

SPAMs: Structured Implicit Parametric Models

Pablo Palafox^{1*} Nikolaos Sarafianos² Tony Tung² Angela Dai¹

¹Technical University of Munich ²Meta Reality Labs Research, Sausalito, USA

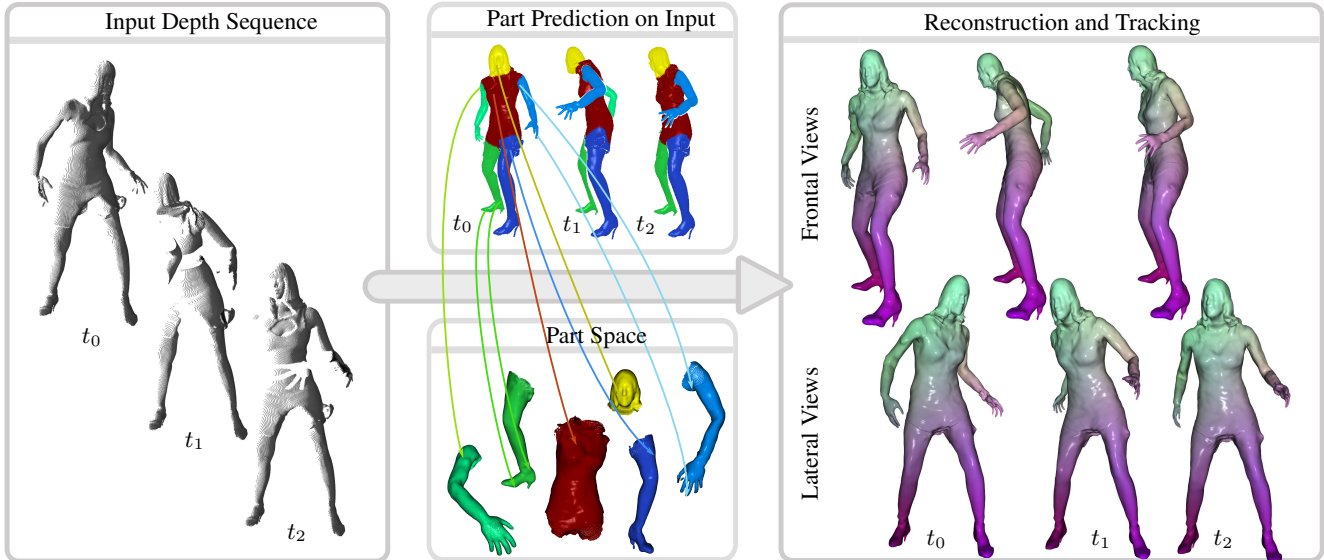


Figure 1. We propose to represent deforming shapes with a structural decomposition into part-based disentangled spaces characterizing of shape and pose, as Structured-implicit Parametric Models (SPAMs). SPAMs learn optimizable local shape and pose spaces, which we can traverse at test time to fit to depth sequence observations of an unseen deforming object. Our structured part decompositions facilitate low-dimensional coarse motion correspondence through part-to-part correlation, which guides a robust joint optimization over the part-based shape and pose spaces for globally consistent, accurate tracking under complex motion sequences.

Abstract

Parametric 3D models have formed a fundamental role in modeling deformable objects, such as human bodies, faces, and hands; however, the construction of such parametric models requires significant manual intervention and domain expertise. Recently, neural implicit 3D representations have shown great expressibility in capturing 3D shape geometry. We observe that deformable object motion is often semantically structured, and thus propose to learn Structured-implicit Parametric Models (SPAMs) as a deformable object representation that structurally decomposes non-rigid object motion into part-based disentangled representations of shape and pose, with each being represented by deep implicit functions. This enables a structured characterization of object movement, with part decomposition characterizing a lower-dimensional space in which we can establish coarse motion correspondence. In particular, we can leverage the part decompositions at test time to

fit to new depth sequences of unobserved shapes, by establishing part correspondences between the input observation and our learned part spaces; this guides a robust joint optimization between the shape and pose of all parts, even under dramatic motion sequences. Experiments demonstrate that our part-aware shape and pose understanding lead to state-of-the-art performance in reconstruction and tracking of depth sequences of complex deforming object motion.

1. Introduction

Understanding non-rigidly deforming shapes is essential for perception, as we live in a 4D world where humans, animals, and other 3D objects move in a non-rigid fashion. Dynamic tracking and reconstruction remains a notable challenge, and while significant advances have been made in 4D reconstruction and tracking, they often require complex multi-view setups [5] or build on a domain-specific,

*This work was conducted during an internship at Meta RL Research.

fixed-topology template [1,4,35,36]. In the latter case, parametric 3D models in particular have made notable impact in modeling domain-specific deformable 3D objects, such as for human bodies [2,20,23], faces [22,34], hands [40], and animals [51]. However, such parametric 3D models require a complex construction process involving domain-specific knowledge and manual efforts, while remaining limited in expressability of local shape details.

Recently, advances in learned continuous implicit representations for modeling 3D shapes have shown impressive representation power for capturing effective static 3D shape geometry at relatively high resolutions [7, 8, 17, 28, 29, 33, 39]. Such approaches have also been extended to represent 4D reconstruction of dynamic objects by efficiently disentangling learned implicit spaces representing shape and dynamic movement [31, 32]. This has proven to be a very promising direction, but these approaches characterize objects as a whole, whereas we observe that the 4D motion of an object typically maintains a strong structured correlation on a lower-level part basis.

