

Fusion-Aware Point Convolution for Online Semantic 3D Scene Segmentation

Jiazhao Zhang^{1,*} Chenyang Zhu^{1,*} Lintao Zheng¹ Kai Xu^{1,2†}
¹National University of Defense Technology ²SpeedBot Robotics Ltd.

Abstract

Online semantic 3D segmentation in company with real-time RGB-D reconstruction poses special challenges such as how to perform 3D convolution directly over the progressively fused 3D geometric data, and how to smartly fuse information from frame to frame. We propose a novel fusion-aware 3D point convolution which operates directly on the geometric surface being reconstructed and exploits effectively the inter-frame correlation for high quality 3D feature learning. This is enabled by a dedicated dynamic data structure which organizes the online acquired point cloud with global-local trees. Globally, we compile the online reconstructed 3D points into an incrementally growing coordinate interval tree, enabling fast point insertion and neighborhood query. Locally, we maintain the neighborhood information for each point using an octree whose construction benefits from the fast query of the global tree. Both levels of trees update dynamically and help the 3D convolution effectively exploits the temporal coherence for effective information fusion across RGB-D frames. Through evaluation on public benchmark datasets, we show that our method achieves the state-of-the-art accuracy of semantic segmentation with online RGB-D fusion in 10 FPS.

1. Introduction

Semantic segmentation of 3D scenes is a fundamental task in 3D vision. The recent state-of-the-art methods mostly apply deep learning on either 3D geometric data solely [25] or the fusion of 2D and 3D data [20]. These approaches, however, are usually offline, working with an already reconstructed 3D scene geometry [5, 14]. Online scene understanding associated with real-time RGB-D reconstruction [13, 22], on the other hand, is deemed to be more appealing due to the potential applications in robot and AR. Technically, online analysis can also fully exploit the spatial-temporal information during RGB-D fusion.

For the task of semantic scene segmentation in company with RGB-D fusion, deep-learning-based approaches commonly adopt the *frame feature fusion paradigm*. Such meth-

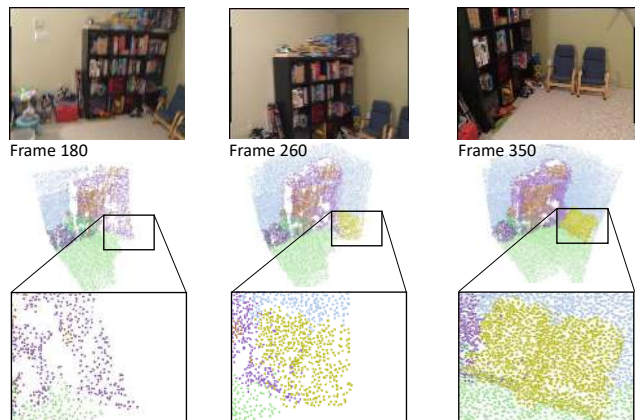


Figure 1: We present fusion-aware 3D point convolution which operates directly over the progressively acquired and online reconstructed scene surface. We show the point-wise labeling is being gradually improved (the chairs are recognized) as more and more frames (first row) are fused in.

ods first perform 2D convolution in the individual RGB-D frames and then fuse the extracted 2D features across consecutive frames. Previous works conduct such feature fusion through either max-pooling operation [14] or Bayesian probability updating [20]. We advocate the adoption of *direct convolution over 3D surfaces* for frame feature fusion. 3D convolution on surfaces learns features of the intrinsic structure of the geometric surfaces [2] that cannot be well-captured by view-based convolution and fusion. During online RGB-D fusion, however, the scene geometry changes progressively with the incremental scanning and reconstruction. It is difficult to perform 3D convolution directly over the time-varying geometry. Besides, to attain a powerful 3D feature learning, special designs are needed to exploit the temporal correlation between adjacent frames.

In this work, we argue that a fast and powerful 3D convolution for online segmentation necessitates an efficient and versatile in-memory organization of dynamic 3D geometric data. To this end, we propose a tree-based global-local dynamic data structure to enable efficient data maintenance and 3D convolution of time-varying geometry. Globally, we organize the online fused 3D points with an incrementally growing coordinate interval tree, which enables fast point

*Joint first authors

†Corresponding author: kevin.kai.xu@gmail.com

insertion and neighborhood query. Locally, we maintain the neighborhood information for each point using an octree whose dynamic update benefits from the fast query of the global tree. The local octrees facilitate efficient point-wise 3D convolution directly over the reconstructed geometric surfaces. Both levels of trees update dynamically along with the online reconstruction.

The dynamic maintenance of the two-level trees supports 3D point convolution with feature fusion across RGB-D frames, leading to so-called *fusion-aware point convolution*. *First*, point correspondence between consecutive frames can be easily retrieved from the global tree, so that both the 2D and 3D features of a point can be efficiently aggregated from frame to frame when the point is observed by multiple frames. *Second*, with the help of per-point octrees, we realize adaptive convolution kernels at each point through weighting its neighboring points based on approximate geodesic distance. This allows a progressive improvement of labeling accuracy across frames.

Through extensive evaluation on the public benchmark datasets, we demonstrate that our method performs online 3D scene semantic segmentation at interactive frame-rate (10 FPS and even higher for key-frame-based processing) while achieving high accuracy outperforming the state-of-the-art offline methods. In particular, the accuracy achieves the top-ranking in the ScanNet benchmark, outperforming many existing approaches including both online and offline ones. Our main contributions include:

- A tree-based global-local dynamic data structure enabling efficient and powerful 3D convolution on time-varying 3D geometry.
- A fusion-aware point convolution which exploits inter-frame correlation for quality 3D feature learning.
- An interactive system implementing our online segmentation with real-time RGB-D reconstruction.

