

# Mesh Convolutional Autoencoder for Semi-Regular Meshes of Different Sizes

Sara Hahner and Jochen Garcke

Fraunhofer Center for Machine Learning and SCAI, Sankt Augustin, Germany  
Institut für Numerische Simulation, Universität Bonn, Germany

sara.hahner@scai.fraunhofer.de

## Abstract

*The analysis of deforming 3D surface meshes is accelerated by autoencoders since the low-dimensional embeddings can be used to visualize underlying dynamics. But, state-of-the-art mesh convolutional autoencoders require a fixed connectivity of all input meshes handled by the autoencoder. This is due to either the use of spectral convolutional layers or mesh dependent pooling operations. Therefore, the types of datasets that one can study are limited and the learned knowledge cannot be transferred to other datasets that exhibit similar behavior. To address this, we transform the discretization of the surfaces to semi-regular meshes that have a locally regular connectivity and whose meshing is hierarchical. This allows us to apply the same spatial convolutional filters to the local neighborhoods and to define a pooling operator that can be applied to every semi-regular mesh. We apply the same mesh autoencoder to different datasets and our reconstruction error is more than 50% lower than the error from state-of-the-art models, which have to be trained for every mesh separately. Additionally, we visualize the underlying dynamics of unseen mesh sequences with an autoencoder trained on different classes of meshes.*

## 1. Introduction

We study three-dimensional data that is discretized by a triangular surface mesh. In particular, we study the deformation of surfaces, which discretize human bodies, animals, or work pieces from computer aided engineering. Surface deformation is locally described by the same physical rules, which motivates the application of convolution to learn translation-invariant localized features.

Convolutional neural networks (CNN) are successful in the analysis and generation of data, especially images, because of their efficient calculation of translation-invariant localized features by sliding filters over the images [26]. Regular pixel grids describe the 2D images. This global grid structure, determined by the two axes of the two-




Figure 1. MNIST dataset sample [25] with a visualization of the underlying grid defined by the pixels aligned along the two axes.

dimensional space (see Figure 1), is essential for CNNs, because it implies properties such as a common system of coordinates, shift invariance, and a fixed neighborhood structure [7]. These characteristics allow for an efficient application of the local kernels, the sliding of the kernels along the two axes, and a constant definition of the pooling operator.

While two-dimensional surfaces embedded in  $\mathbb{R}^3$  are locally homeomorphic to the two-dimensional space, the surfaces are of non-Euclidean nature. Therefore, they generally lack the global grid structure, which is so essential for the efficient application of CNNs. Furthermore, the meshes are usually heterogeneous in the number of vertices, faces, their connectivity, and size, which hinders the direct application of 2D-convolution.

In our novel approach, we calculate an alternative discrete approximation of the surface data based on semi-regular meshes. Semi-regular meshes have regular regional patches, which means that every vertex inside the patch has exactly six neighbors. Also semi-regular meshes have by definition a multi-scale structure that allows us to look at the meshes in different resolutions [34].

To this semi-regular approximation of our data, we apply spatial convolution, which follows the idea of 2D-convolution and defines kernels on local neighborhoods of the vertices [7]. Since the neighborhoods of most of the vertices of the semi-regular mesh are regular, we use convolutional kernels that can be straightforwardly implemented and preserve the orientation of the neighborhood as well as the permutation of neighbors. Additionally, the multi-

scale structure of the semi-regular mesh allows us to define a general pooling operation that works on all semi-regular meshes. Since the CNNs learn local features, we feed the regional patches separately to the network. This allows us to apply the network to meshes of different sizes and topology. The global context is not lost, but fed to the network via padding.

The research objectives can be summarized as a) a remeshing approach to preprocess surface meshes into a representation by semi-regular meshes, which turns out to be more beneficial for CNNs, b) the definition of an autoencoder that handles semi-regular meshes of different size and topology, and by this means c) creating a possibility to transfer and apply trained models to different classes of surface data<sup>1</sup>.

Further on in section 2, we discuss related work. In section 3, we present some theoretical background of surface meshes and relevant characteristics for CNNs, followed by the definition of our convolution and pooling for semi-regular meshes in section 4. In section 5, we introduce our mesh convolutional autoencoder for semi-regular meshes. Results for different datasets are presented in section 6.

## 2. Related Work

### Semi-Regular Meshes

To the best of our knowledge, triangular semi-regular meshes have not been used in the context of convolutional neural networks for graph data of different shape. Their piecewise regular structure has been in focus for multi-resolution analysis, because the iterative refinement allows analysis at different levels of resolution, which is especially interesting to adapt wavelets to surfaces and geometry compression [22, 28, 34]. They are generally obtained by remeshing of irregular meshes. In [21, 34] an overview of semi-regular triangle remeshing algorithms is given.

The authors of [2] train a neural network on a quadrilateral surface mesh that is mapped to a box. The box's sides coincide with a general grid-based mesh as for images in 2D allowing the application of 2D convolution. Although called semi-regular quadrilateral meshes by the authors, their definition of semi-regular is different than ours, since the regular sides of the boxes are not created by iterative subdivision, which allows us to define the pooling operator. This remeshing approach can only be applied to shapes without boundaries, that can be mapped to a box. In [20] a semi-regular mesh structure is used for efficient pooling and unpooling, but the network is not independent of the mesh size and limited to meshes without boundaries.

