

InsetGAN for Full-Body Image Generation

Anna Frühstück^{1,2} Krishna Kumar Singh² Eli Shechtman²
Niloy J. Mitra^{2,3} Peter Wonka¹ Jingwan Lu²
¹ KAUST ² Adobe Research ³ University College London

anna.fruehstueck@kaust.edu.sa, {krishsin, elishe}@adobe.com, {niloym, pwonka}@gmail.com, jlu@adobe.com



Figure 1. **InsetGAN applications.** Our full-body human generator is able to generate reasonable bodies at state-of-the-art resolution (1024×1024 px) (a). However, some artifacts appear in the synthesized results, most visibly in extremities and faces. We make use of a second, specialized generator to seamlessly improve the face region (b). We can also use a given face as an input for unconditional generation of bodies (c). Furthermore, we can select both specific faces and bodies and compose them in a seamlessly merged output (d).

Abstract

While GANs can produce photo-realistic images in ideal conditions for certain domains, the generation of full-body human images remains difficult due to the diversity of identities, hairstyles, clothing, and the variance in pose. Instead of modeling this complex domain with a single GAN, we propose a novel method to combine multiple pretrained GANs, where one GAN generates a global canvas (e.g., human body) and a set of specialized GANs, or insets, focus on different parts (e.g., faces, shoes) that can be seamlessly inserted onto the global canvas. We model the problem as jointly exploring the respective latent spaces such that the generated images can be combined, by inserting the parts from the specialized generators onto the global canvas, without introducing seams. We demonstrate the setup by combining a full body GAN with a dedicated high-quality face GAN to produce plausible-looking humans. We evaluate our results with quantitative metrics and user studies.

1. Introduction

Generative adversarial networks (GANs) have emerged as a very successful image generation paradigm. For example, StyleGAN [14] is now the method of choice for creating

near photorealistic images for multiple classes (e.g., human faces, cars, landscapes). However, for classes that exhibit complex variations, creating very high quality results becomes harder. For example, full-body human generation still remains an open challenge, given the high variability of human pose, shape, and appearance.

How can we generate results at both high resolution and high quality? One approach is to break the target image into tiles and train a GAN to sequentially produce them [7]. Such methods, however, are unsuited for cases where the coupling between the (object) parts are nonlocal and/or cannot easily be statistically modeled. An alternate approach is to aim for collecting very high resolution images and train a single GAN, at full resolution. However, this makes the data collection and training tasks very expensive, and variations in object configuration/poses cause further challenges. To the best of our knowledge, neither such a high resolution dataset, nor a corresponding high resolution GAN architecture has been published.

We propose InsetGAN towards solving the above problems. Specifically, we propose to combine a generator to provide the global context in the form of a canvas, and a set of specialized part generators that provide details for different regions of interest. The specialized results are then