2. Related work

3D scene segmentation. Scene segmentation is a long-standing problem in computer vision. Here, we only review offline approaches handling 3D geometric data obtained either by RGB-D fusion or LiDAR acquisition. For fusion-based 3D reconstruction, many works [16, 18] show that the 2D labeling of the RGB-D frames can be incorporated into the volumetric or surfel map, resulting in stable 3D labeling. The labeling can be further improved with MRF or CRF inference over the 3D map. These works enjoy the advances in image-based CNN for 2D segmentation. Taking the advantage of direct 3D geometric feature learning, 3D deep learning approaches become increasingly popular, where an efficient 3D convolution operation is the key. In these methods, 3D labeling is attained with CNNs operating

directly on point clouds [25] or their voxelization [7]. Several other approaches conduct object detection over 3D reconstruction and then predict a segmentation mask for each detection, leading to instance segmentation [14, 35, 40].

Online scene segmentation. Apart from the majority of offline batch methods, online and incremental mapping and labeling starts to gain renewed interest lately due to the big success of multi-view deep learning [24, 32]. Since the early attempts on 3D semantic mapping from RGB-D sequences [10, 28, 30], a notable recent work of such kind is SemanticFusion [20]. It performs CNN-based 2D labeling for individual RGB-D frames and then probabilistically fuses the 2D predictions into a semantic 3D map. Instead of fusing prediction results, some methods [5, 14] adopt feature map fusion based on max-pooling operation which is more deep learning friendly. In our method, we advocate the use of 3D convolution to aggregate 2D features where a major challenge is how to handle time-varying 3D geometric data. The DA-RNN method [38] aggregates frame features using recurrent neural networks with dynamically routed connections. It smartly utilizes the data association from SLAM to connect recurrent units on the fly. Our method, on the other hand, pursues effective direct 3D convolution through exploiting the data association between frames.

Point cloud convolution. Since 3D convolution is naturally performed on 3D Euclidean grids, early practice opts to first convert 3D point clouds to 2D images [32] or 3D volumes [19] and then perform Euclidean convolution. For the task of semantic segmentation, most approaches choose to extract features in 2D and then perform segmentation in 3D based on the 2D features [18]. Volumetric convolution is limited by resolution due to computational cost, which can be relieved with efficient data structure [15, 27]. These acceleration, however, cannot handle dynamically changing point clouds like ours. Atzmon et al. [1] propose a unique volume-based point convolution which consists of two operators, extension and restriction, mapping point cloud functions to volumetric ones and vice-versa. Point cloud convolution is defined by the extension and restriction of volumetric convolution against the point cloud.

Since the pioneering work of PointNet [25], there have been many works focusing on direct convolution on 3D point clouds. Existing works aim either to improve the neighborhood structure [7, 17, 26, 31, 33] or to enhance the convolutional filters [11, 29, 36, 37, 39]. These methods are designed to process over fixed neighborhood on static point clouds. We design a new point convolution for time-varying geometric data with dedicated designs targeting both aspects. First, we maintain and update a surface-aware neighborhood structure based on a tree-based dynamic data structure. Second, we learn adaptive convolutional filters via exploiting the temporal coherence between consecutive frames.

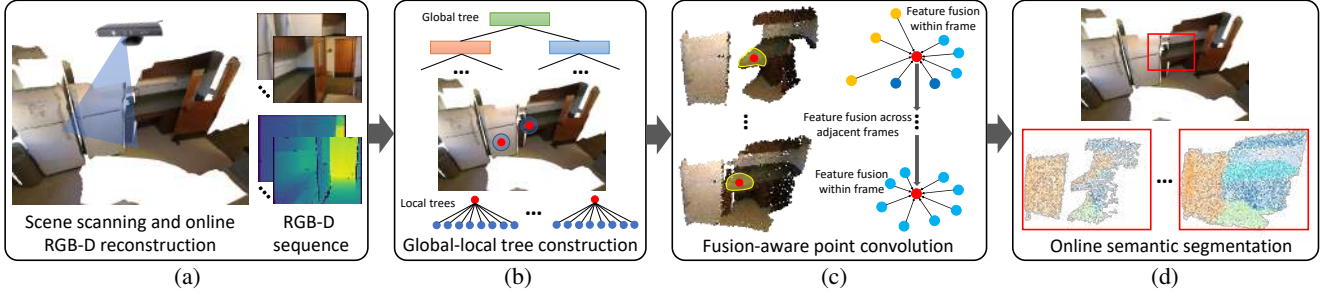


Figure 2: An overview of our pipeline. The input to our method is an online acquired RGB-D sequence being reconstructed in real-time (a). Based on the online reconstruction, we construct global-local trees to maintain a global spatial organization of the reconstructed point cloud as well as per-point local neighborhood (b). The dynamic data structure supports fusion-aware point convolution encompassing intra-frame and inter-frame feature fusion (c). Finally, the point-wise features are used for point label prediction, leading to a semantic segmentation (d).

3. Method

Overview. Figure 2 provides an overview of the proposed online 3D scene segmentation method. The input to our method is an online acquired RGB-D sequence being reconstructed with real-time depth-fusion [6]. Let us denote the RGB-D sequence by $f^k = \{(c_m^k, p_m^k)\}_{m=0}^M, k = 0, 1, \dots, K$, where c_m^k and p_m^k store the RGB-D information of pixel m of frame k and the coordinates of its corresponding 3D point, respectively. Given a reconstruction represented by a point set \mathcal{P} , we construct a global tree \mathcal{T}_G maintaining the spatial organization of all points, as well as a per-point local trees $\{\mathcal{T}_L(p)\}_{p \in \mathcal{P}}$ storing the 1-ring neighborhood for each point (Section 3.1). The dynamic data structure supports fusion-aware point convolution encompassing intra-frame and inter-frame feature fusion (Section 3.2). The point-wise features are used for point label prediction, resulting in a semantic segmentation (Section 3.3).

