

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/346668849>

Multiscale Mesh Deformation Component Analysis with Attention-based Autoencoders

Preprint · December 2020

CITATIONS

0

READS

282

6 authors, including:



Jie Yang

Chinese Academy of Sciences

22 PUBLICATIONS 468 CITATIONS

[SEE PROFILE](#)



Lin Gao

Chinese Academy of Sciences

96 PUBLICATIONS 2,086 CITATIONS

[SEE PROFILE](#)



Qingyang Tan

University of the Arts London

11 PUBLICATIONS 381 CITATIONS

[SEE PROFILE](#)



Shihong Xia

Chinese Academy of Sciences

104 PUBLICATIONS 1,850 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Virtual Human Synthesis [View project](#)



Shape Deformation and Morphing [View project](#)

Multiscale Mesh Deformation Component Analysis with Attention-based Autoencoders

Jie Yang, Lin Gao*, Qingyang Tan, Yihua Huang, Shihong Xia and Yu-Kun Lai

Abstract—Deformation component analysis is a fundamental problem in geometry processing and shape understanding. Existing approaches mainly extract deformation components in local regions at a similar scale while deformations of real-world objects are usually distributed in a multi-scale manner. In this paper, we propose a novel method to extract multiscale deformation components automatically with a stacked attention-based autoencoder. The attention mechanism is designed to learn to softly weight multi-scale deformation components in active deformation regions, and the stacked attention-based autoencoder is learned to represent the deformation components at different scales. Quantitative and qualitative evaluations show that our method outperforms state-of-the-art methods. Furthermore, with the multiscale deformation components extracted by our method, the user can edit shapes in a coarse-to-fine fashion which facilitates effective modeling of new shapes.

Index Terms—Multi-Scale, Shape Analysis, Attention Mechanism, Sparse Regularization, Stacked Auto-Encoder



1 INTRODUCTION

WITH the development of 3D scanning and modeling technology, 3D mesh collections are becoming much more popular. These mesh models usually use fixed vertex connectivity with variable vertex positions to characterize different shapes. Analyzing these mesh model collections to extract meaningful components and using these components for new model generation are key research problems in these areas. Some works [1], [2], [3], [4] propose to extract deformation components from mesh data sets. They mainly focus on extracting local deformation components with sparse regularization at a uniform scale. However, real-world objects deform at multiple scales. For example, a person may have different facial expressions which are more localized deformations on the face, but the whole body can also be bent, which is a larger scale deformation.

Multi-scale techniques are getting increasingly popular in various fields. In Finite Element Methods, multiscale analysis is widely used [5], [6], [7]. In the spectral geometry field, research works apply multiscale technology on the deformation representation [8], physics-based simulation of deformable objects [9], and surface registration [10] by analyzing the non-isometric global and local (multiscale) deformation. Moreover, for shape editing, multiscale technology also enables modeling rich facial expressions on human faces [11].

Such multiscale deformation components are especially useful to support model editing from coarse level to fine level. One motivation of this work is to achieve editing consistent with perceptual semantics by modifying shapes at suitable scales. The

user would be able to make rough editing of the overall shape at a large scale, as well as localized modifications to surface details at a small scale. Inspired by the recent advances in image processing with attention mechanism [12], the attention is formulated to focus on specific regions in our approach.

We propose a novel autoencoder architecture to extract multiscale local deformation components from shape deformation datasets. Our network structure is based on the mesh-based convolutional autoencoder architecture and also uses an effective representation of the shapes [13] as input which is able to encode large-scale deformations. In this work, a stacked autoencoder architecture is proposed such that the network can encode the residual value of the former autoencoder with the attention mechanism, which helps to separate the deformations into different scales and extract multiscale local deformation components. The network architecture is shown in Fig. 1. We further utilize a sparsity constraint on the parameters of the fully connected layers to keep the deformation components localized. The autoencoder architecture ensures the extracted deformation components are suitable for multiscale shape editing and helps reconstruct high quality shapes with less distortion.

Our contributions are twofold:

- To the best of our knowledge, this is the first work that automatically extracts multiscale deformation components from a deformed shape collection. With these extracted components, the user can edit the 3D mesh shape in a coarse-to-fine fashion, which makes 3D modeling much more effective.
- To achieve this, we propose a novel deep architecture involving attentional stacked autoencoders. The attention mechanism is designed for learning the soft weights that help extract multiscale deformation components in the shape analysis and the stacked autoencoders are used to decompose the deformation of shape collections into different shape components with different scales.

All the components of the network are tightly integrated and help each other. The attention mechanism makes the follow-

* Corresponding author is Lin Gao.

- J. Yang, L. Gao, Y. Huang and S. Xia are with the Beijing Key Laboratory of Mobile Computing and Pervasive Device, Institute of Computing Technology, Chinese Academy of Sciences and also with the University of Chinese Academy of Sciences, Beijing, China.
E-mail: {yangjie01, gaolin, huangyihua20g, xsh}@ict.ac.cn
- Q. Tan is with University of Maryland, College Park, Maryland, USA.
E-mail: qytan@cs.umd.edu
- Y.-K. Lai is with Visual Computing Group, School of Computer Science and Informatics, Cardiff University, Wales, UK.
E-mail: LaiY4@cardiff.ac.uk

up autoencoders focus on a specific region, to allow extracting smaller-scale local deformation components. Extensive comparisons prove that our method extracts more meaningful deformation components than state-of-the-art methods.

In Sec. 2, we review the most related work. We then give a brief description of the input features used in our method in Sec. 3, and present detailed description of our novel autoencoder with the attention mechanism including implementation details in Sec. 4. Finally, we present experimental results, including extensive comparisons with state-of-the-art methods in Sec. 5 and draw conclusions in Sec. 6.

2 RELATED WORK

Mesh deformation component analysis has attracted significant interest in the research of shape analysis and data-driven mesh deformation. Many data-driven methods for editing of either man-made objects [14], [15], [16] or general deformable surfaces [13], [17], [18] benefit from extracted deformation components. Our work shares the same interest as theirs, aiming to assist users to edit shapes efficiently. Although our focus is to extract more meaningful multiscale deformation components automatically, it can be incorporated into existing data-driven deformation methods. In the following, we will review work most related to ours.

3D shape deformation component extraction. With the increasing use of 3D models, the need for analyzing their intrinsic characteristics becomes mainstream. Early work [19] extracts principal components from the mesh data set by Principal Component Analysis (PCA), but the extracted components contain global deformations, which is not effective for users to make local edits. Some works [20] demonstrate that sparse constraints are effective for achieving localized deformation results. However, the classical sparse PCA [21] does not take the spatial information into consideration. By promoting sparsity in the spatial domain, many works extract localized deformation components with a sparsity constraint [1], [22], which outperforms the standard PCA variants such as Clustering-PCA [23] with respect to choosing suitable compact basis modes, especially for producing more localized meaningful deformation. Moreover, the pioneering work [1] represents meshes with Euclidean coordinates, but this representation is sensitive to rigid and non-rigid transformations. Later work [2] extends the previous work [1] to better deal with rotations by using deformation gradients to represent shapes, but the method still cannot cope with larger rotations greater than 180° due to their inherent ambiguity. Based on deformation gradients, Neumann et al. [24] learn the arm-muscle deformation using a small set of intuitive parameters. The work [3] extends [1] by using a rotation invariant representation based on edge lengths and dihedral angles [25], so can handle large-scale deformations. However, the representation [25] is not suitable for extrapolation as this would result in negative edge lengths. This limits the capability of [3] for deformation component analysis, as extrapolation is often needed e.g. when utilizing the extracted deformation components for data-driven shape editing. Recent work [4] proposes a convolutional autoencoder based on an effective shape representation [13] to learn the localized deformations of a shape set, but their architecture is not suitable for extracting multiscale deformation components. Overall, different from these works [1], [2], [3], [4], our method can produce meaningful and multiscale localized deformation components.

Deep learning on 3D Shapes. With the development of artificial intelligence, deep learning and neural networks have made great progress in many areas, in particular 2D image processing. Hence, some researchers transform the non-uniform geometry signals defined on meshes of different topologies into a regular domain, while preserving shape information as much as possible, which enables powerful uniform Cartesian grid based CNN (Convolutional Neural Network) backbone architectures to be used on problems such as cross-domain shape retrieval [26], surface reconstruction [27] and shape completion [28]. DDSL [29] was recently proposed, which is a differentiable layer compatible with deep neural networks for learning geometric signals. However, due to the irregularity of 3D shapes, deep learning is difficult to apply straightaway. Inspired by image processing, some works [30], [31], [32] apply deep learning to the voxel representation with regular connectivity. However, voxel representation incurs significant computation and memory costs, which limits the resolution such methods can handle. Wang et al. [33], [34] improve the performance of voxel-based convolutions by proposing an adaptive octree structure to represent 3D shapes, and apply it to shape completion [35]. Meanwhile, recent works [36], [37] define the convolution on point clouds by using K-nearest-neighbors (KNN) and spherical convolutions. More recently, the work [38] proposed the EdgeConv operator for learning on the point cloud to improve the performance of segmentation and classification. In addition to the voxel and point cloud representations, shapes can also be represented as multiview projection images to perform 2D-CNNs [31], [39], [40] for 3D object recognition and classification. Such approaches are used to learn the local shape descriptors for shape segmentation and correspondence [41]. For applications that take meshes as input or generate meshes as output, turning meshes to alternative representations can lead to useful topology information to be lost.

Alternatively, as a mesh can be represented as a graph, CNNs can be extended to graph CNNs in the spatial domain [42], [43] or spectral domain [44], [45], [46] for mesh-based deep learning. In the spatial domain, works [47], [48], [49], [50] apply variational autoencoders on 3D meshes for various applications such as reconstruction, interpolation, completion and embedding. The work [51] uses autoencoders to analyze deformable solid dynamics. However, none of the existing methods can extract multiscale deformation components, which we address by using a novel attention-based mesh convolutional network architecture.

