Instant 3D Human Avatar Generation using Image Diffusion Models

Nikos Kolotouros[®], Thiemo Alldieck[®], Enric Corona[®], Eduard Gabriel Bazavan, and Cristian Sminchisescu[®]

Google Research* {kolotouros,alldieck,egbazavan,enriccorona,sminchisescu}@google.com



Fig. 1: We present AvatarPopUp, a new method for the automatic generation of 3D human assets. AvatarPopUp can generate rigged 3D models from text or from single images and has control over body pose and shape. In this example, we show 77 models generated from various text prompts in 12 minutes on a single GPU.

Abstract. We present AvatarPopUp, a method for fast, high quality 3D human avatar generation from different input modalities, such as images and text prompts and with control over the generated pose and shape. The common theme is the use of diffusion-based image generation networks that are specialized for each particular task, followed by a 3D lifting network. We purposefully decouple the generation from the 3D modeling which allow us to leverage powerful image synthesis priors, trained on billions of text-image pairs. We fine-tune latent diffusion networks with additional image conditioning for image generation and back-view prediction, and to support qualitatively different multiple 3D hypotheses. Our partial fine-tuning approach allows to adapt the networks for each task without inducing catastrophic forgetting. In our experiments, we demonstrate that our method produces accurate, high-quality 3D avatars with diverse appearance that respect the multimodal text, image, and body control signals. Our approach can produce a 3D model in as few as 2 seconds, a four orders of magnitude speedup w.r.t. the vast majority of existing methods, most of which solve only a subset of our tasks, and with fewer controls. AvatarPopUp enables applications that require the controlled 3D generation of human avatars at scale. The project website can be found at https://www.nikoskolot.com/avatarpopup/.

^{*} Now at Google DeepMind.

1 Introduction

We present AvatarPopUp, a method for instant generation of rigged full-body 3D human avatars, with multimodal controls in the form of text, images, and/or human pose and shape. The remarkable recent progress in image synthesis [15, 24, 54, 56, 58, 63] acted as a catalyst for a wide range of media generation applications. In just a few years, we have witnessed rapid developments in video generation [9,23,34,68,76], audio synthesis [49,77] or text-to-3D object generation [40, 41, 51, 53, 65], among others. Pivotal to success of all these methods is their probabilistic nature, however this requires very large training sets. While inspiring efforts have been made [14], training set size is still a problem in many domains, and particularly for 3D. In an attempt to alleviate the need for massive 3D datasets, DreamFusion [51] leverages the rich priors of text-to-image diffusion models in an optimization framework. The influential DreamFusion ideas were also quickly adopted for 3D avatar creation [32, 37], a field previously dominated by image- or video-based reconstruction solutions [6, 8, 60, 61]. Text-to-avatar methods enabled novel creative processes, but came with a significant drawback. While image-based methods typically use pretrained feed-forward networks and create outputs in seconds, existing text-to-avatar solutions are optimizationbased and take minutes to several hours to complete, per instance.

In this paper, we are closing this gap and present, for the first time, a methodology for *instant*, text-controlled, rigged, full-body 3D human avatar creation. Our AvatarPopUp is purely feed-forward, can be conditioned on images and textual descriptions, allows fine-grained control over the generated body pose and shape, can generate multiple hypotheses, and runs in 2-10 seconds per instance.

Key to success is our pragmatic decoupling of the two stages of probabilistic text-to-image generation and 3D lifting. Decoupling 2D generation and 3D lifting has two major advantages: (1) We can leverage the power of pretrained text-toimage generative networks, which have shown impressive results in modeling complex conditional distributions. Trained with large training sets of images, both generation quality and diversity are very high. (2) We alleviate the need for very large 3D datasets required to train state-of-the-art generative 3D models. Our method generates diverse plausible image configurations that contain rich enough information that can be lifted in 3D with minimal ambiguity. In other words, we distribute the workload between two expert systems: a pretrained probabilistic generation network fine-tuned for our task to produce front and back image views of the person, and a state-of-the-art unimodal, feed-forward image-to-3D model that can be trained using comparably small datasets.

Our proposed decoupling strategy allows us to maximally exploit available data sources with different levels of supervision. We first fine-tune a pretrained Latent Diffusion network to generate images of people based on textual descriptions and with additional control over the desired pose and shape. This step does not require any ground truth 3D data for supervision and enables scaling our image generator to web scale data of images of people in various poses. Next, we leverage a small-scale dataset of scanned 3D human assets and fine-tune a second latent diffusion network to learn the distribution of back side views con-

| 1 from Text Control | | | | | | | | |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|------|-------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|---------------|--------------------------------------------|--|
| Contr | Contr | Body | Gener | Genera | Editable | Runtin | | |
| Image: A second s | 1 | 1 | 1 | Image: A second s | ✓/X | several hours | Optimization-based text-to-3D [26, 32, 37] | |
| × | 1 | X | 1 | ✓ / X | × | seconds | Image-to-3D [4, 8, 60, 61] | |
| 1 | × | 1 | 1 | X | × | minutes | Hybrid [31] | |
| × | × | 1 | 1 | 1 | × | seconds | 3D GANs [16, 25] | |
| Image: A set of the set of the | 1 | 1 | 1 | 1 | 1 | seconds | AvatarPopUp (Ours) | |

Table 1: AvatarPopUp generates 3D assets with texture from text prompts or input images of a target subject and can be controlled with body pose and shape. In contrast to baselines that require up to hours per prompt, our model takes under five seconds and can de facto by used in interactive applications. In the experimental section we demonstrate the large diversity of the model, and applications in cloth editing.

