

Deep Parametric Continuous Convolutional Neural Networks

Shenlong Wang^{1,3,*} Simon Suo^{2,3,*} Wei-Chiu Ma³ Andrei Pokrovsky³ Raquel Urtasun^{1,3}
¹University of Toronto, ²University of Waterloo, ³Uber Advanced Technologies Group
{slwang, suo, weichiu, andrei, urtasun}@uber.com

Abstract

Standard convolutional neural networks assume a grid structured input is available and exploit discrete convolutions as their fundamental building blocks. This limits their applicability to many real-world applications. In this paper we propose Parametric Continuous Convolution, a new learnable operator that operates over non-grid structured data. The key idea is to exploit parameterized kernel functions that span the full continuous vector space. This generalization allows us to learn over arbitrary data structures as long as their support relationship is computable. Our experiments show significant improvement over the state-of-the-art in point cloud segmentation of indoor and outdoor scenes, and lidar motion estimation of driving scenes.

1. Introduction

Discrete convolutions are the most fundamental building block of modern deep learning architectures. Its efficiency and effectiveness relies on the fact that the data appears naturally in a dense grid structure (e.g., 2D grid for images, 3D grid for videos). However, many real world applications such as visual perception from 3D point clouds, mesh registration and non-rigid shape correspondences rely on making statistical predictions from non-grid structured data. Unfortunately, standard convolutional operators cannot be directly applied in these cases.

Multiple approaches have been proposed to handle non-grid structured data. The simplest approach is to voxelize the space to form a grid where standard discrete convolutions can be performed [29, 24]. However, most of the volume is typically empty, and thus this results in both memory inefficiency and wasted computation. Geometric deep learning [3, 15] and graph neural network approaches [25, 16] exploit the graph structure of the data and model the relationship between nodes. Information is then propagated through the graph edges. However, they either have difficulties generalizing well or require strong feature representations as input to perform competitively. End-to-end learning is typically performed via back-propagation through time,

but it is difficult to learn very deep networks due to the memory limitations of modern GPUs.

In contrast to the aforementioned approaches, in this paper we propose a new learnable operator, which we call *parametric continuous convolution*. The key idea is a parameterized kernel function that spans the full continuous vector space. In this way, it can handle arbitrary data structures as long as its support relationship is computable. This is a natural extension since objects in the real-world such as point clouds captured from 3D sensors are distributed unevenly in continuous domain. Based upon this we build a new family of deep neural networks that can be applied on generic non-grid structured data. The proposed networks are both expressive and memory efficient.

We demonstrate the effectiveness of our approach in both semantic labeling and motion estimation of point clouds. Most importantly, we show that very deep networks can be learned over raw point clouds in an end-to-end manner. Our experiments show that the proposed approach outperforms the state-of-the-art by a large margin in both outdoor and indoor 3D point cloud segmentation tasks, as well as lidar motion estimation in driving scenes. Importantly, our outdoor semantic labeling and lidar flow experiments are conducted on a very large scale dataset, containing 223 billion points captured by a 3D sensor mounted on the roof of a self-driving car. To our knowledge, this is 2 orders of magnitude larger than any existing benchmark.

2. Related Work

Deep Learning for 3D Geometry: Deep learning approaches that exploit 3D geometric data have recently become popular in the computer vision community. Early approaches convert the 3D data into a two-dimensional RGB + depth image [17, 10] and exploit conventional convolutional neural networks (CNNs). Unfortunately, this representation does not capture the true geometric relationships between 3D points (i.e. neighboring pixels could be potentially far away geometrically). Another popular approach is to conduct 3D convolutions over volumetric representations [29, 21, 24, 9, 18]. Voxelization is employed to convert point clouds into a 3D grid that encodes the geo-

metric information. These approaches have been popular in medical imaging and indoor scene understanding, where the volume is relatively small. However, typical voxelization approaches sacrifice precision and the 3D volumetric representation is not memory efficient. Sparse convolutions [9] and advanced data structures such as oct-trees [24] have been used to overcome these difficulties. Learning directly over point clouds has only been studied very recently. The pioneer work of PointNet [20], learns an MLP over individual points and aggregates global information using pooling. PointNet++ [22], the follow-up, improves the ability to capture local structures through a multi-scale grouping strategy.