Attention mechanism on convolutional networks. Deep neural networks have proved their superiority in extracting high-level semantics and highly discriminative features on various image datasets. Researchers now pay more attention to using convolution features more effectively on fine-grained datasets to improve the performance. Such works can be widely seen in different areas of computer vision and natural language processing, such as image translation (DA-GAN) [52], person re-identification [53], [54], document classification [55], object detection [56], [57], video classification [58], etc. The work [59] proposes a “soft attention” mechanism which predicts soft weights and computes a weighted combination of the items in machine translation. In [60], a hierarchical co-attention method is proposed to learn the conditional representation of the image given a problem. Following [60], [61] extends the co-attention model to higher orders. Some works effectively utilize attention as a way to focus on specific regions for learning. Wang et al. [62] demonstrate the benefit of guiding the feature learning by using residual

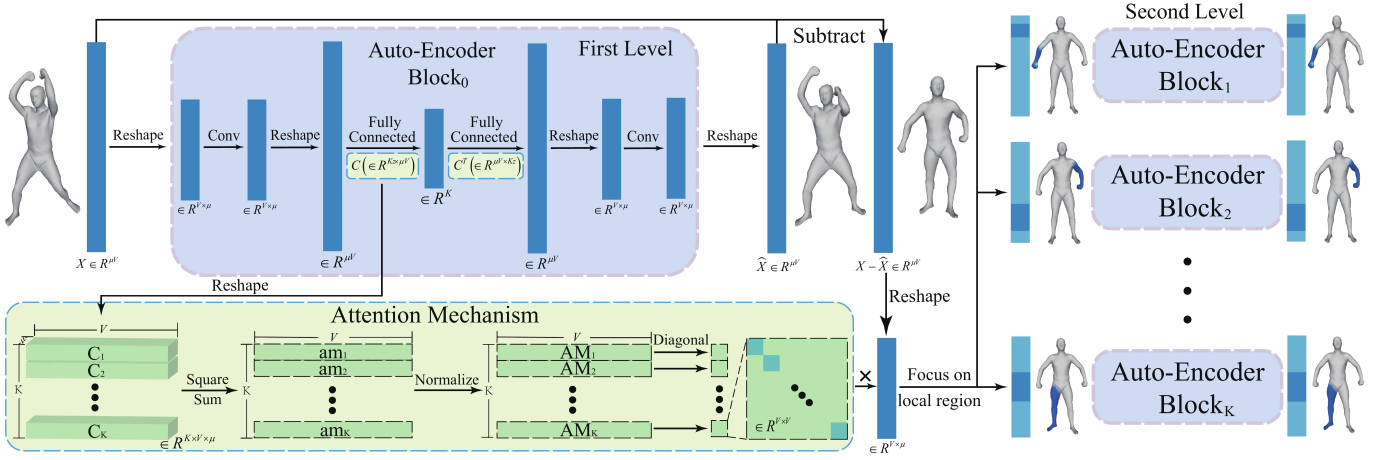


Figure 1. Our network architecture with attention mechanism and stacked autoencoders. We obtain large-scale deformation components and attention masks from the first-level autoencoder AE_0 . For the second-level autoencoders $AE_k, 1 \leq k \leq K$, we put the residual value $X - \hat{X}$ of the first-level autoencoder AE_0 with the attention masks into K autoencoders focusing on different sub-regions of the shape. We can then obtain small-scale deformation components by these autoencoders. This architecture can be further extended to include more scale levels. The autoencoder has mirrored encoder and decoder structure. The encoder has one convolution layer and a fully connected layer, and the encoder and decoder share the same trainable parameters. K, V and μ are the dimension of the latent space (the number of attention masks), the number of vertices and the dimension of the vertex features, respectively. By making the latent vector as a one-hot vector, we can extract the K attention masks $AM_k, 1 \leq k \leq K$ from the parameter C as the top-left corner of the figure shows. Please refer to Sec. 4.2 for details.

attention learning for improving the recognition performance. Another example is the attention-focused CNN (RA-CNN) [12] based on the Attention Proposal Network (APN), which actively identifies the effective region and uses bi-linear interpolation to adjust the scale, and then the enlarged region of interest is used for improved fine-grained classification. However, the above works apply the attention mechanism in the 2D domain. Our work extends the attention mechanism to the 3D domain to extract multiscale deformation features based on an effective shape representation [13].

3 DEFORMATION REPRESENTATION AND CONVOLUTION OPERATOR

The input of our overall network is based on the recently proposed as-consistent-as-possible (ACAP) deformation representation [13], which can cope with large-scale deformations of shapes and is defined only on vertices, making mesh-based convolutions easier to implement. To validate it is a good choice, we compare this with a recently proposed general-purpose mesh autoencoder (AE) architecture DEMEA [49]. For fair comparison, we compare with DEMEA [49] on the COMA [48] dataset with the same setting as used in [49]. We also use the same training set and test set split: the training set contains 10 various expressions and 17794 meshes in total. The test set contains 2 challenging expressions (high smile and mouth extreme), with 2671 meshes in total. We set the same latent space dimension (32) and our network only uses a single level of AE to make the network architecture comparable. The results on the test set in terms of average per-vertex errors show that our base AE network with 0.822mm error outperforms DEMEA with 1.05mm error. Such benefits can be more substantial when shapes undergo more substantial deformations, such as various human body poses, so we build our network architecture based on this.

For a given shape set with N shapes that share the same connectivity each with V vertices, without loss of

generality, we choose the first shape as the reference shape. For the patch which consists of the i^{th} vertex and its 1-ring neighbor vertices, we can calculate its deformation gradient $\mathbf{T}_{m,i} \in \mathbb{R}^{3 \times 3}$. The deformation gradient of the patch is defined on the i^{th} vertex of the m^{th} shape, which describes the local deformation w.r.t. the reference shape. $\mathbf{T}_{m,i}$ of shape m is obtained by minimizing the following formula: $\arg \min_{\mathbf{T}_{m,i}} \sum_{j \in \mathcal{N}_i} c_{ij} \|(\mathbf{p}_{m,i} - \mathbf{p}_{m,j}) - \mathbf{T}_{m,i}(\mathbf{p}_{1,i} - \mathbf{p}_{1,j})\|_2^2$ where \mathcal{N}_i is the 1-ring neighbor vertices of vertex i and $c_{ij} = \alpha_{ij} + \beta_{ij}$ is the cotangent weight [63], where α_{ij} and β_{ij} are the angles in the two faces that share a common edge (i, j) . Then $\mathbf{T}_{m,i}$ can be decomposed as $\mathbf{T}_{m,i} = \mathbf{R}_{m,i} \mathbf{S}_{m,i}$ using the polar decomposition, where $\mathbf{S}_{m,i} \in \mathbb{R}^{3 \times 3}$ is a symmetry matrix that describes the scaling/shear deformation and $\mathbf{R}_{m,i} \in \mathbb{R}^{3 \times 3}$ is an orthogonal matrix that describes the rotation. For the rotation matrix $\mathbf{R}_{m,i}$, it can be represented by the rotation axis $\omega_{m,i}$ and rotation angle $\omega_{m,j}$. But the mapping from the rotation matrix to rotation axis and angle is one-to-many. For shapes with large-scale rotations, the rotation axis and rotation angle of adjacency vertices may become inconsistent, which results in artifacts when synthesizing new shapes, as shown in [13]. Gao et al. propose a two-step integer optimization to solve the problem which makes the rotation angle and rotation axis of adjacent vertices as consistent as possible. For the details, please refer to [13].

Next, for each vertex i on the m^{th} shape, we can obtain the feature $q_{m,i} = \{r_{m,i}, s_{m,i}\} \in \mathbb{R}^9$ by extracting non-trivial elements $r_{m,i} \in \mathbb{R}^3$ and $s_{m,i} \in \mathbb{R}^6$ from the logarithm of rotation matrix $\mathbf{R}_{m,i}$ and scaling/shear matrix $\mathbf{S}_{m,i}$ respectively. Finally, the ACAP feature of the m^{th} shape can be represented by $\{q_{m,i} | 1 \leq i \leq V\}$. Due to the usage of the \tanh activation function [4], we further linearly scale each element in $r_{m,i}$ and $s_{m,i}$ to $[-0.95, 0.95]$ separately. Then we concatenate $q_{m,i}, 1 \leq i \leq V$ together in the vertex order to form a long vector $X_m \in \mathbb{R}^{\mu V \times 1}$ as the feature of the m^{th} shape, where $\mu = 9$ is the dimension of the ACAP feature of each vertex.

We further introduce the graph convolutional operator used in our architecture. As illustrated in [4], the output of the convolution operator for every vertex is a linear combination of the inputs of the vertex and its 1-ring neighbor vertices, along with a bias. The output y_i for the i^{th} vertex is defined as follows:

$$y_i = \mathbf{W}_{point}\mathbf{x}_i + \mathbf{W}_{neighbor} \frac{1}{D_i} \sum_{j=1}^{D_i} \mathbf{x}_{n_{ij}} + \mathbf{b} \quad (1)$$

where \mathbf{x}_i is the input feature vector of the i^{th} vertex, D_i is the degree of the i^{th} vertex, and n_{ij} ($1 \leq j \leq D_i$) is the j^{th} neighbor vertex of the i^{th} vertex. \mathbf{W}_{point} , $\mathbf{W}_{neighbor} \in \mathbb{R}^{\mu \times \mu}$ and $\mathbf{b} \in \mathbb{R}^{\mu}$ are the trainable parameters of the graph convolutional layer. All these weights are shared by all the vertex and their neighborhoods in the same convolutional layer and learned during the training of the network.

4 METHODOLOGY

In this section, we introduce our network architecture from these three main aspects: our novel autoencoder structure, attention mechanism and redundant component removal. Firstly, we introduce the novel autoencoder structure. Then we describe our attention mechanism, which is applied to help autoencoders to extract multi-scale deformation components. Lastly, we explain how we remove redundant components, followed by implementation details of our network. This architecture is flexible to support multiple levels of scales. In most cases, two levels of deformation scales are sufficient to represent deformation in the dataset of this paper. Please refer to Sec. 5.3.3 for details.

4.1 Autoencoder Block

As Fig. 1 illustrates, we have achieved the multiscale structure by stacking autoencoder blocks. In the first (coarsest) level, we have one autoencoder AE_0 , and in the second level, K autoencoders AE_k ($k = 1, 2, \dots, K$) are built, each focusing on one local region through an attention mechanism, which will be detailed later. The number of second level AEs is determined by the dimension of the latent space of AE_0 .

For each shape m , $1 \leq m \leq N$, we represent it by the pre-processed ACAP feature $X_m \in \mathbb{R}^{V \times \mu}$, as described in Sec. 3. We use an encoder to map the feature to a 128-dimensional latent code and a decoder which reconstructs the shape ACAP feature from a latent code z . Both of the encoder and decoder have one mesh-based graph convolutional layer and a fully-connected layer and their network structure is symmetrical, where the learnable parameters of the fully connected layer are defined as $\mathbf{C} \in \mathbb{R}^{K_z \times \mu V}$, K_z is the dimension of the latent space. Especially, the fully connected layers of encoder and decoder share same learnable parameter \mathbf{C} without bias. The latent vector z for all the N shapes form a matrix $\mathbf{Z} \in \mathbb{R}^{N \times K_z}$. Similar to [4], all the layers use the \tanh activation function. Figure 1 illustrates the autoencoder architecture in the left top corner.

