

Neural 3D Morphable Models: Spiral Convolutional Networks for 3D Shape Representation Learning and Generation

Giorgos Bouritsas^{*1} Sergiy Bokhnyak^{*2} Stylianos Ploumpis^{1,3}

Michael Bronstein^{1,2,4} Stefanos Zafeiriou^{1,3}

Imperial College London, UK¹ Universita Svizzera Italiana, Switzerland² FaceSoft.io³ Twitter⁴

¹{g.bouritsas18, s.ploumpis, m.bronstein, s.zafeiriou}@imperial.ac.uk ²bokhns@usi.ch

Abstract

Generative models for 3D geometric data arise in many important applications in 3D computer vision and graphics. In this paper, we focus on 3D deformable shapes that share a common topological structure, such as human faces and bodies. Morphable Models and their variants, despite their linear formulation, have been widely used for shape representation, while most of the recently proposed non-linear approaches resort to intermediate representations, such as 3D voxel grids or 2D views. In this work, we introduce a novel graph convolutional operator, acting directly on the 3D mesh, that explicitly models the inductive bias of the fixed underlying graph. This is achieved by enforcing consistent local orderings of the vertices of the graph, through the spiral operator, thus breaking the permutation invariance property that is adopted by all the prior work on Graph Neural Networks. Our operator comes by construction with desirable properties (anisotropic, topology-aware, lightweight, easy-to-optimize), and by using it as a building block for traditional deep generative architectures, we demonstrate state-of-the-art results on a variety of 3D shape datasets compared to the linear Morphable Model and other graph convolutional operators.

1. Introduction

The success of deep learning in computer vision and image analysis, speech recognition, and natural language processing, has driven the recent interest in developing similar models for 3D geometric data. Generalisations of successful architectures such as convolutional neural networks (CNNs) to data with non-Euclidean structure (e.g. manifolds and graphs) is known under the umbrella term *Geometric deep learning* [10]. In applications dealing with 3D data, the key challenge of geometric deep learning is a meaningful definition of intrinsic operations analogous to

convolution and pooling on meshes or point clouds. Among numerous advantages of working directly on mesh or point cloud data is the fact that it is possible to build invariance to shape transformations (both rigid and nonrigid) into the architecture, as a result allowing to use significantly simpler models and much less training data. So far, the main focus of research in the field of geometric deep learning has been on *analysis* tasks, encompassing shape classification and segmentation [35, 37], local descriptor learning, correspondence, and retrieval [31, 9, 27].

On the other hand, there has been limited progress in representation learning and generation of geometric data (*shape synthesis*). Obtaining descriptive and compact representations of meshes and point clouds is essential for downstream tasks such as classification and 3D reconstruction, when dealing with limited labelled training data. Additionally, geometric data synthesis is pivotal in applications such as 3D printing, computer graphics and animation, virtual reality, and game design, and can heavily assist graphics designers and speed-up production. Furthermore, given the high cost and time of acquiring quality 3D data, geometric generative models can be used as a cheap alternative for producing training data for geometric ML algorithms.

Most of the previous approaches in this direction rely on intermediate representations of 3D shapes, such as point clouds [1], voxels [44] or mappings to a flat domain [32, 4] instead of direct surface representations, such as meshes. Despite the success of such techniques, they either suffer from high computational complexity (e.g. voxels) or absence of smoothness of the data representation (e.g. point clouds), while usually pre- and post-processing steps are needed in order to obtain the output surface model. Learning directly on the mesh was only recently explored in [25, 38, 43, 22] for shape completion, non-linear facial morphable model construction and 3D reconstruction from single images, respectively.

In this paper, we propose a novel representation learning and generative framework for fixed topology meshes. For this purpose, we formulate an ordering-based graph convo-

^{*}Equal Contribution

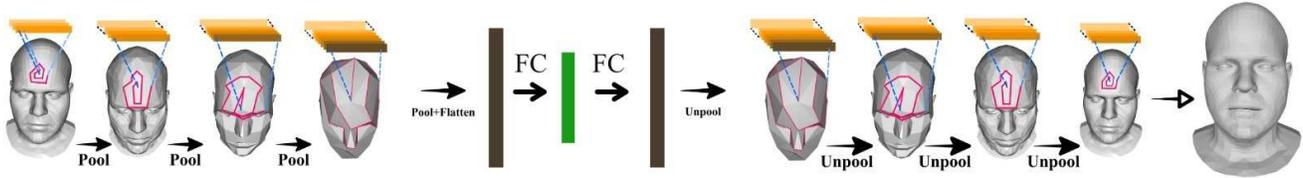


Figure 1: Illustration of our Neural3DMM architecture

lutional operator, contrary to the permutation invariant operators in the literature of Graph Neural Networks. In particular, similarly to image convolutions, for each vertex on the mesh, we enforce an explicit ordering of its neighbours, allowing a “1-1” mapping between the neighbours and the parameters of a learnable local filter. The order is obtained via a spiral scan, as proposed in [24], hence the name of the operator, *Spiral Convolution*. This way we obtain anisotropic filters without sacrificing computational complexity, while simultaneously we explicitly encode the fixed graph connectivity. The operator can potentially be generalised to other domains that accept implicit local orderings, such as arbitrary mesh topologies and point clouds, while it is naturally equivalent to traditional grid convolutions. Via this equivalence, common CNN practices, such as dilated convolutions, can be easily formulated for meshes.

