

# Pop-Out Motion: 3D-Aware Image Deformation via Learning the Shape Laplacian


Jihyun Lee<sup>\*1</sup>Minhyuk Sung<sup>\*†1</sup>Hyunjin Kim<sup>1</sup>Tae-Kyun Kim<sup>1,2</sup><sup>1</sup> KAIST<sup>2</sup> Imperial College London

Figure 1. **3D-aware image deformation of a 2D human image.** Compared to the images generated by directly deforming a mesh reconstructed using PIFu [39] and Marching Cubes [32] (blue), our method (red) can produce more plausible 3D-aware image deformations via inferring an additional intrinsic shape property: the shape Laplacian.

## Abstract

We propose a framework that can deform an object in a 2D image as it exists in 3D space. Most existing methods for 3D-aware image manipulation are limited to (1) only changing the global scene information or depth, or (2) manipulating an object of specific categories. In this paper, we present a 3D-aware image deformation method with minimal restrictions on shape category and deformation type. While our framework leverages 2D-to-3D reconstruction, we argue that reconstruction is not sufficient for realistic deformations due to the vulnerability to topological errors. Thus, we propose to take a supervised learning-based approach to predict the shape Laplacian of the underlying volume of a 3D reconstruction represented as a point cloud. Given the deformation energy calculated using the predicted shape Laplacian and user-defined deformation handles (e.g., keypoints), we obtain bounded biharmonic weights to model plausible handle-based image deformation. In the experiments, we present our results of deforming 2D character and clothed human images. We also quantitatively show that our approach can produce more accurate deformation weights compared to alternative methods (i.e., mesh reconstruction and point cloud Laplacian methods).

## 1. Introduction

The capability of photo editing, which had been confined in 2D space, has recently been popped out into 3D

space. For example, predicting depth from a 2D image enables the composition of objects in an image [41, 57]. Object segmentation allows the projection of an image to a new view [27, 36, 52]. Lighting and ground plane estimation make it possible to relight objects and generate a new shadow in an image [9, 14, 15, 54]. Such techniques for 3D-aware image editing have allowed the user to manipulate an image in a more intuitive manner — as if the object exists in 3D space — and opened new opportunities in downstream applications.

As a missing piece of the existing 3D-aware image manipulation methods, we focus on *3D-aware image deformation*. Unlike the aforementioned techniques, 3D-aware deformation does not just alter the scene information (e.g., camera parameters, lighting conditions) or modify the 2.5D information. Instead, it allows the user to directly manipulate the 3D geometry and appearance of an object. More relevant topics to 3D-aware deformation would be (1) human pose transfer [6, 29, 33], which works only for human bodies, (2) novel view synthesis [13, 37, 51], which is limited to altering a viewpoint of an image, and (3) 3D model-based manipulation [26], which requires the exact 3D model of the object in an image. To address these limitations, we aim to enable 3D-aware image deformation with minimal restrictions on shape category and deformation type.

For 3D-aware deformation, it is *necessary* to reconstruct the object in a 2D image to 3D space; however, it is *not sufficient* in general. Deformation requires either a surface or volume information [21, 48]. However, most of existing

\* equal contributions, † corresponding author

methods of image-based 3D reconstruction do not directly output a surface or volume [8, 45] or produce a surface without proper consideration about intrinsic shape properties [12, 23, 28, 39, 40, 46, 49, 53] – which can largely affect the deformation result. See an example in the red branch of Figure 1. The inaccurate topological prediction connecting legs causes undesired visual artifacts in the deformation. Indeed, aiming for topological correctness in 3D reconstruction is a difficult task due to its nature defined with both continuous and discrete quantities.

In this paper, given a 3D point cloud of the object in an input image (whose 2D-to-3D reconstruction is performed by PIFu [39]), we propose to enable 3D-aware image deformation through learning the additional intrinsic geometric property: the shape Laplacian. The shape Laplacian is the *essential* information encoding the geometry intrinsics. In particular, bounded biharmonic weights [21] – which has been widely used as a standard technique for deformation in computer graphics – compute the linear blending weights associated with deformation handles as the minimizers of the deformation energy defined using the shape Laplacian. In our framework, we utilize the *estimated* shape Laplacian of a 3D reconstruction to obtain bounded biharmonic weights to plausibly model handle-based image deformation.

To this end, we introduce a neural network that can predict the shape Laplacian of the underlying volume of a 3D point cloud reconstructed from a 2D image — without directly converting the point cloud to a volume. Considering that the deformation energy can be discretized with the standard linear FEM Laplacian  $LM^{-1}L$  (where  $L$  is a symmetric cotangent Laplacian matrix and  $M$  is a diagonal lumped mass matrix), we design our network to learn the matrices  $L$  and  $M^{-1}$  from the supervision obtained from a ground truth 3D mesh. The elements in the inverse mass matrix  $M^{-1}$  are predicted for each individual point, while the elements of the cotangent Laplacian matrix  $L$  are predicted by taking *pairs* of the input points. We use a symmetric feature aggregation function for such pairs and also a weight module to enforce the output matrix  $L$  to be symmetric and sparse. In test time, we recover the deformation energy from the predicted  $L$  and  $M^{-1}$  to compute bounded biharmonic weights with user-specified deformation handles. Since our method learns the shape Laplacian instead of the handle-dependent deformation weights, it can generalize well to arbitrary handle configurations.

In the experiments, we show our results of 3D-aware deformation on 2D character and clothed human images. We also showcase an user-interactive image editing scenario, where the user produces intuitive 3D deformations based on the specified control points. For quantitative evaluation, we test our method on a large-scale 3D point cloud dataset (i.e., DFAUST [4]), in which our method is shown to pro-

duce more accurate deformation weights compared to the alternative methods on mesh reconstruction and point cloud Laplacian.

Our main contributions can be summarized as follows:

- We propose a method for 3D-aware deformation of 2D images, which can be applied with minimal restrictions on shape category and deformation type.
- We introduce a novel network architecture that can learn the shape Laplacian with several desired properties (i.e., positive semi-definiteness, symmetry and sparsity) from a 3D reconstruction. To the best of our knowledge, this is the first study to demonstrate that a learning-based approach can be effective in predicting the shape Laplacian of the underlying volume of a point cloud.
- We empirically demonstrate that our *learning*-based approach leads to more plausible deformations compared to the alternative cases of *calculating* the approximation of the shape Laplacian using mesh reconstruction or point cloud Laplacian methods.