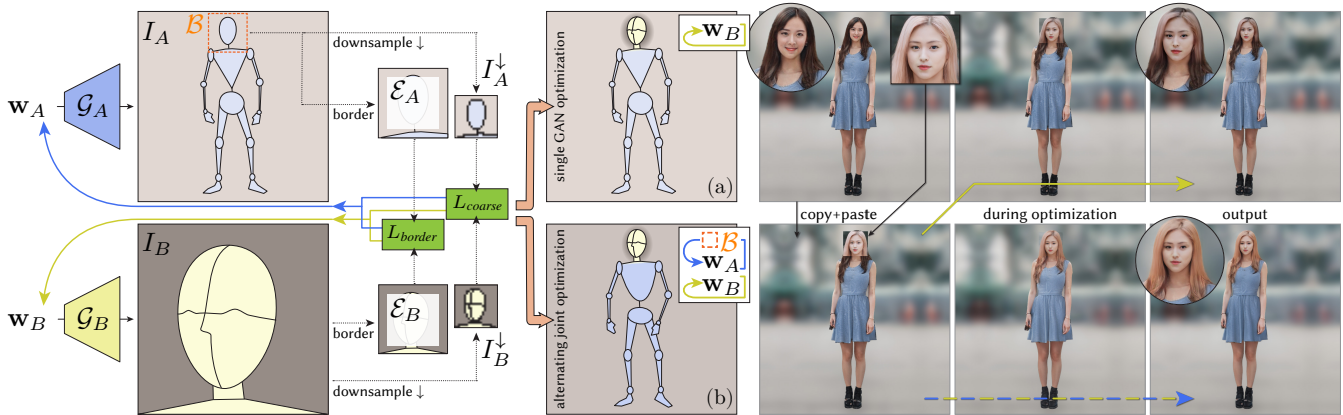


Figure 2. **InsetGAN Pipeline.** Given two latents w_A and w_B , along with pretrained generators \mathcal{G}_A and \mathcal{G}_B , that generate two images $I_A := \mathcal{G}_A(w_A)$ and $I_B := \mathcal{G}_B(w_B)$, respectively, we design a pipeline that can optimize either only w_A (a), or iteratively optimize both w_A and w_B (b) in order to achieve a seamless output composition of face and body. We use a set of losses L_{coarse} and L_{border} to describe the conditions we want to minimize during optimization. On the right, we show that given an input body, mere copy and pasting of a target face yields boundary artifacts. We show an application of one-way optimization (*top right*) and two-way optimization (*bottom right*) to create a seamlessly merged result. Note that when the algorithm can optimize in both inset-face and canvas-body generator spaces, it produces more natural results at the seam boundary – notice how the hair and skin tone blend from the head to the body region. The joint optimization is challenging as the bounding box $\mathcal{B}(I_A)$ is conditioned on the variable w_A .

pasted, as insets, on to the canvas to produce a final generation. Such an approach has multiple advantages: **(I)** the canvas GAN can be trained on medium quality data, where the object parts are not necessarily aligned. Although this results in the individual parts in the canvas being somewhat blurry (e.g., fuzzy/distorted faces in case of human bodies), this is sufficient to provide global coordination for later specialized parts to be inserted; **(II)** the specialized parts can be trained on part-specific data, where consistent alignment can be more easily achieved; and **(III)** different canvas/part GANs can be trained at different resolutions, thus lowering the data (quality) requirements. CollageGAN [20] has explored a similar idea in a conditional setting. Given a semantic map which provides useful shape and alignment hints, they create a collage using an ensemble of outputs from class-specific GANs [20]. In contrast, our work focuses on the unconditional setting, which is more challenging since our multiple generators need to collaborate with one another to generate a coherent shape and appearance together without access to a semantic map for hints.

The remaining problem is how to coordinate the canvas and the part GANs, such that adding the insets to the canvas does not reveal seam artifacts at the inset boundaries. This aspect is particularly challenging when boundary conditions are nontrivial and the inset boundaries themselves are unknown. For example, a face, when added to the body, should have consistent skin tone, clothing boundaries, and hair flow. We solve the problem by jointly seeking latent codes in (pretrained) canvas and part GANs such that the final image, formed by inserting the part insets on the canvas, does not exhibit any seams. In this paper, we investigate this problem in the context of human body generation, where the

human faces are created by a face-specific GAN.

We evaluate InsetGAN on a custom dataset, compare with alternative approaches, and evaluate the quality of the results with quantitative metrics and user studies. Fig. 1 shows human body generation applications highlighting both seamless results, across face insets, as well as having diversity of solutions across face insertion boundaries.

Contributions. **(1)** We propose a multi-GAN optimization framework that jointly optimizes the latent codes of two or more collaborative generators such that the overall composed result is coherent and free of boundary artifacts when the generated parts are inserted as insets into the generated canvas. **(2)** We demonstrate our framework on the highly challenging full-body human generation task and propose the first viable pipeline to generate plausible-looking humans unconditionally at 1024×1024 px resolution.

2. Related Work

Unconditional Image Generation via Generative Adversarial Networks (GANs) [8] has shown a lot of promise in recent years. In this context, the StyleGAN architecture was developed over a sequence of papers [11–14] and is widely considered the state of the art for synthesizing individual object classes. For class-conditional image generation on the ImageNet dataset, BigGAN [5] is often the architecture of choice. In our work, we are building on StyleGAN2-ADA, since this architecture yields better FID [9] and Precision&Recall [17] scores on our domain when compared to StyleGAN3. In addition, generating complete human body images using StyleGAN2 is a baseline we would like to improve upon in our work.



Figure 3. **Unconditional generation results.** Examples are created using our adaptive truncation approach (described in the supplementary material) and are cropped horizontally. At first glance, the results look realistic, but face regions show noticeable artifacts (please zoom).

Image Outpainting refers to image completion problems where the missing pixels are not surrounded by available pixels. Recent papers build on the ideas to use generative adversarial networks [29] and the explicit modeling of structure [19, 34]. Though these two papers specialize in human bodies, we find that the GAN architecture CoModGAN [33] has even more impressive results for image outpainting (see the comparison in Section 5).

Conditional Generation of Full-Body Humans has two possible advantages. First, the conditional generation enables more control. Second, conditional generation can help in controlling the variability and improve the visual quality. In the context of humans, a natural idea is to condition the generation on the human pose [4, 15, 16, 22, 23, 26, 28] or segmentation information [10].

As many conditional architectures are not able to handle the same high resolution (1024×1024 px) of unconditional StyleGAN, an alternative to developing new architectures is conditional embedding into an unconditional generator’s latent space. Two approaches used in this context are StyleGAN embedding using optimization [1, 2] or StyleGAN embedding using an encoder architecture [3, 25, 30]. Our work also makes use of embedding algorithms.

