

Point Convolutional Neural Networks by Extension Operators

Matan Atzmon*

Haggai Maron*

Yaron Lipman

Weizmann Institute of Science




Figure 1: A new framework for applying convolution to functions defined over point clouds: First, a function over the point cloud (in this case the constant one) is *extended* to a continuous volumetric function over the ambient space; second, a continuous volumetric *convolution* is applied to this function (without any discretization or approximation); and lastly, the result is *restricted* back to the point cloud.

Abstract

This paper presents *Point Convolutional Neural Networks (PCNN)*: a novel framework for applying convolutional neural networks to point clouds. The framework consists of two operators: *extension* and *restriction*, mapping point cloud functions to volumetric functions and vice-versa. A point cloud convolution is defined by pull-back of the Euclidean volumetric convolution via an extension-restriction mechanism.

The point cloud convolution is computationally efficient, invariant to the order of points in the point cloud, robust to different samplings and varying densities, and translation invariant, that is the same convolution kernel is used at all points. PCNN generalizes image CNNs and allows readily adapting their architectures to the point cloud setting.

Evaluation of PCNN on three central point cloud learning benchmarks convincingly outperform competing point cloud learning methods, and the vast majority of methods working with more informative shape representations such as surfaces and/or normals.

1 Introduction

The huge success of deep learning in image analysis motivates researchers to generalize deep learning techniques to work on 3D shapes. Differently from images, 3D data has several popular representation, most notably surface meshes and points clouds. Surface-based methods exploit connectivity information for 3D deep learning based on rendering [39], local and global parameterization [24, 38, 23], or

spectral properties [45]. Point cloud methods rely mostly on points' locations in three-dimensional space and need to implicitly infer how the points are connected to form the underlying shape.

The goal of this paper is to introduce Point Cloud Convolutional Neural Networks (PCNN) generalizing deep learning techniques, and in particular Convolutional Neural Networks (CNN) [22], to point clouds. As a point cloud $X \subset \mathbb{R}^3$ is merely an approximation to some underlying shape S , the main challenges in building point cloud networks are to achieve: (i) Invariance to the order of points supplied in X ; (ii) Robustness to sampling density and distribution of X in S ; and (iii) Translation invariance of the convolution operator (*i.e.*, same convolution kernel is used at all points).

Invariance to point order in X was previously tackled in [31, 34, 33, 46] by designing networks that are composition of euquivariant layers (*i.e.*, commute with permutations) and a final symmetric layer (*i.e.*, invariant to permutations). As shown in [34], any linear equivariant layer is a combination of scaled identity and constant linear operator and therefore missing many of the degrees of freedom existing in standard linear layers such as fully connected and even convolutional.

Volumetric grid methods [43, 25, 32, 35] use 3D occupancy grid to deal with the point order in X and provide translation invariance of the convolution operator. However, they quantize the point cloud to a 3D grid, usually producing a crude approximation to the underlying shape (*i.e.*, piecewise constant on voxels) and are confined to a fixed 3D grid structure.

Our approach toward these challenges is to define CNN

*equal contribution

on a point cloud X using a pair of operators we call *extension* \mathcal{E}_X and *restriction* \mathcal{R}_X . The extension operator maps functions defined over the point cloud X to volumetric functions (*i.e.*, functions defined over the entire ambient space \mathbb{R}^3), where the restriction operator does the inverse action. Using $\mathcal{E}_X, \mathcal{R}_X$ we can translate operators such as Euclidean volumetric convolution to point clouds, see Figure 1. In a nutshell, if O is an operator on volumetric functions then its restriction to the point cloud X would be

$$O_X = \mathcal{R}_X \circ O \circ \mathcal{E}_X. \quad (1)$$

We take \mathcal{E}_X to be a Radial Basis Function (RBF) approximation operator, and \mathcal{R}_X to be a sampling operator, *i.e.*, sample a volumetric function at the points in X . As O we take continuous volumetric convolution operators with general kernels κ represented in the RBF basis as-well. In turn (1) is calculated using a sparse linear tensor combining the learnable kernel weights k , function values over the point cloud X , and a tensor connecting the two, defined directly from the point cloud X .

Since our choice of \mathcal{E}_X is invariant to point order in X , and \mathcal{R}_X is an equivariant operator (w.r.t. X) we get that O_X in (1) is equivariant. This construction leads to new equivariant layers, in particular convolutions, with more degrees of freedom compared to [34, 31, 46]. The second challenge of robustness to sampling density and distribution is addressed by the *approximation power* of the extension operator \mathcal{E}_X . Given a continuous function defined over a smooth surface, $f : S \rightarrow \mathbb{R}$, we show that the extension of its restriction to X approximates the restriction of f to S , namely

$$\mathcal{E}_X \circ \mathcal{R}_X[f] \approx f|_S.$$

This means that two different samplings $X, X' \subset S$ of the same surface function are extended to the *same volumetric function*, up to an approximation error. In particular, we show that extending the simplest, constant one function over the point cloud, $\mathcal{E}_X[1]$, approximates the indicator function of the surface S , while the gradient, $\nabla \mathcal{E}_X[1]$, approximates the mean curvature normal field over the surface. Then, the translation invariance and robustness of our convolution operator naturally follows from the fact that the volumetric convolution is translation invariant and the extension operator is robust.

PCNN provides a flexible framework for adapting standard image-based CNNs to the point cloud setting, while maintaining data only over the point cloud on the one hand, and learning convolution kernels robust to sampling on the other. We have tested our PCNN framework on standard classification, segmentation and normal estimation datasets where PCNN outperformed all other point cloud methods and the vast majority of other methods that use more informative shape representations such as surface connectivity.

2 Previous Work

We review different aspects of geometric deep learning with a focus on the point cloud setting. For a more comprehensive survey on geometric deep learning we refer the reader to [7].

Deep learning on point clouds. PointNet [31] pioneered deep learning for point clouds with a Siamese, per-point network composed with a symmetric max operator that guarantees invariance to the points’ order. PointNet was proven to be a universal approximator (*i.e.*, can approximate arbitrary continuous functions over point clouds). A follow up work [33] suggests a hierarchical application of the PointNet model to different subsets of the point cloud; this allows capturing structure at different resolutions when applied with a suitable aggregation mechanism. In [16] the PointNet model is used to predict local shape properties from point clouds. In a related work [34, 46] suggest to approximate set function, with equivariant layers composed with a symmetric function such as max. Most related to our work is the recent work of [21] that suggested to generalize convolutional networks to point clouds by defining convolutions directly on kd-trees built out of the point clouds [3], and [36] that suggested a convolutional architecture for modeling quantum interactions in molecules represented as point clouds, where convolutions are defined by multiplication with continuous filters. The main difference to our work is that we define the convolution of a point cloud function using an exact volumetric convolution with an extended version of the function. The approximation properties of the extended function facilitate a robust convolution on point clouds.

Volumetric methods. Another strategy is to generate a tensor volumetric representation of the shape restricted to a regular grid (*e.g.*, by using occupancy indicators, or a distance function) [43, 25, 32]. The main limitation of these methods is the approximation quality of the underlying shape due to the low resolution enforced by the three dimensional grid structure. To overcome this limitation a few methods suggested to use sparse three dimensional data structures such as octrees [40, 35]. Our work can be seen as a generalization of these volumetric methods in that it allows replacing the grid cell’s indicator functions as the basis for the network’s functions and convolution kernels with more general basis functions (*e.g.*, radial basis functions).

