DEMEA: Deep Mesh Autoencoders for Non-Rigidly Deforming Objects
— Supplementary Material —

Edgar Tretschk[1]   Ayush Tewari[1]   Michael Zollhöfer[2]   Vladislav Golyanik[1]   Christian Theobalt[1]

[1]Max Planck Institute for Informatics, Saarland Informatics Campus   [2]Stanford University

In this supplementary material, we expand on several points from the main paper. In Sec. 1, we offer additional analysis of CoMA [6] on larger latent spaces in comparison to DEMEA and our baselines. Sec. 2 explains how we obtain meshes from networks trained with the graph loss (GL). Sec 3 compares different rotation computations for the embedded deformation layer (EDL). In Sec. 4, we describe how we apply temporal smoothing in latent space. Sec. 5 contains details about skinning template meshes to embedded graphs. Sec. 6 provides low-level details of our architecture. In Sec. 7, we describe how we normalize depth maps and meshes (for reconstruction from real depth data). In Sec. 8, we describe how to incorporate the embedded graph into the mesh hierarchy and show the employed mesh hierarchy for all four datasets. Finally, in Sec. 9, we give further details on the graph convolutions.

1. Additional Comparisons to CoMA

In addition to latent dimension 8, we also trained CoMA [6] on latent dimensions 32 and 128 on all four datasets. Since CoMA uses a batch size of 16, while we use a batch size of 8 for our method, we report two versions of CoMA: one that is matched in the number of epochs to our method and one that is matched in the number of iterations to our method. *I.e.*, the iteration-matched version is trained for twice the number of epochs as our method. Note that we report the iteration-matched numbers for CoMA in the main paper. See Table 1 for an expanded version of Table 3 from the main paper.

Again, the fully-connected baselines outperform the graph-convolutional networks for latent dimension 128. For latent dimension 8, DEMEA gives better quantitative results than CoMA on all datasets except for CoMA, which does not include large non-rigid deformations. For latent dimension 32, the results are more mixed: DEMEA, again, has better results on Dynamic FAUST and on Cloth, but performs on par on SynHand5M. On the CoMA dataset, DEMEA is slightly worse. While these numbers show that our approach compares favorably to CoMA for large non-rigid deformations, the advantage of our architecture which

includes the EDL is more evident qualitatively. We avoid many artifacts present in the results of CoMA and our baselines, see Fig. 4 in the main paper.

2. Graph Loss

To obtain mesh results from a network trained for the graph loss (GL), we need to apply embedded deformation at test time. Although the trained network predicts graph node positions t_l, it does not regress graph node rotations R_l. We compute the missing rotation for each graph node l as follows: assuming that each node's neighborhood transforms roughly rigidly, we solve a small Procrustes problem that computes the rigid rotation between the 1-ring neighborhoods of l in the template graph and in the regressed network output. We directly use this rotation as R_l.

3. Rotations in the EDL

In the main paper, the graph node rotations R_l are regressed by convolutional layers. Instead, practitioners might consider using the local Procrustes (LP) formulation of Sec. 2 *inside the network*. For this alternative, the network only needs to regress graph node positions t_l. From these, the graph node rotations can be inferred. The EDL can then take these graph node transformations to compute the deformed mesh, on which we put the reconstruction loss.

We extend Table 2 from the main paper by this experiment in Table 2. Note that we do not back-propagate through the rotation computation as we found this to lead to training instability. In addition to being faster, we find that regressing rotations is quantitatively better than LP.

