

# PFCNN: Convolutional Neural Networks on 3D Surfaces Using Parallel Frames

Yuqi Yang<sup>\*1,3</sup> Shilin Liu<sup>\*2,3</sup> Hao Pan<sup>†3</sup> Yang Liu<sup>3</sup> Xin Tong<sup>3</sup>

<sup>1</sup>Tsinghua University <sup>2</sup>University of Science and Technology of China <sup>3</sup>Microsoft Research Asia  
 yangyq18@mails.tsinghua.edu.cn freelin@mail.ustc.edu.cn  
 {haopan, yangliu, xtong}@microsoft.com

## Abstract

*Surface meshes are widely used shape representations and capture finer geometry data than point clouds or volumetric grids, but are challenging to apply CNNs directly due to their non-Euclidean structure. We use parallel frames on surface to define PFCNNs that enable effective feature learning on surface meshes by mimicking standard convolutions faithfully. In particular, the convolution of PFCNN not only maps local surface patches onto flat tangent planes, but also aligns the tangent planes such that they locally form a flat Euclidean structure, thus enabling recovery of standard convolutions. The alignment is achieved by the tool of locally flat connections borrowed from discrete differential geometry, which can be efficiently encoded and computed by parallel frame fields. In addition, the lack of canonical axis on surface is handled by sampling with the frame directions. Experiments show that for tasks including classification, segmentation and registration on deformable geometric domains, as well as semantic scene segmentation on rigid domains, PFCNNs achieve robust and superior performances without using sophisticated input features than state-of-the-art surface based CNNs.*

## 1. Introduction

Applying CNNs to 3D geometric domains is critical for deep learning beyond the 2D images. Unlike regular 2D images, 3D geometric data can be represented in different forms, posing challenges to standard CNNs. For example, volumetric grids regularly sample  $\mathbf{R}^3$  on which CNNs can be trivially deployed, but they are memory consuming and inflexible for capturing fine geometric details. For representation efficiency, 3D objects and scenes are frequently encoded by their boundary surfaces discretized as triangle meshes. However, the curved and irregularly sampled meshes do not admit the standard CNNs designed for flat image domains with regular pixel grids. While several surface based CNNs have been proposed to tackle this problem, in this paper we use parallel frame fields that con-

tain pointwise  $N$ -direction frames (Fig. 1) to define a novel PFCNN framework whose convolution mimics standard image convolutions more faithfully.

Similar to standard CNNs, the PFCNN convolution works on a local surface patch each time and maps it onto the flat tangent space where the convolution kernel is parameterized, as done by many previous surface based CNNs [28, 3, 30, 32]. Different from the previous approaches, however, we also align the tangent spaces of different surface points such that they locally form a flat Euclidean structure, on which the surface-based feature maps and convolution kernels can be moved as in the standard image domain. For images, such translation operations are formally captured by the translation equivariance property of convolution [4], which is a key factor contributing to the effectiveness of CNNs by enabling shared trainable weights and thus significantly reducing the amount of network parameters to avoid overfitting and achieve generalization [23, 13]; our surface based convolution is shown to reproduce the image-domain translation equivariance locally.

We adopt the tool of locally flat connections from discrete differential geometry [45] to align the tangent spaces. The locally flat connection is encoded by the field of pointwise tangential  $N$ -direction frames (Fig. 1) that is efficiently computed to be parallel and aligned to salient geometric features to better capture semantics. In addition, because there exists no canonical axis on a surface, we sample the axes using the same  $N$  frame directions and organize the resulting feature maps with an  $N$ -cover space of the domain surface [9]; on each sheet of the cover space, the canonical axis is selected and the convolution is readily defined. Furthermore, to handle the irregular mesh vertices, for each patch we resample with a regular grid and apply standard shaped convolution kernels on it.

The PFCNNs resemble standard CNNs so that efficient network structures can be leveraged accordingly. Through experiments of deformable shape classification, segmentation and matching as well as rigid scene segmentation, we show that PFCNNs using only raw input signals achieve superior performances than competing surface CNN frameworks. In addition, we do extensive ablation studies to validate the components of our framework.

<sup>\*</sup>Joint first author. Work done during internship at Microsoft.

<sup>†</sup>Corresponding author.

## 2. Related work

We briefly review 3D neural networks by classifying them according to the forms of domain representation, and focus on the most related works that use surface meshes.

**3D neural networks for volumetric grids, point clouds and multi-view representations.** The earliest works for 3D deep learning directly extend CNNs to 3D volumetric grids [47, 29], which are later improved for computational efficiency by using adaptive grids like octrees that use high resolution only around the boundary surfaces [37, 46]. Point sets also conveniently encode 3D shapes, for which the set-based PointNet [33] is proposed and later extended by PointNet++ [35] to take advantage of the local surface patch structure. Similarly, more works utilize the local patch structures of 3D point clouds, by e.g. tangent plane projection [42], localization with lattice structure [40], or localized kernel functions [1, 43]. Multi-view representations encode 3D data with a set of 2D images [41, 34], on which standard CNNs are applied to extract intermediate features and aggregated for final output. The PFCNNs presented in this paper work on surface meshes, which have been used pervasively for 3D representation due to their high efficiency for capturing geometry to the fine details.

