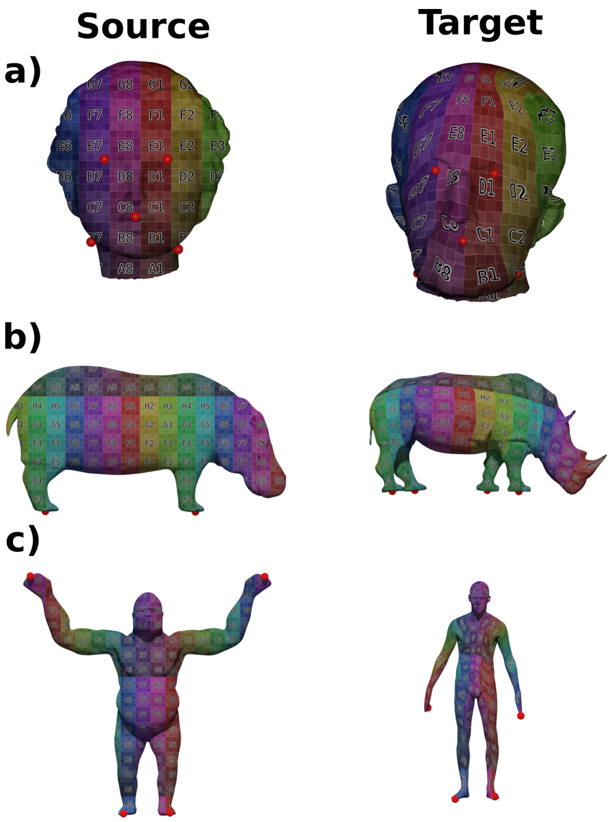

Figure A1. Surface to Surface maps. Red spheres depicts landmarks.

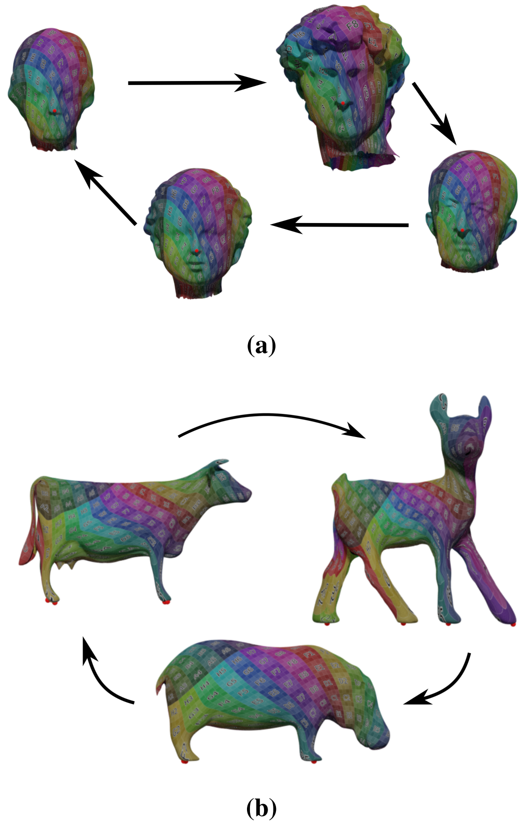

Figure A2. Cycle-consistent maps over a collection. We minimize the isometric distortion of the maps between all pairs of shapes. Landmarks are visualized as red sphere.

A. Supplementary material

A.1. Surface-to-surface Maps

As described in the paper, we can represent and optimize surface-to-surface maps by composing two neural atlases ϕ and ψ with h that maps one atlas to the other.

Fig. A1 depicts more examples of surface maps. In all cases we use landmarks to guide the distortion to the correct minimum.

A.2. Cycle-consistent Mapping for Collections of Surfaces

As discussed in the paper, thanks to our formulation it is possible to optimize all maps between a collection of shapes. As previously illustrated, a collection of neural maps is inherently cycle-consistent, and we can, for the first time, optimize the distortion of all maps in the collection. Fig. A2 depicts surface maps over a collection of heads (a) and a 4-legged models (b). In both cases we employ landmarks highlighted as red spheres.

A.3. Approximating Surfaces

As described in the main paper, a *neural map* ψ can approximate a given surface S by approximating the atlas $f : \mathbb{R}^2 \to S$ accurately. To further demonstrate the effectiveness of our formulation, in Fig. A6 we show more overfitted models for categories such as Stanford Bunny (a), hand (b), heads (c to d), busts (e and f), animals (g and h), and human (i and j). Each neural surface map consists of a ten-layer residual fully-connected network with Softplus activation.

In Fig. A3 we compare different activation functions such as LeakyRelu and Relu. These non-smooth functions introduce artefacts, such as unwanted wrinkles Fig. A3 (b and c). Overlooking this behavior might bear negative effects in map composition as these introduced details can be amplified or inject distortion in the final map. The use Softplus alleviates these artefacts, but biases the map towards smooth surfaces, hence the hairs in Fig. A6(d and f), will have minor discrepancies from the ground truth, as some areas will be smoothed out.

