

Point-Based Modeling of Human Clothing

Ilya Zakharkin^{1,2*}, Kirill Mazur^{1*}, Artur Grigorev¹, Victor Lempitsky^{1,2}

¹ Samsung AI Center, Moscow

² Skolkovo Institute of Science and Technology (Skoltech), Moscow



Figure 1: Our approach models the geometry of diverse clothing outfits using point clouds (top row; random point colors). The point clouds are obtained by passing the SMPL meshes (shown in grey) and latent outfit code vectors through a pretrained deep network. Additionally, our approach can model clothing appearance using *neural point-based graphics* (bottom row). The outfit appearance can be captured from a video sequence, while a single frame is sufficient for point-based geometric modeling.

Abstract

We propose a new approach to human clothing modeling based on point clouds. Within this approach, we learn a deep model that can predict point clouds of various outfits, for various human poses, and for various human body shapes. Notably, outfits of various types and topologies can be handled by the same model. Using the learned model, we can infer the geometry of new outfits from as little as a single image, and perform outfit retargeting to new bodies in new poses. We complement our geometric model with appearance modeling that uses the point cloud geometry

as a geometric scaffolding and employs neural point-based graphics to capture outfit appearance from videos and to re-render the captured outfits. We validate both geometric modeling and appearance modeling aspects of the proposed approach against recently proposed methods and establish the viability of point-based clothing modeling.

1. Introduction

Modeling realistic clothing is a big part of the overarching task of realistic modeling of humans in 3D. Its immedi-

ate practical applications include virtual clothing try-on as well as enhancing the realism of human avatars for telepresence systems. Modeling clothing is difficult since outfits have wide variations in geometry (including topological changes) and in appearance (including wide variability of textile patterns, prints, as well as complex cloth reflectance). Modeling interaction between clothing outfits and human bodies is an especially daunting task.

In this work, we propose a new approach to modeling clothing (Figure 1) based on point clouds. Using a recently introduced synthetic dataset [7] of simulated clothing, we learn a joint geometric model of diverse human clothing outfits. The model describes a particular outfit with a latent code vector (the *outfit code*). For a given outfit code and a given human body geometry (for which we use the most popular SMPL format [34]), a deep neural network (the *draping network*) then predicts the point cloud that approximates the outfit geometry draped over the body.

The key advantage of our model is its ability to reproduce diverse outfits with varying topology using a single latent space of outfit codes and a single draping network. This is made possible because of the choice of the point cloud representation and the use of topology-independent, point cloud-specific losses during the learning of the joint model. After learning, the model is capable of generalizing to new outfits, capturing their geometry from data, and to drape the acquired outfits over bodies of varying shapes and in new poses. With our model, acquiring the outfit geometry can be done from as little as a single image.

We extend our approach beyond geometry acquisition to include the appearance modeling. Here, we use the ideas of differentiable rendering [36, 51, 31] and neural point-based graphics [2, 40, 60]. Given a video sequence of an outfit worn by a person, we capture the photometric properties of the outfit using neural descriptors attached to points in the point cloud, and the parameters of a rendering (decoder) network. The fitting of the neural point descriptors and the rendering network (which capture the photometric properties) is performed jointly with the estimation of the outfit code (which captures the outfit geometry) within the same optimization process. After fitting, the outfit can be transferred and re-rendered in a realistic way over new bodies and in new poses.

In the experiments, we evaluate the ability of our geometric model to capture the deformable geometry of new outfits using point clouds. We further test the capability of our full approach to capture both outfit geometry and appearance from videos and to re-render the learned outfits to new targets. The experimental comparisons show the viability of the point-based approach to clothing modeling. We will publish our code and model at <https://saic-violet.github.io/point-based-clothing/>.

2. Related work on clothing modeling

Modeling clothing geometry. Many existing methods model clothing geometry using one or several pre-defined garment templates of fixed topology. DRAPE [1], which is one of the earlier works, learns from Physics-based simulation (PBS) and allows for pose and shape variation for each learned garment mesh. Newer works usually represent garment templates in the form of offsets (displacements) to SMPL [35] mesh. ClothCap [48] employs such a technique and captures more fine-grained details learned from the new dataset of 4D scans. DeepWrinkles [30] also addresses the problem of fine-grained wrinkles modeling with the use of normal maps generated by a conditional GAN. GarNet [15] incorporates two-stream architecture and makes it possible to simulate garment meshes at the level of realism that almost matches PBS, while being two orders of magnitude faster. TailorNet [46] follows the same SMPL-based template approach as [48, 8] but models the garment deformations as a function of pose, shape and style simultaneously (unlike the previous work). It also shows greater inference speed than [15]. The CAPE system [38] uses graph ConvNet-based generative shape model that enables to condition, sample, and preserve fine shape detail in 3D meshes.

Several other works recover clothing geometry simultaneously with the full body mesh from image data. BodyNet [57] and DeepHuman [66] are voxel-based methods that directly infer the volumetric dressed body shape from a single image. In SiCloPe [44] the authors use similar approach, but synthesize the silhouettes of the subjects in order to recover more details. HMR [26] utilizes SMPL body model to estimate pose and shape from an input image. Some approaches such as PIFu [53] and ARCH [19] employ end-to-end implicit functions for clothed human 3D reconstruction and are able to generalise to complex clothing and hair topology, while PIFuHD [54] recovers higher resolution 3D surface by using two-level architecture. However, these SDF approaches can only represent closed connected surfaces, whereas point clouds may represent arbitrary topologies.

