LatentHuman: Shape-and-Pose Disentangled Latent Representation for Human Bodies



Figure 1: **LatentHuman.** We propose a novel representation for the human body which leverages piecewise neural implicit functions to learn shape and pose disentangled latent embeddings. With the fully differentiable and optimizable design as well as the separated latent spaces, LatentHuman can be applied to several 3D learning tasks, including 3D model fitting (model at the top row in second column) and pose tracking (models at the bottom two rows in second column) from sparse input point clouds, shape swapping, pose animation, and motion retargeting.

Abstract

3D representation and reconstruction of human bodies have been studied for a long time in computer vision. Traditional methods rely mostly on parametric statistical linear models, limiting the space of possible bodies to linear combinations. It is only recently that some approaches try to leverage neural implicit representations for human body modeling, and while demonstrating impressive results, they are either limited by representation capability or not physically meaningful and controllable. In this work, we propose a novel neural implicit representation for the human body, which is fully differentiable and optimizable with disentangled shape and pose latent spaces. Contrary to prior work, our representation is designed based on the kinematic model, which makes the representation controllable for tasks like pose animation, while simultaneously allowing the optimization of shape and pose for tasks like 3D fitting and pose tracking. Our model can be trained and fine-tuned directly on non-watertight raw data with welldesigned losses. Experiments demonstrate the improved 3D reconstruction performance over SoTA approaches and show the applicability of our method to shape interpolation, model fitting, pose tracking, and motion retargeting.

1. Introduction

Modeling of 3D human bodies has been studied extensively. The mesh representation is adopted in most of existing methods including traditional parametric models with PCA decomposed shape space and kinematic treebased pose space [29, 38, 42, 49] as well as more sophisticated learning-based models with deformed mesh-based templates [55, 21, 59]. However, the mesh representation is normally limited by the fixed topology.

Recently, the articulated neural implicit representation [14, 35] has been proposed for human bodies. However, these methods either model human poses for a single human shape [14], or rely on the SMPL [29] model shape parameters [35], which limits the representation capability of these models. As a concurrent work, Palafox *et al.* [39] propose to represent the shape and the pose of human bodies with disentangled latent codes, which enable optimization to fit new observations at test time. Nevertheless, their latent pose code encodes a deformation field rather than a physically meaningful human pose and thus is not suitable for controlling human bodies. Moreover, all these methods require watertight 3D models for training which is difficult to generate for real-world scans.

In this paper, we introduce LatentHuman, a neural representation for human bodies with disentangled meaningful pose and shape representations like traditional parametric models, which is fully differentiable and facilitates a vari-

[†] Corresponding author

^{*} Equal contributions

Project webpage: https://latenthuman.github.io/

ety of applications, e.g., 3D model fitting, 3D pose tracking, shape swapping, pose animation, and motion retargeting, etc (see Fig. 1). LatentHuman fully leverages the power of neural implicit functions and represents pose and shape as differentiable and optimizable latent codes. First, in contrast to the methods using one global latent code for the whole body [17, 1], the disentanglement of pose and shape spaces allows us to optimize them separately. Furthermore, the representation is differentiable end-to-end and can be used in deep learning pipelines. This property, together with the embedding of the shape and pose as latent codes, enables tasks which require optimization - either for the shape, the pose, or both - which is critical, e.g., in 3D model fitting and pose tracking. Additionally, we allow the pose to be parameterized by traditional pose parameters, *i.e.*, used by SMPL [29], to make the representation controllable, which is an important property for tasks like motion retargeting to cartoon characters. Lastly, we make use of implicit geometric regularization [17] in order to train our models on raw input point clouds. This allows us to fine-tune our model on unprocessed raw-scan datasets with point-cloud or depth map input without requiring the watertightness property.

However, disentangling the shape and pose factors is hard because the shape of a human body usually deforms with different poses. In order to fix this problem, we guide the learning of pose, shape properties and skeleton joints by exploiting the kinematic model which introduces a constraint on pose and skeleton. More specifically, the VPoser module [42] is adopted to learn the latent pose representation, while a novel VJointer module is proposed for learning the latent skeleton joint representation as one part of the shape representation. By taking kinematic models into consideration, we can generate meaningful human poses and skeletons. In order to model detailed human shapes even with clothes, we utilize the piecewise deformable model proposed by Deng et al. [14] to learn pose-dependent deformations and the other part of the shape representation responsible for shape identities. A novel dual-weighting mechanism is proposed to guarantee smooth and gradual supervision for part connections, which results in a seamless body reconstruction. Furthermore, we adopt similar loss definitions as proposed in Gropp *et al.* [17] and SIREN [53] to learn the disentangled shape and pose representation directly on non-watertight data. Different from the unstructured model used in [17, 53], our piecewise deformable model is more complex and harder to train since one body part model may also influence other parts. In order to prevent the individual human body parts from growing into neighbouring parts, we propose a novel one-sided loss that exploits the nature of body part placement and guides the piecewise functions to focus on their own body part.

We demonstrate the viability of LatentHuman on 3D learning tasks, including human body representation, shape

interpolation, 3D model fitting, 3D pose tracking, motion retargeting, and fine-tuning on partial raw scans. With the learnt latent spaces for shape and pose, LatentHuman is able to effectively produce valid shape predictions or pose configurations for interpolation while being readily optimizable for model fitting and pose tracking.

The main contributions of this paper can be summarized as follows:

- We propose LatentHuman, a novel shape-posedisentangled representation for human bodies. Our representation uses a neural implicit function per body part and encapsulates pose and shape in latent embeddings based on the kinematic model, so the pose is controllable through SMPL-based pose parameters.
- Our representation is fully differentiable end-to-end and can be optimized during inference in deep learning pipelines.
- We propose a novel dual-weighting mechanism to guarantee a smooth piecewise deformable model, and a novel one-sided loss to enable training directly on nonwatertight data.
- Extensive experiments show that LatentHuman has better representation ability and robustness compared to existing methods.

