Diffusion

ACM Reference Format:

Bowei Chen, Brian Curless, Ira Kemelmacher-Shlizerman, and Steve Seitz. 2023. Total Selfie: Generating Full-Body Selfies. *ACM Trans. Graph.* 1, 1 (August 2023), 12 pages. https://doi.org/10.1145/nnnnnnnnnn

1 INTRODUCTION

The prevalence of selfies has skyrocketed in recent years, with an estimated 93 million taken each day. Despite their popularity, they suffer from multiple shortcomings: (1) they capture only the upper portion of the subject, (2) the close-up camera viewpoint distorts faces and requires awkward poses (e.g., with arm reaching out) and

2 · Chen, Curless. et al

(3) it is difficult to compose a shot that optimally captures both the subject and the scene.

Instead, what if you could capture the full-body image that *some-one else would take of you* in the scene? We call this a *total selfie*. As input, we require a normal selfie in the desired scene, and a photo of the background that you would like to be composited into (Fig. 1 center). This information, however, is not quite enough to compose the total selfie, as it does not contain (1) the lower half of your body and (2) your desired body pose. We therefore additionally require a pre-capture of (1) and (2). Based on this information, we generate convincing full-body photos of you in the specified pose, in the desired scene. In practice, this process requires taking only two photos at each selfie site, and a single pre-capture of your outfit that you can do once per day.

Solving this problem requires addressing a number of challenges. First, we must reconstruct the bottom part of your body, piecing together elements (e.g., a photo of your shoes, another of your pants) from the pre-capture phase. Second, we need to reproject you to a virtual viewpoint from several feet away – far enough to compose your full body within the scene. And third, we need to render you with a target pose, specified in a full-body photo of you – this can be any image from your photo collection, potentially in different clothes, lighting, and in any scene, unrelated to the selfie. Importantly, the resulting composite must retain your identity, expression, and clothing, but composited realistically into the target scene with the desired pose.

Our approach builds on Stable Diffusion [Rombach et al. 2022] and has three components. First, we employ a multi-concept Dream-Booth fine-tuning approach [Ruiz et al. 2022] to train a personspecific diffusion model for the subject based on the pre-capture video. Our approach is pose-aware (via ControlNet [Zhang and Agrawala 2023]) and enables composition into new backgrounds. Second, we employ an *appearance refinement* method, to preserve the facial identity and expression from the normal selfie, while removing perspective distortion. This step also refines the inaccurate clothes generated from the previous step. And third, we introduce an *image harmonization* step to improve the composition of the subject and background.

We demonstrate results for five individuals in various scenes with a wide range of expressions. Our experimental results demonstrate that Total Selfie outperforms existing methods in generating realistic and accurate full-body images.

2 RELATED WORK

2.1 Full-Body Image Generation

Extensive research has been conducted on generating full-body images using Generative Adversarial Networks (GANs) [Goodfellow et al. 2020] or diffusion models [Croitoru et al. 2023]. [Fu et al. 2022] trained a StyleGAN-based model for unconditional full-body generation by collecting a large-scale and high-quality dataset. To further constraint the appearance of the generated images, additional input signals were also considered, such as face [Frühstück et al. 2022; Wang et al. 2022], partial body [Frühstück et al. 2022; Kulal et al. 2023; Sarukkai et al. 2023; Wang et al. 2022], and full body [AlBahar et al. 2021; Bhunia et al. 2022; Karras et al. 2023; Knoche et al. 2020; Kurupathi et al. 2020; Ma et al. 2017; Men et al. 2020; Ren et al. 2022; Sanyal et al. 2021; Siarohin et al. 2019; Zhao and Zhang 2022]. For example, PIDM [Bhunia et al. 2022] used a diffusion-based approach to create a full-body image, taking into consideration a target pose and a full-body image that features the same outfit but in a different pose. [Frühstück et al. 2022] developed a set of specialized GANs for each body part to synthesize full-body images. [Yang et al. 2022] proposed Paint-By-Example, which was capable of inpainting a masked image given a full-body image as the source. However, all these papers assumed the conditional signals come from third-person view images, but not selfies as we desired.

2.2 Selfie-Related Techniques

Numerous studies have explored the use of selfies for various applications including reposing [Ma et al. 2020], face recognition [Botezatu et al. 2022; Kumarapu et al. 2023], style transfer [Li et al. 2021; Torbunov et al. 2023], novel view synthesis [Athar et al. 2023; Bian et al. 2021; Kania et al. 2022; Park et al. 2021a,b], relighting [Capece et al. 2019], and video Stabilization [Yu and Ramamoorthi 2018; Yu et al. 2021]. [Athar et al. 2023] combined neural radiance fields (NeRF) with a 3D morphable face model (3DMM) to create selfie images with diverse expressions and perspectives using a selfie video as input. But this model utilized the face model as prior and thus cannot be extended to the full-body scenario. [Park et al. 2021b] proposed HyperNeRF to model the topological change of the subject when taking a selfie video for novel view synthesis. However, this method cannot be adopted to our task since the background and expression cannot be changed as we desired, which is required in our task when the person is presented in a new location. In summary, none of these techniques are able to solve our task.

2.3 Diffusion Model

Diffusion models have recently been used for text-to-image task [Rassin et al. 2022; Rombach et al. 2022; Saharia et al. 2022] and achieved a great success by training on an extremely large dataset. However, these pretrained models are limited in their ability to generate personalized images of subjects in specific contexts. To tackle this issue, [Ruiz et al. 2022] introduced DreamBooth, a model designed to enable personalization in text-to-image tasks. It accomplished this by fine-tuning the text-to-image model using a small number of image-text pairs. We believe that pretrained text-to-image model has already learned the prior of selfie and full-body shots. Thus we build Total Selfie on top of DreamBooth to generate personalized full-body images.