Thus, we propose Structured-implicit PArametric Models (SPAMs), which learn a structured, part-based, disentangled representation of deformable 3D objects. Given a set of observations of various shape identities in different poses (including a canonical pose) with coarse part annotations, we learn part-based latent spaces characterizing each part’s geometry and motion. Note that we do not require comprehensive surface correspondence throughout the dataset, nor complex domain-specific knowledge (e.g., skeleton, kinematic chain). We leverage continuous implicit function representations for each part’s geometry, represented as a signed distance field in its canonical space, and pose, represented as a local deformation relative to the canonical space.

At test time, we traverse the learned latent part spaces to fit to new depth sequences. Crucially, our part-based representation allows leveraging predicted part segmentation of the new observation to establish global correspondences with our part-based latent representations. By establishing correspondence through our part priors, we can robustly track sequences with significant motion changes by discovering high-level part correspondence and leveraging it to guide our joint optimization over part-based shape and pose. Experiments on non-rigid tracking and reconstruction of single-camera depth sequences of humans from the RenderPeople dataset [12] show that the part-aware reasoning of our SPAMs can outperform the state of the art by an order of magnitude on reconstruction (Chamfer distance) and by 43% on tracking (3D End-Point-Error). In summary, we present the following contributions:

- We learn a part-based disentanglement of shape and pose, capturing local characteristics of deformable 3D objects in latent spaces representing shape and pose of each part.

- Our learned, optimizable spaces enable part-based reasoning to guide joint optimization over parts to fit unseen test sequences. By establishing high-level part correspondence between a new observation and our learned part spaces, we can robustly guide a joint optimization over part geometry and pose, resulting in a more globally consistent non-rigid reconstruction and tracking.

2. Related Work

Parametric and Neural Parametric Models. Parametric body models [23, 48, 50] enable the representation of the human body variations with a limited number of parameters (*i.e.* a low-dimension descriptor). One of the most influential works in this line of research is SMPL [23], a parametric model widely-used to describe the body shape and pose with deformation blend shapes learned from a dataset of diverse 3D body scans. Parametric body models have enabled research work for modeling soft tissue [36] and clothing [1, 25, 45]. However, parametric body models [23, 48] employ deformations based on vertex-based skinning models which have limited resolution and cannot represent non-linear surface deformations of clothed bodies (*e.g.*, wrinkles). To alleviate these limitations, Neural Parametric Models [32] learn to disentangle 4D dynamics into latent-space representations of shape and pose using implicit functions and can fit to new observations by optimizing over the pose and shape codes. However, they treat the human body as a single entity, which results in somewhat unrealistic motions. Genova *et al.* [16, 17] addressed this problem by introducing LDIF, a 3D representation that implicitly describes a shape as the sum of local 3D functions. LDIF outputs a structured decomposition into shape elements by encoding 3D points within each shape using PointNet [15, 37, 38] and hence enable shape encoding such as human bodies in local regions arranged in a global structure. The 3D decompositions of LDIF tend towards temporal consistency, but as tracking is not explicitly considered, surface tracking tends to become inconsistent in more challenging motion scenarios. In contrast, we learn a semantically-driven part decomposition which guides joint part-based shape and pose optimization that results in robust, consistent tracking over dynamic sequences.

Continuous Implicit Deformable Representations. Implicit-based and generative representations for humans or clothing has been an active topic of research for the past few years [9, 10, 30, 33, 42, 43, 47, 49]. Several recent works have focused on learning identity-specific implicit representations to animate clothed people [11, 13, 19, 30, 41, 46]. Neural-GIF [46] introduced a framework to animate clothed people from scans as a function of pose directly, without the need of registration. SCANimate [41] and LEAP [30] learn a pose or shape representation of the human body surface. While promising, these methods learn