## 2. Related Work

**Handle-Based Deformation** In geometry processing, methods for handle-based deformation have been investigated for decades. A common workflow for handle-based deformation consists of two phases: bind time and pose time. In bind time, the source shape is bound to a user-defined set of control handles. In pose time, the handles are manipulated to produce shape deformations in a way that the transformations at each handle are smoothly propagated to the rest of the shape. Most of existing methods compute handle-based deformation weights based on solving an optimization problem to minimize a shape fairness functional (e.g., discretized Laplacian energy [21], linearly precise smoothness energy [48]). In this work, we utilize bounded biharmonic weights [21] to model handle-based deformation for 2D images.

**Learning Shape Deformation** Most of recent methods in 3D vision and graphics utilize neural networks to learn shape deformations. Typically, the main objective of such existing work is to fit the source to the target shape. These target-driven deformation methods have been shown to be effective in various tasks, such as 3D reconstruction [20, 46, 47], shape auto-encoding [44], deformation transfer [43], and data augmentation [1, 35]. Unlike these methods, the main goal of our work is to compute deformation weights associated with *user-defined control handles* – which can serve as an intuitive deformation interface – for a 3D reconstruction whose ground truth topology is unknown. We also note that our framework does not require any source and target shape pairs or semantic labels for training.

**Mesh Reconstruction Methods** One of the alternative approaches to compute handle-based deformation weights for a 3D point cloud is to compute the shape Laplacian using the mesh topology estimated from a surface reconstruction method, an optional manifold conversion algorithm [17, 18]), and a tetrahedral meshing method [16]. Explicit surface reconstruction methods [3, 5, 30] directly estimate the connectivity information for input points. Implicit surface reconstruction methods [24, 25, 38] predict a field function that can be used for iso-surface meshing, which is typically performed via Marching Cubes [32]. However, most of these existing methods are vulnerable to topological errors when ambiguous structures (e.g., spatially adjacent surfaces, high-curvature surface) exist. In this paper, we propose to bypass such explicit mesh conversion procedure and to directly learn the shape Laplacian of a point cloud from the ground truth topology supervision. In Section 4, we empirically demonstrate that our method can produce more accurate deformation weights than the aforementioned scenario.

**Point Cloud Laplacians** Another alternative approach to calculate handle-based deformation weights for a reconstructed 3D point cloud is to directly construct a point cloud Laplacian. While there are existing methods that approximate the shape Laplacian of a point cloud [2, 42], the main challenge is to infer the correct topology of the underlying structure of a point cloud – which is infeasible using the given point cloud only. The construction of a point cloud Laplacian is commonly based on the triangulation on the tangent planes estimated from the  $k$  nearest neighbors of each point in the *Euclidean* space. Therefore, the resulting shape Laplacian is erroneous with respect to the two neighboring points with a long underlying *geodesic* distance. To address this issue, we propose to leverage the prediction power of a neural network, which is trained using the *ground truth* shape Laplacian as a supervisory signal. To the best of our knowledge, this is the first study to take a learning-based approach to predict the shape Laplacian of the underlying volume from a point cloud.

### 3. Learning for Handle-Based Deformation

We propose a learning-based method for 3D-aware image deformation. First, we use a 3D reconstruction method (i.e., PIFu [39]) to construct a point cloud that represents the 3D shape of the object in an input image. Next, we use a carefully-designed neural network to predict the shape Laplacian of the underlying volume of the reconstructed 3D point cloud. The estimated shape Laplacian is then used to compute the deformation blending weights (i.e., bounded biharmonic weights [21]) corresponding to user-defined handles to model 3D-aware image deformation. We argue that our approach to *learning* the shape Laplacian al-

lows more robust and accurate deformations than directly *calculating* the approximation of the shape Laplacian using a mesh reconstruction or point cloud Laplacian method (please refer to Section 4 for experimental results). In addition, since we predict the shape Laplacian – not directly the deformation blending weights that are dependent on deformation handles, our method can model deformation robust to an arbitrary deformation handle configuration.

In what follows, we first briefly review bounded biharmonic weights [21], which is one of the existing methods to compute handle-based deformation weights for a *mesh*, whose ground truth topology is available. We then introduce our method to obtain robust bounded biharmonic weights for a 3D reconstruction.

#### 3.1. Background: Bounded Biharmonic Weights

Bounded biharmonic weights [21] are linear blending weights that propagate affine transformations defined at arbitrary control handles to the other points in the shape. Given a tetrahedral mesh<sup>1</sup>  $\mathcal{M} = \{\mathcal{V}, \mathcal{F}\}$  representing the volume of a shape with the sets of vertices and faces  $\mathcal{V}$  and  $\mathcal{F}$ , respectively, a control handle  $\mathcal{H}$  is defined as a point ( $\mathcal{H} \in \mathcal{V}$ ) or a region ( $\mathcal{H} \subset \mathcal{V}$ ) on the mesh<sup>2</sup>. Specifically, given (1) the source shape  $\mathcal{M}$  with  $n$  vertices ( $\mathcal{V} = \{\mathbf{v}_i\}_{i=1 \dots n}$ ) (2)  $m$  number of control handles  $\{\mathcal{H}_k\}_{k=1 \dots m}$ , and (3) the affine transformations  $\{\mathbf{T}_k\}_{k=1 \dots m}$  defined on each control handle, the new position of  $i$ -th vertex  $\mathbf{v}_i \in \mathcal{V}$  can be calculated using the following linear formulation:

$$\mathbf{v}'_i = \sum_{k=1}^m w_{k,i} \mathbf{T}_k \mathbf{v}_i, \quad (1)$$

where  $w_{k,i}$  is a deformation weight associated with the  $k$ -th control handle  $\mathcal{H}_k$  and the  $i$ -th vertex  $\mathbf{v}_i$ .

