

Learning Compositional Representation for 4D Captures with Neural ODE

Boyan Jiang^{1*} Yinda Zhang^{2*} Xingkui Wei¹ Xiangyang Xue¹ Yanwei Fu¹
¹ Fudan University ² Google

Abstract

Learning based representation has become the key to the success of many computer vision systems. While many 3D representations have been proposed, it is still an unaddressed problem how to represent a dynamically changing 3D object. In this paper, we introduce a compositional representation for 4D captures, i.e. a deforming 3D object over a temporal span, that disentangles shape, initial state, and motion respectively. Each component is represented by a latent code via a trained encoder. To model the motion, a neural Ordinary Differential Equation (ODE) is trained to update the initial state conditioned on the learned motion code, and a decoder takes the shape code and the updated state code to reconstruct the 3D model at each time stamp. To this end, we propose an Identity Exchange Training (IET) strategy to encourage the network to learn effectively decoupling each component. Extensive experiments demonstrate that the proposed method outperforms existing state-of-the-art deep learning based methods on 4D reconstruction, and significantly improves on various tasks, including motion transfer and completion.

1. Introduction

Shape representation is one of the core topics in 3D computer vision, especially in the era of deep learning. Early work uses explicit representation, e.g. volume [13, 17, 58], point cloud [46, 16, 45, 1], and mesh [18, 23, 57] for 3D related tasks, such as shape reconstruction, synthesis, and completion. Recently, deep implicit representation [33, 40, 22] shows promising performance in producing accurate geometry with appealing surface details. However, arguably, we, humans, stay in a 3D world with an additional temporal dimension, and the majority of data we perceive everyday are moving or deforming 3D objects and scenes. Many existing applications also require understand-

* indicates equal contributions.

Boyan Jiang Xingkui Wei and Xiangyang Xue are with the School of Computer Science, Fudan University.

Yanwei Fu is with the School of Data Science, MOE Frontiers Center for Brain Science, and Shanghai Key Lab of Intelligent Information Processing, Fudan University.




Figure 1: We present a compositional representation for 4D object dynamics, through which the input point cloud sequence is disentangled into semantically meaningful representations in three latent spaces for geometry template, initial state, and temporal deformation.

ing or reconstruction of 4D data, such as autonomous driving, robotics, and virtual or augmented reality. But the deep representation for 4D data, i.e. a deforming 3D object over a time span, is barely missing in the literature. As a pioneer work, Niemeyer *et al.* [37] propose to predict velocity field of the 3D motion via a Neural ODE [11]. However, the method mainly focuses on recovering and integrating local flow for 4D reconstruction, which might accumulate error and thus produce sub-optimal quality.

In this work, we propose a novel deep compositional representation for 4D captures. This representation can be used to reconstruct 4D captures, and it also extracts key understanding that supports high-level tasks, such as motion transfer, 4D completion, or future prediction. This is achieved by an encoder that takes a 4D capture as input and produces latent codes representing the geometry template,

initial state, and temporal deformation respectively. Taking human as an example, these three key factors are commonly understood as the identity, initial body pose, and motion¹.

To reconstruct the 4D capture, we design a novel architecture taking three latent codes as inputs. First, we keep the geometry template code (i.e. the identity) unchanged over time since it is not affected by the motion. Then, we propose a novel conditional latent Neural ODE to update the initial state code (i.e. the initial body pose) conditioned on the deformation code (i.e. the motion). The temporally varying state code is further concatenated with the geometry template code, and fed into a decoder to reconstruct an implicit occupancy field for each time frame, which recovers the 3D shape over time. Mostly similar to us, Occupancy Flow [37] also use Neural ODE [11] to update the position of each 3D point for 4D reconstruction. In contrast, our method applies the Neural ODE to update the latent state code that controls the shape globally, which is empirically more stable.

To learn our compositional representation, we propose a training strategy to enable the encoder to decouple the geometry template and deformation, inspired by He *et al.* [19]. Specifically, we take two 4D captures from different subjects and extract their latent codes respectively. We then swap their geometry template code and train the network to reconstruct the motion with swapped geometry template. The training is fully supervised by synthetic data, where the parametric model is used to generate 4D captures with the same motion but different geometry template, e.g. SMPL model [28] for humans. We found this training strategy is effective in separating geometry template from the motion, which naturally supports motion transfer. The representation also enables 4D completion from captures with either missing frames or partial geometry by solving an optimization to update the latent codes until the partial observation is best explained.

Our contributions can be summarized as follows. First, we design a novel deep representation for 4D captures that understands the geometry template, initial state, and temporal deformation, and propose a novel training strategy to learn it. Second, we propose a novel decoder to reconstruct 4D captures from the learned representation, which includes, as a key component, a conditional Neural ODE to recover varying pose codes under the guidance of the motion code; and these codes are then translated into an occupancy field in implicit representation to recover the varying shape. Finally, we show that our model outperforms state-of-the-art methods on 4D reconstruction, and our compositional representation is naturally suitable for various applications, including motion transfer and 4D completion.

¹Since the experiment is mainly conducted on 4D human captures, we use these terms interchangeably.

2. Related Work

There is a large body of work that focuses on 3D representation, 4D capture, 3D pose and motion transfer, and compositional/disentangled representation. We discuss the most related techniques in the context of our work.

3D Representation There has been a lot of work aiming at reconstructing a continuous surface from various type of inputs, such as color images [49, 57, 25, 10, 38], point clouds [7, 3, 26], etc. Recently, great success has been achieved for 3D shape reconstruction using deep learning techniques. In early works, 3D volumes [13, 17, 58] and point clouds [46, 16, 45, 1] are adopted as the outputs of the networks, which suffer from the problems of losing surface details or limited resolutions. With the development of the graph convolution network, recent methods [18, 23, 57, 27] take the triangle mesh as the output representation, most of which regress the vertices and faces directly and require an initial template and fixed topology. Most recently, there has been significant work [33, 40, 22, 12, 15, 8] on learning an implicit field function for surface representation, which allow more flexible output topology and network architectures. Among those methods, Occupancy Network [33] represents 3D shapes using continuous indicator functions by specifying which subset of 3D space the object occupies, and the iso-surface can be extracted by utilizing Marching Cube algorithm [29].