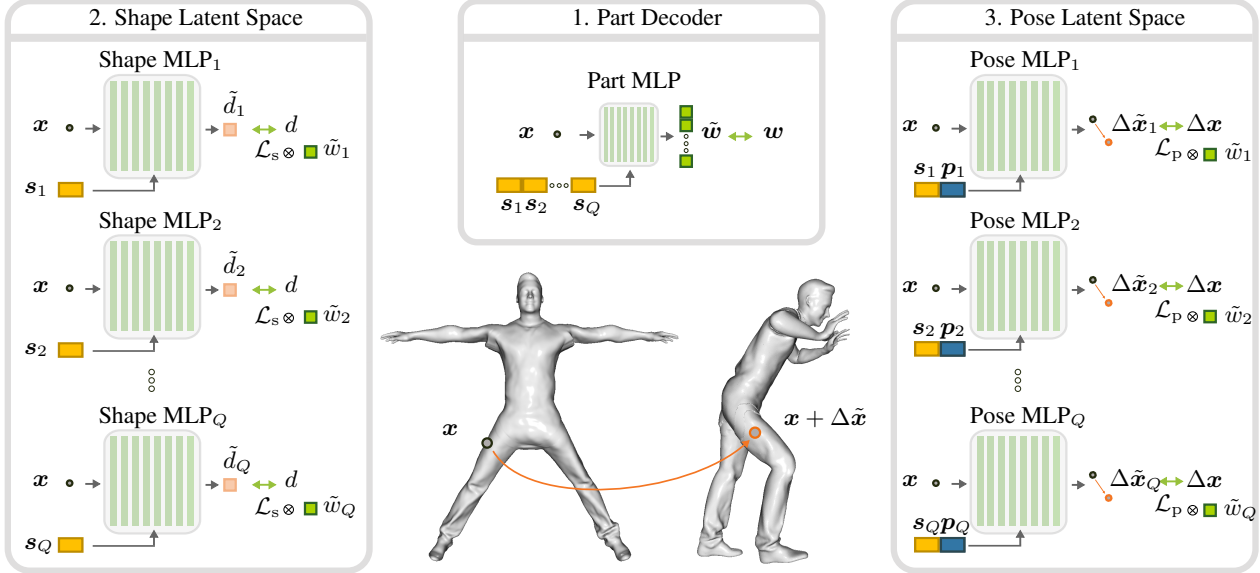


Figure 2. **Structured Implicit Parametric Model Construction.** SPAMs learn a structured decomposition of shape and pose, where local shape codes $\{s_q\}$ can be decoded to represent local part geometry, and both $\{s_q\}$ and a part decoder inform local pose movement. To construct a SPAM, we first learn structural decomposition in the canonical space as a part decoder that predicts part classes, conditioned on all local shape codes $\{s_q\}$; we use this learned part decomposition to guide local (shape and pose) MLPs into focusing on their corresponding space partitioning. We learn a structured, latent shape space in the canonical space by conditioning a set of local shape MLPs on local shape codes $\{s_q\}$ assigned to every identity and part q . We learn local deformation fields around the canonically-posed shape with local pose MLPs conditioned on a local, latent shape and pose codes s_q and p_q to predict Δx for x in the canonical space.

subject or outfit-specific models, and thus lack general characterization of deformable objects. Given a posed but unclothed body model, POP [27] represents a clothed human with a set of points in order to create pose-dependent clothing animations. New animations can then be created from an unseen scan of a person in clothing. Bhatnagar *et al.* [3] proposed IP-Net, a method to combine learned implicit functions and traditional parametric models to produce controllable models of humans. IP-Net predicts correspondences to SMPL and leverages a double-layered surface to represent inner and outer surfaces to better represent a clothed body. Our proposed approach also builds on the expressability of learned implicit representations, and we propose to learn from a dataset without strong requirements on surface correspondence annotations to construct a general, semantically-driven decomposition that provides strong high-level guidance for robust pose tracking under challenging motion scenarios.

3. Structured Implicit Parametric Models

We introduce Structured-implicit Parametric Models (SPAMs), a learned approach to build part-aware, implicit, parametric 3D models from a dataset of posed identities of a given object class. We do not require the dataset to have surface correspondence between each instance nor annotations for physical domain-specific properties such as the skeleton or the kinematic chain. Our SPAMs structurally decom-

pose non-rigid object motion into part-based disentangled implicit spaces representing each part’s shape and pose.

SPAMs consist of (1) learned latent spaces that characterize part geometries in the canonical shape space, (2) a part decoder conditional on the part shapes that guides a part-based structural partitioning of the canonical space, and (3) structured, latent pose spaces conditional on both part shape and pose. At test time, SPAMs not only allow for the joint optimization over the learned, local spaces of shape and pose to fit to a new observation, but, crucially, our part-based formulation enables establishing high-level correspondence between predicted parts and the part-based latent representations. This part correspondence enables robust global optimization over the part shape and pose spaces, resulting in robust, consistent non-rigid reconstruction and tracking of unseen sequences of deforming objects.

3.1. Overview

Given a dataset of shape identities from the same class category in different poses, with coarse part segmentation in the canonical space, our goal is to learn a part-based parametric model that structurally disentangles shape and pose. We leverage this structured part decomposition to fit to unseen depth sequences of new identities, where our part-based representation provides a lower-dimensional space to establish semantic part correspondences that provide strong guidance for our part-based shape and pose fitting under

complex motions.

Our SPAMs characterize disentangled shape and pose as sets of local, part-based shape spaces over shape codes $\{\mathbf{s}_q\}$, and their corresponding local pose spaces over $\{\mathbf{p}_q\}$. To obtain the structured partitioning of these local spaces, we first learn a part decoder that predicts a segmentation of the canonical space into a set of local parts (Sec. 3.2). Given the learned space partitioning, we can then construct the structured shape space as a set of local implicit geometric representations (Sec. 3.3). We then build the structured pose space as a set of local implicit deformation fields that characterize the motion of each local shape (Sec. 3.4). Finally, to fit to new depth sequences of unobserved identities at test time, we leverage part correspondence to the set of shape and pose spaces to obtain a robust, joint optimization that accurately represents the observed shape (Sec. 3.5).

3.2. Part Decoder

In order to establish a structural space partitioning, we learn a part decoder that segments the canonical space and guides the learning of the part-based shape and pose latent spaces. The part decoder f_{θ_q} is characterized by an MLP that predicts part confidences for a query point $\mathbf{x} \in \mathbb{R}^3$ in canonical space. We implement our part decoder in an auto-decoder fashion [33] and condition on the concatenation of D_s -dimensional local latent shape codes \mathbf{s}_q (Sec. 3.3). Let Q denote the number of parts, leading to a QD_s -dimensional concatenation of local shape codes. Formally, we have:

$$f_{\theta_q} : \mathbb{R}^{QD_s} \times \mathbb{R}^3 \rightarrow \mathbb{R}^Q, \quad ([\mathbf{s}_q], \mathbf{x}) \mapsto f_{\theta_q}([\mathbf{s}_q], \mathbf{x}) = \tilde{\mathbf{w}}. \quad (1)$$

where $[\mathbf{s}_q]$ is the concatenation of all Q local shape codes $\{\mathbf{s}_q\}$. We train f_{θ_q} with a binary cross entropy loss between predicted and ground-truth part labels as one-hot vectors.