3. Methodology

We propose a method for the unconditional generation of full-body human images using one or more independent pretrained unconditional generator networks. Depending on the desired application and output configuration, we describe different ways to coordinate the multiple generators.

3.1. Full-Body GAN

The naive approach to generate a full-body human image is to use a single generator trained on tens of thousands of example humans (see Section 4 about the dataset). We adopt the state-of-the-art StyleGAN2 architecture proposed by Karras et al. [11]. Most previous full-body generation or editing work [4, 18, 20, 35] generate images at 256×256 px or 512×512 px resolution. We made the first attempt to unconditionally generate full-body humans at 1024×1024 px resolution. Due to the complex nature of our target domain, the results generated by a single GAN sometimes exhibit artifacts such as weirdly-shaped body parts and non-

photorealistic appearance. These artifacts are most visible in faces and extremities, as shown in Fig. 1(a). Due to the vast diversity of human poses and appearances and the associated alignment difficulty, hands and feet appear in many possible locations in the training images, making them harder for a single generator to learn. Faces are especially hard since we humans are ultra-sensitive to artifacts in these areas. They therefore deserve dedicated networks and special treatment. Fig. 3 shows a variety of unconditional generation results. Our results exhibit correct human body proportions, consistent skin tones across face and body, interesting garment variations and plausible-looking accessories (e.g. handbags and sunglasses) whereas artifacts can be present when viewed in detail.

3.2. Multi-GAN Optimization

To improve the problematic regions generated by the full-body GAN, we use other generators trained on images of specific body regions to generate pixels to be pasted, as insets, into the full-body GAN result. The base full-body GAN and the dedicated body part GANs can be trained using the same or different datasets. In either case, the additional network capacity contained in the multiple GANs can better model the complex appearance and variability of the human bodies. As a proof of concept, we show that a face GAN trained with the face regions cropped from our full-body training images can be used to improve the appearance of the body GAN results. Alternatively, we can also leverage a face generator trained on other datasets such as FFHQ [14] for face enhancement as well. Similarly, specialized hands or feet generators can also be used in our framework to improve other regions of the body. We show that we can also use multiple part generators together in a multi-optimization process, as depicted in Fig. 4.

The main challenge is how to coordinate multiple unconditional GANs to produce pixels that are coherent with one another. In our application, we have a \mathcal{G}_A that generates the full-body human where $I_A := \mathcal{G}_A(\mathbf{w}_A)$ and another \mathcal{G}_B that generates a sub-region or inset within the human body where $I_B := \mathcal{G}_B(\mathbf{w}_B)$. In order to coordinate the specialized part GAN with the global/canvas GAN, we need a bounding box detector to identify the region of I_A that corresponds to the region our part GAN generates. We crop I_A with the detected bounding box and denote the cropped



Figure 4. **Two Insets.** These results are improved using a dedicated shoe generator trained on shoe crops from our full-body humans, and also using our face generator. All three generators (full-body canvas and two insets) are jointly optimized to produce a seamless coherent output. The circular closeups show the shoes before (*top*) and after (*bottom*) improvement (please zoom).

pixels as $\mathcal{B}(I_A)$. The problem of inserting a separately-generated part I_B into the canvas I_A is equivalent to finding a latent code pair $(\mathbf{w}_A, \mathbf{w}_B)$ such that the respective images I_A and I_B can be combined without noticeable seams in the boundary regions of $\mathcal{B}(I_A)$ and I_B . To generate the final result, we directly replace the original pixels inside the bounding box $\mathcal{B}(I_A)$ with the generated pixels from I_B ,

$$\min_{\mathbf{w}_A, \mathbf{w}_B} \int_{\Omega} \mathcal{L}(\mathcal{G}_A(\mathbf{w}_A), \mathcal{G}_B(\mathbf{w}_B)) \quad (1)$$

where, $\Omega := \mathcal{B}(\mathcal{G}_A(\mathbf{w}_A))$ and, with slight abuse of notation, \mathcal{L} captures the loss both along the boundary of Ω measuring seam quality and inside the region Ω measuring similarity of I_A and I_B inside the respective faces. The full optimization is complex as the region of interest Ω depends on \mathbf{w}_A .

Our multi-GAN optimization framework can support various human generation and editing applications. Depending on the application scenario, we optimize either \mathbf{w}_A or \mathbf{w}_B or jointly optimize both for the best results.

