
This is the **accepted version** of the book part:

Laczkó, Hunor; Madadi, Meysam; Escalera, Sergio; [et al.]. «A Generative multi-resolution pyramid and normal-conditioning 3D cloth draping». A: 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). 2024, p. 8694-8703. 10 pàg. IEEE. DOI 10.1109/WACV57701.2024.00851

This version is available at <https://ddd.uab.cat/record/311810>

under the terms of the  IN COPYRIGHT license

A Generative Multi-Resolution Pyramid and Normal-Conditioning 3D Cloth Draping

Hunor Laczkó^{*,†}

Meysam Madadi^{†,‡}

Sergio Escalera^{†,‡}

Jordi Gonzalez^{*,†}

hunor.laczko@uab.cat, mmadadi@cvc.uab.es, sescalera@ub.edu, jordi.gonzalez@uab.cat

^{*}Universitat Autònoma de Barcelona, [†]Computer Vision Center, [‡] Universitat de Barcelona
Barcelona, Spain

Abstract

RGB cloth generation has been deeply studied in the related literature, however, 3D garment generation remains an open problem. In this paper, we build a conditional variational autoencoder for 3D garment generation and draping. We propose a pyramid network to add garment details progressively in a canonical space, i.e. unposing and unshaping the garments w.r.t. the body. We study conditioning the network on surface normal UV maps, as an intermediate representation, which is an easier problem to optimize than 3D coordinates. Our results on two public datasets, CLOTH3D and CAPE, show that our model is robust, controllable in terms of detail generation by the use of multi-resolution pyramids, and achieves state-of-the-art results that can highly generalize to unseen garments, poses, and shapes even when training with small amounts of data. The code can be found at: <https://github.com/HunorLaczko/pyramid-drape>

1. Introduction

Cloth generation and reconstruction is an active topic in computer vision with applications in virtual try-on, fashion, gaming, and movie editing, just to name a few [12].

Unfortunately, garment generation is still an open problem due to challenges such as variable topology, shape, size, style, fabric, and texture, as well as highly dynamic behavior. The problem has been tackled in the literature from different points of view: 2D image generation [13, 16, 17, 23, 43], image-based 3D reconstruction [1, 2, 20, 22, 42, 46–48] and 3D garment draping and manipulation [3, 4, 7, 18, 27, 29, 32, 37, 39, 44]. This paper is focused on the latter category.

Proposed solutions mainly depend on the garment representation. Meshes have been widely used to represent 3D surfaces [3, 4, 7, 8, 28, 32, 37]. However, designing architectures to effectively deal with variable mesh topologies is

	Gen.	GarmUns.	Interp.	Repr.
SCALE [27]	✗	✗	✗	Patch point cloud
Neural-GIF [39]	✗	✗	✗	Implicit function
TailorNet [32]	✗	✗	✓	Mesh
POP [29]	✗	✗	✗	Point cloud
PBNS [5]	✗	✗	✓	Mesh
Santesteban <i>et al.</i> [37]	✗	✗	✗	Mesh
DeePSD [7]	✗	✓	✗	Mesh
Zakharkin <i>et al.</i> [44]	✗	✓	✗	Point cloud
HOOD [18]	✗	✓	✓	Mesh
CAPE [28]	✓	✓	✗	Mesh
CLOTH3D [3]	✓	✓	✗	Mesh
Ours	✓	✓	✓	UV map

Table 1. We compare our approach to state-of-the-art 3D garment generation and draping w.r.t. to their solution’s properties: 1) is it Generative (Gen.)?, 2) is it Garment Unspecific (GarmUns.) i.e. does not require training on many different cloth categories?, 3) does it produce Interpretable (Interp.) results?, and 4) what is the Representation (Repr.)? Most of the available approaches train their specific models per each particular cloth type. Also, there are very few generative models that can output different cloth dynamics for the same input pose. Lastly, some approaches provide explainable predictions either w.r.t. local dynamics (TailorNet and ours) or physically based losses (PBNS and HOOD). Our approach covers all these points while outperforming state-of-the-art and generalizing well to unseen poses.

not a trivial task. Recently, point cloud representation has received special attention thanks to the use of implicit functions [19, 20, 22, 39, 42] or point representations [27, 29, 44]. Although these functions generate high-fidelity details (after converting point clouds to mesh surfaces) with an arbitrary number of output points, their high computational cost and limited quality of the generated surface when applied to unseen cases are major concerns. Lastly, UV maps representations, from texture to 3D body and cloth, have shown successful results in different application scenarios [10, 11, 24, 31, 33, 38, 41, 45], since they can be wrapped to a mesh at no cost while handling arbitrary topologies. We apply UV map representations in our work.

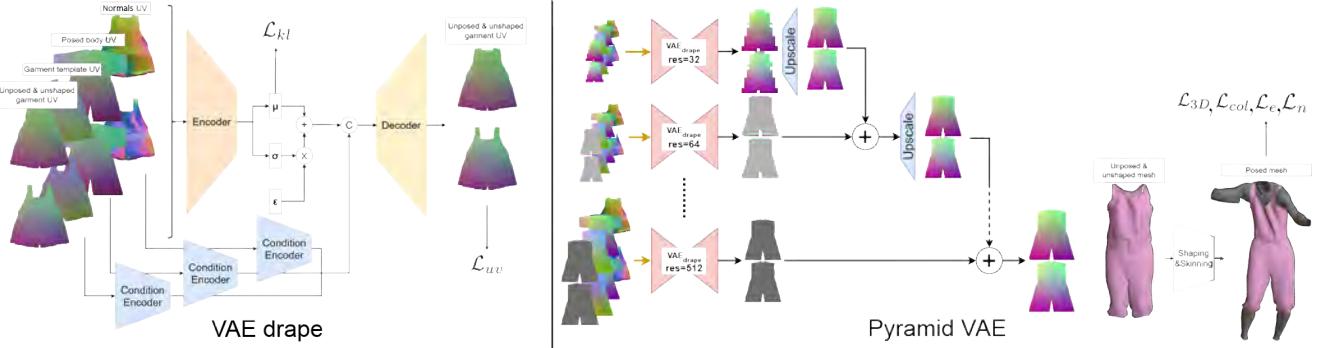


Figure 1. The proposed pyramid pipeline (right) contains basic VAE modules for each draping level (VAE_{drape} , left). VAE_{drape} receives conditioning inputs and garment offsets and reconstructs the unposed and unshaped garment offsets as UV image. In the case of the first level instead of offsets, absolute coordinates are used (as shown on the left) as this will serve as a base for subsequent levels. The conditioning variables (normals, posed body, and garment template UV images) are given into three pre-trained and frozen encoders to fuse with the VAE_{drape} latent code. These conditioning encoders are trained separately in an autoencoder manner (note that normals are trained through VAE_{norm}). Finally, the reconstructed UV image is converted to a mesh and passed to the skinning module after reshaping. Then, in the pyramid module, the lowest resolution level predicts low-frequency garments while the other levels are learned as offsets over their previous level. Each level output is upsampled with the proposed upscaling network and summed to the next level. At inference time, we sample from VAE_{norm} and VAE_{drape} and pass the template garment and posed body UV images.

Regardless of the chosen representation, available techniques are commonly applied to (i) generate 3D garment templates in a canonical pose or (ii) drape/animate a given 3D template w.r.t. the body pose. Most state-of-the-art approaches, like ours, are based on the latter case [3, 4, 6, 7, 27, 29, 32, 36, 37, 39, 44]. However, there are some common limitations in most of these techniques. Firstly, cloth draping is a one-to-many problem, i.e. for a given pose and template garment there are many valid animated garments, and each garment type has different dynamical behavior. This is while these techniques overfit to the training set to predict just one single solution. Although this may be valid for some applications, it does not provide any control over the prediction. Secondly, either every single garment has its own trained model (can be categorized as supervised or unsupervised neural simulator) [5, 27, 29, 36, 37] or one model is trained per each garment category [8, 32], which limit applicability for real-world scenarios.

In order to cope with the aforementioned limitations, one of the main aims of this work is to train all sorts of garments in a **minimal** number of generative models. Generative cloth architectures have been studied before in [3, 28] using convolutional graph networks. In this paper, we develop a variational autoencoder (VAE) on CNN architecture backbones, thanks to the use of UV maps, which have shown great results in generating plausible RGB images. Additionally, we condition the VAE on normal maps, as intermediate representations, which increases stochasticity and has been shown to be effective in gaining high-fidelity results [42].

Finally, there have been substantial efforts in convolutional data processing to learn features specific to dif-

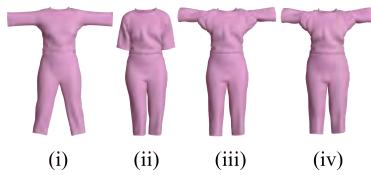
ferent frequency domains [14, 15, 21, 34]. Following this trend, we build a pyramid conditional VAE, with a series of VAEs each trained on a different image resolution. Thus, the lowest resolution VAE is responsible for learning low-frequency cloth dynamics while the other VAEs learn high-frequency details as offsets, progressively added to the previous level.