We consider six coarse parts: head, torso, right arm, right leg, left arm and left leg (see Fig. 1-center). A point $\mathbf{x} \in \mathbb{R}^3$ in canonical space can belong to any of these part classes, where the Q -dimensional $\tilde{\mathbf{w}}$ denotes the likelihoods of \mathbf{x} belonging to each respective part. In particular, we allow points near the boundary between two parts to belong to both parts at the same time, by indicating both parts to be one in the ground truth part vector. This allows regularization of boundary regions to ensure smooth transitions between implicit functions for shape and pose latent spaces.

Note that f_{θ_q} predicts part probabilities for all points $\mathbf{x} \in \mathbb{R}^3$ in canonical space, not only surface points. This strategy provided more accurate part predictions for off-surface locations during our test-time optimization (Sec. 3.5).

Groundtruth for supervising the part decoder is obtained as follows: as will be shown in Sec.3.3, to train our shape latent spaces we sample dense surface points with normals in addition to off-surface points randomly sampled in $[-1, 1]^3$.

We re-use these same points to train our part decoder, with part labels for each surface point obtained directly from the original scan or mesh and by nearest neighbor query for off-surface points.

Implementation details. We use a 6-layer SIREN [44] MLP with a hidden dimension of 256 and a frequency $\omega = 30$ in the sinus activation (following [44]). As shown by [6], conditioning-by-concatenation is sub-optimal for implicit neural representations with period activations, and therefore we employ their proposed FiLM conditioning, where a mapping network takes in a latent code \mathbf{z} and outputs frequencies and phase to condition each layer of the SIREN MLP. The mapping network is a 4-layer LeakyReLU MLP with hidden dimension of 128. We use the Adam optimizer [21] and a learning rate of 1×10^{-4} for both the decoder and the mapping network.

3.3. Structured Shape Space

Our multi-part shape space is learned by a dictionary of Q local MLPs, each learning to represent a local shape part in its canonical pose, characterizing its geometry as the zero iso-surface decision boundary of a signed distance field. The structural decomposition of the Q local MLPs is guided by the part decomposition predicted by the part decoder f_{θ_q} . To extract the complete shape from our structured shape space, we query all Q shape MLPs for every point in a 3D grid, average out their SDF contributions based on predicted part confidence for the given query point, and finally use Marching Cubes [24] to extract a mesh.

Each shape MLP is trained in auto-decoder fashion [33], similar to the part decoder. We directly optimize over a latent code \mathbf{s}_q , which is particular to the shape latent space of its assigned part q . Each part q of a canonically-posed shape identity i in the training set is then encoded in a D_s -dimensional latent shape code \mathbf{s}_q^i . In turn, each shape MLP $f_{\theta_q^i}$ learns to map an input point $\mathbf{x} \in \mathbb{R}^3$ in the canonical space, conditioned on the local shape code \mathbf{s}_q^i , to a predicted SDF value \tilde{d}_q^i :

$$f_{\theta_q^i} : \mathbb{R}^{D_s} \times \mathbb{R}^3 \rightarrow \mathbb{R}, \quad (\mathbf{s}_q^i, \mathbf{x}) \mapsto f_{\theta_q^i}(\mathbf{s}_q^i, \mathbf{x}) = \tilde{d}_q^i. \quad (2)$$

As train data typically do not contain watertight meshes, we train directly on oriented point clouds sampled from the (potentially incomplete) train meshes, following [18]. This is accomplished by solving for an Eikonal boundary value problem that constrains the norm of spatial gradients of the SDF to be 1 almost everywhere [18,44]. In practice, we use the Eikonal loss as presented in [44]. As training data, we only require the surface points with associated normals for each identity, and randomly sampled coordinates in the unit cube. For every identity in a batch, we sample N_s points, half of which are surface points subsampled from the given identity, and the remaining half as coordinates randomly

sampled in space. Then for every part q , we train the corresponding shape MLP by minimizing the following reconstruction energy over the S canonically-posed shape identities of the dataset with respect to the local shape codes $\{\mathbf{s}_q^i\}_{i=1}^S$ and the set of shape MLP weights θ_s^q :

$$\arg \min_{\theta_s^q, \{\mathbf{s}_q^i\}_{i=1}^S} \sum_{q=1}^Q \sum_{i=1}^S \left(\sum_{k=1}^{N_s} w_q^{i,k} \mathcal{L}_s(f_{\theta_s^q}(\mathbf{s}_q^i, \mathbf{x}_i^k), d_i^k) + \mathcal{L}_r \right). \quad (3)$$

Here \mathcal{L}_s is the Eikonal loss version proposed in SIREN [44], enforcing that (1) SDF predictions for surface points are 0, (2) groundtruth surface normals match the estimated normals (computed as the spatial gradient of the SDF function at a given position), (3) the norm of the SDF gradient is 1 almost everywhere and (4) off-surface points do not have SDF values close to 0. We refer to [44] for a more detailed explanation. Additionally, we enforce a zero-mean multivariate-Gaussian distribution with spherical covariance σ_s over the latent shape codes, as was proposed in [33]: $\mathcal{L}_r = \|\mathbf{s}_q^i\|_2^2 / \sigma_s^2$.

