GNPM: Geometric-Aware Neural Parametric Models

Mirgahney Mohamed University College London Lourdes Agapito University College London



Figure 1. **Geometric-Aware Neural Parametric Model:** Given a set of input 3D point-clouds of different identities in a variety of poses, GNPM disentangles shape and deformations by mapping each pose to its canonical t-pose and learning dense temporally consistent correspondences, without the need for any ground-truth annotations.

Abstract

We propose Geometric Neural Parametric Models (GNPM), a learned parametric model that takes into account the local structure of data to learn disentangled shape and pose latent spaces of 4D dynamics, using a geometricaware architecture on point clouds. Temporally consistent 3D deformations are estimated without the need for dense correspondences at training time, by exploiting cycle consistency. Besides its ability to learn dense correspondences, GNPMs also enable latent-space manipulations such as interpolation and shape/pose transfer. We evaluate GNPMs on various datasets of clothed humans, and show that it achieves comparable performance to state of the art methods that require dense correspondences during training.

1. Introduction

Reconstructing temporally consistent shape deformations from visual inputs remains a challenging open problem in computer vision with many applications in AR/VR and content creation. While model-agnostic approaches that exploit local shape smoothness priors [17] have been extremely successful, the remarkable recent advances in machine learning driven capture of domain-specific 3D parametric models such as for human bodies [1, 8, 13]), hands [25], faces [11, 23, 24] or animals [31] have made them an attractive alternative. However, parametric models suffer from two important drawbacks. First, since they do not exploit local geometry priors they often fail to capture local geometry details and have a tendency to over-regularize. Secondly, their construction typically requires manual intervention to aid alignment, or fully dense annotations across all samples to provide dense correspondences at training time.

Neural Parametric Models (NPMs) [20] were recently proposed to tackle some of the challenges faced by traditional parametric models by leveraging deep neural networks and implicit functions to learn a disentangled representation of shape and pose. NPMs are learnt from data alone without the need for any class specific knowledge and, once trained, can be fit to new observations via test-time optimization. While NPMs offer an appealing alternative to traditional 3D parametric models, since they only require the same identity to be seen in different poses (including a canonical pose) dropping the need for registration across different identities, they still require known dense correspondences during training between different poses of the same identity and their canonical pose. Moreover, NPMs do not learn any local geometry regularity priors given that flow predictions are conditioned on global shape and pose codes and local geometry priors

are not exploited.

To tackle both limitations we propose Geometric Neural Parametric Models (GNPMs). We represent surfaces as point clouds and exploit the ability of dynamic graph neural network architectures [28] to explicitly take into account the local structure around deformed points to learn local features and enforce local geometric regularity. In addition, we relax the need for any correspondences at all during training by disentangling 4D dynamics into shape and pose latent spaces via a cycle consistency loss.

Geometric priors are enforced by adapting the Edge-Conv [28] graph convolution operators to design our model as an auto-decoder [21]. Our intuition is that points tend to deform coherently with their neighbours [27]. Furthermore, local features learnt via graph convolutions can also help the model to learn semantic relations between non-neighbouring points, which is potentially useful in deformations such as dancing where distant body parts move synchronously. Inspired by [22], we show that the learned features can be additionally used to segment shapes into semantically meaningful parts which remain consistent across identities, in a completely unsupervised way (see Figure 7). While it is known that geometric-based methods can be inefficient [29], we provide an efficient implementation that allows our model to run at a comparable cost to MLP-based models.

To learn without known correspondences we exploit the observation that transformations between posed and canonical spaces should be cycle consistent. This loss allows us to infer a dense deformation field, by jointly learning the weights of two networks that perform a bi-directional mapping between posed and canonical spaces and the respective shape and pose latent spaces (see Fig. 2). To fit the model to new unseen identities and/or deformations, we use testtime optimisation to minimise the cycle consistency loss to recover shape and pose latent vectors.

In summary, our contributions are:

- GNPM is a geometric-aware neural parametric model that learns to disentangle pose and shape exploiting the local structure of the data via edge convolutions.
- GNPM learns shape and pose embeddings and longterm dense correspondences without the need for ground truth annotations during training.
- GNPMs learn rich geometric features useful for downstream tasks such as unsupervised part segmentation.

