

SMPLitex: A Generative Model and Dataset for 3D Human Texture Estimation from Single Image

Dan Casas
dan.casas@urjc.es

Universidad Rey Juan Carlos
Madrid, Spain

Marc Comino-Trinidad
marc.comino@urjc.es

Abstract

We propose SMPLitex, a method for estimating and manipulating the complete 3D appearance of humans captured from a single image. SMPLitex builds upon the recently proposed generative models for 2D images, and extends their use to the 3D domain through pixel-to-surface correspondences computed on the input image. To this end, we first train a generative model for complete 3D human appearance, and then fit it into the input image by conditioning the generative model to the visible parts of subject. Furthermore, we propose a new dataset of high-quality human textures built by sampling SMPLitex conditioned on subject descriptions and images. We quantitatively and qualitatively evaluate our method in 3 publicly available datasets, demonstrating that SMPLitex significantly outperforms existing methods for human texture estimation while allowing for a wider variety of tasks such as editing, synthesis, and manipulation.

dancasas.github.io/projects/SMPLitex

1 Introduction



Figure 1: From a single image where a human is partly visible (left), SMPLitex automatically estimates a complete 3D texture map that can be applied to SMPL [52] body mesh sequences.

Creating photorealistic 3D virtual humans is a long-standing goal in Computer Graphics and Computer Vision, with important applications in many areas including telecommunications, entertainment, online shopping, and medicine. Among the many tasks required to produce digital humans, *appearance synthesis* is the fundamental step to achieving photorealism. To model virtual humans at scale, methods to create realistic-looking 3D human textures are needed.

Ideally, a method to model human appearance should be generative (to create new complete textures), easy to fit into partial observations (to recover textures from images), and compatible with traditional animation pipelines (to use the textures in popular commercial tools). Unfortunately, despite the impressive advances for digital humans, existing models for human synthesis typically focus on solving one of these challenges in isolation. For example, recent methods have shown impressive results in *posed image synthesis* tasks [9, 12, 30, 36, 45, 52, 53] using generative strategies such as GANs [10, 26] or VAEs, guided with 2D pose representations or text [24]. However, these methods directly output camera-space 2D images of humans but do not generate a complete texture map. This precludes their use in standard 3D animation pipelines, where UV texture maps for texturing 3D meshes are used. Additionally, they are usually not ready to be used to recover textures from in-the-wild images. Similarly, other works leverage neural rendering pipelines [23, 41, 42, 44, 58] to generate view-dependent posed avatars but do not output 3D texture maps either.

Alternatively, and closer to ours, other works focus on *complete texture estimation* from single image [22, 25, 29, 40, 59, 63] and are able to recover 3D texture maps from casual images. Most of these methods used CNN architectures [22, 29] to infer the complete 3D texture map from a single image, sometimes using multi-view supervision [63] or transformer-based architectures [59]. However, they are limited to generating textures from images, which precludes their use in applications that require the synthesis of unseen virtual humans (*e.g.*, via text prompts). Most importantly, they typically output low-detail textures due to the limited expressivity of the latent space of the network. Additionally, we argue that these undesired properties prevent the use of existing methods to build large and high-quality public datasets of 3D human textures which, we believe, is a major shortfall in the field.

In this paper, we address these limitations and propose SMPLitex, a generative method for complete 3D human texture synthesis. SMPLitex enables the estimation of 3D human textures from single images that can be directly applied to SMPL meshes as shown in Figure 1. Additionally, since it is a generative method, it also allows for the synthesis of new textures via text prompts or image editing, as shown in Figures 2 and 3. Under the hood, SMPLitex leverages the recently proposed diffusion models for 2D image synthesis [47, 49] built from a hierarchy of denoising processes. Despite the impressive results of such models in 2D tasks [8, 17, 18, 55, 56], extending diffusion models to 3D humans requires addressing two important challenges: spatial regularization, to enforce multi-view consistency; and 3D awareness, to enable 2D-to-3D tasks. Our key observation is that, for human-related tasks, a proxy of the 3D geometry visible in any image is coarsely modeled with human body models such as SMPL [54], which opens the door to 2D diffusion models for 3D humans. To this end, we first learn a domain-specific diffusion model that is trained to generate *unwrapped* 3D textures of humans, which implicitly learns multi-view consistency. Then, to provide the 3D awareness required for 2D-to-3D tasks such as 3D texture from monocular input, we estimate pixel-to-surface correspondences [13] to project image pixels to an incomplete 3D texture map. Leveraging the 2D structure implicitly enforced in our domain-specific diffusion model, we are able to inpaint the incomplete 3D texture map.