Deep learning on Graphs. Shapes can be represented as graphs, namely points with neighboring relations. In spectral deep learning the convolution is being replaced by a diagonal operator in the graph-Laplacian eigenbasis

[9, 14, 18]. The main limitation of these methods in the context of geometric deep learning is that different graphs have different spectral bases and finding correspondences between the bases or common bases is challenging. This problem was recently targeted by [45] using the functional map framework.

Deep learning on surfaces. Other approaches to geometric deep learning work with triangular meshes that possess also connectivity and normal information, in addition to the point locations. One class of methods use rendering and 2D projections to reduce the problem to the image setting [39, 19]. Another line of works uses local surface representations [24, 5, 26] or global parameterizations of surfaces [38, 23] for reducing functions on surfaces to the planar domain or for defining convolution operators directly over the surfaces.

RBF networks. RBF networks are a type of neural networks that use RBF functions as an activation layer, see [29, 28]. This model was first introduced in [8], and was used, among other things, for function approximation and time series prediction. Usually, these networks have three layers and their output is a linear combination of radial basis functions. Under mild conditions this model can be shown to be a universal approximator of functions defined on compact subsets of \mathbb{R}^d [30]. Our use of RBFs is quite different: RBFs are used in our extension operator solely for the purpose of defining point cloud operators, whereas the ReLU is used as an activation.

3 Method

Notations. We will use tensor (*i.e.*, multidimensional arrays) notation, *e.g.*, $a \in \mathbb{R}^{I \times I \times J \times L \times M}$. Indexing a particular entry is done using corresponding lower-case letters, $a_{ii'jlm}$, where $1 \leq i, i' \leq I$, $1 \leq j \leq J$, etc. When summing tensors $c = \sum_{ijl} a_{ii'jlm} b_{ijl}$, where $b \in \mathbb{R}^{I \times J \times L}$ the dimensions of the result tensor c are defined by the free indices, in this case $c = c_{i'm} \in \mathbb{R}^{I \times M}$.

Goal. Our goal is to define convolutional neural networks on point clouds $X = \{x_i\}_{i=1}^I \in \mathbb{R}^{I \times 3}$. Our approach to defining point cloud convolution is to extend functions on point clouds to volumetric functions, perform standard Euclidean convolution on these functions and sample them back on the point cloud.

We define an extension operator

$$\mathcal{E}_X : \mathbb{R}^{I \times J} \rightarrow C(\mathbb{R}^3, \mathbb{R}^J), \quad (2)$$

where $\mathbb{R}^{I \times J}$ represents the collection of functions $f : X \rightarrow \mathbb{R}^J$, and $C(\mathbb{R}^3, \mathbb{R}^J)$ volumetric functions $\psi : \mathbb{R}^3 \rightarrow \mathbb{R}^J$.

Together with the extension operator we define the restriction operator

$$\mathcal{R}_x : C(\mathbb{R}^3, \mathbb{R}^M) \rightarrow \mathbb{R}^{I \times M}. \quad (3)$$

Given a convolution operator $O : C(\mathbb{R}^3, \mathbb{R}^J) \rightarrow C(\mathbb{R}^3, \mathbb{R}^M)$ we adapt O to the point cloud X via (1). We will show that a proper selection of such point cloud convolution operators possess the following desirable properties:

1. *Efficiency:* O_X is computationally efficient.
2. *Invariance:* O_X is indifferent to the order of points in X , that is, O_X is equivariant.
3. *Robustness:* Assuming $X \subset S$ is a sampling of an underlying surface S , and $f \in C(S, \mathbb{R})$, then $\mathcal{E}_X \circ \mathcal{R}_X[f] \in C(\mathbb{R}^3, \mathbb{R})$ approximates f when sampled over S and decays to zero away from S . In particular $\mathcal{E}_X[\mathbf{1}]$ approximates the volumetric indicator function of S , where $\mathbf{1} \in \mathbb{R}^{I \times 1}$ is the vector of all ones; $\nabla \mathcal{E}_X[\mathbf{1}]$ approximates the mean curvature normal field over S . The approximation property in particular implies that if $X, X^* \subset S$ are different samples of S then $O_X \approx O_{X^*}$.
4. *Translation invariance:* O_X is translation invariant, defined by a stationary (*i.e.*, location independent) kernel.

In the next paragraphs we define these operators and show how they are used in defining the main building blocks of PCNN, namely: convolution, pooling and upsampling. We discuss the above theoretical properties in Section 4.

Extension operator

The extension operator $\mathcal{E}_X : \mathbb{R}^{I \times J} \rightarrow C(\mathbb{R}^3, \mathbb{R}^J)$ is defined as an operator of the form,

$$\mathcal{E}_X[f](x) = \sum_i f_{ij} \ell_i(x), \quad (4)$$

where $f \in \mathbb{R}^{I \times J}$, and $\ell_i \in C(\mathbb{R}^3, \mathbb{R})$ can be thought of as basis functions defined per evaluation point x . One important family of bases are the Radial Basis Functions (RBF), that were proven to be useful for surface representation [4, 11]. For example, one can consider interpolating bases (*i.e.*, satisfying $\ell_i(x_{i'}) = \delta_{ii'}$) made out of an RBF $\Phi : \mathbb{R}_+ \rightarrow \mathbb{R}$. Unfortunately, computing (4) in this case amounts to solving a dense linear system of size $I \times I$. Furthermore, it suffers from bad condition number as the number of points is increased [42]. In this paper we will advocate a novel approximation scheme of the form

$$\ell_i(x) = c\omega_i \Phi(|x - x_i|), \quad (5)$$




Figure 2: Applying the extension operator to the constant 1 function over three airplane point clouds in different sampling densities: 2048, 1024 and 256 points. Note how the extended functions resemble the airplane indicator function, and hence similar to each other.

where c is a constant depending on the RBF Φ and ω_i can be thought of the amount of shape area corresponding to point x_i . A practical choice of ω_i is

$$\omega_i = \frac{1}{c \sum_{i'} \Phi(|x_{i'} - x_i|)}. \quad (6)$$

Note that although this choice resembles the Nadaraya-Watson kernel estimator [27], it is in fact different as the denominator is independent of the approximation point x ; this property will be useful for the closed-form calculation of the convolution operator.

As we prove in Section 4, the point cloud convolution operator, O_X , defined using the extension operator, (4)-(5), satisfies the properties (1)-(4) listed above, making it suitable for deep learning on point clouds. In fact, as we show in Section 4, robustness is the result of the extension operator \mathcal{E}_X approximating a continuous, sampling independent operator over the underlying surface S denoted \mathcal{E}_S . This continuous operator applied to a function f , $\mathcal{E}_S[f]$, is proved to approximate the restriction of f to the surface S .

Figure 2 demonstrates the robustness of our extension operator \mathcal{E}_X ; applying it to the constant one function, evaluated on three different sampling densities of the same shape, results in approximately the same shape.