4D Capture Research on 4D capture has been advancing significantly in the past decades [42, 53, 32, 2]. However, most works are developed based on strong assumptions [54, 42, 53, 31, 56], demand the costly multi-view inputs [36, 52, 35, 14]. Behl *et al.* [4] provide the 4D scene flow estimation leveraging object localization or semantic priors from deep networks, while the motion of scenarios is assumed to be in a tiny range, fixed pattern, rigid or linear, and high quality multi-view inputs are required. This greatly limits the ease of use and stability. Meanwhile, some methods exploit guided transformations on predefined templates to capture the time-dependent 3D flow [5, 28, 60, 43, 24]. Such methods usually focus on specific shape categories, and the performance is restricted by the characteristic and the generalization ability of the template model.

Recently, Occupancy Flow [37] is presented to learn a temporally continuous field to model the motion of every point in space and time with Neural ODE [11] and the continuous implicit occupancy representation. Nevertheless, since the network is trained to model the continuous flow of the initial occupancy space, the quality of 4D reconstruction results relies on the initial frame heavily.

3D Pose and Motion Transfer Conventional methods solving the 3D pose transfer problem via discrete deformation transfer. Learning-based mesh deformation is presented in [55], which leverages the spatially adaptive instance nor-

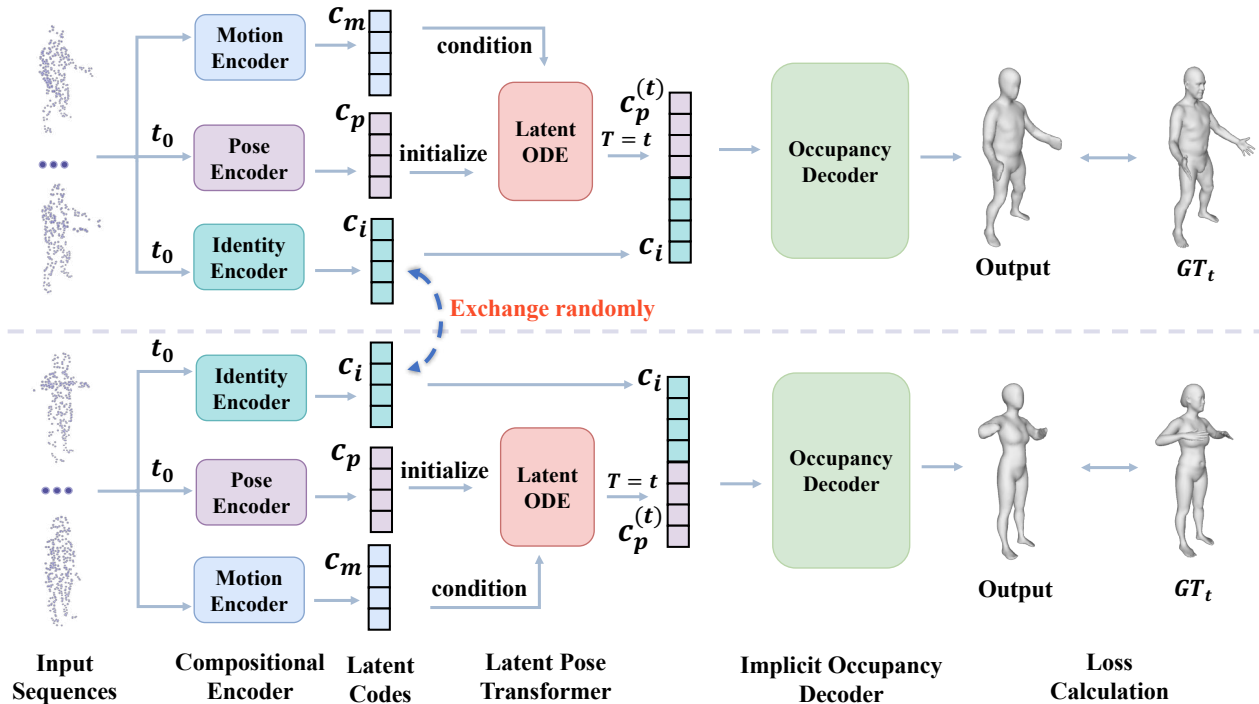


Figure 2: **Overview of our model.** Our full model contains three building blocks, namely, compositional encoder, latent pose transformer and implicit occupancy decoder. During each training step, two point cloud sequences are chosen randomly from the training set as a pair and fed into the three encoders successively (note that the motion encoder is provided with the whole sequence, while the other two encoders are only provided with the first time step of sequence). After that, there is a 50% probability that the identity codes of the two sequences are exchanged before continuing the forward propagation. Note that if the identity codes are exchanged, the ground truth meshes used for loss calculation will also be updated correspondingly.

malization [21] in the network. Nevertheless, dense triangle mesh is required and the modeling of both spatial and temporal motion continuous flow is unavailable.

3D motion transfer aims at producing a new shape sequence given a pair of source and target shape sequences, making the target shape sequence do the same temporal deformation as the source, which focuses on the continuous pose transformation among shape sequences. By applying vector field-based motion code to target shape, Occupancy Flow [37] transfers motion among human model sequences. Essentially, since Occupancy Flow does not explicitly disentangle the representations of pose and shape as done in our work, we notice the good motion transfer results of Occupancy Flow are mostly achieved in the cases that the identities and initial poses from source and target are similar.