The deformation weights  $\mathbf{w}_k = \{w_{k,1}, \dots, w_{k,n}\}^T$  for each handle are computed as a minimizer with respect to the deformation energy subject to several constraints for the desired properties of deformation (e.g., partition of unity, non-negativity). The minimization problem can be written as follows:

$$\operatorname{argmin}_{\{\mathbf{w}_k\}_{k=1 \dots m}} \sum_{k=1}^m \frac{1}{2} \mathbf{w}_k^T A \mathbf{w}_k \quad (2)$$

$$\begin{aligned} \text{subject to: } & w_{k,i} = 1 \quad \forall i \quad \text{s.t.} \quad \mathbf{v}_i \in \mathcal{H}_k \\ & w_{k,i} = 0 \quad \forall i \quad \text{s.t.} \quad \mathbf{v}_i \in \mathcal{H}_{l, l \neq k} \\ & \sum_{k=1}^m w_{k,i} = 1, \quad i = 1, \dots, n, \\ & 0 \leq w_{k,i} \leq 1, \quad k = 1, \dots, m, \quad i = 1, \dots, n, \end{aligned}$$

<sup>1</sup>While it is also possible to compute the shape Laplacian from a *surface* mesh, we focus on the *volume* Laplacian in this work.

<sup>2</sup>The handle can also be defined *out of* the mesh as a skeleton bone or a vertex of a cage. Here, we only describe the case when the handle is defined as a point or a region on the mesh.

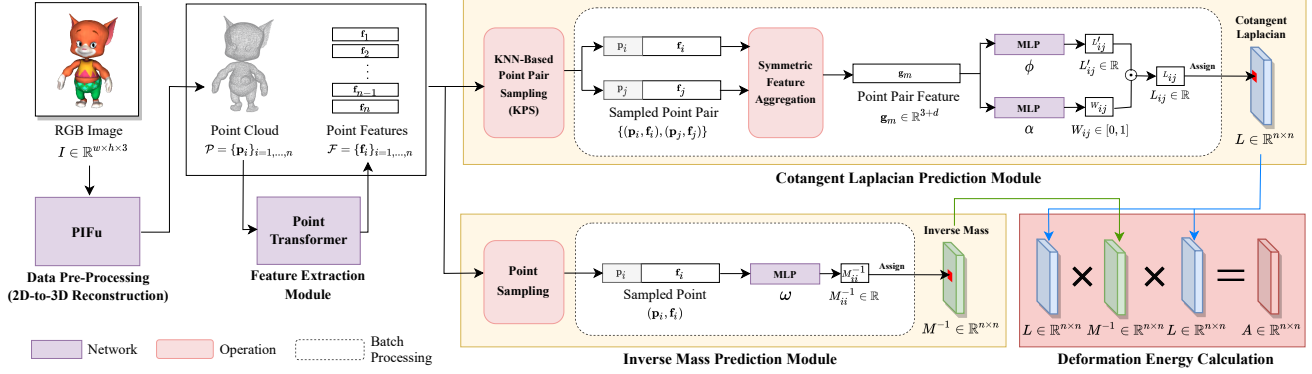


Figure 2. **Architecture overview.** Given a point cloud  $\mathcal{P} = \{\mathbf{p}_i\}_{i=1\dots n}$  (i.e., a 3D reconstruction of an RGB image  $\mathcal{I}$ ), our method learns the shape Laplacian to compute the deformation energy matrix  $A \in \mathbb{R}^{n \times n}$  of the underlying volume of  $\mathcal{P}$ . The proposed framework mainly consists of three modules: (1) Feature Extraction Module, (2) Cotangent Laplacian Prediction Module, and (3) Inverse Mass Prediction Module. Feature Extraction Module first extracts per-point features  $\mathcal{F} = \{\mathbf{f}_i\}_{i=1\dots n}$  of  $\mathcal{P}$ . Given  $\mathcal{P}$  and  $\mathcal{F}$ , Cotangent Laplacian Prediction Module estimates the cotangent Laplacian matrix  $L \in \mathbb{R}^{n \times n}$  and Inverse Mass Prediction Module predicts the inverse mass matrix  $M^{-1} \in \mathbb{R}^{n \times n}$  of  $\mathcal{P}$ . The final deformation energy  $A$  can be obtained as  $LM^{-1}L$ .

where  $A = LM^{-1}L$  is the deformation energy matrix of the source mesh  $\mathcal{M}$ ,  $L \in \mathbb{R}^{n \times n}$  is the cotangent Laplacian matrix of  $\mathcal{M}$ , and  $M \in \mathbb{R}^{n \times n}$  is a lumped mass matrix whose diagonal elements represent the volume associated with each of the vertices in  $\mathcal{M}$ . Both  $L$  and  $M$  are defined on *volume* information of a shape and thus cannot be directly computed from a point cloud. Although it is possible to recover its volume using a mesh reconstruction method (e.g., [3, 5, 16, 24, 25, 30, 38]), it may introduce topological noises that result in erroneous deformation weights. To address this issue, we argue that *learning* for the deformation energy  $A$  enables more accurate and robust deformation.

### 3.2. Laplacian Learning Network

We now introduce our network that can learn the deformation energy matrix  $A \in \mathbb{R}^{n \times n}$  of the underlying volume of a 3D reconstruction represented as a point cloud  $\mathcal{P} = \{\mathbf{p}_i\}_{i=1\dots n}$ . Instead of directly estimating the matrix  $A$ , we propose to learn the cotangent Laplacian matrix  $L \in \mathbb{R}^{n \times n}$  and the inverse mass matrix  $M^{-1} \in \mathbb{R}^{n \times n}$  such that the deformation energy can be later recovered as  $A = LM^{-1}L$ . The first motivation behind this design is to automatically guarantee the positive semi-definiteness of the predicted  $A$ ; it is non-trivial to enforce this property when directly predicting a squared matrix using a neural network. The second motivation is to better learn and enforce the symmetry and sparsity structure of the cotangent Laplacian matrix  $L$  to allow a more accurate prediction.

In what follows, we explain the detailed architecture of our Laplacian learning network, which consists of three modules: (1) Feature Extraction Module, (2) Cotangent Laplacian Prediction Module, and (3) Inverse Mass Prediction Module.

**Feature Extraction Module** Given a point cloud  $\mathcal{P} =$

$\{\mathbf{p}_i\}_{i=1\dots n}$ , Feature Extraction Module extracts a point cloud feature  $\mathcal{F} = \{\mathbf{f}_i\}_{i=1\dots n}$ , where  $\mathbf{f}_i \in \mathbb{R}^d$  denotes a per-point feature vector corresponding to  $\mathbf{p}_i$ . For network architecture, we adopt Point Transformer [55] that can extract permutation and cardinality-invariant feature of a point cloud via self-attention operations.

