

Design2Cloth: 3D Cloth Generation from 2D Masks

Jiali Zheng, Rolandos Alexandros Potamias, Stefanos Zafeiriou
 Imperial College London

{jiali.zheng, r.potamias, s.zafeirou}@imperial.ac.uk



Figure 1. We proposed Design2Cloth: a high fidelity 3D generative model for garment generation from simple 2D masks, trained on a large scale 3D cloth dataset of real-world scans. Figure illustrates interpolation between shirt (Left), trousers (Right) and shapes.

Abstract

In recent years, there has been a significant shift in the field of digital avatar research, towards modeling, animating and reconstructing clothed human representations, as a key step towards creating realistic avatars. However, current 3D cloth generation methods are garment specific or trained completely on synthetic data, hence lacking fine details and realism. In this work, we make a step towards automatic realistic garment design and propose Design2Cloth, a high fidelity 3D generative model trained on a real world dataset from more than 2000 subject scans. To provide vital contribution to the fashion industry, we developed a user-friendly adversarial model capable of generating diverse and detailed clothes simply by drawing a 2D cloth mask. Under a series of both qualitative and quantitative experiments, we showcase that Design2Cloth outperforms current state-of-the-art cloth generative models by a large margin. In addition to the generative properties of our network, we showcase that the proposed method can be used to achieve high quality reconstructions from single in-the-wild images and 3D scans. Dataset, code and pre-

*trained model will become publicly available.*¹

1. Introduction

The advent of 3D digital human avatars has facilitated the generation of realistic humans that continuously progress the gaming and filming industries [29, 36, 53]. Within this context, digitized garments emerge as a pivotal component to enhance realism. However, the diverse and perplexing nature of human garments makes modeling and generation of realistic clothed humans remains non-trivial task.

While 2D cloth generation has been extensively studied obtaining remarkable results [10, 11, 19, 21, 26, 31, 38, 47], the need for precise virtual try-ons necessitates the adoption of 3D cloth models that accurately simulate fabric properties, and enable more realistic digital experience. Recently, several methods have attempted to extend cloth modeling to 3D human avatars [14, 40, 54]. Nevertheless, these approaches prove inadequate when it comes to accurately reconstructing garments in diverse, real-world settings. The primary cause lies in the lack of large 3D garment datasets, compared to the abundance of 2D garment

¹Project Page: <https://jiali-zheng.github.io/Design2Cloth/>

datasets. In particular, real-world cloth scans [25, 30, 50] are composed of only a few number of subjects under pre-defined poses wearing a limited number of garment types. To address the lack of clothed human data, several studies employed artists and simulation engines to curate synthetic data [3, 12, 18, 33, 44, 55, 57] that cover a wider range of garment styles. However, synthetic cloth data have two notable inherited limitations. Firstly, simulation engines fail to accurately replicate the deformations and the wrinkles of real-world garments that leads to clothes that are characterized by an excessive degree of smoothness. Secondly, garment styles are usually generated by deforming a limited number of predefined garment categories, which constrains the model’s generalization to real-world garments.

To overcome the limitations of generative garment models, we have collected a large-scale dataset with high resolution scans from 2010 individuals, spanning a wide range of genders, ages, heights, and weights, wearing more than 2000 unique garments. Leveraging this dataset, we trained a 3D garment generative network, that enables modeling and generation of diverse and highly detailed clothes that follow the real-world distribution. Additionally, to ease the generation process and enable a straightforward utilization of our method for individuals with no prior knowledge in the field, we introduce an inclusion mask representation to describe each garment. Compared to prior works that required UV maps [9, 44] or point clouds [13, 24] to condition the garment generation, the proposed method allows users to effortlessly draw a garment mask to control the generation process. Finally, given that the proposed model is fully differentiable, it can be used as a plug-and-play solution for inverse problems, including 3D garment reconstruction from in-the-wild images and scans.

To summarize, the main contributions of our work are:

- The first large-scale cloth dataset of real-world scans, that contains more than 2k cloths from 2k subjects with various styles and body shapes.
- A 3D cloth generative model, that is able to generate diverse garments with different types, styles and shapes. Design2Cloth is founded on a user-friendly mask encoder that facilitates cloth design from 2D visibility masks. To the best of our knowledge, this is the first cloth generative model trained with over 2000 real-world cloths.
- Being fully differentiable, Design2Cloth can advance the challenging task of 3D garment reconstruction, producing highly detailed 3D clothes. Our cloth reconstruction are far more realistic compared to previous state-of-the-art reconstruction methods, that fail to capture natural cloth creases and produce overly smooth results.
- Leveraging the proposed cloth mask representation, we provide a simple but effective approach for 3D cloth reconstruction from scans. We experimentally show that the proposed method can retrieve high quality 3D garments

from incomplete and corrupted inputs.

2. Related Work

Generative Garment Model for 3D clothes have emerged in recent years due to their impact in fashion design and virtual try-on applications. Early methods approached cloth modeling using Principal Component Analysis (PCA) on fixed topology garments [4, 12, 16, 25, 33, 46]. Nevertheless, such approaches are garment-specific i.e. they can not generalize to clothes with different templates and require training different models for each garment category. To tackle this, Ma *et al.* [30] proposed to model garments as a set of displacements on top of SMPL [29] body topology. Following that, several methods have extended [30], to enrich clothes diversity and train more expressive models [3, 44]. A major limitation, under this setting, is that only a limited number of cloths, that strictly follow the the contour of the body, can be represented as body displacements. Recently, several methods have been proposed to utilize *implicit functions* to train garment agnostic generative models [9, 13, 27], leveraging the property of neural implicit functions to represent topology-agnostic 3D surfaces. SMPLicit [9] pioneered 3D garment generative model using unsigned distance functions (UDF) to represent clothes. Most relevant to our method, DrapeNet [13] has proposed to utilize MeshUDF [17] to achieve better reconstruction results. However, the aforementioned methods rely on low frequency decoders that can only generate overly smooth synthetic data that lack high frequency details, such as wrinkles and creases. In addition, they can only project clothes to the latent space of the model using point clouds and UV maps, which undoubtedly limits their applicability in real-world scenarios. Different from previous methods, we propose a high-fidelity user-friendly generative model, trained on a large scale real-world dataset, that is able to model diverse clothes with high frequency details.