Compositional/Disentangled Representation Learning compositional/disentangled representations has been extensively studied in previous work [51, 34, 48, 61, 41, 50]. One attractive property of human intelligence is to learn novel concepts from a few or even a single example by composing known primitives [51], which is lacking in the current deep learning system. Prior work utilize compositional/disentangled representations to address various tasks. Zhu *et al.* [61] disentangle shape, viewpoint, and texture

and present an end-to-end adversarial learning framework to generate realistic images. Tewari *et al.* [50] learn a face model from in-the-wild video with a novel multi-frame consistency loss, the proposed approach represents the facial geometry and appearance in different spaces and achieves realistic 3D face reconstruction. Park *et al.* [41] propose a fully unsupervised method to learn a swapping autoencoder for deep image manipulation task, which disentangles texture from structure. Most recently, Rempe *et al.* [47] propose CaSPR to learn a 4D representation of dynamic object point cloud sequences in Temporal-NOCS using latent Neural ODE and enable multiple applications. By dividing the latent feature into static and dynamic parts, it realizes shape and motion disentanglement. Unlike methods mentioned above, our goal is to learn a deep compositional representation for 4D captures with conditional latent Neural ODE, which decouples geometry template, initial state, and temporal deformation into different latent spaces, and supports various high-level tasks.

3. Method

In this section, we introduce our compositional representation for 4D captures and the training strategy to learn it from data. The full pipeline of our framework is illustrated

in Fig. 2. Taking a 3D model doing non-rigid deformation in time span $[0, 1]$, we extract the sparse point cloud from the 3D model in K uniformly sampled time stamps and feed them to the network as inputs. Our goal is to learn separate compact representations for identity \mathbf{c}_i , initial pose \mathbf{c}_p , and motion \mathbf{c}_m , and reconstruct the 3D model at any continuous time stamp from them. On the encoder side, we train three PointNet [46] based networks to extract \mathbf{c}_i and \mathbf{c}_p from the first frame, and \mathbf{c}_m from the whole sequence. To reconstruct mesh in target time t , we first update the initial pose code \mathbf{c}_p to $\mathbf{c}_p^{(t)}$ encoding the pose of the 3D model in target time, which is achieved via a Neural ODE [11] conditioned on motion code \mathbf{c}_m . The \mathbf{c}_i and $\mathbf{c}_p^{(t)}$ are then concatenated and fed into a network to produce an implicit occupancy field indicating whether a 3D location is inside or outside the 3D shape, and the 3D mesh surface can be reconstructed via the Marching Cube algorithm [30].

3.1. Compositional Encoder

We utilize three separate encoders to extract 128-d latent codes for identity, pose, and motion respectively. Inspired by Occupancy Flow [37], we use a PointNet-based [46] network architecture with ResNet blocks [20] as the backbone, although the input to each encoder is different according to the semantic meaning of each code. The initial pose only depends on the 3D shape in the first frame, therefore the corresponding encoder only takes the point cloud of the first frame, i.e. $t = 0$, as input. In contrast, the motion encoder takes the whole point cloud sequence as the input since the motion code needs to encode the deformation throughout the whole time span. To get the identity code, the encoder can take the whole sequence as input, but we empirically found that using only the first frame is enough and achieves similar performance.

3.2. Latent Pose Transformer

After obtaining the \mathbf{c}_i , \mathbf{c}_p , and \mathbf{c}_m from the compositional encoder, the next step is to update the pose code for target time t , i.e. $\mathbf{c}_p^{(t)}$, which are used to reconstruct the 3D shape in corresponding time stamp. Intuitively, the target pose code should start from the initial pose code, i.e. $\mathbf{c}_p^{(t=0)} = \mathbf{c}_p$, and varies continuously over time conditioned on the motion code \mathbf{c}_m . To this end, we propose a novel latent pose transformer, which is achieved by a conditional latent Neural ODE.

Neural ODE is used to reconstruct continuous temporal signals $S(t)$. Instead of directly estimating the target value, Neural ODE $f_\theta(t)$ predicts differential elements at each time stamp, which can be integrated to reconstruct the signal, i.e. $S(T) = S(0) + \int_0^T f_\theta(t, S(t)) dt$. For our specific scenario, we train a Neural ODE to predict the variation of the latent pose code over time. Different from the

original Neural ODE, our model is further conditioned on the motion code, which allows the same network to update initial poses in different manners according to the motion exhibit in the input sequence. Therefore, the pose code in target time T is obtained by

$$\mathbf{c}_p^{(T)} = \mathbf{c}_p + \int_0^T f_\theta(\mathbf{c}_p^{(t)}, t | \mathbf{c}_m) dt, \quad (1)$$

where $f_\theta(\cdot)$ is modeled by a neural network of 5 residual blocks with θ as the parameters. Following the advice of [11], we obtain the gradient using the adjoint sensitivity method [44]. For more details, please refer to the Supplementary Material.

3.3. Implicit Occupancy Decoder

The last stage of our model translates the identity code \mathbf{c}_i and the pose code in target time $\mathbf{c}_p^{(t)}$ into 3D shape. Inspired by high geometry quality from recent work using implicit representation [40, 33, 59], we train an Occupancy Network (ONet) [33] to predict for each 3D location if they were inside or outside the object surface:

$$\mathbf{o}_p^{(t)} := \Phi_\eta(\mathbf{p} | \mathbf{c}_i \oplus \mathbf{c}_p^{(t)}), \quad (2)$$

where \mathbf{p} is a 3D location, Φ_η is an ONet parameterized by η , \oplus denotes the concatenation operation between codes, $\mathbf{o}_p^{(t)}$ is the occupancy of location \mathbf{p} in time t . Note that the identity code \mathbf{c}_i remains the same as it should not change over time.

