

Recovering Geometric Information with Learned Texture Perturbations

JANE WU, Stanford University, USA

YONGXU JIN, Stanford University, USA

ZHENGLIN GENG, Epic Games, USA

HUI ZHOU, JD.com, USA

RONALD FEDKIW, Stanford University, Epic Games, USA

Regularization is used to avoid overfitting when training a neural network; unfortunately, this reduces the attainable level of detail hindering the ability to capture high-frequency information present in the training data. Even though various approaches may be used to re-introduce high-frequency detail, it typically does not match the training data and is often not time coherent. In the case of network inferred cloth, these sentiments manifest themselves via either a lack of detailed wrinkles or unnaturally appearing and/or time incoherent surrogate wrinkles. Thus, we propose a general strategy whereby high-frequency information is procedurally embedded into low-frequency data so that when the latter is smeared out by the network the former still retains its high-frequency detail. We illustrate this approach by learning texture coordinates which when smeared do not in turn smear out the high-frequency detail in the texture itself but merely smoothly distort it. Notably, we prescribe perturbed texture coordinates that are subsequently used to correct the over-smoothed appearance of inferred cloth, and correcting the appearance from multiple camera views naturally recovers lost geometric information.

CCS Concepts: • **Computing methodologies** → **Animation**; **Neural networks**; **Computer vision representations**.

Additional Key Words and Phrases: cloth, folds, wrinkles, neural networks

ACM Reference Format:

Jane Wu, Yongxu Jin, Zhenglin Geng, Hui Zhou, and Ronald Fedkiw. 2021. Recovering Geometric Information with Learned Texture Perturbations. *Proc. ACM Comput. Graph. Interact. Tech.* 4, 3, Article 38 (September 2021), 18 pages. <https://doi.org/10.1145/3480137>

1 INTRODUCTION

Since neural networks are trained to generalize to unseen data, regularization is important for reducing overfitting, see e.g. [Goodfellow et al. 2016; Scholkopf and Smola 2001]. However, regularization also removes some of the high variance characteristic of much of the physical world. Even though high-quality ground truth data can be collected or generated to reflect the desired complexity of the outputs, regularization will inevitably smooth network predictions. Rather than attempting to directly infer high-frequency features, we alternatively propose to learn a low-frequency space in which such features can be embedded.

Authors' addresses: Jane Wu, janehwu@stanford.edu, Stanford University, USA; Yongxu Jin, yxjin@stanford.edu, Stanford University, USA; Zhenglin Geng, zhenglin.geng@epicgames.com, Epic Games, USA; Hui Zhou, hui.zhou@ieee.org, JD.com, USA; Ronald Fedkiw, fedkiw@cs.stanford.edu, Stanford University, Epic Games, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

2577-6193/2021/9-ART38 \$15.00

<https://doi.org/10.1145/3480137>

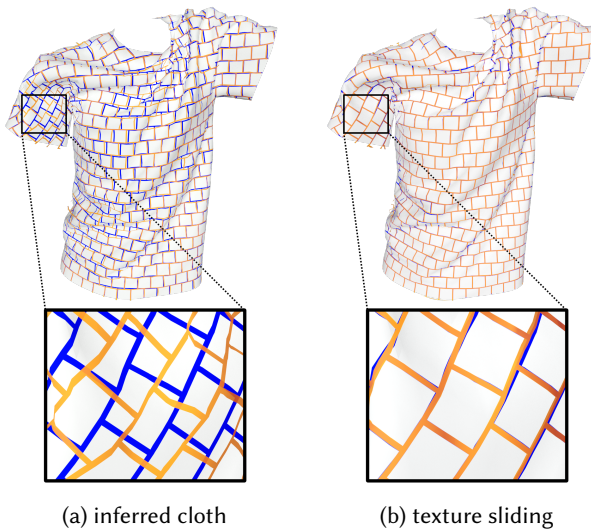


Fig. 1. Texture coordinate perturbations (texture sliding) reduce shape inference errors: ground truth (blue), prediction (orange).

We focus on the specific task of adding high-frequency wrinkles to virtual clothing, noting that the idea of learning a low-frequency embedding may be generalized to other tasks. Because cloth wrinkles/folds are high-frequency features, existing deep neural networks (DNNs) trained to infer cloth shape tend to predict overly smooth meshes [Alldieck et al. 2019a; Daněřek et al. 2017; Guan et al. 2012; Gundogdu et al. 2019; Jin et al. 2020; Lahner et al. 2018; Natsume et al. 2019; Patel et al. 2020; Santesteban et al. 2019; Wang et al. 2018]. Rather than attempting to amend such errors directly, we perturb texture so that the rendered cloth mesh *appears* to more closely match the ground truth. See Figure 1. Then given texture perturbations from at least two unique camera views, 3D geometry can be accurately reconstructed [Hartley and Sturm 1997] to recover high-frequency wrinkles. Similarly, for AR/VR applications, correcting visual appearance from two views (one for each eye) is enough to allow the viewer to accurately discern 3D geometry.

Our proposed texture coordinate perturbations are highly dependent on the camera view. Thus, we demonstrate that one can train a separate texture sliding neural network (TSNN) for each of a finite number of cameras laid out into an array and use nearby networks to interpolate results valid for any view enveloped by the array. Although an approach similar in spirit might be pursued for various lighting conditions, this limitation is left as future work since there are a great deal of applications where the light is ambient/diffuse/non-directional/etc. In such situations, this further complication may be ignored without significant repercussion.

2 RELATED WORK

Cloth: While physically-based cloth simulation has matured as a field over the last few decades [Baraff and Witkin 1998; Baraff et al. 2003; Bridson et al. 2002, 2003; Selle et al. 2008], data-driven methods are attractive for many applications. There is a rich body of work in reconstructing cloth from multiple views or 3D scans, see e.g. [Bradley et al. 2008b; Franco et al. 2006; Vlasic et al. 2008]. More recently, optimization-based methods have been used to generate higher resolution reconstructions [Huang et al. 2015; Pons-Moll et al. 2017; Wu et al. 2012; Yang et al. 2016]. Some of the most interesting work focuses on reconstructing the body and cloth separately [Bálan and

Black 2008; Neophytou and Hilton 2014; Yang et al. 2018; Zhang et al. 2017]. With advances in deep learning, one can aim to reconstruct 3D cloth meshes from single views. A number of approaches reconstruct a joint cloth/body mesh from a single RGB image [Alldieck et al. 2019a,b; Jiang et al. 2020; Natsume et al. 2019; Onizuka et al. 2020; Saito et al. 2019, 2020], RGB-D image [Yu et al. 2019], or video [Alldieck et al. 2018a,b; Habermann et al. 2019; Xu et al. 2018]. To reduce the dimensionality of the output space, DNNs are often trained to predict the pose/shape parameters of human body models such as SCAPE [Angelov et al. 2005] or SMPL [Loper et al. 2015] (see also [Pavlakos et al. 2019]). [Habermann et al. 2019; Natsume et al. 2019; Varol et al. 2018] leverage predicted pose information to infer shape. When only the garment shape is predicted, a number of recent works output predictions in UV space to represent geometric information as pixels [Daněřek et al. 2017; Jin et al. 2020; Lahner et al. 2018], although others [Gundogdu et al. 2019; Patel et al. 2020; Santesteban et al. 2019] define loss functions directly in terms of the 3D cloth vertices.