We use spiral convolution as a basic building block for hierarchical intrinsic mesh autoencoders, which we coin *Neural 3D Morphable Models*. We quantitatively evaluate our methods on several popular datasets: human faces with different expressions (COMA [38]) and identities (Mein3D [7]) and human bodies with shape and pose variation (DFAUST [6]). Our model achieves state-of-the-art reconstruction results, outperforming the widely used linear 3D Morphable Model [5] and the COMA autoencoder [38], as well other graph convolutional operators, including the initial formulation of the spiral operator [24]. We also qualitatively assess our framework showing ‘shape arithmetic’ in the latent space of the autoencoder and by synthesising facial identities via a spiral convolution Wasserstein GAN.

2. Related Work

Generative models for arbitrary shapes: Perhaps the most common approaches for generating arbitrary shapes are **volumetric CNNs** [45, 36, 28] acting on 3D voxels. For example, voxel regression from images [19], denoising autoencoders [40] and voxel-GANs [44] have been proposed. Among the key drawbacks of volumetric methods are their inherent high computational complexity and that they yield coarse and redundant representations. **Point clouds** are a simple and lightweight alternative to volumetric representation recently gaining popularity. Several methods have been proposed for representation learning of fixed-size point clouds [1] using the PointNet [35] architec-

ture. In [46], point clouds of arbitrary size can be synthesised via a 2D grid deformation. Despite their compactness, point clouds are not popular for realistic and high-quality 3D geometry generation due to their lack of an underlying smooth structure. **Image-based** methods have also been proposed, such as multi-view [3] and flat domain mappings such as UV maps [32, 4], however they are computationally demanding, require pre- and post-processing steps and usually produce undesirable artefacts. It is also worth mentioning the recently introduced **implicit-surface** based approaches [29, 13, 33], that can yield accurate results, though with the disadvantage of slow inference (dense sampling of the 3D space followed by marching cubes).

Morphable models: In the case of deformable shapes, such as faces, bodies, hands *etc.*, where a fixed topology can be obtained by establishing dense correspondences with a template, the most popular methods are still statistical models given their simplicity. For **Faces**, the baseline is the PCA-based 3D Morphable Model (3DMM) [5]. The Large Scale Face Model (LSFM) [7] was proposed for facial identity and made publicly available, [12, 23] were proposed for facial expression, while for the entire head a large scale model was proposed in [34]. For **Body & Hand**, the most well known models are the skinned vertex-based models SMPL [26] and MANO [39], respectively. SMPL and MANO are non-linear and require (a) joint localisation and (b) solving special optimisation problems in order to project a new shape to the space of the models. In this paper, we take a different approach introducing a new family of differentiable Morphable Models, which can be applied on a variety of objects, with strong (*i.e.* body) and less strong (*i.e.* face) articulations. Our methods have better representational power and also do not require any additional supervision.

Geometric Deep Learning is a set of recent methods trying to generalise neural networks to non-Euclidean domains such as graphs and manifolds [10]. Such methods have achieved promising results in geometry processing and computer graphics [27, 9], computational chemistry [16, 18], and network science [21, 31]. Multiple approaches have been proposed to construct convolution-like operations, including spectral methods [11, 15, 21, 47], local charting based [27, 9, 31, 17, 24] and soft attention [41, 42]. Finally, graph or mesh coarsening techniques [15, 48] have been proposed, equivalent to image pooling.

3. Spiral Convolutional Networks

3.1. Spiral Convolution

For the following discussion, we assume to be given a manifold, discretised as a triangular mesh $\mathcal{M} = (\mathcal{V}, \mathcal{E}, \mathcal{F})$ where $\mathcal{V} = \{1, \dots, n\}$, \mathcal{E} , and \mathcal{F} denote the sets of vertices, edges, and faces respectively. Furthermore, let $f : \mathcal{V} \rightarrow \mathbb{R}$, a function representing the vertex features.

One of the key challenges in developing convolution-like operators on graphs or manifolds is the lack of a global system of coordinates that can be associated with each point. The first intrinsic mesh convolutional architectures such as GCNN [27], ACNN [9] or MoNet [31] overcame this problem by constructing a *local* system of coordinates $\mathbf{u}(x, y)$ around each vertex x of the mesh, in which a set of local weighting functions w_1, \dots, w_L is applied to aggregate information from the vertices y of the neighborhood $\mathcal{N}(x)$. This allows to define ‘patch operators’ generalising the sliding window filtering in images:

$$(f \star g)_x = \sum_{\ell=1}^L g_\ell \sum_{y \in \mathcal{N}(x)} w_\ell(\mathbf{u}(x, y)) f(y) \quad (1)$$

where $\sum_{y \in \mathcal{N}(x)} w_\ell(\mathbf{u}(x, y)) f(y)$ are ‘soft pixels’ (L in total), f are akin to pixel intensity in images, and g_ℓ the filter weights. The problem of the absence of a global coordinate system is equivalent to the absence of canonical ordering of the vertices, and the patch-operator based approaches can be also interpreted as attention mechanisms, as in [41] and [42]. In particular, the absence of ordering does not allow the construction of a “1-1” mapping between neighbouring features $f(y)$ and filter weights g_ℓ , thus a “all-to-all” mapping is performed via learnable *soft-attention* weights $w_\ell(\mathbf{u}(x, y))$. In the Euclidean setting, such operators boil down to the classical convolution, since an ordering can be obtained via the global coordinate system.