3 TOTAL SELFIE

We define our task setting more formally here. First, a user wearing a specific outfit captures different parts of the user's body through a selfie video. From this video, we extract a series of selfie images. Specifically, we categorize the images of overhead view, cloth, pants, and shoes into the sets V_o , V_c , V_p , and V_s , respectively. Each of these sets contains 15 images (see Fig. 2 blue box for sample images). In any new site, we can generate a full-body image I_h of the individual, given an on-site selfie I_s (wearing the same outfit), a background image I_b , and an image I_t with the person in the desired target pose. Note that I_t may have the person in a different outfit, captured intentionally or just drawn from a personal photo collection.

Our approach relies on the implicit prior knowledge of Stable Diffusion to link the selfie images of different body parts, allowing for the reconstruction of a full view of an individual wearing a specific outfit. As depicted in Fig. 2, our method includes three main modules: (1) Region-Aware Generation trains a multi-concept version of DreamBooth [Ruiz et al. 2022] using V_o , V_c , V_p , and V_s . Then it composites the person into the background I_b with target pose I_t using ControlNet, producing an initial full-body image I_a . (2) Appearance Refinement corrects clothing inaccuracies in I_q and integrates the expression and look of on-site selfie I_s into I_q using local (automatic) refinement via DreamBooth and ControlNet. Furthermore, we introduce a selfie undistortion method to correct perspective distortion of the face present in I_s . The output of this stage is defined as I_r . (3) Image harmonization projects refined image Ir back to data manifold learned by Stable Diffusion to remove boundary artifacts and incorrect shading introduced during the refinement process. It also improves the quality of image details. The final output is defined as I_h .

3.1 Preliminaries

We start with a quick overview of Stable Diffusion [Rombach et al. 2022] and establish notation that we use throughout. Stable Diffusion contains two parts: (1) a variational autoencoder, which includes an encoder E that compresses an image I into a latent vector z = E(I) and a decoder D that reconstructs an image from a latent vector, *i.e.*, I' = D(z), and (2) a denoising U-Net, represented as $\epsilon_{\theta}(z_t, t, c)$, where z_t is the denoised latent vector at timestep t and *c* represents conditioning information, typically in the form of a text prompt embedding. Recently, Zhang and Agarwala [2023] proposed ControlNet to incorporate various control signals (e.g., Canny edges, human pose, etc.) without requiring retraining Stable Diffusion. Consequently, we expand *c* to be a set of control signals, extending beyond just the text prompt embedding. Note that the text prompt embedding is invariably a component of *c*, as Stable Diffusion operates as a text-to-image model. Besides, when c contains control signal(s) (e.g., Canny edge) other than text prompt embedding, we integrate the corresponding ControlNet model(s) (e.g., ControlNet Canny edge model) into $\epsilon_{\theta}(z_t, t, c)$. We retain this original notation for the sake of simplicity.

To generate an image, we sample a random noise vector and the U-Net progressively denoises it in T timesteps. At timestep t, the predicted noise is computed via classifier-free guidance [Ho and Salimans 2022]:

$$\epsilon_{\theta}(z_t, t, c, \emptyset) = \epsilon_{\theta}(z_t, t, \emptyset) + m \cdot (\epsilon_{\theta}(z_t, t, c) - \epsilon_{\theta}(z_t, t, \emptyset)), \quad (1)$$

where *m* is the guidance scale set to 7.5 in our paper. \emptyset is a copy of *c* but with the embedding of an empty text prompt. In other words, any input signal related to ControlNet in \emptyset (if exists) is the same as the one in *c*. Different schedulers [Ho et al. 2020; Liu et al. 2022; Song et al. 2020] can be used to obtain the cleaner (one-step denoised) vector z_{t-1} . In this paper, we opt for a commonly used scheduler

DDIM [Song et al. 2020], which first predicts denoised data point:

$$\hat{z}_0 = \frac{z_t - \sqrt{1 - \alpha_t} \epsilon_\theta(z_t, t, c, \emptyset)}{\sqrt{\alpha_t}},$$
(2)

where α_t is the noise scale. Then we can sample z_{t-1} based on \hat{z}_0 . Please see [Song et al. 2020] for details of α_t and the sampling of z_{t-1} . To ease the notation burden, we define the sampling process as $z_{t-1} = S(z_t, \epsilon_{\theta}(z_t, t, c, \emptyset), t)$.

3.2 Region-Aware Generation

In this section, we describe how we train a personalized Stable Diffusion model (*i.e.*, DreamBooth fine-tuned on images from the selfie video) to generate an initial full-body shot I_g against background I_b with a reasonable camera viewpoint. The advantages of using DreamBooth are two-fold: (1) the implicit prior knowledge in Stable Diffusion can help learn the correspondence among selfie images of different body parts to build a complete "understanding" of a person's outfit without explicit geometric information, and (2) Stable Diffusion can implicitly leverage the lighting conditions in the desired scene, leading to the generation of convincing shading when integrating the person into the scene.

Given a pre-captured selfie video, we first extract image sets V_o , V_c , V_p , and V_s using a pretrained human parsing model [Yang et al. 2023]. Then we train a multi-concept version of DreamBooth with V_o for "photo of [X] person, selfie"; V_c for "[Y] cloth"; V_p for "[V] pants"; and V_s for "[M] shoes". (We use "[Y] cloth" to refer to upper body clothing). Specifically, we estimate body poses in the training images with OpenPose [Cao et al. 2017] and use them as additional input to train the DreamBooth with a frozen ControlNet human pose model. Formally, we only fine-tune the U-Net in $\epsilon_{\theta}(z_t, t, \{p_v, I_v^P\})$, where I_v^p is the detected pose of a training image and p_v is the embedding of the corresponding text prompt. Please note, as we defined in Sec. 3.1, $\epsilon_{\theta}(z_t, t, \{p_v, I_v^p\})$ contain not only the U-Net but also the ControlNet human pose model. We observe that providing additional pose information helps DreamBooth handle different body parts more effectively.