Graph Neural Networks: Graph neural networks (GNNs) [25] are generalizations of neural networks to graph structured data. Early approaches apply neural networks either over the hidden representation of each node or the messages passed between adjacent nodes in the graph, and use back-propagation through time to conduct learning. Gated graph neural networks (GGNNs) [16] exploit gated recurrent units along with modern optimization techniques, resulting in improved performance. In [23], GGNNs are applied to point cloud segmentation, achieving significant improvements over the state-of-the-art. One of the major difficulties of graph neural networks is that propagation is conducted in a synchronous manner and thus it is hard to scale up to graphs with millions of nodes. Inference in graphical models as well as recurrent neural networks can be seen as special cases of graph neural networks.

Graph Convolution Networks: An alternative formulation is to learn convolution operations over graphs. These methods can be categorized into spectral and spatial approaches depending on which domain the convolutions are applied to. For spectral methods, convolutions are converted to multiplication by computing the graph Laplacian in Fourier domain [4, 2, 30]. Parameterized spectral filters can be incorporated to reduce overfitting [4]. These methods are not feasible for large scale data due to the expensive computation, since there is no FFT-like trick over generic graph. Spatial approaches directly propagate information along the node neighborhoods in the graph. This can be implemented either through low-order approximation of spectral filtering [6, 15, 7], or diffusion in a support domain [19, 2, 27, 30, 26]. Our approach generalizes spatial approaches in two ways: first, we use more expressive convolutional kernel functions; second, the output of the convolution could be any point in the whole continuous domain.

Other Approaches: Edge-conditioned filter networks [27] use a weighting network to communicate between adjacent nodes on the graph [13] conditioned on edge labels, which is primarily formulated as relative point locations. In

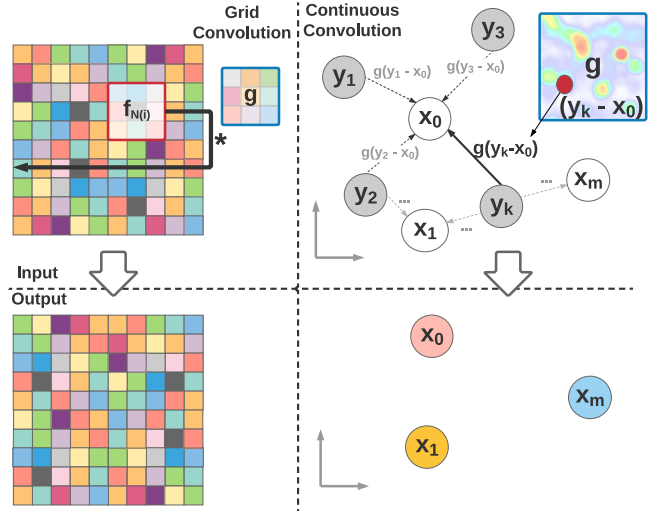


Figure 1: Unlike grid convolution, parametric continuous convolution uses kernel functions that are defined for arbitrary points in the continuous support domain. As a result, it is possible to output features at points not seen in the input.

contrast, our approach is not constrained to a fixed graph structure, and has the flexibility to output features at arbitrary points over the continuous domain. In a concurrent work, [26] uses similar parametric function form $f(\mathbf{x}_i - \mathbf{x}_j)$ to aggregate information between points. However, they only use shallow isotropic gaussian kernels to represent the weights, while we use expressive deep networks to parameterize the continuous filters.

3. Deep Parametric Continuous CNNs

3.1. Parametric Continuous Convolutions

Standard CNNs use discrete convolutions (i.e., convolutions defined over discrete domain) as basic operations.

$$h[n] = (f * g)[n] = \sum_{m=-M}^M f[n-m]g[m]$$

where $f : \mathcal{G} \rightarrow \mathbb{R}$ and $g : \mathcal{S} \rightarrow \mathbb{R}$ are functions defined over the support domain of finite integer set: $\mathcal{G} = \mathcal{Z}^D$ and $\mathcal{S} = \{-M, -M + 1, \dots, M - 1, M\}^D$ respectively.

In contrast, continuous convolutions can be defined as

$$h(\mathbf{x}) = (f * g)(\mathbf{x}) = \int_{-\infty}^{\infty} f(\mathbf{y})g(\mathbf{x} - \mathbf{y})d\mathbf{y} \quad (1)$$

where both the kernel $g : \mathcal{S} \rightarrow \mathbb{R}$ and the feature $f : \mathcal{G} \rightarrow \mathbb{R}$ are defined as continuous functions over the support domain $\mathcal{G} = \mathbb{R}^D$ and $\mathcal{S} = \mathbb{R}^D$ respectively.