2. Related work

Neural Implicit Representations. Recent neural implicit representations, *i.e.*, deep coordinate-based MLP networks, have been proposed for 3D modeling as a lightweight alternative to traditional 3D representations like surfel based [51, 58], mesh-based [3, 4] and voxel-based methods [12, 37, 16]. While original pioneering works were limited to small-scale objects with a lack of details [40, 32, 9, 34], follow-up works have addressed the issues with self-supervision [2, 1, 17], open surfaces [11], room-scale reconstruction [7, 44, 22, 28, 33], training procedure [15, 52] or high-frequency details [53].

Classical Parametric Models. In contrast to generic 3D shapes, articulated shapes have traditionally been learned from registered meshes, accomplished with a skinning algorithm that deforms vertices of a mesh surface as the joints of an underlying skeleton change [20, 57, 24], outperforming the straight-forward and widely used linear blend skinning (LBS) algorithm. Loper *et al.* [29] then improved significantly upon prior works by introducing a skinned multi-person linear model (SMPL). While SMPL only models the human body, related or follow-up works have focused on other shape parts like faces [5, 13, 27, 43, 46] and hands [49], a combination of those [23, 42, 49], animals [62] or improving on issues, *e.g.*, long-distance interdependencies across the skeleton [38].

Learnt Mesh-based Articulated Representations. Recent works have started to adopt deep learning-based approaches for modeling articulated shapes. Some works focus on using template meshes with auto-encoder like structures [59, 55, 60, 21]. Such methods have the advantage of always predicting topologically correct outputs but are restricted to the predefined topology of the template. Other works based on NeRF [36] rely on meshes for novel view synthesis [45] or rendering of avatars [48], although the work from Peng et al. [45] can be seen as a hybrid: On one hand, using a mesh for anchoring of latent features and on the hand using those features as input for neural implicit functions. Our method focuses on human shape representation from a geometric perspective as we specifically model various human shapes and enable shape-related tasks such as 3D model fitting, pose tracking and motion retargeting.

Articulated Neural Implicit Representations. While some early neural implicit representation methods demonstrated compelling results on human body scans [1, 17, 10], Deng *et al.* [14] pointed out that a straightforward application of auto-decoders for encoding posed shapes into one global latent code [17, 1] is often not enough to cope with the large variety of possible human body poses. Instead they propose NASA, a piecewise deformable model which naturally handles large deformations of body poses. They use joint transformation matrices of a kinematic tree to warp input query points prior to running them through a neural implicit function network. Notably, they use one network per joint to learn pose-dependent, canonical body part shapes.

Subsequent works can be distinguished by the the way they warp between a canonical shape space and a posed shape space. The concurrent work SNARF [8], which learns a forward skinning weight field, improves upon NASA, especially for out-of-distribution examples. They find a many-to-one mapping of possible canonical point candidates for a given point in the deformed space through iterative optimization of the implicit forward skinning equation. However, for both methods, their model is conditioned on poses only and tied to the same shape identity, needing retraining for new subjects. Different from NASA and SNARF, we aim to represent human bodies of arbitrary poses and shapes. SCANimate [50] omits per-joint implicit functions and rather trains one big neural implicit decoder for representing the shape in canonical space. In addition, they learn forward and inverse skinning networks with cycle-consistency constraints which they use for posing the shape. The idea of learning skinning weights was also simultaneously suggested by Mihajlovic et al. [35]. They propose LEAP, a method for learning the articulated occupancy of human bodies. They build upon SMPL [29] and provide efficient occupancy checks on query points. Query points are mapped to the canonical space through a trained inverse LBS network which learns valid skinning weights

		Test	Shape Optim.	Pose		Non-	Raw	Mutli.
Method	Out	Input		Optim.	Contr.	WaTi.	Scan	Shape
NASA [14]	Occ.	θ		\checkmark	\checkmark			
LEAP [35]	Occ.	β , $ heta$, \hat{V}			\checkmark		\checkmark	\checkmark
SCANimate [50]	SDF	θ			\checkmark	\checkmark	\checkmark	
NPMs [39]	SDF	$\mathbf{z}_s, \mathbf{z}_p$	\checkmark	\checkmark			\checkmark	\checkmark
Ours	SDF	\mathbf{z}_s, θ or \mathbf{z}_p	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

Table 1: **Overview of implicit representations for human bodies.** We denote whether shape and pose parameters are optimizable (*e.g.*, for pose-tracking) and/or controllable (*e.g.*, for animation) and compare whether the model can be trained on non-watertight data, generalize to raw scans and be reused for multiple subjects without re-training. (β , θ): SMPL-based shape/pose parameters, ($\mathbf{z}_s, \mathbf{z}_p$): Learnt shape/pose latent codes, \hat{V} : Posed SMPL vertices.

for every query point, even non-manifold points. However, they heavily rely on SMPL shape parameters which restricts the representation capability. In contrast, our method learns the shape latent space from scratch.

The concurrent work NPMs [39] also learns to disentangle shape and pose in a fully implicit manner with several latent codes. Their shape latent space resides in the canonical space while their learned deformation field maps to the posed space. In contrast to NPMs, our representation additionally allows the use of well-known SMPL pose parameters to control the pose latent code which is useful for animation tasks. Further, while NPMs need supplementary initialization networks and depth-projected SDF grids to infer and optimize the pose and shape latent codes, we instead rely on implicit geometric regularization [17] allowing us to work with raw, non-watertight point clouds at training time. A summary of recent related works is shown in Table 1.