One advantage of ControlNet is that its model is independent of the fine-tuned U-Net in DreamBooth. This means that we can flexibly combine any ControlNet model into the U-Net according to our needs. In our case, while $\epsilon_{\theta}(z_t, t, \{p_v, I_v^P\})$ is used for training, during the inference, I_v^p can be substituted with any other ControlNet input signal(s). To this end, we define the trained DreamBooth model more generally as $\epsilon_{\theta}(z_t, t, c)$.

To synthesize a full-body image, one naive approach is to condition the sampling at each timestep, $S(z_t, \epsilon_\theta(z_t, t, \{p_g\}, \emptyset), t)$, with just a text prompt embedding p_g , where the text prompt is set to "photo of [X] person standing, full body | wearing [Y] cloth | wearing [V] pants | wearing [M] shoes". However, as shown in Fig. 3 (b), this approach tends to generate images with (1) camera viewpoint is from above, similar to the overhead training images in V_o , and (2) the face is missing. And of course, the background is not I_b (not used to this point). To solve these issues, we develop a region-aware inference strategy to combine the background image I_b with the foreground person we intend to generate, with a more natural camera viewpoint. The target pose I_t provides the body shape of the

4 · Chen, Curless. et al



Fig. 2. Pipeline of Total Selfie. Given selfie video frames of different body parts (blue box), Region-Aware Generation (green box) trains a multi-concept DreamBooth to generate an initial full body image I_g in the background I_b with the target pose I_t . Appearance Refinement (orange box) refines face region of I_g by incorporating the expression from the on-site selfie I_s with perspective undistortion. In addition, other body parts (e.g. cloth) are also refined using a similar idea with slight modifications. The refined image is defined as I_r . Image Harmonization (purple box) harmonizes the refined image to improve unnatural regions using diffusion prior with appropriate guidance, generating final output I_h .

foreground, which is essential because it is challenging to infer total body shape solely from selfie images of body parts.

The idea of region-aware inference is to perform each step of the denoising process with ControlNet for foreground and background separately and then merge them to obtain a combined latent before taking the next denoising step. We adopt ControlNet's Canny edge model because it can guide the model to generate the person with pose provided by I_t . Specifically, we first detect Canny edge map I_t^c and foreground mask M_g from the *semantic map* of I_t , where the semantic map is obtained from the pretrained human parsing model [Yang et al. 2023], as illustrated in Fig. 2 ①. Empirically, we observe that dilating M_g by 10 pixels leads to better generation results for pixels around the detected Canny edge. By detecting I_t^c from the semantic map rather than directly from I_t , we ensure that the pipeline is not influenced by details of the random outfit in I_t (e.g., cloth texture). Then at timestep t, the foreground latent z_{t-1}^f is sampled using the fine-tuned DreamBooth model now combined

with the pre-trained ControlNet Canny edge model instead of the pose model:

$$z_{t-1}^f = S(z_t, \epsilon_\theta(z_t, t, \{p_g, I_t^c\}, \emptyset), t).$$
(3)

The background latent z_{t-1}^b can be obtained by adding noise to background image I_b using DDIM (see [Song et al. 2020] for details).

To ensure a natural transition between foreground and background, we propose to blend z_{t-1}^f and z_{t-1}^b . We compute the blending weight $W_g \in [0, 1]$ as:

$$W_g = s_w \frac{dst(M_g)}{\|dst(M_g)\|_{\infty}},\tag{4}$$

where s_w is the strength for W_g (see supplementary), and $dst(M_g)$ is the distance to M_g 's boundary inside the mask and 0 for pixels outside the mask. Finally, we calculate z_{t-1} by softly blending z_{t-1}^f

and z_{t-1}^b , given by:

$$z_{t-1} = \begin{cases} z_{t-1}^{f} \cdot W_g + z_{t-1}^{b} \cdot (1 - W_g), & \text{if } t \le s_g T \\ z_{t-1}^{b}, & \text{if } t > s_g T \end{cases},$$
(5)

where $s_g = 0.9$ controls the strength of the foreground generation – a larger s_g implies stronger foreground. The latent z_{t-1} can then be used to compute the next denoised foreground latent, and the iterative process continues to completion, followed by decoding the final latent to yield image I_g . This process is illustrated in Fig. 2 gray box. I_q is visualized in the top left of Fig. 2 orange box and Fig. 3 (c).

3.3 Appearance Refinement

As demonstrated in Fig. 3 (c), region-aware generation can create the full-body image with roughly correct outfit and reasonable shading. However, it does not preserve the identity or expression of the subject or the fine details of the clothes; i.e., the model is not able to fully bridge the gap in distribution between the captured selfie imagery and the full-body image we want. For example, the upper cloth in the training images typically occupies more pixels than the upper cloth in the full-body image we want to generate. Thus the goal of appearance refinement is two-fold: (1) refine the clothing to produce fine details, and (2) incorporate the expression and look of the on-site selfie into the initial full-body photo I_g . Our idea is to utilize local refinement to better represent different body parts in selfies using DreamBooth.

We start by describing the refinement of face region to incorporate the facial features and expressions from the on-site selfie I_s . Before describing local refinement, we first introduce the selfie undistortion module, designed specifically for face region refinement.