Our contributions are:

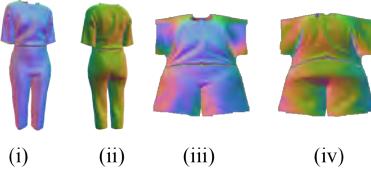
- We develop a pyramid VAE for 3D cloth draping using a ConvNeXt backbone architecture [25] and 3D surface coordinates represented as UV images.
- We condition the network on the surface normals unwrapped in UV maps.
- We show state-of-the-art (SOTA) 3D generative cloth draping results by progressively improving the error and surface quality on two public datasets, CAPE [28] and CLOTH3D [3].

2. Related works

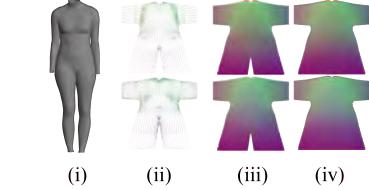
We compare our proposed architecture to the state-of-the-art in Tab. 1. Most of the available works are deterministic approaches [5, 7, 18, 27, 29, 32, 37, 39, 44] in which a single solution is predicted for the given inputs. There has been a set of approaches [5, 27, 29, 37, 39], among deterministic works, in which a single garment instance is trained over many given body poses, called neural simulation, depicting cloth simulation in graphic engines. The works of [27, 29, 35] convert the given body surface into



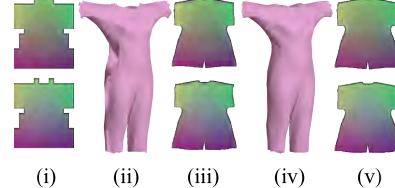
(a) The template garment in T-pose (i) is draped on a given posed body (ii) which is unposed (iii) through an optimization process. Later it is unshaped (iv) by reversing the SMPL shaping process. Notice how the unshaped garment is wider, as the underlying body was thinner than the SMPL template.



(b) Through UV unwrapping a mapping is defined between the 3D mesh vertices (i) and the UV image pixels. After projecting the 3D mesh to the UV space the result is sparse (ii), so we inpaint the missing values by interpolation for non-skirts (iii) and skirts (iv) in two different UV templates.



(c) We calculate normal vectors for each vertex, visible on the front (i) and back (ii) of the mesh which are converted to front (iii) and back (iv) UV maps.



(d) Upscaling from a low-resolution UV image (i) results in artifacts around the boundaries (ii-iii) using regular upscaling. By the proposed upscaling (iv-v) this can be avoided.

Figure 2. Examples of preprocessed data. a) Cloth unposing and unshaping. b) 3D mesh to UV map. c) Surface normals calculation. d) UV image down/upscaling.

a UV map and learn 2D body pose features through CNN architectures, similar to us. In [29] an additional garment geometry feature tensor is learned through auto-decoding. Later 2D sampled features along with positional encodings are given to a shared MLP to predict an offset vector added to the sampled body point. Authors in [39] use a combination of several MLPs as an implicit function to first unpose sampled points and add offset vectors to them. The updated points along with pose encodings are given to a signed distance function to predict the point occupancy. Learning features in a canonical pose has been a common technique in state-of-the-art. Similar to us, [37] preprocess the garments to form a collision-free unposed and unshaped garment. Later a VAE is trained to learn garment-specific deformations. Although they use a VAE for training, their approach is not designed to be generative at inference time. Rather, they predict the encoded features from the body shape parameters and motion information. In a different solution, TailorNet [32] decomposes the garment into high and low-frequency surfaces in a canonical pose. Then the wrinkles are formed by a mixture of high frequencies. All the aforementioned approaches were trained supervisedly. PBNS [4] is the first unsupervised approach that uses physically based losses to train a lightweight neural network. They apply an MLP to extract features from the pose which is multiplied to an optimizable garment geometry matrix. The outputs are offsets added to the template garment in a canonical pose. Later, this type of solution has been extended to include temporal information [6, 36].

Among deterministic approaches, several works develop garment unspecific models in which a single model is

trained to handle all sorts of garments. DeePSD [7] is inspired by the physically-based simulation formulation and can handle variable topologies during training without registration. They train the network by a combination of supervised and unsupervised losses. Zakharkin *et al.* [44] transforms the posed body point cloud into the final garment by conditioning on the garment latent code. HOOD [18] develops a new message-passing mechanism in graph neural networks and trains the network with self-supervised physically based losses.

To the best of our knowledge, there are a few generative methodologies in 3D garment draping [3, 28]. These works apply VAE graph neural networks on registered meshes in a canonical pose. CLOTH3D [3] proposes a new topology for the body to model skirts-like clothes. CAPE [28] trains the network by an additional adversarial loss. Our approach decomposes the surface into different frequencies and learns the details progressively showing accurate reconstruction capabilities.

3. Generative 3D cloth draping

The proposed 3D cloth draping pipeline is: 1) generative and stochastic, 2) efficient, 3) general, as it does not require separate models for each garment type, and 4) able to handle both high and low-frequency details. We assume a) SMPL body pose and shape parameters [26] plus a garment template in canonical pose (T-pose for example) are available for the draping, b) the garment size is matched with the body shape¹, c) the garment has a low elasticity,

¹Although garment resizing is possible as a preprocessing to match the body shape, we consider it as out-of-the-scope of this paper.

thus the mesh edge lengths should not change drastically after the draping, and d) the garment has a single surface layer (e.g. pockets are not considered). Also, we apply our experiments to daily life clothes. Next, we explain the details and architectural choices to fulfill the aforementioned goals. The architecture pipeline is shown in Fig. 1.

3.1. Preprocessing

3.1.1 Cloth unposing and unshaping

We follow the trends in the literature [3, 27–29, 32] that force the network to make predictions in a canonical pose space to improve the results. We explore this idea in several aspects: 1) eliminating the global rotation, 2) unposing, and 3) unshaping. Let $\theta \in \mathbb{R}^{24 \times 3 \times 3}$ and $\beta \in \mathbb{R}^{1 \times 10}$ be the SMPL pose and shape parameters, $B \in \mathbb{R}^{10 \times (m \times 3)}$ and $W \in \mathbb{R}^{m \times 24}$ the SMPL blend shape and blend weights, $D \in \mathbb{R}^{m \times 3}$ the shaped and posed SMPL body vertices, $G_t \in \mathbb{R}^{n \times 3}$ a template garment with n vertices registered on SMPL mesh (see next section for details) and $M \in \{0, 1\}^m$ its registration mask (note $n = \sum M$), and $G \in \mathbb{R}^{n \times 3}$ the G_t garment draped on D . To eliminate the global orientation, we first update the garment and body vertices by $G \cdot \theta^0$ and $D \cdot \theta^0$, and the θ root joint rotation (θ^0) by an identity matrix. Then, we transform G to a canonical pose space by:

$$G_c = \arg \min_{G_c} \|S(G_c + [\beta \cdot B]^M, \theta; W) - G\|, \quad (1)$$

where S is the skinning function and $[.]^M$ is the indexing based on mask M . An example can be seen in Fig. 2(a).

3.1.2 Cloth registration

We follow CLOTH3D [3] protocol for registration, that is 1) increasing SMPL resolution and removing head, hands, and feet vertices ($m = 14475$ after this update), and 2) creating an extra body mesh for skirt-like garments. Additionally, we apply laplacian smoothing to the body and template garment and use a non-rigid ICP to warp the template garment to the body surface this way achieving a less noisy registration. The mask M is extracted by taking the nearest body vertices. Finally, we compute the registration as follows.

The simplest registration solution is by taking body vertices and finding the nearest vertex on the warped garment. However, it is possible that a garment vertex is assigned to two or more body vertices. This causes null faces on the garment which degrades the garment resolution and leads to noisy loss functions. Instead of the nearest mesh vertex, we find the nearest point on the garment surface through a fast numerical approximation. To do so, we generate a random vector $\omega^{k \times 3}$ for each garment face in which $1 = \sum_{j=1}^3 \omega_{ij}$ where $i \in \{1..k\}$. We use ω as a weighting over the face vertices coordinates to sample points on each garment face.

Finally, we find the nearest points from this randomly sampled set and save the ω values for the nearest point and its corresponding face index. This is used to reconstruct the garment in any frame represented by SMPL topology.