3. Method

We introduce LatentHuman, a fully differentiable and disentangled implicit representation for human bodies. Following prior works [40, 2, 17, 53], we express the body shape through signed distance functions (SDF) with feedforward neural networks. As shown in Fig. 2, we consider two spaces which influence the posed human body model and represent both with latent codes respectively: The pose code controls the joint movement of the body while the shape code determines the canonical skeleton joints and specific body shapes (Sec 3.1). For a given query point x_i , LatentHuman predicts per-joint canonical SDF values with a series of piecewise implicit functions, which are conditioned on shape and pose-dependent deformation features (Sec 3.2). Motivated by Gropp et al. [17], we adopt implicit geometric regularization and related auxiliary loss functions to enable training and fine-tuning on non-watertight raw scans. To make the training procedure stable and generate seamless body shapes, we introduce a novel one-sided non-manifold loss and a dual-weighting strategy (Sec 3.3).



Figure 2: **Overview.** For learning SDF functions per body part, we use a piecewise deformable model [14] conditioned on a shape code \mathbf{z}_s , a per-joint feature \mathbf{z}_p^b describing pose-dependent deformations and canonical query points \mathbf{x}^i . The bone transformations needed for obtaining \mathbf{z}_p^b can be computed with SMPL pose joint rotations and skeleton joints in canonical space (*i.e.* T-pose). For the former, we adopt VPoser [42] while for the latter, we introduce our novel VJointer module. The per-joint SDF predictions are combined with a SoftMin function to obtain the SDF values of the final mesh.

3.1. Disentanglement of Pose and Shape

As the human shape normally deforms with different poses, we exploit the kinematic model and prior knowledge to learn the disentanglement of human pose and shape.

Pose Space. Ideally, a human pose encoding should allow valid poses only, be human-readable and reusable in downstream tasks. Latent embeddings [40, 39], in contrast to kinematic tree skeletons, are usually not human-readable and not friendly for downstream tasks like motion animation. To this end, we adopt the kinematic pose prior VPoser [42] as our pose space decoder and use the joint rotations (*i.e.*, $\theta \in \mathbb{R}^{B \times 3}$, where *B* is the number of body parts) of the kinematic tree [29, 18] as output in order to compute bone transformations used by our piecewise deformable model. This allows us to either use SMPL pose parameters for rigging human bodies or optimizing the pose in latent space, which yields a more robust performance on downstream tasks like pose tracking.

Shape Space. Unlike the human pose space, the shape space needs to store more information since human shapes can be more expressive especially when dealing with clothed bodies. Meanwhile, it is known that the human shape structure, including overall height, length of body parts, and relative positioning of body parts with respect to each other, follows the kinematic model and is independent of surface-related details. If a single latent code is used to encode the whole shape, the human shape structure will be deeply intertwined with other shape properties, e.g., the human's body weight index or other surfacerelated properties like wrinkles, and thus we cannot learn it separately. For this reason, we split the shape encoding into two components, one for representing the joints of the skeleton in canonical space, *i.e.*, the T-pose, and one for representing the remaining shape properties. We assume that the distribution of canonical skeleton joints

lies on a lower-dimensional subspace and therefore, inspired by VPoser [42], we propose a data-driven kinematic prior named VJointer, which learns a human joint prior with a variational autoencoder (VAE) [26] on the SUR-REAL dataset [56]. By introducing VJointer, we can benefit from regularized optimization of skeleton joints in a 32dimensional latent space. For the surface-related shape features, we use a 64-dimensional learnable shape code \mathbf{z}_s to encode different human body identities.

3.2. Implicit Human Body Representation

Inspired by NASA [14], we use a series of piecewise implicit functions to encode the human body. Different from NASA that adopts occupancy functions, we represent human bodies with implicit SDF functions which encode the distance to the closest body surface for a 3D query point x, with the sign indicating if the point is inside (negative) or outside (positive) the body. Each SDF function outputs a body part SDF value conditioned on a pose-dependent deformation feature \mathbf{z}_p^b and the corresponding shape code \mathbf{z}_s . All predicted part SDF values are finally blended to obtain the overall SDF value for the whole body.

Pose-dependent Deformation. To encode the posedependent deformation feature \mathbf{z}_p^b , we follow NASA [14] and transform the root bone location \mathbf{t}_0 to the per-bone local spaces and project to a part-specific feature vector with a learnable linear layer $\Pi_b : \mathbb{R}^{B \times 3} \to \mathbb{R}^{32}$:

$$\mathbf{z}_{p}^{b} = \Pi_{b} \left(\oplus_{b}^{B} \mathbf{B}_{b}^{-1} \mathbf{t}_{0} \right), \tag{1}$$

where \oplus concatenates each transformed root.

Piecewise Model. We consider *B* body parts and for each employ a small coordinate-based MLP conditioned on the pose-dependent deformation code \mathbf{z}_p^b and the shape latent code \mathbf{z}_s . We transform query points \mathbf{x}_i to the canonical T-pose according to the per-bone transformations \mathbf{B}_b^{-1} and

use those canonical query coordinates $\hat{\mathbf{x}}_i = \mathbf{B}_b^{-1} \mathbf{x}_i$ as input for each part MLP. To further improve high-frequency surface details, we also apply a positional encoding $\gamma(\cdot)$ [54] to the canonical input coordinates and adopt a progressive training scheme [41] to ensure a stable convergence. For each implicit function, we output SDF values for the respective body part $b \in \{1, ..., B\}$:

$$\hat{s}_b^i = \text{SDF}_b(\gamma(\hat{\mathbf{x}}_i) | \mathbf{z}_p^b, \mathbf{z}_s).$$
(2)

The overall SDF value \hat{s}_i can then be obtained as:

$$\hat{s}_i = \text{SoftMin}_b^B(\lambda_b \hat{s}_b^i), \tag{3}$$

where λ_b controls the blending weight of each body part. We empirically set $\lambda_b = 50$, but a larger λ_b also works well in our experiments.