3.4. Identity Exchange Training

Naively training our network with 4D reconstruction is not sufficient to learn the compositional representation that isolates identity, initial pose, and motion. We introduce a simple yet effective training strategy (shown in Fig. 2), where the network is asked to reconstruct the same motion with different identities. Specifically, we extract latent codes for two sequences, $\{\mathbf{c}_{i_1}, \mathbf{c}_{m_1}, \mathbf{c}_{p_1}\}$ and $\{\mathbf{c}_{i_2}, \mathbf{c}_{m_2}, \mathbf{c}_{p_2}\}$, from different subjects \mathbf{s}_1 and \mathbf{s}_2 . We then swap their identity codes and supervise the model to reconstruct ground truth 4D captures of the same motion performed by the other subjects, i.e. \mathbf{s}_1 performing motion of \mathbf{s}_2 and vice versa. Since \mathbf{c}_{m_1} and \mathbf{c}_{p_1} has no visibility to \mathbf{s}_2 , all the identity information for \mathbf{s}_2 has to be encoded in \mathbf{c}_{i_2} for successful reconstruction. In practice, we perform this identity exchange training strategy for 50% of the iterations, and find it effective in disentangling identity and motion.

3.5. Loss Function

Our model is trained by minimizing the binary cross entropy error (BCE) on the occupancy of 3D locations. Inspired by Occupancy Network [33], we randomly sample a

time step τ and a set of 3D query points \mathcal{S} , and compute the loss $\mathcal{L}^{(\tau)}$ between the predicted occupancy value $\mathbf{o}_p^{(\tau)}$ and the ground truth $\hat{\mathbf{o}}_p^{(\tau)}$:

$$\mathcal{L}^{(\tau)} = \sum_{p \in \mathcal{S}} \text{BCE} \left(\hat{\mathbf{o}}_p^{(\tau)}, \mathbf{o}_p^{(\tau)} \right). \quad (3)$$

To get \mathcal{S} , we normalize all the meshes to $[-0.5, 0.5]$ with the global scale and translation calculated from the dataset, and sample 50% points uniformly in a bounding volume and 50% points near the surface of the mesh. We find the definition of the bounding volume affects the training performance, and experiment with two ways: 1) a fixed volume with the length of 1; 2) a tight bounding volume around the mesh, in our experiments.

During training, we also supervise the predicted occupancy value at time step 0 to ensure a high quality initialization. Therefore, the complete loss function is defined as:

$$\mathcal{L} = \lambda_1 \mathcal{L}^{(0)} + \lambda_2 \mathcal{L}^{(\tau)}, \quad (4)$$

where $\lambda_1 = \lambda_2 = 1.0$ in our experiment. We use Adam optimizer with the learning rate as 10^{-4} . The model is trained with batch size of 16 on a single NVIDIA RTX 2080Ti GPU.

4. Experiments

In this section, we perform extensive experiments to evaluate our method. We first show the ability for 4D reconstruction, and then apply our compositional representation to various tasks like motion transfer, 4D completion and future prediction.

4.1. Data Preparation

We use two datasets to train and evaluate our proposed method. The first dataset is Dynamic FAUST (D-FAUST) [6], which contains 129 mesh sequences of 10 real human subjects performing 14 different motions and all meshes are registered with SMPL [28] model. We augment D-FAUST to meet the needs of our Identity Exchange Training strategy (Sec. 3.4). We first fit the SMPL shape and pose parameters for all the data. Then, the ground truth mesh sequences of all the combinations of identities and motions are generated, extending the number of mesh sequences to about 1000.

We also build a Warping Cars dataset using the approach introduced in Occupancy Flow [37] to investigate the performance of our method on non-human objects. Specifically, we randomly choose 10 car models from ShapeNet [9] Car category and generate 1000 warplings. To generate a warping field, Gaussian displacement vectors are sampled in a $3 \times 3 \times 3 \times 5$ grid and the RBF [39] interpolation is used to obtain a continuous displacement field. We combine different car shapes and warplings and finally get the dataset with total number of mesh sequences to 10000, each of which has 50 time steps.

4.2. 4D Reconstruction

We first verify the reconstruction ability of our model following the setting in Occupancy Flow (OFlow) [37]. The network consumes 300 sparse point trajectories as input, each of which consists of 3D locations at $L = 17$ equally divided time stamps, and the goal is to reconstruct compact mesh at these time stamps even though the model is able to produce mesh at any particular time. For human model, we use the same train/test split on D-FAUST [6] as OFlow, including data on subjects seen and unseen during the training respectively. For Warping Cars dataset, we test on our own testing set as it was not released.

The quantitative results on the D-FAUST dataset and the Warping Cars dataset are summarized in Tab. 1, where we report the average IoU and Chamfer Distance over 17 frames of all testing sequences. As our baseline, ‘‘PSGN 4D’’ is a 4D extension of Point Set Generation Network [16] by predicting a set of trajectories instead of single points, and ‘‘ONet 4D’’ is an extension of Occupancy Network (ONet) [33], which predicts occupancy value for points sampled in 4D space and reconstructs each frame of the sequence separately. OFlow uses Neural ODE to learn a continuous motion vector field for every point in space and time. While OFlow explicitly transform the 3D coordinates of each point, we transform the pose code in the latent space. The results of baselines on the D-FAUST dataset are cited from OFlow [37], and the results on Warping Cars dataset are produced by a retrained OFlow. Overall, our method performs comparable or better than other methods on D-FAUST and Warping Cars datasets, indicating that our model is able to reconstruct accurate surfaces.