Importantly, the contribution to the loss of each point is weighted by its predicted part confidence $w_q^{i,k}$ from f_{θ_q} , enabling each local shape MLP to focus on its respective local region. That is, if for point \mathbf{x} , $w_q^{i,k}$ for a class q is close to 1, then \mathbf{x} likely belongs to q . Then gradients from the contribution of point \mathbf{x} will back-propagate with high weight to the shape MLP of part q , θ_s^q . On the contrary, if $w_q^{i,k}$ is near 0, then \mathbf{x} likely does not belong to q and f_{θ_q} will not learn to generate geometry at \mathbf{x} .

Implementation details. Each shape latent space is implemented as a 6-layer SIREN [44] MLP with a hidden dimension of 256 and a frequency $\omega = 30$. We use 128-dimensional shape latent codes for each part ($D_s = 128$). Similar to our part decoder, we employ FiLM conditioning [6] to condition the SDF predictions on the local latent shape code, with the mapping network implemented as a 4-layer LeakyReLU MLP with 128 units per layer. We use the Adam optimizer and learning rates of 1×10^{-4} and 1×10^{-3} for the local shape MLPs $f_{\theta_s^q}$ and the local shape codes $\{\mathbf{s}_q^i\}_{i=1}^S$, respectively. The latent shape codes are initialized randomly from $\mathcal{N}(0, 0.01^2)$.

3.4. Structured Pose Space

Similar to our structured shape latent space, our structured pose space is learned by a dictionary of Q local MLPs, each optimized to represent a local deformation field that maps query points \mathbf{x} in the canonical space of an identity i to a deformed space j , by predicting a flow vector $\Delta \tilde{\mathbf{x}}_q$. This prediction is conditional on both a D_p -dimensional latent pose code \mathbf{p}_q^j as well as on the latent shape code \mathbf{s}_q^i of the corresponding part q , as pose deformations change with respect to shape. Deformation fields are only defined for a thin layer around the shape surface, since flow vectors become less informative when further away from the surface.

Formally, we have:

$$f_{\theta_p^q} : \mathbb{R}^{D_s} \times \mathbb{R}^{D_p} \times \mathbb{R}^3 \rightarrow \mathbb{R}^3, \\ (\mathbf{s}_q^i, \mathbf{p}_q^j, \mathbf{x}) \mapsto f_{\theta_p^q}(\mathbf{s}_q^i, \mathbf{p}_q^j, \mathbf{x}) = \Delta \tilde{\mathbf{x}}_q.$$

Our local pose MLPs are trained with up to P deformed instances of an identity; note that this does not require different identities to appear in the same pose.

To train these local pose spaces, we sample dense, near-surface correspondences between the canonical and posed frame, which amounts to sampling the canonical and posed raw meshes at the same barycentric coordinates, similar to [32]. Learning one of the local pose spaces θ_p^q amounts to minimizing the following energy term over all P deformation fields with respect to the individual (and local) pose codes $\{\mathbf{p}_q^j\}_{j=1}^P$ and pose MLP weights θ_p^q :

$$\arg \min_{\theta_p^q, \{\mathbf{p}_q^j\}_{j=1}^P} \sum_{i(j)}^P \left(\sum_{q=1}^Q \sum_{k=1}^{N_p} w_q^{i,k} \mathcal{L}_p + \mathcal{L}_r \right), \quad (4)$$

where $i(j)$ is a mapping from the index j of a posed shape to the corresponding index i of its canonical shape and \mathcal{L}_p is an ℓ_2 loss between predicted and ground truth flow vectors:

$$\mathcal{L}_p := \mathcal{L}_p(f_{\theta_p^q}(\mathbf{s}_q^i, \mathbf{p}_q^j, \mathbf{x}_i^k), \Delta \mathbf{x}_{ij}^k). \quad (5)$$

Similar to the structured local shape spaces, we encourage corresponding structural decomposition of the local part spaces by leveraging the predicted part confidences $w_q^{i,k}$ from f_{θ_q} . We also enforce an analogous zero-mean multivariate-Gaussian distribution over the latent pose codes with \mathcal{L}_r . While learning these pose latent spaces, we do not optimize over the local shape codes.

Implementation details. We implement our structured pose latent spaces as 4-layer SIREN MLPs with hidden dimension equal to 256, $\omega = 15$, and set $D_p = 64$ as the dimension of our local pose codes. We employ the same training scheme as used for training the shape spaces.

3.5. Test-time Optimization

Our structured latent representations of shape and pose can be optimized over at test time to fit a SPAM to accurately reconstruct and track an input sequence of L depth maps. This is achieved by solving for the set of Q local latent shape codes, denoted as $[\tilde{\mathbf{s}}]$, and the $Q \times L$ latent pose codes (Q pose codes per frame), $\{[\tilde{\mathbf{p}}_j]\}_{j=1}^L$, that best explain the input sequence.

Each depth map in the sequence is interpreted as a 256^3 -SDF grid of its back-projected values. A volumetric mask M_o is also extracted to mask out regions that are further than 0.02 (in normalized units) from the observed surface. We additionally predict part labels for every point




Figure 3. Our part-aware \mathcal{L}_{icp} can establish high-level correspondences between input points (left) and posed estimates (right), even under difficult motions with poor initial pose estimates (e.g., right arm in green), thus guiding optimization to accurate tracking.

in the input depth map using PointNet++ [38], which we pre-train to predict part labels for the Q parts.