Optimization Objectives. When optimizing the latent codes \mathbf{w}_A , \mathbf{w}_B or both, we consider multiple objectives: **(I)** the face regions generated by the face GAN and body GAN should have similar appearance at a coarse scale so that when the pixels generated by the face GAN are pasted onto the body GAN canvas, attributes match (e.g., the skin tone of the face matches that of the neck); **(II)** the boundary pixels around the face crops match up so that a simple copy-and-paste operation does not result in visible seams; and **(III)** the final composed result looks realistic. To match the face appearance, we downsample the face regions and calculate a combination of $L1$ and perceptual loss [32] L_{lips} :

$$L_{coarse} := \lambda_1 L_1(I_A^\downarrow, I_B^\downarrow) + \lambda_2 L_{lips}(I_A^\downarrow, I_B^\downarrow), \quad (2)$$

where $I_A^\downarrow = \mathcal{D}_{64}(\mathcal{B}(I_A))$ and $I_B^\downarrow = \mathcal{D}_{64}(I_B)$ and \mathcal{D}_{64} refers to downsampling the image to 64×64 px resolution.

For the boundary matching, we also apply a $L1$ and perceptual loss to the border pixels at full resolution:

$$L_{border} := \lambda_3 L_1(\mathcal{E}_8(\mathcal{B}(I_A)), \mathcal{E}_8(I_B)) + \lambda_4 L_{lips}(\mathcal{E}_8(\mathcal{B}(I_A)), \mathcal{E}_8(I_B)) \quad (3)$$



Figure 5. **Face Refinement.** Given generated humans, we use a dedicated face model trained on the same dataset to improve the quality of the face region. We jointly optimize both the face and the human latent codes so that the two generators coordinate with each other to produce coherent results. The two inset face crops show the initial face generated by the body GAN (*bottom*) and the final face improved by dedicated face GAN (*top*).

where $\mathcal{E}_x(I)$ is the border region of I of width x pixels.

To maintain realism during the optimization, we also add two regularization terms:

$$L_{reg} := \lambda_{r1} \|\mathbf{w}^* - \mathbf{w}_{avg}\| + \lambda_{r2} \sum_i \|\delta_i\| \quad (4)$$

The first term prevents the optimized latent code from deviating too far from the average latent. We compute \mathbf{w}_{avg} by randomly sampling a large number of latents in Z space, mapping them to W space, and computing the average. The second term is to regularize the latent code in \mathbf{w}^+ latent space. During StyleGAN2 inference, the same 512 dimensional latent code \mathbf{w} is fed into each of the n generator layers (n is dependent on the output resolution). Many GAN inversion methods optimize in this $n \times 512$ dimensional \mathbf{w}^+ latent space [31] instead of the 512 dimensional \mathbf{w} latent space. We follow recent work to decompose the \mathbf{w}^+ latent into a single base \mathbf{w}^* latent and n offset latents δ_i . The latent used for layer i is $\mathbf{w} + \delta_i$. We use the L_2 norm as regularizer to ensure that the δ_i s remain small. Based on our visual analysis of the results, we use larger weights for the body generator than the face generator for this regularizer.

We mix and match the various losses depending on the specific application at hand.

Face Refinement versus Face Swap. Given a randomly generated human body $\mathcal{G}_A(\mathbf{w}_A)$, we can keep \mathbf{w}_A fixed and optimize for \mathbf{w}_B such that $\mathcal{G}_B(\mathbf{w}_B)$ looks similar to $\mathcal{B}(\mathcal{G}_A(\mathbf{w}_A))$ at a coarse scale and matches the boundary pixels at a fine scale (Fig. 2 top right). We have:

$$\min_{\mathbf{w}_B} (L_{coarse} + L_{border}) \quad (5)$$

While this almost produces satisfactory results, boundary discontinuities show up at times. For further improvement,




Figure 6. **Multimodal Body Generation for an Existing Face.** For each face generated by the pretrained FFHQ model (*middle column*), we use joint optimization to generate three different bodies while maintaining the facial identities from the input faces.

we can optimize both \mathbf{w}_A and \mathbf{w}_B so that both generators coordinate with each other to generate a coherent image free of blending artifacts (Fig. 2 bottom right). To keep the body appearance unchanged during the optimization of \mathbf{w}_A , we introduce an additional loss term:

$$L_{body} := \lambda_5 L_1(R^O(I_A), R^O(I_{ref})) + \lambda_6 L_{lips}(R^O(I_A), R^O(I_{ref})) \quad (6)$$

where I_{ref} is the input reference body generated by \mathcal{G}_A that should remain unchanged during the optimization, R^O defines the body region outside of the face bounding box. We also use the mean latent regularization term L_{reg} to prevent generating artifacts. The final objective function becomes:

$$\min_{\mathbf{w}_A, \mathbf{w}_B} (L_{coarse} + L_{border} + L_{reg} + L_{body}) \quad (7)$$

Fig. 1(b) and Fig. 5 show face refinement results using a dedicated face model trained on faces cropped from the same data used for training the body generator. Our refinement results when using the pretrained FFHQ face model exhibit similar visual quality (see supplementary material).