2. Related Work

Parametric Models: Parametric 3D models have become a prevalent tool to model deformable 3D shapes. They learn to disentangle deformations into several factors of interest and have been applied to various domains such as human bodies [1,8,13,30], hands and faces [25] [11,23,24] and animals [31]. Despite their success, they struggle to capture fine-grained details like wrinkles or to model clothes. Also, their construction can be tedious as it often requires domain knowledge or manual tuning. Neural based methods [20] offer a compelling alternative to learn directly from data without the need for manual tuning or domain-specific knowledge.

Supervised Neural Deformation Models: Building on 3D OccNet [15], OFlow [19] learns a continuous spatiotemporal representation of 4D dynamics that assigns motion vectors to every location in space-time. However, it degrades when capturing long sequences. Inspired by Dynamic Fusion [17], Bozic et al. [4] learn a globally consistent deformation graph while learning dense surface details via local MLPs. However, it is limited to be sequence-specific and cannot be used for shape or pose transfer. Palafox et al. [20] recently introduced Neural Parametric Models (NPMs) leveraging the representation power of implicit functions to disentangle latent spaces of shape and pose. However, NPMs disregard the local geometric structure of 3D shapes and require dense correspondences for training. While our geometric model also learns to disentangle between shape and pose, unlike [20] we take into account the geometric structure and can learn long term correspondences in a self-supervised manner without the need for dense ground truth annotations.

Self-Supervised Neural Deformation Models: LoopReg [2] was the first end-to-end learning framework to solve scan registration with a self-supervision loop. Backward and forward maps are learnt to predict correspondences between every input scan point and the model surface. SCANimate [26] also uses a self-supervision cycle to learn an implicit dense field of skinning weights to map surface points to a canonical pose. Although both [2,26] can model shape deformation in a semi-supervised or unsupervised way, they do not learn latent shape/deformation spaces and rely on SMPL as the underlying body model. In contrast, our model, GNPM, learns disentangled shape and deformation latent spaces. More recently, Neuromorph [6] jointly solves shape interpolation and dense correspondences in an unsupervised way using edge convolutions [28] and a single forward pass. Although Neuromorph can estimate dense correspondences between shapes of different categories in different poses the latent interpolation is modelled via time, limiting its ability to learn a parameterised representation.

3. Method

We introduce Geometric Neural Parametric Models (GNPM), a geometric-aware model that disentangles 4D dynamics into latent spaces of shape identity and pose and can be learnt without the need for correspondences across shapes.

We choose point clouds as our shape representation given their lightweight nature and that they naturally match the raw



Figure 2. **Self-supervision cycle.** Our model is trained to map an input scan to their canonical t-pose and then back to the deformed shape, imposing cycle consistency. In this way, we learn dense correspondences without any ground truth annotations. In addition, we disentangle shape and pose by learning two embeddings which are jointly optimised with the weights of the graph neural networks.

output of commodity depth cameras. Moreover, when combined with the EdgeConv architecture [28], local geometric structure can be exploited to capture both local and global shape properties, unlike the more recently popular implicit representations [20]. We adopt EdgeConv layers (see 3.1.1) and modify the architecture to an auto-decoder [21].

Given a dataset of multiple shape identities in various poses, but without any registration within or across identities, we jointly learn: (*i*) the weights of a network that predicts globally consistent dense point correspondences, and (*ii*) shape and pose latent embeddings. Instead of requiring persequence dense correspondences as supervision signal [20],

we use a self-supervision cycle, in similar spirit to [2,26], to learn the shape/pose latent spaces without any need for registration during training. Given an input posed shape, the forward network learns to deform it to its canonical t-pose. The backward network then learns to deform the t-posed shape back to the input posed shape. Together, forward and backward networks form an identity mapping, and the distance between predicted and input shapes is used as supervision signal. Our only requirement on the training dataset is that each shape identity should be observed in a canonical pose (e.g., T-pose), which is satisfied on a variety of datasets [12, 14, 25]).

At test time, given new observations of an unseen identity, the weights of the network are frozen and shape and pose embedding vectors are jointly optimised by minimizing the same cycle-consistency loss.