MouldingHumans [12] predicts the final surface from estimated “visible” and “hidden” depth maps. MonoClothCap [61] demonstrates promising results in video-based temporally coherent dynamic clothing deformation modeling. Most recently, Yoon et al. [64] design relatively simple yet effective pipeline for template-based garment mesh retargeting.

Our geometric modeling differs from previous works through the use of a different representations (point clouds), which gives our approach topological flexibility, the ability to model clothing separately from the body, while also providing the geometric scaffold for appearance modeling with neural rendering.

Modeling clothing appearance. A large number of work focus on direct image-to-image transfer of clothing bypassing 3D modeling. Thus, [23, 16, 58, 62, 21] address the task of transferring a desired clothing item onto the corresponding region of a person given their images. CAGAN [23] is one of the first works that proposed to utilize image-to-image conditional GAN to tackle this task. VITON [16] follows the idea of image generation and uses non-parametric geometric transform which makes all the procedure two-stage, similar to SwapNet [50] with the difference in the task statement and training data. CP-VTON [58] further improves upon [16] by incorporating a full learnable thin-plate spline transformation, followed by CP-VTON+ [42], LAVITON [22], Ayush et al. [6] and ACGPN [62]. While the above-mentioned works rely on pre-trained human parsers and pose estimators, the recent work of Issenhuth et al. [21] achieves competitive image quality and significant speed-up by employing a teacher-student setting to distill the standard virtual try-on pipeline. The resulting student network does not invoke an expensive human parsing network at inference time. Very recently introduced VOGUE [32] train a pose-conditioned StyleGAN2 [28] and find the optimal combination of latent codes to produce high-quality try-on images.

Some methods make use of both 2D and 3D information for model training and inference. Cloth-VTON [41] employs 3D-based warping to realistically retarget a 2D clothing template. Pix2Surf [43] allows to digitally map the texture of online retail store clothing images to the 3D surface of virtual garment items enabling 3D virtual try-on in real-time. Other relevant research extend the scenario of single template cloth retargeting to multi-garment dressing with unpaired data [45], generating high-resolution fashion model images wearing custom outfits [63], or editing the style of a person in the input image [17].

In contrast to the referenced approaches to clothing appearance retargeting, ours uses explicit 3D geometric models, while not relying on individual templates of fixed topology. On the downside, our appearance modeling part requires video sequence, while some of the referenced works use one or few images.

Joint modeling of geometry and appearance. Octopus [3] and Multi-Garment Net (MGN) [8] recover the textured clothed body mesh based on the SMPL+D model. The latter method treats clothing meshes separately from the body mesh, which gives it the ability to transfer the outfit to another subject. Tex2Shape [5] proposes an interesting framework that turns the shape regression task into an image-to-image translation problem. In [55], a learning-based parametric generative model is introduced that can support any type of garment material, body shape, and most garment topologies. Very recently, StylePeople [20] ap-

proach integrates polygonal body mesh modeling with neural rendering, so that both clothing geometry and the texture are encoded in the neural texture [56]. Similarly to [20] our approach to appearance modeling also relies on neural rendering, however our handling of geometry is more explicit. In the experiments, we compare to [20] and observe the advantage of a more explicit geometric modeling, especially for loose clothing.

Finally, we note that in parallel with us, the SCALE system [37] explored very similar ideas (point-based geometry modeling and its combination with neural rendering) for modeling clothed humans from 3D scans.

3. Method

We first discuss the point cloud draping model. The goal of this model is to capture the geometry of diverse human outfits draped over human bodies with diverse shapes and poses using point clouds. We propose a latent model for such point clouds that can be fitted to a single image or to more complete data. We then describe the combination of the point cloud draping with neural rendering that allows us to capture the appearance of outfits from videos.

3.1. Point cloud draping

Learning the model. We learn the model using generative latent optimization (GLO) [10]. We assume that the training set has a set of N outfits, and associate each outfit with d -dimensional vector z (the *outfit code*). We thus randomly initialize $\{z_1, \dots, z_N\}$, where $z_i \in Z \subseteq \mathbb{R}^d$ for all $i = 1, \dots, N$.

During training, for each outfit, we observe its shape for a diverse set of human poses. The target shapes are given by a set of geometries. In our case, we use synthetic CLOTH3D dataset [7] that provides shapes in the form of meshes of varying topology. In this dataset, each subject is wearing an outfit and performs a sequence of movements. For each outfit i for each frame j in the corresponding sequence, we sample points from the mesh of this outfit and obtain the point cloud $x_i^j \in X$, where X denotes the space of point clouds of a fixed size (8192 is used in our experiments). We denote the length of the training sequence of the i -th outfit as P_i . We also assume that the body mesh $s_i^j \in S$ is given, and in our experiments we work with the SMPL [34] mesh format (thus S denotes the space of SMPL meshes for varying body shape parameters and body pose parameters). Putting it all together, we obtain the dataset $\{(z_i, s_i^j, x_i^j)\}_{i=1..N, j=1..P_i}$ of outfit codes, SMPL meshes, and clothing point clouds.