Body Generation for an Existing Face. Given a real face or a randomly-generated face $\mathcal{G}_B(\mathbf{w}_B)$, we can keep \mathbf{w}_B fixed and optimize for \mathbf{w}_A such that $\mathcal{G}_A(\mathbf{w}_A)$ produces a body that looks compatible with the input face in terms of pose, skin tone, gender, hair style, etc. In practice, we find that to best maintain boundary continuity, especially when generating bodies to match faces of complex hair styles, it is often to discourage large changes in \mathbf{w}_B , such that the

face identity is mostly preserved but the boundary and background pixels can be slightly adjusted to make the optimization of \mathbf{w}_A easier. To preserve the face identity during the optimization, we use an additional face reconstruction loss:

$$L_{face} := \lambda_7 L_1(R^I(I_B), R^I(I_{ref})) + \lambda_8 L_{lips}(R^I(I_B), R^I(I_{ref})) \quad (8)$$

where R^I defines the interior region of the face crop and I_{ref} denotes the referenced input face. For more precise control, face segmentation can be used instead of bounding boxes. Our objective function becomes:

$$\min_{\mathbf{w}_A, \mathbf{w}_B} (L_{coarse} + L_{border} + L_{reg} + L_{face}) \quad (9)$$

With different initialization for \mathbf{w}_A , we can generate multiple results per face as shown in Fig. 6. Note that our model can generate diverse body appearances compatible with the input face. The generated body skin tone generally match the input face skin tone (e.g., the women of African descent in the top and bottom rows of Fig. 6). Fig. 1(c) shows another example.

Face Body Montage. We can combine any real or generated face with any generated body to produce a photo montage. With a real face, we need to first encode it into the latent space of \mathcal{G}_B as \mathbf{w}_B using an off-the-shell encoder [30]. Similarly, a real body could be encoded into the latent space of \mathcal{G}_A , but due to the high variability of human bodies it is difficult to achieve a low reconstruction error. All montage results are created from synthetic bodies generated by \mathcal{G}_B . We use the following objective function:

$$\min_{\mathbf{w}_A, \mathbf{w}_B} (L_{coarse} + L_{border} + L_{reg} + L_{face} + L_{body}) \quad (10)$$

Fig. 7 shows the result of combining faces (*top row*) generated by the pretrained FFHQ model with bodies (*leftmost column*) generated by our full-body generator \mathcal{G}_A . With minor adjustments of both the face and body latent codes, we achieve composition results that are coherent and identity-preserving. While we do not have any explicit loss encouraging skin tone coherence, given faces with different skin tones, our joint optimization slightly adjusts the skin tone of the body’s neck and hand pixels to minimize appearance incoherence and boundary discrepancy in the final results. Fig. 1(d) shows two more examples. Our joint optimization is able to slightly adjust the shoulder region of the lady to extend her hair to naturally rest on her right shoulder. The rightmost column in Fig. 2 shows the improvement the joint optimization makes to the final result quality (*bottom*) compared to only optimizing \mathbf{w}_B given an input body (*top*).

Optimization Details. While the difference is subtle, we observe a slightly better visual performance when using L_1 over L_2 losses. We apply many of our losses to down-sampled versions of the images $\mathcal{D}_{64}(\mathcal{B}(I_A))$ and $\mathcal{D}_{64}(I_B)$



Figure 7. **Face Body Montage.** Given target faces (*top row*) generated by the pretrained FFHQ model and bodies (*leftmost column*) generated by our full-body human generator, we apply *joint* latent optimization to find compatible face and human latent codes that can be combined to produce coherent full-body humans. Notice how face and skin colors get synchronized, and zoom in to observe the (lack of) seams around face insets.

to allow for more flexibility during optimization and to reduce the risk of overfitting to artifacts from the source image (e.g., the body GAN’s face region, which lacks realistic high-frequency details) in a strategy similar to PULSE [24].

One challenge in the joint optimization of w_A and w_B is that the boundary condition Ω depends on the variable w_A . We address this by alternately optimizing for w_A and w_B , and reevaluating the boundary after each update of w_A . We



Figure 8. **Multimodal Face Improvement.** To improve humans generated by a full-body model trained on DeepFashion, we use the pretrained FFHQ model to synthesize a variety of seamlessly merged result faces that all look compatible with the input body.

stop the process when the updates converge.