Wrinkles and Folds: Cloth realism can be improved by introducing wrinkles and folds. In the graphics community, researchers have explored both procedural and data-driven methods for generating wrinkles [De Aguiar et al. 2010; Guan et al. 2012; Hahn et al. 2014; Müller and Chentanez 2010; Rohmer et al. 2010; Wang et al. 2010]. Other works add real-world wrinkles as a postprocessing step to improve smooth captured cloth: [Popa et al. 2009] extracts the edges of cloth folds and then applies space-time deformations, [Robertini et al. 2014] solves for shape deformations directly by optimizing over all frames of a video sequence. [Buffet et al. 2019] improves cloth realism by untangling multiple garment layers, thereby producing a collision-free state that is ready for animation. Recently, [Lahner et al. 2018] used a conditional Generative Adversarial Network [Mirza and Osindero 2014] to generate normal maps as proxies for wrinkles on captured cloth.

Geometry: More broadly, deep learning on 3D meshes falls under the umbrella of *geometric deep learning*, which was coined by [Bronstein et al. 2017] to characterize learning in non-Euclidean domains. [Scarselli et al. 2008] was one of the earliest works in this area and introduced the notion of a Graph Neural Network (GNN) in relation to CNNs. Subsequent works similarly extend the CNN architecture to graphs and manifolds [Boscaini et al. 2016; Maron et al. 2017; Masci et al. 2015; Monti et al. 2017]. [Kostrikov et al. 2018] introduces a latent representation that explicitly incorporates the Dirac operator to detect principal curvature directions. [Hertz et al. 2020; Tan et al. 2018] train mesh generative models to infer novel meshes outside an original dataset. Returning to the specific application of virtual cloth, [Jin et al. 2020] embeds a non-Euclidean cloth mesh into a Euclidean pixel space, making it possible to directly use CNNs to make non-Euclidean predictions.

Texture: In the computer graphics community, textures have historically been used to capture both geometric and material details lost by using simplified models [Foley et al. 1996; Marschner and Shirley 2015], which is similar in spirit to our approach. Though, to the best of our knowledge, we are the first to propose learning texture coordinate perturbations to facilitate the accurate reconstruction of lost geometric details. For completeness, we briefly note a few works that use learning for texture synthesis and/or style transfer [Dumoulin et al. 2016; Gatys et al. 2015, 2016; Gupta et al. 2017; Johnson et al. 2016; Sanakoyeu et al. 2018].

3 METHODS

We define *texture sliding* as the changing of texture coordinates on a per-camera basis such that any point which is visible from some stereo pair of cameras can be triangulated back to its ground truth position. Other stereo reconstruction techniques can also be used in place of triangulation because the *images* we generate are consistent with the ground truth geometry. See e.g. [Bradley

et al. 2008a; Hartley and Sturm 1997; Seitz et al. 2006]. Thus, texture sliding recovers high-frequency 3D geometry embedded in 2D material space perturbations.

3.1 Per-Vertex Discretization

Since the cloth mesh is discretized into vertices and triangles, we take a per-vertex, not a per-point, approach to texture sliding. Our proposed method (see Section 4.1) computes per-vertex texture coordinates on the inferred cloth that match those of the ground truth as seen by the camera under consideration. Then during 3D reconstruction, barycentric interpolation is used to find the subtriangle locations of the texture coordinates corresponding to ground truth cloth vertices. This assumes linearity, which is only valid when the triangles are small enough to capture the inherent nonlinearities in a piecewise linear sense; moreover, folds and wrinkles can create significant nonlinearity. See Figure 2.

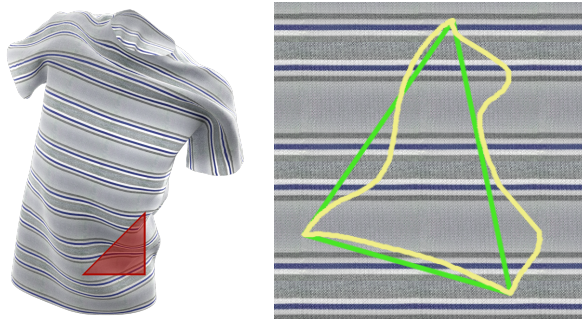


Fig. 2. Consider an extreme case, where the inferred cloth has a quite large triangle (shown in red). That triangle should encompass the nonlinear texture region outlined in yellow (shown in pattern space). Note: the yellow curve was generated by sampling the ground truth cloth's texture coordinates along the projected edges of the red triangle. The linearity assumption implied by barycentric interpolation instead uses the region outlined in green.

3.2 Occlusion Boundaries

Accurate 3D reconstruction requires that a vertex of the ground truth mesh be visible from at least two cameras *and* that camera projections of the vertex to the inferred cloth exist and are valid. However, occlusions can derail these assumptions.

First, consider things from the standpoint of the inferred cloth. For a given camera view, some inferred cloth triangles will not contain any visible pixels, and we denote a vertex as occluded when none of its incident triangles contain any visible pixels. Although we do not assign perturbed texture coordinates to occluded vertices (i.e. they keep their original texture coordinates, or a perturbation of zero), we do aim to keep the texture coordinate perturbation function smooth (see Section 4.2). In addition, there will be so called non-occluded vertices in the inferred cloth that do not project to visible pixels of the ground truth cloth. This often occurs near silhouette boundaries where the inferred cloth silhouette is sometimes wider than the ground truth cloth silhouette. These vertices are also treated as occluded, similar to those around the back side of the cloth behind the silhouette, essentially treating some extra vertices near occlusion boundaries as also being occluded. See Figure 3a.

Next, consider things from the standpoint of the ground truth cloth. For example, consider the case where all the cameras are in the front, and vertices on the back side of the ground truth cloth

are not visible from any camera. The best one can do in reconstructing these occluded vertices is to use the inferred cloth vertex positions; however, care should be taken near occlusion boundaries to smoothly taper between our texture sliding 3D reconstruction and the inferred cloth prediction. A simple approach is to extrapolate/smooth the geometric difference between our texture sliding 3D reconstruction and the inferred cloth prediction to occluded regions of the mesh. Once again, the definition of occluded vertices needs to be broadened for silhouette consideration. Not only will vertices not visible from at least two cameras have to be considered occluded, but vertices that don't project to the interior of an inferred cloth triangle with *valid* texture coordinate perturbations will also have to be considered occluded. See Figure 3b.