ditioned on a front view image of the person. We optionally also condition on a textual description that can naturally complement the evidence available in the front view image. Furthermore, we propose a novel fine-tuning strategy that prevents overfitting to the new datasets. Finally, we design and train a 3D reconstruction network that predicts a textured 3D shape in the form of an implicit signed distance field given the pair of front and back views and optionally 3D body signals. The resulting cascaded method, which we call AvatarPopUp supports a wide range of 3D generation and reconstruction tasks: First, it enables fast and interactive 3D generation of assets at scale, see Fig. 1. Second, we can repurpose parts of the cascade for image-based 3D reconstruction at state-of-theart quality. Finally, we demonstrate how AvatarPopUp can be used for creative editing tasks exemplified in 3D virtual try-on with body shape preservation. To summarize, our main contributions are:

- A method for controllable 3D human avatar generation, based on multimodal text, pose, shape and image input signals, that outputs a detailed human mesh instance in 2-10 seconds.
- We propose a simple yet effective way to fine-tune pretrained diffusion models on small-scale datasets, without inducing catastrophic forgetting.
- While not our primary goal, our approach achieves state-of-the art results in single-image 3D reconstruction and enables 3D creative editing applications.

2 Related Work

Table 1 summarizes the characteristics of AvatarPopUp in comparison to previous work, along several important property axes.

Text-to-3D generation. The success of text-to-image models [55, 57, 59] was quickly followed by a significant amount of work on text-to-3D content generation [38, 45, 51, 66]. Due to limited training data, methods typically use optimization approaches, where a neural representation [47] is optimized per instance by minimizing a distillation loss [51] derived from large text-to-image models. This idea has been extended to generate human avatars [18, 27, 29, 32, 37, 69, 73, 78–80] or heads [19, 36, 39] enabling the text-based creation of 3D human assets that

are diverse in terms of shape, appearance, clothing and various accessories. In these works, the optimization process is often regularized using a 3D body model [32, 37, 78], which also enables animation. However, such approaches generally take hours per instance, and rendering is slow. With the appearance of Gaussian Splatting [30], other works [2, 42, 82] reduced rendering time at the expense of accurate geometry. In any case, creating an avatar still takes a significant amount of time, making such methods unsuitable for interactive applications. In this work we propose an alternative direction, which also builds upon the success of text-to-image models and combines them with 3D reconstruction pipelines. Also related are 3D human generation methods. AG3D [16] and EVA3D [25] are GAN-based methods learned from 2D data, that allow to sample 3D humans anchored in a 3D body model. CHUPA [31] generates dual normal maps based on text and then fits a body model to obtain a full 3D representation. While generation is similar in spirit to our method, CHUPA requires optimization per instance and does not generate texture.

Photorealistic 3D Human Reconstruction. Our framework generates 3D human assets and builds on top of state-of-the-art 3D reconstruction techniques. This has been widely explored in the past and can be roughly categorized by its use of explicit or implicit representations. An important line of work leverages 3D body models [43,50,72] and reconstructs their associated parameters, in some cases extended with vertex offsets to represent some clothing and hair detail [4-7, 48, 85]. Other efforts have considered voxels [67, 84], depth maps [17] and more recently implicit representations [8,13,21,28,60–62,70,71,74,75,83]. Being topology free, the latter allow the representation of loose clothing more easily. They typically provide more detail and enable high-resolution reconstruction, often conditioned on local pixel-aligned features [60]. On the other hand, these methods yield reconstructions with no semantic labels that cannot be easily animated. To solve this problem, some work combined body models with implicit representations [13,21,27,28,70,71], but this is prone to errors when the pose is noisy at inference time. In contrast, we drive the synthesis process with guidance from an input body model – sampled or estimated – so that the generated image is well aligned with the body prior. As we show in Sec. 4, this allows rigging the 3D avatar without post-processing and natively supports 3D animation.

Given a single input image of a person, previous work aims to generate realistic reconstructions for the non-visible parts. However, this often leads to blurry textures and lack of geometric detail, *e.g.* no wrinkles. Some methods [60, 61, 70, 71] generate back normal maps to enhance details, or consider probabilistic reconstructions [3, 62]. However, all these methods cannot be prompted from text or other modalities and still yield limited diversity. In contrast, we guide the synthesis process by means of generated front and back images, yielding high-quality 3D reconstructions. Another challenge in previous work is limited training data. Most prior methods rely on a few hundred 3D scans, due to the pricey and laborious process of good quality human capture. We alleviate the need for large scale 3D training data by proposing a framework that can quickly generate humans with a given clothing, pose and shape.



Fig. 2: AvatarPopUp method. (Top) AvatarPopUp builds on the capacity of textto-image models to generate highly detailed and diverse input images. First, a Latent Diffusion network takes a text prompt and a target body pose and shape \mathcal{G} , and generates a highly detailed front image I_f of a person. Next, a second network generates a consistent back view I_b in the same pose and clothing. (Bottom) We perform pixelaligned 3D reconstruction given the generated front/back views I_f , I_b and optionally the given 3D body pose and shape \mathcal{G} . This decoupling enables the generation of 3D avatars from text, images or a combination of the two.

3 Method

We learn a distribution p(X|c) of textured 3D shapes X conditioned on a collection of signals c that we factorize as follows

$$p(X|c) = \int \int p(X|I_f, I_b, c) \cdot p(I_b|I_f, c) \cdot p(I_f|c) dI_f dI_b,$$
(1)

where $p(X|I_f, I_b, c)$ is the probability of 3D shape X given c and front and back image observations I_f and I_b respectively, $p(I_b|I_f, c)$ is the probability of the back view image given the front image I_f and conditioning signals c, and $p(I_f|c)$ the conditional probability of front view images of the person given c.