Continuous convolutions require the integration in Eq. (1) to be analytically tractable. Unfortunately, this is not

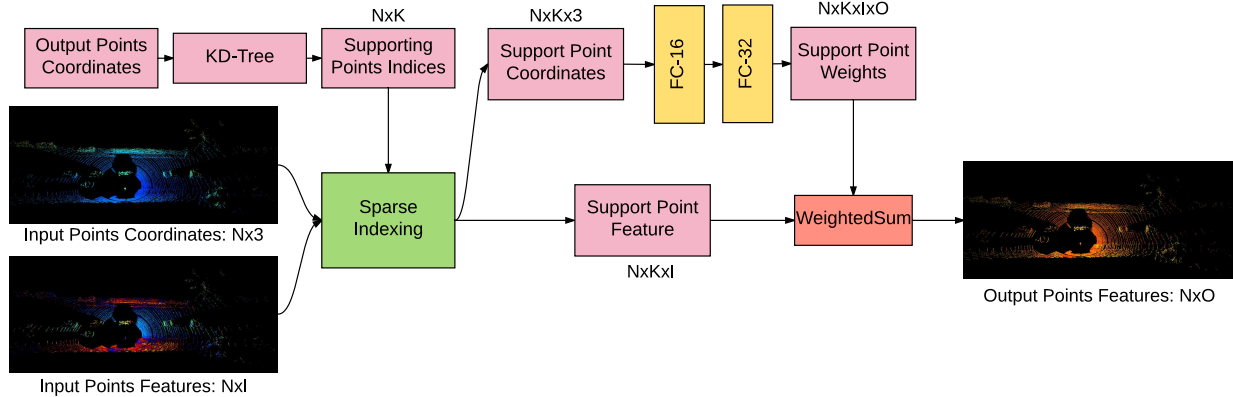


Figure 2: Detailed Computation Block for the Parametric Continuous Convolution Layer.

possible for real-world applications, where the input features are complicated and non-parametric, and the observations are sparse points sampled over the continuous domain.

Motivated by monte-carlo integration [5] we derive our continuous convolution operator. In particular, given continuous functions f and g with a finite number of input points \mathbf{y}_i sampled from the domain, the convolution at an arbitrary point \mathbf{x} can be approximated as:

$$h(\mathbf{x}) = \int_{-\infty}^{\infty} f(\mathbf{y})g(\mathbf{x} - \mathbf{y})d\mathbf{y} \approx \sum_i^N \frac{1}{N} f(\mathbf{y}_i)g(\mathbf{x} - \mathbf{y}_i)$$

The next challenge we need to solve is constructing the continuous convolutional kernel function g . Conventional 2D and 3D discrete convolution kernels are parameterized in a way that each point in the support domain is assigned a value (*i.e.* the kernel weight). Such a parameterization is infeasible for continuous convolutions, since the kernel function g is defined over an infinite number of points (*i.e.*, has infinite support). Instead, in this paper we propose to use parametric continuous functions to model g . We name our approach *Parametric Continuous Convolutions*. In particular, we use a multi-layer perceptron (MLP) as the approximator. With reference to the universal approximation theorem of [12], MLPs are expressive and capable of approximating continuous functions over \mathbb{R}^n . Thus we define:

$$g(\mathbf{z}; \theta) = MLP(\mathbf{z}; \theta)$$

The kernel function $g(\mathbf{z}; \theta) : \mathbb{R}^D \rightarrow \mathbb{R}$ spans the full continuous support domain while remaining parametrizable by a finite number of parameters. Note that other choices such as polynomials are possible, however low-order polynomials are not expressive, whereas learning high-order polynomials can be numerically unstable for back-propagation.

3.2. From Convolutions to Deep Networks

In this section, we first design a new convolution layer based on the parametric continuous convolutions derived in

the previous subsection. We then propose a deep learning architecture using this new convolution layer.