**Patch-based surface CNNs.** A series of works extend standard CNNs to curved surface domains by applying convolution operations on localized geodesic patches; they differ mainly in the specific ways of convolution computation. Masci et al. [28] parameterize each geodesic patch in polar coordinates, upon which the convolution operation is computed by rotating the kernel function for a set of discrete angles and convolving it with input features; the convolved features for different angles are further pooled for output. With such an approach, it is hard to capture anisotropic or directional signals. Later Boscaini et al. [3] propose anisotropic CNNs that extend [28] by aligning the convolution kernels to frames of principal curvature directions, thus removing angular pooling and ambiguity, and show improved performances on various tasks. Xu et al. [48] use a similar convolution on  $n$ -ring neighboring faces with fixed cardinality for shape segmentation. MoNet [30] extends the geodesic convolutions by modeling the convolution kernel as a mixture of Gaussians whose bases and coefficients are fully trainable rather than functions of fixed parameterizations. TextureNet [16] imposes locally rectangular grids define by 4-directional fields on the geodesic patches, and extract the features for center or corner grid points separately to handle the grid orientation ambiguity. Multi-directional CNNs [32] make the further step of resolving the orientation alignment of geodesic patches by using parallel transport to match the directional convolution responses for different surface points, which enables effective propagation of directional signals. Different from these

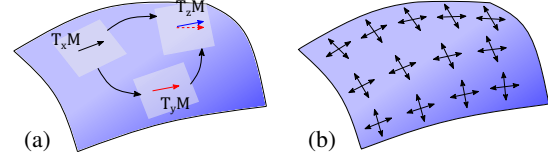


Figure 1. For patch-based surface CNNs, the key problem is how to align the tangent spaces of different surface points. (a) the parallel transport is path-dependent and maps the vector in  $T_x \mathcal{M}$  directly to the blue one in  $T_z \mathcal{M}$  but to the red dashed one by going through  $T_y \mathcal{M}$ . (b) by building a flat connection encoded by the parallel 4-direction frame field, our approach has path-independent translation as in image domain.

manifold based works, SplineCNN [11] defines 3D spline convolution kernels for extracting features on surface and is inherently a volumetric approach focusing on handling the irregular sampling of meshes.

Our PFCNN follows the geodesic convolution paradigm, but differs from others in the convolution computation. Indeed, our framework closely relates to the latest parallel transport approach of [32], but we align the tangent spaces with locally flat connections that not only approximate the parallel transport but also induce a locally Euclidean structure suitable for defining convolutions as for images. In addition, the locally flat connections can be adapted to capture salient geometric features like sharp creases, which further improves performance. As a result, our PFCNNs show superior performances than the previous patch-based surface CNNs on diverse tasks (Sec. 6).

**Surface CNNs using atlas maps.** Another series of works deal with a surface domain by mapping it to a 2D atlas image, on which standard convolutions are applied. Sinha et al. [39] use the geometry image to map a 3D surface of sphere topology to the planar domain and feed the map to CNNs for shape recognition. Maron et al. [27] note the geometry images have gaps between charts of the atlas map and propose to parameterize a surface of sphere topology conformally to the flat image with a toric topology, where standard convolutions with cross boundary cyclic padding are applied. Such convolutions are shown to be conformally translation equivariant but the conformal scaling distortion is uneven for different surface regions. Li et al. [25] handle the gaps of an atlas map by modulating the convolution to jump across the gaps, while the mapping distortion is loosely constrained by subdividing the charts. In comparison, our framework works with surfaces of general topology and automatically preserves the original signals with minimal distortion due to the local patch paradigm.

## 3. Overview

To represent the boundary of a 3D object, we consider a surface mesh  $\mathcal{M} = (V, F)$ , with  $V = \{v_i\}$  the vertices with embedding  $\mathbf{v}_i \in \mathbf{R}^3$ , and  $F = \{f_i = (v_{i0}, v_{i1}, v_{i2})\}$  the faces with corners indexing the vertices. Denote the unit

normal vector at vertex  $v_i$  as  $\mathbf{n}_{v_i} \in \mathbf{R}^3$ , and tangent plane as  $T_{v_i}\mathcal{M}$  on which we can project the local geodesic patch and apply standard image-like convolutions. As reviewed in Sec. 2, while most patch-based surface CNNs follow this general approach, the key challenge is how to coordinate the convolutions for tangent planes of different vertices (Fig. 1). We resolve this challenge by building locally flat connections that align the tangent planes into locally flat Euclidean domains, thus enabling effective weight sharing and translation equivariance that mimic behavior on 2D images.

In Sec. 4 we briefly review the standard Euclidean convolutions with their translation equivariance property, the notion of connections from differential geometry, locally flat connections encoded by  $N$ -direction frame fields and  $N$ -cover spaces for organizing convolution and feature maps. In Sec. 5 we present the extended convolution on surfaces using parallel frames that achieves local translation equivariance and handles irregular vertex sampling on meshes, and the new layers that constitute a PFCNN model.

## 4. Background

### 4.1. Convolution on Euclidean domains

The convolution operation of a CNN exploits the translation equivariance of 2D images [22, 23, 4]. Let  $f, k : \Omega \subset \mathbf{R}^2 \rightarrow \mathbf{R}$  be two functions defined on the image  $\Omega$ , and  $k$  is the convolution kernel usually with a local spatial support. Define the convolution operator  $\star$  as  $f \star k(x) = \int_{y \in \Omega} k(y-x)f(y)dy$ . A planar translation of the image-based function by a vector  $v \in \mathbf{R}^2$  is  $\tau_v(f(x)) = f(x-v)$ . Translation equivariance simply means that the planar translation commutes with convolution, i.e.

$$\tau_v(f \star k) = \tau_v(f) \star k. \quad (1)$$

CNNs parameterize the convolution kernel with trainable weights, which can be shared for different image regions and therefore lead to less overfitting and more generality. As will be discussed next, on curved surface domains the notion of translation is only locally meaningful, which poses difficulty for effective weight sharing of the convolution kernels.

### 4.2. Connections and locally flat connections

Connections generalize the notion of translation onto curved manifolds with non-Euclidean metric [24]. Intuitively, a (linear) connection  $\nabla : T\mathcal{M} \times T\mathcal{M} \rightarrow T\mathcal{M}$  measures the linear differentiation of moving tangent plane  $T_x\mathcal{M}$  along a vector  $v \in T_x\mathcal{M}$  infinitesimally. Therefore, a geodesic curve  $\gamma : [0, 1] \rightarrow \mathcal{M}$  as the “straight line” on a surface has  $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$ , i.e. the curve tangent vector moves straightly along itself. Indeed, the patch-based multi-directional geodesic CNN(MDGCNN) [32] connects the convolutions for two surface patches by translating the

tangent planes along the geodesic curve connecting the two patch centers, which provides a natural extension of translation on 2D images.