As our goal is to learn to predict the geometry in new poses and for new body shapes, we introduce the *draping function* $G_\theta : Z \times S \rightarrow X$ that maps the latent code and the SMPL mesh (representing the naked body) to the outfit point cloud. Here, θ denotes the learnable parameters of the

function. We then perform learning by the optimization of the following objective:

$$\min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N \frac{1}{P_i} \sum_{j=1}^{P_i} L_{3D} \left(G_{\theta}(z_i, s_i^j), x_i^j \right) \quad (1)$$

In (1), the objective is the mean reconstruction loss for the training point clouds over the training set. The loss L_{3D} is thus the 3D reconstruction loss. In our experiments, we use the approximate Earth Mover’s Distance [33]. Note, that as this loss measures the distance between point clouds and ignores all topological properties, our learning formulation is naturally suitable for learning outfits of diverse topology.

We perform optimization jointly over the parameters of our draping function G_{θ} and over the latent outfit code z_i for all $i = 1, \dots, N$. Following [10], to regularize the process, we clip the outfit codes to the unit ball during optimization. The optimization process thus establishes the outfit latent code space and the parameters of the draping function.

Draping network. We implement the draping function $G_{\theta}(z, s)$ as a neural network that takes the SMPL mesh s and transforms this point cloud into the outfit point cloud. Over the last years, point clouds have become (almost) first-class citizens in the deep learning world, as a number of architectures that can input and/or output point clouds and operate on them have been proposed. In our work, we use the recently introduced *Cloud Transformer* architecture [39] due to its capability to handle diverse point cloud processing tasks.

The cloud transformer comprises of blocks, each of which sequentially rasterizes, convolves, and de-rasterizes the point cloud at the learned data-dependent positions. The cloud transformer thus deforms the input point cloud (derived from the SMPL mesh as discussed below) into the output point cloud x over a number of blocks. We use a simplified version of the cloud transformer with single-headed blocks to reduce the computational complexity and memory requirements. Otherwise, we follow the architecture of the generator suggested in [39] for image-based shape reconstruction, which in their case takes the point cloud (sampled from the unit sphere) and a vector (computed by the image encoding network) as an input and outputs the point cloud of the shape depicted in the image.

In our case, the input point cloud and the vector are different and correspond to the SMPL mesh and the outfit code respectively. More specifically, to input the SMPL mesh into the cloud transformer architecture, we first remove the parts of the mesh corresponding to the head, the feet and the hands. We then consider the remaining vertices as a point cloud. To densify this point cloud, we also add the mid-points of the SMPL mesh edges to this point cloud. The resulting point cloud (which is shaped by the SMPL mesh

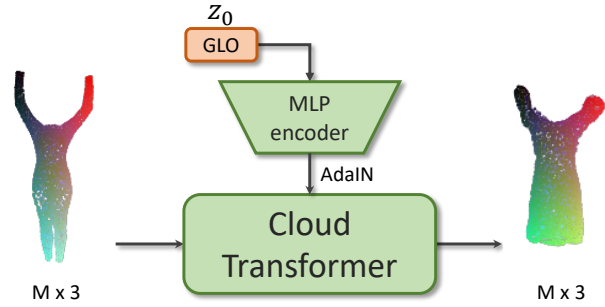


Figure 2: Our draping networks morphs the body point cloud (left) and the outfit code (top) into the outfit point cloud that is adapted to the body pose and the body shape.

and reflects the change of pose and shape) is input into the cloud transformer.

Following [39], the latent outfit code z is input into the cloud transformer through AdaIn connections [18] that modulate the convolutional maps inside the rasterization-derasterization blocks. The particular weights and biases for each AdaIn connection are predicted from the latent code z via a perceptron, as is common for style-based generators [27]. We note that while we have obtained good results using the (simplified) cloud transformer architecture, other deep learning architectures that operate on point clouds (e.g. PointNet [49]) can be employed.

We also note that the morphing implemented by the draping network is strongly non-local (i.e. our model does not simply compute local vertex displacements), and is consistent across outfits and poses (Figure 3).

Estimating the outfit code. Once the draping network is pre-trained on a large synthetic dataset [7], we are able to model the geometry of a previously unseen outfit. The fitting can be done for a single or multiple images. For a single image, we optimize the outfit code z^* to match the segmentation mask of the outfit in the image.

In more detail, we predict the binary outfit mask by passing given RGB image through Graphonomy network [13] and combining all semantic masks that correspond to clothing. We also fit the SMPL mesh to the person in the image using the SMPLify approach [9]. We then minimize the 2D Chamfer loss between the outfit segmentation mask and the projection of the predicted point-cloud onto the image. The projection takes into account the occlusions of the outfit by the SMPL mesh (e.g. the back part of the outfit when seen from the front). In this case, the optimization is performed over the outfit code z^* while the parameters of the draping network remain fixed to avoid overfitting to a single image.