3.1. Shape/Pose Network and Embeddings

To circumvent the lack of correspondences we use a selfsupervision cycle that conducts a bi-directional mapping between input and canonical poses. Our model is composed of two networks, both conditioned on the shape and pose latents (see Section 2). Given an input posed shape, the forward network predicts a dense deformation field $\delta \tilde{x}$ that maps it onto the canonical t-pose while the backward network learns to deform the canonical shape back to the input pose.

Figure 2 shows an overview of the self-supervision cycle. Forward and backward networks are implemented as autodecoders, with parameters θ_a and θ_b , and EdgeConv layers that predict dense deformation fields.

Forward Network: The forward network learns a geometric-aware deformation field that deforms the input posed shape to its canonical t-pose shape. More formally, the input to the network is a set of N observed points on the surface of the shape associated with the f^{th} frame and c^{th} identity, $X^{cf} = \{x_i^{cf}\}_{i=1}^N \in \mathbb{R}^3$ to which we apply positional encoding $\gamma(\cdot)$ [16]. The forward network forms a dynamic graph \mathcal{G}^l , via k-NN, conditional on a learnable per-frame D_p -dimensional latent pose code p_f , and a fixed D_s -dimensional latent shape code s_c . A series of Edge-Conv layers are applied, each using the output features of the previous layer to re-estimate the graph \mathcal{G}^l (see 3.1.1). Perpoint pooling is performed after the last layer and a shallow MLP is applied to predict dense point-wise deformations $\{\delta x_i^{cf}\}_{i=1}^N \in \mathbb{R}^3.$

$$f^{a}_{\theta_{a}} : \mathbb{R}^{N \times 51} \times \mathbb{R}^{D_{s}} \times \mathbb{R}^{D_{p}} \to \mathbb{R}^{N \times 3}$$
$$f^{a}_{\theta_{a}}(\gamma(x_{i}), s_{c}, p_{f}) = \delta \tilde{x}.$$
(1)

Backward Network: The backward network learns the inverse dense deformation field that deforms the canonical t-pose shape to the posed shape: $\{\delta y_i^{cf}\}_{i=1}^N \in \mathbb{R}^3$. While the architecture is equivalent to the forward network, the parameters are not shared.

$$f^{b}_{\theta_{b}} : \mathbb{R}^{N \times 51} \times \mathbb{R}^{D_{s}} \times \mathbb{R}^{D_{p}} \to \mathbb{R}^{N \times 3}$$
$$f^{b}_{\theta_{b}}(\gamma(\tilde{y}_{i}), s_{c}, p_{f}) = \delta \tilde{y}$$
(2)

Losses: Combining (1) and (2) we can close the cycle and have a self-supervision loop. We define the loss \mathcal{L}_{loop} as the l_1 distance between the shape predicted by the cycle and the input shape.

$$f^{a}_{\theta_{a}}(\gamma(x_{i}), s_{c}, p_{f}) = \delta \tilde{x_{i}}$$

$$\tilde{y_{i}} = x_{i} + \delta \tilde{x_{i}}$$

$$f^{b}_{\theta_{b}}(\gamma(\tilde{y_{i}}), s_{c}, p_{f}) = \delta \tilde{y_{i}}$$

$$\tilde{x_{i}} = \tilde{y_{i}} + \delta \tilde{y_{i}}$$
(3)

$$\mathcal{L}_{loop}(\tilde{x}_i^{cf}, x_i^{cf}) = \parallel \tilde{x}_i^{cf} - x_i^{cf} \parallel$$
(4)

To prevent the network from learning an identity mapping by setting $\delta \tilde{x}$ and $\delta \tilde{y}$ to zero, we use an additional symmetric ICP loss that minimizes the distance between the ground



Figure 3. Results of time-consistent dense correspondence estimation for 3D point-cloud inputs from the CAPE [14] (top) and DFAUST [3] (bottom) datasets. We show comparisons with the ground truth dense correspondence maps.

truth t-pose shape points and the t-pose predicted by the forward network

$$\mathcal{L}_{icp}(\tilde{y}_i^{cf}) = \parallel \tilde{y}_i^{cf} - NN_{\mathcal{R}}(\tilde{y}_i^{cf}) \parallel_2^2 \tag{5}$$

Where $NN_{\mathcal{R}}(.)$ is a function that queries the nearest neighbour of a 3D point in the set \mathcal{R} of input points. Which we found important to prevent correspondence flipping in the presence of extreme motion (see 4.6).