In Fig. 3, we show qualitative comparison on the D-FAUST dataset with OFlow. Our method is able to capture more details, such as the shape of the opening hands and the outline of the muscles on the body. In particular, OFlow fails to track the motion of hands in the last frame of the left sequence, while our method produces stable results during the whole sequence time. This is presumably because our method reconstructs each frame of the whole sequence individually with the transformed pose latent code, while OFlow only reconstructs the first frame and deforms it with the learned transformation flow. Furthermore, the qualitative results on the Warping Cars dataset are shown in Fig. 4, in which our method shows better capability of recovering motion than OFlow.

We provide an ablation study on 4D reconstruction task based on the D-FAUST dataset to evaluate the effectiveness of ODE in different aspects. 1) *ODE in feature v.s. 3D*. We found that training an ODE directly in 3D space, i.e. exactly an OFlow, with identity exchange is hard to converge. We trained the model for a week on a single NVIDIA RTX 2080Ti GPU, and found the model learns barely any motion. In contrast, applying ODE in feature space, i.e.

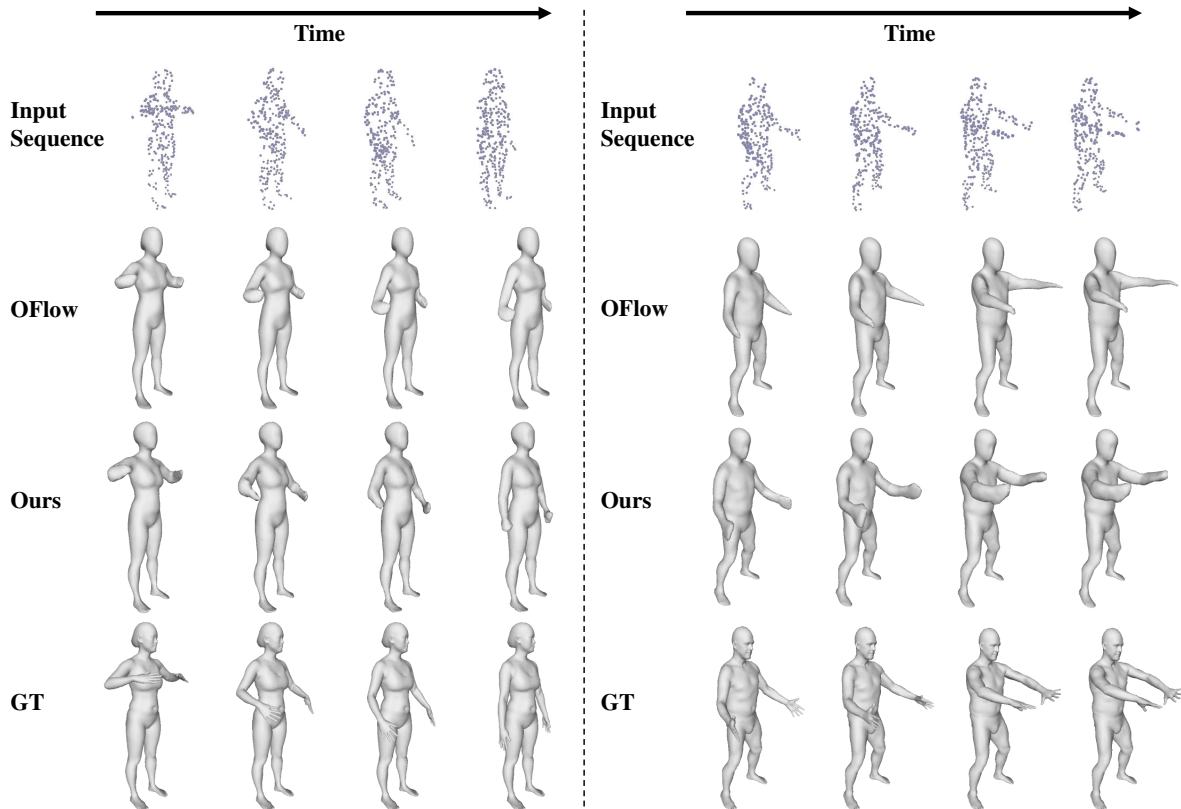


Figure 3: **4D reconstruction from point cloud sequence (D-FAUST)**. GT is short for Ground Truth.

Methods	Seen Individuals (D-FAUST)		Unseen individual (D-FAUST)		Warping Cars	
	IoU \uparrow	Chamfer Distance \downarrow	IoU \uparrow	Chamfer Distance \downarrow	IoU \uparrow	Chamfer Distance \downarrow
PSGN 4D	-	0.108	-	0.127	-	-
ONet 4D	77.9%	0.084	66.6%	0.140	-	-
OFlow	79.9%	0.073	69.6%	0.095	70.0%	0.166
Ours	81.8%	0.068	68.2%	0.100	70.9%	0.154
Ours*	81.5%	0.068	69.9%	0.094	-	-

Table 1: **4D reconstruction from point cloud sequence (D-FAUST and Warping Cars)**. Ours* indicates that when preparing training data, we sample query points in a fixed cubic bounding volume with the length of 1 (see Sec. 3.5).

our method, may benefit from regularization provided by the compact 1-d latent vector and converges well. 2) *ODE* v.s. *MLP*. We train a model replacing ODE to an MLP that directly produces the pose code for a specified time, and get IoU=80.4% and Chamfer Distance=0.073 for reconstruction on D-FAUST, while our ODE model achieves IoU=81.8% and Chamfer Distance=0.068. This indicates that ODE performs better than MLP in reconstructing pose sequence for motion, but the MLP model trains and runs relatively faster.

4.3. Pose and Motion Transfer

Our compositional representation also naturally supports motion transfer. Consider two subjects performing different

motions, namely $id_1 + motion_1$ and $id_2 + motion_2$, and our goal is to generate 4D sequence with $id_2 + motion_1$. To do so, we first extract the latent representations with our compositional encoder for each input sequence, namely $(c_i^1, c_{p_0}^1, c_m^1)$, $(c_i^2, c_{p_0}^2, c_m^2)$, and then feed $(c_i^2, c_{p_0}^1, c_m^1)$ to the latent pose transformer and implicit occupancy decoder.