---

<sup>1</sup>Source code available at: [https://github.com/Fraunhofer-SCAI/conv\\_sr\\_mesh\\_autoencoder](https://github.com/Fraunhofer-SCAI/conv_sr_mesh_autoencoder)

### Convolutional Networks for Graphs and Surfaces

Generally, there are spectral and spatial convolutional networks for graphs, of which [7, 45] give an overview.

At first, [9] exploited the connection of the graph Laplacian and the Fourier basis and they project vertex features to the Laplacian eigenvector basis. Instead of explicitly computing Laplacian eigenvectors, the authors of [12] use truncated Chebyshev polynomials and in [24] they use only first-order Chebyshev polynomials. These spectral methods require a fixed connectivity of the graph. If not, the basis functions change and the features that the network learns are not guaranteed to be meaningful.

Spatial methods for convolution on graphs aggregate features from the neighbors of the vertices. At first, this idea was presented under the name Neural Network for Graphs [30]. Spatial methods allow generalization across different domains and because of their flexibility and efficiency these methods are very popular [45]. Since surface meshes lack a general underlying grid, the orientation of kernels has to be defined with respect to their neighborhood. To avoid this difficulty, kernels often calculate rotation invariant features, which are sometimes referred to as orientation invariant. The authors of [29, 31] calculate rotation invariant features by averaging the result of different anisotropic kernels that are sensitive to orientation. In [5] the kernels are aligned with the principal curvature direction and [10, 11] introduce anisotropic gauge equivariant kernels that encode orientation information in the features, but are computationally expensive. The authors of [10] demonstrate that for special meshes (in their case the icosahedron having areas with regular connectivity) the anisotropic gauge equivariant convolution can be implemented efficiently.

The work of [14] sums up many spatial approaches by their Message Passing Neural Network. It interprets graph convolutions as a message passing process, in which information is passed from one node to another along the edges. We want to point out that usually the topology and the number of vertices is fixed, because the pooling [6, 36, 46] or the spectral convolutional layers require a fixed connectivity [12, 24, 36].

It is possible to consider the vertices of meshes describing surfaces as point clouds, which would allow the application of convolutional architectures for 3D point clouds as [38, 42, 44] handling pointwise inputs that consider neighborhoods via kernels. Albeit being the more flexible representation, it lacks an underlying structure that describe the surface, whose deformation we want to analyze [6]. Also, the mesh provides native connectivity information [18].

### Mesh Convolutional Autoencoders

The authors of [41] present a variational autoencoder for deforming 3D meshes that does not handle meshes but feature representations of the deforming meshes.

In [36] a first convolutional mesh autoencoder (CoMA) has been introduced that handles surface meshes directly. The authors introduced mesh downsampling and mesh up-sampling layers, which have a similar effect as pooling and unpooling. They are combined with spectral convolutional filters using truncated Chebyshev polynomials as in [12]. The Neural 3D Morphable Models (Neural3DMM) network presented in [6] improves those results using the same down and upsampling layers in combination with spiral convolutional layers that break the permutation invariance. The authors of [46] apply the CoMA [36] to different datasets and improve the down and upsampling layers slightly.

The presented mesh convolutional autoencoders work only for meshes of the same size and connectivity, since the downsampling and upsampling layers as well as spectral convolutional layers depend on the adjacency matrix.

The authors of [18] present the MeshCNN architecture that allows an implementation of an encoder and decoder. The edge collapsing based pooling is feature dependent and therefore the low-dimensional mesh embeddings of the deforming meshes can be of different significance.

### 3. Surface Meshes and Their Characteristics

Surfaces in  $\mathbb{R}^3$  are generally discretized by triangular polygonal surface meshes.

**Definition 1 (Triangular Polygonal Mesh)** A *triangular polygonal mesh*  $\mathcal{M}$  is defined by a set of vertices  $V \subset \mathbb{R}^d$  and a set of triangular faces  $F \subset V \times V \times V$ , which describe the shape and point to the vertices they use. The edges  $E = \{\{v_1, v_2\} \in V \times V \mid \exists f \in F \text{ s.t. } v_1 \in f \text{ and } v_2 \in f\}$  are undirected, i.e. if  $(v, w) \in E$ ,  $(w, v)$  is also in  $E$ .

If the mesh is a surface or manifold mesh, it holds that every edge  $e \in E$  is adjacent to at most two faces in  $F$ . A boundary edge  $e \in E$  is adjacent to exactly one face in  $F$ . The vertices  $v \in V$  of a mesh  $\mathcal{M}$  store information which is defined by a function  $f : V \rightarrow \mathbb{R}^m$ . For each vertex  $v \in V$ , we define the  $r$ -ring neighborhood  $N_r(v)$  as all the vertices  $w \in V$  that are connected to  $v$  by at most  $r$  edges in  $E$ . The degree of a vertex  $v \in V$  is the size of its one-ring neighborhood  $N_1(v)$ .

