Tex2Shape: Detailed Full Human Body Geometry From a Single Image

Thiemo Alldieck[1,2]   Gerard Pons-Moll[2]   Christian Theobalt[2]   Marcus Magnor[1]

[1]Computer Graphics Lab, TU Braunschweig, Germany
[2]Max Planck Institute for Informatics, Saarland Informatics Campus, Germany
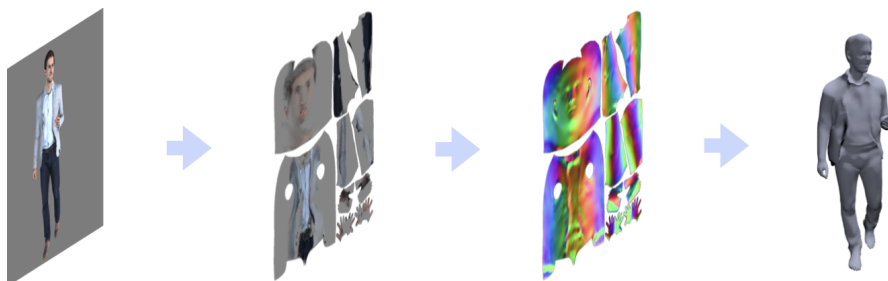{alldieck,magnor}@cg.cs.tu-bs.de {gpons,theobalt}@mpi-inf.mpg.de

Figure 1: We present an image-to-image translation model for detailed full human body geometry reconstruction from a single image. An input image is transformed into an incomplete texture, then our Tex2Shape network translates the texture into normal and displacement maps. The maps augment a smooth body model with details, hair, and clothing. Result visualized in ground truth pose.

Abstract

We present a simple yet effective method to infer detailed full human body shape from only a single photograph. Our model can infer full-body shape including face, hair, and clothing including wrinkles at interactive frame-rates. Results feature details even on parts that are occluded in the input image. Our main idea is to turn shape regression into an aligned image-to-image translation problem. The input to our method is a partial texture map of the visible region obtained from off-the-shelf methods. From a partial texture, we estimate detailed normal and vector displacement maps, which can be applied to a low-resolution smooth body model to add detail and clothing. Despite being trained purely with synthetic data, our model generalizes well to real-world photographs. Numerous results demonstrate the versatility and robustness of our method.

1. Introduction

In this paper, we address the problem of automatic *detailed* full-body human shape reconstruction from a single image. Human shape reconstruction has many applications in virtual and augmented reality, scene analysis, and virtual try-on. For most applications, acquisition should be quick and easy, and visual fidelity is important. Reconstructed geometry is most useful if it shows hair, face, and clothing folds and wrinkles at sufficient detail – what we re-

fer to as detailed shape. Detail adds realism, allows people to feel identified with their self-avatar and their interlocutors, and often carries crucial information.

While a large number of papers focus on recovering pose, and rough body shape from a single image [35, 25, 36, 9], much fewer papers focus on recovering detailed shapes. Some recent methods recover pose and non-rigid deformation from monocular video [57], even in real-time [15]. However, they require a pre-captured static template of each subject. Other recent works [4, 2] recover static body shape, and clothing as displacements on top of the SMPL body model [32] (model-based), or use a voxel representation [50, 33]. Voxel-based methods [50, 33] often produce errors at the limbs of the body and require fitting a model post-hoc [50]. Model-based methods are more robust, but results tend to lack fine detail. We hypothesize there are three reasons for this. Firstly, they rely mostly on silhouettes for either fitting [4], or CNN-based regression plus fitting [2], ignoring the rich illumination and shading information contained in RGB values. Secondly, the regression from image pixels directly to 3D mesh displacements is hard because inputs and outputs are *not aligned*. Furthermore, prediction of high-resolution meshes requires mesh-based neural networks, which are very promising but are harder to train than standard 2D CNNs. Finally, they rely on 3D pose estimation, which is hard to obtain accurately.

Based on these observations, our idea is to turn the shape regression into an *aligned* image-to-image translation prob-

lem (see Fig. 1). To that end, we map input and output pairs to the pose-independent UV-mapping of the SMPL model. The UV-mapping unfolds the body surface onto a 2D image such that every pixel corresponds to a 3D point on the body surface. Similar to [34], we map the visible image pixels to the UV space using DensePose [5] obtaining a partial texture map image, which we use as input. Instead of regressing details directly on the mesh, we propose to regress shape as UV-space displacement and normal maps. Every pixel stores a normal and a displacement vector from a smooth shape (in the space of SMPL) to the detailed shape. We call our model to Tex2Shape.

We train Tex2Shape with a dataset of 2043 3D scans of people in varying clothing, poses, and shapes. To map all scan shapes to the UV-space, we non-rigidly register SMPL to each scan, optimizing for model shape parameters and free-form displacements, and store the latter in a displacement map. Registration is also useful for augmentation; using SMPL, we render multiple images of varying pose and camera view. We further augment the renderings with realistic illumination, which is a strong cue in this problem. Assuming a Lambertian reflectance model, we know that color forms from the dot product of light direction and the surface normal times albedo. Shape-from-shading [60] allows to invert the process and estimate the surface from shading, which was used before to refine geometry of stereo-based [55] or multi-view-based human performance capture results [56, 29]. After synthesizing image pairs, we train a Pix2Pix network [19] to map from partial texture maps to complete normal and displacement maps and a second small network for estimating SMPL body shape parameters.

Several experiments demonstrate that our proposed data pre-processing undoubtedly pays-off. Trained only from synthetic images, our model can robustly produce, in one shot, *full 3D shapes* of people with varied clothing, shape, and hair. In contrast to models that produce normals or shading only for the visible image part, Tex2Shape hallucinates the shape also for the *occluded* part – effectively performing translation and completion together. In summary, our contributions are:

- We turn a hard full-body shape reconstruction problem into an easier 3D pose-independent image-to-image translation one. To the best of our knowledge, this is the first method to infer detailed body shape as image-to-image translation.

- From a single image, our model can regress full 3D clothing, hair and facial details in 50 milliseconds.