Kernel model

We consider a continuous convolution operator $O : C(\mathbb{R}^3, \mathbb{R}^J) \rightarrow C(\mathbb{R}^3, \mathbb{R}^M)$ applied to vector valued function $\psi \in C(\mathbb{R}^3, \mathbb{R}^J)$,

$$O[\psi](x) = \psi * \kappa(x) = \int_{\mathbb{R}^3} \sum_j \psi_j(y) \kappa_{jm}(x-y) dy, \quad (7)$$

where $\kappa \in C(\mathbb{R}^3, \mathbb{R}^{J \times M})$ is the convolution kernel that is also represented in the same RBF basis:

$$\kappa_{jm}(z) = \sum_l k_{ljm} \Phi(|z - y_l|), \quad (8)$$

where with a slight abuse of notation we denote by $k \in \mathbb{R}^{L \times J \times M}$ the tensor representing the continuous

kernel in the RBF basis. Note, that k represents the network’s learnable parameters, and has similar dimensions to the convolution parameters in the image case (*i.e.*, spatial dimensions \times input channels \times output channels).

The translations $\{y_l\}_{l=1}^L \subset \mathbb{R}^3$ are also a degree of freedom and can be chosen to generate a regular $3 \times 3 \times 3$ grid or any other point distribution such as spherical equispaced points. The translations can be predefined by the user or learned (with some similarity to [13]). See inset for illustration of some possible translations.



Restriction operator

Our restriction operator $\mathcal{R}_X : C(\mathbb{R}^3, \mathbb{R}^J) \rightarrow \mathbb{R}^{I \times J}$ is the sampling operator over the point cloud X ,

$$\mathcal{R}_X[\psi] = \psi_j(x_i), \quad (9)$$

where $\psi \in C(\mathbb{R}^3, \mathbb{R}^J)$. Note that $\mathcal{R}_X[\psi] \in \mathbb{R}^{I \times J}$.

Sparse extrinsic convolution

We want to compute the convolution operator $O_X : \mathbb{R}^{I \times J} \rightarrow \mathbb{R}^{I \times M}$ restricted to the point cloud X as defined in (1) with the convolution operator O from (7). First, we compute $\mathcal{E}_X[f] * k$

$$\mathcal{E}_X[f] * k(x) = c \sum_{ijl} f_{ij} k_{ljm} \omega_i \int_{\mathbb{R}^3} \Phi(|y - x_i|) \Phi(|x - y - y_l|) dy$$

Applying our restriction operator finally gives our point cloud convolution operator:

$$O_X[f] = c \sum_{ijl} f_{ij} k_{ljm} \omega_i q_{iil}, \quad (10)$$

where $q = q(X) \in \mathbb{R}^{I \times I \times L}$ is the tensor defined by

$$q_{iil} = \int_{\mathbb{R}^3} \Phi(|y - x_i|) \Phi(|x_i - y - y_l|) dy. \quad (11)$$

Note that $O_X[f] \in \mathbb{R}^{I \times M}$, as desired. Equation (10) shows that the convolution’s weights k_{ljm} are applied to the data f_{ij} using point cloud-dependent weights w , q that can be seen as “translators” of k to the point cloud geometry X . Figure 3 illustrates the computational flow of the convolution operator.

Choice of RBF

Our choice of radial basis function Φ stems from two desired properties: First, we want the extension operator




Figure 3: Point cloud convolution operator, computational flow.

(4) to have approximation properties; second, we want the computation of the convolution of a pair of RBFs in (11) to have an efficient closed-form solution. A natural choice satisfying these requirements is the Gaussian:

$$\Phi_{\sigma}(r) = \exp\left(-\frac{r^2}{2\sigma^2}\right) \quad (12)$$

To compute the tensor $q \in \mathbb{R}^{I \times I \times L}$ in (11) we make use of the following convolution rule for Gaussians, proved in Appendix A for completeness:

Lemma 1. *Let Φ denote the Gaussian as in (12). Then,*

$$\Phi_{\alpha}(|\cdot - a|) * \Phi_{\beta}(|\cdot - b|) \propto \Phi_{\gamma}(|\cdot - a - b|), \quad (13)$$

where $\gamma = \sqrt{\alpha^2 + \beta^2}$.

3.1 Up-sampling and pooling

Aside from convolutions, there are two operators that perform spatially and need to be defined for point clouds: up-sampling $U_{X, X^*} : \mathbb{R}^{I \times J} \rightarrow \mathbb{R}^{I^* \times J}$, and pooling $P_{X, X^*} : \mathbb{R}^{I \times J} \rightarrow \mathbb{R}^{I^* \times J}$, where $X^* \subset \mathbb{R}^3$ is superset of X (i.e., $I^* > I$) in the upsampling case and subset of X (i.e., $I^* < I$) in the pooling case. The upsample operator is defined by

$$U_{X, X^*}[f] = \mathcal{R}_{X^*} \circ \mathcal{E}_X[f]. \quad (14)$$

Pooling does not require the extension/restriction operators and (similarly to [33]) is defined by

$$P_{X, X^*}[f](x_i^*) = \max_{i \in \mathcal{V}_i^*} f_{ij}, \quad (15)$$

where $\mathcal{V}_i^* \subset \{1, 2, \dots, I\}$ denotes the set of indices of points in X that are closer in Euclidean distance to x_i^* than any other point in X^* . The point cloud $X^* \subset X$ in the next

layer is calculated using farthest point sampling of the input point cloud X .

Lastly, similarly to [20] we implement *deconvolution* layers by an upsampling layer followed by a regular convolution layer.

4 Properties

In this section we discuss the properties of the point cloud operators we have defined above.

4.1 Invariance and equivariance

Given a function $f \in \mathbb{R}^{I \times J}$ on a point cloud $X \in \mathbb{R}^{I \times 3}$, an equivariant layer $L : \mathbb{R}^{I \times J} \rightarrow \mathbb{R}^{I \times M}$ satisfies

$$L(\pi f) = \pi L(f),$$

where $\pi \in \Pi_I \subset \mathbb{R}^{I \times I}$ is an arbitrary permutation matrix. Equivariant layers have been suggested in [31, 34, 46] to learn data in the form of point clouds (or sets in general). The key idea is that equivariant layers can be used to represent a set function $F : 2^{\mathbb{R}^3} \rightarrow \mathbb{R}$. Indeed, a set function restricted to sets of fixed size (say, I) can be represented as a symmetric function (i.e., invariant to the order of its arguments). A rich class of symmetric functions can be built by composing equivariant layers and a final symmetric layer.

The equivariance of our point cloud operators O_X stems from the invariance property of the extension operator and equivariance property of the restriction operator. We will next show these properties.