Selfie Undistortion: The selfie image I_s exhibits strong perspective distortion that would not be present in a full-body image taken from a distance. Existing methods reduce selfie distortion by either optimizing a single image [Shih et al. 2019; Wang et al. 2023] or training on a combination of a synthetic dataset, which is rendered with unrealistic texture and lighting, and a small real dataset [Zhao et al. 2019]. For test-time efficiency, we follow the idea of large dataset training, but instead of using an unrealistic dataset like previous work, we render a large paired dataset using a state-of-the-art method that generates realistic textured 3D heads (with backgrounds) using 3D GANs [Chan et al. 2021]. Then we fine-tune a talking-head synthesis network [Wang et al. 2021] to perform perspective undistortion using the rendered dataset. We define the undistorting process of I_s as:

$$I'_{s} = T(A(I_{s})), \tag{6}$$

where $A(\cdot)$ and $T(\cdot)$ are the face alignment operation and the trained undistortion network, respectively (see Fig. 2 ②). Please see supplementary for more details and results. The undistorted image I'_s is then used as the source content in the local refinement.

Local Refinement: The goal now is to transfer the appearance of the face in I'_s to I_g . Directly pasting I'_s into the face region of I_g does not work well because the head pose and lighting are different. Instead, we train a face-specific, "on-site DreamBooth" $\epsilon'_{\theta}(z_t, t, c)$ with prompt "[*Z*] face" (Fig. 2, step ③), fine-tuned on just one image, I'_s , which will then be used to generate the improved face combined

with I_g , through blended diffusion (similar to Sec. 3.2). As I'_s is itself a face crop, we find that similarly cropping and aligning I_g before diffusion blending yields best results. Specifically, we first crop $C(\cdot)$ and align $A(\cdot)$ image I_g , resulting in $A(C(I_g))$ that is closer to I'_s (see Fig. 2 ④ and ⑤). For simplicity, we refer to $A(C(I_q))$ as $AC(I_q)$.

Given $AC(I_g)$, we can now refine it using $\epsilon'_{\theta}(z_t, t, c)$ to obtain the refined image $EAC(I_g)$. The refinement operation (Fig. 2 ⑦) is similar to region-aware inference with small modifications. Specifically, at timestep *t*, we compute the latent from on-site DreamBooth as:

$$z_{t-1}^{r} = S(z_{t}, \epsilon_{\theta}'(z_{t}, t, \{p_{r}, AC(I_{t})^{c}, AC(I_{t})^{l}\}, \emptyset), t),$$
(7)

where p_r is the embedding of text prompt "[Z] face". $AC(I_t)^c$ is the canny edge detected from the semantic map of cropped and aligned target pose $AC(I_t)$. $AC(I_t)^l$ denotes pre-filtered landmarks and pose detected from $AC(I_t)$, as shown in Fig. 2 (6). The latent from $AC(I_g)$ can be obtained using DDIM, denoted as z_{t-1}^g . Finally, z_{t-1} is given by:

$$z_{t-1} = \begin{cases} z_{t-1}^r \cdot AC(M_r) + z_{t-1}^g \cdot (1 - AC(M_r)), & \text{if } t \le s_i T \\ z_{t-1}^g, & \text{if } t > s_i T \end{cases},$$
(8)

where M_r is the mask of face region in I_g , and $s_i = 0.8$ is the strength. After refinement, we map the refined face image $EAC(I_g)$ back into full I_g image coordinates and paste it in. All images are resized to Stable Diffusion's default 512x512 resolution.

Remaining Body Parts: We also apply the same idea to refine the cloth, pants, shoes, and hands with minor modifications. To ease the notation burden, we reuse M_r and p_r for these regions. For example, for cloth, M_r is the cloth mask and p_r is the embedding of prompt "[Y] cloth". In addition, for hands, we simply set p_r as the embedding of "hand". The modifications include: (1) In Eq. 7, we use DreamBooth $\epsilon_{\theta}(z_t, t, c)$ trained in Section 3.2 rather than on-site DreamBooth. In other words, the on-site selfie is not used. (2) In Eq. 7, we replace control signals $\{p_r, AC(I_t)^c, AC(I_t)^l\}$ with $\{p_r, C(I_t)^c\}$, where $C(I_t)^c$ is canny edge of semantic map of the cropped I_t based on M_r . (3) In Eq. 8, $AC(M_r)$ is replaced by $C(M_r)$. Also, $AC(I_q)$ is replaced by $C(I_q)$.

We perform this refinement for each part of the body and paste the results over I_g to yield a new composite image I_r , as visualized in the top left of the Fig. 2 purple box and Fig. 3 (d).

3.4 Image Harmonization

As highlighted in Fig. 3 (d), although the outfit and identity in I_r are mostly correct, it still has obvious artifacts: (1) The shading for the cloth, pants, and shoes is incorrect (highlighted by blue arrow), because the DreamBooth utilized in appearance refinement is trained from selfie video frames, taken in another place under different lighting conditions. (2) The boundary has artifacts between the foreground and background (highlighted by purple arrow) where appearance refinement struggles to preserve the structure of body parts while generating the correct outfit. (3) Some small details have artifacts; e.g., the hands in Fig. 3 (c) are unnatural despite being generated by DreamBooth; even with local refinement in Fig. 3(d), highlighted by the green arrow, the hand region still has noticeable artifacts.

6 · Chen, Curless. et al



(a) On-Site Selfie (Input)

(b) Naive Generation

eration (c) Region

(c) Region-Aware Generation (d) Appearance Ref + (c)

(e) Full Pipeline

Fig. 3. Results for different modules of our pipeline. (a) shows the on-site selfie. The image sample of selfie video is in Fig. 4 row 1. (b) shows the results of trained dreambooth that only uses text prompt as condition for inference. (c)-(e) show the zoom-in of the outputs. With only region-aware generation, output (c) has incorrect identity and clothing. With appearance refinement and region-aware generation, output (d) has better identity but contains boundary artifacts (purple arrow), incorrect shading (blue arrow), and bad image details (green arrow). In contrast, the full pipeline (e) produces a realistic and faithful full-body photo.