For complex outfits we observed instability in the optimization process, which often results in undesired local




Figure 3: More color-coded results of the draping networks. Each row corresponds to a pose. The leftmost image shows the input to the draping network. The remaining columns correspond to three outfit codes. Color coding corresponds to spectral coordinates on the SMPL mesh surface. Color coding reveals that the draping transformation is noticeably non-local (i.e. the draping network does not simply compute local displacements). Also, color coding reveals correspondences between analogous parts of outfit point clouds across the draping network outputs.

minima. To find a better optimum, we start from several random initializations $\{z_1^*, \dots, z_T^*\}$ independently (in our experiments, $T = 4$ random initializations are used). After several optimization steps we take the average outfit vector $\bar{z} = \frac{1}{T} \sum_{t=1}^T z_t^*$ and then continue the optimization from \bar{z} until convergence. We observed that this simple technique provides consistently accurate outfit codes. Typically we make 100 training steps while optimizing T hypothesis. After the averaging the optimization takes 50 – 400 steps depending on the complexity of the outfit’s geometry.

3.2. Appearance modeling

Point-based rendering. Most applications of clothing modeling go beyond geometric modeling and require to model the appearance as well. Recently, it has been shown that point clouds provide good geometric scaffolds for neural rendering [2, 60, 40]. We follow the neural point-based graphics (NPBG) modeling approach [2] to add appearance modeling to our system (Figure 4).

Thus, when modeling the appearance of a certain outfit with the outfit code z , we attach p -dimensional latent appearance vectors $T = \{t[1], \dots, t[M]\}$ to each of the M points in the point cloud that models its geometry. We

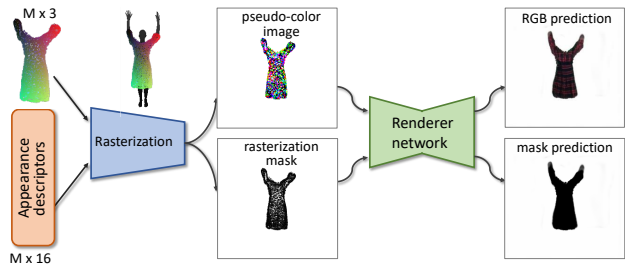


Figure 4: We use neural point-based graphics to model the appearance of an outfit. We thus learn the set of neural appearance descriptors and the renderer network that allow to translate the rasterization of the outfit point cloud into its realistic masked image (right).

also introduce the rendering network R_ψ with learnable parameters ψ . To obtain the realistic rendering of the outfit given the body pose s and the camera pose C , we then first compute the point cloud $G_\theta(z, s)$, and then rasterize the point cloud over the image grid of resolution $W \times H$ using the camera parameters and the neural descriptor $t[m]$ as a pseudo-color of the m -th point. We concatenate the result of the rasterization, which is a p -channeled image, with the rasterization masks, which indicates non-zero pixels, and then process (translate) them into the outfit RGB color image and the outfit mask (i.e. a four-channel image) using the rendering network R_ψ with learnable parameters ψ .

During the rasterization, we also take into account the SMPL mesh of the body and do not rasterize the points occluded by the body. For the rendering network we use a lightweight U-net network [52].

Video-based appearance capture. Our approach allows to capture the appearance of the outfit from video. To do that we perform two-stage optimization. In the first stage, the outfit code is optimized, minimizing the Chamfer loss between the point cloud projections and the segmentation masks, as described in the previous section. Then, we jointly optimize latent appearance vectors T , and the parameters of the rendering network ψ . For the second stage we use (1) the perceptual loss [25] between the masked video frame and the RGB image rendered by our model, and (2) the Dice loss between the segmentation mask and the rendering mask predicted by the rendering network.

Appearance optimization requires a video of a person with whole surface of their body visible in at least one frame. In our experiments training sequences consist of 600 to 2800 frames for each person. The whole process takes roughly 10 hours on NVIDIA Tesla P40 GPU.

After the optimization, the acquired outfit model can be rendered for arbitrarily posed SMPL body shapes, providing RGB images and segmentation masks.

4. Experiments

We evaluate the geometric modeling and the appearance modeling within our approach and compare it to prior art. Please also refer to the supplementary video on the [project page](#)¹ for a more convenient demonstration of qualitative comparison.

Datasets. We use the Cloth3D [7] dataset to train our geometric meta-model. The Cloth3D dataset has 11.3K garment elements of diverse geometry modeled as meshes draped over 8.5K SMPL bodies undergoing pose changes. The fitting uses physics-based simulation. We split the Cloth3D dataset into 6475 training sequences and 1256 holdout sequences, where sequences differ by SMPL parameters and outfit mesh.

We evaluate both stages - geometry and appearance - using two datasets of human videos. These datasets do not contain 3D data and were not used during the draping network training. The *PeopleSnapshot* presented in [4] contains 24 videos of people in diverse clothes rotating in A-pose. In terms of clothing, it lacks examples of people wearing skirts and thus does not reveal the full advantage of our method. We also evaluate on a subset from *AzurePeople* dataset introduced in [20]. This subset contains videos of eight people in outfits of diverse complexity shot from 5 RGBD Kinect cameras. For both datasets we generate cloth segmentation masks with Graphonomy method [13] and SMPL meshes using SMPLify [9]. To run all approaches in our comparison, we also predict Openpose [11] keypoints, DensePose [14] UV renders and SMPL-X [47] meshes. For appearance modeling, we follow StylePeople’s procedure and use the data from four cameras as a training set and validate by the fifth (the leftmost) camera.