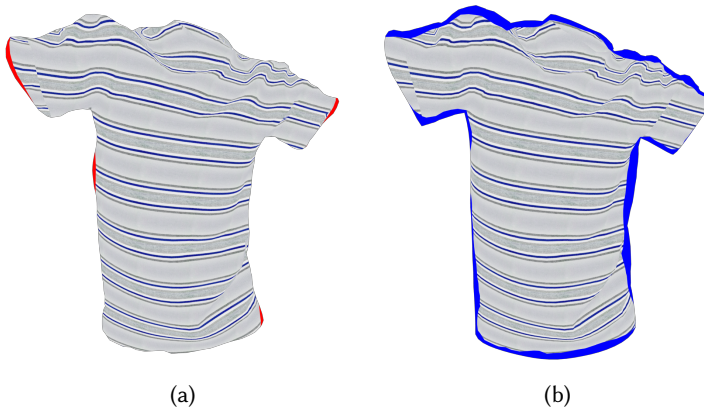


Fig. 3. The method discussed in Section 4.1 can fail near silhouettes of the inferred and ground truth cloth meshes, in which case smoothness assumptions are used (see Section 4.2). In (a), inferred triangles with at least one vertex falling outside the silhouette of the ground truth mesh are colored red. In (b), ground truth triangles with at least one vertex falling outside the silhouette of the inferred mesh are colored blue.

4 DATASET GENERATION

Let $C = \{X, T\}$ be a cloth triangulated surface with n vertices $X \in \mathbb{R}^{3n}$ and texture coordinates $T \in \mathbb{R}^{2n}$. We assume that mesh connectivity remains fixed throughout. The ground truth cloth mesh $C_G(\theta) = \{X_G(\theta), T_G\}$ depends on the pose θ . Given a pre-trained DNN (we use the network from [Jin et al. 2020]), the inferred cloth $C_N(\theta) = \{X_N(\theta), T_G\}$ is also a function of the pose θ . Our objective is to replace the ground truth texture coordinates T_G with perturbed texture coordinates $T_N(\theta, v)$, i.e. to compute $C'_N(\theta, v) = \{X_N(\theta), T_N(\theta, v)\}$ where $T_N(\theta, v)$ depends on both the pose θ and the view v . Even though $T_N(\theta, v)$ is in principle valid for all v using interpolation (see Section 6.3), training data $T_N(\theta, v_p)$ is only required for a finite number of camera views v_p . For each camera p , we also only require training data for finite number of poses θ_k , i.e. we require $T_N(\theta_k, v_p)$, which is computed from T_G using $X_G(\theta_k)$, $X_N(\theta_k)$, and v_p .

4.1 Texture Coordinate Projection

We project texture coordinates to the inferred cloth vertices $X_N(\theta_k)$ from the ground truth cloth mesh $C_G(\theta_k)$ using ray intersection. For each inferred cloth vertex in $X_N(\theta_k)$, we cast a ray from camera p 's aperture through the vertex and find the first intersection with the ground truth mesh $C_G(\theta_k)$; subsequently, T_G is barycentrically interpolated to the point of intersection and assigned to the inferred cloth vertex as its $T_N(\theta_k, v_p)$ value. See Algorithm 1 and Figure 4. Rays are only cast

for inferred cloth vertices that have at least one incident triangle with a nonzero area subregion visible to camera p . Also, a ground truth texture coordinate value is only assigned to an inferred cloth vertex when the point of intersection with the ground truth mesh is visible to camera p . We store and learn texture coordinate displacements $d_{v_p}(\theta_k) = T_N(\theta_k, v_p) - T_G$. After this procedure, any remaining vertices of the inferred cloth that have not been assigned $d_{v_p}(\theta_k)$ values are treated as occluded and handled via smoothness considerations as discussed in Section 4.2.

Algorithm 1: Ray Intersection

```

1 for vertex  $v$  in inferred mesh  $X_N(\theta_k)$  do
2    $r \leftarrow \text{ray}(\text{cameraAperture}(p), v)$ ;
3    $x_{\text{intersection}} \leftarrow \text{firstIntersection}(r, \text{ground truth mesh } C_G(\theta_k))$ ;
4    $T_N(\theta_k, v_p)(v) \leftarrow \text{barycentricInterpolation}(x_{\text{intersection}}, C_G(\theta_k), T_G)$ ;

```

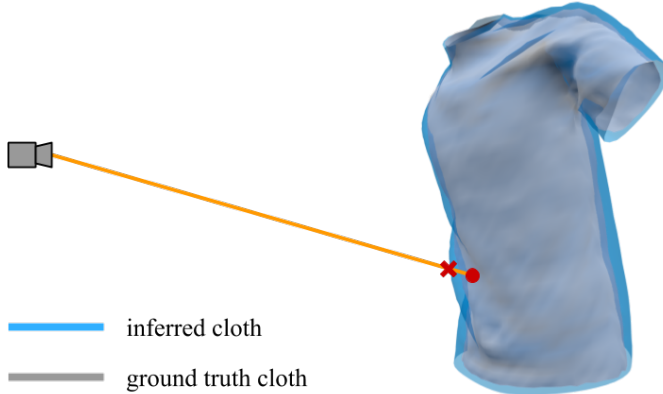


Fig. 4. Illustration of the ray intersection method for transferring texture coordinates to the inferred cloth from the ground truth cloth. Texture coordinates for the inferred cloth vertex (red cross) are interpolated from the ground truth mesh to the point of ray intersection (red circle).

4.2 Occlusion Handling

Some vertices of the inferred cloth mesh remain unassigned with $d_{v_p}(\theta_k) = 0$ after executing the algorithm outlined in Section 4.1. This creates a discontinuity in $d_{v_p}(\theta_k)$ which excites high frequencies that require a more complex network architecture to capture. In order to alleviate demands on the network, we smooth $d_{v_p}(\theta_k)$ as follows. First, we use the Fast Marching Method on triangulated surfaces [Kimmel and Sethian 1998] to generate a signed distance field. Then, we extrapolate $d_{v_p}(\theta_k)$ normal to the distance field into the unassigned region, see e.g. [Osher and Fedkiw 2002]. Finally, a bit of averaging can be used to provide smoothness, while keeping the assigned values of $d_{v_p}(\theta_k)$ unchanged. Alternatively, we solve a Poisson equation as in [Cong et al. 2015] while using the assigned $d_{v_p}(\theta_k)$ as Dirichlet boundary conditions.

4.3 Topological Considerations

There are some edge cases that require additional topological consideration. In particular, the collar, sleeves, and waist are areas where a ray cast to an inferred cloth vertex can intersect with a back-facing triangle on the inside of the ground truth shirt. We aim to define texture coordinates on inferred cloth vertices so that barycentric interpolation can be used to find the texture coordinates of a ground truth vertex for 3D reconstruction. However, mixing texture coordinates from the inside and outside of the shirt in a single triangle causes dramatic interpolation error. In fact, as shown in Figure 2, large errors may occur for any triangle that mixes texture coordinates from geodesically far-away regions. Thus, we omit such triangles from consideration by omitting a vertex from any edge that connects two geodesically far-away regions.

As a further improvement to our method, one can treat the inside and outside of the shirt as separate meshes, applying texture sliding twice and training two separate networks; moreover, one may take a patch-based approach, applying TS and training a TSNN for each (slightly overlapping) patch of the shirt.