However, the problem with parallel transporting along the geodesic curves is that the mapping is path dependent. Consider three nearby points  $x, y, z \in \mathcal{M}$ , and denote the transport of tangent planes along the geodesic curve between  $x, y$  as  $\tau_{x,y} : T_x\mathcal{M} \rightarrow T_y\mathcal{M}$ . In general, we have  $\tau_{y,z} \circ \tau_{x,y} \neq \tau_{x,z}$ , where  $\circ$  is composition, with the difference caused by the curvature of the triangular surface patch bounded by the geodesic curves (Fig. 1(a)).

In this paper, we propose to use a construction called locally flat (or trivial) connections [6, 36] to achieve the path-independent tangent space mapping for all surface patches except at a few singular points. The idea of locally flat connections is to concentrate the surface curvature onto a sparse set of singular points and leave the rest majority of surface area with tangent space mappings as in a Euclidean domain, which in turn paves the way for convolutions as on images.

### 4.3. $N$ -direction frame fields and cover space

One way of encoding the locally flat connections for meshes is through  $N$ -direction frame fields [36, 45]. An  $N$ -direction field at  $x \in \mathcal{M}$  gives a frame of  $N$  rotationally symmetric directions  $\mathbf{u}_x^i \in T_x\mathcal{M}, i = 1, \dots, N$ ; thus two consecutive vectors in the sequence differ by an angle  $\frac{2\pi}{N}$ . A transport (or matching)  $\tau_{x,y}$  between two tangent planes of  $x, y$  can thus be defined by identifying  $\mathbf{u}_x^i$  with  $\mathbf{u}_y^j$ , which simply amounts to a change of bases. In particular, we use the *principal matching* which chooses  $j$  such that  $\|\tau'_{x,y}(\mathbf{u}_x^i) - \mathbf{u}_y^j\|$  is minimal, where  $\tau'_{x,y}$  is the parallel transport between  $x, y$  along geodesics.

In addition, a vertex  $x$  is singular if and only if it has a loop of neighboring vertices  $[p_1, \dots, p_n]$  such that  $\mathbf{u}_{p_1}^i$  mapped by  $\tau_{p_n,p_1} \circ \tau_{p_{n-1},p_n} \circ \dots \circ \tau_{p_1,p_2}$  does not return to itself (Fig. 2(c)). Therefore on a patch containing no singular vertex, the transport  $\tau_{x,y}$  remains the same regardless of the path taken between  $x, y$  [6] (Fig. 1(b)). On the other hand, the concentrated curvature at a singular vertex can only be multiples of  $\frac{2\pi}{N}$ , which explains the usage of  $N$  symmetric directions: larger  $N$  allows for more flexible singularities and flat connections. We discuss the choice of  $N$  later.

By solving for *smooth* (or *parallel*) frame fields that deviate minimally from the parallel transport and align to salient geometric features of the underlying surface [36, 6, 45] (see Appendix B for details), we obtain locally flat connections that closely approximate the linear connection while also having consistency among deformed shapes, therefore supporting improved feature learning by the extended surface convolutions.

While now we can translate tangent planes, another challenge unique to a surface domain rather than 2D images is the lack of canonical axes for the tangent planes. By

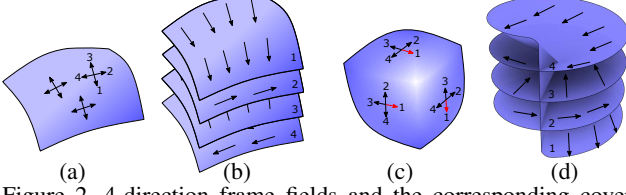


Figure 2. 4-direction frame fields and the corresponding cover spaces. (a)&(b): a field without singular vertex and the four separate sheets of the cover space. (c)&(d): a field with a singular vertex on the cube-corner shaped surface and the four sheets of cover space that are connected and coincide at the singular vertex.

randomly fixing an axis on one tangent plane, we risk significantly biasing the feature learning. Instead, a more robust approach is to sample several directions on the tangent planes as axes, and properly aggregate the learned features for the final output. Fortunately, the  $N$ -direction frames provide a uniform sampling of tangent directions, which motivates introducing their associated  $N$ -cover space that allows to organize the feature learning over multiple axes.

**$N$ -cover space.** A frame field induces an  $N$ -cover space over the domain surface [18, 9]. Intuitively, the cover space consists of  $N$  copies  $\mathcal{M}_i$ ,  $i = 1, \dots, N$  of the base surface, with each copy  $\mathcal{M}_i$  having a unit vector field  $\mathbf{u}_x^{\sigma_x(i)}$ , where  $\sigma_x(i)$  indexes the vector of the  $N$ -direction frame at  $x$  for the sheet  $\mathcal{M}_i$ ; in addition,  $\mathbf{u}_x^{\sigma_x(i)}$  and  $\mathbf{u}_y^{\sigma_y(i)}$  are connected by  $\tau_{x,y}$  the principal matching (Fig. 2). The unit vector field is well-defined everywhere on the cover space, except at singular vertices where different sheets of the cover space coincide. In this paper, we use the vector field as the canonical axes and compute surface convolution on the cover space; around singular vertices, our framework degenerates to a strategy similar to the parallel transport method [32].