To aid temporal smoothness we constrain the current and next pose latents to be close by adding an l_1 loss between them \mathcal{L}_{lt} . Also we found it very useful to enforce temporal regularization for both networks output which we called spacial temporal loss \mathcal{L}_{st} .

$$\mathcal{L}_{lt}(p_f, p_{f+1}) = \parallel p_f - p_{f+1} \parallel$$
(6)

Inspired by the intuition points in consecutive frames deform coherently [27]. And as we do not have access to dense correspondences, we use the input points and query the nearest neighbour of the next frame input points. Then we use l_1 loss between the current network prediction and the prediction of the nearest points in the next frame.

$$\mathcal{L}_{st}^{a}(\tilde{x}_{i}^{cf}) = \| \delta \tilde{x}_{i}^{cf} - \delta \tilde{x}_{i}^{c(f+1)} \| \\
\delta \tilde{x}_{i}^{c(f+1)} = f_{\theta_{a}}^{a}(NN_{\mathcal{Q}^{a}}(x_{i}^{c(f+1)})) \\
\mathcal{L}_{st}^{b}(\tilde{y}_{i}^{cf}) = \| \delta \tilde{y}_{i}^{cf} - \delta \tilde{y}_{i}^{c(f+1)} \| \\
\delta \tilde{y}_{i}^{c(f+1)} = f_{\theta_{a}}^{b}(NN_{\mathcal{Q}^{b}}(y_{i}^{c(f+1)}))$$
(7)

Where $NN_Q(.)$ is a function that queries the nearest neighbour of a 3D point in the set Q of input points. Our final temporal regularization loss is:

$$\mathcal{L}_{temp} = \mathcal{L}_{lt} + \mathcal{L}_{st}^a + \mathcal{L}_{st}^b \tag{8}$$

Finally, we minimize this objective over all possible F deformation fields across all shape identities C with respect to the individual pose and shape codes $\{p_f\}_{f=1}^F$ and $\{s_c\}_{c=1}^C$ respectively and the forward and backward network weights θ_a, θ_b . To enforce a compact pose and shape manifold we also regularize the both codes with σ_p and σ_s .

$$\underset{\theta_{a},\theta_{b},\{s_{c}\}_{c=1}^{C},\{p_{f}\}_{f=1}^{F}}{\operatorname{arg\,min}} \sum_{c=1}^{C} \sum_{f=1}^{F} \sum_{i=1}^{N} \mathcal{L}_{loop} + \lambda_{icp} \mathcal{L}_{icp} + \lambda_{icp} \mathcal{L}_{icp} + \lambda_{temp} \mathcal{L}_{temp} + \frac{\|p_{f}\|_{2}^{2}}{\sigma_{p}^{2}} + \frac{\|s_{c}\|_{2}^{2}}{\sigma_{s}^{2}} \quad (9)$$

3.1.1 EdgeConv Layers

In contrast to previous neural parametric models [20] that treat each data point independently, we exploit the power of edge convolutions in DGCNNs [28] to learn local geometric structure in deformed 3D point clouds. EdgeConv [28] alleviates the lack of topology in point clouds by proposing a differential neural network module suitable for CNNbased high-level tasks. Local geometric structure is exploited by constructing a local neighbourhood graph and applying convolution-like operations on the edges in a dynamic setting, where the connectivity is learned and changes throughout the training. The learned feature space not only captures semantic relations within a local neighbourhood, but also between distant points, which is particularly useful in the context of deformations where different body parts might move synchronously.

Given a set of N uniformly sampled 3D surface points, denoted by $X = \{x_1, ..., x_N\} \subseteq \mathbb{R}^D$, where D is the point cloud dimensionality, a directed graph G = (V, E) of vertices V and edges E, is first built to represent the local point cloud structure via k-nearest neighbours (k-NN). In our setting, input points are represented by their 3D coordinates $x_i = (x_i, y_i, z_i)$, but D more generally refers to the dimensionality at each layer. Edge features are then defined as $e_{ij} = h_{\Theta}(x_i, x_j)$, where $h_{\Theta} : \mathbb{R}^D \times \mathbb{R}^D \to \mathbb{R}^{D'}$ is a multi-layer perceptron (MLP) with learnable parameters Θ . Channel-wise symmetric aggregation \Box (e.g., \sum or max) is then applied on the edge features associated with each vertex *i* with its output given by

$$x'_{i} = \underset{j:(i,j)\in\mathcal{E}}{\Box} h_{\Theta}(x_{i}, x_{j}).$$
(10)

Unlike GCNNs that work on a fixed input graph, Edge-Conv recomputes the graph neighbourhood structure at each layer. In practice, a pairwise distance matrix is computed in feature space, and the closest k points are selected for each vertex point at each layer.