Besides the lack of a global coordinate system, another motivation for patch-operator based approaches when working on meshes, is the need for *insensitivity to meshing* of the continuous surface, *i.e.* ideally, each patch operator should be independent of the underlying graph topology.

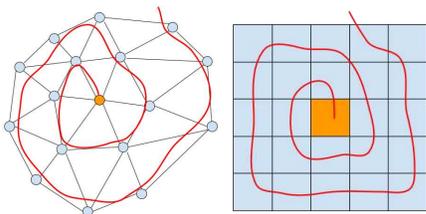


Figure 2: Spiral ordering on a mesh and an image patch

However, all the methods falling into this family, come

at the cost of high computational complexity and parameter count and can be hard to optimise. Moreover, patch-operator based methods specifically designed for meshes, require hand-crafting and pre-computing the local systems of coordinates. To this end, in this paper we make a crucial observation in order to overcome the disadvantages of the aforementioned approaches: the issues of the absence of a global ordering and insensitivity to graph topology are irrelevant when dealing with *fixed topology* meshes. In particular, one can locally order the vertices and keep the order fixed. Then, graph convolution can be defined as follows:

$$(f \star g)_x = \sum_{\ell=1}^L g_\ell f(x_\ell). \quad (2)$$

where $\{x_1, \dots, x_L\}$ denote the neighbours of vertex x ordered in a fixed way. Here, in analogy with the patch operators, each patch operator is a single neighbouring vertex.

In the Euclidean setting, the order is simply a raster scan of pixels in a patch. On meshes, we opt for a simple and intuitive ordering using spiral trajectories inspired by [24]. Let $x \in \mathcal{V}$ be a mesh vertex, and let $R^d(x)$ be the *d*-ring, *i.e.* an ordered set of vertices whose shortest (graph) path to x is exactly d hops long; $R_j^d(x)$ denotes the j th element in the d -ring (trivially, $R_1^0(x) = x$). We define the *spiral patch operator* as the ordered sequence

$$S(x) = \{x, R_1^1(x), R_2^1(x), \dots, R_{|R^h|}^h(x)\}, \quad (3)$$

where h denotes the patch radius, similar to the size of the kernel in classical CNNs. Then, *spiral convolution* is:

$$(f \star g)_x = \sum_{\ell=1}^L g_\ell f(S_\ell(x)). \quad (4)$$

The uniqueness of the ordering is given by fixing two degrees of freedom: the direction of the rings and the first vertex $R_1^1(x)$. The rest of the vertices of the spiral are ordered inductively. The direction is chosen by moving clockwise or counterclockwise, while the choice of the first vertex, the *reference point*, is based on the underlying geometry of the shape to ensure the robustness of the method. In particular, we fix a reference vertex x_0 on a template shape and choose the initial point for each spiral to be in the direction of the shortest geodesic path to x_0 , *i.e.*

$$R_1^1(x) = \arg \min_{y \in R^1(x)} d_{\mathcal{M}}(x_0, y), \quad (5)$$

where $d_{\mathcal{M}}$ is the geodesic distance between two vertices on the mesh \mathcal{M} . In order to allow for fixed-sized spirals, we choose a fixed length L as a hyper-parameter and then either truncate or zero-pad each spiral depending on its size.

Comparison to Lim *et al.* [24]: The authors choose the starting point of each spiral at random, for every mesh sample, every vertex, and every epoch during training. This

choice prevents us from explicitly encoding the fixed connectivity, since corresponding vertices in different meshes will not undergo the same transformation (as in image convolutions). Moreover, single vertices also undergo different transformations every time a new spiral is sampled. Thus, in order for the network to obtain robustness to different spiral samples, it inevitably has to become invariant to different rotations of the neighbourhoods, thus it has reduced capacity. To this end, we emphasise the need of consistent orderings across different meshes.

Moreover, in [24], the authors model the vertices on the spiral via a recurrent network, which has higher computational complexity, is harder to optimise and does not take advantage of the stationary properties of the 3D shape (local statistics are repeated across different patches), which are treated by our spiral kernel with weight sharing.

Comparison to spectral filters: Spectral convolutional operators developed in [15, 21] for graphs and used in [38] for mesh autoencoders, suffer from the fact that they are inherently *isotropic*. This is a side-effect when one, under the absence of a canonical ordering, needs to design a permutation-invariant operator with small number of parameters. In particular, spectral filters rely on the *Laplacian operator*, which performs a weighted averaging of the neighbour vertices :

$$(\Delta f)_x = \sum_{y:(x,y) \in \mathcal{E}} w_{xy} (f(y) - f(x)), \quad (6)$$

where w_{xy} denotes an edge weight. A polynomial of degree r with learnable coefficients $\theta_0, \dots, \theta_r$ is then applied to Δ . Then, the graph convolution amounts to filtering the Laplacian eigenvalues, $p(\Delta) = \Phi p(\Lambda) \Phi^\top$. Equivalently:

$$(f * g) = p(\Delta) f = \sum_{\ell=0}^r \theta_\ell \Delta^\ell f, \quad (7)$$

While a necessary evil in general graphs, spectral filters on meshes are rather weak given that they are locally rotationally-invariant. On the other hand, spiral convolutional filters leverage the fact that on a mesh one can canonically order the neighbours. Thus, they are anisotropic by construction and as will be shown in the experimental section 4 they are expressive by using just one-hop neighbourhoods, contrary to the large receptive fields used in [38]. In Fig 3 we visualise the impulse response (centred on a vertex on the forehead) of a selected laplacian polynomial filter from the architecture of [38] (left) and from a spiral convolutional filter with $h = 1$ (right).