Lemma 2. *The extension operators defined in (4) is invariant to permutations, i.e., $\mathcal{E}_{\pi(X)}[f] = \mathcal{E}_X[f]$, for all permutations $\pi \in \Pi_I$. The restriction operator (9) is equivariant to permutations, $\mathcal{R}_{\pi(X)}[\psi] = \pi \mathcal{R}_X[\psi]$, for all $\pi \in \Pi_I$.*

Proof. The properties follow from the definitions of the operators.

$$\mathcal{E}_{\pi(X)}[f] = \sum_i f_{\pi(i)j} \ell_{\pi(i)} = \sum_i f_{ij} \ell_i = \mathcal{E}_X[f],$$

and

$$\mathcal{R}_X[\psi] = \psi_j(x_{\pi(i)}) = \Pi \mathcal{R}_X[\psi]. \quad \square$$

A consequence of this lemma is that any convolution O acting on volumetric functions in \mathbb{R}^3 translates to an equivariant operator O_X ,

Theorem 1. *Let $O : C(\mathbb{R}^3, \mathbb{R}^J) \rightarrow C(\mathbb{R}^3, \mathbb{R}^M)$ be a volumetric function operator. Then $O_X : \mathbb{R}^{I \times J} \rightarrow \mathbb{R}^{I \times M}$ defined by (1) is equivariant. Namely,*

$$O_{\pi(X)}[f] = \pi O_X[f].$$

Theorem 1 applies in particular to convolutions (7), and therefore our point cloud convolutions are all equivariant by construction. Note that this model provides "data-dependent" equivariant operator that are more general than those suggest in [34, 46].

4.2 Robustness

Overview. Robustness is the key property that allows applying the same convolution kernel to functions over different irregular point clouds. The key idea is to make the extension operator produce approximately the same volumetric function when applied to different samplings of the same underlying shape function. To make things concrete, let $X \in \mathbb{R}^{I \times 3}$, $X^* \in \mathbb{R}^{I^* \times 3}$ be two different point clouds samples of a compact smooth surface $S \subset \mathbb{R}^3$. Let $f \in C(S, \mathbb{R}^J)$ be some function over S and $\mathcal{R}_X[f]$, $\mathcal{R}_{X^*}[f]$ its sampling on the points clouds X , X^* , respectively.

We will show the following:

1. We introduce a continuous extension operator E_S from surface functions to volumetric functions. We show that E_S has several favorable properties.
2. We show that (under mild assumptions) our extension operator \mathcal{E}_X , defined in (4)-(5) converges to \mathcal{E}_S ,

$$\mathcal{E}_X \circ \mathcal{R}_X[f] \approx \mathcal{E}_S[f]. \quad (16)$$

3. We deduce that (under mild assumptions) the properties of \mathcal{E}_S are inherited by \mathcal{E}_X and in particular we have:

$$\mathcal{E}_X \circ \mathcal{R}_X[f] \approx \mathcal{E}_{X^*} \circ \mathcal{R}_{X^*}[f]. \quad (17)$$

Continuous extension operator. We define $\mathcal{E}_S : C(S, \mathbb{R}^J) \rightarrow C(\mathbb{R}^3, \mathbb{R}^J)$ which is an extension operator from surface functions to volumetric functions so that $\mathcal{E}_S[f]|_S \approx f$ and $\mathcal{E}_S[f] \rightarrow 0$ away from S :

$$\mathcal{E}_S[f](x) = \frac{1}{2\pi\sigma^2} \int_S f(y) \Phi_\sigma(|x - y|) da(y), \quad (18)$$

where da is the area element of the surface S .

The operator \mathcal{E}_S enjoys several favorable approximation properties: First,

$$\mathcal{E}_S[f](x) \xrightarrow{\sigma \rightarrow 0} \begin{cases} f(x) & x \in S \\ 0 & \text{otherwise} \end{cases}. \quad (19)$$

That is, $\mathcal{E}_S[f]$ approximates f over S and decays to zero away from S . In particular, this implies that the constant one function, $\mathbf{1} : S \rightarrow \mathbb{R}$, satisfies

$$\mathcal{E}_S[\mathbf{1}](x) \rightarrow \chi_S(x), \quad (20)$$

where $\chi_S(x)$ is the volumetric indicator function of $S \subset \mathbb{R}^3$. Interestingly, $\mathcal{E}_S[\mathbf{1}]$ provides also higher-order geometric information of the surface S ,

$$\nabla \mathcal{E}_S[\mathbf{1}] \Big|_S \rightarrow -H \cdot n, \quad (21)$$

where $H : S \rightarrow \mathbb{R}$ is the mean curvature function of S and $n : S \rightarrow S^2$ ($S^2 \subset \mathbb{R}^3$ is the unit sphere) is the normal field to S .

We prove that the approximation quality in (16) improves as the point cloud sample $X \subset S$ densifies S , and the operator \mathcal{E}_X becomes more and more consistent. In that case $\mathcal{E}_X[\mathbf{1}]$ furnishes an approximation to the indicator function of the surface S and its gradient, $\nabla \mathcal{E}_X[\mathbf{1}]$, to the mean curvature vectors of S . This demonstrates that given the simplest, all ones input data $\mathbf{1} \in \mathbb{R}^{I \times 1}$, the network can already reveal the indicator function and the mean curvature vectors of the underlying surface by simple linear operators corresponding to specific choices of the kernel k in (8).

These results are summarized in the following theorem which is proved in appendix A.

Theorem 2. *Let $f \in C(S, \mathbb{R}^J)$ be a continuous function defined on a compact smooth surface $S \subset \mathbb{R}^3$. The extension operator (4),(5) with*

$$c = \frac{1}{2\pi\sigma^2}, \quad (22)$$

and

$$\omega_i = \text{area}(\Omega_i), \quad (23a)$$

$$\Omega_i = \{y \in S \mid d_S(y - x_i) \leq d_S(y - x_{i'}), \forall i'\}, \quad (23b)$$

the Voronoi cell of $x_i \in S$, where d_S denotes the distance function of points on S , satisfies

$$\mathcal{E}_X \circ \mathcal{R}_X[f](x) \rightarrow \mathcal{E}_S[f](x), \quad (24)$$

where $X \subset S$ is a δ -net and $\delta \rightarrow 0$. Furthermore, \mathcal{E}_S satisfies the approximation and mean curvature properties as defined in (19), (20), (21).

4.3 Revisiting image CNNs

Our model is a generalization of image CNNs. Images can be viewed as point clouds in regular grid configuration, $X = \{x_i\} \subset \mathbb{R}^3$, with image intensities \mathcal{I}_i as functions over this point cloud,

$$\mathcal{E}_X(\mathcal{I}) = \sum_i \mathcal{I}_i \Phi(x - x_i),$$

where Φ is the indicator function over one square grid cell (i.e., pixel). In this case the extension operator reproduces the image as a volumetric function over \mathbb{R}^2 . Writing the convolution kernel also in the basis Φ with regular grid translations leads to (1) reproducing the standard image discrete convolution.

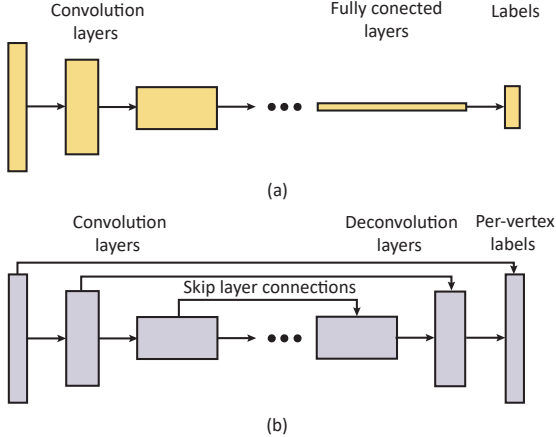


Figure 4: Different architectures used in the paper: (a) classification network; and (b) segmentation network.