4.4 Smoothness Considerations

When training a neural network, more predictable results are obtained when the inferred cloth vertex data is smoother. Thus, there exists tradeoffs between smoothness and accuracy when assigning texture coordinates. An edge that connects two geodesically far-away regions introduces a jump discontinuity in the texture coordinates leading to high frequencies in the ground truth data that place increased demands on the network. Although subdivision adds degrees of freedom along such edges to better sample the high frequency, it is often better to delete such edges entirely by removing one of the edge’s vertices. Recall that any vertex not assigned a ground truth texture coordinate is instead defined via smoothness considerations (see Section 4.2) reducing demands on the network.

5 NETWORK ARCHITECTURE

A separate texture sliding neural network (TSNN) is trained for each camera p ; thus, we drop the v_p notation in this section. The loss is defined over all poses θ_k in the training set

$$\mathcal{L} = \sum_{\theta_k} \left\| d(\theta_k) - \hat{d}(\theta_k) \right\|_2 \quad (1)$$

to minimize the difference between the desired displacements $d(\theta_k)$ and predicted displacements $\hat{d}(\theta_k)$. The inferred cloth data we chose to correct are predictions of the T-shirt meshes from [Jin et al. 2020], each of which contains about 3,000 vertices. The dataset spans about 10,000 different poses, which were sampled from a uniformly random distribution over the body joints’ natural range of motion. The cloth mesh was generated from a scanned garment and simulated in a physics engine [phy [n.d.]] and includes texture coordinates for each vertex. To improve the resolution, we up-sampled each cloth mesh by subdividing each triangle into four subtriangles. Notably, our texture sliding approach can be used to augment the results of any dataset for which ground truth and inferred training examples are available. Moreover, it is trivial to increase the resolution of any such dataset simply by subdividing triangles. Note that perturbations of the subdivided geometry are unnecessary, as we merely desire more sample points (to address Figure 2). Finally, we applied an 80-10-10 training-validation-test set split.

Similar to [Jin et al. 2020], the displacements $d(\theta_k)$ are stored as pixel-based cloth images for the front and back sides of the T-shirt, though we still output per-vertex texture coordinate displacements in UV space. See Figure 5 for an overview of the network architecture. Given input

joint transformation matrices of shape $1 \times 1 \times 90$, TSNN applies a series of transpose convolution, batch normalization, and ReLU activation layers to upsample the input to $512 \times 512 \times 4$. The first two dimensions of the output tensor represent the predicted displacements for the front side of the T-shirt, and the remaining two dimensions represent those for the back side.

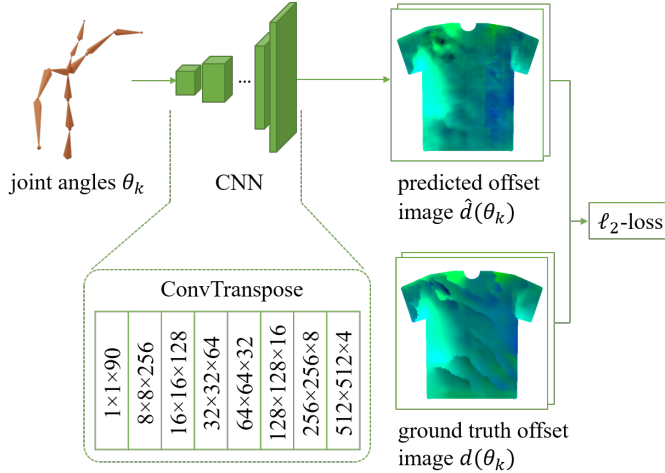


Fig. 5. Texture sliding neural network (TSNN) architecture.

6 EXPERIMENTS

In Section 6.1, we quantify the data generation approach of Section 4 and highlight the advantages of mesh subdivision for up-sampling. In Section 6.2, we evaluate the predictions made by our trained texture sliding neural network (TSNN). In Section 6.3, we demonstrate the interpolation of texture sliding results to novel views between a finite number of cameras. Finally, in Section 6.4, we use multi-view texture sliding to reconstruct 3D geometry.

6.1 Dataset Generation and Evaluation

We aim to have the material coordinates of the cloth be in the correct locations as viewed by multiple cameras, so that the material can be accurately 3D reconstructed with point-wise accuracy. As such, our error metric is a bit more stringent than that commonly used because our aim is to reproduce the actual material behavior, not merely to mimic its look (e.g., by perturbing normal vectors to create shading consistent with wrinkles in spite of the cloth being smooth, as in [Lahner et al. 2018]). In order to elucidate this, consider a two-step approach where one first approximates a smooth cloth mesh and then perturbs that mesh to add wrinkles (similar to [Santesteban et al. 2019]). In order to preserve area and achieve the correct material behavior, material in the vicinity of a newly forming wrinkle should slide laterally towards that wrinkle as it is formed. Merely non-physically stretching the material in order to create a wrinkle may look plausible, but does not admit the correct material behavior. In fact, the texture would be unrealistically stretched as well, although this is less apparent visually when using simple textures.

Since texture coordinates provide a proxy surface parameterization for material coordinates, we measure texture coordinate errors in a per-pixel fashion comparing between the ground truth and inferred cloth at the center of each pixel. Figure 6a shows results typical for cloth inferred using the network from [Jin et al. 2020], and Figure 6b shows the highly improved results obtained on

the same inferred geometry using our texture sliding approach (with 1 level of subdivision). Note that the vast majority of the errors in Figure 6b occur near the wrinkles where the nonlinearities illustrated in Figure 2 are most prevalent. In Figure 6c, we deform the vertices of the inferred cloth mesh so that they lie exactly on the ground truth mesh in order to mimic a two-step approach (as discussed above). Note how our error metric captures the still rather large errors in the material coordinates (and thus cloth vertex positions) in spite of the mesh in Figure 6c appearing to have the same wrinkles and folds as the ground truth mesh. Figure 7 compares the local compression and extension energies of the ground truth mesh (Figure 7a), the inferred cloth mesh (Figure 7b), and the result of this two-step process (Figure 7c). In spite of the untextured mesh in Figure 7c bearing visual similarity to the ground truth in Figure 7a, it still has rather large errors in deformation energy.

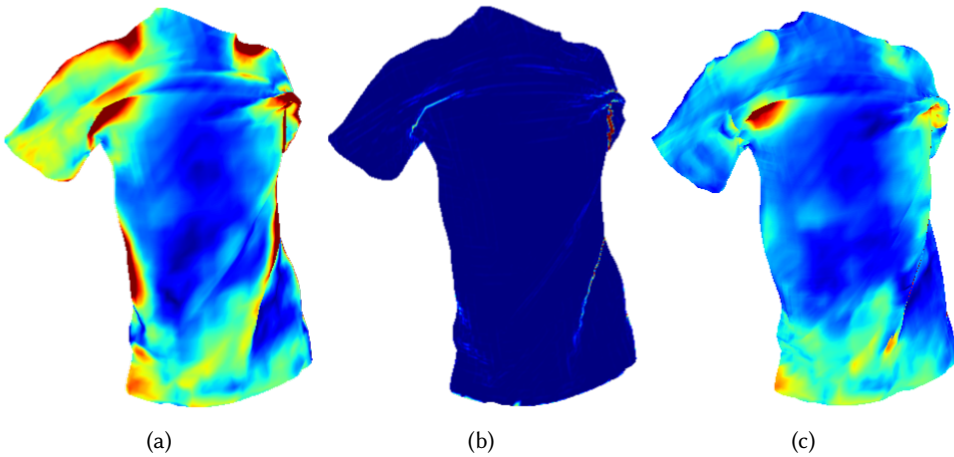


Fig. 6. Per-pixel texture coordinate errors before (a) and after (b) applying texture sliding to the inferred cloth output by the network of [Jin et al. 2020]. The result of a two-step process (c) may well match the ground truth in a visual sense, whilst still having quite large errors in material coordinates. Blue = 0, red ≥ 0.04 .