**Cotangent Laplacian Prediction Module** Given the point cloud  $\mathcal{P}$  and the point cloud feature  $\mathcal{F}$ , Cotangent Laplacian Prediction Module learns the cotangent Laplacian matrix  $L \in \mathbb{R}^{n \times n}$  of the underlying volume of  $\mathcal{P}$ . Following the definition of the cotangent Laplacian,  $L$  is desired to be a symmetric matrix whose element  $L_{ij}$  is non-zero only if  $\mathbf{p}_i$  and  $\mathbf{p}_j$  are topologically connected by an edge. Since  $L$  is highly sparse, and the point cloud cardinality  $n$  is typically large (e.g., in the order of thousands), it is inefficient to learn the relationship between *all* point pairs. Therefore, we use a Euclidean prior to select the initial point pair candidates which are probable to have local connectivity – thus correspond to *non-zero* entries of  $L$ . More specifically, for each point  $\mathbf{p}_i \in \mathcal{P}$ , we compute  $k$  nearest neighbor points based on Euclidean distance. We then couple each of the neighbor points with the source point  $\mathbf{p}_i$  to form point pair candidates. We denote this sampling strategy by *KNN-Based Point Pair Sampling (KPS)*. We note that using KPS not only improves the inference time but also helps the network training by alleviating an imbalanced regression problem, as a significant number of zero-valued regression targets can be filtered out before training. In our experiments, we empirically set  $k$  as 32.

We now extract a feature for each point pair candidate as follows:

$$\mathbf{g}_m = (\gamma_1(\mathbf{p}_i, \mathbf{p}_j), \gamma_2(\mathbf{f}_i, \mathbf{f}_j)). \quad (3)$$

where  $\{\mathbf{p}_i, \mathbf{p}_j\}$  is the  $m$ -th point pair candidate, and  $\gamma_1(\cdot)$

and  $\gamma_2(\cdot)$  are *symmetric* functions used for the pairwise feature aggregation. Since the matrix  $L$  is desired to be *symmetric*, the pairwise feature aggregations must be symmetric to produce the identical features for inputs  $(\mathbf{p}_i, \mathbf{p}_j)$  and  $(\mathbf{p}_j, \mathbf{p}_i)$  – thus guaranteeing the equality between the later predictions for  $L_{ij}$  and  $L_{ji}$ . We empirically choose absolute difference and element-wise multiplication for  $\gamma_1$  and  $\gamma_2$ , respectively. (Refer to Section 4 for the ablation study.)

Next, we estimate the entry of the cotangent Laplacian  $L$  that corresponds to each point pair candidate as follows:

$$L_{ij} = \alpha(\mathbf{g}_m) \odot \phi(\mathbf{g}_m), \quad (4)$$

where  $i$  and  $j$  are the indices of the  $m$ -th point pair candidate,  $\phi(\cdot)$  is a function that outputs a real-valued scalar, and  $\alpha(\cdot)$  is a function predicting a *weight*  $W_{ij} \in [0, 1]$  indicating whether the element  $L_{ij}$  in the cotangent Laplacian is a non-zero value. The addition of the weight prediction is our *key* to improving the accuracy of the deformation. As previously mentioned,  $L$  has non-zero elements only for the point pairs that are connected through edges; this indicates that  $L$  is a very sparse matrix. Capturing its sparsity structure is essential, since it encodes information about the topology of the shape. To better model such sparsity structure, we introduce the function  $\alpha(\cdot)$  and predict the weight indicating whether each element in  $L$  is zero or not. We empirically find that this additional sparsity structure prediction significantly improves the quality of deformations, as shown in Section 4. We also remark that our architecture of the Cotangent Laplacian Prediction Module — taking point pairs as input — has the advantage of learning well even from a scarce training dataset since the number of training examples increases by  $k$  with the number of sampled points.

In test time, the entire cotangent Laplacian matrix  $L$  can be computed by predicting the off-diagonal elements  $\{L_{ij}\}_{i \neq j}$  in parallel via batch processing and calculating diagonal elements  $\{L_{ii}\}_{i=1, \dots, n}$  as the minus sum of the off-diagonal elements in a row – to follow the definition of the cotangent Laplacian.

**Inverse Mass Prediction Module** Given the point cloud  $\mathcal{P}$  and the point cloud feature  $\mathcal{F}$ , Inverse Mass Prediction Module learns the inverse mass matrix  $M^{-1} \in \mathbb{R}^{n \times n}$  of the underlying volume of  $\mathcal{P}$ . For each point in  $\mathcal{P}$ , we first concatenate its 3D coordinate  $\mathbf{p}_i$  and the corresponding feature  $\mathbf{f}_i$  to generate a new per-point feature  $\mathbf{f}'_i = [\mathbf{p}_i; \mathbf{f}_i] \in \mathbb{R}^{3+d}$ . Then, the inverse mass of each point  $\mathbf{p}_i$  can be predicted as  $\omega(\mathbf{f}'_i)$ , where  $\omega$  is a function instantiated as an MLP. In test time, the entire inverse mass matrix  $M^{-1}$  can be computed by predicting the diagonal elements in a batch and padding the off-diagonal entries as zero.

**Loss Functions.** We train our network using the supervision obtained from the ground truth mesh corresponding to

the input point cloud  $\mathcal{P}$ . The overall loss function for our framework can be written as follows:

$$\mathcal{L} = \sum_{i,j \in \mathcal{K}(\mathcal{P})} \frac{1}{|\mathcal{K}(\mathcal{P})|} \|L_{ij} - L_{ij}^{\text{gt}}\|_1 + \sum_{i,j \in \mathcal{K}(\mathcal{P})} \frac{\lambda_W}{|\mathcal{K}(\mathcal{P})|} \|W_{ij} - W_{ij}^{\text{gt}}\|_1 + \sum_i \frac{\lambda_{M^{-1}}}{|\mathcal{P}|} \|M_{ii}^{-1} - M_{ii}^{-1\text{gt}}\|_1, \quad (5)$$

where  $L^{\text{gt}}$ ,  $W^{\text{gt}}$ , and  $M^{-1\text{gt}}$  are ground truth matrices for  $L$ ,  $W$  and  $M^{-1}$ , respectively.  $\mathcal{K}(\mathcal{P})$  denotes a set of point pair indices sampled by KPS. For the weights of the loss terms  $\lambda_W$  and  $\lambda_{M^{-1}}$ , we empirically choose 100 and 1, respectively.