4. Depth-to-Mesh Tracking

We can apply temporal smoothing to the reconstruction of a sequence of real depth images $\{\mathbf{D}_i\}_i$, by decoding a running (causal) exponential average of the latent vectors of this sequence. First, we encode the sequence into latent vectors $\{\mathcal{D}_i\}_i$. We then define a smoothed sequence of latent vectors $\{\mathcal{D}'_i\}_i$ as follows: let $\mathcal{D}'_0 = \mathcal{D}_0$ and set

	DFaust [2]			SynHand5M [5]			Cloth [1]			CoMA [6]		
	8	32	128	8	32	128	8	32	128	8	32	128
CoMA (matched epochs)	8.0	3.4	3.2	10.3	4.72	4.10	1.41	1.02	0.89	1.30	0.99	0.89
CoMA (matched iterations)	8.2	3.4	2.6	9.76	4.55	3.40	1.38	1.05	0.88	**1.25**	**0.93**	0.85
CA	6.7	3.0	2.6	10.30	**4.49**	3.76	1.61	0.90	0.72	1.57	0.97	0.87
MCA	8.6	3.4	2.5	9.33	4.55	3.67	1.75	0.82	0.70	1.61	0.99	0.87
Ours	**6.6**	**2.9**	2.4	**8.97**	4.67	3.53	**1.34**	0.83	0.71	1.49	1.05	0.94
FCA	9.3	3.4	**2.2**	20.96	7.22	**1.44**	1.71	0.71	**0.44**	3.19	1.39	**0.75**
FCED	8.2	3.1	**2.2**	20.49	9.13	1.60	1.89	**0.69**	**0.44**	3.61	3.19	1.08

Table 1. Average per-vertex errors on the test sets of DFaust (in cm), SynHand5M (in mm), textureless cloth (in cm) and CoMA (in mm.)

	FC IG	FC IM	GC IG	GC IM
GL	2.6	8.9	2.4	2.4
LP	2.3	2.6	2.4	2.9
Ours	2.2	2.3	2.3	2.4

Table 2. Evaluation of different settings of our network on the test set of DFaust [2] using the latent code of length 128. The numbers are the average vertex errors in cm.

	DFaust	SynHand5M	Cloth	CoMA
$\|\mathcal{N}_{\mathbf{v}_i}\|$	12	7	5	6
σ	0.05	0.009	0.015	0.01

Table 3. Skinning parameters for each dataset.

$\mathcal{D}'_i = \alpha \cdot \mathcal{D}_i + (1 - \alpha) \cdot \mathcal{D}'_{i-1}$ for $i > 0$ for some $\alpha \in [0, 1]$. The smoothed sequence of meshes $\{\mathbf{M}_i\}_i$ is obtained by decoding $\{\mathcal{D}'_i\}_i$.

5. Skinning

We compute the skinning weight $w_l(\mathbf{v}_i)$ of vertex $\mathbf{v}_i$ to one of its approximately closest skinning nodes $l \in \mathcal{N}_{\mathbf{v}_i}$ as:

$$w_l(\mathbf{v}_i) = \exp\left(\frac{-\|\boldsymbol{g}_l - \mathbf{v}_i\|^2}{2 \cdot \sigma^2}\right), \qquad (1)$$

where $\sigma \in \mathbb{R}$ depends on the dataset. Table 3 contains our choice of parameters for each dataset.

6. Architecture

Fig. 1 contains our low-level architecture. *GC(f)* is a Chebychev graph-convolutional layer with f output features. *DS* is a down-sampling layer and *US* is an up-sampling layer. *Conv2D(f,k,s)* is a 2D convolution with f output features, kernel size $k \times k$ and stride s. We modified ResNetV2 50 by removing its first convolutional layer and its final non-convolutional layers. We use ReLU non-linearities after every graph-convolutional, 2D convolutional and fully-connected layer except for the first 2D convolutional layer in the depth encoder and the last graph-convolutional layer. The third up-sampling module (*i.e.* up-sampling layer followed by a graph convolution) is only used for higher-resolution embedded graphs. All graph con-

volutions use $K = 6$, except for the last two, which use $K = 2$ for local refinement.

7. Normalization

Depth All depth-to-mesh networks rescale the depth values of the input depth map from between $0.3m$ and $7m$ to $[-1, 1]$.