Optimization Initialization. The default choice of initialization for either w_A or w_B is their corresponding average latent vector w_{avg} . This typically leads to reasonable results quickly. However, it is desirable to generate a variety of results for applications like finding matching bodies I_A for an input face I_B . In this case, we start from truncated latent codes $w_{trunc} = w_{rand} * (1 - \alpha) + w_{avg} * \alpha$. Due to the introduced randomness and the interpolation with the average latent code, we can generate diverse yet realistic results (see Fig. 6). In Fig. 8, given humans generated by our full-body model trained on DeepFashion, we use the pretrained FFHQ face model to swap in multiple better looking faces. Different initialization of w_B yields different results. In the cases where either the face region or the body region should remain fixed during the joint optimization of both latent codes, we initialize the optimization with the latent code initially used to generate the synthetic reference image or the latent code encoded from a real image.

4. Dataset and Implementation

We curate a proprietary dataset of 83,972 high-quality full-body human photographs at 1024×1024 px resolution. These images stem from a dataset of 100,718 diverse photographs in the wild purchased from a third-party data vendor. The dataset includes hand-labeled ground-truth segmentation masks. We apply a human pose-detection network [6] on the original images and filter out those that contain extreme poses causing pose detection results to have low confidence. Fig. 9 shows some sample training images.

Feature alignment plays an important role in high-quality image generation, as can be seen in the qualitative difference of models trained on FFHQ data vs. other face datasets. Therefore, we carefully align the humans using their pose skeletons. We define an upper body axis based on the position of the neck and hip joints. We position humans so that the upper body axis is aligned in the center of the image. As the variance in perspective and pose is very large, choosing the appropriate scale for each person within the context of their image frame is challenging. We



Figure 9. **Full-body Human Dataset.** We create a dataset from photographs of humans in the wild. The images are automatically preprocessed, aligned, and cropped to 1024×1024 px resolution using the ground-truth segmentation masks and detected pose skeletons.

scale the humans based on their upper-body length and then evaluate the extent of the face region as defined by the segmentation mask. If the face length is smaller (larger) than a given minimum (maximum) value, we rescale so that the face length is equal to the minimum (maximum).

Lastly, we enlarge the backgrounds using reflection padding and heavily blur them using a Gaussian kernel of size 27 to focus the generator capacity on modeling only the foreground humans. The huge variation in background appearance in these in-the-wild photos poses extreme challenges for the GAN, especially with limited data quantity.

We also considered completely removing the background, but did not do it for two reasons: (1) Human-labeled segmentation masks are still imperfect around boundaries and (2) We observe that current GAN architectures do not handle large areas of uniform color well.

We also show our method on DeepFashion [21], which consists of 66,607 fashion photographs, including garment pieces and garments on humans. Using the same alignment strategy as above, we extract 10,145 full-body images at 1024×768 resolution. Since the backgrounds are already uniform, we do not blur them.

Training Details We trained our main human body generator network at 1024×1024 px resolution using the StyleGAN2-ADA architecture using all augmentation schemes proposed in the paper [11] for 28 days and 18 hours on 4 Titan V GPUs, using a batch size of 4, processing a total of 42M images. After experimenting with different $R1$ γ values between 0.1 and 20, we chose a value of 13. Similarly, we trained our DeepFashion human generator network at 1024×768 px resolution for 9 days on 4 v100 GPUs, using a batch size of 8, processing a total of 18M images. We use a pretrained FaceNet [27] to detect and align the bounding boxes of the face regions in our generated bodies and faces. The running time of our optimization algorithm for jointly optimizing two generator latents at 1024×1024 px output resolution is about 75 seconds on a Titan RTX GPU. If \mathcal{G}_B has a smaller resolution of 256×256 px, the optimization time decreases to around 60 seconds.

5. Evaluation and Discussion

Quantitative Evaluation. We follow the standard practice to calculate FID (Fréchet Inception Distance) to measure how closely our generated full-body results follow the training distribution. Many previous papers including Co-

ModGAN point out that FID statistics are noisy and do not correlate well with human perception about visual quality. We also observe that FID is more sensitive to result diversity than quality and increases significantly as we truncate the generated results, which reduces variation but is crucial for generating natural looking images with fewer artifacts. While the FID for untruncated results is 13.96, it rises to 26.67 for $t=0.7$ and 71.90 for $t=0.4$ (more truncation). We compare FID values of several alternative approaches for our face refinement application. We use two different truncation settings, $t=0.7$ and $t=0.4$ and evaluate on both the full-body images and image crops that include the refined face and the boundary pixels after copy&pasting.

	t=0.7		t=0.4	
FID Score (lower is better)	body	face	body	face
Unconditional Generation	26.67	27.14	71.90	66.61
InsetGAN	25.33	31.61	69.58	61.57

The differences in FID are small. This indicates that the face refinement using joint optimization does not modify the distribution learned by the unconditional generator and therefore does not decrease the result diversity. However, large differences in perceptual quality are still possible despite similar FID values as demonstrated in our user study.