- Experiments demonstrate that, while very simple, Tex2Shape is very effective and is capable of regressing full 3D clothing, hair and facial details in a static reference pose in one shot.

- Tex2Shape is available for research purposes [1].

2. Related Work

Human shape reconstruction is a wide field of research, often jointly approached with pose reconstruction. In the following, we review methods for human pose and shape reconstruction from monocular image and video. Full body methods are often inspired by methods for face geometry estimation. Hence, we include face reconstruction in our review. When it comes to detailed reconstruction, clothing plays an important role. Therefore, we conclude with a brief overview of garment reconstruction and modeling.

Pose and shape reconstruction. Methods for monocular pose and shape reconstruction often utilize parametric body models to limit the search space [6, 16, 32, 39, 23], or use a pre-scanned static template to capture pose and non-rigid surface deformation [57, 15]. To recover pose and shape, the 3D body model is fitted against 2D poses. In early works 2D poses have been entirely or partially manually clicked [14, 62, 21, 42], later the process was automated [9, 28] with 2D landmark detections from deep neural networks [37, 18, 11]. In recent work, the SMPL [32] model has been integrated into network architectures [25, 36, 35, 48]. This further automates and robustifies the process. All these works focus mostly on robust pose detection. Shape estimation is often limited to surface correlations with bone lengths. Most importantly, the shape is limited to the model space. In contrast, we focus only on shape and estimate geometry details beyond the model space.

Clothing and hair can be obtained by optimization-based methods [4, 3]. From a video of a subject turning around in A-pose, silhouettes are fused in canonical pose. In the same setting, the authors in [2] present a hybrid learning and optimization-based method, that makes the process completely automatic, fast, and dependent only on a handful of images. However, all these methods can only process A-poses and depend on robust pose detection. The method in [54] loosens this restriction and creates humanoid shapes from a single image via 2D warping of SMPL parameters, but only partially handles self-occlusion. Another recent line of research estimates pose and shape in form of a voxel representation [49, 20, 33], which allows for more complex clothing but limits the level of detail. In [61] the authors alleviate this limitation by augmenting the visible parts with a predicted normal map. In contrast, we present 3D pose-independent shape estimation in a reference pose with high-resolution details also on non-visible parts.

Several previous methods exploited shading cues in high-frequency texture to estimate high-frequency detail. For instance, they estimated lighting and reflectance to compute shape-from-shading-refined geometry of a human template from stereo [55] or multi-view imagery [56, 29].

Face reconstruction. Several recent monocular face reconstruction and performance capture methods use shading-

based refinement for geometry improvement, e.g., in analysis-by-synthesis fitting [43] or refinement, or in a trained neural network [44, 17]. Also related to our approach are recent works integrating a differentiable face renderer in a neural network to estimate instance correctives of geometry and albedo relative to a base model [47], or learn an identity geometry and albedo basis from scratch from video [46].

Garment reconstruction and modeling. Body shape under clothing has been estimated without [59] and jointly with a separate clothing layer [38] from 3D scans and from RGB-D [45]. [58] introduces a technique, which allows complex clothing to be modeled as offsets from the naked body. The work in [52] describes a model that encodes shape, garment sketch, and garment model, in a single shared latent code, which enables interactive garment design. High frequency wrinkles are predicted as a function of pose either in UV space using a CNN [27, 22] or directly in 3D using a data-driven optimization method [40]. All these methods [27, 58, 22] target realistic *animation* of clothing and can only predict garments in isolation [27, 22]. Learning based normals and depth recovery [7] or meshes [12] has been demonstrated but again only for single garments. In contrast, our approach is the first to reconstruct the detailed shape of a *full-body* from a *single image* by learning an image-to-image mapping.

3. Method

The goal of this work is to create an animatable 3D model of a subject from a single photograph. The model should reflect the subject's body shape and contain details such as hair and clothing with garment wrinkles. Details should be present also on body parts that have not been visible in the input image, e.g. on the back of the person. In contrast to previous work [33, 54, 2] we aim for fully automatic reconstruction which does not require accurate 3D pose. To this end, we train a Pix2Pix-style [19] convolutional neural network to infer normals and vector displacement (*UV shape-images*) on top of the SMPL body model [32]. To align the input image with the output UV-shape images, we extract a partial UV texture map of the visible area using off-the-shelf methods [5, 25]. An overview is given in Fig. 2. A second small CNN infers SMPL shape parameters from the image (see Sec. 5.1). In Sec. 3.1 we describe the parametric body model used in this work, and in Sec. 3.2 we explain our parameterization of appearance, normals, and displacements.

3.1. Parametric body model

SMPL is a parameterized body model learned from scans of subjects in minimal clothing. It is defined as a function of pose θ and shape β returning a mesh of $N = 6890$ vertices and $F = 13776$ faces. Shape β corresponds to the first 10 principal components of the training data subjects. Since scale is an inherent ambiguity in monocular images, we made β independent of body height in this work. Our method estimates β with a standardized height and is independent of pose θ. Details that go beyond the SMPL shape space are added via UV displacement and normal maps (UV shape-images), as described in Sec. 3.2. During the dataset generation (see Sec. 4), we use SMPL to synthesize images of humans posing in front of the camera.

3.2. UV parameterization

The SMPL model describes body shapes with a mesh containing 6890 vertices. Unfortunately, this resolution is not high enough to explain fine details, such as garment wrinkles. Another problem is that meshes do not live on a regular 2D grid like images, and consequently require taylored solutions [10] that are not yet as effective as standard CNNs on the image domain. To leverage the power of standard CNNs, we propose to use a well-established parameterization of mesh surfaces: UV mapping [8]. A UV map unwraps the surface onto an image, allowing to represent functions defined on the surface as images. Hereby, U and V denote the 2 axes of the image. The mapping is defined once per mesh topology and assigns every pixel in the map to a point on the surface via barycentric interpolation of neighboring vertices. By using a UV map, a mesh can be augmented with geometric details of a resolution proportional to the UV map resolution.