Bodies: Depth For our depth-to-mesh network on bodies, we employ a number of additional normalization steps to focus on non-rigid reconstruction. First, we assume to be given a segmentation mask that filters out the background. The depth value of background pixels is set to 2. We crop the foreground tightly and use bilinear sampling to isotropically rescale the crop to 256×256. Given such a depth crop, we compute the average (foreground) depth value and subtract it from the input. Such normalization necessitates normalizing the network output, as we will describe next.

Bodies: Meshes We first normalize out the global translation from the meshes by subtracting from each mesh vertex the average vertex position. Since scale information is also lost, we fix the scale of the meshes by normalizing their approximate spine length. To that end, we compute the approximate spine length of the template mesh and of each mesh in the dataset. We then isotropically rescale all the meshes to the same spine length as the template mesh. The depth-to-mesh body reconstruction errors in the main paper are reported for these normalized meshes.

8. Mesh Hierarchy

We use the code of [6] to generate the mesh hierarchy. However, we need to preserve the nodes of the embedded graph throughout all hierarchy levels between the embedded graph and the full-resolution mesh. The code is based on QSlim [4], which uses quadric edge errors. For every edge, it computes the cost of removing one of the vertices from the mesh. For all levels that are at least as fine as the embedded graph, we simply set the cost of removing a node of the embedded graph to infinity.
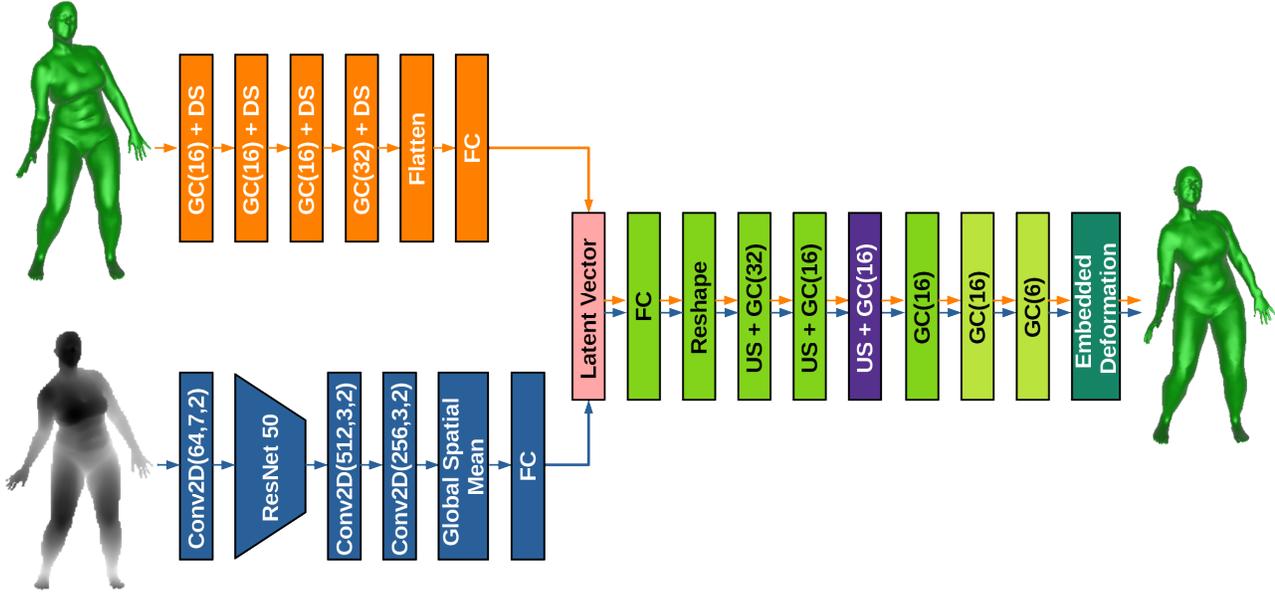
Figure 1. The low-level architecture of DEMEA (orange path) and the depth-to-mesh network (blue path). Note that the two paths are not trained simultaneously.
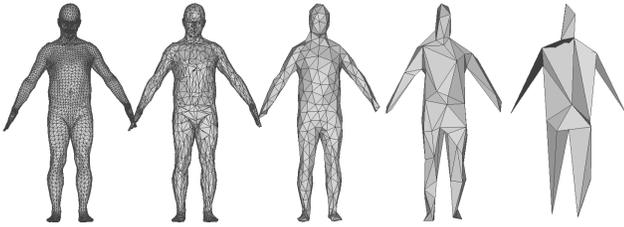
Figure 2. Body hierarchy.