3.1. Dynamic Global-Local Tree Organization

To support both intra-frame and inter-frame feature learning with point-based convolution, we require a data structure to organize the dynamically reconstructed, unstructured point cloud. There are several considerations in designing such a dynamic data structure. *Firstly*, to facilitate point-based convolution, we need to construct the local neighborhood of any given point. *Second*, the data structure should support fast update of the local neighborhoods under time-varying geometry. *Thirdly*, to realize 2D-to-3D and frame-to-frame feature fusion, the data structure should allow us to find correspondence between image pixels and reconstructed points. This way, pixels across different frames can be matched through the shared corresponding 3D point.

To meet those requirements, we design a two-level tree-based data structure. Globally, we construct a coordinate-based tree organization of points which supports fast neigh-

borhood query for any given point. This allows us to find pixel-to-point correspondence and point-based neighborhood efficiently. Based on the global tree, we build for each point an octree from which multi-scale local neighborhood can be found quickly for point-based convolution.

Global coordinate interval tree. We maintain three coordinate interval trees $\mathcal{T}_G^x, \mathcal{T}_G^y$ and \mathcal{T}_G^z , one for each dimension. Without loss of generality, we take \mathcal{T}_G^x for example to describe the tree construction. Each node $n_i \in \mathcal{T}_G^x$ records a set of point $\mathcal{P}_i^x \subset \mathcal{P}$ in which each point has its x-coordinate lie in the interval $[x_{\min}(n_i), x_{\max}(n_i)]$. $x_{\min}(n_i)$ and $x_{\max}(n_i)$ are the minimum and maximum threshold for node n_i . We stipulate the adjacent nodes in a coordinate interval tree complies with the following interval constraints:

$$x_{\max}(n_l) < x_{\min}(n_p), x_{\max}(n_p) < x_{\min}(n_r),$$

with n_l and n_r being the left and right child of node n_p . The entire 3D scene is then split into slices along x-dimension.

The coordinate interval tree is constructed dynamically as more 3D points are reconstructed and inserted. The point insertion of coordinate interval tree is conducted as follows. Given a 3D point $p = (x_p, y_p, z_p)$, we first find a node $n_i \in \mathcal{T}_G^x$ satisfying $x_p \in [x_{\min}(n_i), x_{\max}(n_i)]$, through a top-down traverse of the tree. If such a node exists, p is added to the corresponding point set \mathcal{P}_i^x of the node. Otherwise, we create a new leaf node whose point set is initialized as $\{p\}$ and coordinate interval as $[x_p - h, x_p + h]$. Here, h is the half size of coordinate intervals. This new node is then attached to the node whose interval is closest to the new node's interval. The detailed explanation of tree construction with balance maintenance can be found in the supplemental material.

After constructing the coordinate interval trees for all three dimensions, we can achieve efficient point correspondence search and neighborhood retrieval for any given query 3D point $q = (x_q, y_q, z_q)$. Through traversing the three trees, we obtain three nodes $n_i \in \mathcal{T}_G^x, n_j \in \mathcal{T}_G^y, n_k \in \mathcal{T}_G^z$ satis-

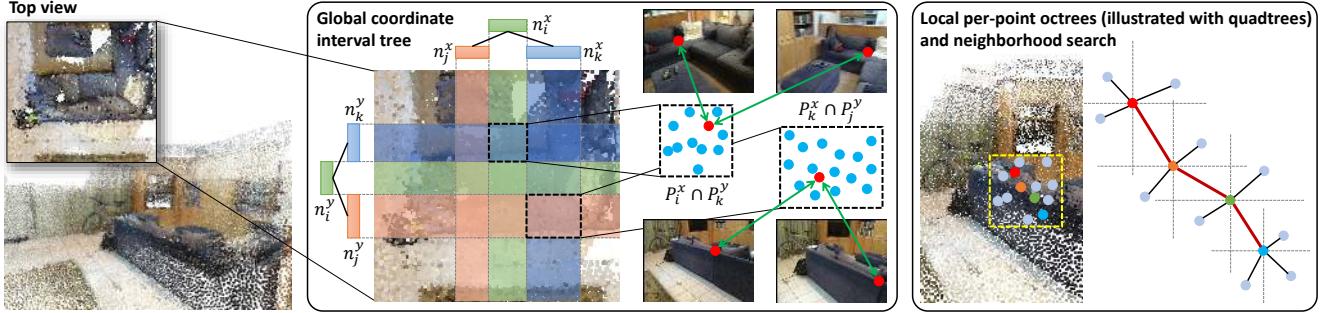


Figure 3: Illustration of global-local trees. The global coordinate interval trees are shown for x- and y-dimension only. With these trees, we can find a local neighborhood for any given point as well as the correspondence between two pixels from different frames. The per-point octrees (illustrated with 2D quadtrees) can be used to find multi-ring neighborhoods.

finding $x_p \in [x_{\min}(n_i), x_{\max}(n_i)]$, $y_p \in [y_{\min}(n_j), y_{\max}(n_j)]$ and $z_p \in [z_{\min}(n_k), z_{\max}(n_k)]$, respectively. The neighboring points of point q is simply the intersection of the three corresponding point sets: $\mathcal{N}(q) = \mathcal{P}_i^x \cap \mathcal{P}_j^y \cap \mathcal{P}_k^z$. Point correspondence can also be found efficiently within the neighborhood point set by using a distance threshold. And the adjacent intervals $\cup \mathcal{N}(q)$ around $\mathcal{N}(q)$ can be retrieved in a similar fashion.