3.2. Test Optimization

Once the network weights and shape and pose latent embeddings have been learnt, they can be used to fit a new sequence to the model by optimizing for the per-identity shape and per-frame pose codes that best explain the observations. We use the pose network and minimize the cycle consistency training objective with respect to the shape and pose latent codes $\{s_c\}_{c=1}^S, \{p_l\}_{f=1}^F$, keeping network weights fixed:

$$\underset{\{s_c\}_{c=1}^C, \{p_f\}_{f=1}^F}{\arg\min} \sum_{c=1}^C \sum_{f=1}^F \sum_{i=1}^N \mathcal{L}_{loop} + \frac{\|p_f\|_2^2}{\sigma_p^2} + \frac{\|s_c\|_2^2}{\sigma_s^2}$$
(11)

The shape and pose codes are initialized to the mean of the respective embedding. When dealing with a sequence, the pose code for each frame can be initialized with the result for the previous frame.

3.3. Implementation Details

Forward and Backward Networks: For both forward and backward networks, we stack three EdgeConv layers with two MLPs to compute edge features. LeakyReLU with 0.2 negative slope was used as the activation function and max as the channel-wise aggregation function for all layers. The graph connectivity is updated for each layer, based on the learned features in previous ones. Features from each EdgeConv are concatenated and passed through the MLPs for the final prediction. The dimensionality of the EdgeConv features is set to D = 128, except for the last layer where D = 256. Shape and pose latent dimensionalities are set to 128. We use the Adam optimizer [10] with learning rates of 1e - 3 for the shape/pose model and latent codes. In addition, we apply a learning rate scheduler with a decay

factor of 0.5 every 30 epochs. And use a cosine annealing our ICP loss weight λ_{icp} with $1e^{-1}$ initial weight and $1e^{-2}$ minimum weight. As for the regularization losses we use $5e^{-2}$ for temporal loss weight λ_{temp} . We use shape and pose code regularization weights $\sigma_s = \sigma_p = 1e - 4$. Despite the symbolic k-NN implementation allowing more point samples, we sample 1024 points per batch from each shape as input to the model. The latent codes are initialized randomly from a normal distribution. We use the positional encoding proposed in [16] to encode query points with 8 frequency bands.

Efficient k-NN Implementation: We use the KeyOps [5,7] python package to symbolically implement the k-NN algorithm. In the supplementary material we provide a comparison with a naive dense k-NN implementation and show a tenfold improvement in terms of training time.

4. Experimental Evaluation

Datasets: We evaluate our model on real-world scans from the CAPE [14] and D-FAUST [3] datasets, which provide real clothed humans and their corresponding SMPL+D registration. We train on 31 different posed shapes from 35 different distinct identities [14] and test on a total of 4 test sequences from 4 unseen characters. We also learn a hand model from the MANO [25] dataset.

Evaluation metrics: We measure reconstruction and 4D tracking performance following the established per-frame evaluation protocol [15, 20]. The l_2 Chamfer distance (C- l_2) offers a measure combining the accuracy and completeness of the reconstructed surface. We also use End-Point Error (EPE), which measures keyframe-to-frame deformation l_2 distance between predicted and ground truth as in [4]. We visualise corresponding points with the same colour.

4.1. Evaluation on CAPE Dataset

Reconstruction and unsupervised dense correspondences: We compare our model to the state of the art on the CAPE [14] dataset which provides ground truth dense correspondences. We compare with NPM [20] (the most related approach to ours), although it requires dense correspondences at training time, unlike our approach.