## 5. Surface based PFCNN

### 5.1. Surface convolution via parallel frames

Given a surface mesh  $\mathcal{M}$  equipped with a parallel  $N$ -direction frame field  $\mathbf{u}_x^i$ , the surface convolution for a vertex  $v_i$  with its feature vector  $\mathbf{F}_{v_i}^j$ ,  $j = 1, \dots, N$  on the  $\sigma_{v_i}^{-1}(j)$ -th cover sheet is computed by the following steps:

1. Choose  $\mathbf{u}_{v_i}^j$  as the  $x$ -axis of the tangent plane  $T_{v_i}\mathcal{M}$ . Thus the local coordinate system is encoded by the  $2 \times 3$  matrix  $F_{v_i}^j = (\mathbf{u}_{v_i}^j, \mathbf{n}_{v_i} \times \mathbf{u}_{v_i}^j)^T$ .
2. For each vertex  $v_k$  in the neighboring geodesic patch  $\mathcal{N}_{v_i}$ , project it onto the tangent plane as  $\mathbf{v}_k'$  under coordinate system  $F_{v_i}^j$ . Let  $\mathbf{u}_{v_k}^l = \tau_{v_i, v_k}(\mathbf{u}_{v_i}^j)$ ; the projected point has feature vector  $\mathbf{F}_{v_k}^l$ . Resample the projected feature map into a regular grid, denoted  $\mathbf{F}_{\mathcal{N}_{v_i}}^j$ .
3. Convolve  $\mathbf{F}_{\mathcal{N}_{v_i}}^j$  with regular kernels  $\mathbf{K}$  defined under  $F_{v_i}^j$ . The responses constitute the feature vector of  $v_i$  for the next network layer.

Before presenting details for step 2, we remark that translation equivariance in the form of (1) indeed holds locally:

$$\tau_{v_i, v_k}(f \star k) = \tau_{v_i, v_k}(f) \star k, \quad (2)$$

where we assume  $f$  is a function defined on the tangent plane  $T_{v_i}\mathcal{M}$ , and  $k$  is the convolution kernel supported on tangent planes. The equality holds because: on the left hand side,  $f \star k$  returns a function on  $T_{v_i}\mathcal{M}$  which is then transported by the flat connection  $\tau_{v_i, v_k}$  to a function on  $T_{v_k}\mathcal{M}$ , while on the right hand side  $f$  is first transported onto  $T_{v_k}\mathcal{M}$  and then convolved with  $k$  on  $T_{v_k}\mathcal{M}$ ; since the transport  $\tau_{v_i, v_k}$  only changes the underlying coordinate system bases, the functions  $f, k$  when defined using local coordinates do not change at all by the transport, which makes the equality trivially true. In addition, on a patch without singular vertices, the transport and equation holds regardless of the path taken between two vertices, which is different from the path-dependent parallel transport [32]; for patches with singular vertices, because the transport minimizes deviation from the parallel transport (Sec. 4.3), our convolution closely resembles the parallel transport approach.

**Projection to tangent space and resampling.** Previous patch-based surface CNNs use various kinds of geodesic curve tracing to impose a polar coordinate system onto the neighborhood patch  $\mathcal{N}_{v_i}$  and map each neighboring point onto the tangent plane [28, 3, 30, 32]. We follow a similar approach adapted from [5] that is simpler to compute and works even for point clouds, thus enabling easy extension of our framework to point clouds. In particular, we modulate the geodesic coordinates computation using the local axes  $F_{v_i}^j$ , and re-triangulate the projected neighboring points with Delaunay triangulation to avoid flipped triangles, over which a regular grid in the shape of convolution kernels is then resampled and feature vectors interpolated. The operation is encoded by a sparse tensor  $S$  that is pre-computed for a surface mesh and can be applied efficiently with standard NN libraries. See Appendix C for details.

### 5.2. PFCNN structures

In this section we present the detailed structures of layers specific to PFCNNs. These layers can be combined with standard CNN layers and stacked into networks such as U-Net [38] and ResNet [14].

**Input layers.** The PFConv takes as input a group of  $N$  feature maps corresponding to the  $N$  cover sheets. These features can be constructed by simply duplicating the original input for  $N$  copies, i.e.  $|V| \times C_{in} \rightarrow |V| \times N \times C_{in}$ , where  $C_{in}$  is the input per-vertex feature length, or can be computed by further utilizing the local coordinate systems for different cover sheets. Indeed, we find that for tasks on deformable domains, e.g. non-rigid shape classification, segmentation and registration, a simple but effective input



feature that is invariant to global rigid transformations is the normal vector and height from tangent plane in local coordinates, i.e.  $\mathbf{F}_{v_k}^l = (F_{v_k}^j \mathbf{n}_{v_k}, \mathbf{n}_{v_k}^T \mathbf{n}_{v_k}, \mathbf{n}_{v_k}^T (\mathbf{v}_k - \mathbf{v}_i))$  for each patch vertex  $v_k \in \mathcal{N}_{v_i}$ . In this case, the input layer constructs an expanded  $|V| \times N \times H \times W \times C_{in}$  feature map directly by sampling the local features with regular grids (Sec. 5.1), where  $H \times W$  is the spatial shape of the subsequent convolution kernel to be applied.

**Output layers.** For the final per-vertex output we need to reduce the grouped feature maps to an aggregation, i.e.  $|V| \times N \times C \rightarrow |V| \times C$ . The reduction operation can be in different forms, e.g. taking the maximum or average among  $N$  parallel channels, or being learned implicitly by a standard  $1 \times 1$  convolution. The outputs can be further aggregated over all vertices into a single output for the whole shape, as in classification tasks.

**Convolution layers.** Given an input feature map  $\mathbf{F}_{in}$  of shape  $|V| \times N \times C_{in}$ , the convolution layer first vectorizes it into  $\text{vec}(\mathbf{F}_{in})$ , multiplies with the sparse matrix  $S$  of shape  $(|V| \times N \times H \times W, |V| \times N)$  that does the feature map resampling (Sec. 5.1), and reshapes the result vector into a tensor of shape  $|V| \times N \times H \times W \times C_{in}$ ; it then multiplies with the convolution kernel of shape  $H \times W \times C_{in} \times C_{out}$  to obtain the output feature map  $|V| \times N \times C_{out}$ . In case the input layer provides an expanded feature map with local features as discussed above, the convolution is a simple multiplication with the kernel. In addition, the special case of  $1 \times 1$  convolution on each cover sheet through a  $C_{in} \times C_{out}$  kernel skips the feature resampling step and is directly multiplied with  $\mathbf{F}_{in}$  to obtain the output. Note the same convolution kernel is shared for all  $N$  cover sheets of feature maps, as the different cover sheets effectively sample the canonical axes over the surface domain.