To address these issues, we propose and apply an image harmonization stage. The key idea is to reproject I_r back onto the data manifold learned by Stable Diffusion. In other words, Stable Diffusion should have a sense of what a natural, full-body shot looks like, and we aim to utilize this knowledge.

First, we fine-tune the pretrained Stable Diffusion decoder D on the image I_r using L1 and perceptual loss [Johnson et al. 2016]. This step is necessary, as small details, especially the face region, are otherwise not reconstructed well in I_r . Then we invert I_r to random noise in the latent space of $\epsilon_{\theta}(z_t, t, c)$ using null-text inversion [Mokady et al. 2022], given by:

$$\{\tilde{\varnothing}_t\}, \tilde{z}_T = inv(I_r) \tag{9}$$

where $\{\tilde{\emptyset}_t\}$ is a set of optimized null text embeddings (one for each timestep), and \tilde{z}_T is the output of DDIM inversion. With \tilde{z}_T as initialization, we sample \tilde{z}_{t-1} at each timestep to obtain \tilde{z}_0 , using the function $S(\tilde{z}_t, \epsilon_{\theta}(\tilde{z}_t, t, \{p_h\}, \tilde{\emptyset}_t), t)$, where p_h is the embedding of empty text prompt. We observe this performs better than using prompts related to "*person*", which overfit to our training image since we use prompts with "*person*" for DreamBooth training. Then we can reconstruct I_r using the fine-tuned decoder, given by $I_r \approx$ $D(\tilde{z}_0)$. Note that, we do not use any ControlNet model in the image harmonization stage.

This inversion technique allows us to project the generation of I_r back into the Stable Diffusion process. Our goal then is to move the denoised image towards the the learned data manifold, all the while maintaining the primary content in I_r . Following [Bansal et al. 2023], we adopt a similar idea of classifier guidance [Dhariwal and Nichol 2021] for the denoising process without any test-time fine-tuning. Specifically, at a specific timestep t, we apply forward guidance on the predicted noise to enforce the denoising process to our desired direction, given by:

$$\hat{\epsilon}_{\theta}(z_t, t, \{p_h\}, \tilde{\varnothing}_t) = \epsilon_{\theta}(z_t, t, \{p_h\}, \tilde{\varnothing}_t) + \nabla_{z_t} L, \qquad (10)$$

ACM Trans. Graph., Vol. 1, No. 1, Article . Publication date: August 2023.

where *L* is the guidance loss. With this predicted noise as guidance, we can compute z_{t-1} , again with the blending of latent, by:

$$z_{t-1} = \begin{cases} S(z_t, \hat{e}_{\theta}(z_t, t, \{p_h\}, \tilde{\varnothing}_t), t) \cdot M_h + \tilde{z}_{t-1} \cdot (1 - M_h) & \text{if } t \le s_h T \\ \tilde{z}_{t-1} & \text{if } t > s_h T \end{cases}$$
(11)

where $s_h = 0.2$ is to control the strength of forward guidance. $M_h = M_g - M_f$, and M_f is mask of face region of I_r . In other words, we keep the face region unchanged during the whole process since it is generated from on-site selfie I_s with desired lighting conditions. Then, instead of proceeding to the next timestep t - 1, we recompute z_t with combination of z_{t-1} and a random gaussian noise, given by:

$$z_t = \sqrt{\alpha_t / \alpha_{t-1}} z_{t-1} + \sqrt{1 - \alpha_t / \alpha_{t-1}} \mathcal{N}(0, 1).$$
(12)

This injection of Gaussian noise enables the denoise process to move toward the generation of more realistic images. We repeat this recomputation N times (we set N = 3), where a larger N will push the image closer to the data manifold, but farther from the content of I_r .

We construct *L* in Eq. 10 from two losses: (1) perceptual loss [Johnson et al. 2016] L_p to preserve the content in I_r , and (2) style loss [Gatys et al. 2015] L_s to help match the appearance (e.g., shading) in I_s . To implement this, at timestep *t*, we compute the predicted denoised latent \hat{z}_0 using Eq. 2 by replacing \emptyset with $\tilde{\emptyset}_t$ and *c* with $\{p_h\}$. Then the predicted clean image can be computed by $\hat{I} = D(\hat{z}_0)$. Finally, we compute $L = w_p \cdot L_p(I_r, \hat{I}) + w_s \sum_{\omega \in \Omega} \cdot L_s(C_{\omega}(I_s), C_{\omega}(\hat{I}))$, where $w_p = 100000$ and $w_s = 10000$ are weights for perceptual and style loss respectively. Ω is a set of body parts presented in both I_s and \hat{I} , excluding *face* as its shading is already well-modeled. As an example, for I_s in Fig. 2, we set $\Omega = \{cloth\}$, corresponding to the shirt region, as no other non-face regions are visible. $C_{\omega}(I_s)$ and $C_{\omega}(\hat{I})$ are images cropped from I_s and \hat{I} based on their masks for body part ω .

In summary, the recomputation of z_t with Gaussian noise encourages the final image to be realistic, perceptual loss preserves the content, and style loss corrects the shading. The final result I_h is shown in Fig.3 (e) and Fig. 2 purple box. Please see the implementation details of the pipeline in the supplementary.



(a) Overhead (b) Cloth (c) Pants (d) Shoes

Fig. 4. Sample images of different body parts of different users extracted from their selfie videos. The appearance of the same outfit can vary across different selfies, depending largely on factors such as spatially variable lighting conditions and diverse camera settings. For instance, when comparing the black top in row 1 (a) to row 1 (b), it is noticeable that the black top appears somewhat lighter in the latter image.