Figure 3 shows how our method is capable of reconstructing the deformed shapes and estimating temporally consistent point correspondences accurately. Despite only using 1024 points at training time, we can densify the points during test time by sampling new points and combining the result to get a dense deformed version of each shape. We use this property to further evaluate the quality of the deformation and use Poisson surface reconstruction [9] on 300K points to get surface reconstruction.

Figure 4 shows a visualisation of the meshes obtained after reconstructing with our method followed by Poisson



Figure 4. **Mesh Visualisation.** We sampled 300K points and use Poisson surface reconstruction [9] to get meshes from our model. Left we have the reconstructed point clouds. Right we have the reconstructed meshes with Poisson. We use a supervised GNPM version of our model where we assume we have access to dense correspondences. And we also shows the result of evaluating NPM [20] on depth and complete 3D scans. As we can see that our method supervised/self-supervised can recover good quality meshes.

Method	Input	$\mathbf{EPE}\downarrow$	\mathbf{C} - $l_2 (imes 10^{-3}) \downarrow$
NPM [20]	3D scan	0.231	0.019
GNPM (K 10)	3D scan	0.287	0.0122

Table 1. Comparison with NPM [20] on real scans of CAPE [14]. Since GNPM (our approach) requires full scans as input, we also evaluated NPM [20] on complete input scans for a fair comparison. Note that, in contrast to our method, NPM [20] can take partial observations (depth maps) as input at test time. However, it requires known dense correspondences at training time, unlike our approach.



Figure 5. Results of time-consistent dense correspondence estimation for 3D point-cloud inputs from the MANO [25] dataset. We show ground truth dense correspondences for comparison.

reconstruction [9]. We note that Poisson surface reconstruction [9] requires surface normals which we estimate. For fair comparison we sample 300K GT points and run them through Poisson reconstruction to obtain the GT meshes. We also visualise the results obtained by NPM [20] with two types of inputs: depth images and 3D scans. As we can see our method can obtain reasonable meshes.

In addition, we trained a supervised version GNPM where GT dense correspondences were given at training time (similarly to NPM). For this model we only use the backward network and we learn to disentangle shape and pose latents with the same network. Figure 4 shows visualisations of the predicted dense point clouds and reconstructed meshes with this 'supervised' version of our model.

Table 1 shows a numerical evaluation against NPM [20] on the CAPE test set. It is important to note that our approach GNPM can only be used with complete scans as input. For that reason we also evaluated NPM [20] using complete input scans even though this method can take partial observations (depth images) as input. But since ours cannot, we evaluated on complete input scans for fairness.

Table 1 shows our method achieves better Chamfer distance $(C-l_2)$ while having a comparable End-Point Error



Figure 6. **Shape and pose transfer in latent space.** We can transfer a given identity to a posed shape (shape transfer). In addition, given a source identity in different poses, we can repose a target identity with the poses of the source (pose transfer).



Figure 7. Unsupervised 3D part segmentation with clustering on learned feature space with EdgeConv. In left column the clusters are computed with Kmeans using 7 clusters on the features of first EdgConv layer. And in right column we experiment with different number of clusters and observed increasing the number of cluster leads to less semantically meaningful parts.

(EPE) to NPM [20], which is trained with ground truth dense correspondences, while our method works without them.



Figure 8. Qualitative evaluation of geometric inductive bias on CAPE [14] datasets. We replace all EdgeConv layers with MLPs following the NPM [20] architecture choices for the MLPs. As we can see the MLP-based architecture struggles to model the deformation of the arms while the EdgConv based architecture does not suffer from this problem.

GNPM can find the shape and pose latents more efficiently during test time. While NPM [20] report that optimizing over an input sequence of 100 frames takes approximately 4 hours on a GeForce RTX 3090, our approach takes 20 minutes on a GeForce RTX 3080 on a similar length sequence.

Method	Input	$\mathbf{EPE}\downarrow$	\mathbf{C} - $l_2 (\times 10^{-3}) \downarrow$
ONet 4D [15]	3D scan	-	0.028
OFlow [19]	3D scan	-	0.031
GNPM (K 10)	3D scan	0.253	0.0094

Table 2. Comparison with state-of-the-art methods on real scans of the DFAUST [3] dataset. Since OFlow [19] works only on sequences of up to 17 frames, we report the average over sub-sequences of such length.