Baseline Comparison. To the best of our knowledge, no other prior work generates full-body humans unconditionally or inpaints/outpaints humans at 1024×1024 px resolution without requiring conditioning other than reference pixels of the known regions. Previous works have attempted to generate plausible human bodies, but they require segmentation masks as input [19, 20]. The best state-of-the-art method that can be repurposed for our body generation and face refinement applications is CoModGAN [33]. Fig. 10 shows that our InsetGAN (*top right*) outperforms CoModGAN (bottom right) in replacing the initial face generated by our body generator (*left*). We trained CoModGAN with square (with small random offsets for generalization) holes around faces using the official implementation, training data and default parameters for two weeks on four V100 GPUs. Similarly, we train CoModGAN with rectangular holes around the bodies to compare with our InsetGAN for the body generation task. In Fig. 11, we show the two best results of CoModGAN obtained by using several random initializations per input face. Compared to our results in Fig. 6, CoModGAN produces less realistic and diverse image completions.



Figure 10. **Face Refinement Comparison with CoModGAN [33].** Given generated humans (*left*), InsetGAN improves the face quality (*top right*), producing more convincing results than CoModGAN (*bottom right*). CoModGAN results are generated by defining rectangular holes around the face regions.

User Study. We performed a user study to better evaluate the perceptual quality of our method. We aggregated 500 generated humans from our full-body generator and 500 random training images. We then applied either our joint optimization method or CoModGAN to replace the face regions in the generated samples. We showed several sets of image pairs to volunteer participants on Amazon Mechanical Turk and asked them to pick “in which of the two images the person looks more plausible and real”. Per image pair, 5 votes were collected and aggregated towards the majority votes. The study shows that in 12.4% of image pairs, users prefer our unrefined results over the training images when given only 1 second to look at each image. This shows that our results get the basic human proportion and pose right, being able to confuse people about being real. In 98% of cases, users prefer our joint optimization results over the unrefined images. In contrast, only 7% of the CoModGAN samples were picked over the unrefined images, which is consistent with our observation from Fig. 10.

Limitations. Our work has multiple limitations that could benefit from improvements. First, the joint optimization approach may change details such as hair style, neckline or clothing details. In many cases the changes are minor, but in some cases changes can be larger, e.g. the woman’s hair in the middle row of Fig. 6 and the man’s collar in the top row of Fig. 7. Second, as shown in Fig. 3, our full-body GAN has other problems not improved by InsetGAN: symmetry, e.g. noticeable in the hands and feet and on outfits (the woman with a light blue shirt), and the consistency of the fabric used for the clothing (the rightmost person). Third, the generated results have limited variations in body type and pose as discussed next in more detail. The vast majority of our generated results have a slender body type due to the training data distribution.

Dataset Bias and Societal Impact. Both DeepFashion and our proprietary dataset contain biases. DeepFashion contains a limited number of unique identities. The same models appear in multiple images. An overwhelming majority of the images are female (around 9:1) and the range of age, ethnicity and body shape does not represent the real human population. As a result, models trained on it can only



Figure 11. **Body Generation with CoModGAN [33].** We show results generated by CoModGAN trained to fill a rectangular hole covering the body in a given image. Inputs with holes are shown in insets. We generate several results per input and show the best looking two here. Please refer to Fig. 6 for our results on the same input faces. We observe that CoModGAN creates seamless content, but worse visual quality compared to ours.

generate limited range of identities (mostly young white females) as shown in Fig. 8. We made our best attempt to look for a diverse dataset and purchased one from a data vendor in East Asia but noticed the over-representation of young Asian females in the images. Also, many images appear to portray slim street fashion models; as a result the vast majority of them contain slender body type and formal attires. Models trained on biased datasets tend to learn biased representations of the human body. Due to the over-representation of Asians, our results on other ethnicities contain more artifacts in the face region (see the rightmost four results in Fig. 3). We encourage future research efforts to diversify the training dataset to better serve our diverse society. The unconditional nature of our generation process combined with the limitation of some optimization losses operating at a low resolution implies that our generated human outputs might not preserve the attributes in the input image exactly. In addition, since the age distribution in our data is almost exclusively adult humans, we are not able to faithfully produce bodies for faces of children. Similar to other human domain generative models, our approach can be exploited by malicious users to produce deep fakes. However, as we have seen in the user study, even in a short second, users identify most of our generated results as fake compared to real human images. As we further improve the result qualities, we hope, and encourage other researchers, to investigate deep fake detection algorithms.

6. Conclusion

We presented InsetGAN, the first viable framework to generate plausible-looking human images unconditionally at 1024×1024 px resolution. The main technical contribution of InsetGAN is to introduce a multi-GAN optimization framework that jointly optimizes the latent codes of two or more collaborative generators. In future work, we propose to extend the multi-generator idea to 3D shape representations, such as 3D GANs or auto-regressive models based on transformers. We also plan to demonstrate the InsetGAN framework on other image domains and investigate coordinated latent editing in the multi-GAN setup.