Finally, the equivalence of spiral convolutions to image convolutions allows the use of long-studied practices in the computer vision community. For example, small patches can be used, leading to few parameters and fast computation. Furthermore, dilated convolutions [49] can also be adapted to the spiral operator by simply sub-sampling the

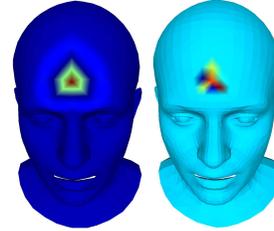


Figure 3: Activations of ChebNet vs spiral convolutions

spiral. Finally, we argue here that our operator could be applied to other domains, such as point clouds, where an ordering of the data points can be enforced.

3.2. Neural 3D Morphable Models

Let $F = [f_0 | f_1 | \dots | f_N]$, $f_i \in \mathbb{R}^{d \times m}$ the matrix of all the signals defined on a set of meshes in dense correspondence that are sampled from a distribution \mathcal{D} , where d the dimensionality of the signal on the mesh (vertex position, texture etc.) and m the number of vertices. A linear 3D Morphable Model [5] represents arbitrary instances $y \in \mathcal{D}$ as a linear combination of the k largest eigenvectors of the covariance matrix of F by making a gaussianity assumption:

$$y \approx \bar{f} + \sum_i^k \alpha_i \sqrt{d_i} \mathbf{v}_i \quad (8)$$

where \bar{f} the mean shape, \mathbf{v}_i is the i th principal component, d_i the respective eigenvalue and α_i the linear weight coefficient. Given its linear formulation, the representational power of the 3DMM is constrained by the span of the eigenvectors, while its parameters scale linearly w.r.t the number of the eigencomponents used, leading to large parametrisations for meshes of high resolution.

In contrast, in this paper, we use spiral convolutions as a building block to build a fully differentiable non-linear Morphable Model. In essence, a Neural 3D Morphable Model is a deep convolutional mesh autoencoder, that learns hierarchical representations of a shape. An illustration of the architecture can be found in Fig 1. Leveraging the connectivity of the graph with spiral convolutional filters, we allow for local processing of each shape, while the hierarchical nature of the model allows learning in multiple scales. This way we manage to learn semantically meaningful representations and considerably reduce the number of parameters. Furthermore, we bypass the need to make assumptions about the distribution of the data.

Similar to traditional convolutional autoencoders, we make use of series of convolutional layers with small receptive fields followed by pooling and unpooling, for the encoder and the decoder respectively, where a decimated or upsampled version of the mesh is obtained each time and

the features of the existing vertices are either aggregated or extrapolated. We follow [38] for the calculation of the features of the added vertices after upsampling, *i.e.* through interpolation by weighting the nearby vertices with barycentric coordinates. The network is trained by minimising the L_1 norm between the input and the predicted output.

3.3. Spiral Convolutional GAN

In order to improve the synthesis of meshes of high resolution, thus increased detail, we extend our framework with a distribution matching scheme. In particular, we propose a mesh Wasserstein GAN [2] with gradient penalty to enforce the Lipschitz constraint [20], that is trained to minimise the wasserstein divergence between the real distribution of the meshes and the distribution of those produced by the generator network. The generator and discriminator architectures, have the same structure as the decoder and the encoder of the Neural3DMM respectively. Via this framework, we obtain two additional properties that are inherently absent from the autoencoder: high frequency detail and a straightforward way to sample from the latent space.

4. Evaluation

In this section, we showcase the effectiveness of our proposed method on a variety of shape datasets. We conduct a series of ablation studies in order to compare our operator to other Graph Neural Networks, by using the same autoencoder architecture. First, we demonstrate the inherent higher capacity of spiral convolutions compared to ChebNet (spectral). Moreover, we discuss the advantages of our method compared to soft-attention based Graph Neural Networks, such as patch-operator based. Finally, we show the importance of the consistency of the ordering by comparing our method to different variants of the method proposed in [24].

Furthermore, we quantitatively show that our method can yield better representations than the linear 3DMM and COMA, while maintaining a small parameter count and frequently allowing a more compact latent representation. Moreover, we proceed with a qualitative evaluation of our method by generating novel examples through vector space arithmetic. Finally, we assess our intrinsic GAN in terms of its ability to produce high resolution realistic examples.

For all the cases, we choose as signal on the mesh the normalised deformations from the mean shape, *i.e.* for every vertex we subtract its mean position and divide with the standard deviation. In this way, we encourage signal stationarity, thus facilitating optimisation. The code is available at <https://github.com/gbouritsas/neural3dmm>.