3.1.3 3D mesh to UV map

UV map representation has been used by several approaches [10, 11, 24, 33, 38, 41, 45]. Its advantage over meshes is that it is a 2D image, which makes it applicable in history-rich GPU-efficient CNN models. However, a UV map is only partially continuous, unlike a mesh. The reason is that to unwrap a 3D mesh into 2D, the surface is torn into several segments to fit in. To minimize the number of segments and distortions, we follow [30] in which seamlines at the sides of the body are manually defined where they divide the body into two segments: front and back. Based on this idea we create UV images with a 512×256 resolution. In the supplementary material, we explain the details of how we create the UV pixel coordinates corresponding to 3D vertices to have a 3D reconstructed surface. An example can be seen in Fig. 2(b). Through this operation we convert D , G , G_t and G_c into D^{uv} , G^{uv} , G_t^{uv} and G_c^{uv} , respectively. Also, the mask M is updated to M^{uv} accordingly.

Note that although vertex indices are identical between skirt-like and trousers-like clothes after registration, they do not follow the body in the same way and also have different dynamic behaviors. Therefore, a single template for UV coordinates for both of them is not valid. For this reason, we group garments into skirts and non-skirts, create two UV templates for each, and train them independently. This allows us to cover all garment types minimally with only two models, unlike other methods that require a separate model for each type.

3.1.4 Surface normals

Surface frequency details can be partially reflected in the surface normals as the direction of the normals and their relative change compared to their neighborhood correlate to the details of the surface. An example can be seen in Fig. 2(c). As the figure shows, small wrinkles and even large deformations are easy to identify. This helps guide the model to be able to recreate these details.

3.1.5 UV image down/upscaling

Regular RGB images can be rescaled by simple bilinear or bicubic algorithms. Applying these techniques to UV images will result in invalid values and artifacts near the boundary areas (reflecting in the seamline between the two sides of the mesh) due to zero background pixels. Fortunately, downsampling can be done by a fast nearest neighbor interpolation to fill the background. However, it is not

accurate for upscaling. To overcome this, we develop a fast and differentiable interpolation block by utilizing a CNN with five convolutional layers only to fill the background. For training, we use MS-SSIM loss on the already available data from the previously downsampled images. Later, to train the pyramid VAE, we freeze this upscaling network. We apply image scaling by a factor of 2 and have UV resolutions in the range 64×32 and 512×256 . The results can be seen in Fig. 2(d).

3.2. Conditional pyramid VAE

In the basic definition of cloth draping, a garment is draped on the body given a pose. However, this problem has unlimited valid solutions. Therefore, models with a single solution do not truly formulate the problem. Generative models are able to deal with this issue, though sometimes they suffer from not being stochastic and become one-to-one functions. Cloths are highly dynamic objects and some approaches either converge to smooth and average dynamics [3, 28] or overfit to specific categories and poses [32]. To cope with these issues, we build our generative model based on a VAE (called VAE_{drape}) that is conditioned on the garment template G_t^{uv} , posed body D^{uv} and normal map N^{uv} , and reconstructs \hat{G}_c^{uv} ². Then, the 3D garment \hat{G} can be obtained by:

$$\hat{G} = \mathcal{S}([\hat{G}_c^{uv}]^{J \odot M} + [\beta \cdot B]^M, \theta, W) \quad (2)$$

where J contains corresponding UV indices to SMPL vertices masked by M and \odot is element-wise product.

Although the normal map is not available at inference time, its generation and sampling is a much easier problem than 3D cloth draping. Therefore, we build a second VAE for normals (called VAE_{norm}) that is again conditioned on G_t^{uv} and D^{uv} , and reconstructs \hat{N}^{uv} . Due to its simplicity, adaptation, and performance, we adapt ConvNeXt architecture [25] as the VAEs' encoders, thanks to the usage of UV images. Decoders are mirrors of encoders (downsampling convolutional layers are replaced with transpose convolutions). We explain the adapted architecture and training details in the supplementary material.

As a common problem for CNNs in 3D space, low-frequency dynamics are learned first and high-frequency details are ignored. To solve this issue, we propose a pyramid network that starts from a low-resolution image to learn the global shape of the cloth. Then, the surface quality is progressively improved in the next higher resolution levels as offsets to the previous level. Note that we use the idea in Sec. 3.1.5 to upscale images between consecutive levels. See Fig. 1 for architecture details.

3.3. Losses

We use the following loss functions to train the network:

² $\hat{\cdot}$ means estimated by the network.

- L2 loss between \hat{G}_c^{uv} and G_c^{uv} masked by M^{uv} (\mathcal{L}_{uv}).
- KL loss (\mathcal{L}_{kl}).
- L2 loss between \hat{G} and G (\mathcal{L}_{3D}).
- Garment against body collision loss (\mathcal{L}_{col}) [7]:

$$\mathcal{L}_{col} = \sum_{(i,j) \in \mathcal{N}} \min(\mathbf{v}_{j,i} \cdot \mathbf{N}_j - \epsilon, 0)^2, \quad (3)$$

where \mathcal{N} is the set of nearest neighbor correspondences between the predicted cloth and body, $\mathbf{v}_{j,i}$ is the connecting vector between the j -th body vertex to the i -th cloth vertex, \mathbf{N}_j is the j -th body vertex normal, and $\epsilon = 4\text{mm}$ is a small threshold. This loss works well when the cloth is close to the body and pushes back collided vertices back to the nearest body vertex.

- L1 loss on the mesh edges lengths between \hat{G} and G (\mathcal{L}_e).
- L1 loss on 3D surface normals between \hat{G} and G (\mathcal{L}_n).

Then, we define the regularization, reconstruction, and final losses as:

$$\mathcal{L}_{reg} = \mathcal{L}_{col} + \mathcal{L}_e + \mathcal{L}_n + \mathcal{L}_{3D}, \quad (4)$$

$$\mathcal{L}_{rec} = \mathcal{L}_{uv} + \lambda \cdot \mathcal{L}_{reg}, \quad (5)$$

$$\mathcal{L} = \mathcal{L}_{rec} + \delta \cdot \mathcal{L}_{kl}, \quad (6)$$

where λ is a coefficient controlling the regularization losses and δ is a balancing term between the reconstruction loss and the KL divergence. Tuning δ is important to preserve the stochasticity of the VAE while still reaching a low reconstruction error. Therefore, we use β -VAE technique [9] to train the network. We start the training by setting $\delta < 10^{-3}$ and once we observe a flat \mathcal{L}_{rec} we linearly increase δ .

4. Experiments

Datasets. We use two common 3D human bodies and garments datasets CLOTH3D [3] and CAPE [28] to train our model and compare against SOTA. First, we use CLOTH3D, a large-scale synthetic dataset of clothed humans. It contains more than two million frames from several thousand sequences with a large variety of garment types, topologies, shapes, and motions. We use 20 percent of the original training dataset to speed up the training and evaluate our method on the defined test set.

Second, we use CAPE, a smaller dataset created with high-resolution 3D scans of real subjects. As it is smaller, it captures a smaller variety of garment types and subjects but still covers a wide range of motions. The scans are registered on top of the SMPL body and encoded as offsets

Data	Experiment	Pyr. 1	Pyr. 2	Pyr. 3	Pyr. 4
Non-skirt (baseline: 7.11)					
Incremental training	11.97	7.62	7.20	7.14	
End-to-end w/o reg.	10.07	5.41	4.33	3.79	
End-to-end w/ reg.	9.68	5.30	4.10	3.78	
Skirt (baseline: 9.79)					
Incremental training	15.78	9.49	9.13	8.90	
End-to-end w/o reg.	12.00	6.74	5.38	4.60	
End-to-end w/ reg.	11.94	6.60	5.00	4.28	
Combined (baseline: 8.01)					
Incremental training	13.24	8.24	7.84	7.72	
End-to-end w/o reg.	10.72	5.85	4.68	4.05	
End-to-end w/ reg.	10.43	5.73	4.40	3.95	

Table 2. Detailed V2V results of the ablation study on CLOTH3D dataset. We can see the effect of each contribution across each level of the pyramid. It is also compared to the baseline, non-pyramid model. Note that the combined error is not the average of the non-skirt and skirt numbers, but the average of the six garment categories combined.

over the minimally clothed body. Unlike CLOTH3D, the bodies and garments are available as a single entity and the model needs to predict them as a single mesh. The dataset is split into two separate datasets for female and male subjects, both containing four cloth types. Each of these datasets contains around 20-30K frames for training and around 5K for testing.

Metrics. Since we have SMPL body-garment registration available for both datasets, we are able to calculate a simple point-to-point Euclidean distance averaged across the garment. Therefore, we follow SOTA protocol to report the *reconstruction error*³ with two common metrics: the average per vertex Euclidean error (V2V) and the Chamfer distance (CD) in mm. Following SOTA, the reported errors are the averages over the average per garment type. We also calculate additional qualitative metrics: 1) the percentage of collided vertices and 2) the absolute difference of the surface area between the draped and template garments in m^2 . In the second metric, we analyze the amount of stretching or compression imposed by the model on the input template garment. While small differences are expected, large changes indicate that the prediction does not conform to the input template.