Parametric Continuous Convolution Layer: Note that, unlike standard discrete convolutions which are conducted over the same point set, the input and output points of our parametric continuous convolution layer can be different. This is important for many practical applications, where we want to make dense predictions based on partial observations. Furthermore, this allow us to abstract information from redundant input points (*i.e.*, pooling). As a consequence, the input of each convolution layer contains three parts: the input feature vector $\mathcal{F} = \{\mathbf{f}_{in,j} \in \mathbb{R}^F\}$, the associated locations in the support domain $\mathcal{S} = \{\mathbf{y}_j\}$, as well as the output domain locations $\mathcal{O} = \{\mathbf{x}_i\}$. For each layer, we first evaluate the kernel function $g_{d,k}(\mathbf{y}_i - \mathbf{x}_j; \theta)$ for all $\mathbf{x}_j \in \mathcal{S}$ and all $\mathbf{y}_i \in \mathcal{O}$, given the parameters θ . Each element of the output feature vector is then computed as:

$$h_{k,i} = \sum_d^F \sum_j^N g_{d,k}(\mathbf{y}_i - \mathbf{x}_j) f_{d,j}$$

Let N be the number of input points, M be the number of output points, and D the dimensionality of the support domain. Let F and O be predefined input and output feature dimensions respectively. Note that these are hyperparameters of the continuous convolution layer analogous to input and output feature dimensions in standard grid convolution layers. Fig. 1 depicts our parametric continuous convolutions in comparison with conventional grid convolution. Two major differences are highlighted: 1) the kernel function is continuous given the relative location in support domain; 2) the input/output points could be any points in the continuous domain as well and can be different.

Deep Parametric Continuous CNNs: Using the parametric continuous convolution layers as building blocks, we

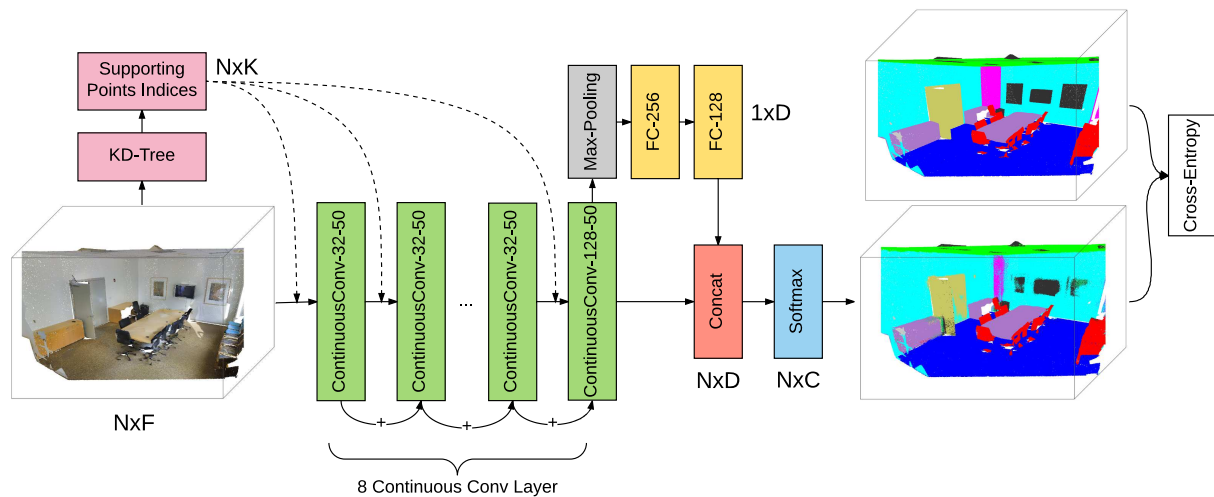


Figure 3: Architecture of the Deep Parametric Continuous CNNs for Semantic Labeling Task.




Figure 4: Semantic Segmentation Results on Stanford Indoor3D Dataset

can construct a new family of deep networks which operates on unstructured data defined in a topological group under addition. In the following discussions, we will focus on multi-dimensional euclidean space, and note that this is a special case. The network takes the input features and

their associated positions in the support domain as input. Then the hidden representations are generated from successive parametric continuous convolution layers. Following standard CNN architectures, we can add batch normalization, non-linearities and residual connections between lay-

ers. Pooling can also be employed over the support domain to aggregate information. In practice, we find adding residual connection between parametric continuous convolution layers is critical to help convergence. Please refer to Fig. 2 for an example of the computation graph of a single layer, and to Fig. 3 for an example of the network architecture employed for our indoor semantic segmentation task.