4.2. Evaluation on DFAUST Dataset

Table 2 shows a quantitative evaluation on the DFAUST dataset [3]. All methods take 3D scans as input and our approach results in the lowest Chamfer distance error. We take the results for ONet [15] and OFlow [18] directly from the publications.

4.3. Qualitative evaluation on Mano Dataset

Hand reconstruction. We demonstrate GNPMs ability to work on the MANO hand dataset [25]. Figure 5 shows our GNPM results showing its ability to model complex deformations of the hand and to establish temporally consistent and dense correspondences. Our comparison with ground truth correspondences shows that GNPM can infer high quality dense 4D temporal tracking, despite the challenging deformations shown by hands.

4.4. Latent Applications

Shape and Pose Transfer: Disentangling shape and pose spaces allows us to transfer shape and/or pose between identities. Given an input identity in a specific pose, we can map the same pose to new identities. Alternatively, we can fix the identity and transfer pose latents. Examples of both are shown in Fig. 6.

4.5. Unsupervised 3D Part Segmentation

In this section we explore the potential of the features learnt by the EdgeConv network. Figure 7, shows the result of performing unsupervised clustering on the learnt features for each frame independently, which leads to consistent part segmentations across identities and poses. We can observe that by grouping points according to their features we discover *locally consistent* groups, to give a semantically meaningful part segmentation. We experimented with different cluster sizes as shown in Fig. 7.

4.6. Ablation Studies

Geometric Inductive Bias: To evaluate the our architecture choices, we replace all EdgeConv layers with MLP-based forward and backward networks using NPM [20] pose model architecture choices and trained the networks on CAPE [14] dataset. The MLP-based self-supervised architecture infers the deformation field independently for each point without considering local neighbourhood structure. As shown in Table 3, we found using EdgeConv layers is advantageous to learning deformations. Figure 8 shows qualitative evaluation. MLP based architecture struggles with modelling arms and can not learn a deformation field that fully reconstructs the shape. On the other hand, the EdgConv based architecture does not suffer from this problem. We conclude this inductive bias is important in learning deformation fields without dense correspondences. And removing this bias results in a degenerate performance.

Method	EPE ↓	\mathbf{C} - $l_2 (\times 10^{-3}) \downarrow$
Self-supervision (MLP)	0.515	0.0635
Self-supervision (EdgConv)	0.287	0.0122

Table 3. Evaluation of geometric inductive bias. We use NPM [20] MLP architecture choices and replace all EdgeConv layer with MLPs. As we can see without considering the local neighbourhood around the points the model fail to learn deformation field.

ICP Loss: We found that for frames with extreme deformation, the forward network has some difficulties dealing with points from similar parts like hands that change sides with respect to the canonical t-pose see supplementary materials. Using an asymmetric ICP loss leads to flipping in the prediction of the forward network. The backward network can handle this to some extent but as we are enforcing a geometric constraint, it fails to completely reverse points flipping resulting in a knot in the waist of the shape. Using a symmetric ICP loss where the nearest neighbours distance is computed from both source and target is important in learning extreme deformation without dense correspondences.

5. Limitations

While GNPM demonstrates potential for learning disentangled pose and shape space without dense correspondences, it still struggles with modelling fine details. For instance, in modelling human deformation often the hand is not modelled in great detail. A potential reason for that could be our uniform sampling approach. Denser sampling in areas like hands and faces could lead to better estimates of detailed deformations. Finally, our method is restricted to full meshes at test time. Future work includes extending our method to deal with partial observations such as depth maps.

6. Conclusion

We have proposed Geometric Neural Parametric Models (GNPM), a learned parametric model that takes as input point-clouds of different identities performing complex motions and disentangles shape/identity and deformations by mapping each pose to its canonical t-pose configuration and learning dense time-consistent correspondences. In addition, our model learns disentangled shape and pose embeddings which can later be used for interpolation, editing or transfer.

In contrast to previous learnt neural parametric models such as NPMs [20], our model does not require ground truth known correspondences at training time, which are costly to obtain. We exploit the concept of cycle-consistency to establish correspondences and show results on real-world scans from the CAPE [14], DFAUST [3] and MANO [25] datasets. Our network uses edge convolutions to extract features, which can be further exploited for downstream tasks such as unsupervised part segmentation. While our current approach cannot deal with partial scans from depth input images we envisage a completion network as future work.