4.1. Datasets

COMA. The facial expression dataset from Ranjan *et al.* [38], consisting of 20K+ 3D scans (5023 vertices) of twelve

unique identities performing twelve types of extreme facial expressions. We used the same data split as in [38].

DFAUST. The dynamic human body shape dataset from Bogo *et al.* [6], consisting of 40K+ 3D scans (6890 vertices) of ten unique identities performing actions such as leg and arm raises, jumps, etc. We randomly split the data into a test set of 5000, 500 validation, and 34,5K+ train.

MeIn3D. The 3D large scale facial identity dataset from Booth *et al.* [8], consisting of more than 10,000 distinct identity scans with 28K vertices which cover a wide range of gender ethnicity and age. For the subsequent experiments, the MeIn3D dataset was randomly split within demographic constraints to ensure gender, ethnic and age diversity, into 9K train and 1K test meshes.

For the quantitative experiments of sections 4.3 and 4.4 the evaluation metric used is **generalisation**, which measures the ability of a model to represent novel shapes from the same distribution as it was trained on. More specifically we evaluate the average per sample and per vertex euclidean distance in the 3D space (in millimetres) between corresponding vertices in the input and its reconstruction.

4.2. Implementation Details

We denote as $SC(h, w)$ a spiral convolution of h hops and w filters, $DS(p)$ and $US(p)$ a downsampling and an upsampling by a factor of p , respectively, $FC(d)$ a fully connected layer, l the number of vertices after the last downsampling layer. The simple Neural3DMM for COMA and DFAUST datasets is the following:

$$\begin{aligned} Enc : SC(1, 16) &\rightarrow DS(4) \rightarrow SC(1, 16) \rightarrow DS(4) \rightarrow SC(1, 16) \rightarrow DS(4) \rightarrow SC(1, 32) \rightarrow DS(4) \rightarrow FC(d) \\ Dec : FC(l * 32) &\rightarrow US(4) \rightarrow SC(1, 32) \rightarrow US(4) \rightarrow SC(1, 16) \rightarrow US(4) \rightarrow SC(1, 16) \rightarrow US(4) \rightarrow SC(1, 3) \end{aligned}$$

For Mein3D, due to the high vertex count, we modified the COMA architecture for our simple Neural3DMM by adding an extra convolution and an extra downsampling/upsampling layer in the encoder and the decoder respectively (encoder filter sizes: [8,16,16,32,32], decoder: mirror of the encoder). The larger Neural3DMM follows the above architecture, but with an increased parameter space. For COMA, the convolutional filters of the encoder had sizes [64,64,64,128] and for Mein3D the sizes were [8,16,32,64,128], while the decoder is a mirror of the encoder. For DFAUST, the sizes were [16,32,64,128] and [128,64,32,32,16] and dilated convolutions with $h = 2$ hops and dilation ratio $r = 2$ were used for the first and the last two layers of the encoder and the decoder respectively. We observed that by adding an additional convolution at the very end (of size equal to the size of the input feature space), training was accelerated. All of our activation functions were ELUs [14]. Our learning rate was 10^{-3} with a decay of 0.99 after each epoch, and our weight decay was

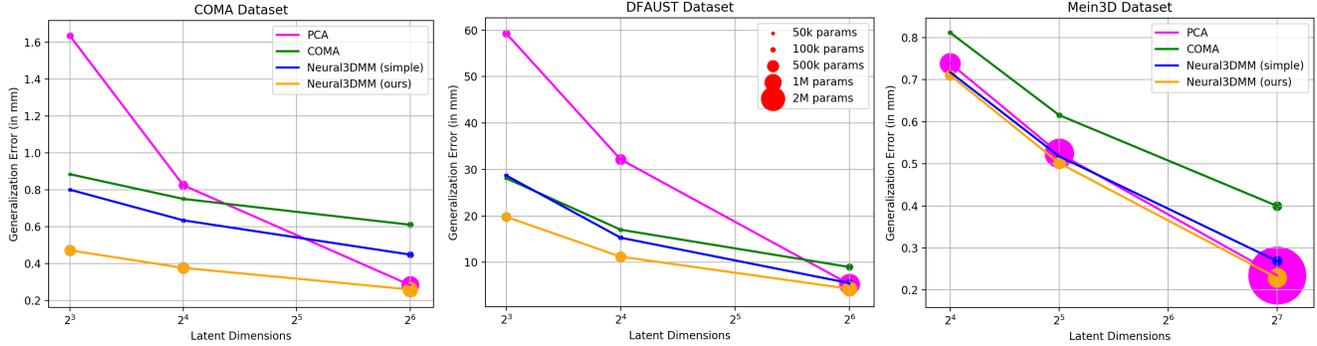


Figure 4: Quantitative evaluation of our Neural3DMM against the baselines, in terms of generalisation and # of parameters

5×10^{-5} . All models were trained for 300 epochs.