Local per-point octrees. Although the global coordinate interval tree can be used to find a local neighborhood for any given point, the neighborhood is a merely unstructured point set. To conduct point convolution, a distance metric between points is required to apply convolutional operations with distance-based kernels [17, 26]. To this end, we need to sort the set of neighboring points into a structured organization based on surface-aware metric. This is achieved by maintaining per-point octrees so that the surface-aware neighborhood in arbitrary scale can be found efficiently.

Given a point $p \in \mathcal{P}$, we first retrieve its local neighborhood $\mathcal{N}(p)$ using the coordinate interval trees. We then divide the extended point set $\cup \mathcal{N}(p)$ and its according to the eight quadrants of the Cartesian coordinate system originated at p . Within each quadrant, we add the point that is the closest to p as the child of the corresponding direction, if the shortest distance is smaller than a threshold d^T . If some quadrant does not contain a point, however, the corresponding child node is left empty. The detailed description of per-point octree construction can be found in the supplemental material. After this process, we compute for each point a 1-ring neighborhood organized in a *direction-aware*. Based on the direction-aware octrees, one can easily expand the 1-ring neighborhood of a point into multiple rings through chaining octree-based neighbor searches; see Figure 3.

The neighborhood search enabled by the per-point octrees has two important characteristics. *First*, for each point, its eight neighbor points (child nodes) are scattered in the eight quadrants of its local Cartesian frame. When finding n -ring neighbors based on the octree-based point connections,

the consecutive searches can roughly follow the eight directions. Consequently, the octrees direct the search to find evenly distributed neighbors in all directions, which can be expensive to realize with naive region growing. *Second*, since the octrees maintain fine-scale local neighborhood, the search path along the octree-based connections approximately follows the 3D surface. This results in surface-aware n -ring neighborhoods. Both the two characteristics benefit substantially point convolution in learning improved 3D features over those working with Euclidean-distance-based neighborhoods as demonstrated in Section 4.

3.2. Fusion-aware Point Convolution

We propose a convolution operation which extends a recent work PointConv [37] with intra-frame and inter-frame feature fusion named fusion-aware point convolution. PointConv introduces a novel convolution operation over point cloud:

$$\text{PC}_p(W, F) = \sum_{\Delta p \in \Omega} W(\Delta p) F(p + \Delta p),$$

where $F(p + \Delta p)$ is the feature of a point in the local region Ω centered at p and W is weight function.

RGB-D frame sequence is a mixed data of rich 2D and 3D information with time stamp. However, PointConv only utilize limited 3D information and the rest does not make contribution to the segmentation task. To improve this method, there are three primary questions that we seek to answer within the section. First, PointConv is mainly about 3D, but how to fuse 2D information properly with 3D? Second, can we construct better local area Ω which ensures the neighborhood would be more relevant? Thirdly, how to utilize the inter-frame information given by the sequence?

2D-3D feature fusion. 3D feature fusion with online RGB-D reconstruction should better utilize the temporal correlation between adjacent frames, which goes beyond simplistic projection-based 2D-3D correlation.

The feature encoding at a 3D point should consider all the matched pixels in different frames if it is observed from multiple views. Pixel correspondence between consecutive frames can be easily retrieved based on \mathcal{T}_G^x , \mathcal{T}_G^y and \mathcal{T}_G^z . This way, each 3D point p in the scene would have a set of corresponding 2D pixels $I(p) = \{c^k | k \in n\}$. We can extract feature for each pixel c^k intra-frame via 2D convolution on image. We adopt a pre-trained FuseNet [9] without multi-scale layers as our 2D feature encoder (Section 3.3):

$$F^{2D}(c^k) = \text{FuseNet}(f_k, c^k),$$

where $\text{FuseNet}(f_k, c^k)$ is the 2D feature given by FuseNet for the pixel c^k in frame f_k . Therefore, each 3D point p in the scene has a set of corresponding 2D features and maxpooling is adopted to fuse them into one feature:

$$F^{2D3D}(p) = \text{maxpooling}\{F^{2D}(c^k) | c^k \in I(p)\}$$

Octree-induced surface-aware 3D convolution. Most deep convolutional neural networks for 3D point clouds gather neighborhood information on the basis of Euclidean distance. Apparently, geodesic distance can better capture the underlying geometry and topology of 3D surfaces. We propose octree-based neighborhood to take advantage of the geodesic approximation offered by our octree structure.

In particular, the local region Ω for each point p is given by its octree $\mathcal{T}_L(p)$. This approach could ensure the neighborhood would only enlarge along the object surface which is surface-aware, but not skip some gaps to reach shortest distance on Euclidean metric. A visual example is shown in Figure 8, the neighborhood searched by our approach would be more semantic related with the central point. This characteristic would benefit the following segmentation task. The construction of n-ring local region $\Omega^n(p)$ follows:

$$\begin{aligned} \Omega^n(p) &= \{\mathcal{T}_L(p) | p \in O(\Omega^{n-1}(p))\}, \\ \Omega^1(p) &= \Omega(p). \end{aligned}$$

We name the n-ring local region found $G^n(p^v)$ as octree based neighborhood, and we can adopt it to improve our fusion-aware point convolution. The convolution formulation $\text{FPC}_p(W, F^{2D3D})$ follows as below:

$$\text{FPC}_p(W, F^{2D3D}) = \sum_{\varphi \in \Omega^n(p)} W(\Delta p) F^{2D3D}(p + \Delta p).$$

Frame-to-frame feature fusion. Beside 2D-3D feature fusion, inter-frame information about segmentation uncertainty can benefit the segmentation task as well. A recent work about active scene segmentation [41] demonstrate that the segmentation entropy or uncertainty is crucial for online processing. We introduce a frame-to-frame feature fusion

which utilize the segmentation results given by previous frames to improve the performance for following frames.