Computing the integral in (1) is intractable, but our goal is to generate samples from the distribution rather than expectations. To do so, we employ ancestral sampling. We first sample a front view I_f given c, we then sample a back view I_b given I_f and c, and last we sample the 3D reconstruction based on the entire context. In practice, $p(I_f|c)$ and $p(I_b|I_f, c)$ are implemented using Latent Diffusion models, whereas $p(X|I_f, I_b, c)$ is a unimodal, neural implicit field generator.

In the case of single-image 3D reconstruction the conditioning signal c is I_f , and consequently we can omit the first step. For text-based generation, c is a text prompt describing the appearance of the person together with a signal encoding the body pose and shape. The conditioning information c may be extended with additional signals, as in the case of 3D editing, cf. Sec. 4.

3.1 Controllable Text-to-Image Generator

Recent advances in diffusion-based text-to-image generation networks [56] have enabled synthesizing high-quality images given only a text prompt as input. However, for certain use cases, such as human generation, it is difficult to inject fine-grained, inherently continuous forms of control, like the 3D pose of people or their precise body shape proportions, when generating with text alone.

Inspired by ControlNet [81], we propose to add simultaneous control over body pose and shape by augmenting a pretrained Latent Diffusion network with an additional image input that jointly encodes both modalities. For control we use GHUM [72] but other models (e.g. [43]) can be used. Specifically given 3D pose and shape parameters $\boldsymbol{\theta}$ and $\boldsymbol{\beta}$ respectively, we render the corresponding mesh $\mathcal{M} = \text{GHUM}(\boldsymbol{\theta}, \boldsymbol{\beta})$ using GHUM's template coordinates and posed vertex locations as 6D vertex colors, obtaining a dense, pixel-aligned pose- and shapeinformed control signal \mathcal{G} .

To fine-tune the network, we generate a dataset of images of people with corresponding GHUM 3D pose and shape parameters and text annotations. This dataset is comprised of a set of scanned assets [1] that are rendered from different viewpoints, as well as a set of real images scraped from the web. For the synthetic part of the dataset, the pose and shape parameters are obtained by fitting GHUM to 3D scans. We additionally use real images to which we fit GHUM using keypoint optimization in the style of [33]. For all images we obtained text annotations using an off-the shelf image captioning system [12] by prompting it to describe the clothing of the people in the image. Since we are interested in generating 3D human assets, we additionally mask out the background in all images, and train the network to output segmented images. This makes the downstream 3D reconstruction task easier, and improves the reconstruction quality because it encourages the network to focus on human appearance, rather than allocating capacity to model complex backgrounds.

We want to exploit the rich priors learned by text-to-image foundation models by fine-tuning a Latent Diffusion model [56] with the dense GHUM rendering as an additional input. For fine-tuning we propose a simpler and more lightweight method than a standard ControlNet. We pad the weights of the input convolutional layer with additional channels initialized with zeros, and we then fine-tune only the weights of the convolutional layers of the encoder network. All the decoder and attention layers are kept frozen. With this simple strategy, even though our model is trained on a relatively small set of images, it can generalize to unseen types of clothing. At the same time, our strategy is more practical than training a ControlNet, which involves keeping in separate copy of the original network weights in memory, and thus enables fine-tuning large models with moderate hardware utilization. We optimize the encoder of the diffusion model by minimizing the simple variant of the diffusion loss [24]

$$\mathcal{L}(\boldsymbol{\psi}_{\text{enc}}) = \mathbb{E}_{\mathcal{E}(x),\epsilon,t,\tau,\mathcal{G}} \left\| \epsilon - \epsilon_{\boldsymbol{\psi}}(z_t, t, \tau, \mathcal{E}(\mathcal{G})) \right\|,$$
(2)

where $t \in \{1, \ldots, T\}$ is the diffusion time step, $\epsilon \sim \mathcal{N}(0, I)$ the injected noise $z_t = \alpha_t \mathcal{E}(x) + \nu_t$ is the noisy image latent, τ is the text encoding, $\mathcal{E}(\mathcal{G})$ the



Fig. 3: Diverse back view hypotheses. Conditioned on the front view, our method is able to generate diverse plausible back views of the person, with different hairstyles, wrinkle patterns, or lighting. Our network can also be controlled with text (second row), to add fine-grained detail to our generated back-side views.

latent encoding of the dense GHUM signal, and $\psi_{\rm enc}$ the encoder subset of the denoising UNet parameters ψ .

3.2 Back View Generation

One could try to lift the generated front views from the previous stage to 3D directly by applying a single-image 3D reconstruction method like PHORHUM [8]. However, because of the inherent ambiguity of the problem this will result in significant loss of geometric detail and blurry textures for the non-visible body surfaces. To avoid this, we propose to fine-tune again a latent diffusion network with the same strategy as in the previous section. This time the additional image conditioning is a front view and optionally a text prompt, and we train the network to learn the distribution of back views conditioned on the front view. The additional text prompt can be used in cases where it is desired to additionally guide the generation by very specific properties. Fig. 3 shows different back sides sampled from the conditional distribution. We also show that the additional text inputs are useful in modulating certain parts of the generation that are not immediately deducible from the front image, such as hairstyles or specific patterns.