We evaluate our method on the D-FAUST testing set, where we randomly select 20 identity and motion pairs, and generate the ground truth 4D sequences after motion transfer using the known SMPL parameters. As baseline, we compare to OFlow which also learns separate codes to represent first frame geometry (i.e. the identity) and velocity field (i.e. the motion) respectively. In addition, we also build a baseline with the recent state-of-the-art neural pose

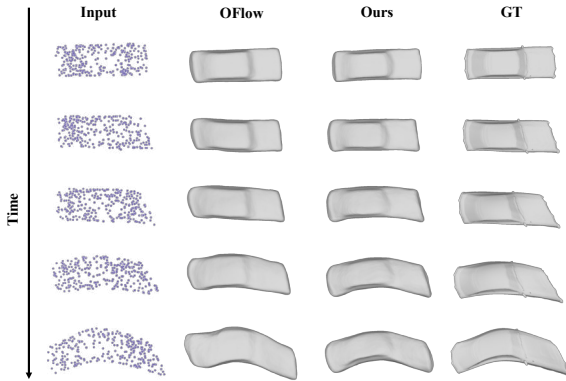


Figure 4: **4D reconstruction on Warping Cars dataset.**

transfer method NPT [55], which utilizes spatially adaptive instance normalization to deform the identity point cloud to each time step of the target motion sequence using pose transfer. The transformed point clouds are then fed into OFlow to generate complete meshes.

The quantitative results are shown in Tab. 2. Our method significantly outperforms other baseline methods with large margins. One qualitative comparison is shown in Fig. 5. The performance of NPT is heavily limited by the density of the input identity and motion sequences, which makes it hard to transfer the continuous motion with sparse inputs. OFlow does not transfer the motion at all, presumably because the pose representations are not decoupled from the shape latent code, which leads to a wrong first frame pose initialization and the failure of the whole generated sequence. In contrast, our method successfully transfers the motion to the new identity, including both the initial pose and following frames. Additional results on Warping Cars dataset are shown in Supplementary Material.

	IoU \uparrow	Chamfer Distance \downarrow
NPT	26.4%	0.498
OFlow	26.7%	0.400
Ours	85.0%	0.055

Table 2: **Motion transfer (D-FAUST).**

We further investigate if the motion code c_m can be transferred without the initial pose code c_p . Even though this is sometimes an ill-posed problem (e.g. forcing a stand-up motion to start with a standing pose), we find, surprisingly, our model is still able to produce reasonable results if the new initial pose is not too different from the original one (See Supplementary Material for results). This indicates that our conditional Neural ODE is robust to some extent against the noise in the initial pose code.

4.4. 4D Completion

Our compositional 4D representation also provides strong prior as the regularization for 4D completion task,

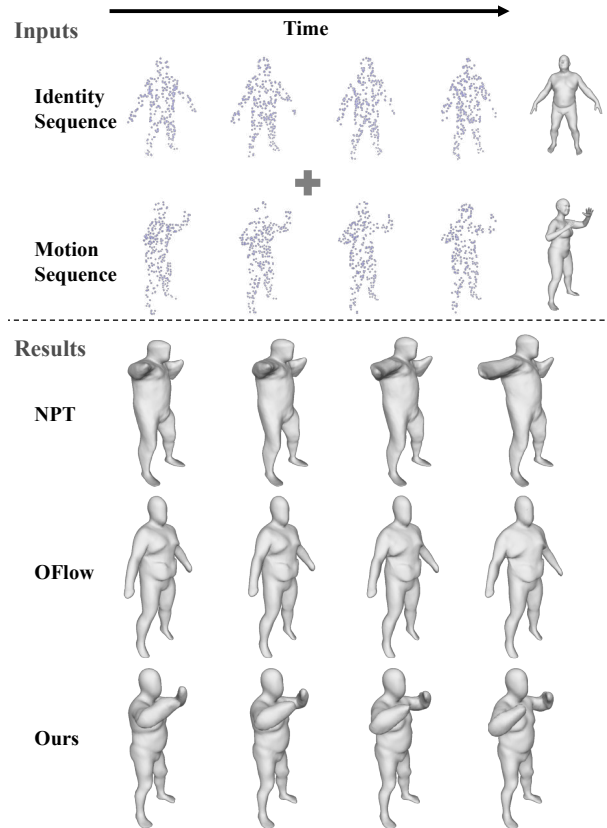


Figure 5: **Motion transfer results on an input sequence pair (D-FAUST).** Our method transfers the motion of the second sequence to the identity of the first sequence successfully, while keeping the shape property of the first sequence unchanged.

Methods	Temporal		Spatial	
	IoU \uparrow	CD \downarrow	IoU \uparrow	CD \downarrow
OFlow	85.1%	0.057	86.0%	0.054
Ours	86.4%	0.056	87.2%	0.051

Table 3: **4D temporal completion and spatial completion (D-FAUST).** CD is short for Chamfer Distance.

in which the goal is to fill in the missing signals in a given 4D capture with only partial observation. This is practically useful when part of the 4D capture is corrupted due to imperfect capturing techniques or challenging scenarios. Specifically, this task can be categorized into two kinds based on the missing data: 1) **Temporal completion**, which recovers the missing frames; 2) **Spatial completion**, which completes partial geometry in each frame. To perform these tasks, we remove the encoder, fix the decoder parameters, and optimize the latent codes with back-propagation until the output 4D sequence matches the partial observation.

The experiments are conducted on the D-FAUST dataset.