We refer to a triangular surface mesh as regular, if the degree of all vertices in  $V$  is 6 [34]. Note, that a regular mesh has limited representation power and not every mesh can be remeshed into a regular one because they cannot approximate all types of curvature (see hedgehog theorem [8]).

Images represented in pixels can also be interpreted as two-dimensional surface meshes of rectangular shape in  $\mathbb{R}^2$  [7]. Every pixel is a vertex, the feature function  $f$  outputs the pixels' color values and every vertex is connected by an edge to the eight neighboring pixels (horizontal, vertical, and diagonal).

	2D Image	Surface Mesh	Semi-Regular Surface Mesh
data	information saved on vertices		
grid structure	global structure	locally Euclidean	locally Euclidean
connectivity	fixed	-	semi-regular
distance to neighbors	fixed	-	-
size of instances	similar	highly different	highly different with similar local patches

Table 1. Characteristics of images and surface meshes in  $\mathbb{R}^3$  relevant for CNNs

### Convolution for Images and Surface Meshes

Although surface meshes in  $\mathbb{R}^3$  and images in  $\mathbb{R}^2$  both fulfill the definition of polygonal meshes, the overall structure and regularity of the data is highly different, which complicates the application of convolution to surface meshes in  $\mathbb{R}^3$  [7].

CNNs in 2D [16, 26] apply the same local filters to local neighborhoods of selected pixels of the image. Because of the global grid structure of the image, the filters can be horizontally and vertically shifted. The filters can be of constant shape and the networks apply them to every local neighborhood. As [10] pointed out, the shifting of the filters is not well-defined for surface meshes because of the lack of a global grid. Also, the local neighborhoods of a surface mesh can have any size and arrangement as long as they are locally Euclidean.

For images, the size of each instance is usually constant. Since the samples are of rectangular shape, they can easily be resized or padded into the desired size. The constant size allows one neural network to handle all the data and allows the use of existing networks as pretrained prototypes for different applications. However, the size of surface meshes varies strongly in general. Table 1 gives an overview of the mentioned mesh characteristics of images and surface meshes. The authors of [7] summarize that both the similar structure for local neighborhoods and the underlying global grid are reasons why CNNs work so efficiently for images.

As mentioned, one cannot enforce a regular mesh discretization for every surface in  $\mathbb{R}^3$ , which would lead to an underlying global grid [8]. We aim to enforce a similar structure in the local neighborhoods by choosing a different approximation of the surface. In this way, an efficient application of convolution on surface meshes becomes possible. Note, that remeshing the polygonal mesh only changes the representation of the objects. The considered surface embedded in  $\mathbb{R}^3$  is the same, but now represented by a different discrete approximation.

## Semi-Regular Meshes

Semi-regular meshes are a flexible representation of surfaces in  $\mathbb{R}^3$ , that have a regular local structure and allow irregularities at selected vertices, whose positions we can control [34]. The overview in Table 1 shows that the mesh characteristics of semi-regular meshes are closer to the ones of 2D images than compared to general surface meshes.

We follow the definition of semi-regular meshes from [34], which gives one condition on the specific structure: iteratively merging four triangular faces of a semi-regular mesh into one leads to a low-resolution mesh. This means that all vertices are regular (i.e. have six neighbors) besides the vertices of the low-resolution mesh (see Figure 2). Therefore, a semi-regular mesh is obtained by regular subdivision of a low resolution mesh that can be irregular. We refer to the faces of the low-resolution mesh that are iteratively subdivided as regional patches.

Note that the iterative subdivision of the low-resolution mesh automatically defines a multi-scale structure. This is why semi-regular meshes are well suited for multi-resolution analysis [27, 34]. Later on, this structure allows us to define a local pooling operator on the semi-regular meshes.

## Remeshing

At first, a coarse approximation of the input mesh is built. To coarsen the surface meshes we employ a slightly adapted Garland-Heckbert-algorithm for surface simplification using quadric error metrics [13]. Afterwards, every face of the coarse base mesh is iteratively subdivided into four new faces for a given number of times. All newly created vertices have six neighbors. When the desired level of refinement is reached, the resulting semi-regular mesh geometry has to be fit to the original irregular mesh in order to describe the surface well. If for one deforming shape the mesh topology is constant over time, it is enough to remesh a template mesh. The semi-regular remeshing result can be transferred to meshes of the same topology.

We provide more details to the remeshing algorithm in the supplementary material.

## 4. Convolution and Pooling for Semi-Regular Meshes

Convolutional kernels consider local features in the neighborhood of the vertices. The size of the considered neighborhood  $N_r(v)$  is given by the kernel size  $r$ . In a convolutional layer, a set of kernels is applied to every vertex of the input. Generally, the number of vertices and the mesh's connectivity have to be constant. Only the vertex features change.

Since the convolutional networks learn local features, we propose to input every regional patch of the semi-regular

mesh separately. This allows us to handle meshes of different size. To the best of our knowledge, we thereby present the first mesh convolutional autoencoder that handles meshes of different size. In order to not lose the embedding of the regional patch in the whole mesh, we consider the environment of the patches via padding.