5 Experiments

We have tested our PCNN framework on the problems of point cloud classification, point cloud segmentation, and point cloud normal estimation. We also evaluated the different design choices and network variations.

5.1 Point cloud classification

We tested our method on the standard ModelNet40 and ModelNet10 benchmarks [43]. ModelNet40 is composed of 9843 train models and 2468 test models in 40 different classes, such as guitar, cone, laptop etc. ModelNet 10 consists 3991 train and 908 test models from ten different classes. The models are originally given as triangular meshes. The upper part of Table 1 compares our classification results versus state of the art learning algorithms that use only the point clouds (*i.e.*, coordinates of points in \mathbb{R}^3) as input: PointNet [31], PointNet++ [33], deep sets [46], ECC[37] and kd-network[21]. For completeness we also provide results of state of the art algorithms taking as input additional data such as meshes and normals. Our method outperforms all point cloud methods and all other non-ensemble methods.

We use the point cloud data of [31, 33] that sampled a point cloud from each model using farthest point sampling. In the training we randomly picked 1024 farthest point sample out of a fixed set of 1200 farthest point sample for each model. As in [21] we also augment the data with random anisotropic scaling in the range $[-0.66, 1.5]$ and uniform translations in the range $[-0.2, 0.2]$. As input to the network we provide the constant one tensor, together with the coordinate functions of the points, namely $(1, x) \in \mathbb{R}^{I \times 4}$. The σ parameter controls the variance of the RBFs (both in the convolution kernels and the extension operator) and is

chosen to be $\sigma = I^{-1/2}$. The translations of the convolution are chosen to be regular $3 \times 3 \times 3$ grid with size 2σ .


At test time, similarly to [21] we use voting: we sample ten different samples of size 1024 from 1200 points on each point cloud, apply anisotropic scaling, propagate it through the net and sum the label probability vectors before taking the label with the maximal probability.

We used standard convolution architecture, see Figure 4:

`conv_block(1024, 256, 64) → conv_block(256, 64, 256)`
`→ conv_block(64, 1, 1024) → fully_connected_block,`

where `conv_block(#points in, #points out, #channels)` consists of a convolution layer, batch normalization, Relu activation and a pooling. The fully connected block is a concatenation of a two fully connected layers with dropout after each one.

Robustness to sampling. The inset compares our method with [31, 33] when feeding a trained 1024 point model on sparser test point clouds of size $k = 1024, 512, 256, 128, 64$. The favorable robustness of our method to sub-sampling can be possibly explained by the fact that our extension operator possess approximation power, even with sparse samples, *e.g.*, for smooth shapes, see Figure 2.



Method variants. We evaluate the performance of our algorithm subject to the variation in: the number of points I , the kernel translations $\{y_l\}$, the input tensor f , different bases ℓ_i in (4), the choice of σ , and number of learnable parameters. Points were randomly sampled by the same ratio as in the above (*e.g.* 512 out of 600).

Table 2 presents the results. Using the constant one as input $f = \mathbf{1} \in \mathbb{R}^{I \times 1}$ provides almost comparable results to using the full coordinates of the points $f = (1, x) \in \mathbb{R}^{I \times 4}$. This observation is partially supported by the theoretical analysis shown in section 4 which states that our extension operator applied to the constant one tensor already provides good approximation to the underlying surface and its normal as well as curvature. Using interpolation basis $\{\ell_i\}$ in the extension operator (4), although heavier computationally, does not provide better results. Applying too small σ provides worse classification result. This can be explained by the observation that small σ results in separated Gaussians centered at the points which deteriorates the approximation (X and σ should be related). Interestingly, using a relatively small network size of 1.4M parameters provides comparable classification result.

algorithm	# points	10 models	40 models
Point cloud methods			
pointnet [31]	1024	-	89.2
pointnet++ [33]	1024	-	90.7
deep sets [46]	1000	-	87.1
ECC [37]	1000	90.8	87.4
kd-network [21]	1024	93.3	90.6
kd-network [21]	32k	94.0	91.8
ours	1024	94.9	92.3
Additional input features			
FusionNet (uses mesh structure) [17]	-	93.1	90.8
VRN single(uses mesh structure) [6]	-	-	92.0
OctNet [35]	-	90.9	86.5
VRN ensemble(uses mesh structure) [6]	-	97.1	95.5
MVCNN (uses mesh structure)[32]	-	-	92.0
MVCNN-MultiRes(uses mesh structure)[32]	-	-	93.8
pointnet++ (uses normals) [33]	5K	-	91.9
OCNN (uses normals) [40]	-	-	90.6

Table 1: Shape classification results on the ModelNet40 and ModelNet10 datasets.

Method variations on ModelNet40		
Variation	# points	accuracy
Less points	256	90.8
Less points	512	91.2
Less points	870	92.2
More points	1800	92.3
Interpolation	1024	92.0
Spherical kernel	1024	91.5
Learned translations	1024	91.3
Indicator input	1024	91.2
xyz only input	1024	91.5
Less parameters	1024	92.2
Small sigma	1024	85.5

Table 2: Classification with variations to the PCNN model.

Feature visualizations. Figure 5 visualizes the features learned in the first layer of PCNN on a few shapes from the ModelNet40 dataset. As in the case of images, the features learned on the first layer are mostly edge detectors and directional derivatives. Note that the features are consistent through different sampling and shapes. Figure 6 shows 9 different features learned in the third layer of PCNN. In this layer the features capture more semantically meaningful parts.

5.2 Point cloud segmentation

Our method can also be used for part segmentation: given a point cloud that represents a shape the task is to label each point with a correct part label. We evaluate PCNN performance on ShapeNet part dataset [44]. ShapeNet contains 16,881 shapes from 16 different categories, and total of 50 part labels.

Table 3 compares per-category and mean IoU(%) scores of PCNN with state of the art point cloud methods: PointNet [31], kd-network [21], and 3DCNN (results taken from [31]). Our method outperforms all of these methods. For completeness we also provide results of other methods that use additional shape features or mesh normals as input. Figure 7 depicts several of our segmentation results.

For this task we used standard convolution segmentation architecture, see Figure 4:

```
conv_block(2048, 512, 64) → conv_block(512, 128, 128)
→ conv_block(128, 16, 256) → deconv_block(16, 128, 512)
→ deconv_block(128, 512, 256) → deconv_block(512, 2048, 256)
→ deconv_block(2048, 2048, 256) → conv_block(2048, 2048, 50),
```

where `deconv_block(#points in,#points out,#features)` consists of an upsampling layer followed by a convolution block. In order to provide the last layers with raw features we also add skip-layers connections, see Figure 4(b). This is a common practice in such architectures where fine details are needed at the output layer (*e.g.*, [12]).




Figure 5: Our point cloud convolution is translation invariant and robust to sample size and density: (a) shows feature activations of two kernels (rows) learned by our network’s first convolution layer on different shapes (columns). The features seem consistent across the different models; (b) shows another pair of kernels (rows) on a single model with varying sampling density (from left to right): 10K points, 5K points (random sampling), 1K points (farthest point sampling) and 1K (random sampling). Note that the convolution captures the same geometric properties on all models regardless of the sampling.

