

DeepWrinkles: Accurate and Realistic Clothing Modeling

Zorah Lähler^{1,2}, Daniel Cremers², and Tony Tung¹

¹ Facebook Reality Labs, USA, tony.tung@fb.com

² Technical University Munich, Germany, {laehner,cremers}@in.tum.de

Abstract. We present a novel method to generate accurate and realistic clothing deformation from real data capture. Previous methods for realistic cloth modeling mainly rely on intensive computation of physics-based simulation (with numerous heuristic parameters), while models reconstructed from visual observations typically suffer from lack of geometric details. Here, we propose an original framework consisting of two modules that work jointly to represent global shape deformation as well as surface details with high fidelity. Global shape deformations are recovered from a subspace model learned from 3D data of clothed people in motion, while high frequency details are added to normal maps created using a conditional Generative Adversarial Network whose architecture is designed to enforce realism and temporal consistency. This leads to unprecedented high-quality rendering of clothing deformation sequences, where fine wrinkles from (real) high resolution observations can be recovered. In addition, as the model is learned independently from body shape and pose, the framework is suitable for applications that require retargeting (e.g., body animation). Our experiments show original high quality results with a flexible model. We claim an entirely data-driven approach to realistic cloth wrinkle generation is possible.

Keywords: 3D surface deformation modeling · cloth simulation · normal maps · deep neural networks.

1 Introduction

Realistic garment reconstruction is notoriously a complex problem and its importance is undeniable in many research work and applications, such as accurate body shape and pose estimation in the wild (i.e., from observations of clothed humans), realistic AR/VR experience, movies, video games, virtual try-on, etc.

For the past decades, physics-based simulations have been setting the standard in movie and video game industries, even though they require hours of labor by experts. More recently methods for full clothing reconstruction using multi-view videos or 3D scan systems have also been proposed [38]. Global deformations can be reconstructed with high fidelity semi-automatically. Nevertheless, accurately recovering geometric details such as fine cloth wrinkles has remained a challenge.

In this paper, we present DeepWrinkles (see Fig. 1), a novel framework to generate accurate and realistic clothing deformation from real data capture. It

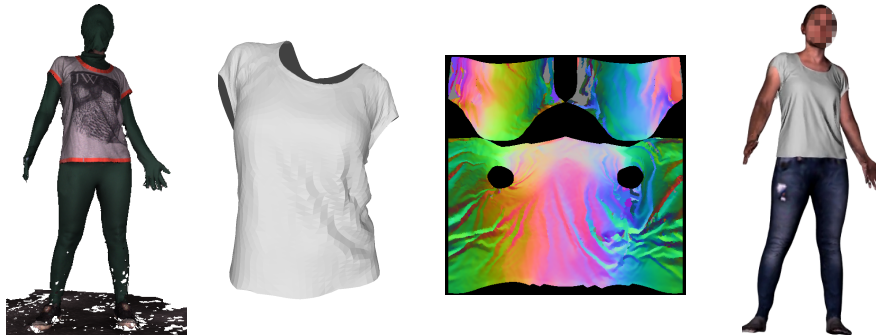


Fig. 1. Accurate and realistic clothing modeling with DeepWrinkles, our entirely data-driven framework. (Left) 4D data capture. (Middle Left) Reconstruction from subspace model. (Middle Right) Fine wrinkles in normal map generated by our adversarial neural network. (Right) 3D rendering and animation on virtual human.

consists of two complementary modules: (1) A statistical model is learned from 3D scans of clothed people in motion, from which clothing templates are precisely non-rigidly aligned. Clothing shape deformations are therefore modeled using a linear subspace model, where human body shape and pose are factored out, hence enabling body retargeting. (2) Fine geometric details are added to normal maps generated using a conditional adversarial network whose architecture is designed to enforce realism and temporal consistency.

To our knowledge, this is the first method that tackles 3D surface geometry refinement using deep neural network on normal maps. With DeepWrinkles, we obtain unprecedented high-quality rendering of clothing deformation, where global shape as well as fine wrinkles from (real) high resolution observations can be recovered, using an entirely data-driven approach. Figure 2 gives an overview of our framework with a T-shirt as example. Additional materials contain videos of results. We show how the model can be applied to virtual human animation, with body shape and pose retargeting.

2 Related Work

Cloth modeling and garment simulation have a long history that dates back to the mid 80s. A general overview of fundamental methods is given in [51]. There are two mostly opposing approaches to this problem. One is using physics-based simulations to generate realistic wrinkles, and the other captures and reconstructs details from real-world data.

Physics-based simulation. For the past decades, models relying on Newtonian physics have been widely applied to simulate cloth behavior. They usually model various material properties such as stretch (tension), stiffness, and weight. For certain types of applications (e.g., involving human body) additional models or