Figure 8 illustrates the efficacy of subdividing the cloth mesh to get more samples for texture sliding. The particular ground truth cloth wrinkle shown in Figure 8e is not captured by the inferred cloth geometry shown in Figure 8a. The texture sliding result shown in Figure 8b better represents the ground truth cloth. Figures 8c and 8d show how subdividing the inferred cloth mesh one and two times (respectively) progressively alleviates errors emanating from the linearity assumption illustrated in Figure 2. Table 1 shows quantitative results comparing the inferred cloth to texture sliding with and without subdivision.

Table 1. Per-pixel square root of mean squared error (SqrtMSE) for the entire dataset.

Method	SqrtMSE ($\times 10^{-3}$)
[Jin et al. 2020]	24.496 ± 6.9536
TS	5.2662 ± 2.2320
TS + subdivision	3.5645 ± 1.6617




Fig. 7. Local compression (blue) and extension (red) energies for a sample pose, comparing the ground truth cloth (a), the inferred cloth (b), and the result of a two-step process (c). In spite of the cloth mesh in (c) bearing visual resemblance to the ground truth in (a), it still has quite erroneous deformation energies.

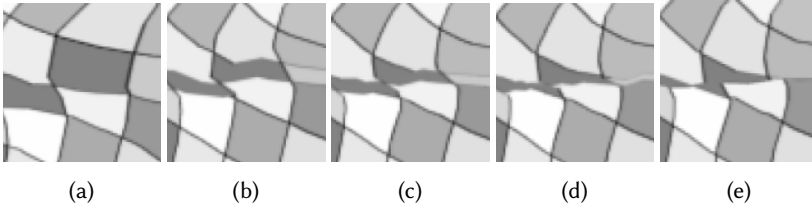


Fig. 8. As the inferred cloth mesh (a) is subdivided, texture sliding (b-d) moves the appearance of the inferred mesh closer to the ground truth (e). This is because mesh subdivision reduces barycentric interpolation errors from the ray intersection method discussed in Section 4.

6.2 Network Training and Inference

The network was trained using the Adam optimizer [Kingma and Ba 2014] with a 10^{-3} learning rate in PyTorch [Paszke et al. 2017]. As mentioned earlier, we subdivided the mesh triangles once. Figure 9 shows a typical prediction on a test set example, including the per-pixel errors in predicted texture coordinates. See the supplemental material for additional test set example predictions. While the TSNN is able to recover the majority of the shirt, it struggles near wrinkles. Figure 10 highlights a particular wrinkle comparing the inferred cloth (Figure 10a) and the results of the TSNN before (Figure 10b) and after (Figure 10c) subdivision to the ground truth (Figure 10d). Table 2 shows quantitative results comparing the inferred cloth to TSNN results with and without subdivision.

6.3 Interpolating to Novel Views

Given a finite number of camera views v_p , one can specify a new view enveloped by the array using a variety of interpolation methods. For the sake of demonstration, we take a simple approach assuming that one can interpolate via $v = \sum_p w_p v_p$, and then use these same weights to compute

$$T_N(\theta_k, v) = \sum_p w_p T_N(\theta_k, v_p) \quad (2)$$

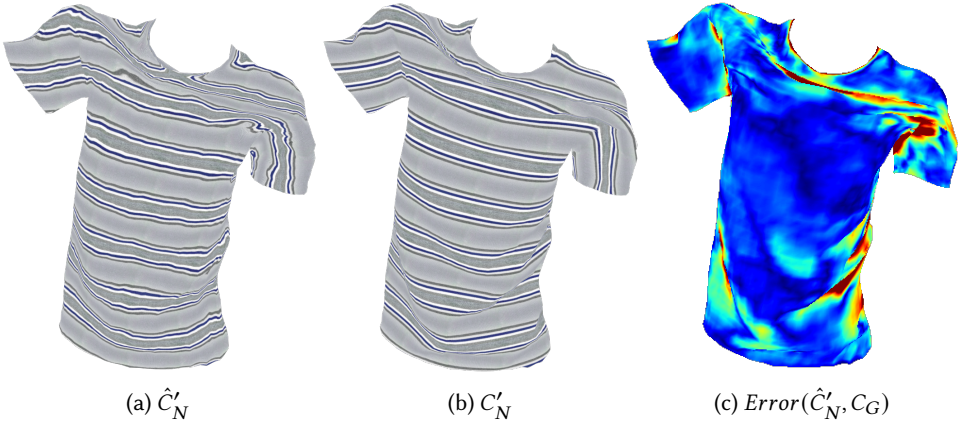


Fig. 9. A typical test set example prediction. The per-pixel errors are shown in (c) (blue = 0, red ≥ 0.04).

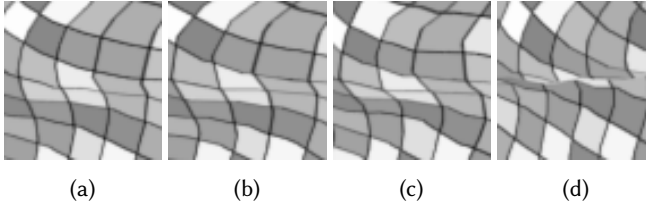


Fig. 10. The results of the TSNN before (b) and after (c) subdivision, as compared to the ground truth (d). In spite of Table 2, some wrinkles are better resolved by the TSNN after subdivision. The inferred mesh with ground truth texture coordinates is shown in (a).

Table 2. Per-pixel SqrtMSE for the test set. In spite of Table 1 demonstrating that subdivision improves the ground truth TS data, the improvements are not uniformly realized by the TSNN (which we discuss in Appendix C).

Network	SqrtMSE ($\times 10^{-3}$)
[Jin et al. 2020]	24.871 ± 7.0613
TSNN	13.335 ± 4.2924
TSNN + subdivision	13.591 ± 4.5194

This same equation is also used for $\hat{T}_N(\theta_k, v)$. Figure 11 shows the results obtained by linearly interpolating between two camera views. Note how the largest errors appear near areas occluded by wrinkles, where one (or both) of the cameras has no valid texture sliding results and instead uses the inferred cloth textures. This can be alleviated by using more cameras placed closer together. Figure 12 quantifies these results for the inferred cloth $C_N(\theta_k)$, texture sliding $C'_N(\theta_k, v)$, and the results of the TSNN $\hat{C}'_N(\theta_k, v)$. In Figure 13, we repeat these comparisons, except using bilinear interpolation between four camera views.