We note that the two evaluation datasets (PeopleSnapshot and AzurePeople) were not seen during the training of the draping network. Furthermore, the comparisons in this section and all the visualizations in the supplementary material are obtained given the previously unseen outfit segmentations. The poses and body shapes were also sampled from the holdout set and were not seen by the draping network and by the rendering network during their training. By this, we emphasize the ability of our approach to generalize to new outfit styles, new body poses, and new body shapes.

4.1. Details of the draping network

To build a geometric prior on clothing, our draping function G_θ is pre-trained on synthetic Cloth3D dataset. We split it into train and validation parts, resulting in $N = 6475$ training video sequences. Since most of the consequent frames share similar pose/clothes geometry, only every tenth frame is considered for the training. As described

¹<https://saic-violet.github.io/point-based-clothing>

	Ours v Tex2Shape	Ours v MGN	Ours v Octopus
<i>PeopleSnapshot</i>	38.1% vs 61.9%	50.9% vs 49.1%	47.8% vs 52.2%
<i>AzurePeople</i>	65.6% vs 34.3%	74.5% vs 25.5%	73.7% vs 26.3%

Table 1: Results of user study, in which the users compared the quality of 3D clothing geometry recovery (fitted to a single image). Our method is preferred on the AzurePeople dataset with looser clothing, while the previously proposed methods work better for tighter clothing of fixed topology.

in Sec. 3.1, we randomly initialize $\{z_1, \dots, z_N\}$, where $z_i \in Z \subseteq \mathbb{R}^d$ for each identity i in the dataset. In our experiments, we set the latent code dimensionality relatively low to $d=8$, in order to avoid overfitting during subsequent single-image shape fitting (as described in Sec. 3.1).

We feed the outfit codes z_i to an MLP encoder consisting of 5 fully-connected layers to obtain a 512-dimensional latent representation. Then it is passed to the AdaIn branch of the Cloud Transformer network. For pose and body information, we feed an SMPL point cloud with hands, feet and head vertices removed, see Figure 1. The draping network outputs three-dimensional point clouds with 8.192 points in all experiments. We choose approximate Earth Mover’s Distance [33] as the loss function and optimize each GLO-vector and the draping network simultaneously using Adam [29].

While our pre-training provides expressive priors on dresses and skirts, the ability of the model to produce tighter outfits is somewhat limited. We speculate that this effect is mainly caused by a high bias towards jumpsuits in the Cloth3D tight clothing categories.

4.2. Recovering outfit geometry

In this series of experiments, we evaluate the ability of our method to recover the outfit geometry from a single photograph. We compare *Ours* point-based approach with the following three methods:

1. The *Tex2Shape* method [5] that predicts offsets for vertices of SMPL mesh in texture space. It is ideally suited for the PeopleSnapshot dataset, while less suitable to AzurePeople sequences with skirts and dresses.
2. The *Octopus* work [3] uses the displacements to the SMPL body model vertices to reconstruct a full-body human avatar with hair and clothing. Though the authors note that it is not ideally suited for reconstruction based on a single photograph.
3. The *Multi-garment net* approach [8] builds upon Octopus and predicts upper and lower clothing as separate meshes. It proposes a virtual wardrobe of pre-fitted garments, and is also able to fit new outfits from a single image.

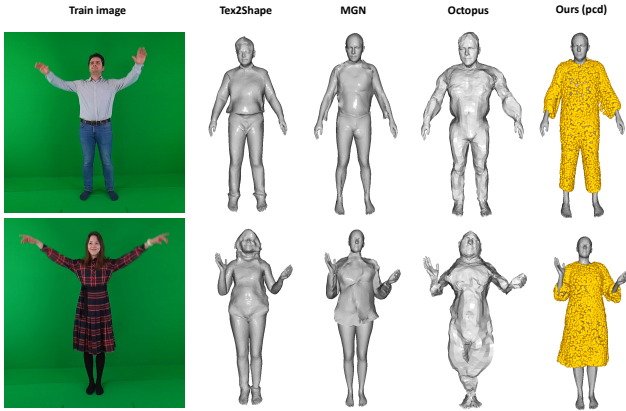


Figure 5: We show the predicted geometries in the validation poses fitted to a single frame (left). For our method (right) the geometry is defined by a point cloud (shown in yellow), while for Tex2Shape and MultiGarmentNet (MGN) the outputs are mesh based. Our method is able to reconstruct the dress, while other methods fail (bottom row). Note, our method is able to reconstruct a tighter outfit too (top row), though Tex2Shape with its displacement-based approach achieves a better result in this case.

We note that the compared systems use different formats to recover clothing (point cloud, vertex offsets, meshes). Furthermore, they are actually solving slightly different problems, as our method and Multi-garment net recover clothing, while Tex2Shape recovers meshes that comprise clothing, body, and hair. All three systems, however, support retargeting to new poses. We therefore decided to evaluate the relative performance of the three methods through a user study that assesses the realism of clothing retargeting.

We present the users with triplets of images, where the middle image shows the source photograph, while the side images show the results of two compared methods (in the form of shaded mesh renders for the same new pose). The result of such pairwise comparisons (user preferences) aggregated over $\sim 1.5k$ user comparisons are shown in Table 1. Our method is strongly preferred by the users in the case of AzurePeople dataset that contains skirts and dresses, while Tex2Shape and MGN are preferred on PeopleSnapshot dataset that has tighter clothing with fixed topology. Figure 5 shows typical cases, while the supplementary material provides more extensive qualitative comparisons. Note, in user study we paint our points in gray to exclude the coloring factor in user’s choice.