3.3 3D Reconstruction Model

Our 3D reconstruction network is inspired by PHORHUM [8], and our design choices are informed by the limitations of typical single-image 3D reconstruction methods. Specifically, given a collection of input image signals $\mathcal{I} = \{I_f, I_b, \mathcal{G}\}$, we first concatenate them, and then use a convolutional encoder G to compute

a pixel-aligned feature map $G(\mathcal{I})$. The 3D body control signal \mathcal{G} is optional and may be omitted, *e.g.* for single-image reconstruction. Then, each point $\mathbf{x} \in \mathbb{R}^3$ in the scene gets projected on this feature map to get pixel-aligned features $\mathbf{z}_{\mathbf{x}} = g(\mathcal{I}, \mathbf{x}; \pi) = b(G(\mathcal{I}), \pi(\mathbf{x}))$ using interpolation, where $b(\cdot)$ is the bilinear sampling operator and $\pi(\mathbf{x})$ is the pixel location of the projection of \mathbf{x} using the camera π . These pixel aligned features are then concatenated with a positional encoding $\gamma(\mathbf{x})$ of the 3D point and are fed to an MLP f that outputs the signed distance from the surface d as well as surface color c. Finally, the 3D shape \mathcal{S} is represented as the zero-level-set of d

$$\mathcal{S}(\mathcal{I}) = \{ \mathbf{x} \in \mathbb{R}^3 | f(g(I, \mathbf{x}; \pi), \gamma(\mathbf{x})) = (0, \mathbf{c}) \}.$$
(3)

 \mathcal{S} can be transformed to a mesh directly, using Marching Cubes [44].

3.4 Animation of Generated Avatars

Our method can generate diverse 3D avatars with various poses, shapes and appearances. Optionally, we may leverage the conditioning body model to rig the estimated 3D shape. As a result of our conditioning strategy, 3D avatars and the conditional body model instances are aligned in 3D. This allows us to anchor the reconstructed 3D shape on the body model surface [10] and repose or re-shape it accordingly. Alternatively, we may also transfer just the LBS skeleton and weights from the body model to the scan. This enables importing and animating the generated 3D assets in various rendering engines. See Fig. 8 and Sup. Mat. for examples and videos of animations of the generated 3D assets.

4 Experiments

Data. We use meshes from RenderPeople [1] for training as well as our own captured data, totaling in ~ 10 K scans with diverse poses, body shapes, and clothing styles. We render each scan with randomly sampled HDRI background, random cloth color augmentations, and lighting using Blender [11]. During this process, we render both front and back views used to train the different stages of our model. For the front image generation network we also use a set of 10K real images on which we fitted the GHUM model using 2D keypoints. For testing we defined a split based on subject identity and held out ~ 1 K scans.

We provide results for 2 different versions of our model. The standard quality model is generated in 2 seconds by running 5 DDIM [64] steps during inference and Marching Cubes at 256^3 resolution. The high quality model is generated in 10 seconds using 50 DDIM steps and Marching Cubes at 512^3 resolution. The timings were recorded on a single 40 GB A100 GPU. Unless otherwise stated, all results we report are obtained using the high quality model.

Metrics and Baselines. We compare our method numerically for two different problems. First we consider the task of text-to-3D human generation, where we sample 100 different text prompts and compare against representative text-to-3D generation methods. For numerical comparisons we evaluate the text-image



Fig. 4: Diversity of our 3D generation. For the same text prompt and the same pose and shape conditioning, our model can generate a diverse set of 3D avatars that respect both the text and the 3D body controls.

alignment using CLIP [52]. Specifically, we use the retrieval accuracy using CLIP, as proposed in [51]. We additionally show qualitative results.

Second, we validate the performance of our 3D reconstruction component against state-of-the-art methods [8, 27, 60, 61, 70, 71] considering both geometry and texture. Pixel-aligned image features dominate in recent work [8, 60, 61], but some methods proposed to combine them with body models [70, 71], which offer advantages in being able to animate reconstructions. We also rely on pixelaligned features, yet propose a method that inherently enables animation. We run AvatarPopUp by generating the back side of the subject and applying the reconstruction network. To evaluate 3D geometry, we report bi-directional Chamfer distance $\times 10^{-3}$, Normal Consistency (NC \uparrow), and Volumetric Intersection over Union (IoU \uparrow) after ICP alignment. However, these metrics do not necessarily correlate with good visual quality, *e.g.* Chamfer distance is minimized by smooth, non-detailed geometry. To measure the quality of reconstructions, we additionally report FID scores [22] of the front/back views for both geometry and texture.

4.1 Text-to-3D Generation

In Fig. 4 we generate different avatars given the same text prompt and driving poses. We can see that our model is able to create a very diverse set of assets, a property not observed in the previous text-to-3D generation methods [32].

In Tab. 2 we use CLIP to evaluate our model against other text-to-3D generation methods. In general, CLIP-based metrics are not indicative of the generated image quality, because they only consider the alignment with the text, and often over-saturated images with extreme details tend to have high CLIP scores. To further demonstrate that our method generates higher quality avatars, we also include a qualitative comparison in Fig. 5.

| Col | or | Col | or | | |
|--------------------|------------------|--------------------|------------------|-------------------------------------|---------------------|
| R-Prec. \uparrow | Top-3 \uparrow | R-Prec. \uparrow | Top-3 \uparrow | $\operatorname{Runtime} \downarrow$ | |
| 0.68 | 0.92 | 0.04 | 0.25 | 8h | DreamHuman [32] |
| 0.56 | 0.82 | 0.03 | 0.15 | 3h | TADA [37] |
| - | _ | 0.03 | 0.08 | 3m | CHUPA [31] |
| 0.58 | 0.73 | 0.08 | 0.17 | ${\sim}2{ m s}$ | Ours |
| 0.62 | 0.77 | 0.11 | 0.17 | $\sim 10s$ | Ours (high quality) |

Table 2: Numerical comparisons with other text-to-3d human generation methods. We mark the **best** and second best results. Our method allows to tradeoff speed and quality and we report results for two different settings.