References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN: How to embed images into the stylegan latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4432–4441, 2019. 3
- [2] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN++: How to edit the embedded images? In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8296–8305, 2020. 3
- [3] Yuval Alaluf, Or Patashnik, and Daniel Cohen-Or. Restyle: A residual-based stylegan encoder via iterative refinement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6711–6720, 2021. 3
- [4] Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. Pose with Style: Detail-preserving pose-guided image synthesis with conditional stylegan. *ACM Transactions on Graphics*, 2021. 3
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations (ICLR)*, 2019. 2
- [6] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 43(01):172–186, 2021. 6
- [7] Anna Frühstück, Ibraheem Alhashim, and Peter Wonka. TileGAN: Synthesis of large-scale non-homogeneous textures. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 38(4):58:1–58:11, 2019. 1
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. 2
- [9] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 2
- [10] Jinfeng Jiang, Guiqing Li, Shihao Wu, Huiqian Zhang, and Yongwei Nie. BPA-GAN: Human motion transfer using body-part-aware generative adversarial networks. *Graphical Models*, 115:101107, 2021. 3
- [11] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Proceedings of the IEEE Conference on Neural Information Processing Systems (NeurIPS)*, 2020. 2, 3, 7
- [12] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. In *Proc. NeurIPS*, 2021. 2
- [13] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410, 2019. 2
- [14] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8107–8116, 2020. 1, 2, 3
- [15] Markus Knoche, Istvan Sarandi, and Bastian Leibe. Reposing humans by warping 3D features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020. 3
- [16] Sheela Raju Kurupathi, Pramod Murthy, and Didier Stricker. Generation of human images with clothing using advanced conditional generative adversarial networks. In *DeLTA*, 2020. 3
- [17] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [18] Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. TryOnGAN: Body-aware try-on via layered interpolation. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 40(4), 2021. 3
- [19] Yijun Li, Lu Jiang, and Ming-Hsuan Yang. Controllable and progressive image extrapolation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2140–2149, January 2021. 3, 7
- [20] Yuheng Li, Yijun Li, Jingwan Lu, Eli Shechtman, Yong Jae Lee, and Krishna Kumar Singh. Collaging class-specific GANs for semantic image synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021. 2, 3, 7
- [21] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 7
- [22] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *Advances in Neural Information Processing Systems*, pages 405–415, 2017. 3
- [23] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed GAN. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 3
- [24] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. PULSE: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2437–2445, 2020. 6
- [25] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in Style: a StyleGAN encoder for image-to-image translation. In *Proceedings of the IEEE/CVF International Conference*

- on *Computer Vision and Pattern Recognition (CVPR)*, June 2021. [3](#)
- [26] Soubhik Sanyal, Alex Vorobiov, Timo Bolkart, Matthew Loper, Betty Mohler, Larry S. Davis, Javier Romero, and Michael J. Black. 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 (ICCV)*, pages 11138–11147, October 2021. [3](#)
- [27] Florian Schroff, Dmitry Kalenichenko, and James Philbin. FaceNet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, June 2015. [7](#)
- [28] Aliaksandr Siarohin, Stéphane Lathuilière, Enver Sangineto, and Nicu Sebe. Appearance and Pose-Conditioned human image generation using deformable GANs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. [3](#)
- [29] Piotr Teterwak, Aaron Sarna, Dilip Krishnan, Aaron Maschinot, David Belanger, Ce Liu, and William T. Freeman. Boundless: Generative adversarial networks for image extension. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. [3](#)
- [30] Omer Tov, Yuval Alaluf, Yotam Nitzan, Or Patashnik, and Daniel Cohen-Or. Designing an encoder for stylegan image manipulation. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 40(4), jul 2021. [3](#), [5](#)
- [31] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. GAN inversion: A survey. *CoRR*, abs/2101.05278, 2021. [4](#)
- [32] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [4](#)
- [33] Shengyu Zhao, Jonathan Cui, Yilun Sheng, Yue Dong, Xiao Liang, Eric I Chang, and Yan Xu. Large scale image completion via Co-Modulated generative adversarial networks. In *International Conference on Learning Representations (ICLR)*, 2021. [3](#), [7](#), [8](#)
- [34] Zibo Zhao, Wen Liu, Yanyu Xu, Xianing Chen, Weixin Luo, Lei Jin, Bohui Zhu, Tong Liu, Binqiang Zhao, and Shenghua Gao. Prior based human completion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7951–7961, June 2021. [3](#)
- [35] Qiang Zhou, Shiyin Wang, Yitong Wang, Zilong Huang, and Xinggong Wang. Human de-occlusion: Invisible perception and recovery for humans. *CoRR*, abs/2103.11597, 2021. [3](#)