The output $\hat{X} \in \mathbb{R}^{\mu V}$ of the whole autoencoder block can be scaled back to the ACAP deformation representation and reconstruct the Euclidean coordinates using [13]. For every autoencoder, we optimize the following loss function that includes three terms: reconstruction loss that ensures accurate reconstruction of the input, sparsity loss $\Omega(\mathbf{C})$ that promotes localized deformation components, and non-trivial regularization term $\mathcal{V}(\mathbf{Z})$ to avoid

creating trivial solutions. The total loss for an autoencoder block AE_k is as follows:

$$\mathcal{L}_{AE_k} = \lambda_1 L_{recon} + \lambda_2 \Omega(\mathbf{C}) + \mathcal{V}(\mathbf{Z}) \quad (2)$$

where AE_k , $0 \leq k \leq K$ represent the k^{th} autoencoder, λ_1, λ_2 are the balancing weights.

The reconstruction loss is the MSE (mean square error) loss, defined as $L_{recon} = \frac{1}{N} \sum_{i=1}^N \|X_i - \hat{X}_i\|_2^2$. For the non-trivial regularization term, $\mathcal{V}(\mathbf{Z}) = \frac{1}{K_z} \sum_{j=1}^{K_z} \max((\max_m |Z_{jm}| - \theta), 0)$, where Z_{jm} is the weight for the j^{th} dimension of the m^{th} shape, and θ is a positive number and we set $\theta = 5$ in our experiments.

The above two loss terms are the same as the previous work [4]. But for the loss $\Omega(\mathbf{C})$, there are some difference: we choose the step function Λ defined above to map geodesic distances to $\{0, 1\}$, rather than the previously used clipped linear interpolation function. This is because our network architecture extracts hierarchical deformation components, so at any level, a fixed component size (rather than a range) is preferred. Autoencoder blocks at different levels will produce localized deformation components of different scales by adjusting the tunable parameter d . Our sparsity loss term $\Omega(\mathbf{C})$ is defined as: $\Omega(\mathbf{C}) = \frac{1}{K_z} \sum_{k=1}^{K_z} \sum_{i=1}^V \Lambda_{ik} \|C_{k,i}\|_2$, where $C_{k,i}$ is the μ -dimensional vector of component k of vertex i , $\|\cdot\|_2$ is group sparsity ($\ell_{2,1}$ norm), and Λ_{ik} is sparsity regularization parameters defined as follows:

$$\Lambda_{ik} = \begin{cases} 0 & d_{ik} < d \\ 1 & d_{ik} \geq d \end{cases} \quad (3)$$

where the Λ_{ik} is a binary function, where d_{ik} denotes the normalized geodesic distance [64] from vertex i to the center point c_k of component k , which is defined as $c_k = \operatorname{argmax}_i \|C_{k,i}\|_2$, and is updated in each iteration of network training. d is a tunable parameter, which controls the size of the deformation region of a component. Larger d corresponds to bigger deformed regions of the shape. For our task, AEs of different levels choose different values of d . Please refer to Sec. 4.4 for the default value of d .

We train the whole network end-to-end by adding all the losses of autoencoder blocks together as $L_{total} = \sum_{k=0}^K \mathcal{L}_{AE_k}$, which includes a first-level AE (AE_0) and K second-level AEs (AE_k , $1 \leq k \leq K$).

4.2 Attention Mechanism for Multiscale Analysis

Similar to 2D images, there are many tasks (image recognition) that focus on the different level of images by attention mechanism, such as [12]. So we want to explore the 3D shape with multiscale fashion by the attention mechanism. We design an attention mechanism that facilitates our autoencoder blocks extracting the multiscale deformations components.

Mostly, the deformation datasets (e.g. human, horse and fabric) have both global scale and local scale deformations. Hence, it is naturally to extract these deformations in multiscale manner. To make the second level AEs focus on sub-regions to extract finer level components and then form a multiscale structure, we extract learnable attention masks from the fully-connected layer of the first-level AE AE_0 . Our attention mechanism is shown in the bottom left corner of Fig. 1. Due to the sparsity constraint $\Omega(\mathbf{C})$, the parameter \mathbf{C} of the fully connected layer represents the sparse deformation components. The $C_k \in \mathbb{R}^{V \times \mu}$, $1 \leq k \leq K$ that is denoted as each vector in \mathbf{C} represents a deformed sub-region

of the shape. The deformed sub-regions can be regarded as the interested mask of the second level AEs. So in every iteration of training, we can extract each row of \mathbf{C} by setting the latent vector to a one-hot vector for second level AEs:

$$C_k = \mathbf{C}^T \times OH_k \quad (4)$$

where the OH_k is a K -dim column vector and k^{th} entry of OH_k is 1 and the rest are 0. Then we reshape C_k to a 2D array with the size $V \times \mu$, denoted as C_k^r . The unnormalized attention mask $am_{k,i} \in \mathbb{R}^{K \times V}$ for the k^{th} component of the i^{th} vertex is defined as

$$am_{k,i} = \sum_{j=1}^{\mu} C_{k,i,j}^r \cdot 2. \quad (5)$$

where the $C_{k,i,j}^r$ is the (i, j) entry of C_k^r . We further normalize it to obtain the normalized attention mask $AM \in \mathbb{R}^{K \times V}$, where

$$AM_{k,i} = \frac{am_{k,i}}{\sum_{k=1}^K am_{k,i}}. \quad (6)$$

So for the first-level autoencoder AE_0 , the residual value of the reconstruction is $X - \hat{X}$ and the normalized attention mask is AM . We reshape $(X - \hat{X})$ to a 2D array $X_{res} \in \mathbb{R}^{V \times \mu}$. For the second-level autoencoder AE_k , $1 \leq k \leq K$ as in Fig. 1, its input is $diag(AM_k) \times X_{res}$, where $diag(\cdot)$ returns a square diagonal matrix with the elements of the vector on the main diagonal. The input of each second-level AE is therefore the weighted (by the corresponding attention mask) residual of the first-level AE. Therefore, this attention mechanism ensures that the second-level AEs can reconstruct smaller scale deformations that cannot be well captured by AE_0 , and each AE_k focuses on an individual local region. The sum of every column of AM is one according to Eq. 6, which ensures the sum of inputs to AE_k , $1 \leq k \leq K$ equals the residual value X_{res} of the first-level autoencoder AE_0 .

Under the supervision of loss function and attention mechanism, the first level autoencoder AE_0 is capable of capturing the large scale deformation and the second level AEs can capture the smaller scale deformations on specified regions of the shape. Consequently, our network can learn the multiscale deformation components of the whole shape set, the multiscale deformation components can be extracted from the parameters of fully connected layers of all AEs.

4.3 Redundant Component Removal

For all autoencoder blocks, we extract the fix number of deformation components for each AE. For fair comparison and capture all deformations, we set $K_z = 10$ for AE_0 , $K_z = 5$ for AE_k , $1 \leq k \leq K$. Because the multiscale analysis and our setting is to extract as much as possible to avoid missing any components, so there are some redundant components in second level AEs, which means the components does not contain any deformation, or only slight deformations compared to the reference mesh. So we remove these components if the contained information is lower than the given threshold value. All results in our paper are processed by the redundant component removal. The Fig. 6 shows the results of the output of network without the process (there are some components with slight deformations compared to the reference mesh). Fig. 11 illustrates some finite components are sufficient enough to represent the whole dataset during the deformation components analysis. Thus, we will remove the

redundant components which contain slight deformations defined in Eq. 7.

The process aims to make our results more brief and reasonable and is done after network training. It can gain trade off between the multiscale decomposition and avoiding overfitting on training data. As a result, we could exclude these redundant deformation components. We define the following deformation strength of deformation components to filter slight or noisy deformation components. The strength I on the features X_m is defined as:

$$I(X_m) = \frac{\sum_{i=1}^V 1(\|(X_{diff})_i\| > \epsilon_1) \|(X_{diff})_i\|}{\sum_{i=1}^V 1(\|(X_{diff})_i\| > \epsilon_1)} \quad (7)$$

where $X_{diff} = X_m - X_r$, and $X_r, X_m \in \mathbb{R}^{V \times \mu}$ are the features of the reference mesh and the extracted deformation components from the autoencoders respectively. $\|\cdot\|$ is the ℓ_2 norm of the vector, and $1(\cdot)$ gives 1 if the condition is true, and 0 otherwise. $\epsilon_1 = 1e - 6$.

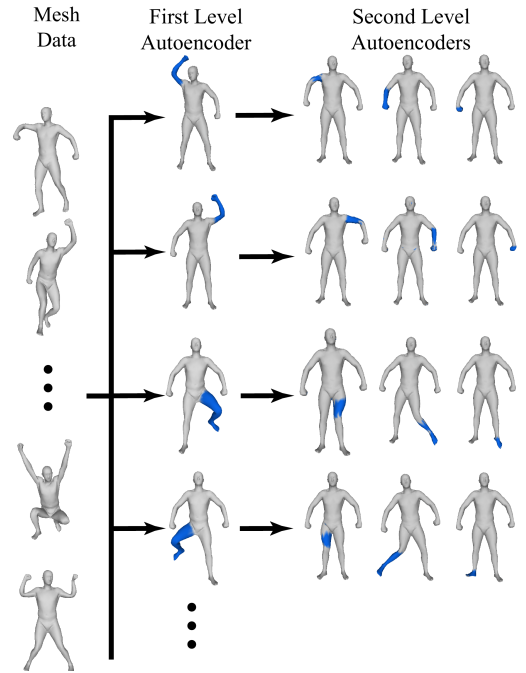


Figure 2. The multiscale structure of deformation components on the shape set SCAPE [65]. In the figure, we filter the redundant components. As a result, our method can learn deformation components of different scales. The first column shows some examples from the SCAPE dataset, the second column presents coarse level deformation components from the first-level AE_0 , and the third column gives the fine level deformation components from the second-level AE_k , $1 \leq k \leq K$.

If the extracted deformation component corresponds to a slight deformation, defined as its strength being smaller than a threshold ϵ_2 , we will remove the component. In our experiments, we set $\epsilon_2 = 1e - 2$. Finally, we obtain the multiscale structure of the deformation components, as shown in Figs. 2, 9, 10, 11, 12 and 13, where deformations at different scales are indicated with arrows.

We also check if our network produces similar components (near duplicates). To achieve this, we test the similarity of the extracted deformation components. Fig. 3 visualizes the cosine similarity matrix of the components extracted from the first-level AE. It shows that the components have low similarity, because our AE applies the localization constraint and reconstruction

error minimization to ensure different components in the latent space represent different parts on the shape; having duplicated components would lead to reduced representation capability, so higher total loss.