We demonstrate that the proposed model, SMPLitex, outperforms state-of-the-art methods [29, 59] for 3D human texture estimation in 3 publicly available datasets [22, 53, 62], generating high-resolution texture maps. Additionally, we exploit the generative capabilities of the proposed model to create a new dataset of high-quality 3D textures. Our new dataset overcomes the quality, diversity, and number of samples of existing datasets such as SURREAL [65]. This paves the way for new data-driven models that require photorealistic human data such as 3D pose estimation, 3D human reconstruction, and neural rendering.

In summary, our main contributions are:

- SMPLitex, a new generative model for 3D human textures that can be used as a drop-in replacement for textures in any SMPL-based pipeline.
- A novel diffusion-based method to infer 3D human textures from single RGB input.
- A new dataset of high-quality 3D human textures that significantly surpasses the detail, diversity, and size of existing datasets.

2 Related Work

Texture recovery from single image. This group of methods attempts to recover the complete appearance of humans [8, 6, 22, 29, 40, 59, 63], faces [61] or category-specific objects [25, 58] from a single image. To tackle the challenges arising from such ill-posed problem, a variety of learning-based solutions relying on neural networks have been proposed, including the use of multi-view supervision [40, 63], transformer-based architectures [59], and differentiable rendering [22, 25].

Similar to ours, some works frame the texture recovery problem as image inpainting problem [29, 40]. For example, Lazova *et al.* [29] computes pixel-to-surface correspondences from the input image and project the pixel information into a partial texture map. The incomplete texture is then inpainted using a GAN-based network, generating a complete 3D texture of the input image. Neverova *et al.* [40] also train a partial texture inpainting network to infer the occluded parts of the body, which is supervised by multi-view ground truth data. Instead of inpainting a texture map, other methods directly predict the full texture given the input image [22, 59, 60, 63]. For example, Xu *et al.* [59] map 3D coordinates of a human body mesh to a UV texture map which, in combination with a 2D part segmentation image, is converted to a texture map using a transformer-based network. Wang *et al.* [22] utilize a distance metric based on a re-identification loss to learn to generate texture maps given a dataset of images of humans taken from different viewpoints and their corresponding 3D pose. Similarly, Zhao *et al.* [63] add part-based segmentation and enforce cross-view consistency to learn to generate complete texture maps. Despite the impressive results, these methods typically output low-detail textures due to the limited expressivity of the latent space of the networks used. Furthermore, they are also limited to generating textures from images, which precludes their use in applications that require the synthesis of new textures or their manipulation, for example, via text prompts.

Closely related to the texture estimation task are the method that aim to reconstruct 3D humans and appearance from single image [9, 15, 64, 39, 65]. These methods produce high-fidelity 3D reconstructions, including fine geometric details, but the estimated appearance is usually not explicitly baked into a consistent UV texturemap. Furthermore, they do not enable the synthesis of unseen appearances.

Posed image synthesis. Instead of predicting the complete 3D texture map, this group of methods aim at generating images of posed humans [9, 11, 12, 24, 30, 36, 45, 52, 53], hence only requiring to synthesize the visible parts of the subject. Under the hood, these methods use generative strategies, such as GANs [11, 26] or VAEs, guided with 2D pose representations [45], dense surface correspondences [11, 11], or text prompts [24] to describe the target pose. For example, Sarkar *et al.* [52] encode the partial (*i.e.*, visible) UV-space appearance

to a global latent vector to modulate a StyleGAN2 [26] image generator. Similarly, AlBahar *et al.* [1] inpaint a correspondence field and transfer local surface details to the target pose.

Despite the high-quality results of these methods, most of them do not output an explicit texture map that can be used in 3D animation pipelines. This limits their use to purely 2D image synthesis use cases. [14, 57] are notable exceptions and predict explicit dynamic texture maps of the subject for pose synthesis, but require per-subject retraining.