Fig. 5: Comparisons with text-to-3d human generation methods. Our method generates high quality results that respect the text prompt well, at a fraction of the others' runtime, *cf*. Tab. 3. TADA's results appear unnatural; DreamHuman failed for

one subject and produces oversaturated colors; CHUPA failed to respect the prompt.

4.2 Single-image 3D Reconstruction

While not specifically designed for 3D reconstruction, our method achieves stateof-the-art performance also for this task. The evaluation setup is the following: given an input image I, we draw one random sample from the back view image generator network, and then feed the pair of front/back images to our 3D reconstruction network. For all methods we extract a textured 3D mesh using Marching Cubes and report numerical results in Tab. 3. Furthermore, we show qualitative results in Fig. 6. Notably, our method not only performs on par numerically and qualitatively on reconstructed front views, but also generates highly detailed back view texture and geometry. Finally, we also compare with the optimization-based method TeCH [27] in Fig. 7. TeCH produces detailed front and back geometry but also also exhibits problems at times, rooted in its 3D pose estimation method. Most importantly, TeCH runs for several hours per instance, while ours computes results in a single feed-forward pass, in only a few seconds.

| 31 | D Metri | cs | | | | | |
|------------------|------------------------|-----------------------|----------|----------|-----------|-----------|--------------|
| Ch. \downarrow | $\mathrm{IoU}\uparrow$ | $\mathrm{NC}\uparrow$ | Color. F | Color B. | Normal F. | Normal B. | |
| 7.1 | 0.50 | 0.72 | _ | _ | 29.6 | 57.2 | PIFu [60] |
| 6.2 | 0.52 | 0.74 | - | - | 26.9 | 52.5 | PIFuHD [61] |
| 7.0 | 0.48 | 0.72 | _ | _ | 36.1 | 49.0 | ICON [71] |
| 3.4 | 0.58 | 0.76 | _ | _ | 18.3 | 36.1 | ECON [70] |
| 0.84 | 0.71 | 0.87 | 11.0 | 40.8 | 11.2 | 38.8 | PHORHUM* [8] |
| 0.91 | 0.70 | 0.86 | 11.0 | 20.0 | 12.8 | 25.8 | Ours |

Table 3: Numerical comparisons with single-view 3D reconstructions methods and ablations of our method. We mark the **best** and second best results. All Chamfer metrics are $\times 10^{-3}$. Not all methods generate colors. For fair comparisons, we retrained PHORHUM [8] using the same data as AvatarPopUp. We observe comparable results in terms of 3D metrics, however ours performs better on generating more realistic and diverse back views and back normals.

4.3 3D Virtual Try On

An immediate application of our method is the option to perform 3D garment edits for a given identity. Given an input image, we first recover 3D pose and shape parameters of the person in the image. Using a text prompt and the identity preservation strategy introduced below, we can generate updated images of the same person wearing different garments or accessories. To preserve the identity of the person, we first locate the person's head in the source image and use then Repaint [46] to out-paint a novel body for the given head while still conditioning on the estimated pose and shape parameters. With this strategy we ensure that our method not only preserves the facial characteristics but also the overall body proportions. In Fig. 9 we illustrate such editing examples. The generated 3D edits present garment details like wrinkles on both front and back views, and preserve the subjects' facial appearance. Also, note that body shape and identity are well preserved for the subjects, even though only one image is given. While there has been a significant amount of 2D virtual try-on research [20, 35, 86], our methodology can generate consistent and highly detailed 3D meshes that can be animated and rendered from other viewpoints.

4.4 Animation

AvatarPopUp enables animation of the generated assets by design, provided that its generation is conditioned on an underlying body model. In Fig. 8, we show an example of a generated avatar that is rigged automatically.

4.5 Ablation Study

Pose and shape encoding in the 3D reconstruction network. We compare the effectiveness of our additional pose and shape encoding inputs \mathcal{G} to our reconstruction network. To do so, we use the same set of 100 text prompts as before and for each text prompt sample a random pose and shape configuration. For each (τ, θ, β) triplet we run inferences and compute 2 meshes: one using



Fig. 6: Qualitative comparisons with state-of-the-art single image 3D reconstruction methods. Our method produces front color and normals on par with state-of-the-art and much more detailed back view hypotheses.

only the front and back images, and another one additionally using the dense GHUM encodings. We then evaluate the Chamfer distance between each mesh and the corresponding GHUM mesh. The model using the additional GHUM signals has an average Chamfer distance of $d_{\text{with}} = 1.4$ whereas the one without $d_{\text{without}} = 8.6$, thus validating our design choice. Not only is the control respected well, but this also allows for animation as discussed previously.

Partial vs. Complete Network Fine-tuning. We finetune two Latent Diffusion networks, one in a standard way by optimizing all parameters, and another one using our proposed strategy where we only finetune the convolutional layers of the encoder. Empirically we observed that the network that was finetuned as a whole experienced catastrophic forgetting, and has poor performance when asked to generate types of garment not seen in the training set. Fig. 10 shows a comparison for text prompts not in the training set.