## Hexagonal Convolution

The regional patches are of the same regular structure. All vertices have exactly six neighbors, only the three corners can be irregular, but we project their neighborhood to a regular one. Additionally, the patches are intrinsically two-dimensional and represent a surface. Therefore, the application of a 2D-convolutional kernel is possible. Since the regional patches are represented in hexagonal grids, the application of hexagonal 2D-convolutional kernels has shown to give better results [19, 39]. Similarly to [2], the consistent degree of the vertices results in better runtimes since similar calculations at these vertices can be moved to GPU.

On the local regular structure, the translation of the convolutional kernels is well-defined. Therefore the kernels preserve the orientation of the neighborhood and are anisotropic. The padded patch based approach assures gauge equivariance of the network. The authors of [10, 11] show how anisotropic kernels that preserve orientation significantly improve the expressivity of models.

Note that the network does not correct differences in the distances to neighbors or angles between neighbors. The edge lengths of the semi-regular meshes are stable, because of the edge length regularization during the remeshing.

## Pooling

The piecewise regular form of the semi-regular meshes has a multi-scale structure, which is created by the iterative subdivision of the faces of the low-resolution mesh. We take advantage of this structure that all semi-regular meshes have in common, and define an average pooling operator, that undoes the subdivision of one into four faces. Herewith, we reduce the dimensions of the features and the number of network parameters, as pooling layers for 2D convolution do as well [16]. Since we apply average pooling, the vertices that are kept during the pooling take the average of their own value and the values of the neighboring vertices in the one-ring neighborhood that are removed.

To increase the resolution of the mesh patches in the decoder, the unpooling operator recreates the multi-scale structure of the semi-regular mesh. Every face is subdivided into four faces. The newly created vertices are assigned the average value of neighboring vertices from the lower-resolution mesh patch. Figure 3 illustrates how the pooling and unpooling layers undo the subdivision of the regular patch or increase its resolution respectively.




Figure 2. Remeshing of the horse template mesh. After coarsening the original mesh to a low resolution base mesh, the faces are subdivided three times. In the semi-regular mesh the faces adjacent to an irregular vertex are highlighted in red.




Figure 3. The pooling layers change the resolution of the semi-regular meshes. We input the padded regular patch and apply pooling twice, which undoes the subdivision of the faces. In the decoder the unpooling increases the resolution again by subdividing the faces.

## Padding

The padding is crucial for the network to consider the regional patches in a larger context. Since the network handles the patches separately, we consider the features of the neighboring patches in the padding.

We apply a padding of the size of the first layer’s kernel size. The padding considers the vertices of the neighboring patches. If the number of neighbors is lower than six, we interpolate the values for empty vertices. If the number is higher than six, we take the closest vertices in both cyclic rotations. If the vertices are boundary vertices of the whole mesh, we decide to pad the patch with the boundary vertices’ features. Figure 3 shows a padded regular patch.

## 5. Autoencoder for Semi-Regular Meshes of Different Size

The network handles the regional patches separately, which allows us to handle meshes of different sizes. The information of the neighboring patches is not lost but included in the padding.

The mesh autoencoder consists of an encoder and a decoder. The encoder consists of 2 hexagonal convolutional layers (implementation of [39]) of kernel sizes 2 and 1. Each of the convolutions is followed by a biased ReLU [15]. The average pooling layers (see section 4) are interleaved

between convolutional layers. The encoder transforms every padded patch, which corresponds to one face of the low-resolution mesh, from  $\mathbb{R}^{111 \times 3}$  to an 8-dimensional latent vector using a fully connected layer at the end. Following the decoder’s fully connected layer, 2 hexagonal convolutional layers (followed by a biased ReLU) with interleaved average unpooling layers reconstruct the patches. Each unpooling layer subdivides every face into 4 faces, following the subdivision process of the remeshing. The last layer is a hexagonal convolutional layer without activation function that reduces the number of features to three dimensions. A detailed structure of the network is supplied as additional material together with the distribution of the 18184 trainable weights.

Note that we are able to handle non-manifold edges of the coarse base mesh because the patches, whose interiors by construction have only manifold-edges, are fed separately. Figure 3 illustrates the patch sizes inside the autoencoder.

## 6. Experiments

We test our convolutional autoencoder for semi-regular meshes on four different datasets and compare the achieved reconstruction errors to state-of-the-art models.

## Datasets

**GALLOP:** The authors of [40] present a dataset containing triangular meshes representing a motion sequence from a galloping horse, elephant, and camel. Each sequence has 48 timesteps. The three animals move in a similar way but the meshes that represent the surfaces of the three animals are highly different in connectivity and in the number of vertices (horse: 8,431, camel: 21,887, elephant: 42,321). This is why the authors of [46] trained three different mesh autoencoders as presented in [36]. The surface approximations are remeshed to semi-regular meshes for each animal. The coarse base meshes of approximately 110 faces are subdivided 3 times. The new meshes are still of different connectivity, but all are made up of regional regular patches. The resulting numbers of vertices are listed in Table 2.