We augment SMPL using two UV maps, namely normal map and vector displacement map. A normal map contains new surface normals, that can add or enhance visual details through shading. A vector displacement map contains 3D vectors that displace the underlying surface. Displacements and normals are defined on the canonical T-pose of SMPL. The input to our neural network is a partial texture map of the visible pixels on the input photograph (see Sec. 5.3).

4. Dataset Generation

To learn our model we synthesize a varied dataset from real 3D scans of people. Specifically, we synthesize images of humans in various poses under realistic illumination paired with normal maps, displacement maps, and SMPL shape parameters β. The large majority of scans (1826) was kindly provided from Twindom (https://web.twindom.com/). We additionally purchased 163 scans from renderpeople.com and 54 from axyz-design.com. These scans do not share the same mesh layout, and therefore we can not directly compute coherent normal and displacement maps. To this end, we non-rigidly register the SMPL model against each of the scans. This ensures that all vertices share the same contextual information across the dataset. Furthermore, we can change the pose of the scans using SMPL. Unfortunately, non-rigid registration
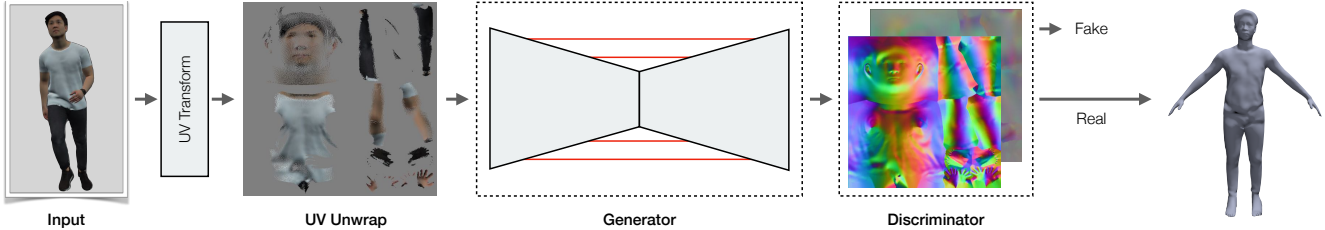
Figure 2. Overview of the key component of our method: A single photograph of a subject is transformed into a partial UV texture map. This map is then processed with a U-Net with skip connections that preserve high-frequent details. A PatchGAN discriminator enforces realism. The generated normals and displacements can be applied to the SMPL model using standard rendering pipelines.

of clothed people is a very challenging problem itself (see Sec. 4.1), and often results in unnatural shapes. Hence, we manually selected 2043 high quality registrations. Unfortunately, our current dataset is slightly biased towards men because registration currently fails more often for women, due to long hair, skirts and dresses. Of the 2043 scans, we reserve 20 scans for validation and 55 scans for testing.

In the following, we explain our non-rigid registration procedure in more detail and describe the synthetization of the paired dataset for training of the models.

4.1. Scan registration

As discussed in Sec. 3.1, $N = 6890$ vertices are not enough to explain fine details. To this end, we sub-divide each face in SMPL into four, resulting in a new mesh consisting of $N = 27554$ vertices and $F = 55104$ faces. This high-resolution mesh can better explain fine geometric details in the scans. While joint optimization is generally desirable, registration is much more robust when done in stages: we first compute 3D pose, then body shape and finally non-rigid details. We start the registration by reconstructing the pose of the scan subject. Therefore, we find 3D landmarks by rendering the scan from multiple cameras and minimizing the 2D re-projection error to 2D joint OpenPose detections [11]. Then we optimize the SMPL pose parameters θ to explain the estimated 3D joint locations. Next, we optimize for shape parameters β to minimize scan to SMPL surface distance. Here, we make sure SMPL vertices stay inside the scan by paying a higher cost for vertices outside the scan since SMPL can only reliable explain the naked body shape. Finally, we recover fine-grained details by optimizing the location of SMPL vertices. The resulting registrations explain high-frequency details of the scans with the subdivided SMPL mesh layout and can be re-posed.

4.2. Spherical harmonic lighting

For a paired dataset, we first need to synthesize images of humans. For realistic illumination, we use spherical harmonic lighting. Spherical harmonics (SH) are orthogonal basis functions defined over the surface of the sphere. For rendering SH are used to describe the directions from where light is shining into the scene [41]. We follow the standard

procedure and describe the illumination with the first 9 SH components per color. To produce a large variety of realistic illumination conditions, we convert images of the *Laval Indoor HDR dataset* [13] into diffuse SH coefficients, similar to [24]. For further augmentation, we rotate the coefficients randomly around the Y-axis.

4.3. UV map synthetization

To complete our dataset, we calculate UV maps that explain details of the 3D registrations. In UV mapping every face of the mesh has a 2D counterpart in the UV image. Hence, UV mapping is essentially defined through a 2D mesh. Given a 3D mesh and a set of per-vertex information, a UV map can be synthesized through standard rendering. Information between vertices is filled through barycentric interpolation. This means, given the high-resolution registrations, we can simply render detailed UV displacement and normal maps. The displacement maps encode the free-form offsets, that are not part of SMPL. The normal maps contain surface normals in canonical T-pose. These maps are used to augment the standard-resolution naked SMPL, which eliminates the need for higher mesh-resolution or per-vertex offsets. We use the standard-resolution SMPL augmented with the UV maps in all our experiments.

5. Model and Training

In the following, we explain the used network architectures, losses, and training schemes in more detail. Further, we explain how a partial texture can be obtained from DensePose [5] results.

5.1. Network architectures

Our method consists of two CNNs – one for normal and displacement maps and one for SMPL shape parameters β. The main component of our method is the Tex2Shape-network as depicted in Fig. 2. The network is a conditional Generative Adversarial Network (Pix2Pix) [19] consisting of a U-Net generator and a PatchGAN discriminator. The U-Net features each seven convolution-ReLU-batchnorm down- and up-sampling layers with skip connections. The discriminator consists of four of such down-sampling layers. We condition on 512×512 partial textures, based on

two observations: First, when mapping pixels from an HD 1024×1024 image to UV, the resolution is high enough to contain most pixels from the foreground, and not too high to prevent large unoccupied regions. Second, using the mesh resolution of the training set, larger UV maps would only contain more interpolated data. See supplemental material for an ablation experiment using smaller UV maps.