Learning: All of our building blocks are differentiable, thus our networks can be learned through back-prop:

$$\frac{\partial h}{\partial \theta} = \frac{\partial h}{\partial g} \cdot \frac{\partial g}{\partial \theta} = \sum_d^F \sum_j^N f_{d,j} \cdot \frac{\partial g}{\partial \theta}$$

3.3. Discussions

Locality Enforcing Continuous Convolution: Standard grid convolution are computed over a limited kernel size M to keep locality. Similarly, locality can be enforced in our parametric continuous convolutions by constraining the influence of the function g to points close to \mathbf{x} , *i.e.*,

$$g(\mathbf{z}) = MLP(\mathbf{z})w(\mathbf{z})$$

where $w(\cdot)$ is a modulating window function. This can be achieved in differently. First, we can constrain the cardinality of its local support domain and only keep non-zero kernel values for its K -nearest neighbors: $w(\mathbf{z}) = \mathbf{1}_{\mathbf{z} \in KNN(\mathcal{S}, \mathbf{x})}$. Alternatively we can keep non-zero kernel values for points within a fixed radius r : $w(\mathbf{z}) = \mathbf{1}_{\|\mathbf{z}\|_2 < r}$.

Efficient Continuous Convolution: For each continuous convolution layer, the kernel function is evaluated $N \times |\mathcal{S}| \times F \times O$ times, where $|\mathcal{S}|$ is the cardinality of the support domain, and the intermediate weight tensor is stored for back-propagation. This is expensive in practice, especially when both the number of points and the feature dimension are large. With the locality enforcing formulation, we can constrain the cardinality of \mathcal{S} . Furthermore, motivated by the idea of separable filters, we use the fact that this computation can be factorized if the kernel function value across different output dimensionality is shared. That is to say, we can decompose the weight tensor $W \in \mathbb{R}^{N \times |\mathcal{S}| \times F \times O}$ into two tensors $W_1 = \mathbb{R}^{F \times O}$ and $W_2 = \mathbb{R}^{N \times |\mathcal{S}| \times O}$, where W_1 is a linear weight matrix and W_2 is evaluated through the MLP. With this optimization, only $N \times |\mathcal{S}| \times O$ kernel evaluations need to be computed and stored. Lastly, in inference stage, through merging the operations of batchnorm and fc layer in MLP, 3x speed boosting can be achieved.

Special Cases: Many previous convolutional layers are special cases of our approach. For instance, if the points are sampled over the finite 2D grid we recover conventional 2D

convolutions. If the support domain is defined as concatenation of the spatial vector and feature vector with a gaussian kernel $g(\cdot)$, we recover the bilateral filter. If the support domain is defined as the neighboring vertices of a node we recover the first-order spatial graph convolution [15].

4. Experimental Evaluation

We demonstrate the effectiveness of our approach in the tasks of semantic labeling and motion estimation of 3D point clouds, and show state-of-the-art performance. We conduct point-wise semantic labeling experiments over two datasets: a very large-scale outdoor lidar semantic segmentation dataset that we collected and labeled in house and a large indoor semantic labeling dataset. To our knowledge, these are the largest real-world outdoor and indoor datasets that are available for this task. The datasets are fully labeled and contain 137 billion and 629 million points respectively. The lidar flow experiment is also conducted on this dataset with ground-truth 3D motion label for each point.

4.1. Semantic Segmentation of Indoor Scenes

Dataset: We use the Stanford large-scale 3D indoor scene dataset [1] and follow the training and testing procedure used in [28]. We report the same metrics, *i.e.*, mean-IOU, mean class accuracy (TP / (TP + FN)) and class-wise IOU. The input is six dimensional and is composed of the xyz coordinates and RGB color intensity. Each point is labeled with one of 13 classes shown in Tab. 1.

Competing Algorithms: We compare our approach to PointNet [20] and SegCloud [28]. We evaluate the proposed end-to-end continuous convnet with eight continuous convolution layers (*Ours PCCN*). The kernels are defined over the continuous support domain of 3D Euclidean space. Each intermediate layer except the last has 32 dimensional hidden features followed by batchnorm and ReLU nonlinearity. The dimension of the last layer is 128. We observe that the distribution of semantic labels within a room is highly correlated with the room type (*e.g.* office, hallway, conference room, *etc.*). Motivated by this, we apply max pooling over all the points in the last layer to obtain a global feature, which is then concatenated to the output feature of each points in the last layer, resulting in a 256 dimensional feature. A fully connected layer with softmax activation is used to produce the final logits. Our network is trained end-to-end with cross entropy loss, using Adam optimizer.