Since our approach uses 2D information to fit outfit code, we decided to omit quantitative comparison by the standard metrics due to the lack of datasets that contain both realistic RGB and realistic 3D data. However, we compare our method to MGN on BCNet [24] dataset. For both methods,

we use projection masks of the outfit meshes to fit clothing geometry. Our approach fits clothing geometry better in terms of Chamfer distance to vertices of ground truth outfit meshes in validation poses (0.00121 vs 0.0025) on 200 randomly chosen samples.




Figure 6: Our method is also capable of modeling the separated top and bottom garment styles. Here two different outfits in two different poses are shown.

4.3. Appearance modeling

We evaluate our appearance modeling pipeline against the StylePeople system [20] (the multiframe variant) that is the closest to ours in many ways. StylePeople fits a neural texture of the SMPL-X mesh alongside the rendering network using a video of a person using backpropagation. For comparison purposes we modify StylePeople to generate clothing masks along with rgb images and foreground segmentations. Both approaches are trained separately on each person from AzurePeople and PeopleSnapshot dataset. We then compare outfit images generated for holdout views in terms of three metrics that measure visual similarity to ground truth images, namely learned perceptual similarity (LPIPS) [65] distance, structural similarity (SSIM) [59] and its multiscale version (MS-SSIM).

The results of the comparison are shown in Table 2, while qualitative comparison is shown in Figure 7. In Figure 1, we show additional results for our methods. Specifically, we show a number of clothing outfits of varying topology and type that are retargeted to new poses from both test datasets. Finally, in Figure 8, we show examples of retargeting of outfit geometry and appearance to new body shapes within our approach.

5. Summary and Limitations

We have proposed a new approach to human clothing modeling based on point clouds. We have thus built a generative model for outfits of various shape and topology that allows us to capture the geometry of previously unseen outfits and to retarget it to new poses and body shapes. The topology-free property of our geometric representation (point clouds) is particularly suitable for modeling clothing due to wide variability of shapes and composition of outfits in real life. In addition to geometric modeling, we use



Figure 7: We compare the appearance retargeting results of our method to new poses unseen during fitting between our method and the StylePeople system (multi-shot variant), which uses the SMPL mesh as the underlying geometry and relies on neural rendering alone to “grow” loose clothes in renders. As expected, our system produces sharper results for looser clothes due to the use of more accurate geometric scaffolding. *Zoom-in is highly recommended.*



Figure 8: Our approach can also retarget the geometry and the appearance to new body shapes. The appearance retargeting works well for uniformly colored clothes, though detailed prints (e.g. chest region in the bottom row) can get distorted.

the ideas of neural point-based graphics to capture clothing appearance, and to re-render full outfit models (geometry + appearance) in new poses on new bodies.

Geometry limitations Our model does not consider cloth dynamics, and to extend our model in that direction some integration of our approach with physics-based modeling (e.g. finite elements) could be useful. Also, our model is limited to outfits similar to those represented in the Cloth3D dataset. Garments not present in the dataset (e.g. hats) can

	LPIPS↓	SSIM↑	MS-SSIM↑
<i>PeopleSnapshot</i>			
Ours	0.031	0.950	0.976
StylePeople	0.0569	0.938	0.972
<i>AzurePeople</i>			
Ours	0.066	0.925	0.937
StylePeople	0.0693	0.923	0.946

Table 2: Quantitative comparisons with the StylePeople system on the two test datasets using common image metrics. Our approach outperforms StylePeople in most metrics thanks to more accurate geometry modeling within our approach. This advantage is validated by visual inspection of quantitative results (Figure 7).

not be captured by our method. This issue could be possibly addressed by using another synthetic datasets on par with Cloth3D, as well as using real-world 3D scan datasets with ground truth clothing meshes.

Appearance limitations Our current approach to appearance modeling requires a video sequence in order to capture outfit appearance, which can be potentially addressed by expanding the generative modeling to the neural descriptors in a way similar to generative neural texture model from [20]. We also found the results of our system to be prone to flickering artifacts, which is a common issue for neural rendering schemes based on point clouds [2]. We believe, those artifacts may be alleviated by introducing more sophisticated rendering scheme or by using denser point clouds.