3.3. Non-rigid Geometric Supervision

To make our representation trainable on non-watertight raw scans, we employ a series of supervision strategies.

Supervision Sampling. For each ground-truth body mesh, we sample a set of body part-weighted query points $\{\mathbf{x}_i\}^{2M+Q}$ consisting of M points on the surface, M near-surface points and Q random points sampled within the bounding boxes of each body part. For our experiments, we use M = 4000 and Q = 800. As our training dataset consists of organized SMPL models, we use the provided skinning weights to assign to each query point the part label with the highest skinning weight. Please refer to the supplementary material for more details.

Geometric Supervision. Recent works [14, 35] mainly use densely supervised data to regress occupancy values, which requires a watertight mesh to ensure a successful ground-truth occupancy sampling. To achieve the supervision directly on partial raw scan data, we employ implicit geometric regularization as introduced by Gropp *et al.* [17] and Sitzmann *et al.* [53] by minimizing the following loss:

$$\mathcal{L}_{geom} = \lambda_{m} \mathcal{L}_{m} + \lambda_{n} \mathcal{L}_{n} + \lambda_{e} \mathcal{L}_{e} + \lambda_{nm} \mathcal{L}_{nm}, \qquad (4)$$

$$\mathcal{L}_{\mathrm{m}} = \sum_{\mathbf{x}_{i} \in \Omega_{0}} \phi(\mathbf{x}_{i}), \quad \phi(\mathbf{x}_{i}) = |\mathrm{SDF}(\mathbf{x}_{i})|, \quad (5)$$

$$\mathcal{L}_{n} = \sum_{\mathbf{x}_{i} \in \Omega_{0}} (1 - \langle \nabla_{\mathbf{x}_{i}} \text{SDF}(\mathbf{x}_{i}), \mathbf{n}(\mathbf{x}_{i}) \rangle), \qquad (6)$$

$$\mathcal{L}_{e} = \sum_{\mathbf{x}_{i} \in \Omega} \left\| \left\| \nabla_{\mathbf{x}_{i}} \text{SDF}(\mathbf{x}_{i}) \right\|_{2} - 1 \right\|, \tag{7}$$

$$\mathcal{L}_{nm} = \sum_{\mathbf{x} \in \Omega \setminus \Omega_0} \exp(-\alpha \cdot |SDF(\mathbf{x})|), \tag{8}$$

where the manifold loss \mathcal{L}_m enforces points on the surface to yield zero SDF value, with Ω_0 denoting the zero-level set. \mathcal{L}_n is the normal loss that constrains the on-surface points to have a consistent gradient with the ground-truth surface normals $\mathbf{n}(\mathbf{x})$. The eikonal loss \mathcal{L}_{e} regularizes the norm of the spatial gradients to 1, where Ω denotes the whole domain of the sampling. \mathcal{L}_{nm} is the non-manifold loss which penalizes close-to-zero SDF values at off-surface points, where α is a hyper parameter and we set $\alpha = 5$. In practice, we add the above losses to both the SDF output of each piecewise SDF function (Eq. 2) and the overall blended SDF values (Eq. 3).

One-sided Non-manifold Loss. The aforementioned geometric supervision works for rigid shape learning [17, 53], but we find it does not ensure good convergence in our piecewise non-rigid case. This is due to the lack of supervision signals on overlapping intersection areas at part connections, *i.e.*, the outputs of the piecewise functions are unconstrained, which might lead one body part to extrapolate surface into neighboring parts. We therefore introduce a one-sided non-manifold loss, through which we can penalize the prediction of negative (inside) SDF values for onsurface sampled points not belonging to the current body part:

$$\mathcal{L}_{\text{osnm}} = \sum_{\mathbf{x}_i \in \Omega_0^{j \neq b}} \text{SoftPlus}_{\beta} \left(\frac{-\text{SDF}_b(\mathbf{x}_i) + \delta}{2\delta} \right), \quad (9)$$

where $\delta = 0.01$ denotes a truncation distance and $\beta = 10$ is the SoftPlus parameter. We use SoftPlus rather than ReLU in order to allow smooth gradient flows.

Dual-weighting. Although the aforementioned part supervision enforces each SDF function to focus on their own body parts, it will inevitably lead to joint artifacts near the part connection areas (see Fig. 8). Therefore, we introduce a dual-weighting mechanism by exploiting the skinning weights of the SMPL model to eliminate those joint artifacts and facilitate seamless body reconstruction. More specifically, for each surface sampling point \mathbf{x}_i , we find two body part labels with the highest skinning weights w_i^0 and w_i^1 . We can then rewrite the part manifold loss in a dual-weighting manner:

$$\mathcal{L}_{\text{dual-m}} = \sum_{\mathbf{x} \in \Omega_0^{j \neq b^0}} w_i^0 \phi(\mathbf{x}_i) + \sum_{\mathbf{x} \in \Omega_0^{j \neq b^1}} w_i^1 \phi(\mathbf{x}_i), \quad (10)$$

Intuitively, the highest and the second-highest skinning weights w_i^0, w_i^1 transition smoothly near the connection areas, and the dual-weighting mechanism takes advantage of this trait to bring gradual supervision to these areas.