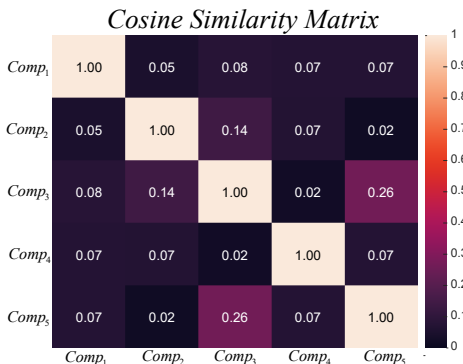


Figure 3. Visualization of the cosine similarity matrix of the components extracted from the first-level AE. It shows that the components have low similarity. The value (0-1) in each grid indicates the similarity between two components. Larger values mean more similar.

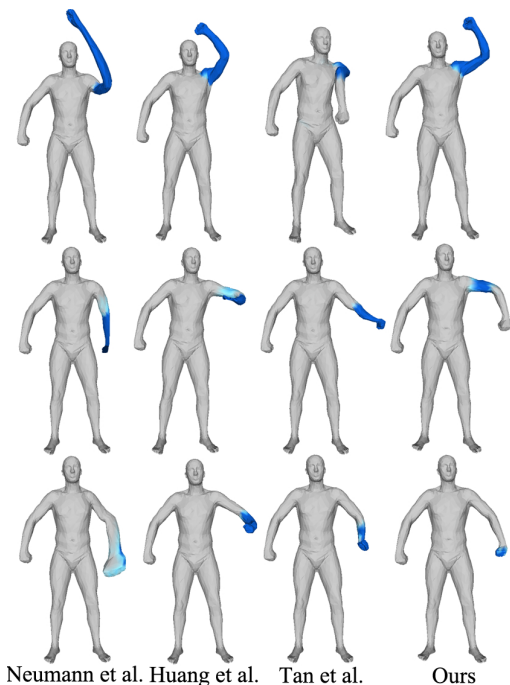


Figure 4. Comparison of deformation components located in the left arm of SCAPE [65], which are extracted by different methods. The deformed region is highlighted in blue. Every row shows the components located in a similar region. It shows that our results are more reasonable.

4.4 Implementation Details

Our experiments were carried out on a computer with an i7-6850K CPU, 16GB RAM and an Nvidia GTX 1080Ti GPU.

Datasets: We compare with the state-of-the-art methods on the SCAPE dataset [65], Horse dataset [66], Face dataset [67], Humanoid dataset [3], Dress dataset, Pants dataset [68], Flag dataset, Skirt Dataset, Fat person (ID:50002) from the Dyna [69] Dataset, Coma [48] Dataset, Swing and Jumping datasets [70]. The Dress,

Flag and Skirt datasets were synthesized by the NVIDIA Clothing Tools on 3ds MAX. Our used data range from rigged deformation like human motion to non-rigged deformation (Face, cloth), from small dataset (several hundreds shapes) to large dataset (thousands of shapes(Dyna)). For the above datasets, they are mainly sequence models which have similar deformation between two neighbour shapes. So for our model’s generaliability, we select one from ten models to training and the rest to testing, the ratio is 9:1 for test and training. For un-sequence data like SCAPE, we split the training set and test set randomly with a ratio 1:1. A special case for Coma dataset, we use same setting as the DEMEA [49] for fair comparison. The statistics of the dataset are shown in Table 1. The table lists the number of shapes that dataset contains, training examples and testing examples. All data is very easy to obtain, some are the public data, some are the synthesized data by the professional software. We will release the synthesized data for future research.

Table 1

Data Statistics. We summarize the data statistics of 10 datasets in our experiments. Each data is split into training set and test set with a ratio of 1:9. For the Coma [48] dataset, we use the same setting with DEMEA [49] for fair comparison.

DataSet	# All Shapes	# Training Shapes	#Testing Shapes
Swing	280	28	252
Scape	71	36	35
Pants	241	25	216
Humanoid	154	16	138
Horse	49	5	44
Flag	500	50	450
Dress	500	50	450
Jump	150	15	135
Fat	4737	469	4274
Face	385	39	346
Skirt	231	23	208
Coma	20465	17794	2671

In our experiments, our network takes the ACAP features of 3D shapes as input, which can describe the large scale deformations and be calculated by the novel work [13]. We have two levels of autoencoders, the first-level only has one AE (AE_0), and the second-level has the same number of AEs ($AE_k, 1 \leq k \leq K$) as the attention masks. Since in most wild data sets, the deformation of the shape is not very exaggerated, two levels of autoencoders are enough to extract the multiscale localized deformation components in our experiments, but this can be extended if necessary. For above different categories, we use fixed hyperparameters. We perform experiments on the SCAPE dataset to demonstrate how we choose the suitable hyper-parameter in the Sec. 5.3. As shown in that section, our stacked AEs has the lowest error with the following default parameters: $\lambda_1 = 10.0, \lambda_2 = 1.0$, where λ_1, λ_2 are the weights of reconstruction error and sparsity constrain term, respectively. $d = [d_1, d_2] = [0.4, 0.2], K_z = [10, 5]$, corresponding to the coarse and fine levels. Here, we train the whole network end-to-end rather than in a separate manner. We set the learning rate as 0.001 with the exponential decay by the ADAM solver [71] to train the network end-to-end until it converges, which takes approximately 10,000 epochs. For all AEs, we set the batch size as 256, which is randomly sampled from the training data set. For a typical dataset, the training of stacked AEs takes about 10 hours. Once our network is trained, the extracted components will output efficiently: outputting one component takes only about 50 milliseconds.

Table 2

Errors of applying our method to generate unseen data from Horse [72], Face [67], Humanoid [3], Pants [68], Flag dataset, Fat person (ID:50002 from the Dyna dataset [69]), Swing and Jumping [70] datasets. From the table, the generation ability of our network is better than the other methods on the E_{rms} error and $STED$ error.

Dataset	Metric	Method					
		Ours	Tan et al.	Wang et al.	Huang et al.	Neumann et al.	Bernard et al.
Horse	E_{rms}	6.9246	12.9605	29.6090	18.0624	7.3682	20.1994
	$STED$	0.0336	0.04004	0.04332	0.05273	0.08074	0.4111
Face	E_{rms}	1.4409	2.9083	8.5620	12.3221	2.9106	2.9853
	$STED$	0.0071	0.007344	0.01320	0.01827	0.008611	0.02662
Jumping	E_{rms}	16.3475	24.4827	44.3362	37.9915	29.3368	49.9374
	$STED$	0.0321	0.04862	0.05400	0.06305	0.1268	0.4308
Humanoid	E_{rms}	3.2127	3.4912	60.9925	16.1995	14.3610	6.6320
	$STED$	0.0226	0.01313	0.03757	0.02247	0.07319	0.04612
Swing	E_{rms}	12.2615	14.0836	29.5329	24.495	15.1942	22.6571
	$STED$	0.0311	0.03789	0.04224	0.04343	0.0830	0.1139
Pants	E_{rms}	6.4083	7.8986	39.2946	10.1880	28.4118	23.6785
	$STED$	0.0372	0.0414	0.0540	0.04958	0.1762	0.06484
Dress	E_{rms}	11.5117	34.0579	35.2816	35.2340	55.5806	12.2239
	$STED$	0.0397	0.0415	0.0635	0.0782	0.0784	0.2167
Fat	E_{rms}	4.3456	4.5609	25.9187	5.3215	7.4522	5.3348
	$STED$	0.0052	0.0053	0.0055	0.0035	0.0372	0.0289
Flag	E_{rms}	20.0627	26.3174	62.2925	51.2551	23.1364	23.1535
	$STED$	0.0157	0.0176	0.0354	0.2183	0.0914	0.0169

As we will later discuss, compared with separate training in Table 4, training the network end-to-end can result in smaller reconstruction errors E_{rms} on all data sets. The main reason is that the network can adjust the attention mask by minimizing the loss function, and conversely, the adjusted attention mask will result in smaller reconstruction errors E_{rms} . Under the collaboration, we get better results as shown in Table 4.

5 EXPERIMENTAL RESULTS & EVALUATION

In this section, we will evaluate our method on the above datasets from the following aspects: Quantitative Evaluation, Qualitative Evaluation and Applications.

5.1 Quantitative Evaluation

We compare the generation ability of our method with the state-of-the-art methods [1], [2], [3], [4], [22] on various datasets. In this experiment, we select one model from every ten models for training and the remaining for testing. After training, we align all the models and scale them into a unit ball. Then we use E_{rms} (root mean square) error [73] and $STED$ error [74] to compare the generalization error on the test data (i.e., the reconstruction error for unseen data) with the various methods. In particular, $STED$ error is designed for motion sequences with a focus on ‘perceptual’ error of models. To ensure fairness, we train each

autoencoder to extract 50 components. As Table 2 shows, the performance of our method is better than the existing methods on the E_{rms} and $STED$. Because the Euclidean coordinate representation is sensitive to rotation, the extracted deformation components of the methods [1], [22] have more artifacts and implausible deformation, leading to larger reconstruction errors. Due to the limitation of the edge lengths and dihedral angle representation, the reconstruction using the method [3] can also be inaccurate and unstable. The method [2] is not capable of encoding large-scale deformation (e.g. folds on the fabric), so it cannot recover the original deformation accurately in such cases. The method [4] uses a large-scale deformation representation to achieve good performance, but it cannot produce multiscale deformation components. In comparison, our method can keep lower reconstruction errors by using stacked AEs and analyze the residual value of the first-level AE, extracting effective multiscale deformation components.

Meanwhile, our extracted components are served for data-driven deformation. In the real world, the user usually edits the shape by a limited number of control points. To demonstrate the ability of each method to reconstruct deformed models by the limited control points, we use the furthest point sampling [73] to sample the control points to ensure that the sampled points distribute on the shape evenly. Under the constraint of the control points, we use the same number of the extracted components to




Figure 5. The reconstruction error using sparse control points to deform the SCAPE data set [65]. The control points are obtained by furthest point sampling. The generalization error is measured by the data-driven deformation with the extracted deformation components of various methods. In this figure, our performance is better than the other methods with lower errors.

perform data-driven deformation on the SCAPE dataset. As shown in Fig. 5, the reconstruction by our extracted components always keep a lower error. In comparison, due to the use of Euclidean coordinate representation, the methods [1], [22] fail to reconstruct shapes accurately. Our multiscale deformation components better characterize the deformation of the shape, leading to reduced errors.

5.2 Qualitative Evaluation

We also provide qualitative evaluation of extracted multiscale deformation components, by comparing the visualization results of our method with the other methods. In our experiments, we have two level autoencoders. For the each level autoencoders, we extract corresponding scale deformation components from the fully connected layer using the same method as [4]. For every autoencoders, we can extract same number components as the dimension of latent space. So we visualize each extracted components in each level. Due to our setting ($K_z = [10, 5]$), the first level and second level autoencoders output 10 and 5 deformation components respectively. For all results, there are some slight deformation components that is not shown for the brief and reasonable visualization. We remove these components by the post-process (Sec.4.3). Furthermore, we colored deformed region compared with the reference mesh.