4 EXPERIMENTS

Evaluation data. We collected a dataset of five people, each of them has one selfie video with three different on-site selfies in two different scenes. This resulted in a total of 30 examples. The target poses were provided from their photo collections. Fig. 4 illustrates the representative images that were derived from these videos. For evaluation, we also collected ground truth full-body photos of each person, i.e., with the same facial expression, clothing, and back-ground.

Results. Fig. 9 demonstrates that Total Selfie can produce full-body shots in diverse backgrounds with reasonable shading, despite of potentially different lighting conditions in pre-captured selfies. To illustrate this point, consider the black top in Fig. 4 (b) row 1, which under a certain light appears somewhat lighter. Our pipeline can effectively discern and resolve this ambiguity by exploiting the clothing appearance from the on-site selfie (explicitly through style loss) and overhead selfies (implicitly through DreamBooth) in Fig. 4 (a) row 1. Furthermore, Total Selfie has the capability to handle complex facial expressions, including wide-open mouths, visible teeth, and even winks, as shown in Fig. 9. See more results in the supplementary.

Method	LPIPS \downarrow	SSIM ↑	PSNR ↑	FID \downarrow
PIDM	0.415	0.472	11.12	350.0
Paint-By-Example	0.367	0.466	12.62	222.5
DreamBooth+ControlNet	0.261	0.614	18.27	149.9
Ours-IH-AR	0.254	0.614	18.47	153.8
Ours-IH	0.214	0.631	18.95	116.1
Ours-SU	0.208	0.650	19.26	116.5
Ours-Style	0.208	0.649	19.26	118.6
Ours-Lpips	0.217	0.643	19.23	132.8
Ours	0.205	0.652	19.28	112.2

Table 1. Comparison with baselines and Total Selfie variants. The metrics are evaluated only in the foreground. To align the foreground in output and ground truth, we use ground truth as the target pose for all tested methods and variants.

Ablation Study. We perform two studies. First, we study the effect of different parts of the pipeline and design three variants: (1) *Ours-IH-AR*: full pipeline without image harmonization and appearance refinement (*i.e.*, region-aware generation only), (2) *Ours-IH*: full pipeline without image harmonization, (3) *Ours-SU*: full pipeline without selfie undistortion. As discussed in Sec. 3, Fig. 3 and Table. 1 show that the Total Selfie performs better than these three variants.

Second, we study the effects of guidance loss in the image harmonization stage by testing two variants: (1) *Ours-Style*: full pipeline without style loss. (2) *Ours-Lpips*: full pipeline without perceptual loss. Table 1 indicates that *Ours-Lpips* performs the worst, as the content is not well preserved without perceptual loss. The full pipeline outperforms *Ours-Style*, demonstrating that the style loss can effectively correct shading with the assistance of on-site selfies.

See supplementary for more results of the ablation study. Comparison to Baselines. To the best of our knowledge, there are no existing papers solving the same tasks. Therefore, we modify three existing methods to adapt to our task: (1) **PIDM** [Bhunia et al. 2022] is originally designed to generate full-body photo given a source image and a target pose. We modify this method by using a selected overhead selfie (from the selfie video) as the source image and I_t as the target pose to generate the initial output. The face region and other body parts of this initial result are further refined sequentially using the on-site selfie and selected body selfies (from the selfie video) as source images. (2) Paint-By-Example [Yang et al. 2022] is originally developed to inpaint a masked image using content from a source image. We develop a two-stage process for our adaptation. The initial output is generated using a selected overhead selfie as the source image and a scene image masked by M_a as the masked image. The face region of this initial output is then further inpainted using the on-site selfie and face mask. Note that, we do not apply this for other body parts or combine the inpainting model with ControlNet as they lead to degraded results. (3) DreamBooth+ControlNet [Ruiz et al. 2022; Zhang and Agrawala 2023] is implemented the same as our region-aware generation. The difference is that both pre-captured and on-site inputs are also used for DreamBooth training.

8 • Chen, Curless. et al



Fig. 5. Qualitative Comparison with baselines. Sample pre-captured images (from selfie video) are shown in the second and last column in Fig. 4. All results are zoomed in for clear visualization. (d) are the results of DreamBooth+ControlNet. In all methods, we use the ground-truth as target pose to constrain the pose. Our pipeline clearly outperforms all baselines in terms of photo realism and faithfulness. Note that, despite being captured nearly at the same time, the color tone of the on-site selfie, background image, and ground-truth may not match due to differences in lighting conditions, auto exposure, and white balance *etc*.



(a) Total Selfie for Pose A





of flexibility and diversity in the generated full-body photographs.

Table 1 and Fig. 5 show the quantitative and qualitative results of all methods, respectively. PIDM underperforms since it is trained on a dataset of fashion images, characterized by a third-person view without background. DreamBooth+ControlNet surpasses Paint-By-Example due to its use of all available input images for model finetuning. In contrast, Paint-By-Example does not perform any testtime optimization on input images and can only consider one source image per forward pass. Total Selfie performs the best by flexibly utilizing DreamBooth and ControlNet throughout the whole pipeline. **Results on Natural Pose** Total Selfie can also generate full-body images with more natural poses, while still preserving the patterns on the clothes, as shown in Fig. 6. These poses can considerably differ from those utilized during the pre-capturing of the selfie video.



(a) Total Selfie

(b) Ground-Truth

Fig. 7. Limitations of Total Selfie, where the output has incorrect shading and unrealistic hands.

This flexibility empowers users to choose their preferred poses and thereby generate a wide variety of full-body photos.