References

- [1] DRAPE: DRessing Any PErson. **2**
- [2] K. Aliev, A. Sevastopolsky, M. Kolos, D. Ulyanov, and V. S. Lempitsky. Neural point-based graphics. In A. Vedaldi, H. Bischof, T. Brox, and J. Frahm, editors, *Proc. ECCV*, volume 12367 of *Lecture Notes in Computer Science*, pages 696–712. Springer, 2020. **2, 5, 8**
- [3] T. Alldieck, M. Magnor, B. L. Bhatnagar, C. Theobalt, and G. Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *Proc. CVPR*, 2019. **3, 6**
- [4] T. Alldieck, M. Magnor, W. Xu, C. Theobalt, and G. Pons-Moll. Video based reconstruction of 3d people models. In *Proc. CVPR*, pages 8387–8397, 2018. **6**
- [5] T. Alldieck, G. Pons-Moll, C. Theobalt, and M. Magnor. Tex2shape: Detailed full human body geometry from a single image. In *Proc. 3DV*, 2019. **3, 6**
- [6] K. Ayush, S. Jandial, A. Chopra, and B. Krishnamurthy. Powering virtual try-on via auxiliary human segmentation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. **3**
- [7] H. Bertiche, M. Madadi, and S. Escalera. Cloth3d: Clothed 3d humans. In *Proc. ECCV*, 2020. **2, 3, 4, 6**
- [8] B. L. Bhatnagar, G. Tiwari, C. Theobalt, and G. Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *Proc. ICCV. IEEE*, oct 2019. **2, 3, 6**
- [9] F. Bogo, A. Kanazawa, C. Lassner, P. V. Gehler, J. Romero, and M. J. Black. Keep it SMPL: automatic estimation of 3d human pose and shape from a single image. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Proc. ECCV*, volume 9909 of *Lecture Notes in Computer Science*, pages 561–578. Springer, 2016. **4, 6**
- [10] P. Bojanowski, A. Joulin, D. Lopez-Paz, and A. Szlam. Optimizing the latent space of generative networks. In *Proc. ICML*, 2019. **3, 4**
- [11] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh. OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):172–186, 2019. **6**
- [12] V. Gabeur, J.-S. Franco, X. Martin, C. Schmid, and G. Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *Proc. ICCV*, 2019. **2**
- [13] K. Gong, Y. Gao, X. Liang, X. Shen, M. Wang, and L. Lin. Graphonomy: Universal human parsing via graph transfer learning. In *Proc. CVPR*, 2019. **4, 6**
- [14] R. A. Güler, N. Neverova, and I. Kokkinos. Densepose: Dense human pose estimation in the wild. In *Proc. CVPR*, pages 7297–7306, 2018. **6**
- [15] E. Gundogdu, V. Constantin, A. Seifoddini, M. Dang, M. Salzmann, and P. Fua. Garnet: A two-stream network for fast and accurate 3d cloth draping. In *Proc. ICCV*, 2019. **2**
- [16] X. Han, Z. Wu, Z. Wu, R. Yu, and L. S. Davis. Viton: An image-based virtual try-on network. In *Proc. CVPR*, 2018. **3**
- [17] W.-L. Hsiao, I. Katsman, C.-Y. Wu, D. Parikh, and K. Grauman. Fashion++: Minimal edits for outfit improvement. In *Proc. ICCV*, 2019. **3**
- [18] X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proc. ICCV*, 2017. **4**
- [19] Z. Huang, Y. Xu, C. Lassner, H. Li, and T. Tung. Arch: Animatable reconstruction of clothed humans. In *Proc. CVPR*, 2020. **2**
- [20] K. Iskakov, A. Grigorev, A. Ianina, R. Bashirov, I. Zakharkin, A. Vakhitov, and V. Lempitsky. StylePeople: A Generative Model of Fullbody Human Avatars. In *Proc. CVPR*, 2021. **3, 6, 7, 8**
- [21] T. Issenhuth, J. Mary, and C. Calauzènes. Do not mask what you do not need to mask: a parser-free virtual try-on. In *Proc. ECCV*, 2020. **3**
- [22] H. Jae Lee, R. Lee, M. Kang, M. Cho, and G. Park. Laviton: A network for looking-attractive virtual try-on. In *Proc. ICCV*, Oct 2019. **3**
- [23] N. Jetchev and U. Bergmann. The conditional analogy gan: Swapping fashion articles on people images. In *Proc. ICCV*, 2017. **3**
- [24] B. Jiang, J. Zhang, Y. Hong, J. Luo, L. Liu, and H. Bao. Bcnet: Learning body and cloth shape from a single image, 2020. **7**
- [25] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Proc. ECCV*, 2016. **5**
- [26] A. Kanazawa, M. J. Black, D. W. Jacobs, and J. Malik. End-to-end recovery of human shape and pose. In *Proc. CVPR*, 2018. **2**
- [27] T. Karras, S. Laine, and T. Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proc. CVPR*, pages 4401–4410. Computer Vision Foundation / IEEE, 2019. **4**
- [28] T. Karras, S. Laine, M. Aittala, J. Hellsten, J. Lehtinen, and T. Aila. Analyzing and improving the image quality of stylegan. In *Proc. CVPR*, 2020. **3**
- [29] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization, 2017. **6**
- [30] Z. Laehner, D. Cremers, and T. Tung. DeepWrinkles: Accurate and Realistic Clothing Modeling. In *Proc. ECCV*, 2018. **2**
- [31] S. Laine, J. Hellsten, T. Karras, Y. Seol, J. Lehtinen, and T. Aila. Modular primitives for high-performance differentiable rendering. *ACM Transactions on Graphics*, 39(6), 2020. **2**
- [32] K. M. Lewis, S. Varadharajan, and I. Kemelmacher-Shlizerman. Vogue: Try-on by stylegan interpolation optimization, 2021. **3**
- [33] M. Liu, L. Sheng, S. Yang, J. Shao, and S.-M. Hu. Morphing and sampling network for dense point cloud completion. *arXiv preprint arXiv:1912.00280*, 2019. **4, 6**
- [34] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. **2, 3**
- [35] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. **2**