**Deformation Weight Computation.** In test time, once we predict the cotangent Laplacian matrix  $L$  and the inverse mass diagonal matrix  $M^{-1}$  to recover the deformation energy matrix  $A = LM^{-1}L$ , the linear deformation blending weights in Equation 3.1 for the set of deformation handles are calculated by solving the quadratic programming problem. Note that, while the deformation blending weights are *dependent* on the given set of deformation handles, the quantities we predict ( $L$  and  $M^{-1}$ ) are *not dependent* on the deformation handles. Hence, our framework can allow the user to choose an arbitrary set of deformation handles and move them freely without any restriction.

### 3.3. Implementation Details

**Network Architectures.** For Feature Extraction Module, we adopt Point Transformer [55] architecture to output a 64-dimensional feature for each point  $\mathbf{p}_i$ . For  $\alpha$  and  $\phi$  in Cotangent Laplacian Prediction Module and  $\omega$  in Inverse Mass Prediction Module, we use multilayer perceptron (MLP) architecture composed of three fully-connected layers, each of them followed by batch normalization, dropout and LeakyReLU – except for the last activation. For  $\alpha$  and  $\omega$ , we adopt sigmoid as the final activation, while we use no activation for the last layer of  $\phi$ . The output feature dimensions for each of the layers are 128, 256, and 1, respectively. We purposely design a lightweight architecture for  $\alpha$ ,  $\phi$ , and  $\omega$ , as they are desired to operate in parallel to batch-compute the entire cotangent Laplacian matrix  $L$  and the inverse mass matrix  $M^{-1}$ . Due to the space limit, please refer to supplementary material for more details of our network training (e.g., learning rate, batch size).

**Deformation Weight Computation.** We use the implementation of libigl [22] to compute bounded biharmonic weights. This implementation solves the quadratic programming problem in Equation 3.1 by executing Mosek [34], which is a software package for solving large-scale optimization problems. We also utilize libigl library for computing ground truth cotangent stiffness matrix  $L$  and mass matrix  $M$  to obtain supervision for our network training.

Table 1. **Quantitative comparison of 3D point cloud deformation results on DFAUST [4] dataset.** The point handles for a source shape are selected using farthest point sampling.

Number of Handles	Metric	PSR [25]	APSS [11]	BPA [3]	DeepSDF [38]	DGP [50]	MIER [30]	PCDLap [2]	NMLap [42]	Ours
16	Weight L1 ( $\times 100$ ) ↓	3.86	3.46	4.32	2.66	4.15	3.26	3.53	3.34	<b>2.10</b>
	Shape CD ( $\times 100$ ) ↓	3.84	3.04	3.83	2.61	4.09	3.16	2.97	4.04	<b>1.81</b>
	Shape HD ( $\times 0.1$ ) ↓	1.81	1.31	1.73	0.48	2.85	1.13	<b>0.42</b>	0.43	<b>0.42</b>
32	Weight L1 ( $\times 100$ ) ↓	3.08	1.53	1.72	1.38	3.01	1.09	1.54	1.49	<b>1.06</b>
	Shape CD ( $\times 100$ ) ↓	5.12	2.39	2.29	2.15	4.55	1.26	2.06	2.18	<b>1.45</b>
	Shape HD ( $\times 0.1$ ) ↓	3.39	0.79	0.83	0.57	4.25	2.10	<b>0.53</b>	<b>0.53</b>	<b>0.53</b>

## 4. Experiments

In this section, we experimentally validate the effectiveness of our method. In Section 4.1, we first evaluate our method on a large-scale 3D point cloud dataset for *quantitative* comparisons with the alternative methods. In Section 4.2, we present our results of 3D-aware image deformation. Lastly, we report our ablation study in Section 4.3.

### 4.1. 3D Point Cloud Deformation

To *quantitatively* evaluate our deformed shape quality, we first test our method on a large-scale 3D point cloud dataset (i.e., DFAUST [4]). We also report our deformation results on 3D *partial* point clouds (i.e., SHREC’16 [7]).

#### 4.1.1 Experiments on 3D Point Clouds

In this subsection, we perform quantitative evaluation on DFAUST [4] dataset, which contains 40,000 3D human body meshes that represent 129 motion sequences of 10 real humans. While no image is used in this experiment, we consider the point clouds generated from the meshes are 3D reconstructions for the sake of *quantitative* comparisons. To prepare the data, we first subsample the shapes in each sequence by a factor of 16 due to the high frame rate of the motion sequences. We then use the shapes of eight humans for training and the shapes of the remaining two humans for testing. For evaluation metrics, we use (1) L1 distance (Weight L1) between the ground truth and the predicted deformation weights, (2) Chamfer distance (Shape CD) and (3) Hausdorff distance (Shape HD) between the ground truth and the predicted deformed point clouds. For more detailed information about the metric computation, please refer to our supplementary section.

We compare our 3D point cloud deformation results with those computed using the mesh reconstruction methods – Screened Poisson Surface Reconstruction (PSR) [25], Algebraic Point Set Surfaces (APSS) [11], Ball-Pivoting Algorithm (BPA) [3], DeepSDF [38], Deep Geometric Prior (DGP) [50], and Meshing Point Clouds with Intrinsic-Extrinsic Ratio (MIER). For implicit reconstruction methods that do not directly output a shape in mesh representation, we additionally apply Marching Cubes [32] to extract a surface mesh. We then convert the surface mesh into a tetra-

hedron mesh using [16] and [17] to allow the computation of the shape Laplacian of the underlying *volume*. We also compare our method with point cloud Laplacians – PCD Laplace (PCDLap) [2] and Nonmanifold Laplacians (NMLap) [42]. In Table 1, our method is shown to outperform all the alternative methods in terms of the three evaluation metrics. This indicates that our method can model handle-based deformations in a more accurate manner by directly *learning* the shape Laplacian from ground truth supervision. In Figure 3, we also provide the visualizations of the deformed point clouds. Our deformations can better preserve the local geometric smoothness of the source shape; they also more accurately match to the ground truth deformations.

