Neural 3D Morphable Models:
Spiral Convolutional Networks for 3D Shape Representation Learning and Generation
Supplementary Material

1. Additional Implementation Details

Padding in Spiral Convolutional Networks:

- **Spiral length:** As mentioned in the paper, one can choose an arbitrary length L and then adjust the spirals on the mesh so that they all share the same size. In this work, the length was chosen to be equal to $\mu + 2\sigma$, where μ and σ the mean and standard deviation of the spiral lengths corresponding to the given number of hops h across the entire mesh. Shorter spirals were padded with a dummy vertex $u \notin V$, such that $f(u) = \mathbf{0}$, where f is the signal on the graph, while larger spirals were truncated in order to match the chosen length.

- **Boundary conditions:**The spiral on the boundaries of the mesh were similarly padded with u. More specifically, if a ring R^k had a disconnect, i.e. edge $(R_i^k, R_{i+1}^k) \notin E$, a dummy vertex u was inserted in the spiral ordering between elements R_i^k and R_{i+1}^k. In this way, the spiral stays consistent across boundary vertices as well.

Spiral Convolutional GAN architecture: For the Spiral Convolutional GAN and its linear 3D Morphable Model counterpart we used a latent space dimension $d = 256$ (99.4% PCA explained variance). The Generator and Critic (Discriminator) networks had the following architectures (we use the same notation as in Section 4.2 – Implementation Details of the main paper):

$Disc: SC(1, 64) \rightarrow DS(4) \rightarrow SC(1, 128) \rightarrow DS(4) \rightarrow SC(1, 128) \rightarrow DS(4) \rightarrow SC(1, 128) \rightarrow DS(4) \rightarrow SC(1, 256) \rightarrow DS(4) \rightarrow FC(1)$

$Gen: FC(l * 256) \rightarrow US(4) \rightarrow SC(1, 128) \rightarrow US(4) \rightarrow SC(1, 128) \rightarrow US(4) \rightarrow SC(1, 128) \rightarrow US(4) \rightarrow SC(1, 64) \rightarrow US(4) \rightarrow SC(1, 3)$

2. Generalization Error Tables

In the Tables 1, 3, and 2, we report the exact results and parameter counts of the methods of Fig. 5 of the main paper.

Latent Size	Explained Variance	Model	# of Params	Generalization (mm)
8	83.1 %	PCA	120k	1.636
	n/a	COMA	28k	0.885
	n/a	Neural3DMM (small)	38k	0.801
	n/a	Neural3DMM (ours)	381k	**0.472**
16	94.6 %	PCA	241k	0.825
	n/a	COMA	39k	0.751
	n/a	Neural3DMM (small)	48k	0.635
	n/a	Neural3DMM (ours)	425k	**0.377**
64	99.1 %	PCA	965k	0.284
	n/a	COMA	100k	0.611
	n/a	Neural3DMM (small)	113k	0.449
	n/a	Neural3DMM (ours)	682k	**0.260**

Table 1: COMA dataset comparison

Latent Size	Explained Variance	Model	# of Params	Generalization (mm)
8	84.8 %	PCA	165k	59.30
	n/a	COMA	32k	28.09
	n/a	Neural3DMM (small)	41k	28.69
	n/a	Neural3DMM (ours)	274k	**19,77**
16	96.1 %	PCA	330k	32.16
	n/a	COMA	46k	17.03
	n/a	Neural3DMM (small)	56k	15.30
	n/a	Neural3DMM (ours)	332k	**11.20**
64	99.8 %	PCA	1.32M	5.28
	n/a	COMA	129k	8.98
	n/a	Neural3DMM (small)	142k	5.51
	n/a	Neural3DMM (ours)	676k	**4.29**

Table 2: DFAUST dataset comparison

Latent Size	Explained Variance	Model	# of Params	Generalization (mm)
16	86.0 %	PCA	1.36M	0.739
	n/a	COMA	53k	0.812
	n/a	Neural3DMM (small)	66k	0.718
	n/a	Neural3DMM (ours)	320k	**0.711**
32	93.0 %	PCA	2.79M	0.525
	n/a	COMA	82k	0.616
	n/a	Neural3DMM (small)	95k	0.518
	n/a	Neural3DMM (ours)	438k	**0.502**
128	98.5 %	PCA	10.91M	0.235
	n/a	COMA	254k	0.400
	n/a	Neural3DMM (small)	274k	0.269
	n/a	Neural3DMM (ours)	1.15M	**0.229**

Table 3: Mein3D dataset comparison

3. Additional Shape Analogies

In Fig. 1 and 2, we show additional shape analogies, similar to Fig. 9 of the main paper. In particular, we show body pose and facial expression transfer in the DFAUST and COMA datasets respectively.