5 Discussion & Conclusions

Limitations. AvatarPopUp inherits limitations from other pixel-aligned methods, *e.g.* with regions parallel to camera rays being not as detailed in the reconstruction. Further, artifacts might be visible in poses or very loose clothing that are underrepresented during training.

Ethical Considerations. We present a generative tool to create 3D human assets, thus reducing the risk of scanning and using real humans for training large scale 3D generative models. AvatarPopUp generates diverse results and can lead to a better coverage of subject distributions.



Fig. 7: Additional comparisons with TeCH. TeCH is optimization-based and takes several hours to complete for a single reconstruction. Our results are obtained in seconds, yet are as detailed and even less noisy.



Fig. 8: Reposing example. We first reconstruct an avatar "*wearing a gray suit*" in the A-pose and then we transfer it to different poses.



Fig. 9: Identity preserving 3D avatar editing. Our method allows for editing the clothing, while preserving the identity of the generated person. In each row we're using the image on the left as input, recover the person's pose and shape parameters and then run our method using the shown prompts. The generated avatars share the same identity and respect the target prompts.



Fig. 10: Partial vs. complete fine-tuning strategy. For each text prompt the first example is generated using our partial fine-tuning strategy whereas the second by fine-tuning the entire network. Empirically, partial fine-tuning is more resilient to catastrophic forgetting, generating more diverse and better text-aligned images.

We have presented AvatarPopUp, a novel framework for generating 3D human avatars controlled by text or images and yielding rigged 3D models in 2-10 seconds. AvatarPopUp is purely feed-forward, allows for fine-grained control over the generated body pose and shape, and can produce multiple qualitatively different hypotheses. AvatarPopUp is composed of a cascade of expert systems, decoupling image generation and 3D lifting. Through this design choice, Avatar-PopUp benefits from both web-scale image datasets, ensuring high generation diversity, and from smaller size accurate 3D datasets, resulting in reconstructions with increased detail and precisely controlled based on text and identity specifications. In the future, we would like to explore other 3D construction strategies besides pixel-aligned features. Longer term, we aim to support highly detailed and controllable 3D human model generation for entertainment, education, architecture and art, or medical applications.

References

- 1. https://renderpeople.com/
- Abdal, R., Yifan, W., Shi, Z., Xu, Y., Po, R., Kuang, Z., Chen, Q., Yeung, D.Y., Wetzstein, G.: Gaussian shell maps for efficient 3d human generation. arXiv preprint arXiv:2311.17857 (2023)
- AlBahar, B., Saito, S., Tseng, H.Y., Kim, C., Kopf, J., Huang, J.B.: Single-image 3d human digitization with shape-guided diffusion. In: SIGGRAPH Asia (2023)
- 4. Alldieck, T., Magnor, M., Bhatnagar, B.L., Theobalt, C., Pons-Moll, G.: Learning to reconstruct people in clothing from a single rgb camera. In: CVPR (2019)
- 5. Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Detailed human avatars from monocular video. In: 3DV (2018)
- Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3D people models. In: CVPR (2018)
- Alldieck, T., Pons-Moll, G., Theobalt, C., Magnor, M.: Tex2shape: Detailed full human body geometry from a single image. In: ICCV (2019)
- 8. Alldieck, T., Zanfir, M., Sminchisescu, C.: Photorealistic monocular 3D reconstruction of humans wearing clothing. In: CVPR (2022)
- Bar-Tal, O., Chefer, H., Tov, O., Herrmann, C., Paiss, R., Zada, S., Ephrat, A., Hur, J., Li, Y., Michaeli, T., et al.: Lumiere: A space-time diffusion model for video generation. arXiv preprint arXiv:2401.12945 (2024)
- Bazavan, E.G., Zanfir, A., Zanfir, M., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: Hspace: Synthetic parametric humans animated in complex environments. arXiv (2021)
- Blender Online Community: Blender a 3D modelling and rendering package. Blender Foundation, Blender Institute, Amsterdam (2020), http://www.blender. org
- Chen, X., Wang, X., Changpinyo, S., Piergiovanni, A., Padlewski, P., Salz, D., Goodman, S., Grycner, A., Mustafa, B., Beyer, L., Kolesnikov, A., Puigcerver, J., Ding, N., Rong, K., Akbari, H., Mishra, G., Xue, L., Thapliyal, A., Bradbury, J., Kuo, W., Seyedhosseini, M., Jia, C., Ayan, B.K., Riquelme, C., Steiner, A., Angelova, A., Zhai, X., Houlsby, N., Soricut, R.: Pali: A jointly-scaled multilingual language-image model (2022). https://doi.org/10.48550/ARXIV.2209.06794, https://arxiv.org/abs/2209.06794
- Corona, E., Zanfir, M., Alldieck, T., Bazavan, E.G., Zanfir, A., Sminchisescu, C.: Structured 3d features for reconstructing relightable and animatable avatars. In: CVPR (2023)
- Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13142–13153 (2023)
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems 34, 8780–8794 (2021)
- Dong, Z., Xu Chen, J.Y., J.Black, M., Hilliges, O., Geiger, A.: AG3D: Learning to generate 3D avatars from 2D image collections. In: International Conference on Computer Vision (ICCV) (2023)
- 17. Gabeur, V., Franco, J.S., Martin, X., Schmid, C., Rogez, G.: Moulding humans: Non-parametric 3D human shape estimation from single images. In: ICCV (2019)
- Gong, C., Dai, Y., Li, R., Bao, A., Li, J., Yang, J., Zhang, Y., Li, X.: Text2avatar: Text to 3d human avatar generation with codebook-driven body controllable attribute. arXiv preprint arXiv:2401.00711 (2024)