The β-network takes 1024×1024 DensePose detections as input. These are then again down-sampled with seven convolution-ReLU-batchnorm layers and finally mapped to 10 β-parameters by a fully-connected layer.

5.2. Losses and training scheme

The goal of our method is to create results with high *perceived quality*. We believe structure is more important than accuracy and therefore experiment with the following loss: The structural similarity index (SSIM) was introduced to predict the perceived quality of images. The multi-scale SSIM (MS-SSIM) [53] evaluates the image on different image scales. We maximize the structural similarity of ground truth and predicted normal and displacement maps by minimizing the dissimilarity (MS-DSSIM): $(1 - \mathrm{MS\text{-}SSIM})/2$. We further train with the well-established L1-loss and the GAN-loss coming from the discriminator. Finally, the β-network is trained with an L2 parameter loss. We train both CNNs with the Adam optimizer [26] and decay the learning-rate once the losses plateau.

5.3. Input partial texture map

The partial texture forming the input to our method is created by transforming pixels from the input image to UV space based on DensePose detections, see Fig. 3. Dense-Pose predicts UV coordinates of 24 body parts of the SMPL body model (Fig. 3 middle). For easier mapping, we pre-compute a look-up table to convert from 24 DensePose UV maps to the single joint SMPL UV parameterization. Each pixel in the DensePose detection now maps to a coordinate in the SMPL UV map. Using this mapping, we compute a partial texture from the input image (Fig. 3 right).

6. Experiments

In the following, we qualitatively and quantitatively evaluate our proposed method. Results on four different datasets and comparisons to state-of-the-art demonstrate the versatility and robustness of our method as well as the quality of results (Sec 6.1). Further, we study the effect of different supervision losses (Sec. 6.2), evaluate different methods for UV mapping (Sec. 6.3), and measure the robustness for different visibility levels (Sec. 6.4). Finally, in Sec. 6.5 we demonstrate a potential application of our proposed method, namely garment transfer between subjects. More experiments and ablation studies can be found in the supplemental material. Due to scale ambiguity in monocular images,
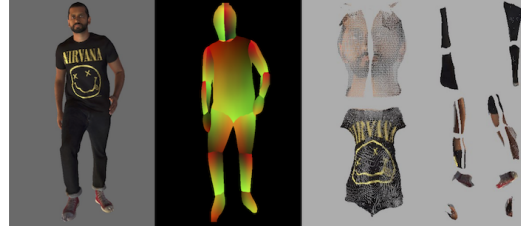

Figure 3. To create the input to our method, we first process the input image (left) with DensePose. The DensePose result (middle) contains UV coordinates, that can be used to map the input image into a partial texture (right).

all results are up to scale. Also, our method does not compute pose. For better inspection, we depict results in ground truth or A-pose. Further, we color-code the results by the used method for UV-mapping (see Sec. 6.3). Results using DensePose mapping are *green*, *blue* marks ground truth mapping, *red* indicates HMR-based [25] texture reprojection, and ground truth shapes are *grey*. All results have been calculated at interactive frame-rates. Precisely, our method takes on average 50 ms for displacement map, normal map, and β-estimation on an NVIDIA Tesla V100. UV mapping using DensePose can be performed in real-time.

6.1. Qualitative results and comparisons

We qualitatively compare our work against four relevant methods for monocular human shape reconstruction on the *PeopleSnapshot* dataset [4]. BodyNet [49] is a voxel-based method to estimate human pose and shape from only one image. SiCloPe [33] is voxel-based, too, but recovers certain details by relying on synthesized silhouettes of the subject. HMR [25] is a method to estimate pose and shape from single image using the SMPL body model. In [4] the authors present the first video-based monocular shape reconstruction method, that goes beyond the parameters of SMPL. They use 120 images of the same subject roughly posed in A-poses and fuse the silhouettes into a canonical representation. However, the method is optimization-based and requires to fit the pose in each frame first, which makes the process very slow. In Fig. 6, we show a side-by-side comparison with our results. Our method clearly features the highest level of detail, even compared to [4] using 120 frames, while our method only takes a single image as input and runs at interactive frame-rates.