Avatars from text. Also related to ours are the methods that aim at synthesizing humans from text descriptions [1, 19, 24, 28, 62]. These methods combine recent image generative models [9] based on GANs, VAEs, or diffusion models, with large vision-language pre-trained models such as CLIP [46] to condition the output. For example, AvatarCLIP [19] is able to generate and animate 3D textured humans directly from text, and Text2Human [24] synthesizes high-quality 2D posed humans given detailed outfit descriptions. Other works focus on more general objects [21, 27, 57]. However, they do not generate consistent texture maps, and cannot easily be fitted into partial observations to estimate textures from images.

3 3D Human Texture Estimation

Given an input RGB image \mathbf{x} , where a person is visible or partially visible, our goal is to estimate a 3D texture map \mathbf{u} that encodes the complete appearance of the subject. This is a challenging ill-posed problem because (1) the 3D geometry of the scene (*i.e.*, the 3D surface of the person) is unknown; and (2) many parts of the image suffer from natural occlusions and self-occlusions (*i.e.*, not all surface points are visible).

To address this problem, a common approach [29, 59] is, first, to use a coarse geometry proxy [54] and infer pixel-to-surface correspondences [13] to build an incomplete texture map; and later use an image-to-image translation framework to inpaint or estimate the incomplete texture. However, the limited expressivity of existing methods leads to low-resolution inpainted textures that lack detail. Furthermore, the resulting models cannot be manipulated with text prompts, which limits their applicability to image-to-image tasks.

In contrast, we propose a pipeline based on image diffusion models [49] that is capable of recovering high-quality textures, and additionally, it naturally allows for text-based manipulations. Figure 2 presents a visualization of our pipeline. In the rest of this section, we first describe how we formulate and train our generative model (Section 3.1), and then describe how we leverage it to estimate complete 3D textures from in-the-wild monocular RGB images (Section 3.2).

3.1 Generative 3D Human Textures

3.1.1 Background

SMPLitex uses a diffusion model as a generative backbone. Diffusion models are a type of generative model that gradually remove noise from an image to learn the distribution space p_θ [47]. To accomplish this, the model starts with a sampled Gaussian noise and performs a step-by-step denoising process over T time steps until it produces a final, noise-free image. During each step of the denoising process, the diffusion model generates noise ϵ_t that is used to create an intermediate denoised image \mathbf{x}_t . The initial noise \mathbf{x}_T corresponds to the final image, while the fully denoised image \mathbf{x}_0 corresponds to the starting image. This denoising

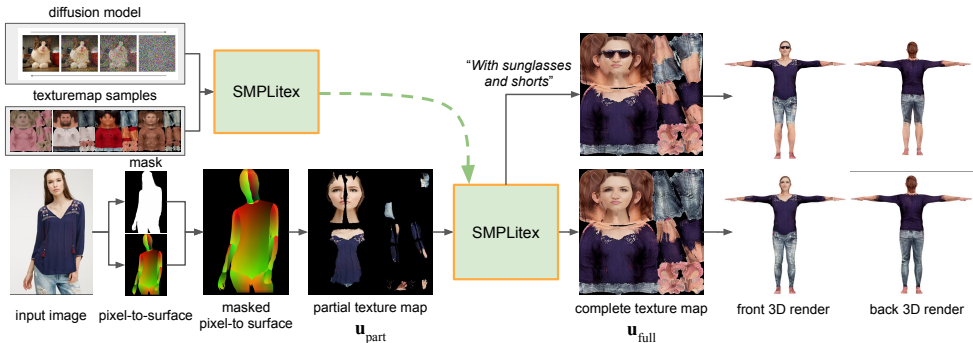


Figure 2: Overview of the proposed approach for texture estimation from a single image.

process is typically modeled as a Markov transition probability as follows

$$p_{\theta}(\mathbf{x}_{T:0}) = p(\mathbf{x}_T) \prod_{t=T}^1 p_{\theta}(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad (1)$$

To improve the efficiency of such diffusion models, *latent* diffusion models (LDM) [49] operate on a lower dimensional latent space \mathbf{z} pre-trained using a variational autoencoder. During training, for an image \mathbf{x} , noise is added to the encoded image $\mathbf{z} = \mathcal{E}(\mathbf{x})$, where \mathcal{E} is the pre-trained encoder, leading to \mathbf{z}_t where the noise level increases with t . Analogous to the original diffusion model, the LDM process can be seen as a sequence of denoising models with shared parameters θ that learn to predict a noisy image $\varepsilon_{\theta}(\mathbf{z}_t, c, t)$, where t is the timestep and c a text condition. LDMs are trained by minimizing the loss term

$$\mathbb{E}_{\mathbf{z}_t \in \mathcal{E}(x), t, c, \varepsilon \sim \mathcal{N}(0,1)} [\|\varepsilon - \varepsilon_{\theta}(\mathbf{z}_t, c, t)\|_2^2] \quad (2)$$

Once trained, the LDM model can generate new samples following the diffusion process in reverse mode, iteratively predicting the noise to be removed from a randomly sampled Gaussian noise (potentially conditioned on text and time step).