Prior to optimization, we initialize our local shape and pose codes. To initialize each local shape code \mathbf{s}_q , we use the average optimized train code for each part q . To initialize local pose codes, we leverage a learned pose encoder that maps the input depth map to a latent code. While we find that our structured part representations of shape and pose can robustly track from a random pose initialization of all codes from $\mathcal{N}(0, 0.01^2)$, we can obtain improved pose tracking by a learned encoder initialization (c.f. Sec. 4).

Given the initial shape code estimates, we can extract an initial canonical shape by querying our structured shape MLPs on a 3D grid, and extracting the iso-surface with Marching Cubes [24]. We use the initial canonical shape to inform sampling near the surface, and sample $N_t = 500\text{k}$ points, $\{\mathbf{x}_k\}_{k=1}^{N_t}$, around this initial estimate of the canonical shape; during optimization, for each frame in a mini-batch, we sub-sample $N_b = 20\text{k}$ points out of the available N_t to minimize the following equation:

$$[\tilde{\mathbf{s}}_q], \{\{\tilde{\mathbf{p}}_q^j\}_{j=1}^L\} = \arg \min_{[\mathbf{s}_q], \{\{\mathbf{p}_q^j\}_{j=1}^L\}} \sum_{j=1}^L \sum_{\forall \mathbf{x}_k} \mathcal{L}_r + \mathcal{L}_c + \mathcal{L}_t + \mathcal{L}_{\text{icp}}. \quad (6)$$

\mathcal{L}_c enforces shape and pose code regularization as in training, and \mathcal{L}_t enforces temporal regularization between the current frame j and its neighboring frames (for more detail, we refer to the supplementary material).

3.5.1 Structurally-Guided Shape & Pose Optimization

In order to inform our reconstruction losses \mathcal{L}_r and \mathcal{L}_{icp} , we employ structural correspondences between parts predicted in the observed depth views as well as the part decomposition from our latent shape and pose spaces (see Fig. 3).

We use a clamped ℓ_1 [14] reconstruction loss \mathcal{L}_r :

$$\mathcal{L}_r = \sum_{q=1}^Q M \mathcal{L}_s \left(f_{\theta_s^q}(\mathbf{s}_q, \mathbf{x}_k), \left[\mathbf{x}_k + f_{\theta_p^q}(\mathbf{s}_q, \mathbf{p}_q^j, \mathbf{x}_k) \right]_{\text{sdf}} \right), \quad (7)$$

where $[\cdot]_{\text{sdf}}$ denotes trilinear interpolation of the SDF grid and $M = M_{\text{part}}^q M_o$ leverages part and occlusion information to inform the reconstruction. M_o denotes a mask of

unoccluded regions, as defined in Sec. 3.5. M_{part} represents a grid of part label predictions, as given by the part decoder. We compute M_{part} by randomly sampling points in the canonical space along with their predicted part labels from f_{θ_q} , warping these points to each frame by $f_{\theta_p^q}$, and finally querying for every grid location the nearest point from the warped points to obtain the part label for the voxel. This enables focusing the reconstruction locally for each estimated part geometry and pose.

Finally, we employ a part-guided ICP-inspired loss \mathcal{L}_{icp} , which plays a key role in ensuring robustness to pose code initialization. That is, we establish part-driven correspondences between each estimated part q from the predicted input depth map part labels and the part decoder predictions. To this end, every I_{resample} iterations we consider the canonically-posed shape from the current state of the shape codes, re-sample a new set of $N_t = 500\text{k}$ points around the mesh and keep those within a distance ϵ_{icp} from the implicitly represented surface. We use our part decoder (Sec. 3.2) to estimate part labels for these canonical points, and then warp them into a posed frame j using our pose decoders. Then for every point in the input depth map, and given its predicted part q (obtained by PointNet++), we find its nearest neighbor in the warped set of points belonging to q based on f_{θ_q} to establish correspondences. Crucially, this provides robust correspondences in challenging motion scenarios where pose initialization may be notably misaligned, as demonstrated in Sec. 4.

4. Experiments

We evaluate SPAMs on the task of model fitting to depth sequences (Sec. 4.1), and analyze the effect of our structured, part-based fitting in Sec. 4.2.

Datasets. We train and evaluate on the public RenderPeople dataset [12], which contains real-world 3D scans of people in clothing, post-processed to be minimally noise-free 3D meshes (i.e., removing holes, self-intersection). We train on 338 identities, each rescaled to a common scale in the unit bounding box as a simple data pre-processing step. To learn the structured pose space, we collected a set of 2,446 3D human motion sequences covering a wide variety of action categories of daily activities and sports, and animated canonically-posed identities. From this set of posed, clothed people, we randomly sample scan-motion pairs and obtain 40K randomly posed instances, without requiring seeing multiple identities in the same pose. We evaluate our method on six unseen test identities performing various dancing moves, comprising to a total of 540 test frames organized in 90-frame sequences per identity.