For each 3D point p , our method would update its segmentation result if it is observed by a new frame f_i . Although we do not know the result is correct or not in the test, predicted segmentation uncertainty $U(p, i)$ given by the convoluted feature $\text{FPC}_p^i(W, F^{2D3D})$ at frame i can be easily retrieved by the network. Our basic idea here is that if p has low segmentation uncertainty in frame f_i , the current form of feature fusion should be useful in the future prediction.

In practice, we record every uncertainty $U(p, i)$ when processing the frame sequence. Further, our method conducts maxpooling operation on $\text{FPC}_p^{\text{current}}(W, F^{2D3D})$ which was updated in the current frame with feature $\text{FPC}_p^i(W, F^{2D3D})$. Thus, we rewrite the feature fusion as follows:

$$\begin{aligned} F^{\text{fused}}(p) &= \\ \text{maxpool}\{ &\text{FPC}_p^{\text{current}}(W, F^{2D3D}), \arg \min_{\text{FPC}_p^i} U(p, i)(W, F^{2D3D})\} \end{aligned}$$

This way, the intra- and inter-frame information of p are fused into $F^{\text{fused}}(p)$, maximally utilizing the information of the input RGB-D frames and significantly improving the performance of semantic segmentation.

3.3. Online Segmentation Network

As shown in Figure 4, the the backbone of our proposed online segmentation network is the global-local tree structure we introduced in Section 3.1. The global coordinates interval tree helps mapping corresponding 2D features for 2D-3D feature fusion and the local octrees help searching surface-aware neighborhood for 3D fusion-aware convolution.

Network architecture. The backbone of our 2D feature encoder is FuseNet [9]. Note, however, we discard the multi-scale layers in our implementation since the re-sampling operation is too time-consuming. Please refer to the supplemental material for details. This modification enables our method to achieve a close-to-interactive performance (10 FPS) with high accuracy.

According to Section 3.2, the 2D features $\text{FuseNet}(f_i, c^k)$ in different frames corresponding to the same 3D point p are fused as $F^{2D3D}(p)$. Fused features are adopted in our proposed fusion-aware convolution, and the convoluted feature $\text{FPC}_p^i(W, F^{2D3D})$ for point p in frame i could be calculated based on its n-ring neighbors $\Omega^n(p)$ along geometry surface. The convoluted feature is sent to a simple fully connected network FC which consists of 3 mlp layers to get the final feature which length is 128. It then be further fused with the selected feature with highest segmentation confidence in previous frames through max-pooling.

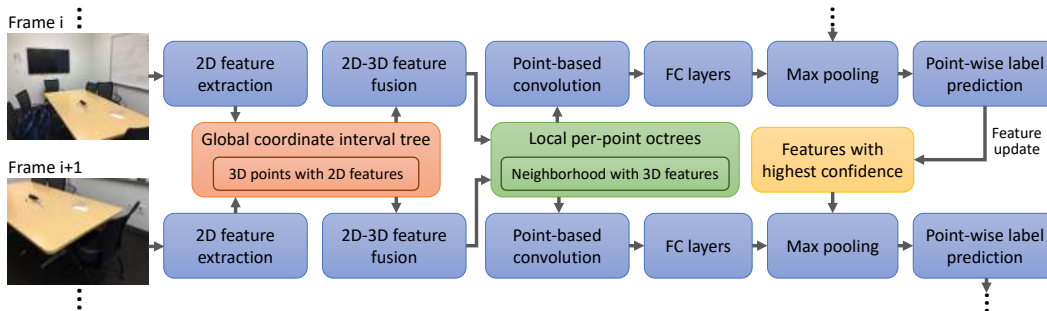


Figure 4: Network architecture. We show the pipeline with two consecutive frames. The global and local trees are dynamic data structure which evolves through time. The network output point-wise labels along with confidence. The labeling prediction is used to update features from the previous frame to be fused into the next frame.

Finally, we use it to predict the semantic label for p with a one layer classifier. Note that, our network updates $\arg \min_{FPC^i(W, F^{2D3D})} U(p, i)$ simultaneously to ensure the feature with the lowest segmentation uncertainty would be adopted in the future fusion-aware processing.

Training details. The batch size in training is 64. For each batch, we randomly select 8 different scenes which each contributes a sequence of 8 frames. The first frame of each sampled sequences are the first 8 data in each batch and the n -th frames are the $(8 * (n - 1))$ -th to $(8 * n - 1)$ -th data in each batch. We back-propagate the training gradients once every 8 forward passes of frames, and update the network weights after the forward pass of the whole batch. The training of our network on ScanNet [4] takes about 40~48 hours on a single Titan Xp GPU. For more details please refer to the supplemental material.

4. Results and evaluations

We first introduce our benchmark dataset and how we setup our experiments. Comparisons with some state-of-the-art alternatives are presented on both online and offline semantic segmentation tasks. We then conduct extensive evaluation for each components of our method. We also demonstrate the advantage of surface-aware characteristic of our method through some experiments.

4.1. Benchmarks

Dataset. We evaluate our method on two datasets: ScanNet [4] and SceneNN [12]. ScanNet contains 1513 scanned scene sequences, out of which we use 1200 sequences for training and the rest 312 for testing. SceneNN contains 50 high-quality scanned scene sequences with semantic label. However, this dataset is not specifically organized for online segmentation task. Some of the scanned scene sequences do not have camera pose information. Color image and depth map are not well aligned in some of the sequences as well. After some filtering work, we select 15 clean se-

Table 1: Accuracy comparison between our method and two state-of-the-art online scene segmentation methods.