**Pooling/Unpooling layers.** Pooling and unpooling layers effectively change the spatial resolution of learned features. For surface meshes the different domain resolutions can be built through a hierarchy of simplified meshes  $\mathcal{M}_i$  with  $\mathcal{M}_1 = \mathcal{M}$  and each coarse vertex  $v \in V_{i+1}$  corresponding to a subset of dense vertices  $\{v'_k\} \subset V_i$ , using e.g. [12, 15]. We adapt the simplification process so that their  $N$ -direction frames are also mapped, i.e.  $F_v^j$  corresponds to  $F_{v'_k}^l$  the closest axes by rotation. Pooling is then defined as  $\mathbf{F}_v^j = \text{Pool}(\{\mathbf{F}_{v'_k}^l\})$ , where  $\text{Pool}(\cdot)$  takes channel-wise maximum or average; the layer has a signature of  $|V_i| \times N \times C \rightarrow |V_{i+1}| \times N \times C$  in terms of feature map shapes. Unpooling is the inverse operation of pooling.

Throughout the paper we assume batch size one, although using larger batch size is trivial as long as each mesh of a batch has the same number of vertices on every domain resolution. We have implemented the above layers with Tensorflow; the code is publicly available<sup>1</sup>.

<sup>1</sup>Code and data are available at <https://github.com/msraig/pfcnn>.

Table 1. Results on SHREC’15 non-rigid shape classification. PN+ is PointNet++[35]; “raw” means using spatial coordinates as input, “en” means using an ensemble of intrinsic shape descriptors. MDG is MDGCNN[32] using SHOT features as input.

	PN+(raw)	PN+(en)	MDG	Ours
Accu.(%)	60.18	96.09	99.5	99.5

Table 2. Results on human body segmentation. Our method outperforms MDGCNN on both original data and the remeshed data.

	Original		Remeshed	
Method	MDGCNN	Ours	MDGCNN	Ours
Accu.(%)	88.2	91.45	89.53	91.79

## 6. Experiments

We test the PFCNN framework and compare it mainly with the state-of-the-art MDGCNN [32] on deformable domain tasks involving shape classification, segmentation and registration where MDGCNN achieves uniformly superior performances than other methods, and with the state-of-the-art TextureNet [16] on the scene semantic segmentation task which has a rigid underlying domain. We further do ablation study on the impact of parallel frames, cover space grouped feature maps and layer normalization, etc.

### 6.1. Deformable domain tasks

For fair comparison, we use  $5 \times 5$  convolution kernel for PFCNN and a larger  $4(\text{radial}) \times 8(\text{angular})$  kernel for MDGCNN, and the same network structure for both methods in each task, except for registration where the same number of convolution layers are adopted. Network and training details are provided in Appendix D.

**Classification.** The SHREC’15 non-rigid shape classification challenge [26] has 1200 shapes represented by surface meshes that belong to 50 categories. We use a network with three levels of resolution and the localized normal vectors as input features (Sec. 5.2) for PFCNN. As shown in Table. 1, our results outperform PointNet++ [35] even when it uses an ensemble of sophisticated input features, e.g. WKS and HKS, that are agnostic to non-rigid deformations. We are on par with MDGCNN that uses as input the SHOT descriptor [44] which is rotation invariant and more sophisticated than our raw input.

**Human body segmentation.** The human body segmentation dataset proposed by [27] contains labeled meshes of diverse human identities and poses from various sources, split by 381/18 for training and testing. The meshes have very different scales which we normalize first. The mesh resolutions are also very different, with the number of vertices varying from 3k to 12k, yet our network works well on these data without remeshing. The network is a U-Net like structure with three levels of domain resolutions.

To compare with MDGCNN, we test on both the original meshes and the resampled meshes generated with its open sourced code. Testing results are reported in Table 2

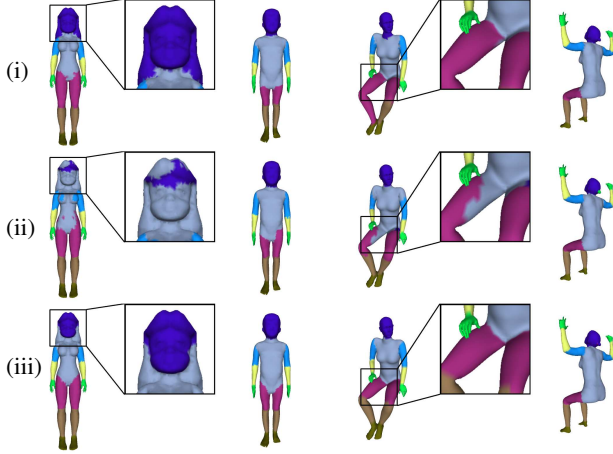


Figure 3. Results of human body segmentation. (i) the ground truth labeling; (ii) the results of MDGCNN; (iii) our results.

and visualized in Fig. 3. Note that the ground truth labeling for different samples are not always consistent, which hinders the possibility of achieving very high accuracy. For example, in Fig. 3 the third column GT mistakenly labels the shank to thigh. But our method correctly segments this part and has better coverage than MDGCNN. Still for some shapes which are dissimilar to the training data, e.g. the first column in Fig. 3 which has exceptional hair, both methods fail to segment the hair properly, although our method captures the face better.