- 16 N. Kolotouros et al.
- Han, X., Cao, Y., Han, K., Zhu, X., Deng, J., Song, Y.Z., Xiang, T., Wong, K.Y.K.: Headsculpt: Crafting 3d head avatars with text. Advances in Neural Information Processing Systems 36 (2024)
- Han, X., Wu, Z., Wu, Z., Yu, R., Davis, L.S.: Viton: An image-based virtual try-on network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7543–7552 (2018)
- 21. He, T., Xu, Y., Saito, S., Soatto, S., Tung, T.: Arch++: Animation-ready clothed human reconstruction revisited. In: CVPR (2021)
- 22. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. Advances in neural information processing systems **30** (2017)
- Ho, J., Chan, W., Saharia, C., Whang, J., Gao, R., Gritsenko, A., Kingma, D.P., Poole, B., Norouzi, M., Fleet, D.J., et al.: Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303 (2022)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. arXiv preprint arxiv:2006.11239 (2020)
- Hong, F., Chen, Z., LAN, Y., Pan, L., Liu, Z.: EVA3d: Compositional 3d human generation from 2d image collections. In: International Conference on Learning Representations (2023), https://openreview.net/forum?id=g7U9jD_2CUr
- Hong, F., Zhang, M., Pan, L., Cai, Z., Yang, L., Liu, Z.: Avatarclip: Zero-shot textdriven generation and animation of 3d avatars. ACM Transactions on Graphics (TOG) 41(4), 1–19 (2022)
- Huang, Y., Yi, H., Xiu, Y., Liao, T., Tang, J., Cai, D., Thies, J.: TeCH: Textguided Reconstruction of Lifelike Clothed Humans. In: International Conference on 3D Vision (3DV) (2024)
- 28. Huang, Z., Xu, Y., Lassner, C., Li, H., Tung, T.: Arch: Animatable reconstruction of clothed humans. In: CVPR (2020)
- Jiang, R., Wang, C., Zhang, J., Chai, M., He, M., Chen, D., Liao, J.: Avatarcraft: Transforming text into neural human avatars with parameterized shape and pose control. arXiv preprint arXiv:2303.17606 (2023)
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics 42(4) (2023)
- Kim, B., Kwon, P., Lee, K., Lee, M., Han, S., Kim, D., Joo, H.: Chupa: Carving 3d clothed humans from skinned shape priors using 2d diffusion probabilistic models. arXiv preprint arXiv:2305.11870 (2023)
- Kolotouros, N., Alldieck, T., Zanfir, A., Bazavan, E., Fieraru, M., Sminchisescu, C.: Dreamhuman: Animatable 3d avatars from text. Advances in Neural Information Processing Systems 36 (2024)
- 33. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 2252–2261 (2019)
- 34. Kondratyuk, D., Yu, L., Gu, X., Lezama, J., Huang, J., Hornung, R., Adam, H., Akbari, H., Alon, Y., Birodkar, V., et al.: Videopoet: A large language model for zero-shot video generation. arXiv preprint arXiv:2312.14125 (2023)
- Lee, S., Gu, G., Park, S., Choi, S., Choo, J.: High-resolution virtual try-on with misalignment and occlusion-handled conditions. In: European Conference on Computer Vision. pp. 204–219. Springer (2022)
- Lei, B., Yu, K., Feng, M., Cui, M., Xie, X.: Diffusiongan3d: Boosting text-guided 3d generation and domain adaption by combining 3d gans and diffusion priors. arXiv preprint arXiv:2312.16837 (2023)

- 37. Liao, T., Yi, H., Xiu, Y., Tang, J., Huang, Y., Thies, J., Black, M.J.: Tada! text to animatable digital avatars. In: 3DV (2023)
- Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 300–309 (2023)
- Liu, H., Wang, X., Wan, Z., Shen, Y., Song, Y., Liao, J., Chen, Q.: Headartist: Text-conditioned 3d head generation with self score distillation. arXiv preprint arXiv:2312.07539 (2023)
- Liu, M., Xu, C., Jin, H., Chen, L., Varma T, M., Xu, Z., Su, H.: One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. Advances in Neural Information Processing Systems 36 (2024)
- Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9298–9309 (2023)
- 42. Liu, X., Zhan, X., Tang, J., Shan, Y., Zeng, G., Lin, D., Liu, X., Liu, Z.: Humangaussian: Text-driven 3d human generation with gaussian splatting. arXiv preprint arXiv:2311.17061 (2023)
- 43. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. ToG (2015)
- 44. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. SIGGRAPH (1987)
- Lorraine, J., Xie, K., Zeng, X., Lin, C.H., Takikawa, T., Sharp, N., Lin, T.Y., Liu, M.Y., Fidler, S., Lucas, J.: Att3d: Amortized text-to-3d object synthesis. arXiv preprint arXiv:2306.07349 (2023)
- Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., Van Gool, L.: Repaint: Inpainting using denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11461–11471 (2022)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. ECCV (2020)
- Onizuka, H., Hayirci, Z., Thomas, D., Sugimoto, A., Uchiyama, H., Taniguchi, R.i.: Tetratsdf: 3D human reconstruction from a single image with a tetrahedral outer shell. In: CVPR (2020)
- Oord, A.v.d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K.: Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499 (2016)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: CVPR (2019)
- Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. In: Int. Conf. Learn. Represent. (2022)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- 53. Raj, A., Kaza, S., Poole, B., Niemeyer, M., Mildenhall, B., Ruiz, N., Zada, S., Aberman, K., Rubenstein, M., Barron, J., Li, Y., Jampani, V.: Dreambooth3d: Subject-driven text-to-3d generation. ICCV (2023)