We use the data from [31] (2048 uniformly sampled points on each model). As done in [31] we use a single network to predict segmentations for each of the object classes and concatenate a hot-one encoding of the object’s label to the bottleneck feature layer. At test time, we use only the part labels that correspond to the input shape (as in [31, 21]).

5.3 Normal estimation

Estimating normals of a point cloud is a central sub-problem of the 3D reconstruction problem. We cast this problem as supervised regression problem and employ segmentation network with the following changes: the output layer is composed of 3 channels instead of 50 which are then normalized and fed into cosine-loss with the ground truth normals.

We have trained and tested our network on the standard train/test splits of the ModelNet40 dataset (we used the data generator code by [33]). Table 4 compares the mean cosine loss (distance) of PCNN and the normal estimation of [31]



Figure 6: High level features learned by PCNN’s third convolution layer and visualized on the input point cloud. As expected, the features are less geometrical than the first layer’s features (see Figure 5) and seem to capture more semantically meaningful shape parts.

and [33]. Figure 8 depicts normal estimation examples from this challenge.

5.4 Training details, timings and network size

We implemented our method using the TensorFlow library [1] in Python. We used the Adam optimization method with learning rate 0.001 and decay rate 0.7. The models were trained on Nvidia p100 GPUs. Table 5 summarizes running times and network sizes. Our smaller classification network achieves state of the art result (see Table 2, previous to last row) and has only 1.4M parameters with a total model size of 17 MB.

6 Conclusions

This paper describes PCNN: a methodology for defining convolution of functions over point clouds that is efficient, invariant to point cloud order, robust to point sampling and density, and posses translation invariance. The key idea is to translate volumetric convolution to arbitrary point clouds using extension and restriction operators.

Testing PCNN on standard point cloud benchmarks show state of the art results using compact networks. The main limitation of our framework compared to image CNN is the extra computational burden due to the computation of farthest point samples in the network and the need to compute

	input	mean	aero	bag	cap	car	chair	ear-p	guitar	knife	lamp	laptop	motor	mug	pistol	rocket	skate	table
Point clouds																		
Ours	2K pnts	85.1	82.4	80.1	85.5	79.5	90.8	73.2	91.3	86.0	85.0	95.7	73.2	94.8	83.3	51.0	75.0	81.8
PointNet	2K pnts	83.7	83.4	78.7	82.5	74.9	89.6	73	91.5	85.9	80.8	95.3	65.2	93	81.2	57.9	72.8	80.6
kd-network	4K pnts	82.3	80.1	74.6	74.3	70.3	88.6	73.5	90.2	87.2	81	94.9	57.4	86.7	78.1	51.8	69.9	80.3
3DCNN		79.4	75.1	72.8	73.3	70	87.2	63.5	88.4	79.6	74.4	93.9	58.7	91.8	76.4	51.2	65.3	77.1
Additional input																		
SyncSpecCNN	sf	84.7																
Yi	sf	81.4	81	78.4	77.7	75.7	87.6	61.9	92.0	85.4	82.5	95.7	70.6	91.9	85.9	53.1	69.8	75.3
PointNet++	pnts, nors	85.1	82.4	79.0	87.7	77.3	90.8	71.8	91.0	85.9	83.7	95.3	71.6	94.1	81.3	58.7	76.4	82.6
OCNN (+CRF)	nors	85.9	85.5	87.1	84.7	77.0	91.1	85.1	91.9	87.4	83.3	95.4	56.9	96.2	81.6	53.5	74.1	84.4

Table 3: ShapeNet segmentation results by point cloud methods (top) and methods using additional input data (sf - shape features; nors - normals). The methods compared to are: PointNet [31]; kd-network [21]; 3DCNN [31]; SyncSpecCNN [45]; Yi [44]; PointNet++ [33]; OCNN (+CRF refinement) [40].



Figure 7: Results of PCNN on the part segmentation benchmark from [44]

Data	algorithm	# points	error
ModelNet40	PointNet	1024	0.47
	PointNet++	1024	0.29
	ours	1024	0.19

Table 4: Normal estimation in ModelNet40.

the “translating” tensor $q \in \mathbb{R}^{I \times I \times L}$ which is a function of the point cloud X . Still, we believe that a sparse efficient implementation can alleviate this limitation and mark it as a future work. Other venues for future work is to learn kernel translations per channel similarly to [13], and apply the method to data of higher dimension than $d = 3$ which seems to be readily possible. Lastly, we would like to test this framework on different problems and architectures.

7 Acknowledgements

This research was supported in part by the European Research Council (ERC Consolidator Grant, “Lift-Match” 771136), the Israel Science Foundation (Grant No. 1830/17). We would like thank the authors of PointNet [31] for sharing their code and data.

References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean,

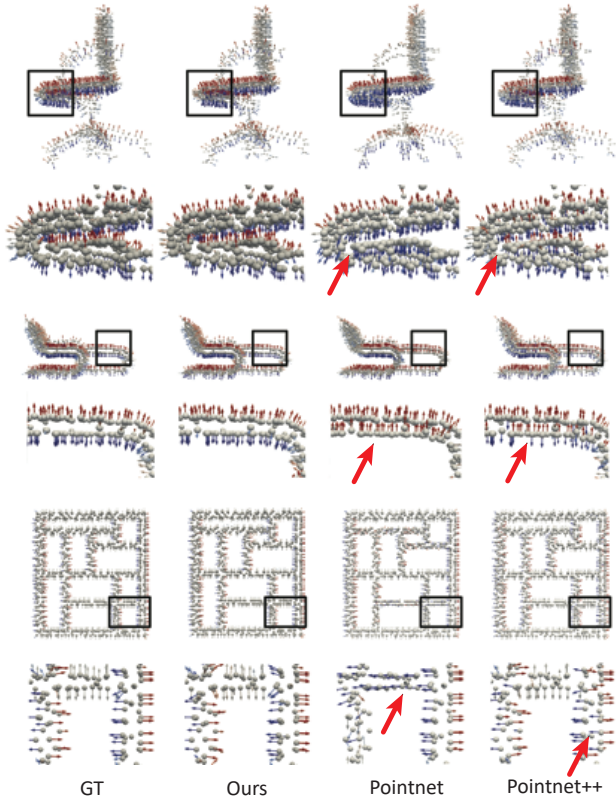


Figure 8: Normal estimation in ModelNet40. We show normal estimation of four models (rows) with blow-ups. Normals are colored by one of their coordinates for better visualization. Note that competing methods sometimes fail to recognize the outward normal direction (examples indicated by red arrows).

	#Param.	Size (mb)	Converge (epochs)	Training (min)	Forward (msec)
Classification					
Large	8.1M	98	250	6	80
Small	1.4M	17	250	5	70
Segmentation					
	5.4M	62	85	37	200

Table 5: Timing and network size. Training time is measured in minutes per epoch.