4.1. Ablation results

We perform an ablation of our approach on the CLOTH3D dataset.

³Note that to do so, we use the full VAE model including encoders and decoders.

	\mathcal{L}_e		\mathcal{L}_n		Collision (%)	
	Non-sk.	Skirt	Non-sk.	Skirt	Non-sk.	Skirt
Baseline	0.580	1.326	0.150	0.161	5.49	2.64
Increm.	0.578	1.581	0.159	0.184	5.45	3.10
E2e w/o reg.	0.706	1.058	0.126	0.133	2.00	0.63
E2e w/ reg.	0.447	0.906	0.124	0.132	1.98	0.61

Table 3. Regularization loss values for the ablation experiments

4.1.1 Contributions

We specifically create and train four different models to validate our contributions:

- **Baseline:** A single level VAE_{drap} at full resolution (512×256).
- **Incremental:** Pyramid architecture built from several VAE_{drap} submodels at different resolutions. Trained incrementally, one level at a time. In this approach, we start by training the lowest resolution level, then freeze it, add the next level, and train again. We repeat until we train all levels.
- **End-to-end:** Final pyramid model where all levels were trained at the same time. The condition encoders are still frozen. See details in the supplementary material.
- **End-to-end with regularization:** Final pyramid model where all levels were trained at the same time with the regularization losses (\mathcal{L}_{reg} in Sec. 3.3) applied during training. A $\lambda = 0.1$ was used.

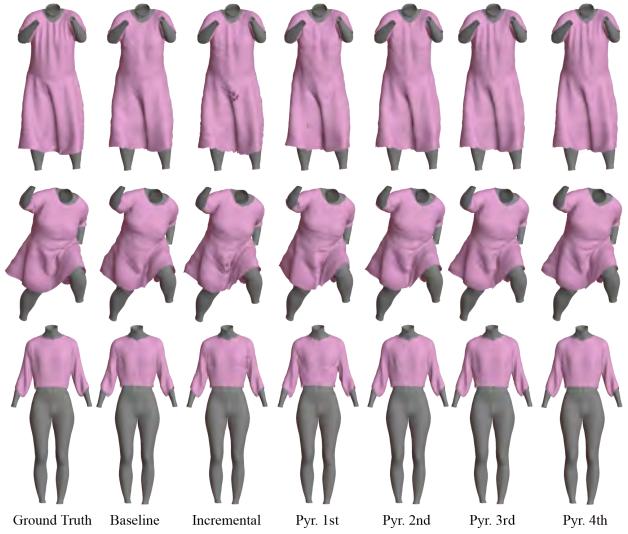


Figure 3. Effectiveness of the pyramid architecture on three examples. Ground truth, with baseline and last output of the incremental method, followed by the output of the final pyramid after each level after skinning and shaping. Notice that each subsequent layer adds details to the previous output.

We compare the contributions quantitatively and qualitatively in Tab. 2 and Fig. 3, respectively. The baseline results are comparable with the pyramid network when trained incrementally. In Fig. 3 it can be seen that the baseline produces good results, but lacks finer details.

Next, we see the output of the incremental method. The aim of this approach is to provide a stable training process and more control over the individual levels. In this case, higher levels do not affect the prediction of lower levels. From Fig. 3, we can see that while this approach achieves a low error, the outputs present some noise along with some artifacts. Since there is no feedback from higher levels if there is an error in one of the lower levels, this needs to be corrected later, which is not always possible.

Furthermore, we present our end-to-end trained model, where each pyramid level was trained simultaneously. By adding supervision at the output of each level, we can still have stable training, where each level outputs a valid prediction. In this approach, there is information flowing both ways between levels. As such, the higher levels can help the lower levels produce more accurate details which in turn lowers the amount of corrections needed to be learned in higher levels. Overall, end-to-end training helps to improve the results by a large margin over the incremental training. Fig. 3 shows that each new level does not only add more details to the garment, but it can also fix small errors or artifacts present in lower levels.

We also explore the use of different regularization losses during training as presented in Sec. 3.3. As we can see in Tab. 2, these losses help to improve the results while they have a higher impact on skirts compared to non-skirts. We further analyze the impact of these losses in Tab. 3. It can be seen that adding the regularization decreases the edge loss and the percentage of the cloth vertices colliding with the body. A lower value on the edge loss means that the model is more robust against compression and stretching of the topology of the mesh.

4.1.2 Sampling

First, we sample the normal maps from VAE_{norm} as these are not available at inference time. Further details about VAE_{norm} can be found in the supplementary material.

After sampling normal maps, we can also study how they can help VAE_{drap} in the generation of drapings dealing with the one-to-many assignment problem. By fixing the template garment and body pose, we randomly sample from VAE_{norm} and VAE_{drap} and generate garment UV maps. The results in Fig. 4 show diverse dynamics in different garment parts. We can also see that the surface areas of the samples only vary around 8%. These minor variations can be attributed to the stretching or compression of the garment. Therefore, the resulting garments negligibly

differ from their template shape.

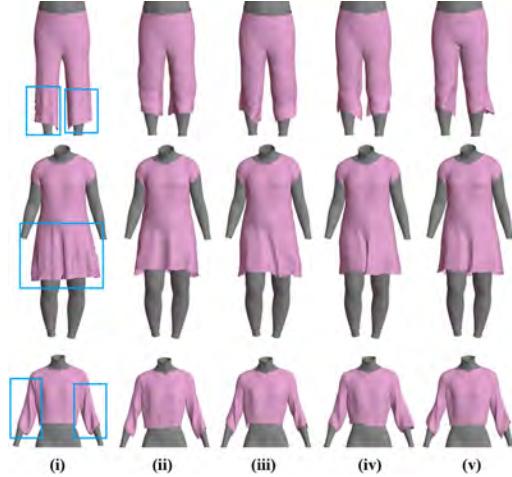


Figure 4. Sampling capabilities of the model. (i) ground truth, (ii-v) results by sampling the VAE latent spaces. Notice the produced changes in the blue rectangles. The shown examples have an average area difference of less than 8% compared to the ground truth.

4.2. SOTA comparison

We compare our final contribution to generative (CLOTH3D) and non-generative (DeePSD and DeepCloth) SOTA on the CLOTH3D dataset in Tab. 4. As can be seen, our methodology outperforms the SOTA by a large margin (79.46% relative improvement). We also use the qualitative metrics introduced at the beginning of Sec. 4 to compare our approach with HOOD and DeePSD. We follow the instructions in their GitHub pages and train these methods on the same CLOTH3D subset we used to train our model. As it can be seen in Tab. 6, we achieve a small area difference similar to DeePSD, while HOOD over-stretches the garment to favor the collision. DeepSD has the worst collision among others. HOOD shows the worst CD error. However, given that it is a self-supervised model, comparing it to the ground truth is not fair. We qualitatively compare with SOTA in Fig. 5. Although HOOD shows realistic draping, it drastically changes the template shape. In our experiments with HOOD, changing the elasticity hyperparameter made very small differences. DeePSD shows the least level of detail among others. Additionally, we observe an over-emphasized presence of physical forces, like gravity, in both HOOD and DeePSD. Overall, our method shows better qualitative results than SOTA.

We also compare our method to CAPE, a generative SOTA, on their dataset. As one can see in Tab. 5, our pyramid model progressively improves the results and we achieve a relative improvement of 40.5% on the male split, and 13.9% on the female split of the CAPE dataset at the third level of the pyramid. We also notice some overfitting towards the last level of the pyramid architecture. This can

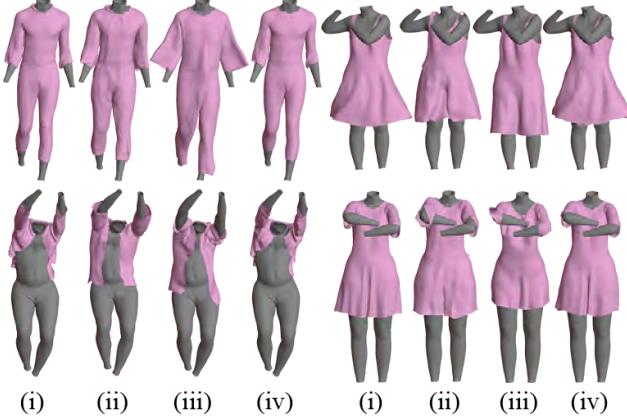


Figure 5. Qualitative comparison of ground truth (i) with predictions from DeePSD (ii), HOOD (iii) and our method (iv). HOOD shows improvement over DeePSD in the level of detail. However, it over-stretches the garment. Our method offers a high level of detail while having the minimum stretching and compression over the template shape.