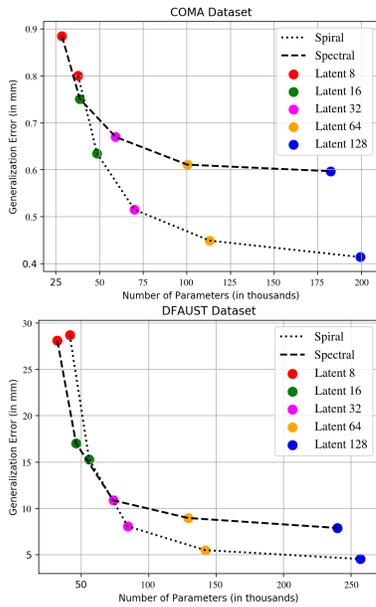


Figure 5: Spiral vs ChebNet (spectral) filters

4.3. Ablation Studies

4.3.1 Isotropic vs Anisotropic Convolutions

For the purposes of this experiment we used the architecture deployed by the authors of [38]. The number of parameters in our case is slightly larger due to the fact that the immediate neighbours, that affect the size of the spiral, range from 7 to 10, while the polynomials used in [38] go up to the 6th power of the Laplacian. For both datasets, as clearly illustrated in Fig 5, spiral convolution-based autoencoders consistently outperform the spectral ones for every latent dimension, in accordance with the analysis made in section 3.1. Additionally, increasing the latent dimensions, our model’s performance increases at a higher rate than its counterpart. Notice that the number of parameters scales

the same way as the latent size grows, but the spiral model makes better use of the added parameters especially looking at dimensions 16, 32, 64, and 128. Especially on the COMA dataset, the spectral model seems to be flattening between 64 and 128 while the spiral is still noticeably decreasing.

4.3.2 Spiral vs Attention based Convolutions

In this experiment we compare our method with certain state-of-the-art soft-attention based Graph Neural Networks: **MoNet**: the patch-operator based model of [31], where the attention weights are the learnable parameters of gaussian kernels defined on a pseudo-coordinate space¹, **FeastNet** [42] and **Graph Attention** [41], where the attention weights are learnable functions of the input features.

In table 1, we provide results on COMA dataset, using the simple Neural3DMM architecture with latent size 16. We choose the number of attention heads (gaussian kernels in [31]) to be either 9 (equal to the size of the spiral in our method, for a fair comparison) or 25 (as in [31], to showcase the effect of over-parametrisation). When it comes to similar number of parameters our method manages to outperform its counterparts, while compared to over-parametrised soft attention networks it either outperforms them, or achieves slightly worse performance. This shows that the spiral operator can make more efficient use of the available learnable parameters, thus being a lightweight alternative to attention-based methods without sacrificing performance. Also, its formulation allows for fast computation; in table 1 we measure per mesh inference time in ms (on a GeForce RTX 2080 Ti GPU).

4.3.3 Comparison to Lim et al. [24]

In order to showcase how the operator behaves when the ordering is not consistent, we perform experiments under four scenarios: the original formulation of [24], where

¹here we display the best obtained results when choosing the pseudo-coordinates to be local cartesian.

kernels	GAT		FeastNet		MoNet		Ours
	9	25	9	25	9	25	-
error	0,762	0,732	0,750	0,623	0,708	0,583	0,635
params	50K	101K	49K	98K	48K	95K	48K
time	12,77	15,37	9,04	9,66	10,55	10,96	8,18

Table 1: Spirals vs soft-attention operators

each spiral is randomly oriented for every mesh and every epoch (*rand mesh & epoch*); choosing the same orientation across all the meshes randomly at every epoch (*rand epoch*); choosing different orientations for every mesh, but keeping them fixed across epochs (*rand mesh*); and fixed ordering (Ours). We compare the LSTM-based approach of [24] and our linear projection formulation (Eq (2)). The experimental setting and architecture is the same as in the previous section. The proposed approach achieves over 28% improved performance compared to [24], which substantiates the benefits of passing corresponding points through the same transformations.

operation	rand mesh & epoch	rand mesh	rand epoch	fixed ordering
LSTM	0.888 [24]	0.880	0.996	0.792
lin. proj.	0.829	0.825	0.951	0.635 (Ours)

Table 2: Importance of the ordering consistency

4.4. Neural 3D Morphable models

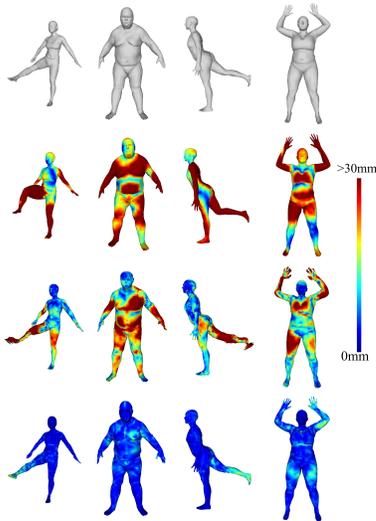


Figure 6: Colour coding of the per vertex euclidean error of the reconstructions produced by PCA (2nd), COMA (3rd), and our Neural3DMM (bottom). Top row is ground truth.

4.4.1 Quantitative results

In this section, we compare the following methods for different dimensions of the latent space: **PCA**, the 3D Mor-

phable Model [5], **COMA**, the ChebNet-based Mesh Autoencoder, **Neural3DMM (small)**, ours spiral convolution autoencoder with the same architecture as in COMA, **Neural3DMM (ours)**, our proposed Neural3DMM framework, where we enhanced our model with a larger parameter space (see Sec. 4.2). The latent sizes were chosen based on the variance explained by PCA (explained variance of roughly 85%, 95% and 99% of the total variance).