- M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [2] M. Belkin and P. Niyogi. Towards a theoretical foundation for laplacian-based manifold methods. In *COLT*, volume 3559, pages 486–500. Springer, 2005.
- [3] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- [4] M. Berger, A. Tagliasacchi, L. M. Seversky, P. Alliez, G. Guennebaud, J. A. Levine, A. Sharf, and C. T. Silva. A survey of surface reconstruction from point clouds. In *Computer Graphics Forum*, volume 36, pages 301–329. Wiley Online Library, 2017.
- [5] D. Boscaini, J. Masci, E. Rodolà, and M. M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *NIPS*, 2016.
- [6] A. Brock, T. Lim, J. M. Ritchie, and N. Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016.
- [7] 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, 2017.
- [8] D. S. Broomhead and D. Lowe. Radial basis functions, multi-variable functional interpolation and adaptive networks. Technical report, Royal Signals and Radar Establishment Malvern (United Kingdom), 1988.
- [9] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.
- [10] Y. D. Burago and V. A. Zalgaller. *Geometry III: theory of surfaces*, volume 48. Springer Science & Business Media, 2013.
- [11] J. C. Carr, R. K. Beatson, J. B. Cherrie, T. J. Mitchell, W. R. Fright, B. C. McCallum, and T. R. Evans. Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 67–76. ACM, 2001.
- [12] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, 2016.
- [13] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. *arXiv preprint arXiv:1703.06211*, 2017.
- [14] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, pages 3837–3845, 2016.
- [15] M. P. Do Carmo, G. Fischer, U. Pinkall, and H. Reckziegel. Differential geometry. In *Mathematical Models*, pages 155–180. Springer, 2017.
- [16] P. Guerrero, Y. Kleiman, M. Ovsjanikov, and N. J. Mitra. Pcpnet: Learning local shape properties from raw point clouds. *arXiv preprint arXiv:1710.04954*, 2017.
- [17] V. Hegde and R. Zadeh. Fusionnet: 3d object classification using multiple data representations. *arXiv preprint arXiv:1607.05695*, 2016.
- [18] M. Henaff, J. Bruna, and Y. LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.
- [19] E. Kalogerakis, M. Averkiou, S. Maji, and S. Chaudhuri. 3d shape segmentation with projective convolutional networks. *arXiv preprint arXiv:1612.02808*, 2016.
- [20] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- [21] R. Klokov and V. Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. *arXiv preprint arXiv:1704.01222*, 2017.
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [23] H. Maron, M. Galun, N. Aigerman, M. Trope, N. Dym, E. Yumer, V. G. KIM, and Y. Lipman. Convolutional neural networks on surfaces via seamless toric covers. SIGGRAPH, 2017.
- [24] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks

- on riemannian manifolds. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 37–45, 2015.
- [25] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 922–928. IEEE, 2015.
- [26] F. Monti, D. Boscaini, J. Masci, E. Rodolà, J. Svoboda, and M. M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. *arXiv preprint arXiv:1611.08402*, 2016.
- [27] E. A. Nadaraya. On estimating regression. *Theory of Probability & Its Applications*, 9(1):141–142, 1964.
- [28] M. J. Orr. Recent advances in radial basis function networks. *Institute for Adaptive and Neural Computation*, 1999.
- [29] M. J. Orr et al. Introduction to radial basis function networks, 1996.
- [30] J. Park and I. W. Sandberg. Universal approximation using radial-basis-function networks. *Neural computation*, 3(2):246–257, 1991.
- [31] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *arXiv preprint arXiv:1612.00593*, 2016.
- [32] C. R. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2016.
- [33] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017.
- [34] S. Ravanbakhsh, J. Schneider, and B. Póczos. Deep learning with sets and point clouds. *arXiv preprint arXiv:1611.04500*, 2016.
- [35] G. Riegler, A. O. Ulusoy, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. *arXiv preprint arXiv:1611.05009*, 2016.
- [36] K. Schütt, P.-J. Kindermans, H. E. S. Felix, S. Chmiela, A. Tkatchenko, and K.-R. Müller. Moleculenet: A continuous-filter convolutional neural network for modeling quantum interactions. In *Advances in Neural Information Processing Systems*, pages 992–1002, 2017.
- [37] M. Simonovsky and N. Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. *arXiv preprint arXiv:1704.02901*, 2017.
- [38] A. Sinha, J. Bai, and K. Ramani. Deep learning 3d shape surfaces using geometry images. In *European Conference on Computer Vision*, pages 223–240. Springer, 2016.
- [39] H. Su, S. Maji, E. Kalogerakis, and E. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 945–953, 2015.
- [40] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):72, 2017.
- [41] E. W. Weisstein. Normal sum distribution. 2000.
- [42] H. Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [43] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. 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.
- [44] L. Yi, V. G. Kim, D. Ceylan, I. Shen, M. Yan, H. Su, A. Lu, Q. Huang, A. Sheffer, L. Guibas, et al. A scalable active framework for region annotation in 3d shape collections. *ACM Transactions on Graphics (TOG)*, 35(6):210, 2016.
- [45] L. Yi, H. Su, X. Guo, and L. Guibas. Syncspeccnn: Synchronized spectral cnn for 3d shape segmentation. *arXiv preprint arXiv:1612.00606*, 2016.
- [46] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. Smola. Deep sets. *arXiv preprint arXiv:1703.06114*, 2017.

A Proofs

A.1 Multiplication law for Gaussians

Proposition 1. *Let*

$$\Phi_{\mu,\sigma}(x) = \exp\left(-\frac{\|x - \mu\|^2}{2\sigma^2}\right)$$

and

$$B(\sigma) = \frac{1}{(2\pi\sigma^2)^{\frac{3}{2}}}$$

then

$$\Phi_{\mu_1, \sigma_1} * \Phi_{\mu_2, \sigma_2} = C(\sigma_1, \sigma_2) \cdot \Phi_{\mu, \sigma}$$

where $\mu = \mu_1 + \mu_2$, $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$ and $C(\sigma_1, \sigma_2) = \frac{B(\sigma_1)B(\sigma_2)}{B(\sigma)}$

Proof. It is well known [41] that the convolution of two normal distributions is again a normal distribution:

$$B(\sigma_1)\Phi_{\mu_1, \sigma_1} * B(\sigma_2)\Phi_{\mu_2, \sigma_2} = B(\sqrt{\sigma_1^2 + \sigma_2^2}) \cdot \Phi_{\mu, \sigma}$$

The result above follows from the linearity of the convolution. \square

A.2 Theoretical properties of the extension operator

Proof of theorem 2. Let us first show (24). Denote $g_x(y) = f(y)\Phi_\sigma(|x-y|)$. By Lemma 4 for arbitrary $\epsilon > 0$ there exists $\delta > 0$ so that $d_S(y, x) < \delta$ implies $|g_x(y) - g_x(x)| < \epsilon$ for all $x \in \mathbb{R}^3$. Taking X to be δ -net of S we get that

$$\left| \mathcal{E}_S[f](x) - \mathcal{E}_X \circ \mathcal{R}_X[f](x) \right| \leq \epsilon \frac{\sum_i \text{area}(\Omega_i)}{2\pi\sigma^2} \leq \epsilon \frac{\text{area}(S)}{2\pi\sigma^2}.$$

To show (19) first let $x \notin S$. Then as $\sigma \rightarrow 0$ we have $\max_{y \in S} \Phi_\sigma(|y-x|) \rightarrow 0$ and therefore $\mathcal{E}_S[f](x) \rightarrow 0$. Next consider $x \in S$. It is enough to show that

$$\frac{1}{2\pi\sigma^2} \int_S \Phi_\sigma(|x-y|) da(y) \xrightarrow{\sigma \rightarrow 0} 1.$$