3.1.2 Diffusion Model for 3D Humans

Many extensions of the original LDM [49] model have been recently proposed for a wide variety of tasks and data modalities, including image inpainting [8, 35, 54] and video generation [17, 18, 56]. In this work, we look into how LDM can be adapted to the specific case of 3D human appearance. However, naively extending LDMs to 3D domains is challenging since multi-view consistency cannot be guaranteed. A few recent works demonstrate promising results in text-to-3D [43] or novel view synthesis [9, 57] tasks, but articulated objects and high-resolution images remain a challenge.

Our key observation is that, for the specific case of 3D humans, these limitations can be circumvented by using a 3D-to-2D parametrization of the surface (*i.e.*, a UV map of the mesh surface). By working directly on the UV map, we are able to encode the 3D appearance of the human directly on a 2D image, opening the door to the potential of LDMs models for appearance synthesis, inpainting, and manipulation.

Therefore, we propose to leverage the highly expressive LDM proposed by Rombach *et al.* [49], and fine-tune it to encode UV textures of humans. The main underlying challenge of this task is how to ensure that the expressivity of the original model is preserved (*i.e.*, the fine-tuned model is able to generate detailed and rich images that are not in the fine-tuning training set) while satisfying the spatial constraints inherent in the UV textures (*i.e.*, the 2D-to-3D parameterization). Crucially, recent text-to-image models [20, 50] address an analogous problem for model personalization, where the goal is to learn to generate images with a specific style or for a particular subject. Under the hood, these models use a class-specific prior preservation loss that enables the synthesis of the target subject or style in arbitrary scenarios, just using an extremely reduced set of training samples. Inspired by this, SMPLitex leverages the work of Ruiz *et al.* [50] to fine-tune the model of Rombach *et al.* [49] such that it is constrained to synthesize SMPL UV texture maps. In practice, we use 10 UV texture maps for SMPL from [0, 29] to fine-tune the model [49] available at [48] for 1,500 iterations. At inference time, our results use 50 denoising steps and a classifier free guidance (CGF) of 2.0.

3.2 Human Texture Estimation from Single Image

The diffusion model described in Section 3.1 enables the synthesis of high-quality UV texture maps of humans. As discussed above, the diffusion model is sampled conditioned to the time step, such that it removes the noise at time t , but it also allows for additional conditioning signals (*e.g.*, condition text c in Equation 2). Our key intuition is that to enable the estimation of 3D human appearance from a single image, we can condition the synthesis of an LDM model for human appearance *to the visible parts of the subject* in the input image. In other words, we are interested in fitting SMPLitex into natural images.

To this end, similar to [29], we leverage the fact that we are under the assumption of modeling 3D humans and use state-of-the-art pixel-to-surface correspondence models to compute a partial texture map \mathbf{u}_{part} . More specifically, given an input image \mathbf{x} , we estimate pixel-to-surface correspondences \mathbf{d} [40] and project the pixels of \mathbf{x} with assigned surface correspondences to a partial UV map \mathbf{u}_{part} , which will be used as a conditional signal.

However, a naive use of the pixel-to-surface correspondence \mathbf{d} can potentially lead to undesired partial UV maps since \mathbf{d} typically coarsely estimates the foreground silhouette. This can lead to background pixels projected into the UV map, which significantly degenerates the condition image. We mitigate this issue by computing an accurate human silhouette \mathbf{s} [6], that we use to mask the pixel-to-surface image \mathbf{d} . More formally, we compute our partial UV map as

$$\mathbf{u}_{\text{part}} = \Pi(\mathbf{x}, \mathbf{d} \odot \mathbf{s}) \quad (3)$$

where \odot is the Hadamard product, Π is the operator that projects all pixels $p \in \mathbf{x}$ to their corresponding surface coordinate according to the map $\mathbf{d} \odot \mathbf{s}$.