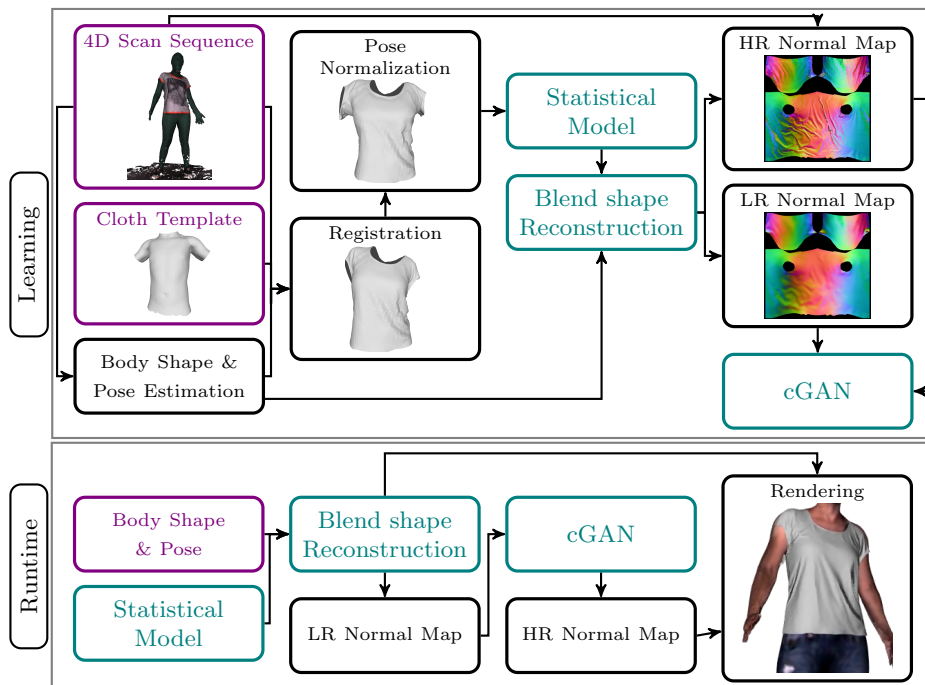


Fig. 2. Outline of DeepWrinkles. During Learning, we learn to reconstruct global shape deformations using a Statistical Model and a mapping from pose parameters to Blend shape parameters using real-world data. We also train a neural network (cGAN) to generate fine details on normal maps from lower resolution ones. During runtime, the learned models reconstruct shape and geometric details given a priori body shape and pose. Inputs are in violet, learned models are in cyan.

external forces have to be taken into account such as body kinematics, body surface friction, interpenetration, etc [9, 8, 4, 16]. Note that several models have been integrated in commercial solutions (e.g., Unreal Engine APEX Cloth/Nvidia NvCloth, Unity Cloth, Maya nCloth, MarvelousDesigner, OptiTex, etc.) [32]. Nevertheless, it typically requires hours or days if not weeks of computation, retouching work, and parameter tuning by experts to obtain realistic cloth deformation effects.

3D capture and reconstruction. Vision-based approaches have explored ways to capture cloth surface deformation under stress, and estimate material properties through visual observations for simulation purpose [53, 34, 33, 29]. As well, several methods directly reconstruct whole object surface from real-world measurements. [54] uses texture patterns to track and reconstruct garment from a video, while 3D reconstruction can also be obtained from multi-view videos without markers [48, 33, 29, 41, 49, 7]. However without sufficient prior,

reconstructed geometry can be quite crude. When the target is known (e.g., clothing type), templates can improve the reconstruction quality [3]. Also, more details can be recovered by applying super-resolution techniques on input images [47, 17, 45, 7], or using photometric stereo and information about lighting [22, 50]. Naturally, depth information can lead to further improvement [36, 14].

In recent work [38], cloth reconstruction is obtained by clothing segmentation and template registration from 4D scan data. Captured garments can be retargeted to different body shapes. However the method has limitations regarding fine wrinkles.

Coarse-to-fine approaches. To reconstruct fine details, and consequently handle the bump in resolution at runtime (i.e., higher resolution meshes or more particles for simulation), methods based on dimension reduction (e.g., linear subspace models) [2, 21] or coarse-to-fine strategies are commonly applied [52, 35, 25]. DRAPE [19] automates the process of learning linear subspaces from simulated data and applying them to different subjects. The model factors out body shape and pose to produce a global cloth shape and then applies the wrinkles of seen garments. However, deformations are applied per triangle as in [42], which is not optimal for online applications. Additionally, for all these methods, simulated data is tedious to generate, and accuracy and realism are limited.

Learning methods. Previously mentioned methods focus on efficient simulation and representation of previously seen data. Going a step further, several methods have attempted to generalize this knowledge to unseen cases. [46] learns bags of dynamical systems to represent and recognize repeating patterns in wrinkle deformations. In DeepGarment [13] the global shape and low frequency details are reconstructed from a single segmented image using a CNN but no retargeting is possible.

Only sparse work has been done on learning to add realistic details to 3D surfaces with neural networks but several methods to enrich facial scans with texture exist [40, 37]. In particular, Generative Adversarial Networks (GANs) [18] are suitable to enhance low dimensional information with details. In [27] it is used to create realistic images of clothed people given a (possibly random) pose.

Outside of clothing, SR-GAN [28] solves the super-resolution problem of recovering photo-realistic textures from heavily downsampled images on public benchmark. The task has similarities to ours in generating high frequency details for coarse inputs but we use a content loss motivated by perceptual similarity instead of similarity in pixel space. [10] uses a data-driven approach with a CNN to simulate highly detailed smoke flows. Instead, pix2pix [24] proposes a conditional GAN that creates realistic images from sketches or annotated regions or vice versa. This design suits our problem better as we aim at learning and transferring underlying image structure.

In order to represent the highest possible level of detail at runtime, we propose to revisit the traditional rendering pipeline of 3D engine with computer vision.

Our contributions take advantage of the normal mapping technique [26, 12, 11]. Note that displacement maps have been used to create wrinkle maps using texture information [5, 15]. However, while results are visually good on faces, they still require high resolution mesh, and no temporal consistency is guaranteed across time. (Also, faces are arguably less difficult to track than clothing which are prone to occlusions and more loose.)

In this work, we claim the first entirely data-driven method that uses a deep neural network on normal maps to leverage 3D geometry of clothing.

3 Deformation Subspace Model

We model cloth deformations by learning a linear subspace model that factors out body pose and shape, as in [19]. However, our model is learned from real data, and deformations are applied per vertex for speed and flexibility regarding graphics pipelines [31]. Our strategy ensures deformations are represented compactly and with high realism. First, we compute robust template-based non-rigid registrations from a 4D scan sequence (Sect. 3.1), then a clothing deformation statistical model is derived (Sect. 3.2) and finally, a regression model is learned for pose retargeting (Sect. 3.3).