3D Garment Datasets have emerged as crucial resources for modeling and animating clothed human avatars. As summarized in Tab. 1, existing garment datasets can be classified as either *real-world* or *synthetic*. The BUFF dataset [50] was among the first high-quality scanned datasets introduced, albeit characterized by a limited number of subjects and garments. Subsequently, several real-world datasets have emerged, with MGN [4], DeepFashion3D [56], and THUman2.0 [49] standing out as they are all composed from over 100 distinct subjects and garments. Due to the challenges associated with capturing large-scale 3D datasets of clothed humans, research efforts have been focused on the creation of synthetic cloth datasets. Leveraging physics-based simulations [3, 33] and artist-curated patterns [44, 57], synthetic datasets have successfully produced extensive collections of garments featuring diverse shapes,

poses, and styles. However, they suffer from a domain gap, exhibiting overly-smoothed surfaces and non-realistic wrinkles. In this work, we take a significant step towards realistic cloth modeling by curating a large-scale real-world dataset, named DigitalMe, comprising 2010 individual subjects wearing more than 2k garments from 31 categories.

Neural Invertible Skinning. Various approaches have been employed to tackle problems in modeling the deformation of non-rigid and articulated 3D objects such as clothed bodies. In [5, 41, 45, 48], the Linear Blend Skinning (LBS) weights are learned separately in deformed and canonical spaces in order to utilize cycle-consistency losses for establishing correspondence. SNARF [7] proposed to establish the correspondences using an iterative method to invert the LBS equation. Recently, Kant *et al.* [22] proposed a pose-conditioned invertible model, to learn a non-linear blend skinning function that preserves the correspondences with the input mesh. However, all of the aforementioned works rely on overfitting on a sequence of scans and fail to animate independent single scans. Unlike these methods, we show that using the powerful Design2Cloth encoder, we can accurately reconstruct and animate garments from single in-the-wild scans while preserving the intrinsic correspondences.

3. Method

3.1. Large Scale Cloth Dataset

To train a high resolution 3D cloth generative model, we collected a large dataset, namely DigitalMe, comprising of high resolution 3D clothed human scans. The scans were captured using a 3dMD multi-camera structured light stereo system, with 14 cameras and 12 uniform lighting LED panels. We captured a total of 2010 distinct subjects spanning different ages, body types and ethnicities. Participants wore a variety of garments and were scanned in numerous poses, starting from the canonical pose. The raw scans have a resolution of approximately 150,000 vertices.

Using an automated pipeline we fitted parametric SMPL [29] model and extracted the 3D clothes from the raw human scans. Specifically, we rendered each scan from 40 different views that span 360° degrees around the subject. For each view, we utilized MediaPipe algorithm [2] to extract 2D joint locations and subsequently lift them to 3D using linear triangulation. To fit SMPL parametric body model [29] to the 3D joint location, we optimized pose θ and shape parameters β of SMPL, by minimizing the joint error \mathcal{L}_j and the Chamfer Distance (CD) \mathcal{L}_{cd} between the body template and the scan. Additionally, to enforce the SMPL body fittings to lie under the raw scan surface, an additional loss function was employed to penalize the SMPL vertices located outside of the scan. To extract the 3D cloth meshes from the scans we initially segmented each render-

Table 1. Comparison of existing 3D garment methods. Methods that utilize real-world data are highlighted in grey, while the rest are based on synthetic data. #Garment: denotes the number of different individual garments, #Subjects: the number of individual identities, #Categories: the number of distinct cloth types, Generation: the generative ability of the network, Garment-Agnostic: whether the network can represent different cloth categories.

Method	#Garment	#Subjects	#Categories	Generation	Garment-Agnostic
DRAPE [16]	420	-	5	✓	×
DeepGarment [12]	2	-	2	✓	×
Wang <i>et al.</i> [46]	8k	-	3	✓	×
Santesteban <i>et al.</i> [43]	17	-	1	✓	×
GarNet [18]	1.5k	-	3	✓	×
TailorNet [33]	207	-	4	✓	×
Cloth3D [3]	11.3k	-	7	✓	×
Shen <i>et al.</i> [44]	104	-	5	✓	✓
SMPLicit [9]	2M	-	11	✓	✓
DIG [27]	200	-	2	✓	✓
DrapeNet [13]	900	-	5	✓	✓
HOOD [15]	15	-	7	×	✓
Zhao <i>et al.</i> [52]	2	-	2	×	×
CLOTH4D [57]	1k	-	6	×	-
ClothesNet [55]	4.4k	-	11	×	-
BUFF [50]	24	6	4	×	✓
ClothCap [34]	-	-	-	×	×
DeepWrinkles [25]	4	2	2	✓	×
MGN [4]	712	356	5	✓	×
CAPE [30]	8	11	5	✓	✓
DeepFashion3D [56]	563	-	10	×	✓
THUman2.0 [49]	-	200	-	×	✓
LGN [1]	140	140	5	×	×
GarmCap [28]	4	-	3	×	✓
Design2Cloth	2k	2k	31	✓	✓

ing using SAM [23] and then projected the 3D mesh to each one of the renderings. Using a majority vote, between all 40 views, we identify the vertices that correspond to cloth regions. The cloth vertices were then cropped from the raw scan and canonicalized using inverse linear blend skinning function (LBS⁻¹). Note that the canonicalization is performed only in the pose space to retain shape identity of the subject. After preprocessing, 2010 unique cloth meshes were generated from the scans. For additional details about the optimization pipeline we refer the reader to the supplementary material.

To build a user-friendly generative model, we represent each cloth using a 2D visibility mask. To enhance the expressivity of the model and include more cloth details, such as V-neck patterns, we constrain the visibility masks to the frontal part of the cloth by rasterizing only the vertices positioned in front of the Z-axis plane.