We normalize the 3D coordinates to  $[-1, 1]$  relative to the coordinates' ratio and translate every patch to zero mean. For the training of CoMA [36] and Neural3DMM [6] every vertex was normalized to zero mean and a standard deviation of one. The patch based approach does not allow a vertex wise normalization if we want to learn and transform the local deformation. This is why for the baseline training of the CoMA and Neural3DMM on the remeshed data, we normalize the whole mesh to zero mean and standard deviation of one.

We use the first 75% of the galloping sequence of the horse and camel for training of the network. The architecture is tested on the remaining 25% and the whole sequence of the elephant, which is never seen during the training.

**FAUST:** We select 100 meshes from the FAUST dataset [3], which are in correspondence to each other, to be able to compare to the other architectures. The dataset consists of 10 different bodies in 10 different poses. The irregular surface meshes have 6890 vertices. We conduct two different experiments: at first we consider known poses of two unseen bodies in the testing set. Then we consider two unknown poses of all bodies in the testing set. In both cases, 20% of the data is included in the testing set.

The meshes are remeshed to a semi-regular mesh representation. The data and the mesh patches are normalized in the same way as for the GALLOP dataset.

**TRUCK and YARIS:** In a car crash simulation the car components, which are generally represented by surface meshes, often deform in different patterns. Every component is discretized by a surface mesh, while the local deformation is described by the same physical rules.

The dataset TRUCK consists of 32 completed frontal crash simulations of a Chevrolet C2500 pick-up truck (from NCAC [32]), using the same truck, but with different material characteristics, which is a similar setup to [4]. For this setup the authors of [17] detect patterns in the deformation of the components using a general representation by oriented bounding boxes, which allowed them to train one au-




Figure 4. Anisotropic filters from the first and second convolutional layers of the encoder trained on the GALLOP dataset.

toencoder for the whole car model. For our analysis, we select 6 components (the front and side beams), whose meshes are remeshed to semi-regular meshes, and 30 equally distributed time steps. The model is trained on 30 simulations and 70% of the timesteps. The trained architecture is tested on 2 complete simulations and on the remaining timesteps for the other 30 simulations.

To study the transfer learning capacities of our architecture, we test the architecture that is trained on the TRUCK dataset on a different dataset, which also contains deforming components from a different car crash simulation. The YARIS dataset consists of 10 completed frontal crash simulations of a detailed model of the Toyota Yaris (from NCAC [32]) with different material characteristics. For our analysis, we select 10 components (the front and side beams plus the crashbox), whose meshes are remeshed to semi-regular meshes, and 26 equally distributed time steps.

We normalize the meshes that discretize car components to zero mean and range  $[-1, 1]$  relative to the coordinates' ratio. Every patch is translated to zero mean.

## Training Details

We train the network (implemented in Pytorch [33]) with the adaptive learning rate optimization algorithm [23] using a learning rate of 0.001. For the GALLOP and the FAUST dataset we train for 500 epochs using a batch size of 100. For the TRUCK data we chose 250 epochs and a batch size of 50, since the variation inside the dataset is higher. We minimize the mean squared error between original and reconstructed regional patches of the surface mesh without considering the padding. To augment the data in the case of the GALLOP and the FAUST dataset we rotate the regional patches by  $0^\circ$ ,  $120^\circ$  and  $240^\circ$ .

Figure 4 shows trained hexagonal anisotropic kernels, which implies sensitivity to orientation.

## Reconstructions of the Meshes

In Table 2 we compare our hexagonal mesh autoencoder to the CoMA [36] and the Neural3DMM [6] network for the GALLOP and FAUST dataset in terms of the mean squared errors of reconstructed unseen shapes, whose 3D coordinates lie in range  $[-1, 1]$ . Albeit training the autoencoder for two different animal meshes of the GALLOP dataset, our network reduces the reconstruction error of unseen data