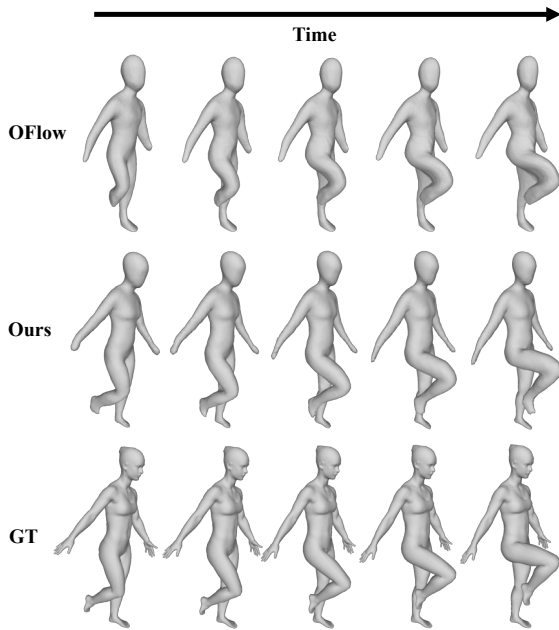


Figure 6: **4D temporal completion (D-FAUST)**. We show the results of 5 missing frames.




Figure 7: **Future prediction (D-FAUST)**. The results on the left of the dotted line are reconstructions for partial observation, and the results on the right are future predictions.

For temporal completion, we select 18 mesh sequences with $L = 30$ frames from the testing set, and randomly withheld half of the frames in each sequence for testing. For spatial completion, we randomly select three points in each frame

and remove the points less than 0.2 away from them.

Comparison to OFlow is shown in Tab. 3. Our method performs comparable or better IoU and Chamfer distance than OFlow on both temporal and spatial completion. Fig. 6 shows a temporal completion result. Our method successfully interpolates correct poses for missing frames with more complete geometry than OFlow. Please refer to Supplementary Material for results on Warping Cars dataset.

4.5. Future Prediction

Not only interpolating internal missing frames, our model can also predict the future of the motion by extrapolating temporal frames onward. To validate this, we select 15 mesh sequences with $L = 20$ frames from the D-FAUST testing set, and always remove the last 10 frames instead of randomly selected ones. Tab. 4 and Fig. 7 show the comparison to OFlow. Though OFlow can also produce reasonable future motion, the magnitudes are usually small which leads to overly slow motion. In contrast, our method predicts much more accurate motion, e.g. with the other leg raised. The results on Warping Cars dataset are shown in Supplementary Material.

	IoU \uparrow	Chamfer Distance \downarrow
OFlow	75.5%	0.099
Ours	80.8%	0.081

Table 4: **Future prediction (D-FAUST)**. We remove the last 10 frames of sequence to investigate the extrapolation ability of our method.

5. Conclusion

This paper introduces a compositional representation for 4D captures by disentangling the geometry template, initial state and temporal deformation with separated compact latent codes, which can reconstruct the deforming 3D object over a temporal span. Furthermore, an identity exchange training strategy is proposed to make geometry template and temporal deformation efficiently decoupled and exchangeable. Extensive experiments on 4D reconstruction, pose and motion transfer, 4D completion, and motion prediction validate the efficacy of our proposed approach.

Acknowledgement

Yanwei Fu is the corresponding author. This work was supported in part by NSFC Projects (U62076067), Science and Technology Commission of Shanghai Municipality Projects (19511120700, 19ZR1471800), Shanghai Research and Innovation Functional Program (17DZ2260900), Shanghai Municipal Science and Technology Major Project (2018SHZDZX01) and ZJLab.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Representation learning and adversarial generation of 3d point clouds. *arXiv preprint arXiv:1707.02392*, 2(3):4, 2017. 1, 2
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8387–8397, 2018. 2
- [3] Nina Amenta, Marshall Bern, and Manolis Kamvyselis. A new voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 415–421, 1998. 2
- [4] Aseem Behl, Omid Hosseini Jafari, Siva Karthik Mustikovela, Hassan Abu Alhaija, Carsten Rother, and Andreas Geiger. Bounding boxes, segmentations and object coordinates: How important is recognition for 3d scene flow estimation in autonomous driving scenarios? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2574–2583, 2017. 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*, pages 187–194, 1999. 2
- [6] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J Black. Dynamic faust: Registering human bodies in motion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6233–6242, 2017. 5
- [7] Jean-Daniel Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Transactions on Graphics (TOG)*, 3(4):266–286, 1984. 2
- [8] Rohan Chabra, Jan Eric Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3d reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 2
- [9] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 5
- [10] ChaoWen, Yinda Zhang, Zhuwen Li, and Yanwei Fu. Pixel2mesh++: Multi-view 3d mesh generation via deformation. In *ICCV*, 2019. 2
- [11] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In *Advances in neural information processing systems*, pages 6571–6583, 2018. 1, 2, 4
- [12] 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*, pages 6970–6981, 2020. 2
- [13] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *ECCV*, 2016. 1, 2
- [14] Junting Dong, Qing Shuai, Yuanqing Zhang, Xian Liu, Xiaowei Zhou, and Hujun Bao. Motion capture from internet videos. In *European Conference on Computer Vision*, pages 210–227. Springer, 2020. 2
- [15] Philipp Erler, Paul Guerrero, Stefan Ohrhallinger, Niloy J. Mitra, and Michael Wimmer. Points2surf: Learning implicit surfaces from point clouds. In *ECCV*, 2020. 2
- [16] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 605–613, 2017. 1, 2, 5
- [17] Rohit Girdhar, David F. Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *ECCV*, 2016. 1, 2
- [18] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. Atlasnet: A papier-mâché approach to learning 3d surface generation. *arXiv preprint arXiv:1802.05384*, 2018. 1, 2
- [19] Jiawei He, Andreas Lehrmann, Joseph Marino, Greg Mori, and Leonid Sigal. Probabilistic video generation using holistic attribute control. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–467, 2018. 2
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [21] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 3
- [22] Chiyu 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*, pages 6001–6010, 2020. 1, 2
- [23] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 371–386, 2018. 1, 2
- [24] Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5614–5623, 2019. 2
- [25] Hiroharu Kato, Yoshitaka Ushiku, and Tatsuya Harada. Neural 3d mesh renderer. In *CVPR*, 2018. 2
- [26] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, volume 7, 2006. 2
- [27] Yiyi Liao, Simon Donne, and Andreas Geiger. Deep marching cubes: Learning explicit surface representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2916–2925, 2018. 2