#### 4.1.2 Experiments on 3D Partial Point Clouds

We also perform experiments on SHREC’16 [7] dataset, which consists of partial shapes of animal objects, such as dog, cat, and wolf. It contains two sub-datasets – *cuts* and *holes* – whose names exemplify different types of partiality. As SHREC’16 is originally released as a benchmark for partial correspondence problem, (1) the ground truth *full* shape for each partial shape and (2) their point-wise correspondences are also available. For our experiment, we randomly split the *cuts* and *holes* samples into the training and test sets with an 8:2 ratio.

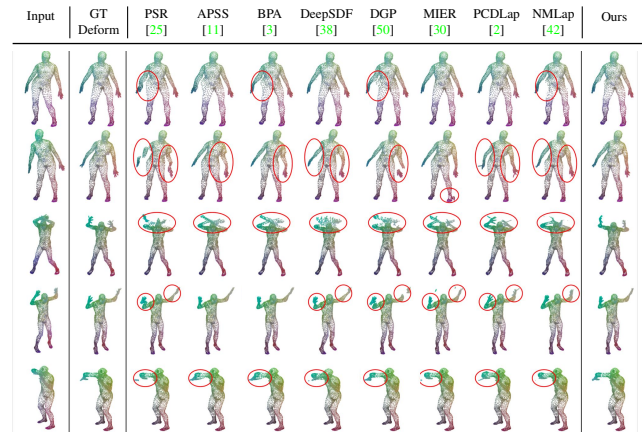


Figure 3. **Qualitative comparison of 3D point cloud deformation results on DFAUST [4] dataset (best viewed with 200% zoom-in).** These deformations are modeled using 32 control handles selected via farthest point sampling.



Figure 4. 3D-aware deformation of Mixamo [19] character images. 32 control points and their new positions are manually chosen.

As our method takes a learning-based approach to estimate the shape Laplacian, it is possible to train our network to predict the shape Laplacian corresponding to the ground truth *full* shape from a *partial* shape. In this way, we can transfer the knowledge of the shape intrinsics of the full shape into the input partial shape to allow more robust shape Laplacian prediction. Similar to the previous experiment, we evaluate our results on the L1 distance between the ground truth and the predicted deformation weights. As shown in Table 2, our method results in more accurate handle-based deformation weights compared to the case of calculating the shape Laplacian from a mesh constructed by APSS [11], Marching Cubes [32] and a volume conversion method [16].

Table 2. **Quantitative comparison of 3D partial point cloud deformation results on SHREC’16 [7] Dataset.** All weights are computed with respect to 16 point handles selected by farthest point sampling. The ground truth deformation weights are calculated using the topology of the *full* meshes that correspond to partial shapes.

Method	Weight L1 ( $\times 100$ ) $\downarrow$
APSS [11] + Marching Cubes [32]	5.15
Ours	4.47

## 4.2. 3D-Aware Image Deformation

We now present our results of 3D-aware image deformation on character and clothed human images.

### 4.2.1 Experiments on Character Images

We present our qualitative results of 3D-aware image deformation on *Mousey* and *Michelle* characters in Mixamo [19] dataset. We remark that images of such characters cannot be manipulated using the existing methods of parametric model fitting [31, 58] or human pose transfer [6, 29, 33]. In this experiment, we directly utilize the Laplacian Learning Network trained on DFAUST [4] (which is used in Section 4.1.1) and evaluate the model on 3D reconstruc-

tions of Mixamo character images – to demonstrate the *generalizability* of our network across different object categories. However, for 2D-to-3D reconstruction, the pre-trained model of PIFu [39] has not been shown to generalize well to character images. Therefore, we populate separate training images by rendering the Mixamo models and train PIFu from scratch. We create images by rendering the 3D character models with randomly sampled animation frames, camera and light source positions. The total number of images in the dataset is 45,000, and the ratio between the training and test sets is 9:1.

Figure 4 shows the results of our 3D-aware image deformation on the images of *Mousey* and *Michelle* characters. Our image deformation method can plausibly model the changes in the pose of the characters *in a 3D-aware manner*. This demonstrates the potential of our method in interactive image editing applications, as it can provide an easy and intuitive interface (i.e., control handles) for image manipulation. We also emphasize again that our Laplacian Learning Network was trained not on Mixamo [19] but on DFAUST [4] that only consists of human body models. Hence, these results show that our network can be *zero-shot generalizable* to the shapes in other categories by learning the *local* geometric information of a shape. Please note that, while two examples are shown in Figure 4, more various examples are presented in the supplementary section.

### 4.2.2 Experiments on Clothed Human Images

We also report our results of 3D-aware image deformation on clothed human images. We use RenderPeople [10] and DeepHuman [56], which are real-world human model datasets containing textured 3D meshes. Similar to the previous experiment, we prepare the data by rendering the textured 3D human meshes with random viewpoints and lighting. Since RenderPeople provides only 9 models in public, we only show the *qualitative* results using it. We instead use DeepHuman dataset for *quantitative* evaluation, which has approximately 7000 models but with low-quality textures.

During the experiment, we again use the Laplacian Learning Network trained on DFAUST [4] dataset, which consists of 3D human body models. For 2D-to-3D reconstruction using PIFu [39], we use the official pre-trained model.




Figure 5. **3D-aware deformation of RenderPeople [10] human images.** Blue dots indicate the user-defined control handles and red arrows show the translations applied to the selected handles.

Figures 1 and 5 show the qualitative results of 3D-aware image deformation on RenderPeople [10] dataset. 32 control points are manually picked and manipulated. In Figure 1, our method is shown to produce more plausible image deformations compared to the case of directly computing the shape Laplacian from a mesh reconstructed by PIFu [39] (whose topology prediction is based on Marching Cubes [32]). Figure 5 shows our image deformation examples in user-interaction scenarios. We also provide the quantitative evaluation results of our learned deformation weights on DeepHuman [56] dataset in Table 3. We again compare our results to the deformation weights obtained using the shape Laplacian computed directly from a mesh reconstructed by PIFu. Our method is shown to estimate bounded biharmonic weights in a more accurate manner.

Table 3. **Quantitative comparison of image deformation results on DeepHuman [56] Dataset.** All weights are computed with respect to the same control handle configurations: 32 point handles selected by farthest point sampling. The ground truth deformation weights are calculated using the geometry intrinsics provided by the ground truth 3D meshes of DeepHuman Dataset.