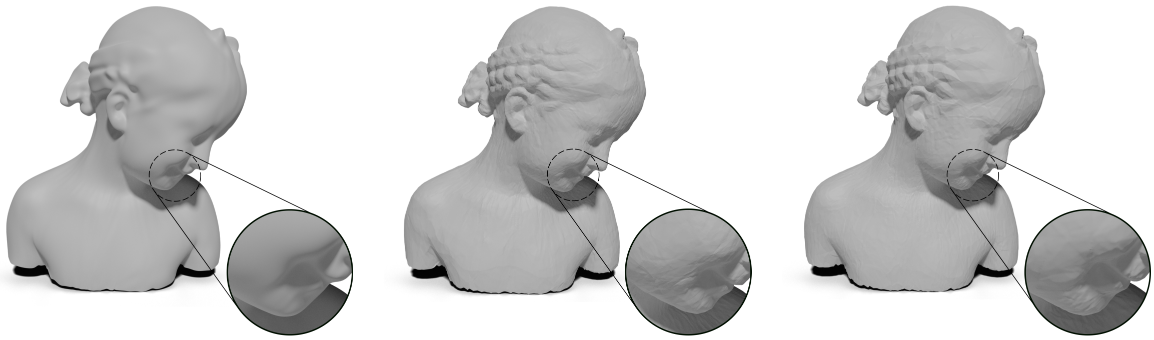

(a) **(b)** **(c)**

Figure A3. A MLP-based neural surface map with different activation functions. Softplus (a) maps are smooth. LeakyRelu (b) and Relu (c) are more oscillatory.

A.4. Surface Parametrization

We evaluate the efficacy of our method in optimizing different energies, by optimizing the distortion of the map $\phi \circ h$ as a function of h. We show a parameterization minimizing conformal-distortion in Fig. A4(b), reducing the conformal distortion from 1.74 to 1.52. Similarly, we show a parametrization minimizing isometric distortion in Fig. A4(c) reducing the median symmetric dirichlet from 11.58 to 9.92.

A.5. Composition with Analytical Maps

As discussed in the paper, we can directly compose a neural map with analytical functions to get surface maps into analytical surfaces. Fig. A5 illustrates several additional analytic surfaces we can map into. The neural map is optimized to minimize the conformal energy.

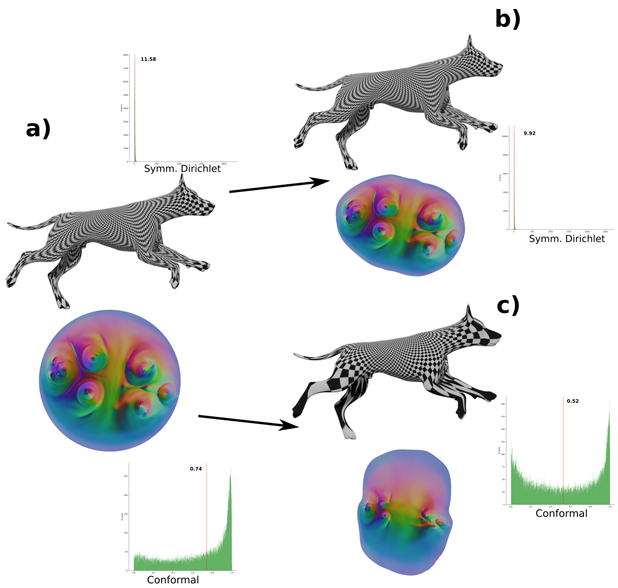

Figure A4. Parametrization with different energies. Starting from a neural surface map (a), we optimize for conformal distortion (b) and isometric distortion boundary free (c).

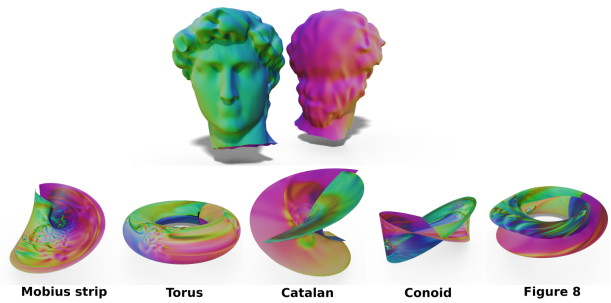

Figure A5. Surface mapping between David and a variety of parametric surfaces. Shape coloring is based on source model normals, *i.e.*, David.