3.1 Data preparation

Data capture. For each type of clothing, we capture 4D scan sequences at 60 fps (e.g., 10.8k frames for 3 min) of a subject in motion, and dressed in a full-body suit with one piece of clothing with colored boundaries on top. Each frame consists of a 3D surface mesh with around 200k vertices yielding very detailed folds on the surface but partially corrupted by holes and noise (see Fig. 1a). This setup allows a simple color-based 3D clothing extraction. In addition, capturing only one garment prevents occlusions where clothing normally overlaps (e.g., waistbands) and clothings can be freely combined with each other.

Body tracking. 3D body pose is estimated at each frame using a method in the spirit of [44]. We define a skeleton with j joints described by p_j parameters representing transformation and bone length. Joint parameters are also adjusted to body shape, which is estimated using [31, 55]. Posed human body is obtained using a linear blend skinning function $S: \mathbb{R}^{3 \times v} \times \mathbb{R}^{p_j} \rightarrow \mathbb{R}^{3 \times v}$ that transforms (any subset of) v vertices of a 3D deformable human template in normalized pose (e.g., T-pose) to a pose defined by j skeleton joints.

Registration. We define a template of clothing $\bar{\mathcal{T}}$ by choosing a subset of the human template with consistent topology. $\bar{\mathcal{T}}$ should contain enough vertices to model deformations (e.g., 5k vertices for a T-shirt), as shown in Fig. 3. The clothing template is then registered to the 4D scan sequence using a variant of non-rigid ICP based on grid deformation [30, 20]. The following objective function




Fig. 3. (Left) Before registration, a template is aligned to clothing scans by skinning. Boundaries are misplaced. (Right) Examples of registrations with different shirts on different people. Texture stability across sequence shows the method is robust to drift.

\mathcal{E}_{reg} , which aims at optimizing affine transformations of grid nodes, is iteratively minimized using Gauss-Newton method:

$$\mathcal{E}_{reg} = \mathcal{E}_{data} + \omega_r \cdot \mathcal{E}_{rigid} + \omega_s \cdot \mathcal{E}_{smooth} + \omega_b \cdot \mathcal{E}_{bound}, \quad (1)$$

where the data term \mathcal{E}_{data} aligns template vertices with their nearest neighbors on the target scans, \mathcal{E}_{rigid} encourages each triangle deformation to be as rigid as possible, and \mathcal{E}_{smooth} penalizes inconsistent deformation of neighboring triangles. In addition, we introduce the energy term \mathcal{E}_{bound} to ensure alignment of boundary vertices, which is unlikely to occur otherwise (see below for details). We set $\omega_r = 500$, $\omega_s = 500$, and $\omega_b = 10$ by experiments. One template registration takes around 15s (using CPU only).

Boundary Alignment. During data capture the boundaries of the clothing are marked in a distinguishable color and corresponding points are assigned to the set \mathcal{B}_S . We call the set of boundary points on the template \mathcal{B}_T . Matching point pairs in $\mathcal{B}_S \times \mathcal{B}_T$ should be distributed equally among the scan and template, and ideally capture all details in the folds. As this is not the case if each point in \mathcal{B}_T is simply paired with the closest scan boundary point (see Fig. 4), we select instead a match $s_t \in \mathcal{B}_S$ for each point $t \in \mathcal{B}_T$ via the following formula:

$$s_t = \max_{s \in \mathcal{C}} \|t - s\| \quad \text{with} \quad \mathcal{C} = \left\{ s' \in \mathcal{B}_S \mid \arg \min_{t' \in \mathcal{B}_T} \|s' - t'\| = t \right\}. \quad (2)$$

Notice that \mathcal{C} might be empty. This ensures consistency along the boundary and better captures high frequency details (which are potentially further away).

3.2 Statistical model

The statistical model is computed using linear subspace decomposition by PCA [31]. Poses $\{\theta_1, \dots, \theta_n\}$ of all n registered meshes $\{\mathcal{R}_1, \dots, \mathcal{R}_n\}$ are factored out from the model by pose-normalization using inverse skinning: $S^{-1}(\mathcal{R}_i, \theta_i) = \bar{\mathcal{R}}_i \in \mathbb{R}^{3 \times v}$. In what remains, meshes in normalized pose are marked with a bar. Each registration $\bar{\mathcal{R}}_i$ can be represented by a mean shape $\bar{\mathcal{M}}$ and vertex offsets o_i , such that $\bar{\mathcal{R}}_i = \bar{\mathcal{M}} + o_i$, where the mean shape $\bar{\mathcal{M}} \in \mathbb{R}^{3 \times v}$ is obtained by

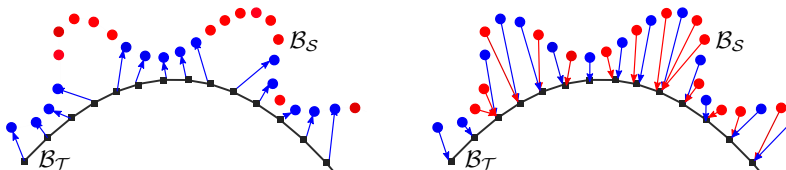


Fig. 4. Strategies for Boundary Alignment. Template boundary points \mathcal{B}_T are denoted in black, scan boundary points \mathcal{B}_S (with significant more points than the template) are in red and blue. Points that are paired with a template point are in blue. (Left) Pairing to the closest neighbor of the template points leads to ignorance of distant details. (Right) Each template point in \mathcal{B}_T is paired with the furthest point (marked in blue) in a set containing its closest points in \mathcal{B}_S .

averaging vertex positions: $\bar{\mathcal{M}} = \sum_{i=1}^n \frac{\bar{\mathcal{R}}_i}{n}$. The n principal directions of the matrix $O = [o_1 \dots o_n]$ are obtained by singular value decomposition: $O = U \Sigma V^T$. Ordered by the largest singular values, the corresponding singular vectors contain information about the most dominant deformations.