Results: As shown in Tab. 1 our approach outperforms the state-of-the-art by 9.3% mIOU and 9.6% mACC. Fig. 4 shows qualitative results. Despite the diversity of geometric structures, our approach works very well. Confusion mainly occurs between columns vs walls and window vs bookcase.




Figure 5: Semantic Segmentation Results on Driving Scene Dataset; Colored: correct prediction; white: wrong prediction.




Figure 6: Semantic Labeling on KITTI Dataset without Retraining

It is also worth noting that our approach captures visual information encoded in RGB channels. The last row shows two failure cases. In the first one, the door in the washroom is labeled as clutter whereas our algorithm thinks is door. In the second one, the board on the right has a window-like texture, which makes the algorithm predict the wrong label.

4.2. Semantic Segmentation of Driving Scenes

Dataset: We first conduct experiments on the task of point cloud segmentation in the context of autonomous driving.

Each point cloud is produced by a full sweep of a roof-mounted Velodyne-64 lidar sensor driving in several cities in North America. The dataset is composed of snippets each having 300 consecutive frames. The training and validation set contains 11,337 snippets in total while the test set contains 1,644 snippets. We report metrics on a subset of the test set which is generated by sampling 10 frames from each snippet to avoid bias brought due to scenes where the ego-car is static (e.g., when waiting at a traffic light). Each point is labeled with one of seven classes defined in Tab. 2. We

Method	mIOU	mAcc	ceiling	floor	wall	beam	column	window	door	chair	table	bookcase	sofa	board	clutter
PointNet [20]	41.09	48.98	88.80	97.33	69.80	0.05	3.92	46.26	10.76	52.61	58.93	40.28	5.85	26.38	33.22
3D-FCN-TI [28]	47.46	54.91	90.17	96.48	70.16	0.00	11.40	33.36	21.12	76.12	70.07	57.89	37.46	11.16	41.61
SEGCloud [28]	48.92	57.35	90.06	96.05	69.86	0.00	18.37	38.35	23.12	75.89	70.40	58.42	40.88	12.96	41.60
Ours PCCN	58.27	67.01	92.26	96.20	75.89	0.27	5.98	69.49	63.45	66.87	65.63	47.28	68.91	59.10	46.22

Table 1: Semantic Segmentation Results on Stanford Large-Scale 3D Indoor Scene Dataset

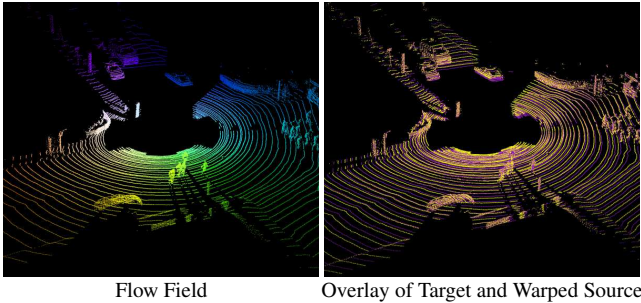


Figure 7: Right: purple shows target frame, yellow shows source frame warped to target frame using ground truth flow

adopt mean intersection-over-union (meanIOU) and point-wise accuracy (pointAcc) as our evaluation metrics.

Baselines: We compare our approach to the point cloud segmentation network (*PointNet*) [20] and a 3D fully convolutional network (*3D-FCN*) conducted over a 3D occupancy grid. We use a resolution of 0.2m for each voxel over a 160m \times 80m \times 6.4m range. This results in an occupancy grid encoded as a tensor of size 800 \times 400 \times 32. We define a voxel to be occupied if it contains at least one point. We use ResNet-50 as the backbone and replace the last average pooling and fully connected layer with two fully convolutional layers and a trilinear upsampling layer to obtain dense voxel predictions. The model is trained from scratch with the Adam optimizer [14] to minimize the class-reweighted cross-entropy loss. Finally, the voxel-wise predictions are mapped back to the original points and metrics are computed over points. We adapted the open-sourced PointNet model onto our dataset and trained from scratch. The architecture and loss function remain the same with the original paper, except that we removed the point rotation layer since it negatively impacts validation performance on this dataset.