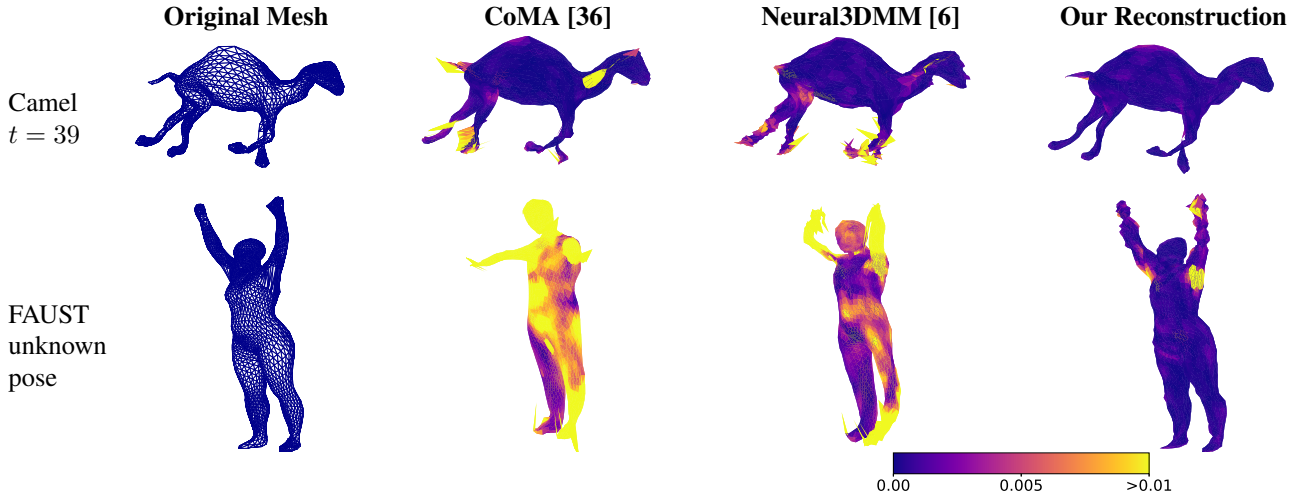


Figure 5. Reconstructed GALLOP and FAUST test samples by CoMA [36], Neural3DMM [6], and our network. The mean squared error of the reconstructed faces is highlighted. More reconstruction examples are given in the supplementary material.

Mesh Class	# Vertices	CoMA [36] per mesh	Neural3DMM [6] per mesh	Ours
FAUST	3501			
known poses		0.0033 + 0.0058	0.00190 + 0.0037	0.00032 + 0.0007
unknown poses		0.1963 + 0.3253	0.04233 + 0.0707	0.00054 + 0.0011
Horse	3601	0.00044 + 0.0015	0.00072 + 0.0023	0.00020 + 0.0003
Camel	3467	0.00051 + 0.0018	0.00186 + 0.0116	0.00020 + 0.0003
Elephant	3781	0.00088 + 0.0055	0.00321 + 0.0267	0.00081 + 0.0160 <sup>†</sup> 0.00155 + 0.0163*

Table 2. Mean squared errors of reconstructed unseen meshes and their standard deviations for two different training runs. We train one autoencoder for all three animals in the GALLOP dataset and one autoencoder for each experiment on the FAUST dataset.

\*: the elephant has not been seen by the network during training. <sup>†</sup>: include elephant in training set.

Mesh Class	CoMA [36]	Neural3DMM [6]	Ours
FAUST	26795	276275	18184
Horse	27339	280499	
Camel	26795	292659	18184
Elephant	27339	296883	

Table 3. Comparison of number of parameters. Our network requires at least 30% fewer parameters.

by more than 50%, if the animal is presented to the autoencoder during the training. We are also able to reconstruct a galloping sequence of an unseen elephant, although the reconstruction error is higher than with a baseline network trained only for this animal. If we include the elephant in the training set, the reconstruction errors are slightly lower compared to CoMA. Especially the reconstruction of the legs is superior with our method in comparison to CoMA and Neural3DMM, as Figure 5 illustrates. Note that the mesh reconstruction is smooth at the patches’ boundaries. This indicates that the padding enforces the incorporation of information from the neighboring patches.

Our network reconstructs known and unknown poses of

Dataset	Train MSE	Test MSE	Eucl. E.
TRUCK	0.0026 + 0.004	0.0027 + 0.004	5.49   3.76
YARIS	–	0.0147 + 0.023	2.23   1.92

Table 4. Mean squared errors of reconstructed unseen meshes and their standard deviations for two different training runs on the TRUCK and YARIS dataset. Additionally, the average Euclidean vertex wise error and its median (in cm) are given.

FAUST with a more than 80% lower error, see Table 2 and Figure 5. Limbs are reconstructed inaccurately by the CoMA and Neural3DMM architectures. Especially if the pose is unknown and not similar to training poses, their reconstruction fails.

In all cases, our architecture requires fewer parameters than the CoMA and Neural3DMM networks (Table 3).

Since the TRUCK and YARIS datasets contain 16 different meshes, we would have to train the baseline architectures 16 times. To validate the reconstruction results, we here calculate the error in cm and have a closer look at the embeddings. In Table 4 we present the mean squared errors and average Euclidean vertex wise error and median



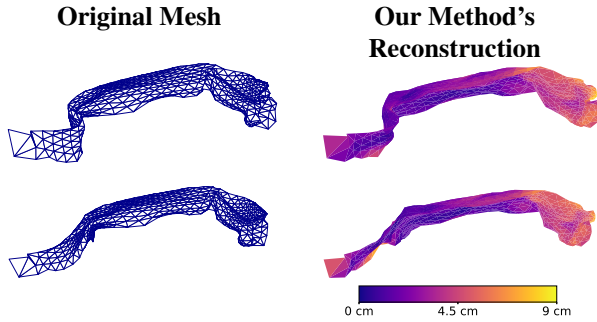


Figure 6. Reconstructed left front beam from the TRUCK (length of 150 cm) at time  $t = 24$  (test sample) from two different crash simulations. The average Euclidean distance (in cm) of the faces is highlighted. In the majority of faces the error is below 4.5 cm and only in highly deformed areas higher.

for the TRUCK dataset, for which the autoencoder has been trained, and for the YARIS dataset that has not been presented to the network during training. The TRUCK components measure between 135 and 370 cm in length, the YARIS components between 21 and 91 cm. If we put the error in relation to the part length, the average Euclidean error on the unseen YARIS dataset is less than 3 times higher in comparison to the testing error on the TRUCK, for which selected reconstruction results are presented in Figure 6.