Method	Weight L1 ( $\times 100$ ) $\downarrow$
Marching Cubes [32]	3.09
Ours	<b>2.14</b>

### 4.3. Ablation Study

In this section, we perform ablation study to compare different design options for our Laplacian Learning Network. We first explore the effectiveness of the KNN-Based Point Pair Sampling (KPS) and the weight function  $\alpha$  that models the sparsity structure of the cotangent Laplacian. As can be seen in Table 4 (columns 2 and 3), we observe significant performance degradation when KPS or  $\alpha$  is removed.

In the same table (columns 4 and 5), we additionally evaluate the effects of using different symmetric functions for  $\gamma_1$  and  $\gamma_2$  in the Cotangent Laplacian Prediction Module. We observe that the best performance can be achieved when the proposed functions are used – in which  $\gamma_1$  and  $\gamma_2$  are instantiated as absolute difference and element-wise multiplication, respectively. Please refer to the supplementary section for (1) the qualitative results of this ablation study and (2) a comparison with the network that directly learns bounded biharmonic weights – instead of learning the shape Laplacian.

Table 4. **Ablation study.**  $-KPS$  and  $-\alpha$  indicate settings where KPS or  $\alpha$  is removed from our method, respectively. EM Only and AD Only denote settings where  $\gamma_1$  and  $\gamma_2$  are both instantiated as element-wise multiplication and absolute difference, respectively. All the other experimental setups (e.g., datasets, control handle configurations) are the same as in the DFAUST experiment in Section 4.1.

Metric	KPS	$-\alpha$	EM Only	AD Only	Ours
Weight L1 ( $\times 100$ ) $\downarrow$	2.79	4.95	3.54	2.35	<b>2.10</b>
Shape CD ( $\times 100$ ) $\downarrow$	2.46	4.07	3.27	2.20	<b>1.81</b>
Shape HD ( $\times 0.1$ ) $\downarrow$	0.43	1.23	0.69	0.47	<b>0.42</b>

## 5. Conclusion and Future Work

We propose a method for 3D-aware deformation of 2D images via learning the shape Laplacian. To this end, we introduce a novel network architecture that can learn the shape Laplacian from a 3D point cloud reconstructed from a 2D image. To the best of our knowledge, this is the first study to demonstrate that a learning-based approach can be effective in predicting the shape Laplacian of the underlying volume of a point cloud.

**Negative Societal Impacts** Our method can be potentially abused to create *deepfake*, while the scope of manipulating images is limited to 3D-aware deformation.

**Limitations and Future Work** Since our method models deformations based on bounded biharmonic weights, it cannot guarantee avoiding self-intersections. In addition, the current pipeline learns the shape Laplacian only from the direct supervisions (the L1-losses of the matrices) but not from the outputs of the deformations. We plan to further investigate the ways to improve the deformation quality.

**Acknowledgments.** We would like to thank Duygu Ceylan for helpful discussions. This work is in part supported by KAIA grant (22CTAP-C163793-02) funded by the Korea government(MOLIT) and NST grant (CRC 21011) funded by the Korea government(MSIT). M. Sung also acknowledges the support by NRF grant (2021R1F1A1045604) funded by the Korea government(MSIT), Technology Innovation Program (20016615) funded by the Korea government(MOTIE), and grants from the Adobe and KT corporations.