In Fig. 4 we show more results of our method. We compare against ground truth on our own dataset and show qualitative results on *3DPW* [51], *DeepFashion* [30, 31], and *PeopleSnapshot* [4] datasets. Our method successfully generalizes to various real-world conditions. Please note how realistic garment wrinkles are hallucinated on the unseen back of the models. In general, we can see our method is able to infer realistic 3D models featuring hair, facial details, and various clothing including garment wrinkles from single image inputs.
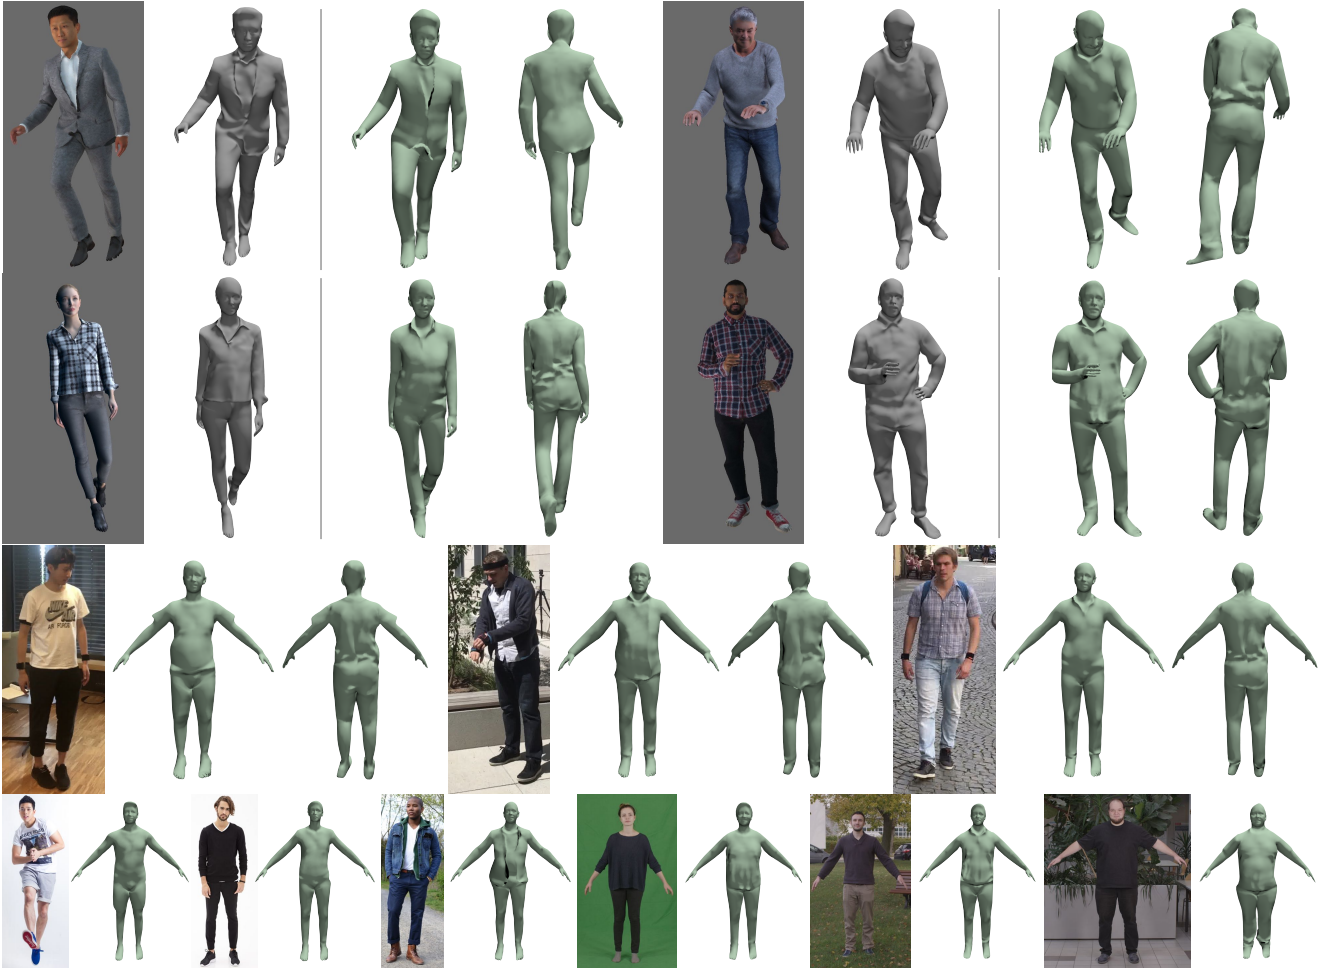
Figure 4. Our 3D reconstruction results (green) on four different datasets. We compare to ground truth (grey) on our synthetic dataset (rows 1 and 2). Qualitative results on 3DPW (3rd row), DeepFashion (4th row left) and PeopleSnapshot (4th row right) demonstrate, that our model generalizes well to real-world footage. Details on the back of the models are hallucinated by our model.
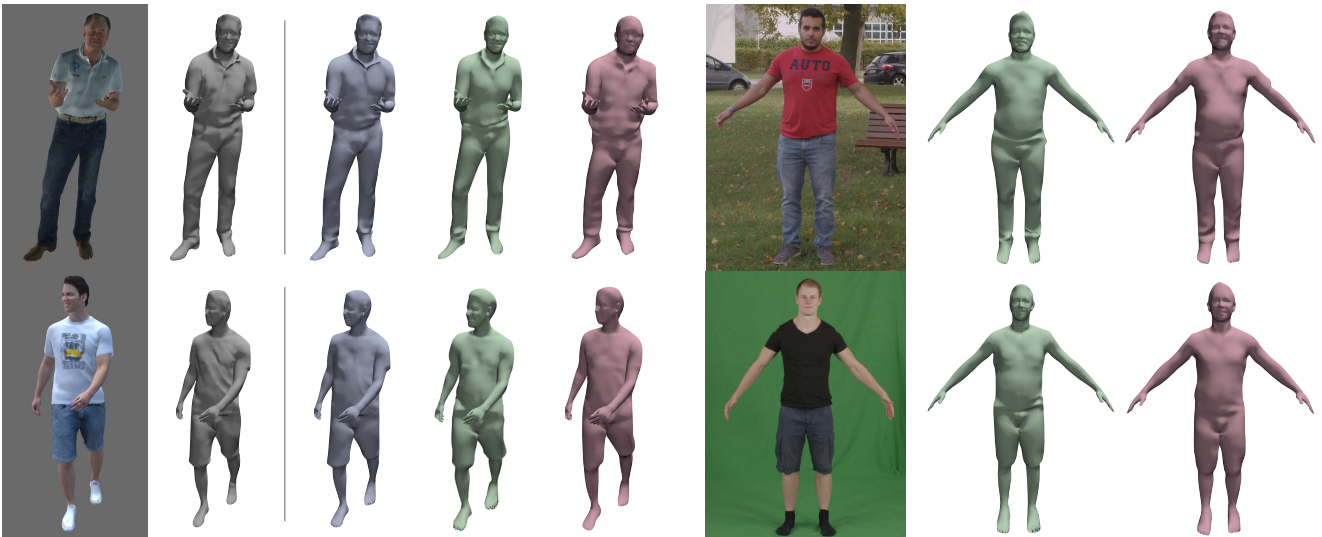

Figure 5. Results using three different methods for partial texture creation compared to input image and ground truth mesh (grey): ground truth UV mapping (blue), DensePose UV mapping (green), HMR-based texture reprojection (red), cf. Fig. 7.
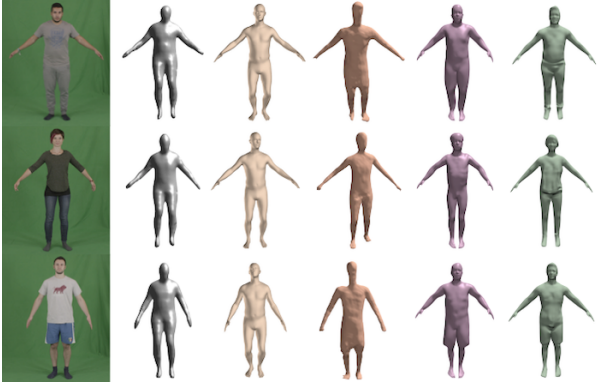
Figure 6. Our method compared to other methods for human shape reconstruction. From left to right: Input image, BodyNet [49], HMR [25], SiCloPe [33], Video Shapes [4], and ours. Our method preserves the highest level of detail.