Finally, each \mathcal{R}_i can be compactly represented by $k \leq n$ parameters $\{\lambda_1^i, \dots, \lambda_k^i\} \in \mathbb{R}^k$ (instead of its $3 \times v$ vertex coordinates), with the linear blend shape function B , given a pose θ_i :

$$B(\{\lambda_1^i, \dots, \lambda_k^i\}, \theta_i) = S \left(\bar{\mathcal{M}} + \sum_{l=0}^k \lambda_l^i \cdot V_l, \theta_i \right) \approx \mathcal{R}_i \in \mathbb{R}^{3 \times v}, \quad (3)$$

where V_l is the l -th singular vector. For a given registration, $\lambda_l^i = V_l^T \bar{\mathcal{R}}_i$ holds. In practice, choosing $k = 500$ is sufficient to represent all registrations with a negligible error (less than 5 mm).

3.3 Pose-to-shape prediction

We now learn a predictive model f , that takes as inputs j joint poses, and outputs a set of k shape parameters Λ . This allows powerful applications where deformations are induced by pose. To take into account deformation dynamics that occur during human motion, the model is also trained with pose velocity, acceleration, and shape parameter history. These inputs are concatenated in the control vector Θ , and f can be obtained using autoregressive models [2, 39, 31].

In our experiments with clothing, we solved for f in a straightforward way by linear regression: $F = \Lambda \cdot \Theta^\dagger$, where F is the matrix representation of f , and \dagger indicates the Moore-Penrose inverse. While this allows for (limited) pose retargeting, we observed loss in reconstruction details. One reason is that under motion, the same pose can give rise to various configurations of folds depending on the direction of movement, speed and previous fold configurations.

To obtain non-linear mapping, we consider the components of Θ and Λ as multivariate time series, and train a deep multi-layer recurrent neural network




Fig. 5. Limits of registration and subspace model. (Left) Global shape is well recovered, but many visible (high frequency) details are missing. (Right) Increasing the resolution of the template mesh is still not sufficient. Note that [38] suffers from the same limitations.

(RNN) [43]. A sequence-to-sequence encoder-decoder architecture with Long Short-term Memory (LSTM) units is well suited as it allows continuous predictions, while being easier to train than RNNs and outperforming shallow LSTMs. We compose Θ with j joint parameter poses, plus velocity and acceleration of the joint root. MSE compared to linear regression are reported in Sect. 5.3.

4 Fine Wrinkle Generation

Our goal is to recover all observable geometric details. As previously mentioned, template-based methods [38] and subspace-based methods [21, 19] cannot recover every detail such as fine cloth wrinkles due to resolution and data scaling limitations, as illustrated in Fig. 5.

Assuming the finest details are captured at sensor image pixel resolution, and are reconstructed in 3D (e.g., using a 4D scanner as in [6, 38]), all existing geometric details can then be encoded in a normal map of the 3D scan surface at lower resolution (see Fig 6). To automatically add fine details *on the fly* to reconstructed clothing, we propose to leverage normal maps using a generative adversarial network [18]. See Figure 8 for the architecture. In particular, our network induces temporal consistency on the normal maps to increase realism in animation applications.

4.1 Data preparation

We take as inputs a 4D scan sequence, and a sequence of corresponding reconstructed garments. The latter can be either obtained by registration, reconstruction using blend shape or regression, as detailed in Sect. 3. Clothing template meshes \mathcal{T} are equipped with UV maps which are used to project any pixel from an image to a point on a mesh surface, hence assigning a property encoded in a pixel to each point. Therefore, normal coordinates can be normalized and stored as pixel colors in normal maps. Our training dataset then consists of pairs of normal maps (see Fig. 7): *low resolution* (LR) normal maps obtained by blend shape reconstruction, and *high resolution* (HR) normal maps obtained




Fig. 6. All visible details from an accurate 3D scan are generated in our normal map for incredible realism. Here, a virtual shirt is seamlessly added on top of an animated virtual human (e.g., scanned subject).




Fig. 7. Examples of our dataset. (Left) Low resolution input normal map, (Middle) High resolution target normal map from scan. Details, and noise, visible on the scan are reproduced in the image. Gray areas indicate no normal information was available on the scan. (Right) T-Shirt on a human model rendered without and with normal map.

from the scans. For LR normal maps, the normal at surface point (lying in a face) is linearly interpolated from vertex normals. For HR normal maps, per-pixel normals are obtained by projection of the high resolution observations (i.e., 4D scan) onto triangles of the corresponding low resolution reconstruction, and then the normal information is transferred using the UV map of $\bar{\mathcal{T}}$. Note that normal maps cannot be directly calculated from scans because neither is the exact area of the garment defined, nor are they equipped with UV map. Also, our normals are represented in global coordinates, as opposed to tangent space coordinates as is standard for normal maps. The reason is that LR normal maps contain no additional information to the geometry and are therefore constant in tangent space. This makes them unfitting for conditioning our adversarial neural network.

4.2 Network architecture

Due to the nature of our problem it is natural to explore network architectures designed to enhance images (i.e., super-resolution applications). From our experiments, we observed that models trained on natural images, including those

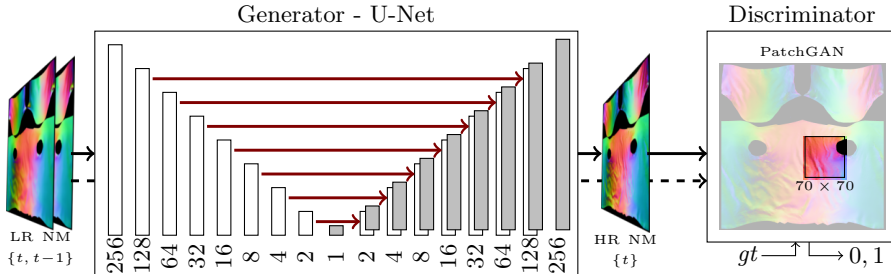


Fig. 8. cGAN for realistic HR normal map generation from LR normal maps as input. Layer sizes are squared. Skip connections (red) in U-Net preserve underlying image structure across network layers. PatchGAN enforces wrinkle pattern consistency.

containing a perceptual loss term fail (e.g., SR-GAN [28]). On the other hand, cloth deformations exhibit smooth patterns (wrinkles, creases, folds) that deform continuously in time. In addition, at a finer level, materials and fabric texture also contain high frequency details.