Our Approaches: We evaluate two versions of our approach. Our first instance conducts continuous convolutions directly over the raw xyz-intensity lidar points (*Ours PCCN*). Our second version (*Ours 3D-FCN+PCCN*) performs continuous convolutions over the features extracted from *3D-FCN*. *Ours PCCN* has 16 continuous conv layers with residual connections, batchnorm and ReLU nonlinearities. We use the spatial support in \mathbb{R}^3 to define our kernel. We train the network with point-wise cross-

entropy loss and Adam [14] optimizer. In contrast, *Ours 3D-FCN+PCCN* model has 7 residual continuous convolutional layers on top of the trained *3D-FCN* model and performs end-to-end fine-tuning using Adam optimizer.

Results: As shown in Tab. 2, by exploiting sophisticated feature via 3D convolutions, *3D-FCN+PCCN* results in the best performance. Fig. 5 shows qualitative comparison between models. As shown in the figure, all models produce good results. Performance differences often result from ambiguous regions. In particular, we can see that the 3D-FCN model oversegments the scene: it mislabels a background pole as vehicle (red above egocar), nearby spurious points as bicyclist (green above egocar), and a wall as pedestrian (purple near left edge). This is reflected in the confidence map (as bright regions). We observe a significant improvement in our *3D-CNN + PCCN* model, with all of the above corrected with high confidence. For more results and videos please refer to the supplementary material.

Model Sizes: We also compare the model sizes of the competing algorithms in Tab. 2. In comparison to the 3D-FCN approach, the end-to-end continuous convolution network’s model size is eight times smaller, while achieving comparable results. And the 3D-FCN+PCCN is just 0.01MB larger than 3D-FCN, but the performance is improved by a large margin in terms of mean IOU.

Complexity and Runtime We benchmark the proposed model’s runtime over a GTX 1080 Ti GPU and Xeon E5-2687W CPU with 32 GB Memory. The forward pass of a 8-layer PCCN model (32 feature dim in each layer with 50 neighbours) takes 33ms. The KD-Tree neighbour search takes 28 ms. The end-to-end computation takes 61ms. The number of operations of each layer is 1.32GFLOPs.

Generalization: To demonstrate the generalization ability of our approach, we evaluate our model, trained with only North American scenes, on the KITTI dataset [8], which was captured in Europe. As shown in Fig. 6, the model achieves good results, with well segmented dynamic objects, such as vehicles and pedestrians.

4.3. Lidar Flow

Dataset: We also validate our proposed method over the task of lidar based motion estimation, referred to as lidar

Method	pACC	mIOU	vehicle	bicyclist	pedestrian	motorcycle	animal	background	road	params size
PointNet [20]	91.96	38.05	76.73	2.85	6.62	8.02	0.0	89.83	91.96	20.34MB
3D-FCN [11]	94.31	49.28	86.74	22.30	38.26	17.22	0.98	86.91	92.56	74.66MB
Ours PCCN	94.56	46.35	86.62	8.31	41.84	7.24	0.00	87.27	93.20	9.34MB
Ours 3D-FCN+PCCN	95.45	58.06	91.83	40.23	47.74	42.91	1.25	89.27	93.18	74.67MB

Table 2: Semantic Segmentation Results on Driving Scenes Dataset




Figure 8: Lidar Flow Results on Driving Scene Dataset

flow. In this task, the input is two consecutive frames of lidar sweep. The goal is to estimate the 3D motion field for each point in the first frame, to undo both ego-motion and the motion of dynamic objects. The ground-truth ego-motion is computed through a comprehensive filter that takes GPS, IMU as well as ICP based lidar alignment against pre-scanned 3D geometry of the scene as input. And the ground-truth 6DOF dynamics object motion is estimated from the temporal coherent 3D object tracklet, labeled by in-house annotators. Combining both we are able to get the ground-truth motion field. Fig. 7 shows the colormapped flow field and the overlay between two frames after undoing per-point motion. This task is crucial for many applications, such as multi-rigid transform alignment, object tracking, global pose estimation, *etc.* The training and validation set contains 11,337 snippets while the test set contains 1,644 snippets. We use 110k frame pairs for training and validation, and 16440 frame pairs for testing. End-point error, and outlier percentage at 10 cm and 20 cm are used as metric.