Evaluation metrics. To quantitatively evaluate model fitting to depth sequences, we measure reconstruction quality as well as tracking performance. To measure reconstruction quality, we follow the evaluation protocol of Occu-

Method	IoU \uparrow	C- ℓ_2 \downarrow	NC \uparrow	EPE \downarrow
IP-Net	0.729	0.00053	0.837	0.168
NPMs*	0.755	0.00163	0.856	0.053
Ours	0.785	0.00032	0.883	0.034

Table 1. Comparison with state-of-the-art NPMs* on the test set of our RenderPeople dataset. Our part-aware decomposition enables notably improved reconstruction and tracking.

pancyNets [28], computing Intersection over Union, Chamfer distance, and normal consistency on a per-frame basis. *Intersection over Union* (IoU) measures overlap between predicted and ground truth meshes, computed over 10^6 randomly sampled points in the unit bounding box. *Chamfer- ℓ_2* (C- ℓ_2) measures the bi-directional distance between the prediction and ground truth with 100k randomly sampled points on the surfaces, giving distance characterization to any potentially mismatched surface reconstruction. *Normal Consistency* (NC) measures surface quality as the mean absolute dot product of the normals of the predicted mesh with the normals of the corresponding nearest neighbors in the ground truth mesh. Tracking performance is evaluated as in prior non-rigid tracking works [5, 32] via the *End-Point Error* (EPE) as the average ℓ_2 -distance between predicted and ground truth deformations.

4.1. Model Fitting to Monocular Depth Sequences

We evaluate our SPAMs model fitting to new monocular depth sequences in comparison with state of the art on sequences rendered from our Renderpeople [12]-constructed dataset, as well as on CAPE [26] (see supplemental). We compare with the state-of-the-art Neural Parametric Models (NPMs) [32] and IP-Net [3]. We train NPMs on our Renderpeople training split. Since NPMs require watertight meshes for training to determine inside/outside, and inside/outside queries on our train data tend to be unreliable in often-articulated regions such as hands, we adapt NPMs as NPMs* which incorporates the Eikonal loss, SIREN activations, and FiLM conditioning of our approach. For IP-Net [3], we use a model checkpoint provided by the authors, which was also used to evaluate on RenderPeople.

Table 1 shows a comparison with NPMs* and IP-Net on fitting to monocular depth sequences. Our structural, part-driven representation of shape and pose results in notably improved reconstruction and tracking performance. SPAM achieves higher IoU and normal consistency, as well as significantly reduced Chamfer distance and 3D End-Point-Error, indicating more globally consistent tracking and reconstruction leveraging part-based fitting. Qualitative comparisons are depicted in Figure 4. Under significant, complex motion in the input sequence, our SPAM maintains robust tracking and consistent geometry, while NPM* fails to capture more dramatic motions (e.g. in the arms).

Method	IoU \uparrow	C- ℓ_2 \downarrow	NC \uparrow	EPE \downarrow
NPMs* (w/o PE)	0.269	0.01397	0.658	0.182
Ours (w/o PE, w/o PGM)	0.671	0.00070	0.836	0.052
Ours (w/o PE)	0.681	0.00065	0.839	0.052
Ours (w/o PGM)	0.766	0.00037	0.874	0.037
Ours	0.785	0.00032	0.883	0.034
Ours (w/ GT PS)	0.809	0.00021	0.894	0.026

Table 2. Ablation studies. We evaluate the effects of a learned pose encoder initialization (PE), using our part grid mask (PGM) in test-time fitting, and we additionally study the impact of ground truth part segmentation (GT PS) of input depth maps. Our part-driven shape and pose spaces provide significant robustness against lack of pose initializations, and leveraging part mask guidance for optimization additionally improves performance.

4.2. Ablations

Robustness to Pose Code Initialization. We observe that our part-aware disentanglement of shape and pose provides significant robustness in pose tracking, and in particular, maintains robustness in the absence of pose encoder initialization (which may not be available in scenarios such as generalizing to different sensor inputs). We demonstrate this in Table 2, comparing with a variant of our approach without pose encoder initialization (w/o PE) and instead using random initialization for pose codes. This results in very poor initial pose estimates, with an effectively random initial set of deformation fields. While this poor initialization results in slightly impaired performance, our structurally-guided SPAMs nonetheless can recover a significant portion of the reconstruction and tracking performance of using the pose encoder initialization. In contrast, we observe that the full-shape shape and pose encoding of NPMs* fails to recover from poor pose initializations. We refer to the supplementary material for further qualitative visualizations.

Part-based Grid Masking. Table 2 additionally evaluates the effect of leveraging part information as a part grid mask (PGM) during test-time optimization; this local part focus enables more accurate reconstruction and tracking.

Limitations. Our structured parametric modeling of deformable objects enables robust model fitting to challenging monocular depth sequences, but maintains several limitations. For instance, our local pose spaces do not characterize potential high-level motion priors given by motion very far from the local region (e.g., one hand moving back more often occurs with the other hand moving forward than not), which could provide additional global context. Additionally, very fine-scale sharp details can become oversmoothed across the global optimization, which could potentially be characterized with perceptually-oriented measures.




Figure 4. Comparison to the state-of-the-art NPMs* [32] and IPNet [3] on the task of model fitting to a monocular depth sequence. Our SPAMs show superior tracking performance, and better capability to capture high-frequency details (see jacket boundary on third column).

5. Conclusion

In this work we have introduced SPAMs, a deformable object representation where non-rigid object motion is structurally decomposed into part-based disentangled representations of shape and pose. Our structured characterization of object movement can be leveraged at test time to fit to input depth sequences of unseen shapes, leveraging part-based correspondences to establish robust optimization to fit to the input sequence. Our experiments show signifi-

cantly improved robustness in reconstruction and tracking, particularly in scenarios of challenging, complex motions in observed depth sequences. We believe that this representation will be useful in a variety of spatio-temporal tasks.

Acknowledgments: This work was done during an internship at Meta. We thank Yuanlu Xu for informative discussions.