As independently extracted deformation components do not correspond to each other, we adopt the visualization method in [2], [3], [4], and manually select two components with similar deformation areas as much as possible. Figs. 4 and 7 show the comparison with various methods on the SCAPE and Horse datasets. We show the corresponding deformation components in a similar area on the shape (left arm of the SCAPE and every major part of the horse). Our results are more meaningful, capturing major deformation modes in a multiscale manner.

We then verify if our extracted components are more semantically meaningful than other methods by a user study. We test it on three datasets (SCAPE, Horse and Pants) and compare our method with five state-of-the-art methods [1], [2], [3], [4], [22].

Table 3

The user study that verifies the semantic meaning of deformation components extracted by various methods. We ask 10 participants and report their average scores of each method on the three datasets.

DataSet	Neumann et al.	Bernard et al.	Wang et al.	Huang et al.	Tan et al.	Our
Horse	46.69	30.37	48.37	56.28	52.41	68.00
SCAPE	33.87	40.87	55.29	48.13	61.33	72.54
Pants	36.37	55.75	23.00	50.50	46.00	73.37

We adopt a scoring method to determine whether a component has a certain semantics meaningfulness. The scoring is based on the participants' perception of semantic appropriateness of each component. For each data set, we first render the components extracted by all methods into the same rendering style and mix them together. We let the users browse all the pictures to have an idea of distribution first, and then let the users score in the range from 0 to 100 for each rendered image using a slider. 10 participants were involved in the user study. Then we count the average score of each method by every participant as shown in Table 3. Our method receives highest scores in all the datasets.

In summary, compared with existing methods, our method can extract plausible and reasonable localized deformation components with semantic meanings, while the other methods have some distortions and cannot extract multiscale deformation components. Note that our extracted components correspond to *deformations*, rather than semantic *parts*, so it is natural that they are not always aligned with semantic segmentation, but instead aligned at the motion sequence level, as observed also in previous works [1], [2], [3], [4], [22].

In addition, we show more visualization results of multiscale deformation components on the following datasets: Swing [70], a fat person (ID: 50002) from Dyna [69], Flag, Dress and Skirt. The Flag, Dress and Skirt dataset are synthesized by physical simulation. The Skirt dataset contains more complex motion and deformations. In Figs. 10, 12, 9, 8, 13, the components extracted by our method are in a multiscale manner and consistent with semantic meanings and our method can learn deformation components of different scales with multiscale structure in a complex data.

We further compare shape editing using various methods by given control points and deformation components they extract. An example is shown in Fig. 14. We use the 8 control points (rendered as green balls) on the 29th shape in the SCAPE dataset [65] manually chosen by the user. Then we apply the data-driven method [13] to reconstruct it with the help of the extracted deformation components by various methods. As the left part of Fig. 14 shows, our result is similar to the ground truth and plausible, while the other results have some artifacts and distortions, such as the right arm of [3] and left arm of [4]. The right part of Fig. 14 shows the three main activated components during data-driven deformation. Since the SCAPE contains much large-scale rotation, the method [4] only focuses on extracting large-scale deformation, but fails to capture important fine details, which results in the serious distortion in the arm due to lack of essential components.

5.3 Parameter Settings and Ablation Study




Figure 6. The multiscale structure of deformation components on the Pants dataset [68]. In the figure, we don't filter the redundant components, the symbol '...' represents there are the similar results which contains slight deformation. This result is comparable with Fig. 11 which is processed by the Sec4.3. As a result, our method can learn deformation components of different scales. In the deformation components, there will be some finite component to represent the whole dataset, many redundant components which contains some slight deformations will appear.

Table 4

Comparison of different training strategies and the influence of the attention mechanism for shape reconstruction. We show results comparing joint training with separate training, and whether or not the attention mechanism is used. For each setting, we test it on several data sets by computing the E_{rms} of the reconstructed shapes for unseen shapes. It shows that joint training with attention mechanism give the best results.

DataSet	Swing	Scape	Pants	Humanoid	Horse	Flag	Dress	Jump	Fat	Face
With Attention Joint Training	12.2615	23.6807	6.4083	4.1538	6.9246	20.0627	11.5744	16.3475	4.3465	1.4409
Without Attention Joint Training	17.1974	46.2339	8.2945	6.0285	20.8894	25.2675	12.5284	23.2829	5.3045	2.500
With Attention Separate Training	20.2918	36.1484	15.0249	6.0469	11.5899	24.4949	14.0248	19.1332	4.4469	2.1345
Without Attention Separate Training	26.7981	43.4439	19.8732	8.0583	17.4221	28.9938	17.3392	24.9437	7.3341	3.9472




Figure 7. Comparison of deformation components located in a similar area on the Horse [75], which are extracted by different methods. The deformed regions are highlighted in blue. The first row shows the results of other methods, and the second row gives the results of our method. Every column shows a component located in a similar region.

In this section, we evaluate the model sensitivity to the parameters, including the weights (λ_1, λ_2) in the loss function, the size of deformation region (d of Λ_{ik}), the effect of attention mechanism on generalization error and the difference between joint training and separate training.

5.3.1 The choice of λ_1, λ_2

We test the influence of parameters λ_1, λ_2 on the generalization ability of the network. We evaluate it by E_{rms} errors of reconstructing unseen shapes on the SCAPE dataset. By fixing λ_1 to the default value 10, we change λ_2 from 0.2 to 30 as the right

curve of Fig. 16 shows. The result shows our network is robust to different choice of λ_2 . With fixed λ_2 , we change λ_1 from 0.2 to 30 as the left curve of Fig. 16 shows. The result justifies that our network can get lower errors when $\lambda_1 = 10$, which is chosen as the default value in our experiments.

5.3.2 The choice of d in Λ_{ik}

The other parameter to choose is d in Λ_{ik} , which is the parameter that determines the scale of the extracted components. d_1, d_2 are cutoffs for normalized geodesic distance with range from 0 to 1, thus, the original size of specific dataset has little affect on them. These parameters only reflect the related size of the localized deformation components compared to the whole model.

In our network, by default we stack two levels of autoencoders, so we need choose two parameters d_1, d_2 for both the first-level and the second-level autoencoders (AE_0 and AE_1) respectively. In order to extract multiscale deformation components, we need ensure $d_1 > d_2$, then AE_1 will cover a more detailed region than AE_0 . As Fig. 17 shows, we test the network generalization ability (E_{rms}) on the SCAPE dataset with different combinations d_1, d_2 by changing values from 0.05 to 0.5 with step 0.05. The figure illustrates that smaller d_1 and d_2 result in larger reconstruction errors. Although lower error can be achieved when d_1 and d_2 are large enough (close to 0.5), we will extract more global deformation components (i.e., not localized) It is a




Figure 8. The multiscale structure of deformation components on the skirt cloth dataset, which is our simulated data by physics engine. The data contains more folds of cloth and complex motion. In the figure, we filter the redundant components. The first row is some samples in the skirt dataset. And the second and third rows are the visualization of deformation components extracted from the first level AE and second level AEs respectively. As a result, our method can learn deformation components of different scales with multiscale structure.

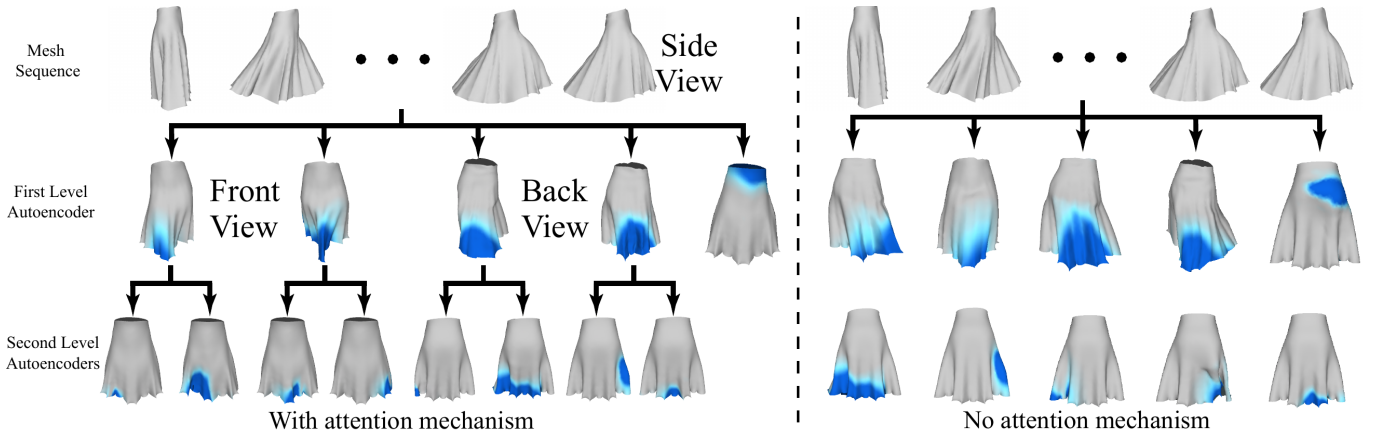


Figure 9. The multiscale structure of deformation components on the Dress data set (a lady walking forward in a skirt). In the figure, we have removed redundant components. In the left part, the result is with the attention mechanism. The first row shows some examples in the dress dataset from the side view. The second row presents the deformation caused by leg movement and the wind, which are extracted by the first-level AE (AE_0). The third row is the detailed deformations of the cloth during the movement, which are extracted by the second-level AEs ($AE_k, 1 \leq k \leq K$). In the second row, the two results on the left are the deformations of the front of the skirt, and the two results on the right are the deformation of the back of the skirt. In the right part, the results is without attention mechanism. It illustrate that the results no longer have a multiscale structure, like the left part shows. The second and third row represent the deformation components that are extracted by first level AE and second level AEs respectively, which is processed by removal of redundant components. The first row is the some sample data in Dress dataset.

trade off progress and we must balance the generalization ability and the extracted deformation components. In our experiments, we set $d = [0.4, 0.2]$ that can make network extract localized components and keep lower reconstruction error simultaneously.