Dataset	SemanticFusion [20]	ProgressiveFusion [23]	Ours
ScanNet	0.518	0.566	0.764
SceneNN	0.628	0.666	0.675

quences from SceneNN with proper scanned information for our evaluation.

Experiment configuration. To evaluate the performance of our method, we adopt accuracy and IOU as two indicators in our experiments. Since different online segmentation methods may adopt different 3D reconstruction approaches, it is really difficult to measure these two indicators in different 3D point clouds. In our experiment, we project the semantic labels of 3D points into their corresponding 2D frames and measure the accuracy and IOU in 2D.

4.2. Semantic Segmentation Comparison

Comparison with other online methods. Our method is compared to two state-of-the-art online segmentation methods for indoor scenes: SemanticFusion [20] and ProgressiveFusion [23]. The comparison is conducted on SceneNN and ScanNet respectively. The mean accuracy of three methods are shown in Table 1. From the results we can clearly see that our method gets the highest segmentation accuracy on both datasets. Note that we only train our method on ScanNet dataset and do not fine-tune it on SceneNN, and our method still outperform other methods on both datasets. This result demonstrate the generality of our method and it can be easily adopted on different dataset.

Comparison with offline methods. To further demonstrate the superiority of our segmentation method, we setup a comparison with three state-of-the-art offline segmentation methods (SparseConvNet [8], PointConv [37] and MinkowskiNet [3]) for indoor scenes as well. Note that, there is a challenge in online segmentation methods when compare to offline alternatives. Partial scene would be more

Table 2: IOU comparison between our method and state-of-the-art offline scene segmentation methods. Our method has the highest mean IOU, outperforming the state-of-the-art methods for nine semantic categories.

model	mean	wall	floor	cabinet	bed	chair	sofa	table	door	window	bookshelf	picture	counter	desk	curtain	fridge	bathshade	toilet	sink	bath tub	others
SparseConvNet	0.685	0.828	0.950	0.620	0.805	0.894	0.825	0.707	0.633	0.588	0.788	0.252	0.601	0.592	0.681	0.428	0.607	0.928	0.596	0.881	0.504
MinkowskiNet	0.715	0.841	0.949	0.641	0.806	0.900	0.845	0.745	0.648	0.608	0.792	0.289	0.637	0.65	0.742	0.509	0.690	0.916	0.689	0.832	0.570
PointConv	0.580	0.741	0.948	0.474	0.672	0.813	0.633	0.651	0.346	0.446	0.713	0.067	0.568	0.525	0.551	0.370	0.520	0.840	0.590	0.750	0.387
Ours	0.720	0.862	0.924	0.615	0.848	0.716	0.804	0.637	0.680	0.698	0.724	0.513	0.617	0.588	0.764	0.734	0.696	0.870	0.681	0.885	0.556

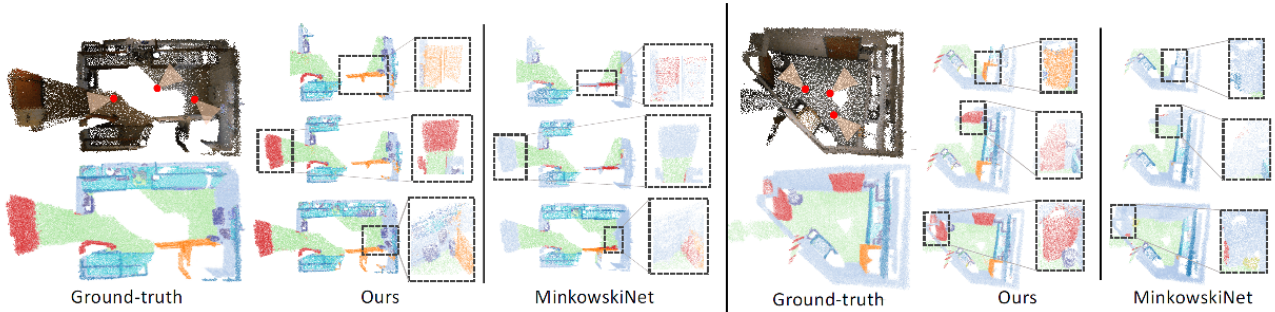


Figure 5: Visual comparison between our method and MinkowskiNet [3]. Our method works better than MinkowskiNet especially on those small and incomplete objects. Live demo is provided in the accompanying video.

Table 3: Ablation study on the two feature fusion operations of our method. The results justify each design.

2D3D feature	frame-to-frame feature	mean IOU
×	×	0.711
✓	×	0.718
×	✓	0.713
✓	✓	0.720

difficult to be segmented than the whole scene. We show some visual comparison on partial scenes in Figure 5. Our method can achieve much better results on these challenging cases which is very crucial for online tasks. We also present a segmentation comparison on the complete scene in Table 2. Our method achieves a comparable performance with the offline alternatives.

4.3. Ablation study

Feature fusion study. We investigate the effect of some crucial designs on segmentation performance. We turn off the 2D-3D feature fusion and frame-to-frame feature fusion in succession to assess how these two components would benefit our method. The results are shown in Table 3. Without 2D-3D feature fusion, our method cannot improve performance anymore with multi-view information in the sequence. The absence of frame-to-frame feature fusion makes our method lose the ability of learning from history in the sequence. We observe significant drop on performance if we turn off these designs, which prove the importance of these two designs. Similar results are also demonstrated in the right plot of Figure 6.