**Human body registration by vertex classification.** We test with the non-rigid human body registration task proposed by the FAUST dataset [2]. In one scenario, the registration is achieved by classifying each input mesh vertex of a body shape into its corresponding vertex on the template mesh, as done in previous works [28, 3, 11]. We use a simple network consisting of a sequence of convolutions in the same level-of-detail for PFCNN, and a two-level network for MDGCNN, following their original setting.

The meshes in the FAUST dataset have the same topology with the template, which may be exploited unfairly to learn correspondences. Following MDGCNN, we remesh them to 5k vertices and different topologies. Using the nearest vertex as the correspondence between original meshes and the remeshed ones, we can get the ground truth vertex correspondence to the remeshed template, to supervise the registration task by classifying each vertex to 5k classes.

We achieved 92.01% accuracy on the remeshed data, as compared to 94.5% accuracy on the original meshes. To fully compare with MDGCNN, we also test variations of their networks with more radial bins and angular directions and different normalizations (more discussions in Sec. 6.3). The accuracies within bounded geodesic errors are plotted in Fig. 4; our results have even better zero-error accuracy than their best with  $4 \times 16$  kernels and instance normalization. The visual results are shown in Fig. 5; we can see that our results have a smoother mapping to the template shape.

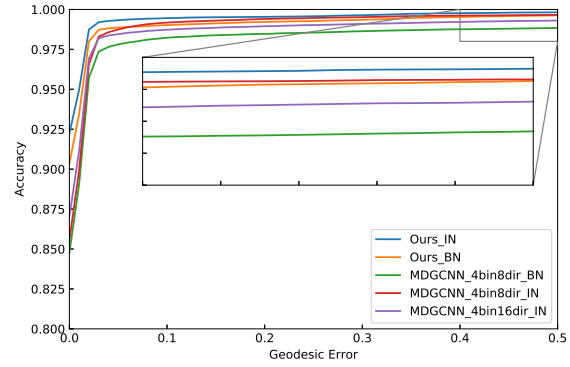


Figure 4. Accuracy within given geodesic error for the non-rigid registration by vertex classification.




Figure 5. Visual comparison of our method and MDGCNN on non-rigid registration. (i) the ground truth mapping; (ii) the best results of MDGCNN with 4 bins, 16 directions; (iii) our result.

In Appendix E, we test with a more challenging scenario of non-rigid registration by regression on noisy real scans with diverse and high genus meshes, where our results are again considerably better and more robust than MDGCNN.

To summarize, compared with the parallel transport based convolution by MDGCNN, using parallel frames that induce locally path-independent transport and the alignment to salient geometric features enables more efficient feature learning for our convolution; the difference is more obvious for finer scale tasks like segmentation and registration.

## 6.2. Semantic scene segmentation

In this section, we evaluate on a widely used indoor scene semantic segmentation task provided by the ScanNet dataset [7]. While indoor scenes generally have rigid geometry dominated by flat walls and floors, PFCNN is still shown to achieve good performances, improving over the state-of-the-art TangentConv [42] and TextureNet [16] that use tangential and local patch convolutions.

We use a network with U-Net structure and three levels of domain resolutions. We follow [16] to prepare the training data by cropping small chunks from a whole scene and

Table 3. Results on ScanNet segmentation task. mIoU is the class mean intersection over union. mA is the class mean accuracy. oA is the overall accuracy, which is significantly biased toward floors and walls that are dominant in scenes. Ours\* uses a network with more convolution layers.

	[42]	[16]	Ours	Ours*
mIoU	0.49	0.58	0.632	0.662
mA(%)	61.4	74.4	75.7	77.92
oA(%)	77.9	80.38	85.01	86.26

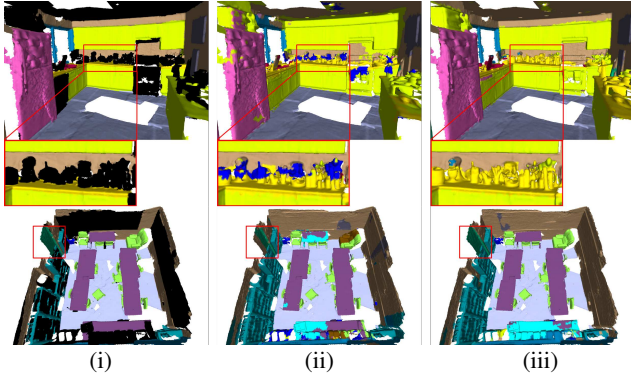


Figure 6. Example indoor scenes of ScanNet segmented by comparing methods. (i) is the ground truth segmentation; (ii) is the results of [16]; (iii) shows our results. Our results have more regular boundaries separating regions of larger consistency.

training on these chunks which are randomly rotated around the upright direction for augmentation. For network input, we follow [42] to include the height above ground, normal vector, color and distance from the local tangent plane for each mesh vertex of a surface patch, rather than the localized normal vector as discussed in Sec. 5.2, while [16] uses additional high resolution texture images as input. For fair comparison, we have used a network with similar amount of trainable parameters to [16]; we also explore the effect of increasing the network size and report a better performance.

The result statistics of comparing methods and ours on validation sets are shown in Table 3; our results have much better mean IoU and mean accuracy than theirs, which shows our network can better distinguish smaller objects than just the dominant segments like floors and walls. Fig. 6 show some visual results. The black regions in (i) are unlabeled data; our method predicts reasonable labels for these regions. The boundaries separating different objects in our results are cleaner than [16], like the boundary between windows and the wall in the first row and the door and wall in the third row; our segments are also more regular and consistent. See Appendix E for more detailed data and visual results on both validation and test sets.

Considering that all three methods use tangent space convolutions, the results demonstrate that our locally translation equivariant convolution as the key difference is more effective in learning features.

Table 4. Testing accuracy of different convolution methods on the non-rigid registration task by vertex classification. PCF means using the principle curvature directions as tangent plane axes. PCF as FF means using the principal directions as the 4-direction frame field for our PFCNN framework.