Our proposed network is based on a conditional Generative Adversarial Network (cGAN) inspired from image transfer [24]. We also use a convolution-batchnorm-ReLu structure [23] and a U-Net in the generative network since we want latent information to be transferred across the network layers and the overall structure of the image to be preserved. This happens thanks to the skip connections. The discriminator only penalizes structure at the scale of patches, and works as a texture loss. Our network is conditioned by low-resolution normal map images (size: 256×256) which will be enhanced with fine details learned from our real data normal maps. See Fig. 8 for the complete architecture.

Temporal consistency is achieved by extending the L_1 network loss term. For compelling animations, it is not only important that each frame looks realistic, but also no sudden jumps in the rendering should occur. To ensure smooth transition between consecutively generated images across time, we introduce an additional loss \mathcal{L}_{loss} to the GAN objective that penalizes discrepancies between generated images at t and expected images (from training dataset) at $t - 1$:

$$\mathcal{L}_{loss} = \underbrace{\|\mathcal{I}_{gen}^t - \mathcal{I}_{gt}^t\|_1}_{\mathcal{L}_{data}} + \underbrace{\left| \sum_{i,j} (\mathcal{I}_{gen}^t - \mathcal{I}_{gt}^{t-1})_{i,j} \right|}_{\mathcal{L}_{temp}}, \quad (4)$$

where \mathcal{L}_{data} helps to generate images near to ground truth in an L_1 sense (for less blurring). The temporal consistency term \mathcal{L}_{temp} is meant to capture global fold movements over the surface. If something appears somewhere, most of the time, it should have disappeared close-by and vice versa. Our term does not take spatial proximity into account though. We also tried temporal consistency based on the L_1 - and L_2 -norm, and report the results in Table 1. See Fig. 9 for a comparison of results with and without the temporal consistency term.




Fig. 9. Examples trained on only 2000 training samples reinforce the effect of the additional loss \mathcal{L}_{temp} . The pairs show twice the same consecutive frames: (left) *without* temporal consistency term geometric noise appears or disappears instantly, and (right) *with* temporal consistency term preserves geometric continuity.

5 Experiments

This section evaluates the results of our reconstruction. 4D scan sequences were captured using a temporal-3dMD system (4D) [1]. Sequences are captured at 60fps. Each frame consists of a colored mesh with 200K vertices. Here, we show results on two different shirts (for female and male). We trained the cGAN network on a dataset of 9213 consecutive frames. The first 8000 images compose the training data set, the next 1000 images the test data set and the remaining 213 images the validation set. Test and validation sets contain poses and movements not seen in the training set. The U-Net auto-encoder is constructed with 2×8 layers, and 64 filters in each of the first convolutional layers. The discriminator uses patches of size 70×70 . \mathcal{L}_{data} weight is set to 100, \mathcal{L}_{temp} weight is 50, while GAN weight is 1. The images have a resolution of 256×256 , although our early experiments also showed promising results on 512×512 .

5.1 Comparison of approaches

We compare our results to different approaches (see Fig. 10). A physics-based simulation done by a 3D artist using MarvelousDesigner [32] returns a mesh imitating similar material properties as our scan and with a comparable amount of folds but containing 53,518 vertices (i.e., an order of magnitude more). A linear subspace reconstruction with 50 coefficients derived from the registrations (Sect. 3) produces a mostly flat surface, while the registration itself shows smooth approximations of the major folds in the scan. Our method, DeepWrinkles, adds all high frequency details seen in the scan to the reconstructed surface. These three methods use a mesh with 5,048 vertices. DeepWrinkles is shown with a 256×256 normal map image.

5.2 Importance of reconstruction details in input

Our initial experiments showed promising results reconstructing details from the original registration normal maps. To show the efficacy of the method it is




Fig. 10. Comparison of approaches. a) Physics-based simulation [32], b) Subspace (50 coefficients) [19], c) Registration [38], d) DeepWrinkles (ours), e) 3D scan (ground truth).

	$L1$ temp	$L2$ temp	Eq. 4	no temp	Registr.	BS 500	BS 200	Regre.
Data	4.63	5.16	3.72	5.06	3.72	6.86	6.40	7.6
Temporal	5.11	4.2	4.2	5.52	4.2	7.05	6.52	7.65

Table 1. Comparison of pixel-wise error values of the neural network for different training types. Data and Temporal are as defined in Eq. 4. (Left) Different temporal consistency terms. $L1$ and $L2$ take the respective distance between the output and target at time $t - 1$. (Right) Different reconstruction methods to produce the input normal map. Registr. refers to registration, BS to the blend shape with a certain number of basis functions and Regre. to regression.

not only necessary to reconstruct details from registration, but also from blend shapes, and after regression. We replaced the input images in the training set by normal maps constructed from the blend shapes with 500, 200 and 100 basis functions and one set from the regression reconstruction. The goal is to determine the amount of detail that is necessary in the input to obtain realistic detailed wrinkles. Table 1 shows the error rates in each experiment. 500 basis functions seem sufficient for a reasonable amount of detail in the result. Probably due to the fact that the reconstruction from regression is more noisy and bumpy, the neural network is not capable of reconstructing long defined folds and instead produces a lot of higher frequency wrinkles (see Fig. 11). This is an indicator that the structures of the inputs are only redefined by the net and important folds have to be visible in the input.

5.3 Retargeting

The final goal is to be able to scan one piece of clothing in one or several sequences and then transferring it on new persons with new movements on the go.



a) No Normal Map b) Ground-Truth c) Reg. with temporal d) 200 with temporal e) 500 with temporal f) Reg. no temporal

Fig. 11. Examples of different training results for high resolution normal maps. Left to right: Global shape, target normal map, learned from registration normal map with temporal consistency, learned from blend shape with 200 basis functions and temporal consistency, as previous but with 500 basis functions, learned from registration normal map without temporal consistency. The example pose is not seen in the training set.