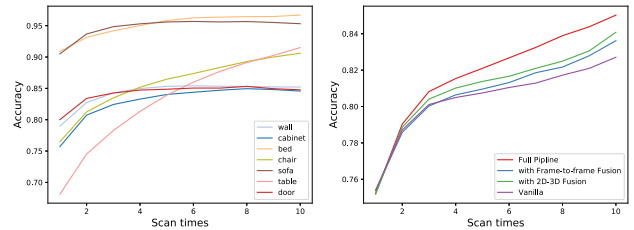


Figure 6: Segmentation accuracy improves over time. The left plot shows per-category accuracy. Objects with complex structures, such as tables and chairs, benefit more from our feature fusion. The right is mean accuracy. Segmentation accuracy improves significantly with increasing scans.

Temporal information study. The segmentation label of a 3D point would be updated if new information is fused in the sequence. To further investigate how our fusion-aware point convolution benefit the segmentation performance, we plot the accuracy of point labeling with increasing feature fusion times in Figure 6. In the plots, “vanilla” refers to our basic model with global-local tree and improved point convolution. We observe a clear increase of segmentation accuracy with increasing scans. Such effect is more observable on objects with complex structures, such as tables and chairs. Note that the performance of the vanilla model also improves with increasing scans. This is because both tree construction and point convolution benefit from a more complete point cloud.

To explore why point with more fusion times would have better segmentation performance, we track feature changes of some points in Figure 7. We plot the feature embedding of the points marked in yellow box, and we found the boundary between points with different semantic labels are

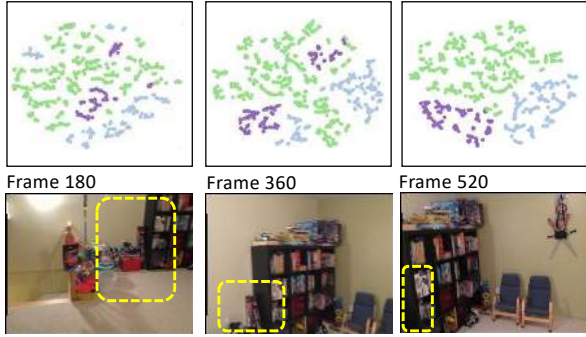


Figure 7: The evolution of point-wise feature embedding with more frames are acquired and fused. We show the t-SNE plots of the feature embedding of the same group of points with semantic labels indicated by color.

getting clear with during the scanning. Our fusion-aware features are more separable in the embedding space which make the segmentation network easier to train.

Surface-aware convolution context. Another important characteristic of our method is the neighborhood used in point convolution is surface-aware. This is enabled by our local octree structure. To verify this claim, we select two points on the point cloud and estimate their distance in Figure 8. We observe that tracing along the local octrees leads to a path whose length is closer to the ground-truth geodesic distance than the Euclidean distance. The visual example also demonstrates that the tree-induced path traces along the surface, making the point neighborhood surface-aware. Figure 8 also visualizes the kernel weight distribution of two points on the point cloud, showing that neighboring points which are geodesically closer have higher weights.

This characteristic makes our method more geometry-aware. To verify the geometric information would benefit the online segmentation, we design a comparison between convolutions based on Euclidean distance and our local-octree-induced distance in Figure 9. The results show that our surface-aware convolution demonstrates a better performance all the way along the increase of convolution area (size of receptive field).

5. Conclusion

For the task of semantic segmentation of a 3D scene being reconstructed with RGB-D fusion, we have presented a tree-based dynamic data structure to organize online fused 3D point clouds. It supports 3D point convolution over time-varying geometry, fusing information between 2D and 3D and from frame to frame. Our method achieves online segmentation at close-to-interactive frame-rate while reaching state-of-the-art accuracy. Our current method has a few limitations on which we plan to investigate in future works.




Figure 8: (a-c): Tracing path between two selected points along the local octrees. We compare the path length to the Euclidean distance and ground-truth geodesic distance. We find that the path traced along the local octrees is a good approximate of the geodesic distance. (d): The kernel weight distribution of two points (green dot) are visualized (red is high and blue is low).

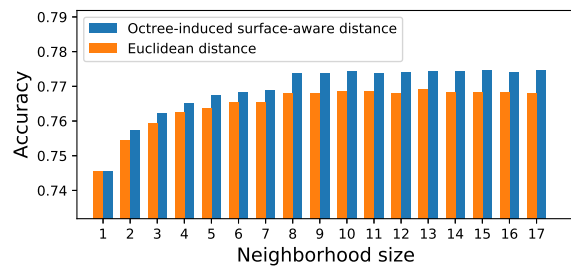


Figure 9: Segmentation accuracy of different neighborhood searching approaches. Our local-octree-induced neighborhood always leads to better segmentation accuracy than Euclidean based. We also find that increasing the neighborhood size more than a certain value would not improve the performance significantly.

First, our current system does not support data streaming, thus confining the per-frame point cloud density we could handle due to memory limit. *Second*, our method still relies on accurate camera poses for high-quality point fusion and convolution. Although it works well with the popular RGB-D fusion methods (such as BundleFusion [6]), it is interesting to investigate how to integrate online semantic segmentation with camera tracking, accomplishing semantic SLAM with a state-of-the-art accuracy for both. Further, we would like to test our method on outdoor scenes [34].

Acknowledgement

We thank the anonymous reviewers for the valuable suggestions. We are also grateful to Quang-Hieu Pham for the help on comparison with ProgressiveFusion. This work was supported in part by NSFC (61572507, 61532003, 61622212), NUDT Research Grants(No.ZK19-30), National Key Research and Development Program of China (No. 2018AAA0102200) and Natural Science Foundation of Hunan Province for Distinguished Young Scientists (2017JJ1002).