Method	V2V error	Method	Female	Male
CLOTH3D [3]	29.00	Pyr. 1 lvl	6.63	8.94
DeePSD [7]	23.78	Pyr. 2 lvl	3.92	5.44
DeepCloth [38]	19.23	Pyr. 3 lvl	3.11	3.66
Ours	3.95	Pyr. 4 lvl	3.36	4.00
		CAPE [28]	3.61	6.15

Table 4. SOTA comparison on CLOTH3D dataset. We show large performance improvements over SOTA.

Table 5. SOTA comparison on CAPE dataset (V2V error). CAPE [28] numbers are taken from the authors’ GitHub page.

Method	CD error	Collision (%)	Area diff
DeePSD [7]	36.85	3.29	0.05429
HOOD [18]	42.99	0.44	0.13211
Ours	11.59	1.04	0.05282

Table 6. Qualitative metrics comparing with SOTA on CLOTH3D dataset.

be attributed to the fact that the meshes are lower resolution compared to CLOTH3D, and they contain fewer details. Qualitative results in Fig. 6 show how our approach reconstructs more cloth details while CAPE resembles the body.

4.3. Performance

Thanks to the convolutional architecture and having relatively light models, we can achieve fast inference times. The whole pyramid architecture has 40M parameters and without VAE encoders it decreases to 26M. Using an Nvidia 3090 GPU, we can achieve an inference time (for sampling and generation without encoders) of 50ms and 60ms using

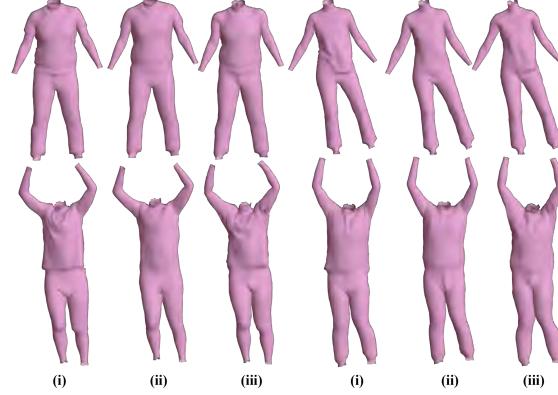


Figure 6. Qualitative examples from CAPE dataset comparing ground truth (i) to CAPE (ii) and ours (iii).

a batch size of 1 and 4, respectively. This allows near real-time inference.

5. Conclusions

We presented an effective approach to solve the problem of 3D garment draping and generation. By projecting the 3D meshes to UV images we were able to leverage established convolutional methods to create a pyramid VAE architecture. We observed how different pyramid levels allow the model to add different levels of detail to the generated cloth meshes. We further conditioned the network on normal maps, which provide additional regularization for the network to better focus on the details. Our results show SOTA performance on two public datasets, CAPE and CLOTH3D. We further showed that we can generalize to several garment types with a minimal number of models.

Limitations and future work. There are some limitations of our work that remain interesting for future research. Our current approach requires garment registration on top of the SMPL model as a preprocessing. The results also depend on the quality of the registration. Our model does not handle multi-layer garments. Finally, although we obtained SOTA results by using default SMPL blend weights, learning garment, and pose-dependent blend weights may provide extra improvements to surface quality.

6. Acknowledgement

This work has been partially supported by the Spanish projects PID2022-136436NB-I00, PDC2022-133305-I00, TED2021-131317B-I00, by ICREA under the ICREA Academia programme. The authors also acknowledge the support of the Spanish Ministry of Economy and Competitiveness (MINECO) and the European Regional Development Fund (ERDF) under Project No. PID2020-120611RB-I00/AEI/10.13039/501100011033.

References

[1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1175–1186, 2019. 1

[2] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1506–1515, 2022. 1

[3] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Cloth3d: clothed 3d humans. In *European Conference on Computer Vision*, pages 344–359. Springer, 2020. 1, 2, 3, 4, 5, 8

[4] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Pbns: Physically based neural simulator for unsupervised garment pose space deformation. *arXiv preprint arXiv:2012.11310*, 2020. 1, 2, 3

[5] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Pbns: Physically based neural simulation for unsupervised garment pose space deformation. *ACM Trans. Graph.*, 40(6), dec 2021. 1, 2

[6] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Neural cloth simulation. *ACM Transactions on Graphics (TOG)*, 41(6):1–14, 2022. 2, 3

[7] Hugo Bertiche, Meysam Madadi, Emilio Tylson, and Sergio Escalera. Deepsd: Automatic deep skinning and pose space deformation for 3d garment animation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5471–5480, 2021. 1, 2, 3, 5, 8

[8] 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. 1, 2

[9] Christopher P Burgess, Irina Higgins, Arka Pal, Loic Matthey, Nick Watters, Guillaume Desjardins, and Alexander Lerchner. Understanding disentangling in β -vae. *arXiv preprint arXiv:1804.03599*, 2018. 5

[10] Bindita Chaudhuri, Nikolaos Sarafianos, Linda Shapiro, and Tony Tung. Semi-supervised synthesis of high-resolution editable textures for 3d humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7991–8000, 2021. 1, 4

[11] Zhiqin Chen, Kangxue Yin, and Sanja Fidler. Auv-net: Learning aligned uv maps for texture transfer and synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1465–1474, 2022. 1, 4

[12] Wen-Huang Cheng, Sijie Song, Chieh-Yun Chen, Shintami Chusnul Hidayati, and Jiaying Liu. Fashion meets computer vision: A survey. *ACM Computing Surveys (CSUR)*, 54(4):1–41, 2021. 1

[13] 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. 1

[14] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 28, 2015. 2

[15] Garoe Dorta, Sara Vicente, Lourdes Agapito, Neill DF Campbell, Simon Prince, and Ivor Simpson. Laplacian pyramid of conditional variational autoencoders. In *Proceedings of the 14th European Conference on Visual Media Production (CVMP 2017)*, pages 1–9, 2017. 2

[16] Ruili Feng, Cheng Ma, Chengji Shen, Xin Gao, Zhenjiang Liu, Xiaobo Li, Kairi Ou, Deli Zhao, and Zheng-Jun Zha. Weakly supervised high-fidelity clothing model generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3440–3449, 2022. 1

[17] Artur Grigorev, Karim Iskakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. Stylepeople: A generative model of fullbody human avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5151–5160, 2021. 1

[18] Artur Grigorev, B. Thomaszewski, Michael J. Black, and Ottmar Hilliges. Hood: Hierarchical graphs for generalized modelling of clothing dynamics. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16965–16974, 2022. 1, 2, 3, 8

[19] 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*, pages 8739–8748, 2019. 1

[20] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11046–11056, 2021. 1

[21] Fangzhou Hong, Liang Pan, Zhongang Cai, and Ziwei Liu. Garment4d: Garment reconstruction from point cloud sequences. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6–14, 2021, virtual*, pages 27940–27951, 2021. 2

[22] Yang Hong, Juyong Zhang, Boyi Jiang, Yudong Guo, Ligang Liu, and Hujun Bao. Stereopifu: Depth aware clothed human digitization via stereo vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 535–545, 2021. 1

[23] Jianbin Jiang, Tan Wang, He Yan, and Junhui Liu. Clothformer: Taming video virtual try-on in all module. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10799–10808, 2022. 1

[24] Zorah Lahner, Daniel Cremers, and Tony Tung. Deepwrinkles: Accurate and realistic clothing modeling. In *Proceedings of the European conference on computer vision (ECCV)*, pages 667–684, 2018. 1, 4

[25] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. [2](#), [5](#), [11](#)

[26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. [3](#)

[27] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J Black. Scale: Modeling clothed humans with a surface codec of articulated local elements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16082–16093, 2021. [1](#), [2](#), [4](#)

[28] 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*, pages 6469–6478, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#)

[29] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10974–10984, 2021. [1](#), [2](#), [3](#), [4](#)

[30] Meysam Madadi, Hugo Bertiche, Wafa Bouzouita, Isabelle Guyon, and Sergio Escalera. Learning cloth dynamics: 3d+ texture garment reconstruction benchmark. In *NeurIPS (Competition and Demos)*, pages 57–76, 2020. [4](#)

[31] Sahib Majithia, Sandeep N Parameswaran, Sadhbhavana Babar, Vikram Garg, Astitva Srivastava, and Avinash Sharma. Robust 3d garment digitization from monocular 2d images for 3d virtual try-on systems. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3428–3438, 2022. [1](#)

[32] Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7365–7375, 2020. [1](#), [2](#), [3](#), [4](#), [5](#)

[33] Albert Rial-Farràs, Meysam Madadi, and Sergio Escalera. Uv-based reconstruction of 3d garments from a single rgb image. In *2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pages 1–8. IEEE, 2021. [1](#), [4](#)

[34] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [2](#)

[35] Jordi Sanchez-Riera, Albert Pumarola, and Francesc Moreno-Noguer. Physxnet: A customizable approach for learning cloth dynamics on dressed people. *2021 International Conference on 3D Vision (3DV)*, pages 879–888, 2021. [2](#)