With the partial texture map \mathbf{u}_{part} computed, we can infer the complete texture map \mathbf{u}_{full} of the input image \mathbf{x} by just sampling the diffused model described in Section 3.1 using \mathbf{u}_{part} as an additional conditioning signal.

4 SMPLitex Dataset

Taking advantage of the generative capabilities of the appearance model described in Section 3.1, we build a dataset of high-quality textures by simply sampling the latent space. More



Figure 3: SMPLitex texture samples, side-by-side to their corresponding front and back SMPL render. Notice that these textures are not used for training SMPLitex, instead, they are part of the dataset that we built by arbitrarily sampling our model. We believe such high-quality textures will open the door to new possibilities in the area of virtual humans.

specifically, we use text conditioning prompts describing arbitrary outfit combinations, costumes, sports apparel, job titles, and facial characteristics. Figure 3 depicts 9 samples of the dataset, showcasing a large variety of garment types, outfits, identities, and facial details. In total, the SMPLitex dataset consists of 100 curated textures, see the supplementary material for more details. To clarify, on top of the dataset we will also release the trained SMPLitex model to generate new arbitrary samples, and it is the core component of our method to recover textures from single images. However, we believe that releasing a fixed set of curated textures will benefit future research.

5 Results and Evaluation

We qualitatively and quantitatively evaluate our method on 3 publicly available datasets [24, 59, 60] and demonstrate that our results compare favorably with state-of-the-art methods for texture estimation. In contrast to existing methods, SMPLitex can deal with both low and high-resolution imagery, and it is robust to multi-view consistency metrics.

Evaluation on DeepFashion-MultiModal [24]. This dataset consists of a large collection of high-resolution fashion images (750×1101 pixels), where subjects wear a wide variety of clothing styles. In Figure 4 we present our qualitative results on this dataset, and compare them to the results of the state-of-the-art closest methods [29, 59]. SMPLitex qualitatively outperforms the method of Lazova *et al.* [29], which is based on a GAN inpainting network that is unable to output high-quality garment details that are visible in the input image. In contrast, SMPLitex textures exhibit fine details such as wrinkles and facial attributes, while also extrapolating well to the occluded parts of the body.

Evaluation on Market-1501 [64]. This dataset consists of a large collection of low-resolution images (64×128 pixels) of 1501 different subjects, captured in an urban scenario. It was originally proposed to evaluate person re-identification tasks, but Xu *et al.* [69] extended




Figure 4: Qualitative comparison with state-of-the-art methods [24, 59] in the DeepFashion-MultiModal dataset [24]. SMPLitex clearly outperforms the texture quality of previous methods, recovering fine details such as garment wrinkles and facial attributes.

their use to evaluate human texture estimation. The idea is to estimate the texture of a subject on an image and test the fidelity of the recovered texture by rendering and comparing it to another image of the same subject captured from another viewpoint. We use the same test set and evaluation protocol defined by Xu *et al.* [59].

In Table 1 we show that our approach SMPLitex quantitatively outperforms the state-of-the-art method by Xu *et al.* [59]. Furthermore, in Figure 5 we present a qualitative comparison that demonstrates that SMPLitex can infer faithful textures despite the extremely low-resolution image input. Importantly, notice that Xu *et al.* is a method trained specifically on Market-1501 (*i.e.*, it does not generalize well to other datasets, as we show above in the evaluation with DeepFashion). SMPLitex does not suffer from such an image-specific domain, while still being competitive in the Market-1501 dataset.

	SSIM \uparrow	LPIPS \downarrow
CMR [24]	0.7142	0.1275
HPBTT [63]	0.7420	0.1168
RSTG [22]	0.6735	0.1778
TexGlo [60]	0.6658	0.1776
TexFormer [59]	0.7422	0.1154
SMPLitex (ours)	0.8648	0.0695

Table 1: Quantitative evaluation on Market-1501 [59], following the protocol and evaluation code defined by Xu *et al.* [59].



Figure 5: Qualitative comparison on Market-1501 dataset [64]. Despite not being trained for this challenging dataset of 64×128 pixel images, SMPLitex is able to infer convincing textures that are at least on par with the dataset-specific method of Xu *et al.* [59].