### Low-dimensional Embedding

For the three different animals from the GALLOP dataset we can visualize the low-dimensional embeddings of the galloping sequences of the three animals. We concatenate the patch wise embeddings for each timestep and project the resulting vectors to the two-dimensional space using Principal Component Analysis (PCA) [35]. The time dependent embeddings for all three animals, whether included in the training or only in the testing set, exhibit a periodic galloping sequence in the two-dimensional space, as seen in Figure 7 for the elephant’s embedding. The embeddings of the cyclic sequences are similar to the ones from [46], who trained one autoencoder for each animal and showed that the direct application of PCA or t-SNE cannot reveal the intrinsic information of the data.

For each car component we create a 2D-visualization of the low-dimensional representation using t-SNE [43] to detect patterns in the deformation, that separates the simulations into clusters. This speeds up the analysis of car crash simulations, since relations between model parameters and the deformation behavior are discovered faster [4, 17]. For the TRUCK we observe that the selected components deform in two different branches for the 32 simulations (see supplementary material). This behavior manifests after approximately half of the time similar to [4, 17]. For the YARIS, which has never been seen by the network during




Figure 7. Embedded cyclic sequence of the galloping elephant, whose mesh has not been seen during training.




Figure 8. Embedded low-dimensional representation of the front beam of the YARIS. For the timesteps  $t = 7, \dots, 26$  we observe two patterns in the deformation behavior. The mesh has not been presented to the autoencoder during training.

training, we visualize the low-dimensional representation in 2D using t-SNE [43]. We detect a deformation pattern in the front beam that splits up the simulation set into two clusters, see Figure 8.

### 7. Conclusion

We have introduced a novel approach for the analysis of deforming 3D surface meshes with a mesh autoencoder for semi-regular meshes. To the best of our knowledge, the remeshing of the triangular meshes into semi-regular ones allows for the first time an analysis of shapes of different size and geometry. The regular local structure makes a re-utilization of the learned convolutional filters and an efficient mesh-independent pooling operation possible.

We evaluate our network successfully on three datasets from different domains and reconstruct the meshes in significantly better quality than the baseline mesh autoencoders. Additionally, we apply our trained autoencoder to unseen meshes of different shape and connectivity and successfully detect the underlying dynamics of unseen time sequences. In future work we plan to apply the architecture to other tasks such as shape matching and segmentation.



## References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. *35th International Conference on Machine Learning, ICML 2018*, 1:67–85, 2018.
- [2] Pierre Baque, Edoardo Remelli, Francois Fleuret, and Pascal Fua. Geodesic convolutional shape optimization. *35th International Conference on Machine Learning*, 2:797–809, 2018.
- [3] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014.
- [4] Bastian Bohn, Jochen Garcke, Rodrigo Iza-Teran, Alexander Paprotny, Benjamin Peherstorfer, Ulf Schepsmeier, and Clemens August Thole. Analysis of car crash simulation data with nonlinear machine learning methods. *Procedia Computer Science*, 18:621–630, 2013.
- [5] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Michael M. Bronstein, and Daniel Cremers. Anisotropic diffusion descriptors. *Computer Graphics Forum*, 35(2):431–441, 2016.
- [6] Giorgos Bouritsas, Sergiy Bokhnyak, Stylianos Ploumpis, Stefanos Zafeiriou, and Michael Bronstein. Neural 3D morphable models: Spiral convolutional networks for 3D shape representation learning and generation. *Proceedings of the IEEE International Conference on Computer Vision*, 2019-Octob:7212–7221, 2019.
- [7] Michael M. Bronstein, Joan Bruna, Yann Lecun, Arthur Szlam, and Pierre Vandergheynst. Geometric Deep Learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.
- [8] Luitzen Egbertus Jan Brouwer. Über Abbildung von Mannigfaltigkeiten. *Mathematische Annalen*, 71(4), dec 1912.
- [9] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann Lecun. Spectral networks and deep locally connected networks on graphs. *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings*, pages 1–14, 2014.
- [10] Taco S. Cohen, Maurice Weiler, Berkay Kicanaoglu, and Max Welling. Gauge equivariant convolutional networks and the icosahedral CNN. *36th International Conference on Machine Learning*, 2019-June:2357–2371, 2019.
- [11] Pim de Haan, Maurice Weiler, Taco Cohen, and Max Welling. Gauge Equivariant Mesh CNNs: Anisotropic convolutions on geometric graphs. *arXiv preprint arXiv:2003.05425*, 2020.
- [12] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, volume 29, pages 3844–3852, 2016.
- [13] Michael Garland and Paul S Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, pages 209–216, 1997.
- [14] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. *34th International Conference on Machine Learning*, 3:2053–2070, 2017.
- [15] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep Sparse Rectifier Neural Networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315—323, 2011.
- [16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [17] Sara Hahner, Rodrigo Iza-Teran, and Jochen Garcke. Analysis and Prediction of Deforming 3D Shapes using Oriented Bounding Boxes and LSTM Autoencoders. In *Artificial Neural Networks and Machine Learning*, pages 284–296. Springer International Publishing, 2020.
- [18] Rana Hanocka, Amir Hertz, Noa Fish, Raja Giryes, Shachar Fleishman, and Daniel Cohen-Or. MeshCNN. *ACM Transactions on Graphics*, 38(4):1–12, jul 2019.
- [19] Emiel Hoogeboom, Jorn W.T. Peters, Taco S. Cohen, and Max Welling. HEXACONV. *arXiv preprint arXiv:1803.02108*, pages 1–11, 2018.
- [20] Shi-Min Hu, Zheng-Ning Liu, Meng-Hao Guo, Jun-Xiong Cai, Jiahui Huang, Tai-Jiang Mu, and Ralph R. Martin. Subdivision-Based Mesh Convolution Networks. *ACM Transactions on Graphics*, 1(1):1–15, 2021.
- [21] Dawar Khan, Alexander Plopski, Yuichiro Fujimoto, Masayuki Kanbara, Gul Jabeen, Yongjie Zhang, Xiaopeng Zhang, and Hirokazu Kato. Surface Remeshing: A Systematic Literature Review of Methods and Research Directions. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [22] Andrei Khodakovsky, Peter Schröder, and Wim Sweldens. Progressive geometry compression. *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, pages 271–278, 2000.
- [23] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, pages 1–15, 2015.
- [24] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*, pages 1–14, 2016.
- [25] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998.
- [26] Yann LeCun, Lionel D. Jackel, Brian Boser, John S. Denker, Henry P. Graf, Isabelle Guyon, Don Henderson, Richard E. Howard, and William Hubbard. Handwritten digit recognition: applications of neural network chips and automatic learning. *IEEE Communications Magazine*, 27(11):41–46, nov 1989.
- [27] Aaron W.F. Lee, Wim Sweldens, Peter Schröder, Lawrence Cowsar, and David Dobkin. MAPS: Multiresolution adaptive parameterization of surfaces. *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1998*, pages 95–104, 1998.
- [28] Stephane G. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions*