Acknowledgements

Research presented here has been supported by the UCL Centre for Doctoral Training in Foundational AI under UKRI grant number EP/S021566/1.

References

- Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. *ACM Trans. Graph*, 24:408–416, 2005.
- [2] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. LoopReg: Self-supervised learning of implicit surface correspondences, pose and shape for 3D human mesh registration. Advances in Neural Information Processing Systems, (NeurIPS), 2020.
- [3] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [4] Aljaž Božič, Pablo Palafox, Michael Zollhöfer, Justus Thies, Angela Dai, and Matthias Nießner. Neural deformation graphs for globally-consistent non-rigid reconstruction. CVPR, 2021.
- [5] Benjamin Charlier, Jean Feydy, Joan Alexis Glaunès, François-David Collin, and Ghislain Durif. Kernel operations on the gpu, with autodiff, without memory overflows. *Journal of Machine Learning Research*, 22(74):1–6, 2021.
- [6] Marvin Eisenberger, David Novotný, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 7473–7483. Computer Vision Foundation / IEEE, 2021.
- [7] Jean Feydy, Joan Glaunès, Benjamin Charlier, and Michael Bronstein. Fast geometric learning with symbolic matrices. Advances in Neural Information Processing Systems, 33, 2020.
- [8] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, pages 8320–8329. Computer Vision Foundation / IEEE Computer Society, 2018.
- [9] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction.
- [10] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. 3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings, pages 1–15, 2015.
- [11] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. ACM Transactions on Graphics, 36(6):194:1– 194:17, Nov. 2017. Two first authors contributed equally.
- [12] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. arXiv preprint arXiv:2105.01905, 2021.
- [13] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIG-GRAPH Asia), 34(6):248:1–248:16, Oct. 2015.

- [14] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to dress 3d people in generative clothing. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, pages 6468– 6477. Computer Vision Foundation / IEEE, 2020.
- [15] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR* 2019, Long Beach, CA, USA, June 16-20, 2019, pages 4460– 4470. Computer Vision Foundation / IEEE, 2019.
- [16] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. 2020.
- [17] Richard A. Newcombe, Dieter Fox, and Steven M. Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *IEEE Conference on Computer Vision* and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015, pages 343–352. IEEE Computer Society, 2015.
- [18] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4D reconstruction by learning particle dynamics. *Proceedings of the IEEE International Conference on Computer Vision*, pages 5378–5388, 2019.
- [19] Michael Niemeyer, Lars M. Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, pages 5378–5388. IEEE, 2019.
- [20] Pablo R. Palafox, Aljaz Bozic, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3d deformable shapes. *ICCV 2021*.
- [21] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 165–174. Computer Vision Foundation / IEEE, 2019.
- [22] Despoina Paschalidou, Angelos Katharopoulos, Andreas Geiger, and Sanja Fidler. Neural parts: Learning expressive 3d shape abstractions with invertible neural networks. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021.
- [23] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In Stefano Tubaro and Jean-Luc Dugelay, editors, Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance, AVSS 2009, 2-4 September 2009, Genova, Italy, pages 296– 301.
- [24] Stylianos Ploumpis, Haoyang Wang, Nick E. Pears, William A. P. Smith, and Stefanos Zafeiriou. Combining 3d morphable models: A large scale face-and-head model. pages 10934– 10943, 2019.
- [25] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies

together. ACM Transactions on Graphics, (Proc. SIGGRAPH Asia), 36(6), Nov. 2017.

- [26] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. pages 2886–2897, 2021.
- [27] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In Proceedings of EUROGRAPHICS/ACM SIG-GRAPH Symposium on Geometry Processing, pages 109–116, 2007.
- [28] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph CNN for learning on point clouds. ACM Trans. Graph., 38(5):146:1–146:12, 2019.
- [29] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A comprehensive survey on graph neural networks. *IEEE Trans. Neural Networks Learn. Syst.*, 32(1):4–24, 2021.
- [30] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T. Freeman, Rahul Sukthankar, and Cristian Sminchisescu. GHUM & GHUML: generative 3d human shape and articulated pose models. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, pages 6183–6192, 2020.
- [31] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6365–6373, 2017.