Interpolating between two cameras, each with TS or TSNN data, has the effect of following a straight-line path. However, by choosing the camera array and subsequent interpolation carefully one can interpolate along curved paths. For example in Figure 14, one can interpolate between the 12 cameras (represented by blue dots) in order to follow the curved camera path.



Fig. 11. Given two camera views (far left and far right images), texture sliding can be linearly interpolated to novel views between them. The top row shows per-pixel errors (blue = 0, red ≥ 0.04), and the bottom row shows the cloth from a *fixed* front-facing view to illustrate how the interpolated texture changes as a function of the chosen novel view.

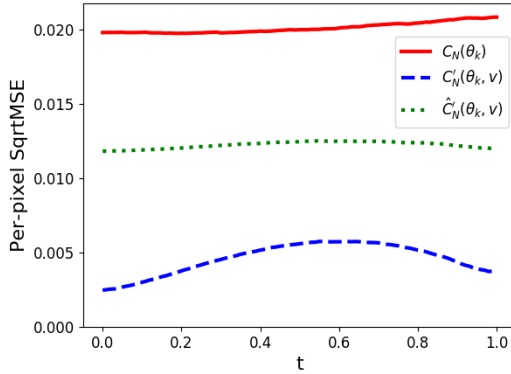


Fig. 12. Per-pixel SqrtMSE for interpolating between two cameras (using a test set example). Note that the inferred cloth does not use any view based information, but that our error metric does depend on the view.

6.4 3D Reconstruction

In order to reconstruct the 3D position of a vertex of the ground truth mesh, we take the usual approach of finding rays that pass through that vertex and the camera aperture for a number of cameras. Then given at least two rays, one can triangulate a 3D point that is minimal distance from all the rays. We can do this without solving the typical image to image correspondence problem because we know the ground truth texture coordinates for any given vertex. Thus, we merely have to find the ray that passes through the camera aperture and the ground truth texture coordinate for the vertex under consideration.

To find a ground truth texture coordinate on a texture corrected inferred cloth mesh $C'_N(\theta_k, v)$, or $\hat{C}'_N(\theta_k, v)$, we first find the triangle containing that texture coordinate. This can be done quickly by using a hierarchical bounding box structure where the base level boxes around each triangle are defined using the min/max texture coordinates at the three vertices. Then one can write the barycentric interpolation formula that interpolates the triangle vertex texture coordinates to obtain

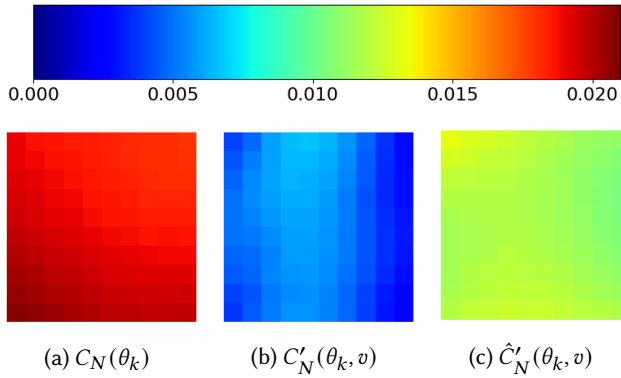


Fig. 13. Per-pixel SqrtMSE for interpolating between four cameras (one at each corner of the square). The pose θ_k is the same as in Figure 12, which plots the values along the bottom edge of the square.

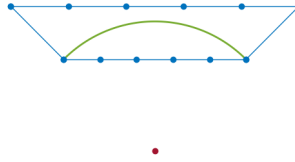


Fig. 14. Let the red dot represent the center of the cloth mesh. One can interpolate between the 12 cameras (blue dots) on the trapezoid in order to follow the curved camera path.

the given ground truth texture coordinate, and subsequently invert the matrix to solve for the weights. These weights determine the sub-triangle position of the vertex under consideration (taking care to note that different answers are obtained in 3D space versus screen space, since the camera projection is nonlinear). Figure 15 shows the 3D reconstruction of a test set example using texture sliding (Figure 15c) and the TSNN (Figure 15d). Figure 16 compares the per-pixel errors and local compression/extension energies of Figures 15c and 15d.

7 DISCUSSION AND FUTURE WORK

There are many disparate applications for clothing including for example video games, AR/VR, Hollywood special effects, virtual try-on and shopping, scene acquisition and understanding, and even bullet proof vests and soft armor. Various scenarios define accuracy or fidelity in vastly different ways. So while it is typical to state that one cares about more than just the visual appearance (or “graphics”), often those aiming for predictive capability still make concessions. For example, wherein [Santesteban et al. 2019] proposes a network that well predicts wrinkles mapped to new body types, the discussion in [Lahner et al. 2018] implies that the horizontal wrinkles predicted by [Santesteban et al. 2019] are more characteristic of inaccurate physical simulation than real-world behavior. Instead, [Lahner et al. 2018] strives for more vertical wrinkles to better match their data, but they accomplish this by predicting lighting to match an image while accepting overly smooth geometry. And as we have shown in Figure 7c, predicting the correct geometry still allows for rather large errors in the deformation (see [Geng et al. 2020]).

In light of this, we state the problem of most interest to us: Our aim is to study the efficacy of using deep neural networks to aid in the modeling of material behavior, especially for those materials for

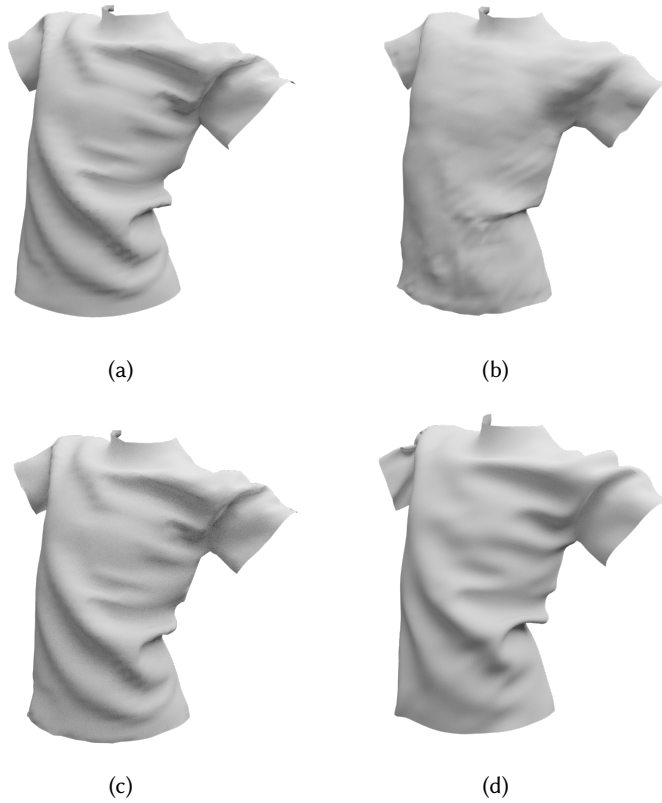


Fig. 15. Comparison of the ground truth cloth (a) and inferred cloth (b) to the 3D reconstructions obtained using texture sliding (c) and the TSNN (d). To remove reconstruction noise generated by network inference errors in (d), we used the physics-based postprocess from [Geng et al. 2020]; although, there are many other smoothing options in the literature that one might also consider.

which predictive methods do not currently exist because of various unknowns including friction, material parameters (for cloth and body), etc. Given this goal, we focus on the accurate prediction of material coordinates, which are a super set of deformation, geometry, lighting, visual plausibility, etc.