Indeed, let $\epsilon > 0$. Since f is uniformly continuous, take $\delta > 0$ sufficiently small so that $|f(x) - f(x')| < \epsilon$ if $d_S(x, x') < \delta$. Take $\sigma > 0$ sufficiently small so that

$$\frac{1}{2\pi\sigma^2} \int_{S \setminus B(x, \delta)} \Phi_\sigma(|x-y|) da(y) \leq \epsilon,$$

where $B(x, \delta) = \{y \mid |y-x| < \delta\}$ and

$$\left| \frac{1}{2\pi\sigma^2} \int_S \Phi_\sigma(|x-y|) da(y) - 1 \right| \leq \epsilon.$$

Hence

$$\left| \frac{1}{2\pi\sigma^2} \int_{S \cap B(x, \delta)} \Phi_\sigma(|x-y|) da(y) - 1 \right| \leq 2\epsilon.$$

Therefore,

$$\left| \frac{1}{2\pi\sigma^2} \int_S f(y)\Phi_\sigma(|x-y|) da(y) - f(x) \right| \leq 3\epsilon(1 + |f|_\infty).$$

Lastly, to show (21) note that

$$\nabla_x \mathcal{E}_S[\mathbf{1}](x) = -\frac{1}{2\pi\sigma^4} \int_S (x-y)\Phi(|x-y|) da(y).$$

Using an argument from [2] (see Section 4.2) where we take $\sigma^2 = 2t$ in their notation and get convergence to $-\frac{1}{2}\Delta_S x$, where Δ_S is the Laplace-Beltrami operator on surfaces S . To finish the proof remember that [10]

$$-\frac{1}{2}\Delta_S x = -H \cdot n.$$

\square

Lemma 3. Let $S \subset \mathbb{R}^3$ be a compact smooth surface. Then,

$$\lim_{\sigma \rightarrow 0} \frac{1}{2\pi\sigma^2} \int_S \Phi_\sigma(|x-y|) da(y) = 1 \quad (25)$$

Proof. Denote $T_x S$ the tangent plane to S centered at $x \in S$. Let $y = y(u) : T_x S \rightarrow S$ be the local parameterization to S over $T_x S$, where u is the local coordinate at $T_x S$. Since S is smooth and compact we have that $\forall u \in \Upsilon_\delta = T_x S \cap B(x, \delta)$,

$$|y(u) - u| = \mathcal{O}(\delta^2) \quad (26)$$

$$||dy(u)| - 1| = \mathcal{O}(\delta), \quad (27)$$

where $|dy(u)|$ is the pulled-back area element of S [15]. We break the error to $\left| \frac{1}{2\pi\sigma^2} \int_S \Phi_\sigma(|x-y|) da(y) - 1 \right|$

$$\leq \frac{1}{2\pi\sigma^2} \int_{S \setminus \Upsilon_\delta} \Phi_\sigma da + \frac{1}{2\pi\sigma^2} \left| \int_{\Upsilon_\delta} \Phi_\sigma da - \int_{\Upsilon_\delta} \Phi_\sigma du \right| \quad (ii)$$

$$\left| \frac{1}{2\pi\sigma^2} \int_{T_x S} \Phi_\sigma du - 1 \right| + \frac{1}{2\pi\sigma^2} \int_{T_x S \setminus \Upsilon_\delta} \Phi_\sigma du. \quad (iv)$$

First, we note that $(iii) = 0$. Now, take $\delta = \sigma^{1-\tau}$, for some fixed $0 < \tau < 1$, where $\sigma > 0$. Then, $(i) = \mathcal{O}(\sigma)$, $(iv) = \mathcal{O}(\sigma)$. Lastly (ii)

$$\leq \frac{1}{2\pi\sigma^2} \int_{\Upsilon_\delta} \left| \Phi_\sigma(|y(u)|) - \Phi_\sigma(|u|) \right| |dy(u)| du + \frac{1}{2\pi\sigma^2} \int_{\Upsilon_\delta} \Phi_\sigma(|u|) \left| |dy(u)| - 1 \right| du$$

$$\leq \frac{1}{2\pi\sigma^2} \frac{\max_{u \in \Upsilon_\delta} \left| |y(u)| - |u| \right|}{\sigma e^{1/2}} (1 + \mathcal{O}(\delta)) \mathcal{O}(\delta^2) + \mathcal{O}(\delta) = \mathcal{O}(\sigma^{1-4\tau}),$$

where we used Lemma 5 in the last inequality. Taking any $\tau < 1/4$ proves the result. \square

Lemma 4. Let $f \in C(S, \mathbb{R})$, with $S \subset \mathbb{R}^3$ a compact surface. The family of functions $\{g_x\}_{x \in \mathbb{R}^3}$ defined by $g_x(y) = f(y)\Phi_\sigma(|x-y|)$, $y \in S$ is uniformly equicontinuous.

Proof. $|g_x(y) - g_x(y')|$

$$\begin{aligned} &\leq |f(y) - f(y')| \Phi_\sigma(|x - y|) + \\ &|f(y')| |\Phi_\sigma(|x - y|) - \Phi_\sigma(|x - y'|)| \\ &\leq |f(y) - f(y')| + |f|_\infty \frac{|y - y'|}{\sigma e^{1/2}} \end{aligned}$$

where in the last inequality we used $||x - y| - |x - y'|| \leq |y - y'|$ and Lemma 5. Since $f, |\cdot|$ are both uniformly continuous over S (as continuous functions over a compact surface), f is bounded, *i.e.*, $|f|_\infty < \infty$, equicontinuity of $\{g_x\}$ is proved. \square

Lemma 5. *The gaussian satisfies $|\Phi_\sigma(r') - \Phi_\sigma(r)| \leq \frac{(r' - r)}{\sigma e^{1/2}}$ for $0 \leq r < r'$.*

Proof.

$$|\Phi_\sigma(r') - \Phi_\sigma(r)| \leq \int_r^{r'} \frac{t}{\sigma^2} e^{-\frac{t^2}{2\sigma^2}} dt \leq \frac{(r' - r)}{\sigma e^{1/2}},$$

where in the last inequality we used the fact that $te^{-\frac{t^2}{2\sigma^2}} \leq \sigma e^{-1/2}$. \square

Lastly, to justify (6) let us use Theorem 2 and consider $f(x) \equiv 1$,

$$1 \approx \mathcal{E}_S[f](x) \approx c \sum_i \omega_i \Phi(|x - x_i|).$$

Plugging $x = x_{i'}$ we get

$$1 \approx c \sum_i \omega_i \Phi(|x_{i'} - x_i|) = c\tilde{\omega}_{i'} \sum_i \Phi(|x_{i'} - x_i|),$$

where $\tilde{\omega}_{i'}$ is an average of values of ω_i (note that $\Phi(|x_{i'} - x_i|)$ are fast decaying weights away from $x_{i'}$). Hence,

$$c\tilde{\omega}_{i'} \approx \frac{1}{\sum_i \Phi(|x_{i'} - x_i|)}.$$