[36] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Snug: Self-supervised neural dynamic garments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8140–8150, 2022. [2](#), [3](#)

[37] Igor Santesteban, Nils Thuerey, Miguel A Otaduy, and Dan Casas. Self-supervised collision handling via generative 3d garment models for virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11763–11773, 2021. [1](#), [2](#), [3](#)

[38] Zhaoqi Su, Tao Yu, Yangang Wang, and Yebin Liu. Deep-cloth: Neural garment representation for shape and style editing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [1](#), [4](#), [8](#)

[39] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11708–11718, 2021. [1](#), [2](#), [3](#)

[40] Z. Wang, E.P. Simoncelli, and A.C. Bovik. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, pages 1398–1402 Vol.2, 2003. [12](#)

[41] You Xie, Huiqi Mao, Angela Yao, and Nils Thuerey. Temporaluv: Capturing loose clothing with temporally coherent uv coordinates. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3450–3459, 2022. [1](#), [4](#)

[42] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13286–13296. IEEE, 2022. [1](#), [2](#)

[43] Chaojie Yang, Hanhui Li, Shengjie Wu, Shengkai Zhang, Haonan Yan, Nianhong Jiao, Jie Tang, Runnan Zhou, Xiaodan Liang, and Tianxiang Zheng. Bodygan: General-purpose controllable neural human body generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7733–7742, 2022. [1](#)

[44] Ilya Zakharkin, Kirill Mazur, Artur Grigorev, and Victor Lempitsky. Point-based modeling of human clothing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14718–14727, 2021. [1](#), [2](#), [3](#)

[45] Meng Zhang, Tuanfeng Wang, Duygu Ceylan, and Niloy J Mitra. Deep detail enhancement for any garment. In *Computer Graphics Forum*, volume 40, pages 399–411. Wiley Online Library, 2021. [1](#), [4](#)

[46] Fang Zhao, Wenhao Wang, Shengcui Liao, and Ling Shao. Learning anchored unsigned distance functions with gradient direction alignment for single-view garment reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12674–12683, 2021. [1](#)

[47] 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 *European Conference on Computer Vision*, pages 512–530. Springer, 2020. [1](#)

[48] Heming Zhu, Lingteng Qiu, Yuda Qiu, and Xiaoguang Han. Registering explicit to implicit: Towards high-fidelity garment mesh reconstruction from single images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3845–3854, 2022. [1](#)

A Generative Multi-Resolution Pyramid and Normal-Conditioning 3D Cloth Draping

Supplementary material

In this supplementary material, we provide further details on the architecture (Sec. A), training (Sec. B), building template UV coordinates (Sec. C), and additional qualitative results (Sec. D).

A. Architecture⁴

A.1. Baseline

As we showed in Fig. 1 in the main paper, VAE_{drap} is a conditional VAE based on ConvNext [25]. ConvNext is made up of stages, where after each stage the image or its representation is downsampled. Each stage can have a different depth, meaning the number of blocks it contains. We empirically update the number of blocks and features at each level. In this case, the depths are 3, 3, 3, 9, and 3 with 32, 64, 128, 256, and 128 ($= F$) features, respectively. The baseline receives an input of resolution 512×256 and the encoder ConvNext has five stages, this way the output features will have a size $8 \times 4 \times F$. The condition encoders work similarly. The dimensionality F is set to 128, 16, 32, and 80 for the VAE_{drap} latent vector, template garment, body, and normals condition encoders, respectively. The combined dimensionality of condition encoders is 128 to balance between conditioning features and VAE_{drap} latent vector. We also use a tanh activation on the conditioning latent codes to ensure we have the same scale in the feature space. Finally, all the latent codes are concatenated before feeding the decoder. For the decoder, we mirror the encoder architecture and replace the downsampling step with upsampling using transposed convolutions. The output size is the original 512×256 resolution. We use tanh activation on the last decoder layer.

A.2. Conditioning autoencoders

We explained conditioning encoders in the baseline architecture. However, these encoders are part of their own autoencoders, pretrained and frozen. That means one autoencoder is used to train each of the three conditioning inputs: template garment, posed body, and normal map. The encoder parts of these networks will act as the condition encoders. When incorporating them in the main pipeline, their weights are frozen to ensure that the quality of the features

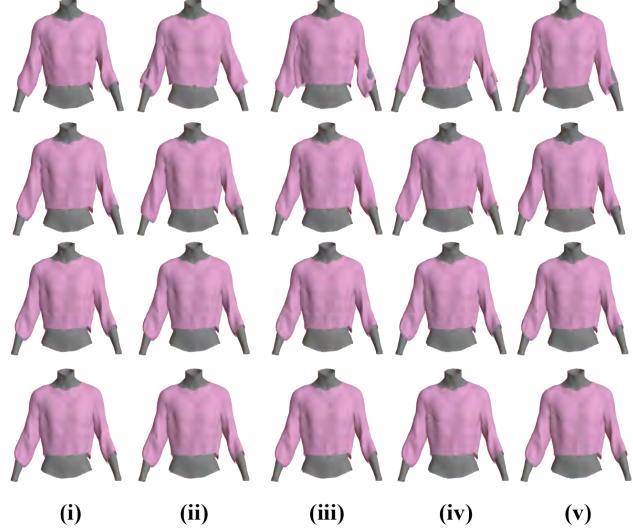


Figure A. We show the pyramid sampling at each level, one level per row. The first element of each row is the base for the next row.

remains intact. We use tanh activation on the last decoder layer.

The conditioning network for the normal map is the VAE_{norm} . The reason for this design is that normals are not available at inference time. Therefore, we create a generative pipeline for them such that we can sample from it at inference time. VAE_{norm} itself is conditioned on the template garment and posed body UV images. This is because we want the normal features to be disentangled from the other conditioning variables.

A.3. Pyramid

The pyramid model is built using several VAE_{drap} models of varying image resolutions. The main difference between them is the number of stages, which have to be decreased for lower levels of the pyramid. As such we use 5, 4, 3, and 2 stages for the resolutions 512×256 , 256×128 , 128×64 , and 64×32 , respectively. Similarly, we decrease the kernel sizes of the convolutions and use kernel sizes of 7, 5, 5, and 3 for the mentioned resolutions. Note that we keep the same feature dimension of $8 \times 4 \times 128$ for all the levels' latent codes.

The first level provides the base for the prediction which

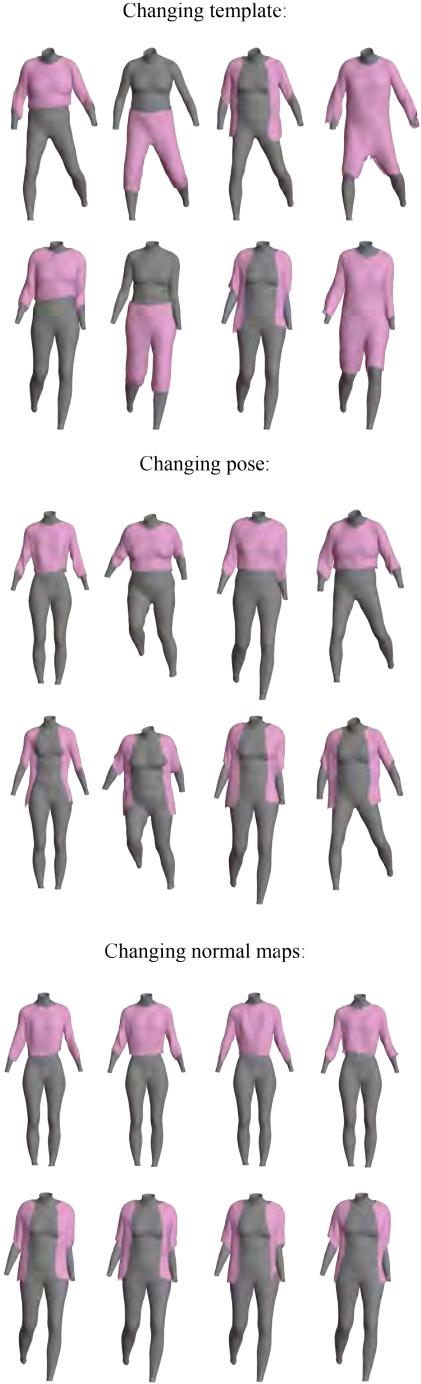


Figure B. We show the sampling process when changing only one condition. In each case, the stated condition is changing, and the other two and the latent code are fixed. For each case we have two examples as the two rows.

is a low-resolution UV image and all subsequent layers are added on top of this as offsets. To achieve this, the first level

receives the low-resolution ground truth as input, while later layers receive the offset between the ground truth and the previous level’s output along the conditioning inputs. The output’s resolution matches that of the input. In order to be able to add the offsets from the next level, this output needs to be upscaled. For this, we use a custom model to upscale the UV images (see sec. A.4). Repeating this at each level, we double the resolution after each level until we reach the final 512×256 resolution.