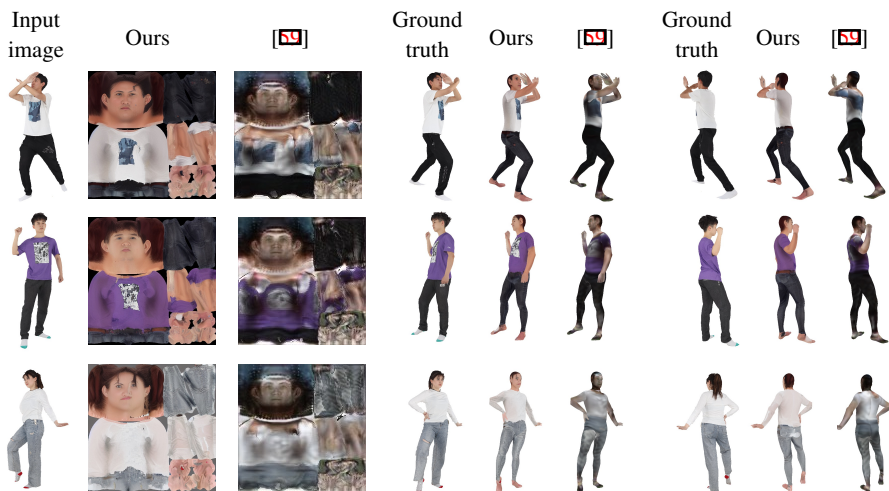


Figure 6: Multi-view consistency evaluation on THUMAN2.0 dataset. Using the input image on the left, we show output textures by SMPLitex (ours) and [59]. Next, we show 2 validation ground truth viewpoints and demonstrate that our renders closely match the ground truth.

Evaluation on THUMAN2.0 [61]. To further evaluate our approach, we contribute with a new evaluation protocol for 3D texture inference from a single image. To this end, we leverage the THUMans2.0 [61] dataset, a high-quality collection of 3D scans with SMPL pose labels, and generate a test set consisting of a render of each scan from a specific viewpoint. We then compare the estimated texture rendered from a different viewpoint with ground truth renders of the scan and compute the pixel similarity. Figure 6 presents qualitative results of this evaluation, demonstrating that our estimated textures closely match the ground truth images from different viewpoints. Since we have 3D ground truth scans, we are able to compute camera-space pixel-based errors of the rendered textures. Table 2 shows that SMPLitex outperforms the state-of-the-art method of Xu *et al.* [59] in both SSIM and LPIPS metrics in this multi-view evaluation protocol on high-resolution images.

	SSIM \uparrow	LPIPS \downarrow
TexFormer [59]	0.8761	0.1223
SMPLitex (ours)	0.8829	0.1067

Table 2: Quantitative evaluation on THUMAN2.0 [61].

6 Conclusions and Limitations

We have presented SMPLitex, a generative model for 3D human appearance that enables the estimation of 3D human textures from single images. SMPLitex leverages recent image diffusion models for 2D image synthesis and uses pixel-to-surface correspondence estimation to bridge the gap between 2D images and 3D surfaces. By conditioning the diffusion model to the visible pixels of a human in a single view, SMPLitex is able to synthesize a complete texture map of the subject, outperforming current methods based on GANs or VAEs.

Despite the convincing quality of the results, SMPLitex suffers from limitations as well. If the subject on the input image is significantly occluded or if surface-to-pixel correspondence fails, SMPLitex sampling is weakly conditioned hence it can potentially generate textures that do not match well the subject. Similarly, when sampling the model with text, if prompts are too general or not related to humans, output textures can exhibit unrealistic facial or body features such as deformed faces or missing limbs.

Acknowledgments. This work has been partially funded by: the European Union’s Horizon 2020 research and innovation program under grant agreement No 899739 (H2020-FETOPEN-2018-2020 CrowdDNA project); and by the Universidad Rey Juan Carlos through the Distinguished Researcher position INVESDIST-04 under the call from 17/12/2020.