Poses. We experimented with various combinations of control vectors Θ , including pose, shape, joint root velocity and acceleration history. It turns out most formulations in the literature are difficult to train or unstable [2, 39, 31]. We restrict the joint parameters to those directly related to each piece of clothing to reduce the dimensionality. In the case of shirts, this leaves the parameters related to the upper body. In general, linear regression generalized best but smoothed out a lot of overall geometric details, even in the training set. We evaluated on 9213 frames for 500 and 1000 blend shapes: $MSE_{500} = 2.902$ and $MSE_{1000} = 3.114$.

On the other hand, we trained an encoder-decoder with LSTM units (4 layers with dimension 256), using inputs and outputs equally of length 3 (see Sect. 3.3). We obtained promising results: $MSE_{rnn} = 1.892$. Supplemental materials show visually good reconstructed sequences.

Shapes. In 3.2 we represented clothing with folds as offsets of a mean shape. The same can be done with a human template for persons with different body shapes. Each person $\bar{\mathcal{P}}_i$ in normalized pose can be represented as an average template plus a vertex-wise offset $\bar{\mathcal{P}}_i = \bar{\mathcal{T}}' + o'_i$. Given the fact that the clothing mean shape $\bar{\mathcal{M}} = \bar{\mathcal{T}}'|_{\mathcal{M}} + o'_{\mathcal{M}}$ contains a subset of vertices of the human template, it can be adjusted to any deformation of the template by taking $\bar{\mathcal{M}}_{o'} = \bar{\mathcal{M}} + o'_{i|\mathcal{M}}$. $|_{\mathcal{M}}$ restricts vertices of the human template to those used for clothing. Then the mean in the blend shape can simply be replaced by $\bar{\mathcal{M}}_{o'}$. Equation 3 becomes:

$$B(\{\lambda_1^i, \dots, \lambda_k^i\}, \theta_i) = S\left(\bar{\mathcal{M}}_{o'} + \sum_{l=0}^k \lambda_l^i \cdot V_l, \theta_i\right) \approx \mathcal{P}_{i|\mathcal{M}}, \quad (5)$$

Replacing the mean shape affects surface normals. Hence, it is necessary to use normal maps in tangent space at rendering time. This makes them applicable to any body shape (see Fig. 12).

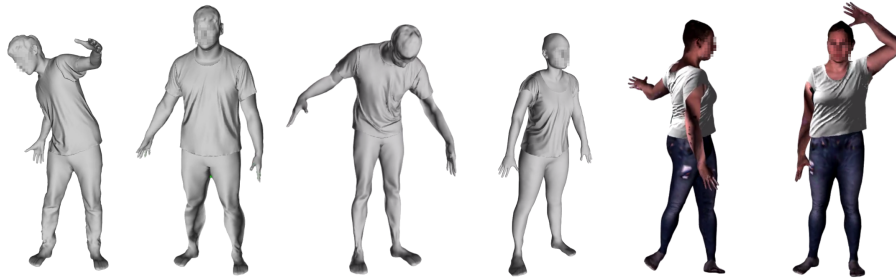


Fig. 12. Body shape retargeting. The first and fourth entry are shirts on the original models, the following two are retargeted to new body shapes.

6 Conclusion

We present DeepWrinkles, an entirely data-driven framework to capture and reconstruct clothing in motion out from 4D scan sequences. Our evaluations show that high frequency details can be added to low resolution normal maps using a conditional adversarial neural network. We introduce an additional temporal loss to the GAN objective that preserves geometric consistency across time, and show qualitative and quantitative evaluations on different datasets. We also give details on how to create low resolution normal maps from registered data, as it turns out registration fidelity is crucial for the cGAN training. The two presented modules are complementary to achieve accurate and realistic rendering of global shape and details of clothing. To the best of our knowledge, our methods exceeds the level of detail of the current state of the art in both physics-based simulation and data-driven approaches by far. Additionally, the space requirement of a normal map is negligible in comparison to increasing the resolution of clothing mesh, which makes our pipeline suitable to standard 3D engines.

Limitations. High resolution normal maps can have missing information in areas not seen by cameras, such as armpit areas. Hence, visually disruptive artifacts can occur although the clothing template can fix most of the issues (e.g., by doing a pass of smoothing). At the moment pose retargeting works best when new poses are similar to ones included in the training dataset. Although the neural network is able to generalize to some unseen poses, reconstructing the global shape from a new joint parameter sequence can be challenging. This should be fixed by scaling the dataset.

Future Work. Scanning setup can be extended to reconstruct all body parts with sufficient details without occlusions, and apply our method to more diverse types of clothing and accessories like coats, scarfs. Normal maps could also be used to add fine details like buttons which are hard to capture in 3D.

Acknowledgements. We would like to thank the FRL teams for their support, and Vignesh Ganapathi-Subramanian for preliminary work on the subspace model.