- on *Pattern Analysis and Machine Intelligence*, 11(7):674–693, jul 1989.
- [29] Jonathan Masci, Davide Boscaini, Michael M. Bronstein, and Pierre Vandergheynst. Geodesic Convolutional Neural Networks on Riemannian Manifolds. *Proceedings of the IEEE International Conference on Computer Vision*, 2015-Febru:832–840, 2015.
- [30] Alessio Micheli. Neural network for graphs: A contextual constructive approach. *IEEE Transactions on Neural Networks*, 20(3):498–511, 2009.
- [31] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodolà, Jan Svoboda, and Michael M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5115–5124, 2017.
- [32] National Crash Analysis Center (NCAC). Finite Element Model Archive (<http://web.archive.org/web/20160110143219/www.ncac.gwu.edu/vml/models.html>). Accessed on: 2016-01-10.
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32, pages 8026–8037, 2019.
- [34] Frédéric Payan, Céline Roudet, and Basile Sauvage. Semi-regular triangle remeshing: A comprehensive study. *Computer Graphics Forum*, 34(1):86–102, 2015.
- [35] Karl Pearson. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, nov 1901.
- [36] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J. Black. Generating 3D Faces Using Convolutional Mesh Autoencoders. *Proceedings of the European Conference on Computer Vision*, pages 704–720, 2018.
- [37] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3D Deep Learning with PyTorch3D. *arXiv preprint arXiv:2007.08501*, 2020.
- [38] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. DiffusionNet: Discretization Agnostic Learning on Surfaces. 2020.
- [39] Constantin Steppa and Tim L. Holch. HexagDLY—Processing hexagonally sampled data with CNNs in PyTorch. *SoftwareX*, 9:193–198, 2019.
- [40] Robert W. Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Transactions on Graphics*, 3:399–405, 2004.
- [41] Qingyang Tan, Lin Gao, Yu Kun Lai, and Shihong Xia. Variational Autoencoders for Deforming 3D Mesh Models. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 5841–5850, 2018.
- [42] Hugues Thomas, Charles R. Qi, Jean Emmanuel Deschaud, Beatriz Marcotegui, Francois Goulette, and Leonidas Guibas. KPConv: Flexible and deformable convolution for point clouds. *Proceedings of the IEEE International Conference on Computer Vision*, 2019-October:6410–6419, 2019.
- [43] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, aug 2008.
- [44] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics*, 38(5), 2019.
- [45] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. A Comprehensive Survey on Graph Neural Networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [46] Yu Jie Yuan, Yu Kun Lai, Jie Yang, Qi Duan, Hongbo Fu, and Lin Gao. Mesh variational autoencoders with edge contraction pooling. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, volume 2020-June, pages 1105–1112. IEEE Computer Society, jun 2020.