As demonstrated by the remarkably accurate 3D reconstruction in Figure 15c (see 16a), our approach to encoding high frequency wrinkles into lower frequency texture coordinates (i.e. texture sliding) works quite well. It can be used as a post-process to any existing neural network to capture lost details (as long as ground truth and inferred training examples are available); moreover, we showed that trivial subdivision could be used to increase the sampling resolution to limit linearization artifacts. The main drawback of our approach is that it relies on triangulation or multi-view stereo in order to construct the final 3D geometry, although this step is not required for AR/VR applications. One needs to take care when training the texture sliding neural network (TSNN) since inference errors can cause reconstruction noise. Thus, as future work, we plan on experimenting with the network architecture, the size of the image used in the CNN, the smoothing methods near occlusion boundaries, the amount of subdivision, etc. It would also be interesting to consider more savvy multiview 3D reconstruction methods (particularly ones that employ DNNs;




Fig. 16. Per-pixel errors (top) and local compression/extension energies (bottom) for Figure 15c (a) and Figure 15d (b).

then, one might train the whole process end-to-end). In addition, as cloth data becomes increasingly available, we would like to extend our method to a variety of garment types.

ACKNOWLEDGMENTS

Research supported in part by ONR N00014-13-1-0346, ONR N00014-17-1-2174, and JD.com. We would like to thank Reza and Behzad at ONR for supporting our efforts into machine learning, as well as Rev Lebaredian and Michael Kass at NVIDIA for graciously loaning us a GeForce RTX 2080Ti to use for running experiments. We would also like to thank Matthew Cong, Yilin Zhu, and Zhengping Zhou for their insightful discussions, and Congyue Deng for generating realistic cloth textures.

REFERENCES

- [n.d.]. PhysBAM: physically based animation. <http://physbam.stanford.edu>.
Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. 2019a. Learning to reconstruct people in clothing from a single RGB camera. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1175–1186.

- Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. 2018a. Detailed human avatars from monocular video. In *2018 International Conference on 3D Vision (3DV)*. IEEE, 98–109.
- Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. 2018b. Video based reconstruction of 3d people models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8387–8397.
- Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. 2019b. Tex2Shape: Detailed Full Human Body Geometry from a Single Image. In *Proceedings of the International Conference on Computer Vision (ICCV)*. IEEE.
- Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. 2005. SCAPE: shape completion and animation of people. In *ACM transactions on graphics (TOG)*, Vol. 24. ACM, 408–416.
- Alexandru O Bălan and Michael J Black. 2008. The naked truth: Estimating body shape under clothing. In *European Conference on Computer Vision*. Springer, 15–29.
- David Baraff and Andrew Witkin. 1998. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. ACM, 43–54.
- David Baraff, Andrew Witkin, and Michael Kass. 2003. Untangling cloth. In *ACM Transactions on Graphics (TOG)*, Vol. 22. ACM, 862–870.
- Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. 2016. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*. 3189–3197.
- Derek Bradley, Tamy Boubekeur, and Wolfgang Heidrich. 2008a. Accurate multi-view reconstruction using robust binocular stereo and surface meshing. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1–8.
- Derek Bradley, Tiberiu Popa, Alla Sheffer, Wolfgang Heidrich, and Tamy Boubekeur. 2008b. Markerless garment capture. In *ACM Transactions on Graphics (TOG)*, Vol. 27. ACM, 99.
- Robert Bridson, Ronald Fedkiw, and John Anderson. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *ACM Transactions on Graphics (ToG)*, Vol. 21. ACM, 594–603.
- Robert Bridson, Sebastian Marino, and Ronald Fedkiw. 2003. Simulation of clothing with folds and wrinkles. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. ACM, 28–36.
- Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* 34, 4 (2017), 18–42.
- Thomas Buffet, Damien Rohmer, Loic Barthe, Laurence Boissieux, and Marie-Paule Cani. 2019. Implicit untangling: A robust solution for modeling layered clothing. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- Matthew Cong, Michael Bao, Jane L E, Kiran S Bhat, and Ronald Fedkiw. 2015. Fully automatic generation of anatomical face simulation models. In *Proceedings of the 14th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. ACM, 175–183.
- R Daněřek, Endri Dibra, Cengiz Öztireli, Remo Ziegler, and Markus Gross. 2017. Deepgarment: 3d garment shape estimation from a single image. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 269–280.
- Edilson De Aguiar, Leonid Sigal, Adrien Treuille, and Jessica K Hodgins. 2010. Stable spaces for real-time clothing. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 106.
- Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. 2016. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629* (2016).
- James D Foley, Foley Dan Van, Andries Van Dam, Steven K Feiner, John F Hughes, J Hughes, and Edward Angel. 1996. *Computer graphics: principles and practice*. Vol. 12110. Addison-Wesley Professional.
- Jean-Sébastien Franco, Marc Lapierre, and Edmond Boyer. 2006. Visual shapes of silhouette sets. In *Third International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT'06)*. IEEE, 397–404.
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2015. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576* (2015).
- Leon A Gatys, Alexander S Ecker, and Matthias Bethge. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2414–2423.
- Zhenglin Geng, Daniel Johnson, and Ronald Fedkiw. 2020. Coercing machine learning to output physically accurate results. *J. Comput. Phys.* 406 (2020), 109099.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep learning*. MIT press.
- Peng Guan, Loretta Reiss, David A Hirshberg, Alexander Weiss, and Michael J Black. 2012. DRAPE: DRessing Any PErson. *ACM Trans. Graph.* 31, 4 (2012), 35–1.
- Erhan Gundogdu, Victor Constantin, Amrollah Seifoddini, Minh Dang, Mathieu Salzmann, and Pascal Fua. 2019. Garnet: A two-stream network for fast and accurate 3d cloth draping. In *Proceedings of the IEEE International Conference on Computer Vision*. 8739–8748.
- Agrim Gupta, Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2017. Characterizing and improving stability in neural style transfer. In *Proceedings of the IEEE International Conference on Computer Vision*. 4067–4076.
- Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. 2019. Livecap: Real-time human performance capture from monocular video. *ACM Transactions on Graphics (TOG)* 38, 2 (2019), 14.