As can be seen from the graphs in Fig 4, our Neural3DMM achieves smaller generalisation errors in every case it was tested on. For the COMA and DFAUST datasets all hierarchical intrinsic architectures outperform PCA for small latent sizes. That should probably be attributed to the fact that the localised filters used allow for effective reconstruction of smaller patches of the shape, such as arms and legs (for the DFAUST case), whilst PCA attempts a more global reconstruction, thus its error is distributed equally across the entire shape. This is well shown in Fig 6, where we compare exemplar reconstructions of samples from the test set (latent size 16). It is clearly visible that PCA prioritises body shape over pose resulting to body parts in the wrong locations (for example see the right leg of the woman on the leftmost column). On the contrary COMA places the vertices in approximately correct locations, but struggles to recover the fine details of the shape leading to various artefacts and deformities; our model on the other hand seemingly balances these two difficult tasks resulting in quality reconstructions that preserve pose and shape.

Comparing to [38], it is again apparent here that our spiral-based autoencoder has increased capacity, which together with the increased parameter space, makes our larger Neural3DMM outperform the other methods by a considerably large margin in terms of both generalisation and compression. Despite the fact that for higher dimensions, PCA can explain more than 99% of the total variance, thus making it a tough-to-beat baseline, our larger model still manages to outperform it. The main advantage here is the substantially smaller number of parameters of which we make use. This is clearly seen in the comparison for the MeIn3D dataset, where the large vertex count makes non-local methods as PCA impractical. It is necessary to mention here, that larger latent space sizes are not necessarily desirable for an autoencoder because they might lead to less semantically meaningful and discriminative representation for downstream tasks.

4.4.2 Qualitative results

Here, we assess the representational power of our models by the common practice of testing their ability to perform linear algebra in their latent spaces.

Interpolation Fig 7: We choose two sufficiently different samples x_1 and x_2 from our test set, encode them in their

latent representations \mathbf{z}_1 and \mathbf{z}_2 and then produce intermediate encodings by sampling the line that connects them *i.e.* $\mathbf{z} = a\mathbf{z}_1 + (1 - a)\mathbf{z}_2$, where $a \in (0, 1)$.



Figure 7: Interpolations between expressions and identities

Extrapolation Fig 8: Similarly, we decode latent representations that reside on the line defined by \mathbf{z}_1 and \mathbf{z}_2 , but outside the respective line segment, *i.e.* $\mathbf{z} = a*\mathbf{z}_1 + (1-a)*\mathbf{z}_2$, where $a \in (-\infty, 0) \cup (1, +\infty)$. We choose \mathbf{z}_1 to be our neutral expression for COMA and neutral pose for DFAUST, in order to showcase the exaggeration of a specific characteristic on the shape.

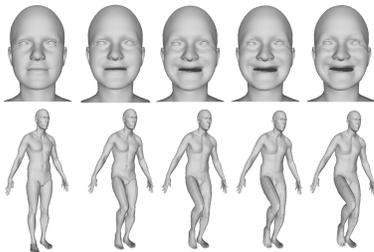


Figure 8: Extrapolation. Left: neutral expression/pose

Shape Analogies Fig 9: We choose three meshes A , B , C , and construct a D such that it satisfies $A:B::C:D$ using linear algebra in the latent space as in [30]: $e(B) - e(A) = e(D) - e(C)$ ($e(*)$ the encoding), where we then solve for $e(D)$ and decode it. This way we transfer a specific characteristic using meshes from our dataset.

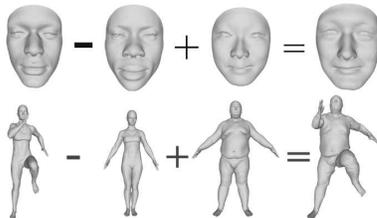


Figure 9: Analogies in MeIn3D and DFAUST

4.5. GAN evaluation

In figure 10, we sampled several faces from the latent distribution of the trained generator. Notice that they are realistically looking and, following the statistics of the

dataset, span a large proportion of the real distribution of the human faces, in terms of ethnicity, gender and age. Compared to the most popular approach for synthesizing faces, *i.e.* the 3DMM, our model learns to produce fine details on the facial structure, making them hard to distinguish from real 3D scans, whereas the 3DMM, although it produces smooth surfaces, frequently makes it easy to tell the difference between real and artificially produced samples. We direct the reader to the supplementary material to compare with samples drawn from the 3DMM’s latent space.



Figure 10: Generated identities from our intrinsic 3D GAN

5. Conclusion

In this paper we introduced a representation learning and generative framework for fixed topology 3D deformable shapes, by using a mesh convolutional operator, spiral convolutions, that efficiently encodes the inductive bias of the fixed topology. We showcased the inherent representational power of the operator, as well as its reduced computational complexity, compared to prior work on graph convolutional operators and show that our mesh autoencoder achieves state-of-the-art results in mesh reconstruction. Finally, we present the generation capabilities of our models through vector space arithmetic, as well as by synthesising novel facial identities. Regarding future work, we plan to extend our framework to general graphs and 3D shapes of arbitrary topology, as well as to other domains that have capacity for an implicit ordering of their primitives, such as point clouds.