	PCF	PCF as FF	Ours
Accu.(%)	83.29	89.80	92.01

Table 5. Accuracy and runtime cost of different frame field symmetry orders  $N$ , on the non-rigid registration task by vertex classification. The costs are measured on an RTX2080 GPU.

$N$	1	2	4	6	8
Accuracy(%)	83.11	91.81	92.01	92.45	93.35
Time(ms)	56.81	87.57	139.10	183.91	227.72
Memory(MB)	148.67	156.15	205.81	371.45	409.18

### 6.3. Ablation study

In this section, we evaluate how the core constructions and hyper parameters of PFCNNs affect performance. We also study the impact of normalization on deformable domain tasks, as well as the behavior around singular vertices.

**Using frames and grouped features.** We evaluate the performances of different configurations that add components of the PFCNN construction one-by-one onto a baseline model. The evaluations are done on the task of human body registration by vertex classification (Sec. 6.1).

- **Baseline model.** When using principal curvature frames as coordinate frames of the tangent plane, we have a baseline model similar to a bunch of recent previous works [3, 30, 48, 42]. Using the network structure similar to PFCNN but without aligning the tangent planes by flat connections or feature map grouping by cover sheets, the trainable convolution kernel parameters are actually 16 times of PFCNN. However, the accuracy for this baseline configuration is 83.29% (Table 4), much lower than PFCNN.
- **Principle curvature frames as 4-direction field.** As a modification to the baseline model, we regard the principal curvature frames as a 4-direction frame field and apply the PFCNN network. The result accuracy is 89.8% (Table 4), much higher than the baseline model, while using only 1/16 trainable parameters. The improvement demonstrates that even if the frame field is not globally optimized to be smooth or aligned to salient features, by using its encoded flat connections that enable local translation equivariance and its induced cover space feature maps that sample tangent directions, the feature learning is significantly improved.
- **Full PFCNN model.** By additionally optimizing for a parallel frame field that aligns to geometric features, the PFCNN framework further improves to 92.01% registration accuracy (Table 4).

**Frame symmetry order.** As discussed in Sec. 4.3, when the rotational symmetry order  $N$  of the frame field gets



Table 6. Classification accuracy with different normalization.

	Ours		MDGCNN	
Normalization	BN	IN	BN	IN
Accuracy(%)	11	99.5	14.0	99.5

larger, the frame field has more flexibility to achieve both smoothness and alignment to salient features. However, an increased  $N$  also leads to larger computational cost, as the size of feature maps to compute increases too.

We tested the different  $N$  values again with the registration by vertex classification task, but modified the network structure to make sure each group of the feature map has the same size (i.e. 64), so that for different  $N$  the amount of trainable convolution kernel parameters remains the same. The performances for different  $N$  are shown in Table 5. We can see that the choice of  $N = 4$  strikes a balance between accuracy and computational overhead: for  $N < 4$  the accuracy is notably lower due to the limited field smoothness, and for  $N > 4$  the computational cost is higher, with the extra runtime roughly in proportion to the number of axes sampled. We have used  $N = 4$  for all the other experiments in this paper.

**Normalization.** It is well known that normalization can speed up the training procedure and make it more stable. Here we study the impact of different normalizations on surface based CNNs more closely. For the registration by classification task, we test our method and MDGCNN with batch normalization (BN) and instance normalization (IN). Note that since the batch size is one, the difference between BN and IN is that, the channel wise statistics of moving mean and average are used in testing stage for BN but not IN. The result is shown in Fig. 4. We find that with IN both our method and MDGCNN achieve better performances. We repeat the experiments on the shape classification task (Sec 6.1); the result is shown in Table 6. From all these experiments, we can see that IN is better than BN for these tasks on deformable domains. We argue that this is because the diversely deformed shapes do not share common statistics of channel-wise mean and variance, akin to the observation in image style transfer [31, 10, 17] that these statistics encode styles rather than content.

**Frame field singularity.** As discussed in Sec. 5.1, near singularities the translation equivariance is no longer path-independent, but our scheme degenerates to being similar to MDGCNN using parallel transport. To find the relationship between the singularity of vertices and prediction error, we compare the distribution of singular vertices and the error map of geodesic distance between predicted vertex and the ground truth correspondence vertex. The distribution is shown in Fig. 7; we can see that the singular vertices mainly distribute on the nose, fingers or toes but the error maps of different shapes do not reflect these similarity. We also compare the accuracy of singular vertices and all ver-




Figure 7. (i) shows the singular vertices in red and (ii) shows the prediction error map. There is no clear correlation between the singular vertices and the erroneous predictions.

tices in the registration by classification task. In particular, on the original dataset and the remeshed dataset, the registration accuracy of singular vertices versus that of all vertices are, 93.4%/94.5% and 90.2%/92.2%, indicating no clear correlation of singular vertices and prediction errors. Such robustness can be attributed to the degenerated convolution with path-dependent translation equivariance, the consistency of singularities across shapes (see Appendix B) and the capability of learned filters.

## 7. Conclusion

We have presented a surface mesh based PFCNN framework that closely mimics the standard image based CNNs and has local translation equivariance for convolutions. It is enabled by using parallel  $N$ -direction frames that both encode flat connections on the surface to define path-independent translation, and sample tangent plane canonical axes to organize the convolutions by the  $N$ -cover spaces. The PFCNNs are shown to be more effective at fine-scale feature learning than previous surface based CNNs. In the future, we would like to investigate how the PFCNN framework can handle surface generation tasks, where the frame field also needs to be generated rather than precomputed.