A.4. Upscaling model

As we mentioned in the main paper we design an upscaling model to deal with the values on the garment mask boundaries within the pyramid network. This model contains a network of 5 convolutional layers with relu activation, batch normalization, and 256, 128, 64, 32, and 3 feature sizes, respectively, and the kernel size follows the corresponding pyramid-level kernel size. The last convolution has a tanh activation instead of relu to predict the 3D coordinates. The upscaling procedure is as follows. First, the image is passed to the above network. This outputs an image that is multiplied with the background mask (inverse of the garment mask) and added to the input image. Finally, we resize this combined image using bilinear interpolation.

B. Training

Due to the number of modules to be trained, we first apply incremental training in which each sub-module is trained independently. Later we will try end-to-end training of the levels (conditioning encoders are still frozen). Next, we will present these sub-modules individually.

B.1. Conditiong autoencoders

The initial set of modules encodes the conditional inputs of the pipeline (as explained in sec. A.2). We train these networks once on the highest resolution. We use Multiscale-SSIM [40] loss for the autoencoders, and add an extra KL loss (using a β -VAE strategy) when training VAE_{norm} . Finally, we optimize the networks with a batch size of 4 and Adam optimizer with a learning rate 10^{-4} .

After the training, we are able to achieve a reconstruction error of 1.1mm with the posed body and 9.3mm with the template garment autoencoders. Note that the available data for the template garments is only 1.6 thousand outfits which are further split into training and validation data. As for the VAE_{norm} , we achieve an MS-SSIM value of 0.993 with an L1 error of 0.04.

B.2. Pyramid

In the pyramid pipeline, each level is trained individually in the incremental approach and together in the end-to-end approach. The advantages of both cases are detailed

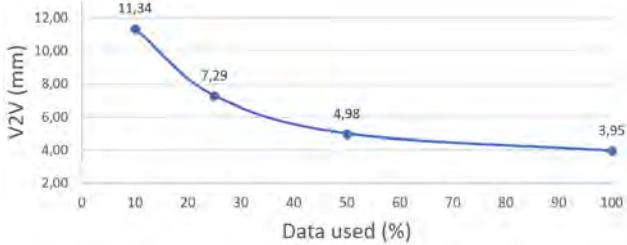


Figure C. Generalization capability of our pyramid network on CLOTH3D dataset as a function of the amount of training data used.

in the main paper Sec. 3.1.1. During training, each level receives the same conditional features generated from the high-resolution conditional input, thereby avoiding the need to train three encoders per level and generate separate conditions for each level. This has significant speed improvements. For the variational encoders, each condition and the unposed garment are downsampled to the given resolution. The decoder outputs an image at the same resolution which is upscaled to be added to the next level’s output.

The losses used are described in detail in the main paper. Note that we only start using the normal loss after 2000 iterations, since otherwise, it makes the initial steps of the training highly unstable.

We optimize the network with a batch size of 4 and Adam optimizer with a learning rate 10^{-4} .

C. Template UV coordinates

A template UV map is obtained by projecting a three-dimensional body mesh to a fixed two-dimensional shape. To construct this mapping, we start by annotating the body mesh vertices according to the front and back of the body. This is followed by defining some keypoints: 1) around the ends of legs and arms, 2) around the neck, and 3) bilaterally for knees, hips, and underarms. We also create a low-resolution 2D mesh as a template for the UV map. We define pairs of previously selected key points and corresponding points on the UV map template. Using these control points we apply thin plate spline interpolation (TPS) to morph the 3D vertices to the 2D template (a 3D flat surface) once for the front and then for the back vertices. The vertices along the border between the front and back belong to both parts, so they have to be duplicated. During the 3D reconstruction, these duplicated pixels will be averaged to obtain the 3D coordinate.

To further improve the construction of UV maps, we define more keypoints automatically. We take the vertices of the previously mentioned border and find the closest vertex on the 2D template. This vertex is obtained by calculating distances between the points and also from their normal vectors, the latter ensuring that the mesh is stretched

outwards from the center to the template. With these new keypoint pairs, we apply TPS again. Finally, we obtain UV mapping that maximizes the data representation by covering a large portion of the given rectangular grid (as shown in Fig. 2(b) in the main paper).

D. Results

D.1. Pyramid sampling

We demonstrate how the sampling process works in the pyramid architecture. Given a set of conditioning variables, which are the template garment, posed body, and normal maps, we sample the VAE latent space, concatenate the encoded conditions, and pass it to the decoder. When we do this at the first pyramid level, we get some potential base garments as seen in Fig. A. After this, we sample from the second level and generate the offsets over the first level, which adds some smaller details. The contribution of each level can be seen in the figure. In the case shown, after the second level, the changes are minimal. To further demonstrate the sampling possibilities, we show examples of only one of the conditions changing to better demonstrate how it works. This can be seen in Fig. B.

D.2. Normal map sampling

The aim of the normals is to provide guidance and variability for the generation of samples. We study the VAE_{norm} model’s ability to generate normal maps. We can see some examples in Fig. D which show how we can generate multiple plausible normal maps for a given input.

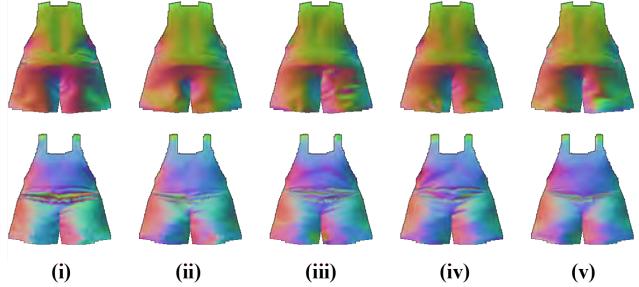


Figure D. Examples of sampling normal maps. The left (i) is the original, and the four on the right (ii-v) are randomly sampled from the VAE_{norm} model. We can see high variance while they still represent realistic deformations.

D.3. Generalization

To show the generalization and representation capacity of our model we evaluate how the error scales when using a smaller portion of the training data. We run the training with the following portions of the data: 0.5, 0.25, and 0.1⁵.

⁵Note these portions are not applied on the whole CLOTH3D dataset but the 20% we initially selected in this paper.

As seen in the error trend in Fig. C, we can achieve state-of-the-art with only 10 percent of the data. It also shows we can scale the method by increasing the amount of data, but the improvements will diminish soon.

D.4. Additional example results

We provide some additional qualitative examples from the CLOTH3D dataset in Fig. E and Fig. F. Our model is able to reconstruct a high level of detail on a variety of different garments. We also show some additional comparison with HOOD in Fig. G.



Figure E. Additional qualitative examples from the CLOTH3D dataset: ground truth and proposed model's reconstruction.