3.2. Design2Cloth: An Implicit Garment GAN

The architecture of the proposed model is composed of three main components: the Cloth and Shape Encoders that take as input a binary visibility mask and a target shape code, respectively, and embed them into a latent representation (Sec. 3.2.1), the Implicit Cloth Generator G_t that decodes the latent code to a UDF (Sec. 3.2.2) and the dual-resolution Discriminator D that enforces the generator to produce highly detailed garments (Sec. 3.2.3). Fig. 2 depicts an overview of the proposed method.

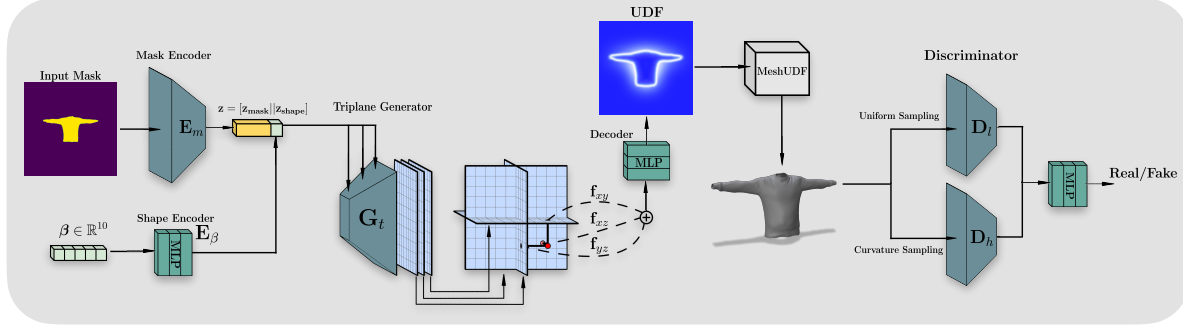


Figure 2. Overview of the proposed Design2Cloth: A binary mask \mathbf{M} along with a shape vector β are fed to the encoder modules, \mathbf{E}_m , \mathbf{E}_β to produce a latent vector \mathbf{z} that is used to drive the triplane generator \mathbf{G}_t . The decoder network takes as input the triplane features and regresses their corresponding unsigned distance function that is then meshed using the differentiable MeshUDF [17]. To enforce the generation of highly detailed clothes we utilize a dual resolution discriminator network \mathcal{D} , that take as input two sparse point clouds sampled from the surface of the generated cloth. The low frequency branch \mathbf{D}_l takes as input a uniformly sampled point cloud whereas the high frequency branch \mathbf{D}_h takes as input a point cloud sampled from coarse areas of the garment surface.

3.2.1 Cloth Encoder

Cloth encoder can be considered as the core component of the proposed method as it learns a compact and smooth latent space which can be then used to condition the generation process and enable interpolation between different clothes. In contrast to previous methods that require complex inputs to project a cloth in the latent space, we propose visibility masks \mathbf{M} as a natural representation of different garments. More specifically, both SMPLicit [9], that uses an occupancy UV image, and DrapeNet [13], that requires a 3D cloth surface to project a cloth in the latent space, apart from being impractical, fail to provide a user-friendly experience. In contrast, visibility masks are an intuitive approach that does not require any expertise knowledge and can turn the proposed method into a scalable tool for user-friendly garment design.

The proposed cloth encoder is divided in two modules the mask encoder \mathbf{E}_m and the shape encoder \mathbf{E}_β . The mask encoder \mathbf{E}_m is based on a lightweight image feature extractor [42] which takes as input a binary mask indicating the cloth style and outputs a compact latent representation of the cloth \mathbf{z}_{mask} . In addition to the clothing style, we use a shape encoder \mathbf{E}_β to condition the garment generation on SMPL shape parameters β . The resulting latent vector \mathbf{z} is given from the concatenation of the two modalities, style and shape, as :

$$\mathbf{z} = [\mathbf{E}_m(\mathbf{M}) || \mathbf{E}_\beta(\beta)] \quad (1)$$

where $||$ denotes the concatenation operator. Using a shape conditioned encoder, our framework learn both style and shape variations of the cloths, in contrast to previous methods that can only generate clothes of the mean shape.

3.2.2 Implicit Cloth Generator

The main objective of the generator \mathcal{G} is to learn a mapping function $\mathcal{G}(\mathbf{z}) \rightarrow C \in \mathbb{R}^{N \times 3}$ from shape and style latent codes \mathbf{z} to a detailed 3D cloth C . Given that cloth registration to a common template is a particularly expensive and challenging task [4], we model clothes as implicit fields.

Inspired from [6], we built a generative model \mathbf{G}_t upon the computationally efficient and expressive hybrid triplane representation. The backbone of the tri-plane generator is a 2D convolutional neural network that outputs three plane feature images $H \times W \times 32$. Using this powerful representation that explicitly learns the features of the grid we can generate detailed clothes without any additional computational requirements. Specifically, using a single multi-layer perceptron (MLP) decoder we can sufficiently predict the *unsigned distance* $d(\mathbf{p})$ of a point \mathbf{p} to the surface of the garment. Finally, to decode the generated UDF to a discrete mesh garment, we utilized MeshUDF [17] since it is fully differentiable and enables the gradient flow in our generator network. During inference, any other decoding method could be used to obtain smoother results [8, 51].

3.2.3 Dual-Resolution Discriminator

To build a high fidelity cloth generative model that is able to reconstruct and generate highly detailed clothes we devise a novel branched discriminator \mathcal{D} . Following [33] that introduced the concept of low and high level cloth generation, we propose an elevated alternative that enforces the generation of wrinkles and high frequency details through adversarial learning. More specifically, instead of using a single branch discriminator to distinguish the generated from the real samples, we divide the workload into two distinct branches that allows the separation of high and low level cloth details. The low-level branch \mathbf{D}_l takes as input a uni-

formly sampled garment and learns structural features of the overall shape. On the contrary, to enforce the generator to preserve the wrinkles of the input, we feed the high-frequency branch \mathbf{D}_h with a point cloud sampled from the surface areas with the maximum Gaussian mean curvature [35]. Both high and low level branches were implemented using PointNet++ [37] encoder that takes as input the sampled point cloud and aggregates their hierarchical features into a compact latent space. The concatenation of the two latent codes are then fed to a trunk network which regresses a real-fake score. Our dual frequency discrimination not only encourages the generated garment to match the distribution of real ones, but also enforces the generation of high frequency details that previous methods failed to model.