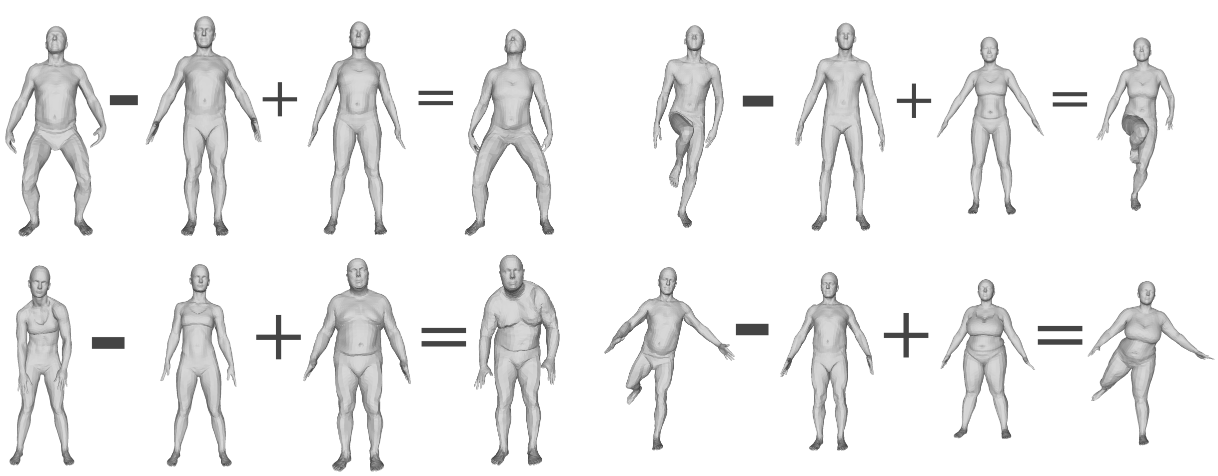

Figure 1: Pose transfer examples through latent space analogies in the DFAUST dataset

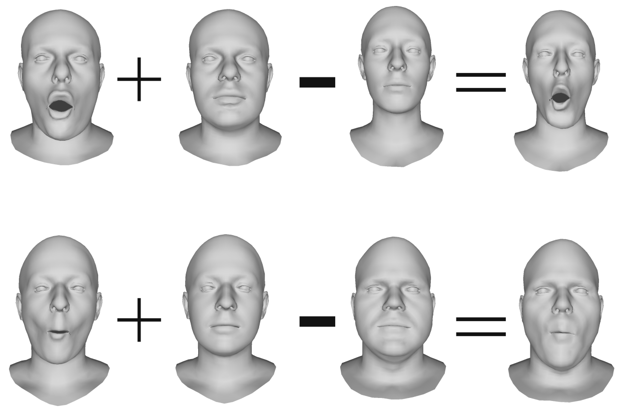

Figure 2: Expression transfer examples through latent space analogies in the COMA dataset

4. Synthetic faces generated by the 3DMM

In order to allow a direct comparison of the Spiral Convolutional GAN with the 3DMM, we included several synthetic faces generated by the 3DMM in Fig. 3. As already mentioned in the main paper, although PCA can produce very smooth and noiseless surfaces, the synthetic faces look artificial, due to the absence of high frequency detail.

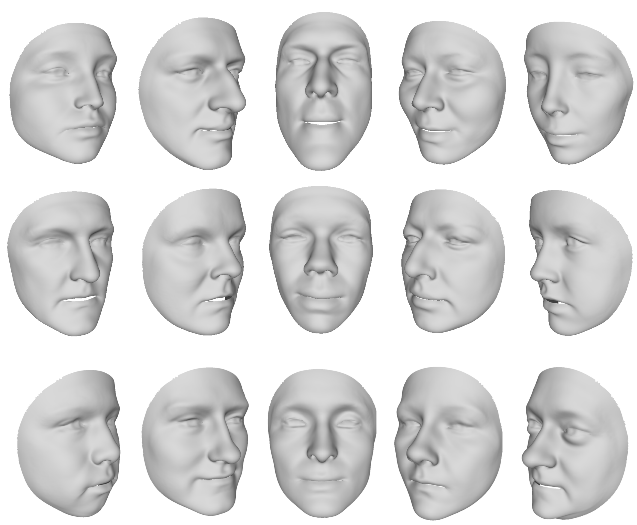

Figure 3: Faces sampled from the PCA-based 3DMM