6.2. Type of supervision

In Sec. 5.2, we have introduced the MS-DSSIM loss. The intuition behind using this loss is that for visual fidelity structure is more important than accuracy. To evaluate this design decision, we train a variant of our Tex2Shape network with L1 and GAN losses only. Since it is not straight forward to quantify better structure, we closely inspect our results on a visual basis. We find, that the variant trained with MS-DSSIM loss is able to reconstruct complex clothing more reliably. Examples are shown in Fig. 8. Note that the results computed with MS-DSSIM loss successfully reconstruct the jackets.

6.3. Impact of UV mapping

Our method requires to first map an input image to a partial UV texture. We propose to use DensePose [5], which makes our method independent of the 3D pose of the subject. In the following, we evaluate the impact of the choice of UV mapping on our method. To this end, we train three variants of our network. Firstly, we train with ground truth UV mappings calculated from the scans. We render the scan's UV coordinates in image space, that are then used for UV mapping, similar to the mapping using DensePose (see Sec. 5.3). We refer to this variant as *GT-UV*. Secondly, we train a variant that can be used with off-the-shelf 3D pose estimators. To this end, we render UV coordinates of the naked SMPL model without free-form offsets. This way only pixels that are covered by the naked SMPL shape are mapped, what simulates texture reprojection from results of 3D pose detectors (*3D pose variant*). Finally, we compare with our standard training procedure using DensePose. A comparison of partial textures created with the three variants is given in Fig. 7. Note how we lose large parts of the texture by using DensePose mapping.

To evaluate the 3D pose variant, we choose HMR [25] as 3D pose detector. Unfortunately, the results of HMR do
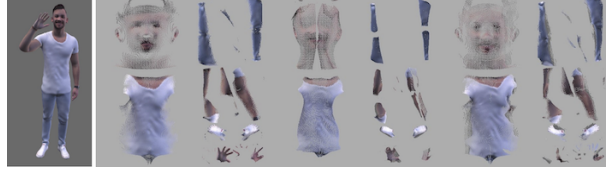

Figure 7. Partial textures computed with different methods. From left to right: Input, ground truth UV mapping, DensePose, HMR.
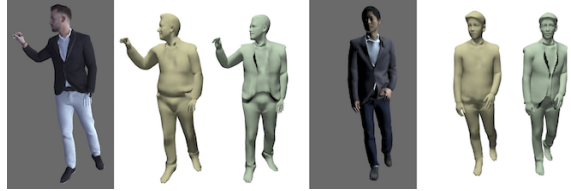

Figure 8. After training with MS-DSSIM loss enabled (green) complex clothing is reconstructed more reliably, than after training with L1 loss only (yellow).

not always align with the input image what produces large errors in the UV space. To this end, we refine the results by minimizing the 2D reprojection error of SMPL joints to OpenPose [11] detections. We choose dogleg optimization and optimize for 20 steps.

In Fig. 5 we show a side-by-side comparison of the three variants. While GT-UV and DensePose variants are almost identical, the 3D pose variant lacks some detail and introduces noise in the facial region. This is caused by the fact, that perfect alignment is still not achieved even after pose-refinement. The GT-UV and DensePose variants differ the most in hairstyle and at the boundary of the shorts, what is not surprising since hair and clothing are only partially mapped by DensePose. However, both variants closely resemble ground truth results. The DensePose and 3D pose mapping variants can directly be used on real-world footage, while only being trained with synthetic data.

6.4. Impact of visibility

In the following, we numerically evaluate the robustness of our method to different visibility settings caused by different poses and distances to the camera. The following results have been computed using GT UV mapping to factor out noise introduced by DensePose. Which pixels can be mapped to the UV partial texture is determined by the subject's pose and distance to the camera. Parts of the body might be not visible (e.g. the subject's back) or occluded by other body parts. If the subject is far away from the camera, it only covers only a small area of the image and thus only a small number of pixels can be mapped.

In Fig. 9 we measure how this influences the accuracy of our results. Over a test-set with 55 subjects, we synthesize images of three different poses with various distances to the camera. The three poses are A-pose, walking towards the camera, and posing sideways with hands touching. We
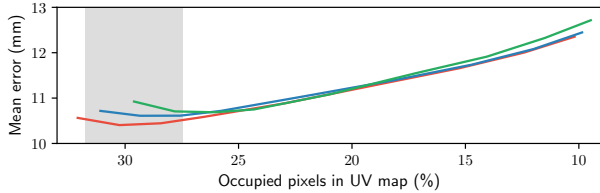
Figure 9. Average displacement error for three different poses (red: A-pose, blue: walking, green: posing sideways with hands touching) and different distances to the camera. The shaded region marks the margin of trained UV map occupancy.
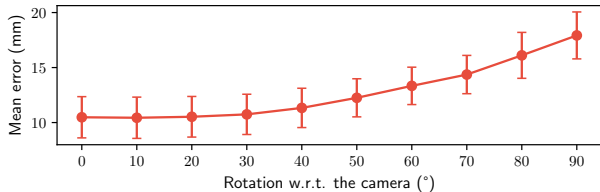

Figure 10. Average displacement error for A-posed subjects and different rotations around Y-axis with respect to the camera. Our model has been trained on rotations $\pm 20°$.

report the mean per-pixel error of 3D displacements maps (including unseen areas) against the percentage of occupied pixels in the partial texture. For all three poses, the error increases linearly, even for untrained texture occupations. Not surprisingly, the minimum of all three poses lies in the margin of trained occupations. Admittedly, for higher occupations, the error slightly goes up what is caused by the fact, that the network was not trained for scenarios where the subject fully covers the input image.