Training Objective: To train our method we use a combination of loss functions that enforce the generator \mathcal{G} to accurately predict the UDF $d(\cdot)$ of randomly sampled points in the grid (\mathcal{L}_{UDF}) and generate highly detailed clothes ($\mathcal{L}_{\mathcal{G}}^{\text{adv}}$). Additionally, using the adversarial loss, we train the discriminator network \mathcal{D} to distinguish real from generated clothes ($\mathcal{L}_{\mathcal{D}}^{\text{adv}}$). Finally, to regularize the MeshUDF triangulation, we penalize the gradients of the predicted UDF to match the ground-truth gradients at each sampled point ($\mathcal{L}_{\text{grad}}$). Formally:

$$\begin{aligned} \mathcal{L}_{\text{UDF}} &= \sum_{\mathbf{p}} \|\mathcal{G}(\mathbf{z}; \mathbf{p}) - d(\mathbf{p})\|_2 \\ \mathcal{L}_{\mathcal{G}}^{\text{adv}} &= \mathbb{E}_{\mathbf{z}}[\log(\mathcal{D}(\mathcal{G}(\mathbf{z})))] \\ \mathcal{L}_{\mathcal{D}}^{\text{adv}} &= \mathbb{E}_x[\log \mathcal{D}(x)] + \mathbb{E}_{\mathbf{z}}[\log(1 - \mathcal{D}(\mathcal{G}(\mathbf{z})))] \\ \mathcal{L}_{\text{grad}} &= \sum_{\mathbf{p}} \|\hat{\mathbf{g}}_{\mathbf{p}} - \mathbf{g}_{\mathbf{p}}\|_2 \end{aligned} \quad (2)$$

where \mathbf{z} denotes a latent code produced from the encoder, $d(\mathbf{p})$ is the ground truth unsigned distance and $\mathbf{g}_{\mathbf{p}}, \hat{\mathbf{g}}_{\mathbf{p}}$ the ground truth and the predicted gradients of the UDF at a point \mathbf{p} . For notation purposes we omit the sampling function used in the discriminator network.

4. Experiments

In this section, we conduct a comprehensive evaluation of the capabilities of the proposed Design2Cloth network to generate, interpolate, and fit ‘in-the-wild’ clothes.

4.1. Cloth Generation

To compare the performance of the proposed method on cloth generation, we trained Design2Cloth and DrapeNet [13] models on DigitalMe dataset with 100 additional clothes from Cloth3D [3] ClothesNet [55], using a common 80%-20% train-test split. In Tab. 2, we report the reconstruction quality of each method trained only on Cloth3D and DigitalMe datasets measured in terms of Chamfer Distance (CD) and Normal Consistency (NC) metrics. As expected, the performance of both models drop on DigitalMe

Table 2. Quantitative comparison of the reconstruction performance of the Proposed and the DrapeNet [13] on the Cloth3D [3] and DigitalMe datasets.

Method	Cloth3D		DigitalMe	
	CD ($\times 10^{-4}$) ↓	NC ↑	CD ($\times 10^{-2}$) ↓	NC ↑
DrapeNet	0.36	0.97	0.56	0.96
Proposed	0.18	0.99	0.12	0.98

dataset which contains more challenging clothes, with high frequency wrinkles and details. Despite the challenges, the proposed method not only outperforms DrapeNet under all metrics and datasets, but also manages to achieve 0.98 NC performance on DigitalMe dataset, which validates the ability of our generator to model wrinkles and creases. The superiority of the proposed method can also be qualitatively verified in Fig. 3, where the proposed method can accurately generate highly detailed clothes compared to the smooth garments produced from DrapeNet. Furthermore, the proposed method inherits many of the desirable generative network properties such as a well-behaved latent space with smooth interpolation between diverse clothes. Specifically, as depicted in Fig. 4, Design2Cloth can not only smoothly interpolate between diverse garments, i.e. from an open jacket to a dress, but also transition from top to bottom clothes. Additionally, given that the proposed method is also conditioned on the cloth shape parameter, one can also interpolate between different shapes while retaining the style fixed.

Finally, to evaluate the realism of the generated garments, we designed an online user study where 50 participants, including 20 people working in the fashion industry, were asked to assess the quality of the generations. Each participant was asked to score, on a scale from 1 (poor) to 10 (very realistic), the quality of the generated garments from the proposed and DrapeNet methods, along with the ground truth meshes from DigitalMe and Cloth3D datasets. As shown in Fig. 5 (Left), DigitalMe clothes proved to be the most realistic, achieving an average score of 8.7. Interestingly, the highly detailed clothes generated by the proposed method manages to produce more realistic clothes compared to the ground truth Cloth3D dataset and DrapeNet generations, achieving an average score of 7.2 compared to 7 and 4.6 respectively, which validates the generative power of our model.

4.2. 3D Garment Reconstruction In-The-Wild

Given that the proposed model is fully differentiable it can be used to fit 3D clothes from in-the-wild images. The first step in recovering cloth from an in-the-wild image is to estimate SMPL pose θ and shape β parameters using an off-the-shelf pose estimation method [39]. Using the detected

