Deep Parametric Surfaces for 3D Outfit Reconstruction from Single View Image

Hugo Bertiche, Meysam Madadi, and Sergio Escalera

hugo_bertiche@hotmail.com

Universitat de Barcelona

CVC

HuPBA

3D Garment Regression from Images

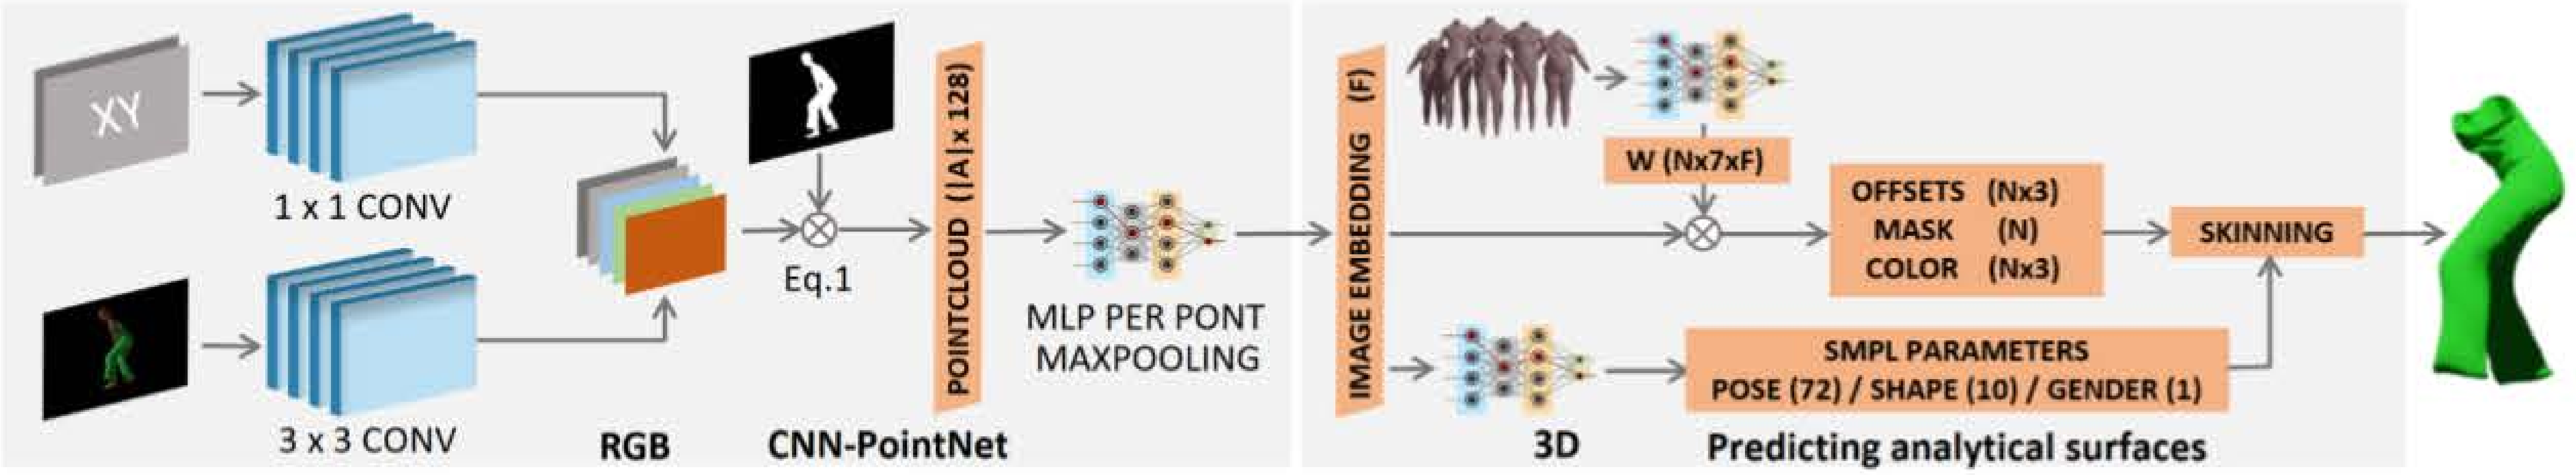

RGBA images are fed to our CNN to obtain an image embedding that encodes 3D information.
We propose a lightweight CNN based on pixel positional encoding and leveraging the image alpha channel as an attention mask.
In our experiments, we show how this architecture provides of a much more compact networks than classical CNNs and faster and better performance.

We propose a novel architecture to predict analytical surfaces instead of point clouds. Traditional 2D-to-3D works usually follow the CNN with a MLP to regress a set of 3D points. We predict functions that map points in the body surface to points in the garment surface.
Then, our predictions consist on surfaces parametrized with the body surface. Since the body is a continuous 2-manifold, so it is the predicted garment.

Parametric Space

We use SMPL body surface as the parametric space for predicted 2-manifolds. Due to shape variability, 3D points are ambiguous. We additionally include surface normals and blend weights to solve this ambiguity.

Skirts and dresses do not follow body topology, for this reason, we propose an alternative body model with the correct topology for such garments. SMPL legs are replaced with a skirt-like shape. Blend shapes are transferred to allow consistent re-shaping along the body.

Differential Geometry: Loss

During training, we sample uniformly from the parametric space to generate point clouds as predictions. Similarly, we sample uniformly from the ground truth garment surface. This allows applying Chamfer loss to predicted point clouds against ground truth point clouds.

$$\mathcal{L}_{CD} = \sum_{\mathbf{y}'_i \in \mathbf{Y}'_M} \min_{\mathbf{y}_j \in \mathbf{Y}} \|\mathbf{y}'_i - \mathbf{y}_j\|_2^2 + \sum_{\mathbf{y}_j \in \mathbf{Y}} \min_{\mathbf{y}'_i \in \mathbf{Y}'_M} \|\mathbf{y}'_i - \mathbf{y}_j\|_2^2$$

$$\mathcal{L}_{smooth} = \sum_{\mathbf{p}}^{\mathbf{P}} \frac{\partial^2 \mathbf{y}'_{\mathbf{U}}}{\partial \omega^2}(\mathbf{p})$$

Chamfer loss is sensitive to wrong matchings, specially for loose garments such as dresses. This yields noisy predictions. To address this issue, we leverage the continuous nature of our predictions to design a novel loss term.
Minimizing the second derivative of predictions w.r.t. the parametric space surface avoids sudden changes in shape, producing smoother surfaces.

Qualitative Results

We display here qualitative results. For each, image (left) and regressed garment (right). As observed, our approach can be adapted to predict a color for the parametric surfaces.
For this work, we used CLOTH3D++ dataset, a huge collection of video sequences paired with 3D data, texturing, lighting. Link: https://chalearnlap.cvc.uab.cat/dataset/38/description/