In Fig. 10, we study the robustness of our method against unseen poses. We trained the network with images of humans roughly facing the camera. Therefore, we randomly sampled poses in our dataset and Y-axis rotations between $\pm 20°$. In this experiment, we rotate an A-pose around the Y-axis and report the mean per-pixel 3D displacement error. From $0°$ to $30°$, the error stays almost identical, after $30°$ it increases linearly. Again this behavior can be explained by the network not being trained for such angles.

Both experiments demonstrate the robustness of our method against scenarios not covered by our training set.

6.5. Garment transfer

In our final experiment, we want to demonstrate a potential application of our method, namely garment transfer or virtual try-on. We take several results of our method and use them to synthesize a subject in new clothing. To achieve this, we keep the SMPL shape parameters β. Then we alter normal and displacement maps according to a different result. Hereby, we keep details in the facial region, to preserve the subject's identity and hair-style. Since we edit in UV space, this operation can simply be done using standard image editing techniques. In Fig. 11 we show a subject in three different synthesized clothing styles.
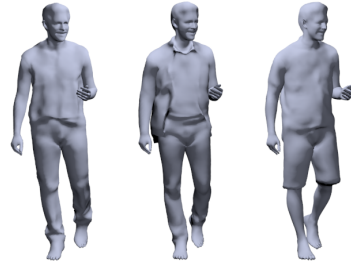

Figure 11. Since all reconstructions share the same mesh layout, we can extract clothing styles and transfer them to other subjects.

Figure 12. Failure cases of our method: The predictor confuses a dress with short pants, a female subject with a male, and hallucinates a hood from a collar.

7. Discussion and Conclusion

We have proposed a simple yet effective method to infer full-body shape of humans from a single input image. For the first time, we present single image shape reconstruction with fine details also on occluded parts. The key idea of this work is to turn a hard full-body shape reconstruction problem into an easier 3D pose-independent image-to-image translation one. Our model Tex2Shape takes partial texture maps created from DensePose as input and estimates details in the UV-space in form of normal and displacement maps. The estimated UV maps allow augmenting the SMPL body model with high-frequent details without the need for high mesh resolution. Our experiments demonstrate that Tex2Shape generalizes robustly to real-world footage, while being trained on synthetic data only.

Our method finds its limitations in hair and clothing that is not covered by the training set. This is especially the case for long hair and dresses since they cannot be modeled as vector displacement fields. Typical failure cases are depicted in Fig. 12. These failures can be explained with garment-type or gender confusion, caused by missing training samples. In future work, we would like to further open up the problem of human shape estimation and explore shape representations that allow all types of clothing and even accessories.

We have shown, that by transferring a hard problem into a simple formulation, complex models can be outperformed. Our method lays the foundation for wide-spread 3D reconstruction of people for various applications and even from legacy material.

References

[1] http://virtualhumans.mpi-inf.mpg.de/tex2shape/. 2

[2] 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 *IEEE Conf. on Computer Vision and Pattern Recognition*, 2019. 1, 2, 3

[3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *International Conf. on 3D Vision*, sep 2018. 2

[4] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3D people models. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2018. 1, 2, 5, 7

[5] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. Densepose: Dense human pose estimation in the wild. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 7297–7306, 2018. 2, 3, 4, 7

[6] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. In *ACM Transactions on Graphics*, volume 24, pages 408–416. ACM, 2005. 2

[7] Jan Bednarik, Pascal Fua, and Mathieu Salzmann. Learning to reconstruct texture-less deformable surfaces from a single view. In *International Conf. on 3D Vision*, pages 606–615, 2018. 3

[8] James F Blinn and Martin E Newell. Texture and reflection in computer generated images. *Communications of the ACM*, 19(10):542–547, 1976. 3

[9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conf. on Computer Vision*. Springer, 2016. 1, 2

[10] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 2017. 3

[11] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2017. 2, 4, 7

[12] R Daněřek, Endri Dibra, C Öztireli, Remo Ziegler, and Markus Gross. Deepgarment: 3d garment shape estimation from a single image. In *Computer Graphics Forum*, volume 36, pages 269–280. Wiley Online Library, 2017. 3

[13] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gambaretto, Christian Gagné, and Jean-François Lalonde. Learning to predict indoor illumination from a single image. *ACM Transactions on Graphics*, 9(4), 2017. 4

[14] Peng Guan, Alexander Weiss, Alexandru O Bălan, and Michael J Black. Estimating human shape and pose from a single image. In *IEEE International Conf. on Computer Vision*, 2009. 2

[15] Marc Habermann, Weipeng Xu, Michael Zollhöfer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM Transactions on Graphics*, 38(2):14:1–14:17, 2019. 1, 2

[16] Nils Hasler, Carsten Stoll, Martin Sunkel, Bodo Rosenhahn, and H-P Seidel. A statistical model of human pose and body shape. In *Computer Graphics Forum*, 2009. 2

[17] Loc Huynh, Weikai Chen, Shunsuke Saito, Jun Xing, Koki Nagano, Andrew Jones, Paul Debevec, and Hao Li. Mesoscopic facial geometry inference using deep neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8407–8416, 2018. 3

[18] Eldar Insafutdinov, Leonid Pishchulin, Bjoern Andres, Mykhaylo Andriluka, and Bernt Schieke. Deepercut: A deeper, stronger, and faster multi-person pose estimation model. In *European Conf. on Computer Vision*, 2016. 2

[19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1125–1134, 2017. 2, 3, 4

[20] Aaron S Jackson, Chris Manafas, and Georgios Tzimiropoulos. 3d human body reconstruction from a single image via volumetric regression. In *European Conference on Computer Vision*, pages 64–77. Springer, 2018. 2