References

- [1] Matan Atzmon, Haggai Maron, and Yaron Lipman. Point convolutional neural networks by extension operators. *arXiv preprint arXiv:1803.10091*, 2018. [2](#)
- [2] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017. [1](#)
- [3] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. *arXiv preprint arXiv:1904.08755*, 2019. [6](#), [7](#)
- [4] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017. [6](#)
- [5] Angela Dai and Matthias Nießner. 3dmv: Joint 3d-multi-view prediction for 3d semantic scene segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–468, 2018. [1](#), [2](#)
- [6] Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (TOG)*, 36(3):24, 2017. [3](#), [8](#)
- [7] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9224–9232, 2018. [2](#)
- [8] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018. [6](#)
- [9] Caner Hazirbas, Lingni Ma, Csaba Domokos, and Daniel Cremers. Fusetnet: Incorporating depth into semantic segmentation via fusion-based cnn architecture. In *Asian Conference on Computer Vision (ACCV)*, 2016. [5](#)
- [10] Alexander Hermans, Georgios Floros, and Bastian Leibe. Dense 3d semantic mapping of indoor scenes from rgb-d images. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2631–2638. IEEE, 2014. [2](#)
- [11] Pedro Hermosilla, Tobias Ritschel, Pere-Pau Vázquez, Àlvar Vinacua, and Timo Ropinski. Monte carlo convolution for learning on non-uniformly sampled point clouds. *ACM Trans. on Graph. (SIGGRAPH Asia)*, page 235, 2018. [2](#)
- [12] Binh-Son Hua, Quang-Hieu Pham, Duc Thanh Nguyen, Minh-Khoi Tran, Lap-Fai Yu, and Sai-Kit Yeung. SceneNN: A Scene Meshes Dataset with aNNotations. In *International Conference on 3D Vision (3DV)*, 2016. [6](#)
- [13] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, and Andrew Fitzgibbon. KinectFusion: Real-time 3D reconstruction and interaction using a moving depth camera. In *UIST*, pages 559–568, 2011. [1](#)
- [14] Hou Ji, Angela Dai, and Matthias Nießner. 3d-sis: 3d semantic instance segmentation of rgb-d scans. In *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 2019. [1](#), [2](#)
- [15] Roman Klokov and Victor Lempitsky. Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 863–872, 2017. [2](#)
- [16] Kevin Lai, Liefeng Bo, and Dieter Fox. Unsupervised feature learning for 3d scene labeling. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3050–3057. IEEE, 2014. [2](#)
- [17] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *Advances in Neural Information Processing Systems*, pages 820–830, 2018. [2](#), [4](#)
- [18] Lingni Ma, Jörg Stückler, Christian Kerl, and Daniel Cremers. Multi-view deep learning for consistent semantic mapping with rgb-d cameras. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 598–605. IEEE, 2017. [2](#)
- [19] 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](#)
- [20] John McCormac, Ankur Handa, Andrew Davison, and Stefan Leutenegger. Semanticfusion: Dense 3d semantic mapping with convolutional neural networks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4628–4635. IEEE, 2017. [1](#), [2](#), [6](#)
- [21] Gaku Narita, Takashi Seno, Tomoya Ishikawa, and Yohsuke Kaji. Panopticfusion: Online volumetric semantic mapping at the level of stuff and things. *arXiv preprint arXiv:1903.01177*, 2019.
- [22] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger. Real-time 3D reconstruction at scale using voxel hashing. *ACM Trans. on Graph. (SIGGRAPH Asia)*, 32(6):169, 2013. [1](#)
- [23] Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Real-time progressive 3d semantic segmentation for indoor scenes. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1089–1098. IEEE, 2019. [6](#)
- [24] Charles Qi, Hao Su, Matthias Niessner, Angela Dai, Mengyuan Yan, and Leonidas Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proc. CVPR*, 2016. [2](#)
- [25] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017. [1](#), [2](#)
- [26] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017. [2](#), [4](#)
- [27] 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
- [28] Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H. J. Kelly, and Andrew J. Davison. SLAM++: Simultaneous localisation and mapping at the level of objects. In *Proc. CVPR*, pages 1352–1359, 2012. 2
- [29] Martin Simonovsky and Nikos Komodakis. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3693–3702, 2017. 2
- [30] Jörg Stückler and Sven Behnke. Multi-resolution surfel maps for efficient dense 3d modeling and tracking. *Journal of Visual Communication and Image Representation*, 25(1):137–147, 2014. 2
- [31] 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 *Proc. CVPR*, pages 2530–2539, 2018. 2
- [32] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3D shape recognition. In *Proc. ICCV*, pages 945–953, 2015. 2
- [33] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3d. In *Proc. CVPR*, pages 3887–3896, 2018. 2
- [34] The SYNTHIA dataset. <https://synthia-dataset.net/>. 8
- [35] Weiyue Wang, Ronald Yu, Qiangui Huang, and Ulrich Neumann. Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2569–2578, 2018. 2
- [36] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *arXiv preprint arXiv:1801.07829*, 2018. 2
- [37] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proc. CVPR*, pages 9621–9630, 2019. 2, 4, 6
- [38] Yu Xiang and Dieter Fox. Da-rnn: Semantic mapping with data associated recurrent neural networks. *arXiv preprint arXiv:1703.03098*, 2017. 2
- [39] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In *Proc. ECCV*, pages 87–102, 2018. 2
- [40] Li Yi, Wang Zhao, He Wang, Minhyuk Sung, and Leonidas J Guibas. Gspn: Generative shape proposal network for 3d instance segmentation in point cloud. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3947–3956, 2019. 2
- [41] Lintao Zheng, Chenyang Zhu, Jiazhao Zhang, Hang Zhao, Hui Huang, Matthias Niessner, and Kai Xu. Active scene understanding via online semantic reconstruction. *arXiv preprint arXiv:1906.07409*, 2019. 5