5.3.3 The number of level of AE

The architecture of this neural net supports multiple levels of deformation scales from coarse to fine. In almost all tested dataset in this paper, the two level AEs architecture is enough to represent. The error between the input ACAP feature X_i and the reconstructed ACAP feature $\hat{X}_i, 1 < i < N$ is divided by the norm of X_i to get the percentage as shown in Eqn. 8.

$$Perc_i = \frac{\|X_i - \hat{X}_i\|_F^2}{\|X_i\|_F^2} \quad (8)$$

where the $\|\cdot\|_F$ is the Frobenius norm, $Perc_i$ is the percentage of error on the i^{th} shape in this dataset. Then, we choose

the maximum percentage of error in the whole dataset as the proportion of deformation that has not been represented. We show the results in the Table 5. The experiment illustrates that the percentage of two-level AEs is very low in all datasets so that our two-level AEs is enough.

5.3.4 The choice of K_z

The K_z describes the dimension of latent space of two level AEs. The hyper-parameter decides the number of deformation components in each autoencoder of each level. Since methods compared in the paper produce 50 deformation components, we let $K_{z_0} \times K_{z_1} = 50$. We have the following combinations: $K_z = [1, 50], [2, 25], [5, 10], [10, 5], [25, 2], [50, 1]$. However, the $[1, 50], [50, 1]$ are the trivial setting. So we test the other four settings of K_z on the SCAPE dataset. The Table 6 shows the results, which illustrate that it perform well when $K_z = [25, 2]$. But, there are only two deformation components for each second




Figure 10. The multiscale structure of deformation components on the fat person (ID: 50002) from the Dyna [69] dataset. The first row shows example shapes in the dataset, the second row presents coarse-level deformation components from the first-level AE_0 , and the third row shows the fine-level deformation components from the second-level AEs.

Table 5

The maximum percentage of the reconstruction error for each shape in every tested dataset in this paper. The values are calculated by the Eqn.8, as shown in table, the two-level AEs could represent the deformation of the very well with the tiny percentage of the reconstruct error, while the reconstruction error with one-level AE is much larger for reconstructing the input feature.

DataSet	Horse	Face	Jumping	Swing	Pants	Dress	Fat	Flag
Percentage of Error (one-level AE)	0.1219	0.06325	0.2312	0.4756	0.1763	0.5471	0.1322	0.1983
Percentage of Error (two-level AEs)	0.005491	0.001065	0.007800	0.06876	0.01745	0.06296	0.005652	0.02674

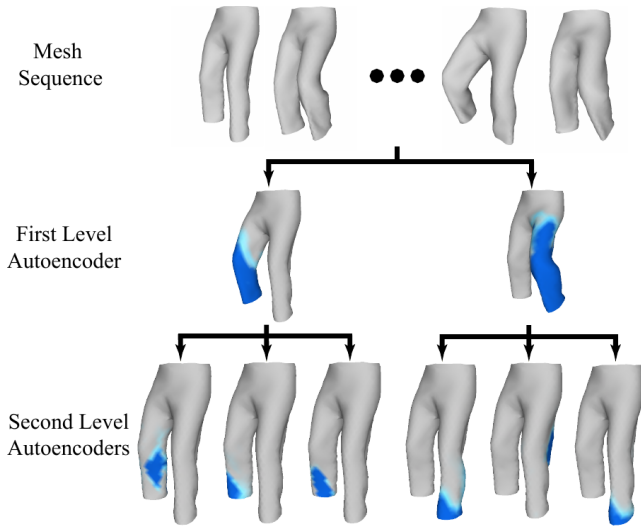


Figure 11. The multiscale structure of deformation components on the running pants data set [68]. In the figure, we have removed redundant components. Our method can learn deformation components at different scales. The first row shows example shapes of the running pants, the second row gives deformation components caused by leg movement, which are extracted by the first-level AE, and the third row presents the detailed deformations of the cloth from the second-level AEs.

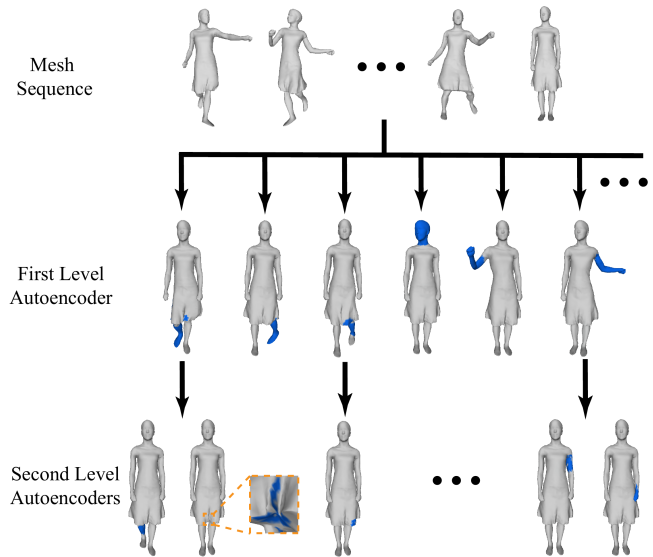


Figure 12. The multiscale structure of deformation components on the Swing [70] dataset. The first row shows some example shapes in the dataset, the second row presents coarse-level deformation components from the first-level AE_0 , and the third row shows the fine-level deformation components from the second-level AEs ($AE_k, 1 \leq k \leq K$).

level AEs, which is not reasonable for most datasets. For example, in Fig. 2, our network extract three meaningful components for some branches on SCAPE dataset. Furthermore, most of our experimental results show that there are more than 2 meaningful components extracted from the second level AEs. So we choose

the second higher performance with the setting $k_z = [10, 5]$.

5.3.5 Training strategies and attention mechanism

Finally, we evaluate the effect of the attention mechanism and different training strategies. The statistics are shown in Table 4




Figure 13. The multiscale structure of deformation components on the Flag dataset extracted by our method. The first row shows some example shapes of the dataset, the second row presents coarse-level deformation components from the first-level AE (AE_0), and the third row shows the fine-level deformation components from the second-level AEs ($AE_k, 1 \leq k \leq K$).

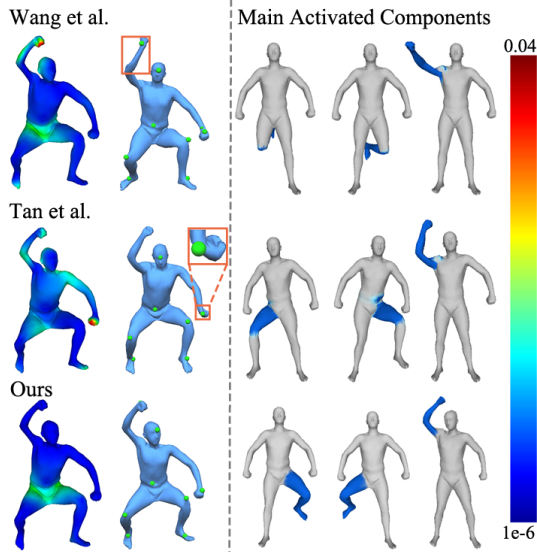


Figure 14. Comparison of shape reconstruction with different methods. The first column shows the error heat maps on the ground truth shape between the editing results and ground truth. The second column presents the editing results of different methods. The right side of the figure shows the three main activated deformation components during data-driven deformation. We use the same control points manually chosen by the user (8 vertices on the 29th shape in the SCAPE [65]) to reconstruct the shape with the data-driven deformation method [13] and the same number of components. The figure shows that our result is more plausible than the existing methods, and it is similar to the ground truth.

based on the experiments on the SCAPE dataset.

For the training strategy, we compare the reconstruction error (E_{rms}) by training the network either jointly or separately. The results are shown in the first and third rows of Table 4, and jointly training the network can get better results.

We also perform the experiment to demonstrate the effect of the attention mechanism qualitatively and quantitatively. The results are shown in the first and second rows of Table 4. Training with attention mechanism will get lower errors. The reason is that each autoencoder of the second-level will focus on a different sub-region to minimize the loss.

Table 6

The reconstruction errors on the unseen data from SCAPE with different settings of K_z . For each setting, it represents the dimension of latent space of two level AEs. From the table, it shows that our method has high performance on the setting $K_z = [25, 2]$. However, there are only up to 2 deformation components for each second level AEs with the setting, which is not reasonable for most of datasets. For example, in figure 2, our network extract three meaningful components for some brunches on Scape dataset. Furthermore, most of our experimental results show that there are more than 2 meaningful components extracted from the second level AEs. So we choose $K_z = [10, 5]$.

K_z	[2, 25]	[5, 10]	[10, 5]	[25, 2]
E_{rms}	34.9763	29.1127	23.6807	21.7749

For the qualitative evaluation, we show the results in Fig. 9. If we train our network without the attention mechanism, our network architecture degenerates into the version of Tan et al. [4] with two level autoencoders. In the second level AEs, there will be only one AE because all of them are same. Despite this, different level autoencoders are also able to extract the different scale deformation components as shown in the right part of Fig. 9. In the left part of Fig. 9, the second and third row represent the deformation components that are extracted by first level AE and second level AE respectively, which is processed by removal of redundant components. The first row is the some sample data in Dress dataset. The results illustrate that the deformation components no longer have a multiscale structure when the second level AEs do not focus on sub-regions to extract the deformations components.

5.4 Multiscale Shape Editing

Multiscale shape editing is an important application in computer graphics. Users usually start with editing of the overall shape, and then focus on adjusting the details. With existing methods, the extracted deformation components either contain some global information [1], [22] thus making the components unsuitable for local editing, or focusing too much on large-scale deformations and failing to capture essential small-scale deformations for faithful reconstruction, leading to distortions like [4], which would affect the users' editing efficiency for 3D animations. Given a shape deformation dataset that contains diverse deformations, our method can produce multiscale localized deformation components which are visually semantically meaningful, corresponding to typical deformation behaviour. Along with data-driven deformation [13], this allows users to edit shapes efficiently and intuitively under the constraints of the control points and subspace spanned by extracted deformation components. Please refer to the work [13] for details of data-driven deformation.