6. Acknowledgements

This research was partially supported by ERC Consolidator Grant No. 724228 (LEMAN), Google Research Faculty awards, and the Royal Society Wolfson Research Merit award. G. Bouritsas was funded by the Imperial College London, Department of Computing, PhD scholarship. Dr. Zafeiriou acknowledges support from a Google Faculty award and EPSRC fellowship Deform (EP/S010203/1). S. Ploumpis was supported by EPSRC Project (EP/N007743/1) FACER2VM.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. *International Conference on Machine Learning (ICML)*, 2018. [1](#), [2](#)
- [2] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017. [5](#)
- [3] Amir Arsalan Soltani, Haibin Huang, Jiajun Wu, Tejas D Kulkarni, and Joshua B Tenenbaum. Synthesizing 3d shapes via modeling multi-view depth maps and silhouettes with deep generative networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#)
- [4] Heli Ben-Hamu, Haggai Maron, Itay Kezurer, Gal Avineri, and Yaron Lipman. Multi-chart generative surface modeling. In *SIGGRAPH Asia 2018 Technical Papers*, page 215. ACM, 2018. [1](#), [2](#)
- [5] Volker Blanz, Thomas Vetter, et al. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, 1999. [2](#), [4](#), [7](#)
- [6] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#), [5](#)
- [7] James Booth, Anastasios Roussos, Allan Ponniah, David Dunaway, and Stefanos Zafeiriou. Large scale 3d morphable models. *International Journal of Computer Vision (IJCV)*, 2018. [2](#)
- [8] James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. A 3d morphable model learnt from 10,000 faces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [5](#)
- [9] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2016. [1](#), [2](#), [3](#)
- [10] 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. [1](#), [2](#)
- [11] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. In *International Conference on Learning Representations (ICLR)*, 2014. [2](#)
- [12] Chen Cao, Yanlin Weng, Shun Zhou, Yiying Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE Transactions on Visualization and Computer Graphics*, 2014. [2](#)
- [13] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. *CVPR*, 2019. [2](#)
- [14] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *CoRR*, 2015. [5](#)
- [15] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. *Advances in Neural Information Processing Systems (NIPS)*, 2016. [2](#), [4](#)
- [16] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. [2](#)
- [17] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous b-spline kernels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#)
- [18] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017. [2](#)
- [19] Rohit Girdhar, David F Fouhey, Mikel Rodriguez, and Abhinav Gupta. Learning a predictable and generative vector representation for objects. In *European Conference on Computer Vision (ECCV)*. Springer, 2016. [2](#)
- [20] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems (NIPS)*, 2017. [5](#)
- [21] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017. [2](#), [4](#)
- [22] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019. [1](#)
- [23] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 2017. [2](#)
- [24] Isaak Lim, Alexander Dielen, Marcel Campen, and Leif Kobbelt. A simple approach to intrinsic correspondence learning on unstructured 3d meshes. *Proceedings of the European Conference on Computer Vision Workshops (ECCVW)*, 2018. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [25] Or Litany, Alex Bronstein, Michael Bronstein, and Ameesh Makadia. Deformable shape completion with graph convolutional autoencoders. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1886–1895, 2018. [1](#)
- [26] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, Oct. 2015. [2](#)
- [27] Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 37–45, 2015. [1](#), [2](#), [3](#)
- [28] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition.

- In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015. 2
- [29] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. *CVPR*, 2019. 2
- [30] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NIPS)*. 2013. 8
- [31] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, 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 (CVPR)*, 2017. 1, 2, 3, 6
- [32] Stylianos Moschoglou, Stylianos Ploumpis, Mihalis Nicolaou, and Stefanos Zafeiriou. 3dfacegan: Adversarial nets for 3d face representation, generation, and translation. *arXiv preprint arXiv:1905.00307*, 2019. 1, 2
- [33] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. *CVPR*, 2019. 2
- [34] Stylianos Ploumpis, Haoyang Wang, Nick Pears, William AP Smith, and Stefanos Zafeiriou. Combining 3d morphable models: A large scale face-and-head model. *CVPR*, 2019. 2
- [35] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2
- [36] Charles R Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems (NIPS)*, pages 5099–5108, 2017. 1
- [38] 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 (ECCV)*, 2018. 1, 2, 4, 5, 6, 7
- [39] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 2017. 2
- [40] Abhishek Sharma, Oliver Grau, and Mario Fritz. Vconv-dae: Deep volumetric shape learning without object labels. In *European Conference on Computer Vision*, pages 236–250. Springer, 2016. 2
- [41] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018. 2, 3, 6
- [42] Nitika Verma, Edmond Boyer, and Jakob Verbeek. Featnet: Feature-steered graph convolutions for 3d shape analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2598–2606, 2018. 2, 3, 6
- [43] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *ECCV*, 2018. 1
- [44] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in Neural Information Processing Systems (NIPS)*, 2016. 1, 2
- [45] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [46] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2
- [47] Li Yi, Hao Su, Xingwen Guo, and Leonidas J Guibas. Syncspecnn: Synchronized spectral cnn for 3d shape segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2
- [48] Zhitao Ying, Jiaxuan You, Christopher Morris, Xiang Ren, Will Hamilton, and Jure Leskovec. Hierarchical graph representation learning with differentiable pooling. In *Advances in Neural Information Processing Systems (NIPS)*, 2018. 2
- [49] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *International Conference on Learning Representations (ICLR)*, 2016. 4