References

- [1] Badour AlBahar, Jingwan Lu, Jimei Yang, Zhixin Shu, Eli Shechtman, and Jia-Bin Huang. Pose with Style: Detail-Preserving Pose-Guided Image Synthesis with Conditional StyleGAN. *ACM Transactions on Graphics (TOG)*, 40(6):1–11, 2021.
- [2] Thimo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video Based Reconstruction of 3D People Models. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 8387–8397, 2018.
- [3] Thimo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 1496–1505, 2022. doi: 10.1109/CVPR52688.2022.00156.
- [4] Miguel Ángel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander T Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, Afshin Dehghan, and Joshua M. Susskind. GAUDI: A Neural Architect for Immersive 3D Scene Generation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [5] Sihun Cha, Kwanggyoon Seo, Amirsaman Ashtari, and Junyong Noh. Generating Texture for 3D Human Avatar from a Single Image using Sampling and Refinement Networks. In *Computer Graphics Forum*, volume 42, pages 385–396, 2023.
- [6] Xiangguang Chen, Ye Zhu, Yu Li, Bingtao Fu, Lei Sun, Ying Shan, and Shan Liu. Robust Human Matting via Semantic Guidance. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 2022.

- [7] Soon Yau Cheong, Armin Mustafa, and Andrew Gilbert. KPE: Keypoint Pose Encoding for Transformer-based Image Generation. In *British Machine Vision Conference (BMVC)*, 2022.
- [8] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Improving Diffusion Models for Inverse Problems using Manifold Constraints. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [9] Jianglin Fu, Shikai Li, Yuming Jiang, Kwan-Yee Lin, Chen Qian, Chen Change Loy, Wayne Wu, and Ziwei Liu. StyleGAN-Human: A Data-Centric Odyssey of Human Generation. In *Proc. of European Conference on Computer Vision (ECCV)*, 2022.
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [11] Artur Grigorev, Artem Sevastopolsky, Alexander Vakhitov, and Victor Lempitsky. Coordinate-based texture inpainting for pose-guided human image generation. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [12] Artur Grigorev, Karim Iskakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. StylePeople: A Generative Model of Full-body Human Avatars. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [13] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 7297–7306, 2018.
- [14] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time Deep Dynamic Characters. *ACM Transactions on Graphics*, 40(4), aug 2021.
- [15] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 11026–11036, 2021.
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020.
- [17] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P. Kingma, Ben Poole, Mohammad Norouzi, David J. Fleet, and Tim Salimans. Imagen Video: High Definition Video Generation with Diffusion Models. *arXiv:2210.02303*, 2022.
- [18] Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [19] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. AvatarCLIP: Zero-Shot Text-Driven Generation and Animation of 3D Avatars. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 41(4), 2022.

- [20] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations (ICLR)*, 2022.
- [21] Ajay Jain, Ben Mildenhall, Jonathan T. Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [22] Wang Jian, Zhong Yunshan, Li Yachun, Zhang Chi, and Wei Yichen. Re-Identification Supervised Texture Generation. *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [23] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural Human Radiance Field from a Single Video. In *Proceedings of the European conference on computer vision (ECCV)*, 2022.
- [24] Yuming Jiang, Shuai Yang, Haonan Qiu, Wayne Wu, Chen Change Loy, and Ziwei Liu. Text2Human: Text-Driven Controllable Human Image Generation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 41(4), 2022.
- [25] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 371–386, 2018.
- [26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and Improving the Image Quality of StyleGAN. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [27] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Popa Tiberiu. Clip-mesh: Generating textured meshes from text using pretrained image-text models. *SIGGRAPH Asia 2022 Conference Papers*, December 2022.
- [28] Nikos Kolotouros, Thiemo Alldieck, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Fieraru, and Cristian Sminchisescu. DreamHuman: Animatable 3D Avatars from Text. *arXiv preprint arXiv:2306.09329*, 2023.
- [29] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *International Conference on 3D Vision (3DV)*, pages 643–653, 2019.
- [30] Kathleen M Lewis, Srivatsan Varadharajan, and Ira Kemelmacher-Shlizerman. TryOnGAN: Body-Aware Try-On via Layered Interpolation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 40(4), 2021.
- [31] Zhe Li, Zerong Zheng, Hongwen Zhang, Chaonan Ji, and Yebin Liu. AvatarCap: Animatable Avatar Conditioned Monocular Human Volumetric Capture. In *Proc. of European Conference on Computer Vision (ECCV)*, 2022.
- [32] Lingjie Liu, Weipeng Xu, Marc Habermann, Michael Zollhöfer, Florian Bernard, Hyeonwoo Kim, Wenping Wang, and Christian Theobalt. Learning Dynamic Textures for Neural Rendering of Human Actors. *IEEE Transactions on Visualization and Computer Graphics*, 27(10):4009–4022, 2021. doi: 10.1109/TVCG.2020.2996594.