Latent Regularization. Similar to DeepSDF [40], we add a regularization term to the learnable shape codes to further constrain the learning problem:

$$\mathcal{L}_{\rm zs} = \|\mathbf{z}_s\|_2^2. \tag{11}$$

Methods		AMASS / DF	aust	AMASS / MoVi			
	IoU ↑	$Chamfer \downarrow$	F-Score ↑	$\mathrm{IoU}\uparrow$	$Chamfer \downarrow$	F-Score \uparrow	
NASA [14]	87.67	0.00719	80.95	84.08	0.00885	73.36	
LEAP [35]	96.09	0.00333	98.24	94.72	0.00352	98.23	
Ours	96.45	0.00304	99.06	95.76	0.00314	98.90	

Table 2: **Quantitative reconstruction results.** Our representation outperforms SoTA approaches [14, 35] on all metrics.

Our final loss is then defined as

$$\mathcal{L}_{\text{total}} = \lambda_{\text{dual-m}} \mathcal{L}_{\text{dual-m}} + \lambda_{n} \mathcal{L}_{n} + \lambda_{e} \mathcal{L}_{e} + \lambda_{nm} \mathcal{L}_{nm} + \lambda_{\text{osnm}} \mathcal{L}_{\text{osnm}} + \lambda_{zs} \mathcal{L}_{zs}.$$
(12)

The values for the loss weights λ can be found in the supplementary material.

3.4. Implementation Details

For each part decoder, we use 4 fully-connected layers with 64 neurons. We initialize the network weights with the geometric initialization introduced by Atzmon *et al.* [2]. Additionally we rely on weight normalization. We use the Adam optimizer [25] with an initial learning rate of 0.0001, and divide the learning rate by half every 500 epochs.

4. Experiments

We evaluate our proposed representation on various tasks, including human body representation (Sec. 4.2), shape interpolation (Sec. 4.3), model fitting (Sec. 4.4), human pose tracking and retargeting (Sec. 4.5). Furthermore, we also demonstrate our capability of fine-tuning with partial raw scan data (Sec. 4.6). At last, we perform ablation studies (Sec. 4.7) to analyze the efficacy of the proposed non-rigid geometric supervision and the latent optimization.

4.1. Dataset

We use the DFaust and MoVi subsets from the AMASS [31] dataset to evaluate the model representation, and follow the split from LEAP [35] for training and testing. For the AMASS / DFaust subset, we use training sequences of 10 subjects and randomly keep 1 sequence per subject for testing. For the AMASS / MoVi subset, we select every 10-th subject (which yields 9 unseen subjects) for testing and use the remaining 76 subjects for training in order to evaluate the generalization ability to unseen shapes. To further evaluate our capability of dealing with non-watertight partial raw scan data, we fine-tune on DFaust [6] body scans and the CAPE [30, 47] clothed human dataset, and randomly select two subjects from each dataset for evaluation. Please refer to the supplementary material for more details.

4.2. Human Body Representation

To evaluate our representation on human bodies, we compare our method with NASA [14] and LEAP [35]¹ on



Figure 3: **Qualitative comparison.** We compare our representation qualitatively with NASA [14] and LEAP [35] on the DFaust and MoVi subsets of AMASS [31]. Our method preserves more shape details than NASA. Compared with LEAP, our method mitigates artifacts, especially in cases where different body parts have great proximity (*e.g.* hand and head on the far right). Best viewed digitally.

AMASS / DFaust and AMASS / MoVi data. We follow LEAP [35] and use the same setup for the competitors. For our method, it is worth noting that, when evaluating on the unseen subjects of AMASS / MoVi data, we only optimize the shape code for the first frame of each sequence and reconstruct the following frames with the fixed shape code. For the AMASS / DFaust data, our method directly uses the learnt shape codes of the specific subjects from the training process. For measuring the representation quality, we report the same metrics used in NASA [14], including IoU (%), Chamfer-L1 (m) and F-Score (%).

As shown in Table 2, our representation consistently outperforms NASA and LEAP on all metrics. We also provide qualitative comparisons on the resulting meshes in Fig. 3. NASA pioneered neural implicit representations for articulated shapes like human bodies but only focused on representing one single subject at a time and therefore struggles to represent multiple subjects with the same network. This can be observed in the first row of Fig. 3 where the details for the face and hands are missing. In contrast, our method accurately represents the details of different subjects.

The recently proposed LEAP also shows promising results for the multi-subject case, but suffers from the limitation of the inverse LBS mechanism. In cases where two body parts are close (*e.g.*, sitting with cross-legs or clapping hands), their reconstruction results show concavities or some parts (*e.g.*, fingers) are even missing. We believe that in such situations, the inverse LBS network fails to accurately map a query point to the correct canonical space, which inevitably distorts the occupancy representation as

¹NPMs [39] is a concurrent work without public code available and SCANimate [50] aims to align raw scans of a particular person to learn an

animatable avatar, so they are not adopted for comparison.



Figure 4: **Shape code interpolation.** We linearly interpolate the shape code of two examples from the AMASS / DFaust dataset (top row) and the AMASS / MoVi dataset (bottom row) respectively.



Figure 5: **Model fitting.** We fit our representation to input points of human bodies in a casual A-pose by jointly optimizing the shape and pose. 1000 points were used as input.

discussed in their paper. In comparison, our method directly expresses the SDF values in the posed space rather than learning an inverse mapping to the canonical space and thus shows less susceptibility to such cases where body parts are close by, leading to fewer artifacts. Note that LEAP requires much more input data (*e.g.* posed SMPL mesh vertices) for each reconstruction while our representation only depends on a learnt shape code for each human identity.

4.3. Shape Interpolation

In order to demonstrate that the learnt shape representation is meaningful and continuous in latent space, we perform a shape code interpolation experiment. Specifically, we choose two shape codes from Sec. 4.2 and interpolate linearly between them while having the pose fixed. As visualized in Fig. 4, the body shape transition is fairly smooth, which demonstrates that our shape code learns a meaningful and continuous representation of the human body.