Figure 3. Hand hierarchy.

Figure 4. Face hierarchy.
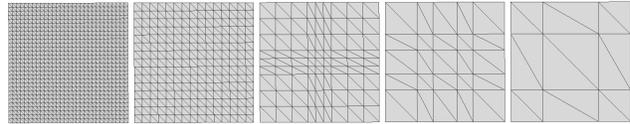
Figure 5. Cloth hierarchy.

Figures 2, 3, 4 and 5 visualize the five levels of the mesh hierarchy used for computing the barycentric up-sampling weights for the graph up-sampling layers.

9. Graph Convolution Details

Our graph encoder-decoder architecture is based on fast localized spectral filtering [3]. Given an F_{in}-channel feature tensor $\mathbf{x} \in \mathbb{R}^{N \times F_{in}}$, where the features are defined at the N graph nodes, and let $\mathbf{x}_i \in \mathbb{R}^N$ denote the i-th input graph feature map, we define the j-th output graph feature map $\mathbf{y}_j \in \mathbb{R}^N$ as follows:

$$\mathbf{y}_j = \sum_{i=1}^{F_{in}} g_{\theta_{i,j}}(\mathbf{L})\mathbf{x}_i \ . \tag{2}$$

Here, $\mathbf{L}$ is the Laplacian matrix of the graph and the filters $g_{\theta_{i,j}}(\mathbf{L})$ are parameterized using Chebyshev polynomials of order K.

Specifically,

$$g_{\theta_{i,j}}(\mathbf{L}) = \sum_{k=0}^{K-1} \theta_{i,j,k} T_k(\widetilde{\mathbf{L}}), \tag{3}$$

where $\theta_{i,j,k} \in \mathbb{R}$ and $\widetilde{\mathbf{L}} = 2\mathbf{L}/\lambda_{max} - \mathbf{I}$. The Chebychev polynomial T_k is defined as $T_k(x) = 2xT_{k-1}(x) - T_{k-2}(x)$, $T_1(x) = x$, and $T_0(x) = 1$.

This leads to K-localized filters that operate on the K-neighbourhoods of the nodes. The complete output feature tensor, that stacks all F_{out} feature maps, is denoted as $\mathbf{y} \in \mathbb{R}^{N \times F_{out}}$. Each filter $g_{\theta_{i,j}}(\mathbf{L})$ is parameterized by K coefficients, which in total leads to $F_{in} \times F_{out} \times K$ trainable parameters for each graph convolution layer, see [3] for more details.

References

[1] J. Bednařík, P. Fua, and M. Salzmann. Learning to reconstruct texture-less deformable surfaces. In *International Conference on 3D Vision (3DV)*, 2018. 2

[2] F. Bogo, J. Romero, G. Pons-Moll, and M. J. Black. Dynamic FAUST: Registering human bodies in motion. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. 2

[3] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *International Conference on Neural Information Processing Systems (NIPS)*, NIPS'16, pages 3844–3852, 2016. 3, 4

[4] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 209–216. ACM Press/Addison-Wesley Publishing Co., 1997. 2

[5] J. Malik, A. Elhayek, F. Nunnari, K. Varanasi, K. Tamaddon, A. Héloir, and D. Stricker. Deephps: End-to-end estimation of 3d hand pose and shape by learning from synthetic depth. *2018 International Conference on 3D Vision (3DV)*, pages 110–119, 2018. 2

[6] A. Ranjan, T. Bolkart, S. Sanyal, and M. J. Black. Generating 3D faces using convolutional mesh autoencoders. In *European Conference on Computer Vision (ECCV)*, pages 725–741, 2018. 1, 2