## References

- [1] Matan Atzmon, Haggai Maron, and Yaron Lipman. Point convolutional neural networks by extension operators. *ACM Trans. Graph.*, 37(4):71:1–71:12, July 2018. 2
- [2] Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. FAUST: Dataset and evaluation for 3D mesh registration. In *CVPR*, June 2014. 6
- [3] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *NIPS*, pages 3189–3197. 2016. 1, 2, 4, 6, 7



- [4] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, July 2017. 1, 3
- [5] Max Budninskiy, Gloria Yin, Leman Feng, Yiyi Tong, and Mathieu Desbrun. Parallel transport unfolding: A connection-based manifold learning approach. *SIAM J. Appl. Algebra Geom.*, 3:266–291, 2018. 4
- [6] Keenan Crane, Mathieu Desbrun, and Peter Schröder. Trivial connections on discrete surfaces. *Computer Graphics Forum*, 29(5):1525–1533, 2010. 3
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3D reconstructions of indoor scenes. In *CVPR*, 2017. 6
- [8] Olga Diamanti, Amir Vaxman, Daniele Panozzo, and Olga Sorkine-Hornung. Designing n-polyvector fields with complex polynomials. *Computer Graphics Forum*, 33(5):1–11, Aug. 2014. 11
- [9] Boris A Dubrovin, Anatolij Timofeevič Fomenko, and Sergei Novikov. *Modern geometry—methods and applications: Part II: The geometry and topology of manifolds*. 1, 4
- [10] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *ArXiv*, abs/1610.07629, 2016. 8
- [11] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. SplineCNN: Fast geometric deep learning with continuous B-spline kernels. In *CVPR*, 2018. 2, 6
- [12] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *SIGGRAPH*, pages 209–216, 1997. 5
- [13] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>. 1
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [15] Hugues Hoppe. Progressive meshes. In *SIGGRAPH*, pages 99–108, 1996. 5
- [16] Jingwei Huang, Haotian Zhang, Li Yi, Thomas Funkhouser, Matthias Niessner, and Leonidas J. Guibas. Texturenet: Consistent local parametrizations for learning from high-resolution signals on meshes. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 5, 6, 7, 13, 15
- [17] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 8
- [18] Felix Kälberer, Matthias Nieser, and Konrad Polthier. Quadcover - surface parameterization using branched coverings. *Computer Graphics Forum*, 26(3):375–384, 2007. 4
- [19] Vladimir Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM Trans. Graph. (SIGGRAPH)*, 30(4), July 2011. 14
- [20] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014. 12
- [21] Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. Globally optimal direction fields. *ACM Trans. Graph. (SIGGRAPH)*, 32(4):59:1–59:10, 2013. 11
- [22] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989. 3
- [23] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998. 1, 3
- [24] John M. Lee. Connections. In *Riemannian Manifolds*, pages 47–64. Springer New York, 1997. 3
- [25] Shiwei Li, Zixin Luo, Mingmin Zhen, Yao Yao, Tianwei Shen, Tian Fang, and Long Quan. Cross-atlas convolution for parameterization invariant learning on textured mesh surface. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6143–6152, 2019. 2
- [26] Z. Lian, J. Zhang, S. Choi, H. ElNaghy, J. El-Sana, T. Furuya, A. Giachetti, R. A. Guler, L. Lai, C. Li, H. Li, F. A. Limberger, R. Martin, R. U. Nakanishi, A. P. Neto, L. G. Nonato, R. Ohbuchi, K. Pevzner, D. Pickup, P. Rosin, A. Sharf, L. Sun, X. Sun, S. Tari, G. Unal, and R. C. Wilson. Non-rigid 3D Shape Retrieval. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2015. 5
- [27] Haggai Maron, Meirav Galun, Noam Aigerman, Miri Trope, Nadav Dym, Ersin Yumer, Vladimir G. Kim, and Yaron Lipman. Convolutional neural networks on surfaces via seamless toric covers. *ACM Trans. Graph. (SIGGRAPH)*, 36(4):71:1–71:10, July 2017. 2, 5
- [28] J. Masci, D. Boscaini, M. M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *ICCV*, pages 832–840, Dec 2015. 1, 2, 4, 6
- [29] 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)*, pages 922–928. IEEE, 2015. 2
- [30] 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 *CVPR*, 2017. 1, 2, 4, 7
- [31] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 8
- [32] Adrien Poulenard and Maks Ovsjanikov. Multi-directional geodesic neural networks via equivariant convolution. *ACM Trans. Graph. (SIGGRAPH ASIA)*, 37(6):236:1–236:14, Dec. 2018. 1, 2, 3, 4, 5, 14
- [33] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2017. 2
- [34] 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*, pages 5648–5656, 2016. 2

- [35] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017. [2](#), [5](#)
- [36] Nicolas Ray, Bruno Vallet, Laurent Alonso, and Bruno Lévy. Geometry Aware Direction Field Processing. *ACM Transactions on Graphics*, 29(1):Article 1, Dec. 2009. [3](#)
- [37] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3D representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3577–3586, 2017. [2](#)
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [4](#)
- [39] Ayan Sinha, Jing Bai, and Karthik Ramani. Deep learning 3D shape surfaces using geometry images. In *ECCV*, pages 223–240, 2016. [2](#)
- [40] Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2530–2539, 2018. [2](#)
- [41] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik G. Learned-Miller. Multi-view convolutional neural networks for 3D shape recognition. In *ICCV*, 2015. [2](#)
- [42] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3D. In *CVPR*, 2018. [2](#), [6](#), [7](#), [15](#)
- [43] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. Kpconv: Flexible and deformable convolution for point clouds. *Proceedings of the IEEE International Conference on Computer Vision*, 2019. [2](#)
- [44] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *European conference on computer vision*, pages 356–369. Springer, 2010. [5](#)
- [45] Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommes, Klaus Hildebrandt, and Mirela Ben-Chen. Directional Field Synthesis, Design, and Processing. *Computer Graphics Forum*, 2016. [1](#), [3](#), [11](#)
- [46] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3D shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):72, 2017. [2](#)
- [47] 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*, pages 1912–1920, 2015. [2](#)
- [48] H. Xu, M. Dong, and Z. Zhong. Directionally convolutional networks for 3D shape segmentation. In *ICCV*, pages 2717–2726, Oct 2017. [2](#), [7](#)