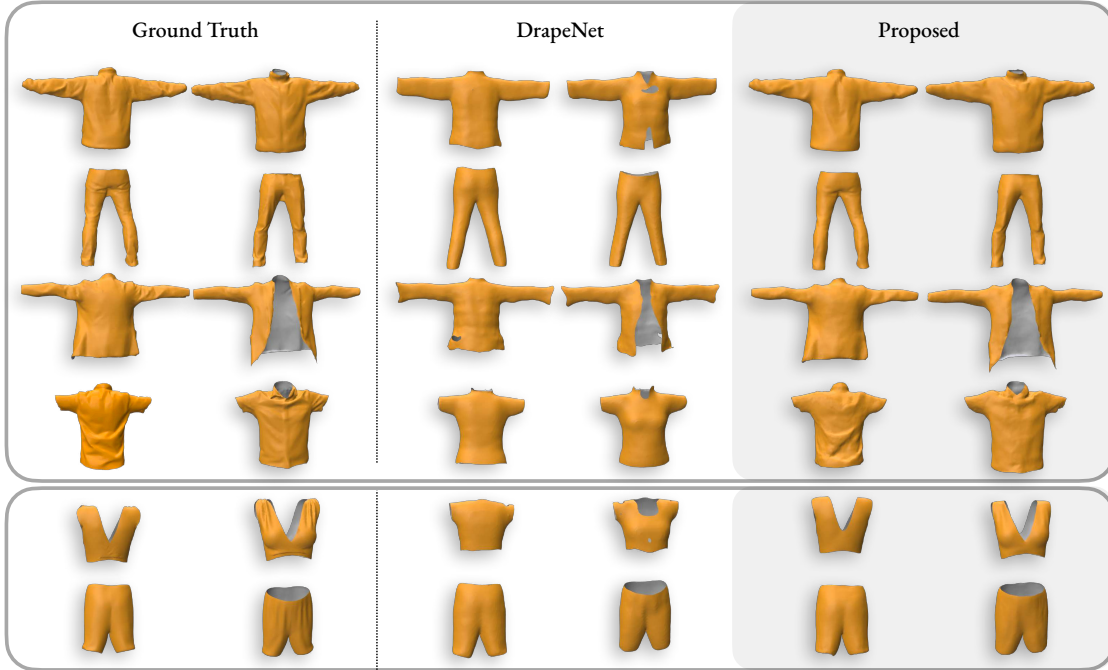


Figure 3. Qualitative comparison between the proposed and DrapeNet methods of reconstruction performance on DigitalMe (Top), Cloth3D [3] and ClothesNet [55] (Bottom) datasets.

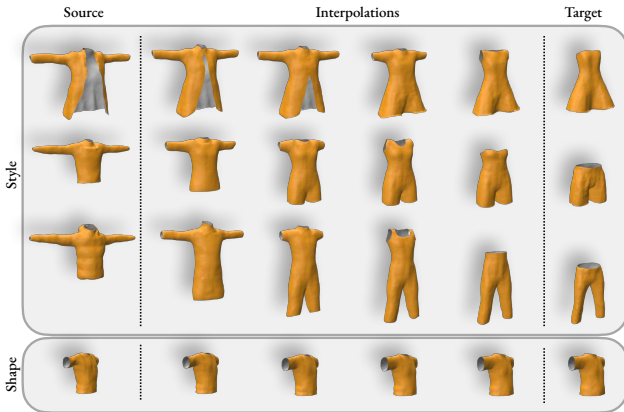


Figure 4. Interpolation between source and target styles (Top) and shapes (Bottom). The proposed method can interpolate between shapes and styles, generating realistic intermediate clothes.

SMPL parameters, we posed the generated cloth and projected it to the image space to obtain a posed mask. We optimized the latent code \mathbf{z} by minimizing the Intersection over Union (IoU) between image cloth segmentation \mathbf{S} extracted using SAM[23] and the projected posed mask.

To facilitate smoother convergence of the model we calculate the latent code statistics of the training set and we initialize \mathbf{z} with the mean value of the dataset. Mathemati-

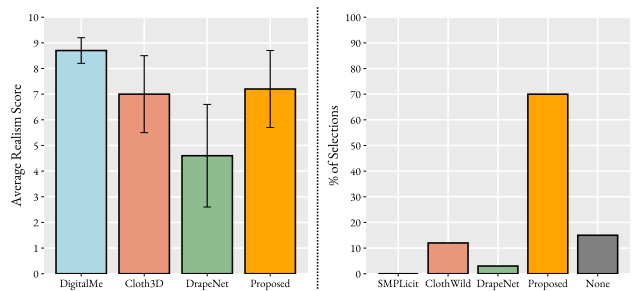


Figure 5. Human evaluation results. Left: Average realism scores of the generated and the ground truth data. Right: Perceptual Evaluation of 3D reconstructions from in-the-wild images between the proposed and the baseline methods.

cally, the optimization scheme can be formulated as:

$$\min_{\mathbf{z}} \mathcal{L}_{\text{IoU}}(\Pi_{\mathbf{M}}(\text{LBS}(\mathcal{G}(\mathbf{z}; \beta), \theta)), \mathbf{S}) + \mathcal{L}_{\text{prior}}(\mathbf{z}) \quad (3)$$

where $\Pi_{\mathbf{M}}$ is the 2D occupancy mask rasterized from a differentiable renderer, \mathbf{S} the extracted cloth mask, $\mathcal{G}(\mathbf{z}; \beta)$ our cloth generator that maps a latent code \mathbf{z} with shape β to a 3D cloth $C \in \mathbb{R}^{N \times 3}$, $\text{LBS}(\cdot, \theta)$ a skinning function that poses a mesh to the target pose θ and $\mathcal{L}_{\text{prior}}(\mathbf{z})$ an \mathcal{L}_2 regularization that constrains \mathbf{z} to feasible values. To compare the reconstruction performance of the proposed method, we utilized three state-of-the-art 3D cloth reconstruction models, namely SMPLicit [9], ClothWild [32] and