References

1. www.3dmd.com, temporal-3dmd systems (4d) (2018)
2. de Aguiar, E., Sigal, L., Treuille, A., Hodgins, J.: Stable spaces for real-time clothing. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. Association for Computing Machinery (ACM) (2010)
3. de Aguiar, E., Stoll, C., Theobalt, C., Ahmed, N., Seidel, H., Thrun, S.: Performance capture from sparse multi-view video. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. vol. 27, pp. 98:1–98:10. Association for Computing Machinery (ACM) (2008)
4. Baraff, D., Witkin, A., Kass, M.: Untangling cloth. In: Hart, J.C. (ed.) *ACM Transactions on Graphics*. vol. 22, pp. 862–870. Association for Computing Machinery (ACM), New York, NY, USA (Jul 2003)
5. Beeler, T., Hahn, F., Bradley, D., Bickel, B., Beardsley, P., Gotsman, C., Sumner, R.W., Gross, M.: High-quality passive facial performance capture using anchor frames. In: Alexa, M. (ed.) *ACM Transactions for Graphics*. vol. 30, pp. 75:1–75:10. Association for Computing Machinery (ACM), New York, NY, USA (August 2011)
6. Bogo, F., Romero, J., Pons-Moll, G., Black, M.J.: Dynamic FAUST: Registering human bodies in motion. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE (Jul 2017)
7. Bradley, D., Popa, T., Sheffer, A., Heidrich, W., Boubekeur, T.: Markerless garment capture. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. vol. 27, pp. 99:1–99:9. Association for Computing Machinery (ACM), New York, NY, USA (Aug 2008)
8. Bridson, R., Marino, S., Fedkiw, R.: Simulation of clothing with folds and wrinkles. In: *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. pp. 28–36. Eurographics Association, San Diego, CA, USA (Jul 2003)
9. Choi, K.J., Ko, H.S.: Stable but responsive cloth. In: *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 604–611. SIGGRAPH, ACM, New York, NY, USA (2002)
10. Chu, M., Thuerey, N.: Data-driven synthesis of smoke flows with cnn-based feature descriptors. In: Alexa, M. (ed.) *ACM Transactions for Graphics*. vol. 36, pp. 69:1–69:14. Association for Computing Machinery (ACM), New York, NY, USA (Jul 2017)
11. Cignoni, P., Montani, C., Scopigno, R., Rocchini, C.: A general method for preserving attribute values on simplified meshes. In: *Proceedings of IEEE Conference on Visualization*. pp. 59–66. IEEE (1998)
12. Cohen, J.D., Olano, M., Manocha, D.: Appearance-perserving simplification. In: *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH*. pp. 115–122. ACM, Orlando, FL, USA (Jul 1998)
13. Danerek, R., Dibra, E., Öztireli, C., Ziegler, R., Gross, M.: Deepgarment : 3d garment shape estimation from a single image. In: Chen, M., Zhang, R. (eds.) *Computer Graphics Forum*. vol. 36, pp. 269–280. Eurographics Association (2017)
14. Dou, M., Khamis, S., Degtyarev, Y., Davidson, P., Fanello, S.R., Kowdle, A., Escolano, S.O., Rhemann, C., Kim, D., Taylor, J., Kohli, P., Tankovich, V., Izadi, S.: Fusion4d: Real-time performance capture of challenging scenes. In: Alexa, M. (ed.) *ACM Transactions for Graphics*. vol. 35, pp. 114:1–114:13. Association for Computing Machinery (ACM), New York, NY, USA (Jul 2016)
15. Fyffe, G., Nagano, K., Huynh, L., Saito, S., Busch, J., Jones, A., Li, H., Debevec, P.: Multi-view stereo on consistent face topology. *Computer Graphics Forum* **36**(2), 295–309 (May 2017)

16. Goldenthal, R., Harmon, D., Fattal, R., Bercovier, M., Grinspun, E.: Efficient simulation of inextensible cloth. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. vol. 26, p. 49. Association for Computing Machinery (ACM) (2007)
17. Goldlücke, B., Cremers, D.: Superresolution texture maps for multiview reconstruction. In: *IEEE 12th International Conference on Computer Vision, ICCV*. pp. 1677–1684. IEEE Computer Society (Sep 2009)
18. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K. (eds.) *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems (NIPS)*. pp. 2672–2680. Curran Associates, Inc., Montreal, Quebec, Canada (Dec 2014)
19. Guan, P., Reiss, L., Hirshberg, D.A., Weiss, A., Black, M.J.: DRAPE: dressing any person. *ACM Transactions for Graphics* **31**(4), 35:1–35:10 (2012)
20. Guo, K., Xu, F., Wang, Y., Liu, Y., Dai, Q.: Robust non-rigid motion tracking and surface reconstruction using l0 regularization. pp. 3083–3091. IEEE Computer Society (2015)
21. Hahn, F., Thomaszewski, B., Coros, S., Sumner, R.W., Cole, F., Meyer, M., DeRose, T., Gross, M.H.: Subspace clothing simulation using adaptive bases. In: Alexa, M. (ed.) *ACM Transactions for Graphics*. vol. 33, pp. 105:1–105:9. Association for Computing Machinery (ACM) (2014)
22. Hernández, C., Vogiatzis, G., Brostow, G.J., Stenger, B., Cipolla, R.: Non-rigid photometric stereo with colored lights. In: *IEEE 11th International Conference on Computer Vision, ICCV*. pp. 1–8. IEEE Computer Society, Rio de Janeiro, Brazil (Oct 2007)
23. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Proceedings of the 32nd International Conference on Machine Learning - Volume 37*. pp. 448–456. ICML, ACM (2015)
24. Isola, P., Zhu, J., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 5967–5976. IEEE Computer Society, Honolulu, HI, USA (Jul 2017)
25. Kavan, L., Gerszewski, D., Bargteil, A.W., Sloan, P.P.: Physics-inspired upsampling for cloth simulation in games. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. vol. 30, pp. 93:1–93:10. Association for Computing Machinery (ACM), New York, NY, USA (Jul 2011)
26. Krishnamurthy, V., Levoy, M.: Fitting smooth surfaces to dense polygon meshes. In: *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH*. New Orleans, LA, USA (Aug.)
27. Lassner, C., Pons-Moll, G., Gehler, P.V.: A generative model of people in clothing. In: *IEEE International Conference on Computer Vision, ICCV*. pp. 853–862. IEEE Computer Society, Venice, Italy (Oct 2017)
28. Ledig, C., Theis, L., Huszar, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A.P., Tejani, A., Totz, J., Wang, Z., Shi, W.: Photo-realistic single image super-resolution using a generative adversarial network. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 105–114. IEEE Computer Society, Honolulu, HI, USA (Jul 2017)
29. Leroy, V., Franco, J., Boyer, E.: Multi-view dynamic shape refinement using local temporal integration. In: *IEEE International Conference on Computer Vision, ICCV*. pp. 3113–3122. Venice, Italy (2017)