- [36] M. M. Loper and M. J. Black. OpenDR: An approximate differentiable renderer. In *Proc. ECCV*, volume 8695 of *Lecture Notes in Computer Science*, pages 154–169. Springer International Publishing, Sept. 2014. [2](#)
- [37] Q. Ma, S. Saito, J. Yang, S. Tang, and M. J. Black. SCALE: Modeling Clothed Humans with a Surface Codec of Articulated Local Elements. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [3](#)
- [38] Q. Ma, J. Yang, A. Ranjan, S. Pujades, G. Pons-Moll, S. Tang, and M. J. Black. Learning to Dress 3D People in Generative Clothing. In *Proc. CVPR*, 2020. [2](#)
- [39] K. Mazur and V. Lempitsky. Cloud transformers. In *Proc. ICCV*, 2021. [4](#)
- [40] M. Meshry, D. B. Goldman, S. Khamis, H. Hoppe, R. Pandey, N. Snavely, and R. Martin-Brualla. Neural re-rendering in the wild. In *Proc. CVPR*, June 2019. [2, 5](#)
- [41] M. R. Minar and H. Ahn. Cloth-vton: Clothing three-dimensional reconstruction for hybrid image-based virtual try-on. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020. [3](#)
- [42] M. R. Minar, T. T. Tuan, H. Ahn, P. Rosin, and Y.-K. Lai. Cp-vton+: Clothing shape and texture preserving image-based virtual try-on. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020. [3](#)
- [43] A. Mir, T. Alldieck, and G. Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *Proc. CVPR*. IEEE, June 2020. [3](#)
- [44] R. Natsume, S. Saito, Z. Huang, W. Chen, C. Ma, H. Li, and S. Morishima. Siclope: Silhouette-based clothed people. In *Proc. CVPR*, 2019. [2](#)
- [45] A. Neuberger, E. Borenstein, B. Hilleli, E. Oks, and S. Alpert. Image based virtual try-on network from unpaired data. In *Proc. CVPR*, June 2020. [3](#)
- [46] C. Patel, Z. Liao, and G. Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proc. CVPR*, 2020. [2](#)
- [47] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proc. CVPR*, pages 10975–10985, 2019. [6](#)
- [48] G. Pons-Moll, S. Pujades, S. Hu, and M. Black. ClothCap: Seamless 4D Clothing Capture and Retargeting. *ACM Transactions on Graphics, (Proc. SIGGRAPH)*, 36(4), 2017. Two first authors contributed equally. [2](#)
- [49] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proc. CVPR*, pages 77–85. IEEE Computer Society, 2017. [4](#)
- [50] A. Raj, P. Sangkloy, H. Chang, J. Hays, D. Ceylan, and J. Lu. Swapnet: Image based garment transfer. *Proc. ECCV*, pages 679–695, 2018. [3](#)
- [51] N. Ravi, J. Reizenstein, D. Novotny, T. Gordon, W.-Y. Lo, J. Johnson, and G. Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. [2](#)
- [52] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. (available on arXiv:1505.04597 [cs.CV]). [5](#)
- [53] S. Saito, Z. Huang, R. Natsume, S. Morishima, A. Kanazawa, and H. Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proc. ICCV*, 2019. [2](#)
- [54] S. Saito, T. Simon, J. Saragih, and H. Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proc. ICCV*, 2020. [2](#)
- [55] Y. Shen, J. Liang, and M. C. Lin. Gan-based garment generation using sewing pattern images. In *Proc. ECCV*, 2020. [3](#)
- [56] J. Thies, M. Zollhöfer, and M. Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 2019. [3](#)
- [57] G. Varol, D. Ceylan, B. Russell, J. Yang, E. Yumer, I. Laptev, and C. Schmid. BodyNet: Volumetric inference of 3D human body shapes. In *Proc. ECCV*, 2018. [2](#)
- [58] B. Wang, H. Zheng, X. Liang, Y. Chen, L. Lin, and M. Yang. Toward characteristic-preserving image-based virtual try-on network. In *Proc. ECCV*, 2018. [3](#)
- [59] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. [7](#)
- [60] O. Wiles, G. Gkioxari, R. Szeliski, and J. Johnson. SynSin: End-to-end view synthesis from a single image. In *Proc. CVPR*, 2020. [2, 5](#)
- [61] D. Xiang, F. Prada, C. Wu, and J. Hodgins. Monoclothcap: Towards temporally coherent clothing capture from monocular rgb video. In *Proc. 3DV*, 2020. [2](#)
- [62] H. Yang, R. Zhang, X. Guo, W. Liu, W. Zuo, and P. Luo. Towards photo-realistic virtual try-on by adaptively generating↔preserving image content. In *Proc. CVPR*, 2020. [3](#)
- [63] G. Yildirim, N. Jetchev, R. Vollgraf, and U. Bergmann. Generating high-resolution fashion model images wearing custom outfits. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. [3](#)
- [64] J. S. Yoon, K. Kim, J. Kautz, and H. S. Park. Neural 3d clothes retargeting from a single image, 2021. [2](#)
- [65] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric, 2018. [7](#)
- [66] Z. Zheng, T. Yu, Y. Wei, Q. Dai, and Y. Liu. Deephuman: 3d human reconstruction from a single image. In *Proc. ICCV*, October 2019. [2](#)