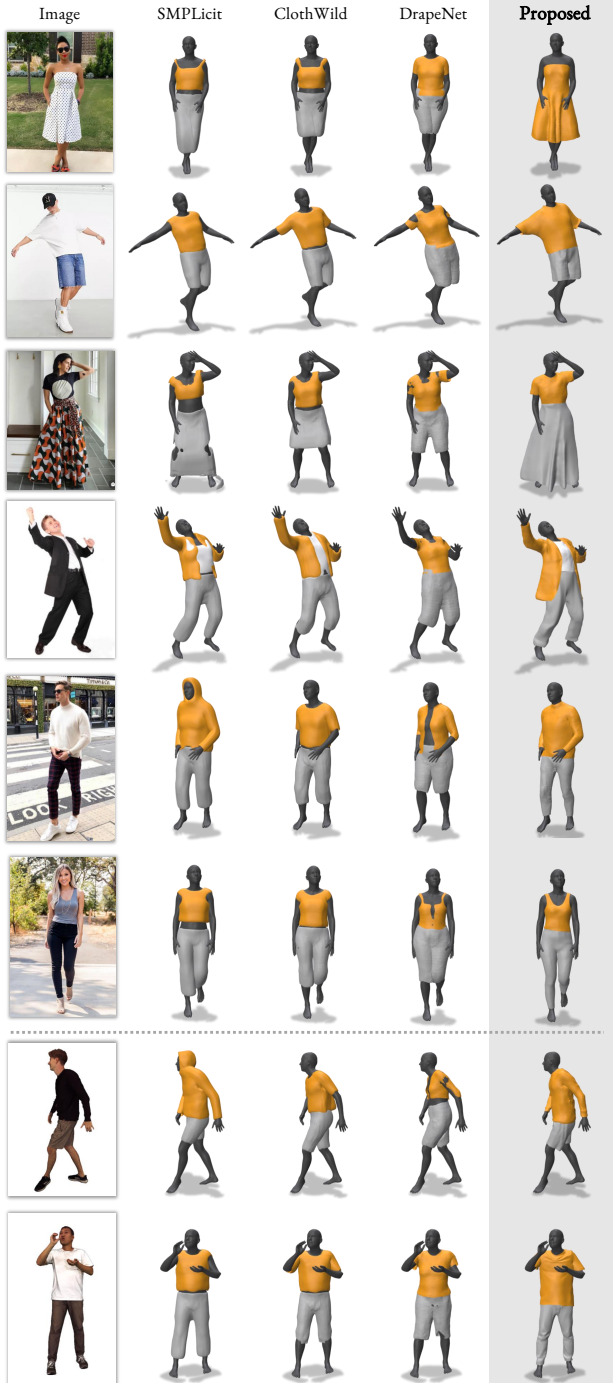


Figure 6. Garment reconstruction from in-the-wild images. Qualitative comparison between the proposed and the baseline method. The proposed method can generate highly detailed clothes compared to the smooth reconstructions of the baseline methods. Figure better viewed with zoom. The last two rows of the figure correspond to renderings from CustomHumans [20] dataset.

DrapeNet [13]. In Fig. 6 we qualitatively compare the proposed and the baseline method in 3D garment reconstruction

Table 3. Cloth Reconstruction Error from single image on CustomHumans Dataset [20].

Method	SMPLicit	ClothWild	DrapeNet	Proposed
CD ($\times 10^{-2}$) ↓	0.40	0.62	1.14	0.12

tion from in-the-wild images. Importantly, in contrast to the baseline methods, the proposed method requires optimization of a single network for both top and bottom garments, which significantly reduces the memory and computational requirements. As can be observed, Design2Cloth can reconstruct detailed garments, with intricate creases and puckering of fabric compared to the smooth reconstructions of the baselines. To quantitatively validate the perceptual reconstruction performance of the proposed method, we have included a second section in the human evaluation study described in Sec. 4.1. This time the participants were given the input image along with the reconstructions from the proposed and the baseline methods, and were asked to choose the one that has the most similarities with the input image. To avoid biases in the evaluation process, users were also given a ‘None’ option. As can be seen in Fig. 5, more than 70% of the participants selected the proposed method as the one achieving the most accurate and realistic reconstructions, compared to less than 15% that selected ClothWild method.

To further quantitatively evaluate the cloth reconstruction performance of the proposed method, we utilized CustomHumans dataset [20] that contains high resolution 3D scans of dressed humans in various poses along with their corresponding SMPL fittings. Specifically, we rendered the posed scans in arbitrary views and measured the CD between the ground truth and the reconstructed garments using the optimization pipeline described above. For a fair comparison, we use the ground truth SMPL parameters. As shown in Tab. 3, the proposed method outperforms the baseline methods in CD reconstruction error by a large margin, which can be also validated from the last two rows of Fig. 6.

4.3. Reconstruction and Animation of Raw Scans

In addition to its user-friendly nature, the proposed mask cloth representation facilitates the reconstruction of corrupted cloth data. As an application case we examined the case of reconstructing and animating raw scans with partial observations. In particular, using a simple LBS we canonicalized posed scans from the CustomHumans dataset [20] and rendered them to obtain their corresponding 2D visibility masks. The visibility masks are then mapped to the real garment distribution using the Mask Encoder E_m . It is noteworthy that, even though both the masks and the canonicalized garments may contain holes and irregular topologies, the cloth reconstructions generated using Design2Cloth are devoid of such artifacts. Finally, the re-