4.4. Model Fitting

In this experiment, we demonstrate how our network can also be applied to model fitting tasks, *e.g.*, fitting our representation to a given point cloud of human bodies in the casual A-pose, which is commonly used when capturing human bodies. For the experimental setup, previously unseen body shapes from the SURREAL [56] dataset are randomly selected, combined with casual A-poses from the AMASS / DFaust dataset [31] and uniformly sampled resulting in 1000 points on the surface. For the fitting, we first initialize the shape code with a random normal distribution



Figure 6: **Qualitative pose tracking results.** We show the input sparse points (first row), tracked human bodies with our reconstruction (second row), and the retargeted characters with the tracked poses (third row).

# Sampled Points	4000	2000	1000	500				
	IoU (%) \uparrow / MPJPE (m) \downarrow							
NASA [14]	92.99 / 0.0150	92.86 / 0.0156	92.59 / 0.0162	92.12 / 0.0178				
Ours	95.88 / 0.0049	95.77 / 0.0051	95.53 / 0.0057	95.20 / 0.0066				

Table 3: **Quantitative pose tracking results.** We report IoU and mean per-joint position error (MPJPE [19]) on four different point densities.

and the pose code with the standard A-pose encoded by VPoser [42]. We then jointly optimize both, the shape and pose code, by minimizing a single overall manifold loss. As shown in Fig. 5, the reconstructed human bodies are intact and close to the ground-truth, which demonstrates the potential usefulness of our proposed representation for real-world applications. Once the disentangled shape and pose codes are recovered, we can easily swap the shape to other identities by exchanging the shape code and animate the human body with novel poses as shown in Fig. 1.

4.5. Human Pose Tracking

As our representation is fully differentiable, it is readily available for human pose tracking on point clouds. We choose three sequences from the AMASS / DFaust subset, and provide a comparison between our method and NASA with the pre-trained model from Sec. 4.2. For our method, we fix the learned shape code from Sec. 4.2 and optimize only the pose code directly within the latent space. It is worth noting that we do not add any temporal or prior constraints like NASA does, but instead only use a single overall manifold loss and initialize the pose code with the optimized code from the previously tracked frame. We evaluate the tracking performance on different point sampling densities, *i.e.*, from 4000 to 500, and also follow NASA's setup by adding normal noise with $\sigma = 0.005$ to the points. The results are shown in Table 3, where we report the IoU (%)and MPJPE (m) (mean per joint position error [19]) metrics. Compared to NASA, we consistently achieve the best IoU and MPJPE in all settings. We also visualize our tracked shapes and show motion retargeting to cartoon characters as an application in Fig. 6, demonstrating the meaningfulness and usefulness of the tracked poses. Please refer to our



Figure 7: **Fine-tuning on raw scans.** Our non-rigid geometric supervision formulation allows us to fine-tune our model on raw partial scans. We are able to better preserve details than SMPL [29] on raw DFaust [6] scans and even can model clothed human from the CAPE [30, 47] dataset.

Methods		AMASS / DF	aust	AMASS / MoVi		
methods	IoU ↑	$Chamfer \downarrow$	F-Score ↑	IoU↑	$Chamfer \downarrow$	F-Score ↑
w/o Dual-Weighting	96.31	0.00310	98.96	94.29	0.00532	96.34
w/o One-Sided Loss	×	×	×	×	×	×
w/o VJoints	/	/	/	89.44	0.04796	91.45
Full Model	96.45	0.00304	99.06	95.76	0.00314	98.90

Table 4: **Ablation study.** We ablate our dual-weighting strategy, the one-sided manifold loss and our proposed VJointer module. All proposed components are necessary to achieve the best performance.

supplementary material for more results.

4.6. Evaluation on Partial Raw Scan

To further inspect the capability of our representation to handle partial raw scan data, we fine-tune the pre-trained model (from Sec. 4.2) on DFaust [6] body scans and the CAPE [30, 47] clothed human dataset. Thanks to our nonrigid supervision, we are able to learn the shape naturally from these partial scans. As shown in Fig. 7, our method better represents details and captures more realistic shapes of human bodies than SMPL on DFaust body scans (e.g., our head shape is closer to the ground-truth), and also achieves better one-sided Chamfer distance D_{s2m} (*i.e.* from scan to the resulting mesh following [17]) than SMPL. More importantly, our representation is able to successfully represent clothed human bodies on the CAPE dataset, even though the accessible training data of CAPE is limited, *i.e.*, only 2 or 3 sequences are available for a few of released subjects. Please refer to the supplementary video for a vivid animation of these results.

4.7. Ablation Study

Dual-weighting. We analyze the efficacy of the dual-weighting mechanism by ablating it during the training process. Even though the proposed strategy only slightly improves the overall metric scores (Table 4), it leads to a significant visual improvement by eliminating boundary artifacts near part connection areas (Fig. 8).

One-sided Loss. To study the impact of the one-sided loss, we remove it from the supervision during training. As



Figure 8: **Dual-weighting.** The dual-weighting strategy significantly reduces visual boundary artifacts at body part connection areas.

# Sampled Points	4000	2000	1000	500					
IoU (%) \uparrow / MPJPE (m) \downarrow									
Ours (w/o Latent) Ours (w Latent)	95.85 / 0.1520 95.88 / 0.0049	94.17 / 0.1762 95.77 / 0.0051	93.28 / 0.2299 95.53 / 0.0057	× 95.20 / 0.0066					

Table 5: **Ablation on latent pose code optimization.** We perform pose tracking with and without latent pose code optimization, *i.e.* instead directly optimizing bone transformations. Optimizing the pose latent space leads to significant improvements.

demonstrated in Table 4, the network does not converge to form a human body due to the lack of polarized SDF supervision as explained in Sec. 3.3, which highlights the importance of this loss for our non-rigid geometric supervision.