Fig. 15 shows some examples. For the SCAPE dataset, we design two actions: raising the left arm (Action 1) and then turning the wrist (Action 2). All the compared methods perform well in Action 1. However, in Action 2, only our method can naturally twist wrist with the help of our extracted multi-scale deformation components. In contrast, all other methods have various distortions. For the Horse dataset, we also design two actions: raising the whole tail (Action 1) and then twisting the end of the tail (Action 2). Our method can bend the end of the tail naturally after raising the whole tail. But other methods have more distortions and even lead to changes on the entire tail, especially the methods [1], [22] based on the Euclidean

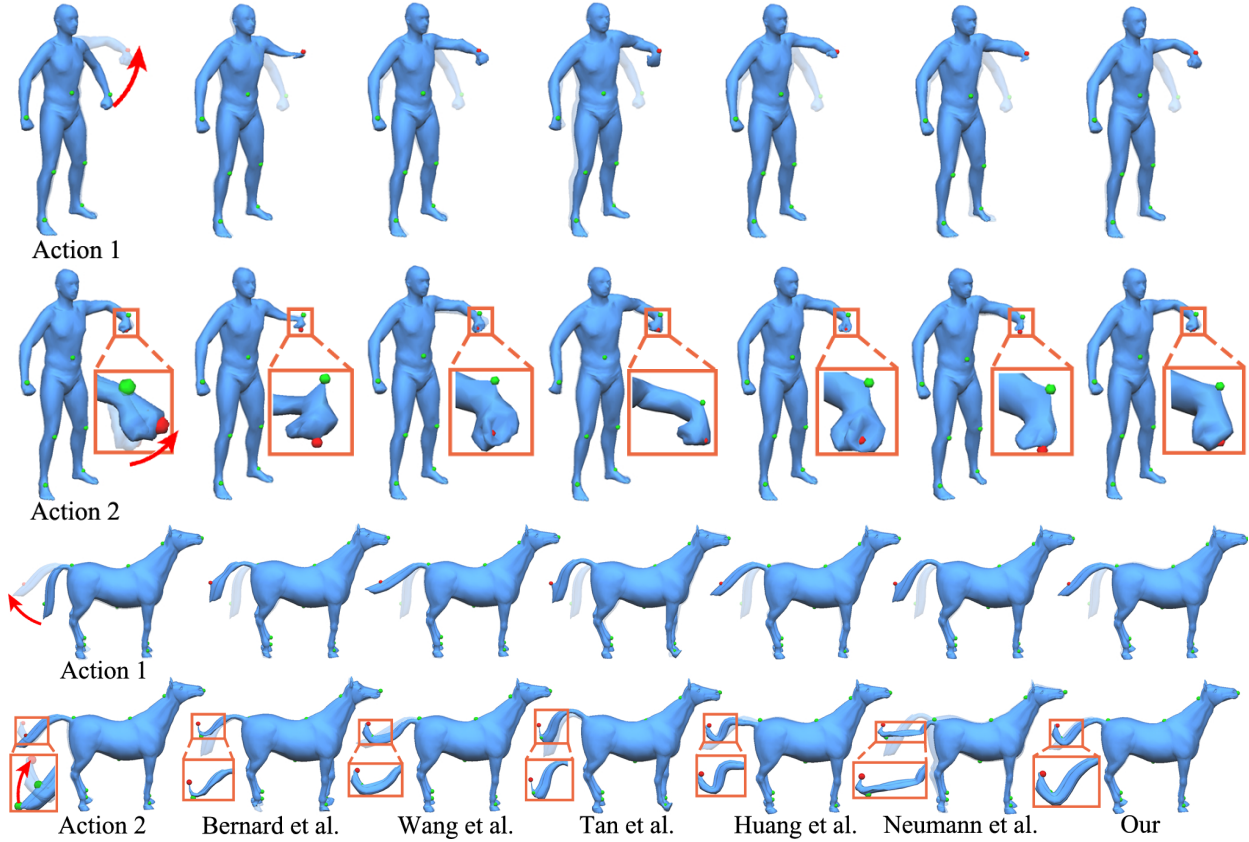


Figure 15. Comparison of multiscale shape editing. We compare the editing results with different methods on the SCAPE [65] and Horse [72] datasets. The results demonstrate that our extracted deformation components are suitable for multiscale shape editing. The first column shows the editing steps. Every row gives the deformed results of the corresponding action of various methods. The differences with other methods are highlighted in the orange rectangle with closeups to show the details. The existing methods have obvious distortions, demonstrating the superiority of our multiscale deformation components.




Figure 16. E_{rms} errors of generating unseen shapes with our overall autoencoder w.r.t. the weights λ_1 and λ_2 . The figure shows that our network can get lower errors when $\lambda_1 = 10$ and is robust to different choice of λ_2 .

coordinate representation. In summary, our extracted multiscale deformation components can perform better than existing methods in multiscale shape editing. See the accompanying video for more results.

6 LIMITATION AND CONCLUSION

In this paper, we propose a novel autoencoder with the attention mechanism to extract multiscale localized deformation components. We use stacked AEs to extract multiscale deformation components. This helps capture richer information for better shape

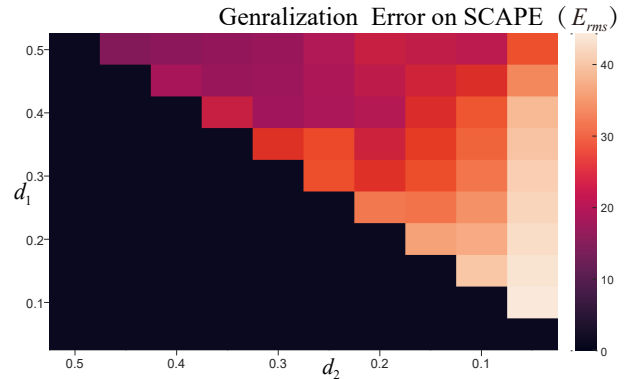


Figure 17. The relationship between E_{rms} of generating the unseen data and (d_1, d_2) is visualized, where d_1 and d_2 are from AE_0 and $AE_k, 1 \leq k \leq K$ on SCAPE data. The noise in the figure is a result of the gradient-descent training procedure. The black (below the main diagonal) means no data, as by definition d_1 should be larger than d_2 .

editing, with better generalization ability (see Table 2). Moreover, the first-level AE extracts some coarse level components and learns attention masks to help the second-level AEs focus on relevant sub-regions to extract fine-level components. Extensive quantitative and qualitative evaluations show that our method is effective, outperforming state-of-the-art methods. The extracted deformation components by our method can be used on multiscale

shape editing for computer animation, which demonstrates that the extracted multiscale localized deformation components are effective and meaningful for reducing the user's efforts. In the future, this work can give more solutions and explorations on applying the attention mechanism on 3D shape analysis and synthesis.

Although our method can analyze the shape dataset with multiscale manner to extract some different-scale components for easy shape editing, there are some limitations. While datasets including ShapeNet contains the general 3D shapes, our methods cannot handle those kinds of data, especially for CAD manufacturing. The data that our method can handle must be with the same connectivity. Also, our method uses a fixed network architecture and attention mechanism to analyze the shape with multiscale manner, so the multiscale structure of shape set will be similar and obey the fixed network structure for all dataset. We cannot capture the multiscale variation between different categories in the dataset. In the future, We are committed to exploit a method that can analyze and capture the variation of the multiscale structure on ShapeNet automatically. Finally, we can also merge the post-process (Sec 4.3) to the neural network which can predict the number of sub-parts automatically.

ACKNOWLEDGMENT

This work was supported by National Natural Science Foundation of China (No. 61872440 and No. 61828204), Beijing Municipal Natural Science Foundation (No. L182016), Royal Society Newton Advanced Fellowship (No. NAF\R2\192151), Youth Innovation Promotion Association CAS, CCF-Tencent Open Fund and Open Project Program of the National Laboratory of Pattern Recognition (No. 201900055).

REFERENCES

- [1] T. Neumann, K. Varanasi, S. Wenger, M. Wacker, M. Magnor, and C. Theobalt, "Sparse localized deformation components," *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, p. 179, 2013.
- [2] Z. Huang, J. Yao, Z. Zhong, Y. Liu, and X. Guo, "Sparse localized decomposition of deformation gradients," in *Computer Graphics Forum*, vol. 33, no. 7. Wiley Online Library, 2014, pp. 239–248.
- [3] Y. Wang, G. Li, Z. Zeng, and H. He, "Articulated-motion-aware sparse localized decomposition," in *Computer Graphics Forum*, vol. 36, no. 8. Wiley Online Library, 2017, pp. 247–259.
- [4] Q. Tan, L. Gao, Y. Lai, J. Yang, and S. Xia, "Mesh-based autoencoders for localized deformation component analysis," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2018, pp. 2452–2459.
- [5] C. N. Alleman, J. W. Foulk, A. Mota, H. Lim, and D. J. Littlewood, "Concurrent multiscale modeling of microstructural effects on localization behavior in finite deformation solid mechanics," *Computational Mechanics*, vol. 61, no. 1-2, pp. 207–218, 2018.
- [6] M. Mathew, A. Ellenberg, S. Esola, M. McCarthy, I. Bartoli, and A. Kotsos, "Multiscale deformation measurements using multispectral optical metrology," *Structural Control and Health Monitoring*, vol. 25, no. 6, p. e2166, 2018.
- [7] M. K. Abeyratne, W. Freeden, and C. Mayer, "Multiscale deformation analysis by cauchy-navier wavelets," *Journal of Applied Mathematics*, vol. 2003, 2002.
- [8] K. C. Lam, T. C. Ng, and L. M. Lui, "Multiscale representation of deformation via beltrami coefficients," *Multiscale Modeling & Simulation*, vol. 15, no. 2, pp. 864–891, 2017.
- [9] Y. Yang, W. Xu, X. Guo, K. Zhou, and B. Guo, "Boundary-aware multidomain subspace deformation," *IEEE transactions on visualization and computer graphics*, vol. 19, no. 10, pp. 1633–1645, 2013.
- [10] H. Hamidian, Z. Zhong, F. Fotouhi, and J. Hua, "Surface registration with eigenvalues and eigenvectors," *IEEE transactions on visualization and computer graphics*, 2019.
- [11] S.-L. Liu, Y. Liu, L.-F. Dong, and X. Tong, "Ras: A data-driven rigidity-aware skinning model for 3d facial animation," in *Computer Graphics Forum*, vol. 39, no. 1. Wiley Online Library, 2020, pp. 581–594.
- [12] J. Fu, H. Zheng, and T. Mei, "Look closer to see better: Recurrent attention convolutional neural network for fine-grained image recognition," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 2017, p. 3.
- [13] L. Gao, Y.-K. Lai, J. Yang, Z. Ling-Xiao, S. Xia, and L. Kobbelt, "Sparse data driven mesh deformation," *IEEE transactions on visualization and computer graphics*, 2019.
- [14] W. Xu, J. Wang, K. Yin, K. Zhou, M. Van De Panne, F. Chen, and B. Guo, "Joint-aware manipulation of deformable models," *ACM Transactions on Graphics (TOG)*, vol. 28, no. 3, pp. 1–9, 2009.
- [15] M. E. Yumer and L. B. Kara, "Co-constrained handles for deformation in shape collections," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 6, p. 187, 2014.
- [16] M. E. Yumer, S. Chaudhuri, J. K. Hodgins, and L. B. Kara, "Semantic shape editing using deformation handles," *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, p. 86, 2015.
- [17] K. Zhou, J. Huang, J. Snyder, X. Liu, H. Bao, B. Guo, and H.-Y. Shum, "Large mesh deformation using the volumetric graph laplacian," in *ACM SIGGRAPH 2005 Papers*, 2005, pp. 496–503.
- [18] L. Gao, Y.-K. Lai, D. Liang, S.-Y. Chen, and S. Xia, "Efficient and flexible deformation representation for data-driven surface modeling," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 5, p. 158, 2016.
- [19] M. Alexa and W. Müller, "Representing animations by principal components," in *Computer Graphics Forum*, vol. 19, no. 3. Wiley Online Library, 2000, pp. 411–418.
- [20] L. Gao, G. Zhang, and Y. Lai, "Lp shape deformation," *Science China Information Sciences*, vol. 55, no. 5, pp. 983–993, 2012.
- [21] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *Journal of Computational and Graphical Statistics*, vol. 15, p. 2006, 2004.
- [22] F. Bernard, P. Gemmar, F. Hertel, J. Goncalves, and J. Thunberg, "Linear shape deformation models with local support using graph-based structured matrix factorisation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 5629–5638.
- [23] J. R. Tena, F. De la Torre, and I. Matthews, "Interactive region-based linear 3D face models," in *ACM Transactions on Graphics (TOG)*, vol. 30, no. 4. ACM, 2011, p. 76.
- [24] T. Neumann, K. Varanasi, N. Hasler, M. Wacker, M. Magnor, and C. Theobalt, "Capture and statistical modeling of arm-muscle deformations," in *Computer Graphics Forum*, vol. 32, no. 2pt3. Wiley Online Library, 2013, pp. 285–294.
- [25] S. Fröhlich and M. Botsch, "Example-driven deformations based on discrete shells," in *Computer graphics forum*, vol. 30, no. 8. Wiley Online Library, 2011, pp. 2246–2257.
- [26] M. Chen, C. Wang, and L. Liu, "Cross-domain retrieving sketch and shape using cycle cnns," *Computers & Graphics*, 2020.
- [27] C. Jiang, D. Wang, J. Huang, P. Marcus, M. Nießner *et al.*, "Convolutional neural networks on non-uniform geometrical signals using euclidean spectral transformation," *arXiv preprint arXiv:1901.02070*, 2019.
- [28] K. Sarkar, K. Varanasi, and D. Stricker, "3d shape processing by convolutional denoising autoencoders on local patches," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 1925–1934.
- [29] C. Jiang, D. Lansigan, P. Marcus, and M. Nießner, "Ddsl: Deep differentiable simplex layer for learning geometric signals," in *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 8769–8778.
- [30] D. Maturana and S. Scherer, "VoxNet: a 3D convolutional neural network for real-time object recognition," in *IEEE Conference on Intelligent Robots and Systems*, 2015, pp. 922–928.
- [31] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller, "Multi-view convolutional neural networks for 3D shape recognition," in *IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 945–953.
- [32] M. E. Yumer and N. J. Mitra, "Learning semantic deformation flows with 3D convolutional networks," in *European Conference on Computer Vision (ECCV)*. Springer, 2016, pp. 294–311.
- [33] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong, "O-cnn: Octree-based convolutional neural networks for 3D shape analysis," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, p. 72, 2017.
- [34] P.-S. Wang, C.-Y. Sun, Y. Liu, and X. Tong, "Adaptive o-cnn: A patch-based deep representation of 3d shapes," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 6, pp. 1–11, 2018.