- [33] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A Skinned Multi-Person Linear Model. *ACM Transactions on Graphics (TOG)*, 34(6):1–16, 2015.
- [35] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 11461–11471, 2022.
- [36] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable Person Image Synthesis with Attribute-Decomposed GAN. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [37] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [38] Aymen Mir, Thiemo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 7023–7034, 2020.
- [39] Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. Siclope: Silhouette-based clothed people. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [40] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense Pose Transfer. In *Proc. of European Conference on Computer Vision (ECCV)*, 2018.
- [41] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural Articulated Radiance Field. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 5762–5772, 2021.
- [42] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable Neural Radiance Fields for Modeling Dynamic Human Bodies. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [43] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D Diffusion. In *International Conference on Learning Representations (ICLR)*, 2023.
- [44] Sergey Prokudin, Michael J Black, and Javier Romero. SMPLpix: Neural Avatars from 3D Human Models. In *Proc. of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1810–1819, 2021.
- [45] Albert Pumarola, Antonio Agudo, Alberto Sanfeliu, and Francesc Moreno-Noguer. Unsupervised Person Image Synthesis in Arbitrary Poses. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 8620–8628, 2018.

- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [47] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [48] Robin Rombach and Patrick Esser. Stable diffusion v1-4 model card. <https://huggingface.co/CompVis/stable-diffusion-v1-4>, 2023. Online; accessed 1-September-2023.
- [49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, 2022.
- [50] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine Tuning Text-to-image Diffusion Models for Subject-Driven Generation. In *cvpr*, 2023.
- [51] Shunsuke Saito, Lingyu Wei, Liwen Hu, Koki Nagano, and Hao Li. Photorealistic facial texture inference using deep neural networks. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [52] Kripasindhu Sarkar, Vladislav Golyanik, Lingjie Liu, and Christian Theobalt. Style and Pose Control for Image Synthesis of Humans from a Single Monocular View. *arXiv:2102.11263*, 2021.
- [53] Kripasindhu Sarkar, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. HumanGAN: A Generative Model of Humans Images. In *International Conference on 3D Vision (3DV)*, pages 258–267. IEEE, 2021.
- [54] Roman Suvorov, Elizaveta Logacheva, Anton Mashikhin, Anastasia Remizova, Arsenii Ashukha, Aleksei Silvestrov, Naejin Kong, Harshith Goka, Kiwoong Park, and Victor Lempitsky. Resolution-robust large mask inpainting with fourier convolutions. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 2149–2159, 2022.
- [55] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [56] Vikram Voleti, Alexia Jolicœur-Martineau, and Christopher Pal. MCVD: Masked Conditional Video Diffusion for Prediction, Generation, and Interpolation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [57] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel View Synthesis with Diffusion Models. In *International Conference on Learning Representations (ICLR)*, 2023.

- [58] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. HumanNeRF: Free-viewpoint Rendering of Moving People from Monocular Video. In *Proc. of Computer Vision and Pattern Recognition (CVPR)*, pages 16210–16220, 2022.
- [59] Xiangyu Xu and Chen Change Loy. 3D human texture estimation from a single image with transformers. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [60] Xiangyu Xu, Hao Chen, Francesc Moreno-Noguer, Laszlo A Jeni, and Fernando De la Torre. 3D Human Pose, Shape and Texture from Low-Resolution Images and Videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 44(9): 4490–4504, 2021.
- [61] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2021)*, June 2021.
- [62] Huichao Zhang, Bowen Chen, Hao Yang, Liao Qu, Xu Wang, Li Chen, Chao Long, Feida Zhu, Kang Du, and Min Zheng. AvatarVerse: High-quality & Stable 3D Avatar Creation from Text and Pose. *arXiv preprint arXiv:2308.03610*, 2023.
- [63] Fang Zhao, Shengcai Liao, Kaihao Zhang, and Ling Shao. Human Parsing Based Texture Transfer from Single Image to 3D Human via Cross-View Consistency. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [64] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proc. of IEEE International Conference on Computer Vision (ICCV)*, pages 1116–1124, 2015.
- [65] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. PaMIR: Parametric Model-Conditioned Implicit Representation for Image-based Human Reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.