References

- [1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1175–1186, 2019. 2
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM SIGGRAPH*, pages 408–416, 2005. 2
- [3] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Combining implicit function learning and parametric models for 3d human reconstruction. In *European Conference on Computer Vision*, 2020. 3, 7, 8
- [4] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *IEEE/CVF International Conference on Computer Vision*, pages 5420–5430, 2019. 2
- [5] Aljaž Božič, Pablo Palafox, Michael Zollhöfer, Justus Thies, Angela Dai, and Matthias Nießner. Neural deformation graphs for globally-consistent non-rigid reconstruction. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 7
- [6] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5799–5809, 2021. 4, 5
- [7] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, 2019. 2
- [8] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6970–6981, 2020. 2
- [9] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3d shape reconstruction and completion. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [10] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, December 2020. 2
- [11] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11875–11885, 2021. 2
- [12] RenderPeople Dataset. <http://renderpeople.com/>. 2, 6, 7
- [13] Boyang Deng, John P Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Nasa neural articulated shape approximation. In *ECCV*, 2020. 2
- [14] Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J Guibas. Curriculum deepsf. In *European Conference on Computer Vision*, pages 51–67. Springer, 2020. 6
- [15] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 605–613, 2017. 2
- [16] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4857–4866, 2020. 2
- [17] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *IEEE/CVF International Conference on Computer Vision*, pages 7154–7164, 2019. 2
- [18] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *arXiv preprint arXiv:2002.10099*, 2020. 4
- [19] Timothy Jeruzalski, David IW Levin, Alec Jacobson, Paul Lalonde, Mohammad Norouzi, and Andrea Tagliasacchi. Nilbs: Neural inverse linear blend skinning. *arXiv preprint arXiv:2004.05980*, 2020. 2
- [20] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8320–8329, 2018. 2
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4
- [22] 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, 2017. 2
- [23] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM Transactions on Graphics*, 34(6):1–16, 2015. 2
- [24] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987. 4, 6
- [25] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6469–6478, 2020. 2
- [26] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 7
- [27] Qianli Ma, Jinlong Yang, Siyu Tang, and Michael J. Black. The power of points for modeling humans in clothing. In *ICCV*, 2021. 3
- [28] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019. 2, 7

- [29] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit surface representations as layers in neural networks. *IEEE/CVF International Conference on Computer Vision*, 2019. [2](#)
- [30] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [2](#)
- [31] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4d reconstruction by learning particle dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5379–5389, 2019. [2](#)
- [32] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3d deformable shapes. *IEEE/CVF International Conference on Computer Vision*, 2021. [2](#), [5](#), [7](#), [8](#)
- [33] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. [2](#), [4](#), [5](#)
- [34] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *IEEE International Conference on Advanced Video and Signal based Surveillance*, pages 296–301, 2009. [2](#)
- [35] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J Black. Clothcap: Seamless 4d clothing capture and retargeting. *ACM Transactions on Graphics*, 36(4):1–15, 2017. [2](#)
- [36] Gerard Pons-Moll, Javier Romero, Naureen Mahmood, and Michael J Black. Dyna: A model of dynamic human shape in motion. *ACM Transactions on Graphics*, 34(4):1–14, 2015. [2](#)
- [37] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. [2](#)
- [38] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017. [2](#), [6](#)
- [39] Eduard Ramon, Gil Triginer, Janna Escur, Albert Pumarola, Jaime Garcia, Xavier Giro-i Nieto, and Francesc Moreno-Noguer. H3d-net: Few-shot high-fidelity 3d head reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5620–5629, 2021. [2](#)
- [40] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics*, 36(6), 2017. [2](#)
- [41] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. [2](#)
- [42] Igor Santesteban, Miguel A Otaduy, and Dan Casas. Learning-based animation of clothing for virtual try-on. In *Computer Graphics Forum*, pages 355–366, 2019. [2](#)
- [43] Igor Santesteban, Nils Thuerey, Miguel A Otaduy, and Dan Casas. Self-supervised collision handling via generative 3d garment models for virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11763–11773, 2021. [2](#)
- [44] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33, 2020. [4](#), [5](#)
- [45] Garvita Tiwari, Bharat Lal Bhatnagar, Tony Tung, and Gerard Pons-Moll. Sizer: A dataset and model for parsing 3d clothing and learning size sensitive 3d clothing. In *European Conference on Computer Vision*, 2020. [2](#)
- [46] Garvita Tiwari, Nikolaos Sarafianos, Tony Tung, and Gerard Pons-Moll. Neural-gif: Neural generalized implicit functions for animating people in clothing. In *International Conference on Computer Vision (ICCV)*, October 2021. [2](#)
- [47] Tuanfeng Y Wang, Tianjia Shao, Kai Fu, and Niloy J Mitra. Learning an intrinsic garment space for interactive authoring of garment animation. *ACM Transactions on Graphics (TOG)*, 38(6):1–12, 2019. [2](#)
- [48] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6184–6193, 2020. [2](#)
- [49] Meng Zhang, Duygu Ceylan, Tuanfeng Wang, and Niloy J Mitra. Dynamic neural garments. *arXiv preprint arXiv:2102.11811*, 2021. [2](#)
- [50] Silvia Zuffi and Michael J. Black. The stitched puppet: A graphical model of 3D human shape and pose. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR 2015)*, pages 3537–3546, June 2015. [2](#)
- [51] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017. [2](#)