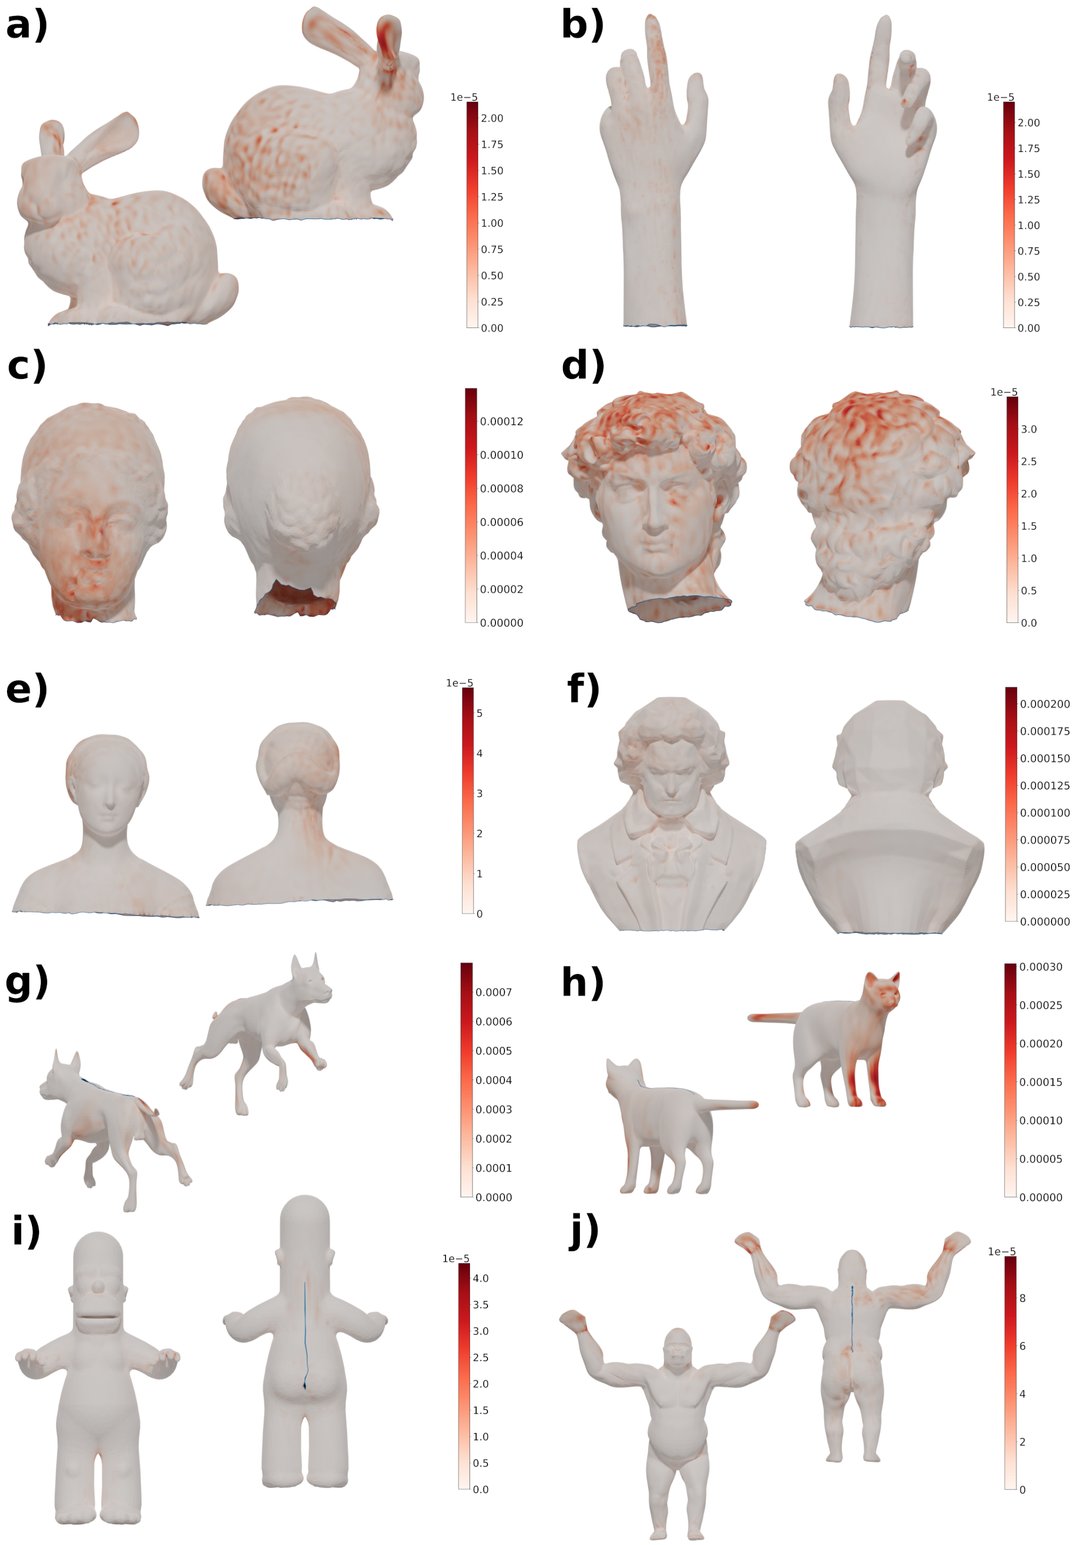

Figure A6. Neural surface maps depicted are color-coded to highlight the error compared to the ground truth mesh. In general, the produced models have low deviation from the original surface.