30. Li, H., Adams, B., Guibas, L.J., Pauly, M.: Robust single-view geometry and motion reconstruction. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. vol. 28, pp. 175:1–175:10. Association for Computing Machinery (ACM), New York, NY, USA (Dec 2009)
31. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. In: Alexa, M. (ed.) *ACM Transactions for Graphics*. vol. 34, pp. 248:1–248:16. Association for Computing Machinery (ACM) (2015)
32. MarvelousDesigner: www.marvelousdesigner.com (2018)
33. Matsuyama, T., Nobuhara, S., Takai, T., Tung, T.: *3D Video and Its Applications*. Springer (2012)
34. Miguel, E., Bradley, D., Thomaszewski, B., Bickel, B., Matusik, W., Otaduy, M.A., Marschner, S.: Data-driven estimation of cloth simulation models. *Computer Graphics Forum* **31**(2), 519–528 (2012)
35. Müller, M., Chentanez, N.: Wrinkle meshes. In: *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. pp. 85–92. SCA '10, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2010)
36. Newcombe, R.A., Fox, D., Seitz, S.M.: Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 343–352. IEEE Computer Society, Boston, MA, USA (Jun 2015)
37. Olszewski, K., Li, Z., Yang, C., Zhou, Y., Yu, R., Huang, Z., Xiang, S., Saito, S., Kohli, P., Li, H.: Realistic dynamic facial textures from a single image using gans. In: *IEEE International Conference on Computer Vision, ICCV*. pp. 5439–5448. IEEE Computer Society, Venice, Italy (Oct 2017)
38. Pons-Moll, G., Pujades, S., Hu, S., Black, M.: Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics (Proc. SIGGRAPH)* **36**(4) (2017)
39. Pons-Moll, G., Romero, J., Mahmood, N., Black, M.J.: Dyna: A model of dynamic human shape in motion. In: Alexa, M. (ed.) *ACM Transactions for Graphics*. vol. 34, pp. 120:1–120:14. Association for Computing Machinery (ACM) (Aug 2015)
40. Saito, S., Wei, L., Hu, L., Nagano, K., Li, H.: Photorealistic facial texture inference using deep neural networks. *IEEE Conference on Computer Vision and Pattern Recognition, CVPR* pp. 2326–2335 (Jul 2017)
41. Starck, J., Hilton, A.: Surface capture for performance-based animation. *IEEE Computer Graphics and Applications* **27**(3), 21–31 (2007)
42. Sumner, R.W., Popović, J.: Deformation transfer for triangle meshes. In: Hart, J.C. (ed.) *ACM Transactions for Graphics*. vol. 23, pp. 399–405. Association for Computing Machinery (ACM), New York, NY, USA (Aug 2004)
43. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to sequence learning with neural networks. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K. (eds.) *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems (NIPS)*. pp. 3104–3112. Curran Associates, Inc., Montreal, Quebec, Canada (Dec 2014)
44. Taylor, J., Shotton, J., Sharp, T., Fitzgibbon, A.: The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 103–110. IEEE Computer Society (Jul 2012)
45. Tsiminaki, V., Franco, J., Boyer, E.: High resolution 3d shape texture from multiple videos. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. pp. 1502–1509. IEEE Computer Society, Columbus, OH, USA (Jun 2014)

46. Tung, T., Matsuyama, T.: Intrinsic characterization of dynamic surfaces. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 233–240. IEEE Computer Society, Portland, OR, USA (Jun 2013)
47. Tung, T., Nobuhara, S., Matsuyama, T.: Simultaneous super-resolution and 3d video using graph-cuts. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society, Anchorage, Alaska, USA (Jun 2008)
48. Tung, T., Nobuhara, S., Matsuyama, T.: Complete multi-view reconstruction of dynamic scenes from probabilistic fusion of narrow and wide baseline stereo. In: IEEE 12th International Conference on Computer Vision, ICCV. pp. 1709–1716. IEEE Computer Society, Kyoto, Japan (Sep 2009)
49. Vlastic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. In: Hart, J.C. (ed.) ACM Transactions for Graphics. pp. 97:1–97:9. Association for Computing Machinery (ACM), New York, NY, USA (2008)
50. Vlastic, D., Peers, P., Baran, I., Debevec, P., Popović, J., Rusinkiewicz, S., Matusik, W.: Dynamic shape capture using multi-view photometric stereo. In: Proceedings of ACM SIGGRAPH Asia. pp. 174:1–174:11. ACM, New York, NY, USA (2009)
51. Volino, P., Magnenat-Thalmann, N.: Virtual clothing - theory and practice. Springer (2000)
52. Wang, H., Hecht, F., Ramamoorthi, R., O’Brien, J.F.: Example-based wrinkle synthesis for clothing animation. In: Alexa, M. (ed.) ACM Transactions for Graphics. vol. 29, pp. 107:1–8. Association for Computing Machinery (ACM), Los Angeles, CA (Jul 2010)
53. Wang, H., Ramamoorthi, R., O’Brien, J.F.: Data-driven elastic models for cloth: Modeling and measurement. In: Alexa, M. (ed.) ACM Transactions for Graphics. vol. 30, pp. 71:1–11. Association for Computing Machinery (ACM), Vancouver, BC Canada (Jul 2011)
54. White, R., Crane, K., Forsyth, D.A.: Capturing and animating occluded cloth. In: Hart, J.C. (ed.) ACM Transactions for Graphics. vol. 26, p. 34. Association for Computing Machinery (ACM) (2007)
55. Zhang, C., Pujades, S., Black, M.J., Pons-Moll, G.: Detailed, accurate, human shape estimation from clothed 3d scan sequences. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5484–5493. IEEE Computer Society, Honolulu, HI, USA (Jul 2017)