VJointer on Shape Code Optimization. Since our experiment on the AMASS / MoVi subset of Sec. 4.2 leverages the network's ability of shape code optimization, we analyze the effectiveness of VJointer by removing it and instead optimizing the $B \times 3$ joint locations directly. As shown in Table 4, the proposed VJointer significantly improves the reconstruction quality by a large margin.

Latent Pose Optimization. We inspect the effectiveness of latent pose optimization on the human pose tracking task of Sec. 4.5 by replacing it with a direct optimization of the bone transformations, *i.e.*, optimizing the rotation with a continuous representation [61]. The tracking results are reported in Table 5. Introducing pose optimization on the latent space significantly enhances the tracking robustness.

5. Conclusion

We introduced LatentHuman, a novel pose-shapedisentangled representation for the human body. We exploit both pose and shape priors by using the kinematic model for guiding the network to disentangle the shape and pose spaces. Our non-rigid geometric supervision also allows to fine-tune LatentHuman on non-watertight raw scans. Experiments demonstrate that LatentHuman outperforms existing methods for 3D human body representation and can be applied to various tasks like 3D model fitting, shape interpolation, and pose tracking. Although LatentHuman can model the clothed human body, it cannot handle the soft cloth dynamics and we leave this as a future work.

Acknowledgments: This work was partially supported by the NSFC (No. 62102356), Zhejiang Lab (2021PE0AC01) and Innosuisse (Grant No. 34475.1 IP-ICT).

References

- Matan Atzmon and Yaron Lipman. SAL: Sign Agnostic Learning of Shapes From Raw Data. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2565–2574, 2020. 2, 3
- [2] Matan Atzmon and Yaron Lipman. SALD: Sign Agnostic Learning with Derivatives. In *International Conference on Learning Representations*, Sept. 2020. 2, 3, 6
- [3] Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, and Yaser Sheikh. Modeling Facial Geometry Using Compositional VAEs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3877–3886, 2018. 2
- [4] Pierre Baque, Edoardo Remelli, Francois Fleuret, and Pascal Fua. Geodesic Convolutional Shape Optimization. In *International Conference on Machine Learning*, pages 472–481.
 PMLR, July 2018. 2
- [5] Volker Blanz and Thomas Vetter. A Morphable Model for the Synthesis of 3D Faces. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, pages 187–194, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co. 2
- [6] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: registering human bodies in motion. In 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, pages 5573–5582. IEEE Computer Society, 2017. 6, 8
- [7] Rohan Chabra, Jan E. Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep Local Shapes: Learning Local SDF Priors for Detailed 3D Reconstruction. In *Computer Vision – ECCV 2020*, Lecture Notes in Computer Science, pages 608–625, Cham, 2020. Springer International Publishing. 2
- [8] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. arXiv preprint arXiv:2104.03953, 2021. 3
- [9] Zhiqin Chen and Hao Zhang. Learning Implicit Fields for Generative Shape Modeling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5939–5948, Dec. 2018. 2
- [10] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit Functions in Feature Space for 3D Shape Reconstruction and Completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 3
- [11] Julian Chibane, Mohamad Aymen Mir, and Gerard Pons-Moll. Neural Unsigned Distance Fields for Implicit Function Learning. Advances in Neural Information Processing Systems, 33:21638–21652, 2020. 2
- [12] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings* of the 23rd Annual Conference on Computer Graphics and Interactive Techniques, pages 303–312. ACM, 1996. 2
- [13] Hang Dai, Nick Pears, William A. P. Smith, and Christian Duncan. A 3D Morphable Model of Craniofacial Shape and

Texture Variation. In *Proceedings of the IEEE International* Conference on Computer Vision, pages 3085–3093, 2017. 2

- [14] Boyang Deng, J. P. Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. NASA Neural Articulated Shape Approximation. In *Computer Vision – ECCV 2020*, 2020. 1, 2, 3, 4, 5, 6, 7
- [15] Yueqi Duan, Haidong Zhu, He Wang, Li Yi, Ram Nevatia, and Leonidas J. Guibas. Curriculum DeepSDF. In *Computer Vision – ECCV 2020*, 2020. 2
- [16] Simon Fuhrmann and Michael Goesele. Fusion of Depth Maps with Multiple Scales. In *Proceedings of the 2011 SIG-GRAPH Asia Conference*, SA '11, pages 148:1–148:8, New York, NY, USA, 2011. ACM. 2
- [17] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems* 2020, pages 3569–3579, 2020. 2, 3, 5, 8
- [18] David A Hirshberg, Matthew Loper, Eric Rachlin, and Michael J Black. Coregistration: Simultaneous alignment and modeling of articulated 3d shape. In *European conference on computer vision*, pages 242–255. Springer, 2012. 4
- [19] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 7
- [20] Alec Jacobson and Yotam Gingold. Skinning: Real-time shape deformation. In *SIGGRAPH Asia 2014 Courses*, SA '14, page 1, New York, NY, USA, Nov. 2014. Association for Computing Machinery. 2
- [21] Boyi Jiang, Juyong Zhang, Jianfei Cai, and Jianmin Zheng. Disentangled Human Body Embedding Based on Deep Hierarchical Neural Network. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–1, 2020. 1, 3
- [22] Chiyu "Max" Jiang, Avneesh Sud, Ameesh Makadia, Jingwei Huang, Matthias Niessner, and Thomas Funkhouser. Local Implicit Grid Representations for 3D Scenes. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2020. 2
- [23] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total Capture: A 3D Deformation Model for Tracking Faces, Hands, and Bodies. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8320–8329, 2018. 2
- [24] Ladislav Kavan and Jiří Žára. Spherical blend skinning: A real-time deformation of articulated models. In *Proceedings of the 2005 Symposium on Interactive 3D Graphics and Games*, I3D '05, pages 9–16, New York, NY, USA, Apr. 2005. Association for Computing Machinery. 2
- [25] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2015. 6
- [26] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun, editors, 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, 2014. 4