Competing Algorithms: We compare against the 3D-FCN baseline using the same architecture and volumetric representation as used in Sec. 4.2. We also adopt a similar 3D-FCN + PCCN architecture with 7 residual continuous convolution layers added as a polishing network. In this task, we remove the ReLU nonlinearity and supervise the PCCN layers with MSE loss at every layer. The training objective function is mean square error loss between the ground-truth flow vector and the prediction.

Method	EPE (cm)	Outlier% ₁₀	Outlier% ₂₀
3D-FCN	8.161	25.92%	7.12 %
Ours 3D-FCN+PCCN	7.810	19.84%	5.97%

Table 3: Lidar Flow Results on Driving Scenes Dataset

Results: Tab. 3 reports the quantitative results. As shown in the table, our 3D-FCN+PCCN model outperforms the 3D-FCN by 0.351cm in end-point error and our method reduces approximately 20% of the outliers. Fig. 8 shows sample flow predictions compared with ground truth labels. As shown in the figure, our algorithm is able to capture both global motion of the ego-car including self rotation, and the motion of each dynamic objects in the scene. For more results please refer to our supplementary material.

5. Conclusions

We have presented a new learnable convolution layer that operates over non-grid structured data. Our convolution kernel function is parameterized by multi-layer perceptrons and spans the full continuous domain. This allows us to design a new deep learning architecture that can be applied to arbitrary structured data, as long as the support relationships between elements are computable. We validate the performance on point cloud segmentation and motion estimation tasks, over very large-scale datasets with up to 200 billion points. The proposed network achieves state-of-the-art performance on all the tasks and datasets.

References

- [1] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese. 3d semantic parsing of large-scale indoor spaces. In *CVPR*, 2016. 5
- [2] D. Boscaini, J. Masci, E. Rodolà, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *NIPS*, 2016. 2
- [3] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 2017. 1
- [4] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. *ICLR*, 2014. 2
- [5] R. E. Caffisch. Monte carlo and quasi-monte carlo methods. *Acta numerica*, 1998. 3
- [6] M. Defferrard, X. Bresson, and P. Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *NIPS*, 2016. 2
- [7] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. In *NIPS*, 2015. 2
- [8] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. 7
- [9] B. Graham and L. van der Maaten. Submanifold sparse convolutional networks. *arXiv*, 2017. 1, 2
- [10] S. Gupta, R. Girshick, P. Arbeláez, and J. Malik. Learning rich features from rgb-d images for object detection and segmentation. In *ECCV*, 2014. 1
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 8
- [12] K. Hornik. Approximation capabilities of multilayer feed-forward networks. *Neural networks*, 1991. 3
- [13] X. Jia, B. De Brabandere, T. Tuytelaars, and L. V. Gool. Dynamic filter networks. In *NIPS*. 2016. 2
- [14] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 7
- [15] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv*, 2016. 1, 2, 5
- [16] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel. Gated graph sequence neural networks. *arXiv*, 2015. 1, 2
- [17] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 1
- [18] D. Maturana and S. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *IROS*, 2015. 1
- [19] 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. *CVPR*, 2017. 2
- [20] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. 2016. 2, 5, 7, 8
- [21] C. R. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. J. Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *CVPR*, 2016. 1
- [22] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. 2017. 2
- [23] X. Qi, R. Liao, J. Jia, S. Fidler, and R. Urtasun. 3d graph neural networks for rgbd semantic segmentation. In *CVPR*, 2017. 2
- [24] G. Riegler, A. O. Ulusoy, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. 2017. 1, 2
- [25] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *TNN*, 2009. 1, 2
- [26] K. T. Schütt, P. Kindermans, H. Sauceda, S. Chmiela, A. Tkatchenko, and K. Müller. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. *arXiv*, 2017. 2
- [27] M. Simonovsky and N. Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. *CVPR*, 2017. 2
- [28] L. P. Tchapmi, C. B. Choy, I. Armeni, J. Gwak, and S. Savarese. Segcloud: Semantic segmentation of 3d point clouds. *arXiv preprint arXiv:1710.07563*, 2017. 5, 7
- [29] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 1
- [30] L. Yi, H. Su, X. Guo, and L. Guibas. Syncspecnn: Synchronized spectral cnn for 3d shape segmentation. *CVPR*, 2017. 2