- Fabian Hahn, Bernhard Thomaszewski, Stelian Coros, Robert W Sumner, Forrester Cole, Mark Meyer, Tony DeRose, and Markus Gross. 2014. Subspace clothing simulation using adaptive bases. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 105.
- Richard I Hartley and Peter Sturm. 1997. Triangulation. *Computer vision and image understanding* 68, 2 (1997), 146–157.
- Amir Hertz, Rana Hanocka, Raja Giryes, and Daniel Cohen-Or. 2020. Deep geometric texture synthesis. *arXiv preprint arXiv:2007.00074* (2020).
- Peng Huang, Margara Tejera, John Collomosse, and Adrian Hilton. 2015. Hybrid skeletal-surface motion graphs for character animation from 4d performance capture. *ACM Transactions on Graphics (ToG)* 34, 2 (2015), 17.
- Boyi Jiang, Juyong Zhang, Yang Hong, Jinhao Luo, Ligang Liu, and Hujun Bao. 2020. BCNet: Learning Body and Cloth Shape from A Single Image. *arXiv preprint arXiv:2004.00214* (2020).
- Ning Jin, Yilin Zhu, Zhenglin Geng, and Ronald Fedkiw. 2020. A pixel-based framework for data-driven clothing. In *Proceedings of the 19th ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, Vol. 39. Association for Computing Machinery.
- Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*. Springer, 694–711.
- Ron Kimmel and James A Sethian. 1998. Computing geodesic paths on manifolds. *Proceedings of the national academy of Sciences* 95, 15 (1998), 8431–8435.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- Ilya Kostrikov, Zhongshi Jiang, Daniele Panozzo, Denis Zorin, and Joan Bruna. 2018. Surface networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2540–2548.
- Zorah Lahner, Daniel Cremers, and Tony Tung. 2018. Deepwrinkles: Accurate and realistic clothing modeling. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 667–684.
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* 34, 6 (2015), 248.
- Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G Kim, and Yaron Lipman. 2017. Convolutional neural networks on surfaces via seamless toric covers. *ACM Trans. Graph.* 36, 4 (2017), 71–1.
- Steve Marschner and Peter Shirley. 2015. *Fundamentals of computer graphics*. CRC Press.
- Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. 2015. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*. 37–45.
- Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. 2017. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5115–5124.
- Matthias Müller and Nuttapon Chentanez. 2010. Wrinkle meshes. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics symposium on computer animation*. Eurographics Association, 85–92.
- Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. 2019. Siclope: Silhouette-based clothed people. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4480–4490.
- Alexandros Neophytou and Adrian Hilton. 2014. A layered model of human body and garment deformation. In *2014 2nd International Conference on 3D Vision*, Vol. 1. IEEE, 171–178.
- Hayato Onizuka, Zehra Hayirci, Diego Thomas, Akihiro Sugimoto, Hideaki Uchiyama, and Rin-ichiro Taniguchi. 2020. TetraTSDF: 3D human reconstruction from a single image with a tetrahedral outer shell. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6011–6020.
- Stanley Osher and Ronald Fedkiw. 2002. *Level Set Methods and Dynamic Implicit Surfaces*. Springer, New York.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. (2017).
- Chaitanya Patel, Zhouyingcheng Liao, and Gerard Pons-Moll. 2020. Tailornet: Predicting clothing in 3d as a function of human pose, shape and garment style. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7365–7375.
- Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 10975–10985.
- Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. 2017. ClothCap: Seamless 4D clothing capture and retargeting. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 73.
- Tiberiu Popa, Quan Zhou, Derek Bradley, Vladislav Kraevoy, Hongbo Fu, Alla Sheffer, and Wolfgang Heidrich. 2009. Wrinkling captured garments using space-time data-driven deformation. In *Computer Graphics Forum*, Vol. 28. Wiley

- Online Library, 427–435.
- Nadia Robertini, Edilson De Aguiar, Thomas Helten, and Christian Theobalt. 2014. Efficient multi-view performance capture of fine-scale surface detail. In *2014 2nd International Conference on 3D Vision*, Vol. 1. IEEE, 5–12.
- Damien Rohmer, Tiberiu Popa, Marie-Paule Cani, Stefanie Hahmann, and Alla Sheffer. 2010. Animation wrinkling: augmenting coarse cloth simulations with realistic-looking wrinkles. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 157.
- Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. 2019. PIFu: Pixel-Aligned Implicit Function for High-Resolution Clothed Human Digitization. In *Proceedings of the International Conference on Computer Vision (ICCV)*. IEEE.
- Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. 2020. PIFuHD: Multi-Level Pixel-Aligned Implicit Function for High-Resolution 3D Human Digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 84–93.
- Artsim Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Bjorn Ommer. 2018. A style-aware content loss for real-time hd style transfer. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 698–714.
- Igor Santesteban, Miguel A Otaduy, and Dan Casas. 2019. Learning-Based Animation of Clothing for Virtual Try-On. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 355–366.
- Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE Transactions on Neural Networks* 20, 1 (2008), 61–80.
- Bernhard Scholkopf and Alexander J Smola. 2001. *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.
- Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. 2006. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, Vol. 1. IEEE, 519–528.
- Andrew Selle, Jonathan Su, Geoffrey Irving, and Ronald Fedkiw. 2008. Robust high-resolution cloth using parallelism, history-based collisions, and accurate friction. *IEEE transactions on visualization and computer graphics* 15, 2 (2008), 339–350.
- Qingyang Tan, Lin Gao, Yu-Kun Lai, and Shihong Xia. 2018. Variational autoencoders for deforming 3d mesh models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5841–5850.
- Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. 2018. Bodynet: Volumetric inference of 3d human body shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 20–36.
- Daniel Vlastic, Ilya Baran, Wojciech Matusik, and Jovan Popović. 2008. Articulated mesh animation from multi-view silhouettes. In *ACM Transactions on Graphics (TOG)*, Vol. 27. ACM, 97.
- Huamin Wang, Florian Hecht, Ravi Ramamoorthi, and James F O'Brien. 2010. Example-based wrinkle synthesis for clothing animation. In *Acm Transactions on Graphics (TOG)*, Vol. 29. ACM, 107.
- Tuanfeng Y Wang, Duygu Ceylan, Jovan Popović, and Niloy J Mitra. 2018. Learning a shared shape space for multimodal garment design. In *SIGGRAPH Asia 2018 Technical Papers*. ACM, 203.
- Chenglei Wu, Kiran Varanasi, and Christian Theobalt. 2012. Full body performance capture under uncontrolled and varying illumination: A shading-based approach. In *European Conference on Computer Vision*. Springer, 757–770.
- Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. 2018. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics (ToG)* 37, 2 (2018), 27.
- Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhler. 2016. Estimation of human body shape in motion with wide clothing. In *European Conference on Computer Vision*. Springer, 439–454.
- Jinlong Yang, Jean-Sébastien Franco, Franck Hétroy-Wheeler, and Stefanie Wuhler. 2018. Analyzing clothing layer deformation statistics of 3d human motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 237–253.
- Tao Yu, Zerong Zheng, Yuan Zhong, Jianhui Zhao, Qionghai Dai, Gerard Pons-Moll, and Yebin Liu. 2019. SimulCap: Single-View Human Performance Capture with Cloth Simulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Chao Zhang, Sergi Pujades, Michael J Black, and Gerard Pons-Moll. 2017. Detailed, accurate, human shape estimation from clothed 3D scan sequences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4191–4200.