- [27] 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. 2
- [28] Sandro Lombardi, Martin R. Oswald, and Marc Pollefeys. Scalable Point Cloud-based Reconstruction with Local Implicit Functions. In 2020 International Conference on 3D Vision (3DV), pages 997–1007, Nov. 2020. 2
- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A Skinned Multiperson Linear Model. ACM Trans. Graph., 34(6):248:1– 248:16, Oct. 2015. 1, 2, 3, 4, 8
- [30] 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 *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6469–6478, 2020. 6, 8
- [31] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of Motion Capture As Surface Shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5442–5451, 2019. 6, 7
- [32] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy Networks: Learning 3D Reconstruction in Function Space. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 4460–4470, 2019. 2
- [33] Zhenxing Mi, Yiming Luo, and Wenbing Tao. Ssrnet: scalable 3d surface reconstruction network. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 970–979, 2020. 2
- [34] Mateusz Michalkiewicz, Jhony K. Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Deep Level Sets: Implicit Surface Representations for 3D Shape Inference. arXiv:1901.06802 [cs], Jan. 2019. 2
- [35] Marko Mihajlovic, Yan Zhang, Michael J. Black, and Siyu Tang. LEAP: Learning Articulated Occupancy of People. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10461–10471, 2021. 1, 3, 5, 6
- [36] 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. In *Computer Vision – ECCV 2020*, 2020. 3
- [37] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-time dense surface mapping and tracking. In *Mixed and Augmented Reality (ISMAR)*, 2011 10th IEEE International Symposium On, pages 127– 136. IEEE, 2011. 2
- [38] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. STAR: sparse trained articulated human body regressor. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part VI,* volume 12351 of *Lecture Notes in Computer Science*, pages 598–613, 2020. 1, 2

- [39] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. NPMs: Neural Parametric Models for 3D Deformable Shapes. arXiv:2104.00702 [cs], page 2104.00702, Apr. 2021. 1, 3, 4, 6
- [40] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 2, 3, 4, 5
- [41] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B. Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. *CoRR*, abs/2011.12948, 2020. 5
- [42] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive Body Capture: 3D Hands, Face, and Body From a Single Image. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019. 1, 2, 4, 7
- [43] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3D Face Model for Pose and Illumination Invariant Face Recognition. In 2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance, pages 296–301, Sept. 2009. 2
- [44] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 523–540. Springer, 2020. 2
- [45] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural Body: Implicit Neural Representations With Structured Latent Codes for Novel View Synthesis of Dynamic Humans. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021. 3
- [46] Stylianos Ploumpis, Haoyang Wang, Nick Pears, William A. P. Smith, and Stefanos Zafeiriou. Combining 3D Morphable Models: A Large Scale Face-And-Head Model. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10934–10943, 2019. 2
- [47] Gerard Pons-Moll, Sergi Pujades, Sonny Hu, and Michael J. Black. Clothcap: seamless 4d clothing capture and retargeting. ACM Trans. Graph., 36(4):73:1–73:15, 2017. 6, 8
- [48] Sergey Prokudin, Michael J Black, and Javier Romero. Smplpix: Neural avatars from 3d human models. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pages 1810–1819, 2021. 3
- [49] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. ACM Transactions on Graphics, 36(6):245:1– 245:17, Nov. 2017. 1, 2
- [50] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly Supervised Learning of Skinned Clothed Avatar Networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2886–2897, 2021. 3, 6

- [51] Thomas Schöps, Torsten Sattler, and Marc Pollefeys. SurfelMeshing: Online Surfel-Based Mesh Reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–1, 2019. 2
- [52] Vincent Sitzmann, Eric R. Chan, Richard Tucker, Noah Snavely, and Gordon Wetzstein. Metasdf: Meta-learning signed distance functions. In *Proc. NeurIPS*, 2020. 2
- [53] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 2, 3, 5
- [54] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. Advances in Neural Information Processing Systems, 33:7537–7547, 2020. 5
- [55] Edgar Tretschk, Ayush Tewari, Michael Zollhöfer, Vladislav Golyanik, and Christian Theobalt. Demea: Deep mesh autoencoders for non-rigidly deforming objects. In *European Conference on Computer Vision*, pages 601–617. Springer, 2020. 1, 3
- [56] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 109–117, 2017. 4, 7
- [57] Xiaohuan Corina Wang and Cary Phillips. Multi-weight enveloping: Least-squares approximation techniques for skin animation. In *Proceedings of the 2002 ACM SIG-GRAPH/Eurographics Symposium on Computer Animation*, SCA '02, pages 129–138, New York, NY, USA, July 2002. Association for Computing Machinery. 2
- [58] Thomas Whelan, Renato F. Salas-Moreno, Ben Glocker, Andrew J. Davison, and Stefan Leutenegger. ElasticFusion: Real-time dense SLAM and light source estimation. *The International Journal of Robotics Research*, 35(14):1697– 1716, 2016. 2
- [59] 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 Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6184–6193, 2020. 1, 3
- [60] Keyang Zhou, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Unsupervised Shape and Pose Disentanglement for 3D Meshes. In *Computer Vision – ECCV 2020*, 2020. 3
- [61] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 5745– 5753, 2019. 8
- [62] 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. 2