## References

- [1] Matan Atzmon, David Novotny, Andrea Vedaldi, and Yaron Lipman. Augmenting implicit neural shape representations with explicit deformation fields. *CoRR*, abs/2108.08931. 2
- [2] Mikhail Belkin, Jian Sun, and Yusu Wang. Constructing laplace operator from point clouds in  $\mathbb{R}^d$ . In *SODA*, 2009. 3, 6
- [3] Fausto Bernardini, Joshua Mittleman, Holly Rushmeier, Claudio Silva, and Gabriel Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE TVCG*, 1999. 3, 4, 6
- [4] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *CVPR*, 2017. 2, 6, 7, 8
- [5] Jean-Daniel Boissonnat and Bernhard Geiger. Three-dimensional reconstruction of complex shapes based on the delaunay triangulation. In *Biomedical image processing and biomedical visualization*, 1993. 3, 4
- [6] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei Efros. Everybody dance now. In *ICCV*, 2019. 1, 7
- [7] Luca Cosmo, Emanuele Rodola, Michael M. Bronstein, Andrea Torsello, Daniel Cremers, and Y. Sahillioglu. SHREC'16: Partial matching of deformable shapes. 2016. 6, 7
- [8] Haoqiang Fan, Hao Su, and Leonidas Guibas. A point set generation network for 3D object reconstruction from a single image. In *CVPR*, 2017. 2
- [9] Marc-André Gardner, Yannick Hold-Geoffroy, Kalyan Sunkavalli, Christian Gagné, and Jean-Francois Lalonde. Deep parametric indoor lighting estimation. In *ICCV*, 2019. 1
- [10] Renderpeople GmbH. RenderPeople. <https://renderpeople.com/>. 7, 8
- [11] Gaël Guennebaud and Marko Gross. Algebraic point set surfaces. In *SIGGRAPH*, 2007. 6, 7
- [12] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-PIFu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *NeurIPS*, 2020. 2
- [13] Derek Hoiem, Alexei A Efros, and Martial Hebert. Automatic photo pop-up. In *SIGGRAPH*, 2005. 1
- [14] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matt Fisher, Emiliano Gambaretto, Sunil Hadap, and Jean-Francois Lalonde. A perceptual measure for deep single image camera calibration. In *CVPR*, 2018. 1
- [15] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Sunil Hadap, Emiliano Gambaretto, and Jean-François Lalonde. Deep indoor illumination estimation. In *CVPR*, 2017. 1
- [16] Yixin Hu, Teseo Schneider, Bolun Wang, Denis Zorin, and Daniele Panozzo. Fast tetrahedral meshing in the wild. In *SIGGRAPH*, 2020. 3, 4, 6, 7
- [17] Jingwei Huang, Hao Su, and Leonidas Guibas. Robust watertight manifold surface generation method for shapenet models. *CoRR*, abs/1802.01698, 2018. 3, 6
- [18] Jingwei Huang, Yichao Zhou, and Leonidas Guibas. ManifoldPlus: A robust and scalable watertight manifold surface generation method for triangle soups. *CoRR*, abs/2005.11621, 2020. 3
- [19] Adobe Systems Inc. Mixamo. <https://www.mixamo.com>. 7
- [20] Dominic Jack, Jhony K Pontes, Sridha Sridharan, Clinton Fookes, Sareh Shirazi, Frederic Maire, and Anders Eriksson. Learning free-form deformations for 3D object reconstruction. In *ACCV*, 2018. 2
- [21] Alec Jacobson, Ilya Baran, Jovan Popovic, and Olga Sorkine. Bounded biharmonic weights for real-time deformation. In *SIGGRAPH*, 2011. 1, 2, 3
- [22] Alec Jacobson, Daniele Panozzo, et al. libigl: A simple C++ geometry processing library. <https://libigl.github.io/>. 5
- [23] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 2
- [24] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *SGP*, 2006. 3, 4
- [25] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM TOG*, 2013. 3, 4, 6
- [26] Natasha Kholgade, Tomas Simon, Alexei Efros, and Yaser Sheikh. 3D object manipulation in a single photograph using stock 3D models. In *SIGGRAPH*, 2014. 1
- [27] Hyunjoon Lee, Eli Shechtman, Jue Wang, and Seungyong Lee. Automatic upright adjustment of photographs with robust camera calibration. *IEEE TPAMI*, 2014. 1
- [28] Jiahui Lei, Srinath Sridhar, Paul Guerrero, Minhyuk Sung, Niloy Mitra, and Leonidas J. Guibas. Pix2Surf: Learning parametric 3D surface models of objects from images. In *ECCV*, 2020. 2
- [29] Yining Li, Chen Huang, and Chen Change Loy. Dense intrinsic appearance flow for human pose transfer. In *CVPR*, 2019. 1, 7
- [30] Minghua Liu, Xiaoshuai Zhang, and Hao Su. Meshing point clouds with predicted intrinsic-extrinsic ratio guidance. In *ECCV*, 2020. 3, 4, 6
- [31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *SIGGRAPH Asia*, 2015. 7
- [32] William E Lorensen and Harvey E Cline. Marching Cubes: A high resolution 3D surface construction algorithm. In *SIGGRAPH*, 1987. 1, 3, 6, 7, 8
- [33] Liqian Ma, Xu Jia, Qianru Sun, Bernt Schiele, Tinne Tuytelaars, and Luc Van Gool. Pose guided person image generation. In *NeurIPS*, 2017. 1, 7
- [34] Mosek. Mosek. <https://www.mosek.com>. 5
- [35] Sanjeev Muralikrishnan, Siddhartha Chaudhuri, Noam Aigerman, Vladimir Kim, Matthew Fisher, and Niloy Mitra. GLASS: Geometric latent augmentation for shape spaces. *CoRR*, abs/2108.03225. 2
- [36] Simon Niklaus, Long Mai, Jimei Yang, and Feng Liu. 3D ken burns effect from a single image. *ACM TOG*, 2019. 1
- [37] Xingang Pan, Bo Dai, Ziwei Liu, Chen Change Loy, and Ping Luo. Do 2D GANs know 3D shape? unsupervised 3D shape reconstruction from 2D image GANs. In *ICLR*, 2021. 1
- [38] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning

- continuous signed distance functions for shape representation. In *CVPR*, 2019. 3, 4, 6
- [39] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 1, 2, 3, 7, 8
- [40] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, 2020. 2
- [41] Lars Schnyder, Manuel Lang, Oliver Wang, and Aljoscha Smolic. Depth image based compositing for stereo 3D. In *3DTV-CON*, 2012. 1
- [42] Nicholas Sharp and Keenan Crane. A laplacian for nonmanifold triangle meshes. In *SGP*, 2020. 3, 6
- [43] Minhyuk Sung, Zhenyu Jiang, Panos Achlioptas, Niloy J Mitra, and Leonidas Guibas. DeformSyncNet: Deformation transfer via synchronized shape deformation spaces. In *SIGGRAPH Asia*, 2020. 2
- [44] Edgar Tretschk, Ayush Tewari, Michael Zollhöfer, Vladislav Golyanik, and Christian Theobalt. DEMEA: Deep mesh autoencoders for non-rigidly deforming objects. In *ECCV*, 2020. 2
- [45] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J. Guibas. Normalized object coordinate space for category-level 6D object pose and size estimation. In *CVPR*, 2019. 2
- [46] 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. 2
- [47] Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. 3DN: 3D deformation network. In *CVPR*, 2019. 2
- [48] Yu Wang, Alec Jacobson, Jernej Barbič, and Ladislav Kavan. Linear subspace design for real-time shape deformation. In *SIGGRAPH*, 2015. 1, 2
- [49] Chao Wen, Yinda Zhang, Zhuwen Li, and Yanwei Fu. Pixel2Mesh++: Multi-view 3D mesh generation via deformation. In *ICCV*, 2019. 2
- [50] Francis Williams, Teseo Schneider, Claudio Silva, Denis Zorin, Joan Bruna, and Daniele Panozzo. Deep geometric prior for surface reconstruction. In *CVPR*, 2019. 6
- [51] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3D objects from images in the wild. In *CVPR*, 2020. 1
- [52] Wenqi Xian, Zhengqi Li, Noah Snavely, Matthew Fisher, Jonathan Eisenman, and Eli Shechtman. UprightNet: Geometry-aware camera orientation estimation from single images. In *ICCV*, 2019. 1
- [53] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. DISN: Deep implicit surface network for high-quality single-view 3D reconstruction. In *NeurIPS*. 2019. 2
- [54] Jinsong Zhang, Kalyan Sunkavalli, Yannick Hold-Geoffroy, Sunil Hadap, Jonathan Eisenman, and Jean-François Lalonde. All-weather deep outdoor lighting estimation. In *CVPR*, 2019. 1
- [55] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, 2021. 4, 5
- [56] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. DeepHuman: 3D human reconstruction from a single image. In *ICCV*, 2019. 7, 8
- [57] Matthias Ziegler, Andreas Engelhardt, Stefan Müller, Joachim Keinert, Frederik Zilly, Siegfried Foessel, and Katja Schmid. Multi-camera system for depth based visual effects and compositing. In *CVMP*, 2015. 1
- [58] Silvia Zuffi, Angjoo Kanazawa, David W. Jacobs, and Michael J. Black. 3D Menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, 2017. 7