5 LIMITATIONS AND FUTURE WORK

Total Selfie has several limitations: (1) The shading in the generated full-body image may not align accurately with the actual photo (see Fig. 7). This happens when the shading in the initial full-body image (generated by DreamBooth) greatly differs from the shading in the on-site selfie. A potential avenue for future exploration could involve harnessing the on-site selfie to guide the region-aware generation. (2) Clothing and hair must be similar in shape between target pose I_t and on-site selfie I_s . For example, wearing pants in I_t but a dress in I_s leads to the failure of our method. For future work, we plan to develop a method capable of addressing the disparity of



(a) Ours-SU

(b) Ours

Fig. 8. Comparison of the pipeline with and without selfie undistortion. With perspective undistortion, the pipeline produces a more natural-looking face.

clothes and hairstyles between the two images. Moreover, generating feasible poses based on a given background presents an exciting prospect for further exploration. (3) The hands and arms region in the generated image is not realistic enough (see Fig. 7), which is a well-known limitation for Stable Diffusion. Fortunately, our pipeline has the flexibility to integrate any pretrained text-to-image model as a plugin, allowing us to enhance its output whenever more advanced models become available.

6 CONCLUSIONS

In this paper, we introduce Total Selfie, a novel approach to generating full-body images from selfies. The method assumes the availability of a target pose, along with a pre-captured selfie video showcasing a new outfit. Then given an on-site selfie featuring the same outfit and a site photo, Total Selfie produces a full-body photo of the person in the given site with the target pose, capturing the facial expression present in the on-site selfie. We address challenges arising from significant disparities between selfies and full-body photos, including substantial differences in camera viewpoints and fields of view, the use of at least one hand to capture the selfie, and the absence of paired selfie and full-body training data for our task. We present results from five individuals who experimented with our approach, as well as ablation studies and comparisons with other potential methods.



(a) Target Pose

(b) On-Site Selfie

(c) Background

(d) Ours (Zoom In)

Fig. 9. Results of Total Selfie. The sample pre-captured images (from the selfie video) are shown in Fig. 1 and the first three rows of Fig. 4. Total Selfie successfully generates authentic and realistic full-body images for diverse individuals, capturing a broad spectrum of expressions set against a variety of backgrounds, while preserving the clothes and producing reasonable shading. Please note, the second row's output has a stripe appearing on the left side of the pants, which is not an artifact but rather a ground pattern in the background.

REFERENCES

- Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. 2021. Pose with Style: Detail-preserving pose-guided image synthesis with conditional stylegan. ACM Transactions on Graphics (TOG) 40, 6 (2021), 1–11.
- ShahRukh Athar, Zhixin Shu, and Dimitris Samaras. 2023. Flame-in-nerf: Neural control of radiance fields for free view face animation. In 2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG). IEEE, 1–8.
- Arpit Bansal, Hong-Min Chu, Avi Schwarzschild, Soumyadip Sengupta, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2023. Universal Guidance for Diffusion Models. arXiv preprint arXiv:2302.07121 (2023).
- Ankan Kumar Bhunia, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Jorma Laaksonen, Mubarak Shah, and Fahad Shahbaz Khan. 2022. Person Image Synthesis via Denoising Diffusion Model. arXiv preprint arXiv:2211.12500 (2022).
- Jia-Wang Bian, Huangying Zhan, and Ian Reid. 2021. NVSS: High-quality Novel View Selfie Synthesis. In 2021 International Conference on 3D Vision (3DV). IEEE, 1085– 1094.
- Cristian Botezatu, Mathias Ibsen, Christian Rathgeb, and Christoph Busch. 2022. Fun selfie filters in face recognition: Impact assessment and removal. *IEEE Transactions* on Biometrics, Behavior, and Identity Science 5, 1 (2022), 91–104.
- Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2017. Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- Nicola Capece, Francesco Banterle, Paolo Cignoni, Fabio Ganovelli, Roberto Scopigno, and Ugo Erra. 2019. Deepflash: Turning a flash selfie into a studio portrait. Signal Processing: Image Communication 77 (2019), 28–39.
- Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. 2021. Efficient Geometry-aware 3D Generative Adversarial Networks. In arXiv.
- Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. 2023. Diffusion models in vision: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023).
- Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. Advances in Neural Information Processing Systems 34 (2021), 8780–8794.
- Anna Frühstück, Krishna Kumar Singh, Eli Shechtman, Niloy J Mitra, Peter Wonka, and Jingwan Lu. 2022. Insetgan for full-body image generation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 7723–7732.
- Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. 2022. Stylegan-human: A data-centric odyssey of human generation. In Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XVI. Springer, 1–19.
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2015. A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576 (2015).
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2020. Generative adversarial networks. *Commun. ACM* 63, 11 (2020), 139–144.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems 33 (2020), 6840–6851.
- Jonathan Ho and Tim Salimans. 2022. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022).
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14. Springer, 694–711.
- Kacper Kania, Kwang Moo Yi, Marek Kowalski, Tomasz Trzciński, and Andrea Tagliasacchi. 2022. Conerf: Controllable neural radiance fields. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 18623–18632.
- Johanna Karras, Aleksander Holynski, Ting-Chun Wang, and Ira Kemelmacher-Shlizerman. 2023. DreamPose: Fashion Image-to-Video Synthesis via Stable Diffusion. arXiv preprint arXiv:2304.06025 (2023).
- Markus Knoche, István Sárándi, and Bastian Leibe. 2020. Reposing humans by warping 3d features. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. 1044–1045.
- Sumith Kulal, Tim Brooks, Alex Aiken, Jiajun Wu, Jimei Yang, Jingwan Lu, Alexei A. Efros, and Krishna Kumar Singh. 2023. Putting People in Their Place: Affordance-Aware Human Insertion into Scenes.
- Laxman Kumarapu, Shiv Ram Dubey, Snehasis Mukherjee, Parkhi Mohan, Sree Pragna Vinnakoti, and Subhash Karthikeya. 2023. WSD: Wild Selfie Dataset for Face Recognition in Selfie Images. arXiv preprint arXiv:2302.07245 (2023).
- Sheela Raju Kurupathi, Pramod Murthy, and Didier Stricker. 2020. Generation of Human Images with Clothing using Advanced Conditional Generative Adversarial Networks. *DeLTA* 3 (2020).
- Yuanming Li, Youngsaeng Jin, Jeonggi Kwak, Dongsik Yoon, David Han, and Hanseok Ko. 2021. Adaptive Content Feature Enhancement GAN for Multimodal Selfie to Anime Translation. (2021).

- Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. 2022. Pseudo numerical methods for diffusion models on manifolds. arXiv preprint arXiv:2202.09778 (2022).
- Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. 2017. Pose guided person image generation. Advances in neural information processing systems 30 (2017).
- Liqian Ma, Zhe Lin, Connelly Barnes, Alexei A Efros, and Jingwan Lu. 2020. Unselfie: Translating selfies to neutral-pose portraits in the wild. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16. Springer, 156–173.
- Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. 2020. Controllable person image synthesis with attribute-decomposed gan. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 5084–5093.
- Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2022. Nulltext Inversion for Editing Real Images using Guided Diffusion Models. arXiv preprint arXiv:2211.09794 (2022).
- Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. 2021a. Nerfies: Deformable neural radiance fields. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 5865–5874.
- Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. 2021b. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. arXiv preprint arXiv:2106.13228 (2021).
- Royi Rassin, Shauli Ravfogel, and Yoav Goldberg. 2022. Dalle-2 is seeing double: flaws in word-to-concept mapping in Text2Image models. arXiv preprint arXiv:2210.10606 (2022).
- Yurui Ren, Xiaoqing Fan, Ge Li, Shan Liu, and Thomas H Li. 2022. Neural texture extraction and distribution for controllable person image synthesis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 13535–13544.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). 10684–10695.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2022. Dreambooth: Fine tuning text-to-image diffusion models for subjectdriven generation. arXiv preprint arXiv:2208.12242 (2022).
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems 35 (2022), 36479–36494.
- Soubhik Sanyal, Alex Vorobiov, Timo Bolkart, Matthew Loper, Betty Mohler, Larry S Davis, Javier Romero, and Michael J Black. 2021. Learning realistic human reposing using cyclic self-supervision with 3d shape, pose, and appearance consistency. In Proceedings of the IEEE/CVF International Conference on Computer Vision. 11138– 11147.
- Vishnu Sarukkai, Linden Li, Arden Ma, Christopher Ré, and Kayvon Fatahalian. 2023. Collage Diffusion. arXiv preprint arXiv:2303.00262 (2023).
- YiChang Shih, Wei-Sheng Lai, and Chia-Kai Liang. 2019. Distortion-free wide-angle portraits on camera phones. ACM Transactions on Graphics (TOG) 38, 4 (2019), 1–12.
- Aliaksandr Siarohin, Stéphane Lathuilière, Enver Sangineto, and Nicu Sebe. 2019. Appearance and pose-conditioned human image generation using deformable gans. IEEE transactions on pattern analysis and machine intelligence 43, 4 (2019), 1156–1171.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020).
- Dmitrii Torbunov, Yi Huang, Haiwang Yu, Jin Huang, Shinjae Yoo, Meifeng Lin, Brett Viren, and Yihui Ren. 2023. Uvcgan: Unet vision transformer cycle-consistent gan for unpaired image-to-image translation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. 702–712.
- Qinghe Wang, Lijie Liu, Miao Hua, Qian He, Pengfei Zhu, Bing Cao, and Qinghua Hu. 2022. HS-Diffusion: Learning a Semantic-Guided Diffusion Model for Head Swapping. arXiv preprint arXiv:2212.06458 (2022).
- Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-Shot Free-View Neural Talking-Head Synthesis for Video Conferencing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.
- Zhixiang Wang, Yu-Lun Liu, Jia-Bin Huang, Shin'ichi Satoh, Sizhuo Ma, Guru Krishnan, and Jian Wang. 2023. DisCO: Portrait Distortion Correction with Perspective-Aware 3D GANs. arXiv preprint arXiv:2302.12253 (2023).
- Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and Fang Wen. 2022. Paint by Example: Exemplar-based Image Editing with Diffusion Models. arXiv preprint arXiv:2211.13227 (2022).
- Lu Yang, Wenhe Jia, Shan Li, and Qing Song. 2023. Deep Learning Technique for Human Parsing: A Survey and Outlook. *arXiv preprint arXiv:2301.00394* (2023).
- Jiyang Yu and Ravi Ramamoorthi. 2018. Selfie video stabilization. In Proceedings of the European Conference on Computer Vision (ECCV). 551-566.
- Jiyang Yu, Ravi Ramamoorthi, Keli Cheng, Michel Sarkis, and Ning Bi. 2021. Real-time selfie video stabilization. In Proceedings of the IEEE/CVF Conference on Computer

Vision and Pattern Recognition. 12036–12044.

Lvmin Zhang and Maneesh Agrawala. 2023. Adding conditional control to text-to-image

diffusion models. *arXiv preprint arXiv:2302.05543* (2023). Jian Zhao and Hui Zhang. 2022. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

3657-3666.

Yajie Zhao, Zeng Huang, Tianye Li, Weikai Chen, Chloe LeGendre, Xinglei Ren, Ari Shapiro, and Hao Li. 2019. Learning perspective undistortion of portraits. In Pro-ceedings of the IEEE/CVF International Conference on Computer Vision. 7849–7859.