- [28] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 2, 5
- [29] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, 1987. 2
- [30] 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
- [31] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. Xnect: Real-time multi-person 3d motion capture with a single rgb camera. *ACM Transactions on Graphics (TOG)*, 39(4):82–1, 2020. 2
- [32] Moritz Menze and Andreas Geiger. Object scene flow for autonomous vehicles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3061–3070, 2015. 2
- [33] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 4, 5
- [34] Ishan Misra, Abhinav Gupta, and Martial Hebert. From red wine to red tomato: Composition with context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1792–1801, 2017. 3
- [35] Armin Mustafa, Hansung Kim, Jean-Yves Guillemaut, and Adrian Hilton. General dynamic scene reconstruction from multiple view video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 900–908, 2015. 2
- [36] Jan Neumann and Yiannis Aloimonos. Spatio-temporal stereo using multi-resolution subdivision surfaces. *International Journal of Computer Vision*, 47(1-3):181–193, 2002. 2
- [37] 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. 1, 2, 3, 4, 5
- [38] Chengjie Niu, Jun Li, and Kai Xu. Im2struct: Recovering 3d shape structure from a single RGB image. In *CVPR*, 2018. 2
- [39] Jooyoung Park and Irwin W Sandberg. Universal approximation using radial-basis-function networks. *Neural computation*, 3(2):246–257, 1991. 5
- [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. 1, 2, 4
- [41] Taesung Park, Jun-Yan Zhu, Oliver Wang, Jingwan Lu, Eli Shechtman, Alexei A Efros, and Richard Zhang. Swapping autoencoder for deep image manipulation. *arXiv preprint arXiv:2007.00653*, 2020. 3
- [42] Yuri Pekelny and Craig Gotsman. Articulated object reconstruction and markerless motion capture from depth video. In *Computer Graphics Forum*, volume 27, pages 399–408. Wiley Online Library, 2008. 2
- [43] Leonid Pishchulin, Stefanie Wuhrer, Thomas Helten, Christian Theobalt, and Bernt Schiele. Building statistical shape spaces for 3d human modeling. *Pattern Recognition*, 67:276–286, 2017. 2
- [44] Lev Semenovich Pontryagin. *Mathematical theory of optimal processes*. Routledge, 2018. 4
- [45] Charles R. Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J. Guibas. Frustum pointnets for 3d object detection from RGB-D data. In *CVPR*, 2018. 1, 2
- [46] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 1, 2, 4
- [47] Davis Rempe, Tolga Birdal, Yongheng Zhao, Zan Gojcic, Srinath Sridhar, and Leonidas J Guibas. Caspr: Learning canonical spatiotemporal point cloud representations. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [48] Austin Stone, Huayan Wang, Michael Stark, Yi Liu, D Scott Phoenix, and Dileep George. Teaching compositionality to cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5058–5067, 2017. 3
- [49] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *ICCV*, 2017. 2
- [50] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. Fml: Face model learning from videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10812–10822, 2019. 3
- [51] Pavel Tokmakov, Yu-Xiong Wang, and Martial Hebert. Learning compositional representations for few-shot recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6372–6381, 2019. 3
- [52] Ali Osman Ulusoy, Octavian Biris, and Joseph L Mundy. Dynamic probabilistic volumetric models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 505–512, 2013. 2
- [53] Ali Osman Ulusoy and Joseph L Mundy. Image-based 4-d reconstruction using 3-d change detection. In *European Conference on Computer Vision*, pages 31–45. Springer, 2014. 2
- [54] Michael Wand, Philipp Jenke, Qixing Huang, Martin Bokeloh, Leonidas Guibas, and Andreas Schilling. Reconstruction of deforming geometry from time-varying point clouds. In *Symposium on Geometry processing*, pages 49–58, 2007. 2
- [55] Jiashun Wang, Chao Wen, Yanwei Fu, Haitao Lin, Tianyun Zou, Xiangyang Xue, and Yinda Zhang. Neural pose transfer by spatially adaptive instance normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5831–5839, 2020. 2, 7

- [56] Kangkan Wang, Jin Xie, Guofeng Zhang, Lei Liu, and Jian Yang. Sequential 3d human pose and shape estimation from point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7275–7284, 2020. [2](#)
- [57] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018. [1](#), [2](#)
- [58] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):72, 2017. [1](#), [2](#)
- [59] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *arXiv preprint arXiv:1905.10711*, 2019. [4](#)
- [60] Qian Zheng, Xiaochen Fan, Minglun Gong, Andrei Sharf, Oliver Deussen, and Hui Huang. 4d reconstruction of blooming flowers. In *Computer Graphics Forum*, volume 36, pages 405–417. Wiley Online Library, 2017. [2](#)
- [61] Jun-Yan Zhu, Zhoutong Zhang, Chengkai Zhang, Jiajun Wu, Antonio Torralba, Josh Tenenbaum, and Bill Freeman. Visual object networks: Image generation with disentangled 3d representations. In *Advances in neural information processing systems*, pages 118–129, 2018. [3](#)