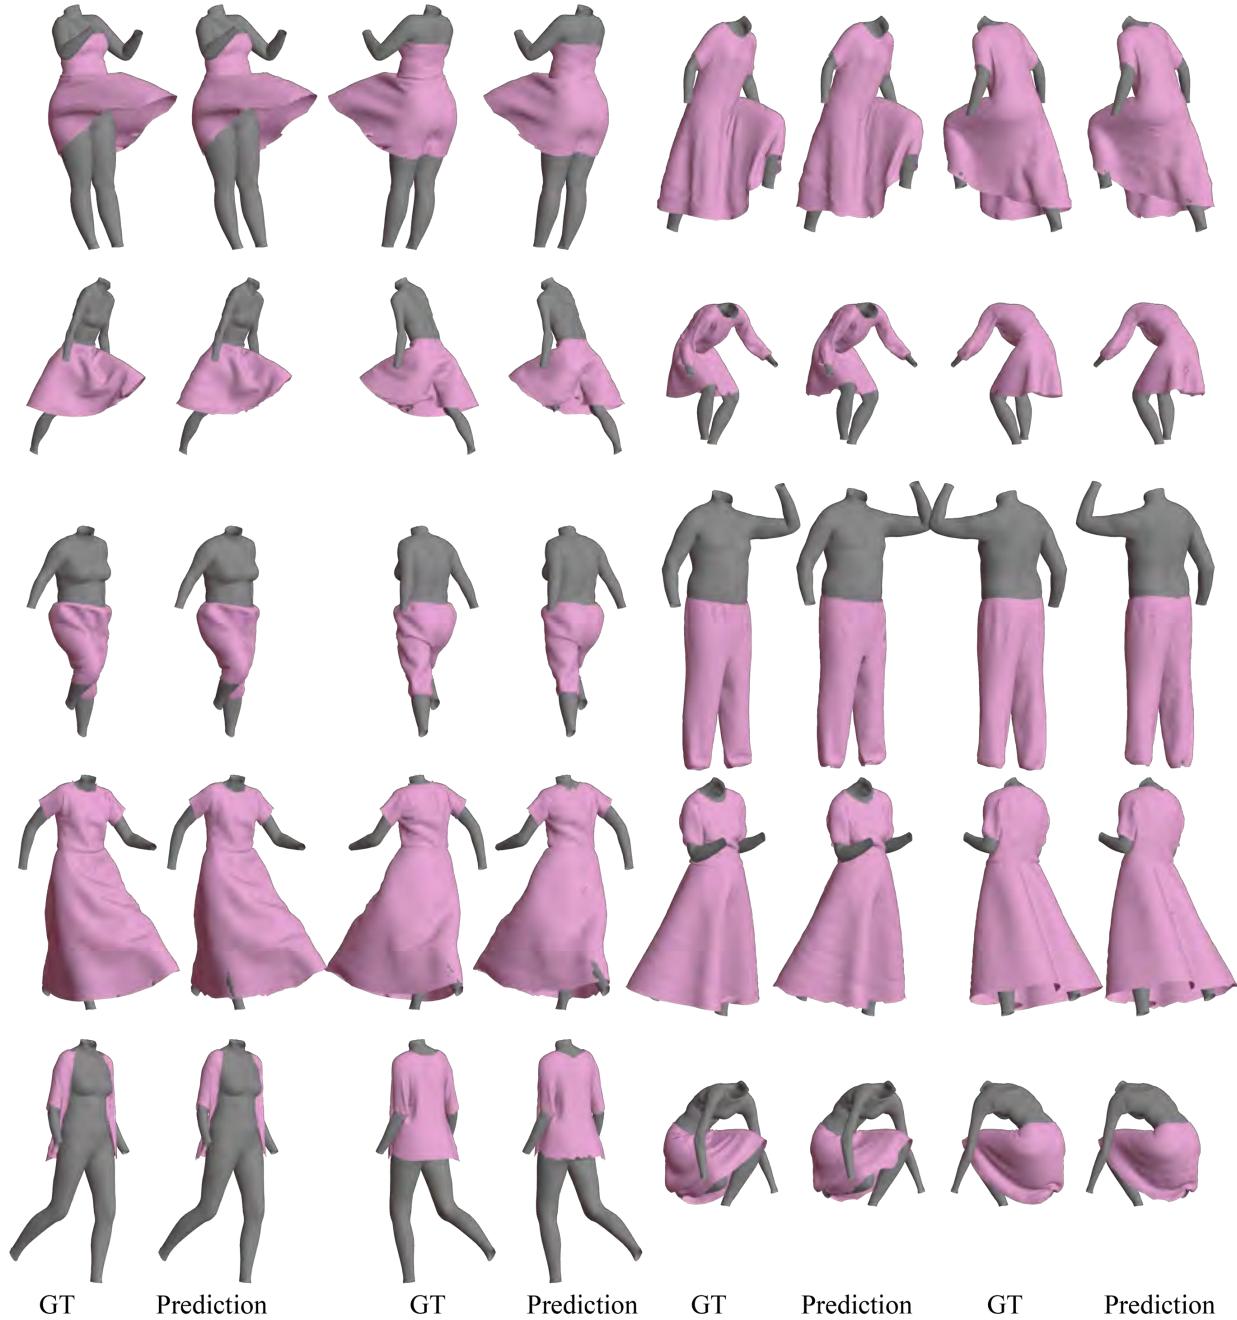


Figure F. Additional qualitative examples from the CLOTH3D dataset: ground truth and proposed model's reconstruction.

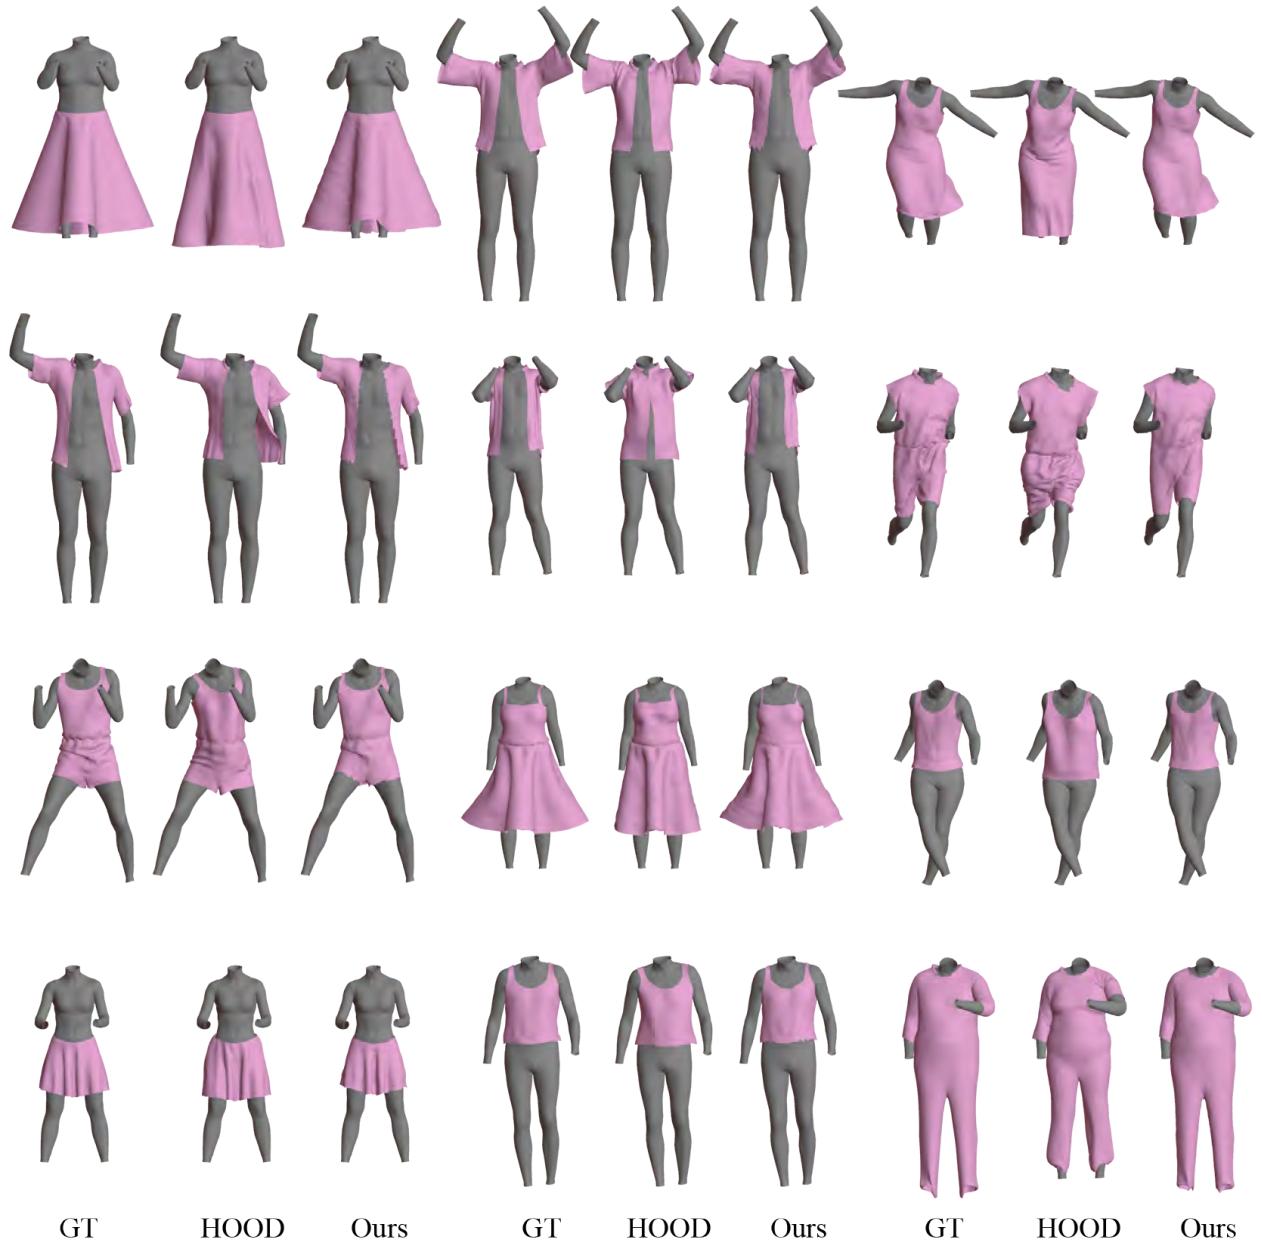


Figure G. Additional qualitative comparison examples with HOOD, where we can see the ground truth, HOOD prediction, and ours.