- 18 N. Kolotouros et al.
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821–8831. PMLR (2021)
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: International Conference on Machine Learning. pp. 8821–8831. PMLR (2021)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 10684–10695 (2022)
- 57. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic textto-image diffusion models with deep language understanding. In: NeurIPS. vol. 35, pp. 36479–36494 (2022)
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic textto-image diffusion models with deep language understanding. Advances in Neural Information Processing Systems 35, 36479–36494 (2022)
- Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: ICCV (2019)
- 61. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In: CVPR (2020)
- Sengupta, A., Alldieck, T., Kolotouros, N., Corona, E., Zanfir, A., Sminchisescu, C.: Diffhuman: Probabilistic photorealistic 3D reconstruction of humans. In: CVPR (2024)
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International conference on machine learning. pp. 2256–2265. PMLR (2015)
- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: Int. Conf. Learn. Represent. (2020)
- Sun, J., Zhang, B., Shao, R., Wang, L., Liu, W., Xie, Z., Liu, Y.: Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. arXiv preprint arXiv:2310.16818 (2023)
- Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653 (2023)
- 67. Varol, G., Ceylan, D., Russell, B., Yang, J., Yumer, E., Laptev, I., Schmid, C.: Bodynet: Volumetric inference of 3D human body shapes. In: ECCV (2018)
- Villegas, R., Babaeizadeh, M., Kindermans, P.J., Moraldo, H., Zhang, H., Saffar, M.T., Castro, S., Kunze, J., Erhan, D.: Phenaki: Variable length video generation from open domain textual description. arXiv preprint arXiv:2210.02399 (2022)
- Wang, J., Liu, Y., Dou, Z., Yu, Z., Liang, Y., Li, X., Wang, W., Xie, R., Song, L.: Disentangled clothed avatar generation from text descriptions. arXiv preprint arXiv:2312.05295 (2023)
- 70. Xiu, Y., Yang, J., Cao, X., Tzionas, D., Black, M.J.: ECON: Explicit Clothed humans Optimized via Normal integration. In: CVPR (June 2023)
- 71. Xiu, Y., Yang, J., Tzionas, D., Black, M.J.: Icon: Implicit clothed humans obtained from normals. In: CVPR (2022)

19

- 72. Xu, H., Bazavan, E.G., Zanfir, A., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: Ghum & ghuml: Generative 3D human shape and articulated pose models. In: CVPR (2020)
- Xu, Y., Yang, Z., Yang, Y.: Seeavatar: Photorealistic text-to-3d avatar generation with constrained geometry and appearance. arXiv preprint arXiv:2312.08889 (2023)
- 74. Yang, Z., Wang, S., Manivasagam, S., Huang, Z., Ma, W.C., Yan, X., Yumer, E., Urtasun, R.: S3: Neural shape, skeleton, and skinning fields for 3D human modeling. In: CVPR (2021)
- 75. Yanting Chan, K., Lin, G., Zhao, H., Lin, W.: Integrated pifu: Integrated pixel aligned implicit function for single-view human reconstruction. In: ECCV (2022)
- Zeng, Y., Wei, G., Zheng, J., Zou, J., Wei, Y., Zhang, Y., Li, H.: Make pixels dance: High-dynamic video generation. arXiv preprint arXiv:2311.10982 (2023)
- 77. Zhang, C., Zhang, C., Zheng, S., Zhang, M., Qamar, M., Bae, S.H., Kweon, I.S.: Audio diffusion model for speech synthesis: A survey on text to speech and speech enhancement in generative ai. arXiv preprint arXiv:2303.13336 (2023)
- Zhang, H., Feng, Y., Kulits, P., Wen, Y., Thies, J., Black, M.J.: Text-guided generation and editing of compositional 3d avatars. arXiv preprint arXiv:2309.07125 (2023)
- Zhang, H., Chen, B., Yang, H., Qu, L., Wang, X., Chen, L., Long, C., Zhu, F., Du, K., Zheng, M.: Avatarverse: High-quality & stable 3d avatar creation from text and pose. arXiv preprint arXiv:2308.03610 (2023)
- Zhang, J., Zhang, X., Zhang, H., Liew, J.H., Zhang, C., Yang, Y., Feng, J.: Avatarstudio: High-fidelity and animatable 3d avatar creation from text. arXiv preprint arXiv:2311.17917 (2023)
- 81. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. IEEE International Conference on Computer Vision (ICCV) (2023)
- Zhao, Z., Bao, Z., Li, Q., Qiu, G., Liu, K.: Psavatar: A point-based morphable shape model for real-time head avatar creation with 3d gaussian splatting. arXiv preprint arXiv:2401.12900 (2024)
- 83. Zheng, Z., Yu, T., Liu, Y., Dai, Q.: Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. PAMI (2021)
- 84. Zheng, Z., Yu, T., Wei, Y., Dai, Q., Liu, Y.: Deephuman: 3D human reconstruction from a single image. In: ICCV (2019)
- 85. Zhu, H., Zuo, X., Wang, S., Cao, X., Yang, R.: Detailed human shape estimation from a single image by hierarchical mesh deformation. In: CVPR (2019)
- Zhu, L., Yang, D., Zhu, T., Reda, F., Chan, W., Saharia, C., Norouzi, M., Kemelmacher-Shlizerman, I.: Tryondiffusion: A tale of two unets. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4606–4615 (2023)