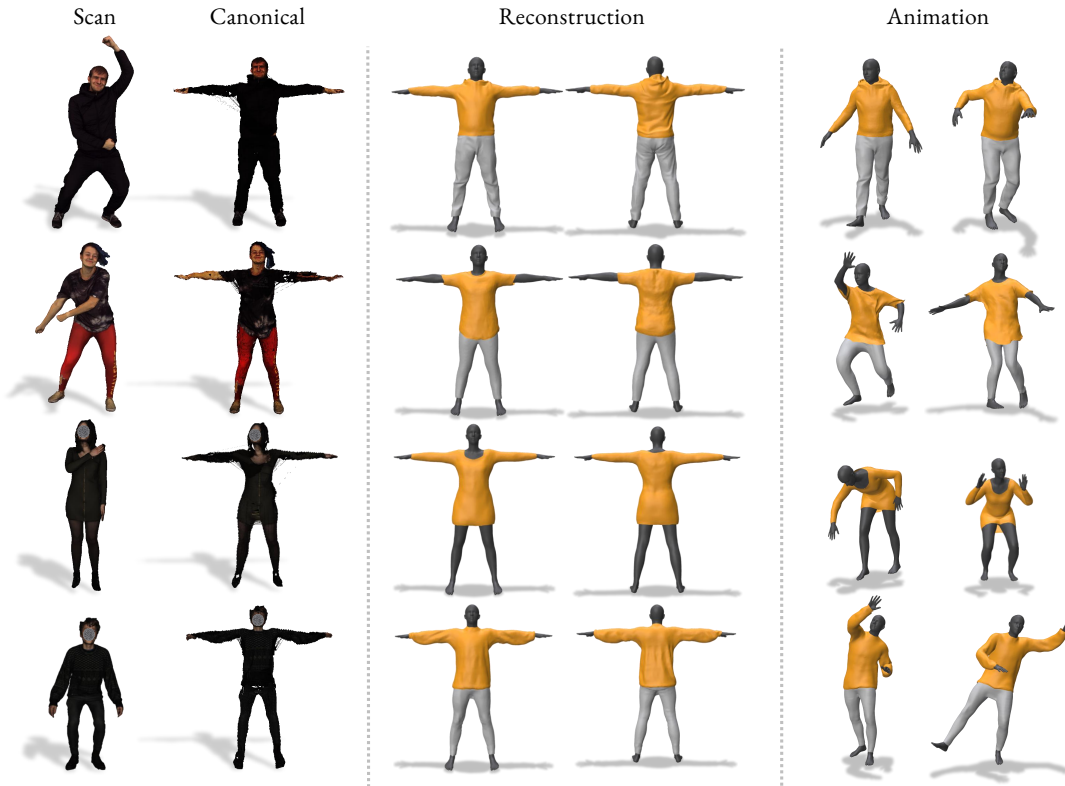


Figure 7. Garment reconstruction from 3D scans. We propose a pipeline to extract clothes from a posed scan, reconstruct them using Design2Cloth and animate them. Using Mask Encoder as a cloth prior we are able to reconstruct corrupted garments with holes and irregular structure.

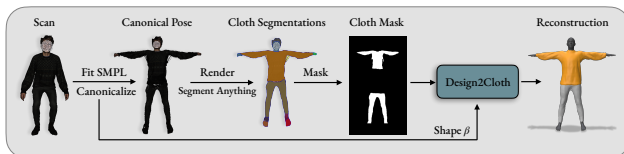


Figure 8. The proposed pipeline used to reconstruct clothes from corrupted 3D scans.

constructed meshes can then be animated using any off-the-shelf draping method [13, 15]. The overall 3D scan reconstruction and animation pipeline is depicted in Fig. 8. Fig. 7 illustrates the garment reconstructions and animations on CustomHumans (top two rows) and DigitalMe (bottom two rows) datasets. The proposed method remains unaffected by the artifacts created during the inverse blend skinning and can accurately reconstruct the 3D garments from the corrupted canonical scans. Under this setting, Design2Cloth can be used to restore and rectify garments from corrupted and occluded clothed human scans.

5. Conclusion

In this study we introduce Design2Cloth, a high fidelity 3D garment generative model that significantly develops current state-of-the-art generative models. Trained on a collected large scale real-world dataset, composed of more than 2,000 garments from 2,010 distinct identities, the proposed method is able to produce highly diverse and realistic garments. Significantly, in contrast to previous methods, Design2Cloth provides a user-friendly technique for garment design from simple 2D masks. In addition, the proposed pipeline is fully differentiable, providing a plug-and-play solution to several inverse problems, including cloth reconstruction from single images and scans. We believe that Design2Cloth can not only aid garment learning using synthetic clothes sampled from latent space, but also abet the realm of 3D design.

Acknowledgements. S. Zafeiriou was supported by EPSRC Project DEFORM (EP/S010203/1) and GNOMON (EP/X011364). R.A. Potamias was supported by EPSRC Project GNOMON (EP/X011364).

References

- [1] Alakh Aggarwal, Jikai Wang, Steven Hogue, Saifeng Ni, Madhukar Budagavi, and Xiaohu Guo. Layered-garment net: Generating multiple implicit garment layers from a single image. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 3000–3017, 2022.
- [2] Valentin Bazarevsky, Ivan Grishchenko, Karthik Raveendran, Tyler Zhu, Fan Zhang, and Matthias Grundmann. Blazepose: On-device real-time body pose tracking, 2020.
- [3] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Cloth3d: clothed 3d humans. In *European Conference on Computer Vision*, pages 344–359. Springer, 2020.
- [4] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5420–5430, 2019.
- [5] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10754–10764, 2021.
- [6] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16123–16133, 2022.
- [7] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11594–11604, 2021.
- [8] Zhiqin Chen, Andrea Tagliasacchi, Thomas Funkhouser, and Hao Zhang. Neural dual contouring. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022.
- [9] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11875–11885, 2021.
- [10] Aiyu Cui, Daniel McKee, and Svetlana Lazebnik. Dressing in order: Recurrent person image generation for pose transfer, virtual try-on and outfit editing. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 14638–14647, 2021.
- [11] Yi Rui Cui, Qi Liu, Cheng Ying Gao, and Zhongbo Su. Fashiongan: display your fashion design using conditional generative adversarial nets. In *Computer Graphics Forum*, pages 109–119. Wiley Online Library, 2018.
- [12] R Daněřek, Endri Dibra, Cengiz Öztireli, Remo Ziegler, and Markus Gross. Deepgarment: 3d garment shape estimation from a single image. In *Computer Graphics Forum*, pages 269–280. Wiley Online Library, 2017.
- [13] Luca De Luigi, Ren Li, Benoît Guillard, Mathieu Salzmann, and Pascal Fua. Drapenet: Garment generation and self-supervised draping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1451–1460, 2023.
- [14] Valentin Gabeur, Jean-Sébastien Franco, Xavier Martin, Cordelia Schmid, and Gregory Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2232–2241, 2019.
- [15] Artur Grigorev, Michael J. Black, and Otmar Hilliges. Hood: Hierarchical graphs for generalized modelling of clothing dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16965–16974, 2023.
- [16] Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. Drape: Dressing any person. *ACM Transactions on Graphics (ToG)*, 31(4):1–10, 2012.
- [17] Benoît Guillard, Federico Stella, and Pascal Fua. Meshudf: Fast and differentiable meshing of unsigned distance field networks. In *European Conference on Computer Vision*, pages 576–592. Springer, 2022.
- [18] Erhan Gundogdu, Victor Constantin, Amrollah Seifoddini, Minh Dang, Mathieu Salzmann, and Pascal Fua. Garnet: A two-stream network for fast and accurate 3d cloth draping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019.
- [19] Xintong Han, Zuxuan Wu, Zhe Wu, Ruichi Yu, and Larry S Davis. Viton: An image-based virtual try-on network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7543–7552, 2018.
- [20] Hsuan-I Ho, Lixin Xue, Jie Song, and Otmar Hilliges. Learning locally editable virtual humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21024–21035, 2023.
- [21] Shion Honda. Viton-gan: Virtual try-on image generator trained with adversarial loss. *arXiv preprint arXiv:1911.07926*, 2019.
- [22] Yash Kant, Aliaksandr Siarohin, Riza Alp Guler, Menglei Chai, Jian Ren, Sergey Tulyakov, and Igor Gilitschenski. Invertible neural skinning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8715–8725, 2023.
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- [24] Maria Korosteleva and Sung-Hee Lee. Neuraltailor: Reconstructing sewing pattern structures from 3d point clouds of garments. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022.
- [25] Zorah Lahner, Daniel Cremers, and Tony Tung. Deepwrinkles: Accurate and realistic clothing modeling. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [26] Hanbit Lee, Jinseok Seol, and Sang-goo Lee. Style2vec: Representation learning for fashion items from style sets. *arXiv preprint arXiv:1708.04014*, 2017.
- [27] Ren Li, Benoît Guillard, Edoardo Remelli, and Pascal Fua. Dig: Draping implicit garment over the human body. In

- Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 2780–2795, 2022.
- [28] Siyou Lin, Boyao Zhou, Zerong Zheng, Hongwen Zhang, and Yebin Liu. Leveraging intrinsic properties for non-rigid garment alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14485–14496, 2023.
- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015.
- [30] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to dress 3d people in generative clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [31] Yifang Men, Yiming Mao, Yuning Jiang, Wei-Ying Ma, and Zhouhui Lian. Controllable person image synthesis with attribute-decomposed gan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5084–5093, 2020.
- [32] Gyeongsik Moon, Hyeongjin Nam, Takaaki Shiratori, and Kyoung Mu Lee. 3d clothed human reconstruction in the wild. In *European Conference on Computer Vision (ECCV)*, 2022.
- [33] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [34] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (ToG)*, 36(4):1–15, 2017.
- [35] Rolandos Alexandros Potamias, Giorgos Bouritsas, and Stefanos Zafeiriou. Revisiting point cloud simplification: A learnable feature preserving approach. In *European Conference on Computer Vision*, pages 586–603. Springer, 2022.
- [36] Rolandos Alexandros Potamias, Stylianos Ploumpis, Stylianos Moschoglou, Vasileios Triantafyllou, and Stefanos Zafeiriou. Handy: Towards a high fidelity 3d hand shape and appearance model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4670–4680, 2023.
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [38] Amir Hossein Raffee and Michael Sollami. Garmentgan: Photo-realistic adversarial fashion transfer. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 3923–3930. IEEE, 2021.
- [39] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. Frankmocap: A monocular 3d whole-body pose estimation system via regression and integration. In *IEEE International Conference on Computer Vision Workshops*, 2021.
- [40] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2304–2314, 2019.
- [41] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2886–2897, 2021.
- [42] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [43] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Learning-based animation of clothing for virtual try-on. In *Computer Graphics Forum*, pages 355–366. Wiley Online Library, 2019.
- [44] Yu Shen, Junbang Liang, and Ming C Lin. Gan-based garment generation using sewing pattern images. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020. Proceedings, Part XVIII 16*, pages 225–247. Springer, 2020.
- [45] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. Metaavatar: Learning animatable clothed human models from few depth images. *Advances in Neural Information Processing Systems*, 34:2810–2822, 2021.
- [46] Tuanfeng Y Wang, Duygu Ceylan, Jovan Popovic, and Niloy J Mitra. Learning a shared shape space for multimodal garment design. *arXiv preprint arXiv:1806.11335*, 2018.
- [47] Qiang Wu, Baixue Zhu, Binbin Yong, Yongqiang Wei, Xuetao Jiang, Rui Zhou, and Qingguo Zhou. Clothgan: generation of fashionable dunhuang clothes using generative adversarial networks. *Connection Science*, 33(2):341–358, 2021.
- [48] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2863–2873, 2022.
- [49] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5746–5756, 2021.
- [50] Chao Zhang, Sergi Pujades, Michael J. Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3d scan sequences. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [51] Congyi Zhang, Guying Lin, Lei Yang, Xin Li, Taku Komura, Scott Schaefer, John Keyser, and Wenping Wang. Surface extraction from neural unsigned distance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 0000–0000, 2023.
- [52] Fang Zhao, Zekun Li, Shaoli Huang, Junwu Weng, Tianfei Zhou, Guo-Sen Xie, Jue Wang, and Ying Shan. Learning anchor transformations for 3d garment animation. In *Pro-*

- ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 491–500, 2023.
- [53] Jiali Zheng, Youngkyoon Jang, Athanasios Papaioannou, Christos Kampouris, Rolandos Alexandros Potamias, Foivos Paraperas Papantoniou, Efstathios Galanakis, Aleš Leonardis, and Stefanos Zafeiriou. IiSh: The imperial light-stage head dataset for human head view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1112–1120, 2023.
- [54] 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*, 44(6): 3170–3184, 2021.
- [55] Bingyang Zhou, Haoyu Zhou, Tianhai Liang, Qiaojun Yu, Siheng Zhao, Yuwei Zeng, Jun Lv, Siyuan Luo, Qiancai Wang, Xinyuan Yu, et al. Clothesnet: An information-rich 3d garment model repository with simulated clothes environment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20428–20438, 2023.
- [56] Heming Zhu, Yu Cao, Hang Jin, Weikai Chen, Dong Du, Zhangye Wang, Shuguang Cui, and Xiaoguang Han. Deep fashion3d: A dataset and benchmark for 3d garment reconstruction from single images. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 512–530. Springer, 2020.
- [57] Xingxing Zou, Xintong Han, and Waikeng Wong. Cloth4d: A dataset for clothed human reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12847–12857, 2023.