[21] Arjun Jain, Thorsten Thormählen, Hans-Peter Seidel, and Christian Theobalt. Moviereshape: Tracking and reshaping of humans in videos. In *ACM Transactions on Graphics*, volume 29, page 148. ACM, 2010. 2

[22] Ning Jin, Yilin Zhu, Zhenglin Geng, and Ronald Fedkiw. A pixel-based framework for data-driven clothing. *arXiv preprint arXiv:1812.01677*, 2018. 3

[23] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 8320–8329, 2018. 2

[24] Yoshihiro Kanamori and Yuki Endo. Relighting humans: occlusion-aware inverse rendering for fullbody human images. *ACM Transactions on Graphics*, 37(270):1–270, 2018. 4

[25] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2018. 1, 2, 3, 5, 7

[26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, volume 5, 2015. 5

[27] Zorah Lahner, Daniel Cremers, and Tony Tung. Deepwrinkles: Accurate and realistic clothing modeling. In *European Conf. on Computer Vision*, pages 667–684, 2018. 3

[28] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2017. 2

[29] Guannan Li, Chenglei Wu, Carsten Stoll, Yebin Liu, Kiran Varanasi, Qionghai Dai, and Christian Theobalt. Capturing relightable human performances under general uncontrolled illumination. In *Computer Graphics Forum*, 2013. 2

[30] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2016. 5

[31] Ziwei Liu, Sijie Yan, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Fashion landmark detection in the wild. In *European Conf. on Computer Vision*, 2016. 5

[32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics*, 2015. 1, 2, 3

[33] Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. Siclope: Silhouette-based clothed people. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2019. 1, 2, 3, 5, 7

[34] Natalia Neverova, Riza Alp Guler, and Iasonas Kokkinos. Dense pose transfer. In *European Conf. on Computer Vision*, 2018. 2

[35] Mohamed Omran, Christop Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *International Conf. on 3D Vision*, 2018. 1, 2

[36] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3D human pose and shape from a single color image. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2018. 1, 2

[37] Leonid Pishchulin, Eldar Insafutdinov, Siyu Tang, Bjoern Andres, Mykhaylo Andriluka, Peter Gehler, and Bernt Schiele. Deepcut: Joint subset partition and labeling for multi person pose estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2016. 2

[38] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael Black. ClothCap: Seamless 4D clothing capture and retargeting. *ACM Transactions on Graphics*, 36(4), 2017. 3

[39] Gerard Pons-Moll, Javier Romero, Naureen Mahmood, and Michael J Black. Dyna: a model of dynamic human shape in motion. *ACM Transactions on Graphics*, 34:120, 2015. 2

[40] Tiberiu Popa, Quan Zhou, Derek Bradley, Vladislav Kraevoy, Hongbo Fu, Alla Sheffer, and Wolfgang Heidrich. Wrinkling captured garments using space-time data-driven deformation. In *Computer Graphics Forum*, volume 28, pages 427–435, 2009. 3

[41] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pages 497–500. ACM, 2001. 4

[42] Lorenz Rogge, Felix Klose, Michael Stengel, Martin Eisemann, and Marcus Magnor. Garment replacement in monocular video sequences. *ACM Transactions on Graphics*, 34(1):6, 2014. 2

[43] Matan Sela, Elad Richardson, and Ron Kimmel. Unrestricted facial geometry reconstruction using image-to-image translation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 1576–1585, 2017. 3

[44] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of facesin the wild'. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 6296–6305, 2018. 3

[45] Yu Tao, Zerong Zheng, Yuan Zhong, Jianhui Zhao, Dai Quionhai, Gerard Pons-Moll, and Yebin Liu. Simulcap : Single-view human performance capture with cloth simulation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, jun 2019. 3

[46] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohammed Elgharib, Hans-Peter Seidel, Patrick Perez, Michael Zollhöfer, and Christian Theobalt. Fml: Face model learning from videos. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2019. 3

[47] Ayush Tewari, Michael Zollhöfer, Pablo Garrido, Florian Bernard, Hyeongwoo Kim, Patrick Pérez, and Christian Theobalt. Self-supervised multi-level face model learning for monocular reconstruction at over 250 hz. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2018. 3

[48] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *Advances in Neural Information Processing Systems*, pages 5236–5246, 2017. 2

[49] Gül Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *European Conf. on Computer Vision*, 2018. 2, 5, 7

[50] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2017. 1

[51] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *European Conf. on Computer Vision*, sep 2018. 5

[52] Tuanfeng Y. Wang, Duygu Ceylan, Jovan Popovic, and Niloy J. Mitra. Learning a shared shape space for multimodal garment design. *ACM Transactions on Graphics*, 37(6):1:1– 1:14, 2018. 3

[53] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *Asilomar Conference on Signals, Systems & Computers*, volume 2, pages 1398–1402, 2003. 5

[54] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Photo wake-up: 3d character animation from a single photo. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2019. 2, 3

[55] Chenglei Wu, Carsten Stoll, Levi Valgaerts, and Christian Theobalt. On-set performance capture of multiple actors with a stereo camera. *ACM Transactions on Graphics*, 32(6):161, 2013. 2

[56] Chenglei Wu, Kiran Varanasi, and Christian Theobalt. Full body performance capture under uncontrolled and varying illumination: A shading-based approach. In *European Conf. on Computer Vision*, pages 757–770, 2012. 2

[57] Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian

Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics*, 2018. 1, 2

[58] Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhrer. Analyzing clothing layer deformation statistics of 3d human motions. In *European Conf. on Computer Vision*, pages 237–253, 2018. 3

[59] Chao Zhang, Sergi Pujades, Michael Black, and Gerard Pons-Moll. Detailed, accurate, human shape estimation from clothed 3D scan sequences. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2017. 3

[60] Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999. 2

[61] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. *arXiv preprint arXiv:1903.06473*, Sept 2019. 2

[62] Shizhe Zhou, Hongbo Fu, Ligang Liu, Daniel Cohen-Or, and Xiaoguang Han. Parametric reshaping of human bodies in images. In *ACM Transactions on Graphics*, volume 29, page 126. ACM, 2010. 2