- [35] P.-S. Wang, Y. Liu, and X. Tong, "Deep octree-based cnns with output-guided skip connections for 3d shape and scene completion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 266–267.
- [36] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen, "Pointcnn: Convolution on x-transformed points," *Advances in Neural Information Processing Systems (NIPS)*, pp. 820–830, 2018.
- [37] H. Lei, N. Akhtar, and A. Mian, "Spherical convolutional neural network for 3D point clouds," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 9631–9640.
- [38] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, "Dynamic graph cnn for learning on point clouds," *Acm Transactions On Graphics (tog)*, vol. 38, no. 5, pp. 1–12, 2019.
- [39] B. Shi, S. Bai, Z. Zhou, and X. Bai, "DeepPano: Deep panoramic representation for 3-D shape recognition," *IEEE Signal Processing Letters*, vol. 22, no. 12, pp. 2339–2343, 2015.
- [40] K. Sarkar, B. Hampiholi, K. Varanasi, and D. Stricker, "Learning 3d shapes as multi-layered height-maps using 2d convolutional networks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 71–86.
- [41] H. Huang, E. Kalegorakis, S. Chaudhuri, D. Ceylan, V. Kim, and E. Yumer, "Learning local shape descriptors with view based convolutional neural networks," *ACM Transactions on Graphics (TOG)*, vol. 2, 2018.
- [42] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, "Convolutional networks on graphs for learning molecular fingerprints," in *Advances in neural information processing systems (NIPS)*, 2015, pp. 2224–2232.
- [43] M. Niepert, M. Ahmed, and K. Kutzkov, "Learning convolutional neural networks for graphs," in *International Conference on Machine Learning (ICML)*, 2016, pp. 2014–2023.
- [44] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in Neural Information Processing Systems (NIPS)*, 2016. [Online]. Available: <https://arxiv.org/abs/1606.09375>
- [45] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, "Spectral networks and locally connected networks on graphs," in *International Conference on Learning Representations (ICLR)*, 2014.
- [46] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," in *Advances in Neural Information Processing Systems (NIPS)*, 2016, pp. 3844–3852.
- [47] Q. Tan, L. Gao, Y. Lai, and S. Xia, "Variational autoencoders for deforming 3D mesh models," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 5841–5850.
- [48] A. Ranjan, T. Bolkart, S. Sanyal, and M. J. Black, "Generating 3D faces using convolutional mesh autoencoders," in *European Conference on Computer Vision (ECCV)*. Springer International Publishing, 2018, pp. 725–741. [Online]. Available: <http://coma.is.tue.mpg.de/>
- [49] E. Tretschk, A. Tewari, M. Zollhöfer, V. Golyanik, and C. Theobalt, "DEMEA: Deep Mesh Autoencoders for Non-Rigidly Deforming Objects," *arXiv e-prints*, 2019.
- [50] O. Litany, A. Bronstein, M. Bronstein, and A. Makadia, "Deformable shape completion with graph convolutional autoencoders," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 1886–1895.
- [51] L. Fulton, V. Modi, D. Duvenaud, D. I. Levin, and A. Jacobson, "Latent-space dynamics for reduced deformable simulation," in *Computer graphics forum*, vol. 38, no. 2. Wiley Online Library, 2019, pp. 379–391.
- [52] S. Ma, J. Fu, C. W. Chen, and T. Mei, "DA-GAN: Instance-level image translation by deep attention generative adversarial networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 5657–5666.
- [53] J. Si, H. Zhang, C.-G. Li, J. Kuen, X. Kong, A. C. Kot, and G. Wang, "Dual attention matching network for context-aware feature sequence based person re-identification," *2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [54] J. Xu, R. Zhao, F. Zhu, H. Wang, and W. Ouyang, "Attention-aware compositional network for person re-identification," *2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [55] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, and E. Hovy, "Hierarchical attention networks for document classification," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016, pp. 1480–1489.
- [56] X. Zhang, T. Wang, J. Qi, H. Lu, and G. Wang, "Progressive attention guided recurrent network for salient object detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 714–722.
- [57] B. Zhuang, Q. Wu, C. Shen, I. Reid, and A. van den Hengel, "Parallel attention: A unified framework for visual object discovery through dialogs and queries," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4252–4261.
- [58] X. Long, C. Gan, G. de Melo, J. Wu, X. Liu, and S. Wen, "Attention clusters: Purely attention based local feature integration for video classification," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7834–7843.
- [59] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [60] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *Advances In Neural Information Processing Systems (NIPS)*, 2016, pp. 289–297.
- [61] P. Wang, Q. Wu, C. Shen, and A. van den Hengel, "The VQA-machine: Learning how to use existing vision algorithms to answer new questions," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 4, 2017, pp. 1173–1182.
- [62] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang, "Residual attention network for image classification," *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3156–3164, 2017.
- [63] O. Sorkine and M. Alexa, "As-rigid-as-possible surface modeling," in *Proceedings of EUROGRAPHICS/ACM SIGGRAPH Symposium on Geometry Processing*, 2007, pp. 109–116.
- [64] K. Crane, C. Weischedel, and M. Wardetzky, "Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow," *ACM Transactions on Graphics (TOG)*, vol. 32, pp. 152:1–152:11, 2013.
- [65] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis, "SCAPE: shape completion and animation of people," in *ACM transactions on graphics (TOG)*, vol. 24, no. 3. ACM, 2005, pp. 408–416.
- [66] R. W. Sumner and J. Popović, "Deformation transfer for triangle meshes," *ACM Transactions on graphics (TOG)*, vol. 23, no. 3, pp. 399–405, 2004.
- [67] L. Zhang, N. Snavely, B. Curless, and S. M. Seitz, "Spacetime faces: High-resolution capture for modeling and animation," in *ACM Transactions on Graphics (SIGGRAPH)*, 2004, pp. 548–558.
- [68] R. White, K. Crane, and D. A. Forsyth, "Capturing and animating occluded cloth," in *ACM Transactions on Graphics (TOG)*, vol. 26, no. 3. ACM, 2007, p. 34.
- [69] G. Pons-Moll, J. Romero, N. Mahmood, and M. J. Black, "Dyna: A model of dynamic human shape in motion," *ACM Transactions on Graphics, (Proc. SIGGRAPH)*, vol. 34, no. 4, pp. 120:1–120:14, Aug. 2015.
- [70] D. Vlasic, I. Baran, W. Matusik, and J. Popović, "Articulated mesh animation from multi-view silhouettes," in *ACM Transactions on Graphics (TOG)*, vol. 27, no. 3. ACM, 2008, pp. 97:1–97:9.
- [71] D. Kingma and J. Ba, "ADAM: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [72] R. W. Sumner and J. Popović, "Deformation transfer for triangle meshes," *ACM transactions on graphics (TOG)*, vol. 23, no. 3, pp. 399–405, 2004.
- [73] L. Kavan, P.-P. Sloan, and C. O'Sullivan, "Fast and efficient skinning of animated meshes," in *Computer Graphics Forum*, vol. 29, no. 2. Wiley Online Library, 2010, pp. 327–336.
- [74] L. Vasa and V. Skala, "A perception correlated comparison method for dynamic meshes," *IEEE transactions on visualization and computer graphics*, vol. 17, no. 2, pp. 220–230, 2011.
- [75] R. W. Sumner, M. Zwicker, C. Gotsman, and J. Popović, "Mesh-based inverse kinematics," in *ACM transactions on graphics (TOG)*, vol. 24, no. 3. ACM, 2005, pp. 488–495.