

SEGCloud: Semantic Segmentation of 3D Point Clouds

<http://segcloud.stanford.edu>

Lyne P. Tchapmi Christopher B. Choy Iro Armeni JunYoung Gwak Silvio Savarese
lynetcha@stanford.edu chrischoy@ai.stanford.edu iarmeni@cs.stanford.edu jgwak@stanford.edu sslivio@stanford.edu

Stanford University

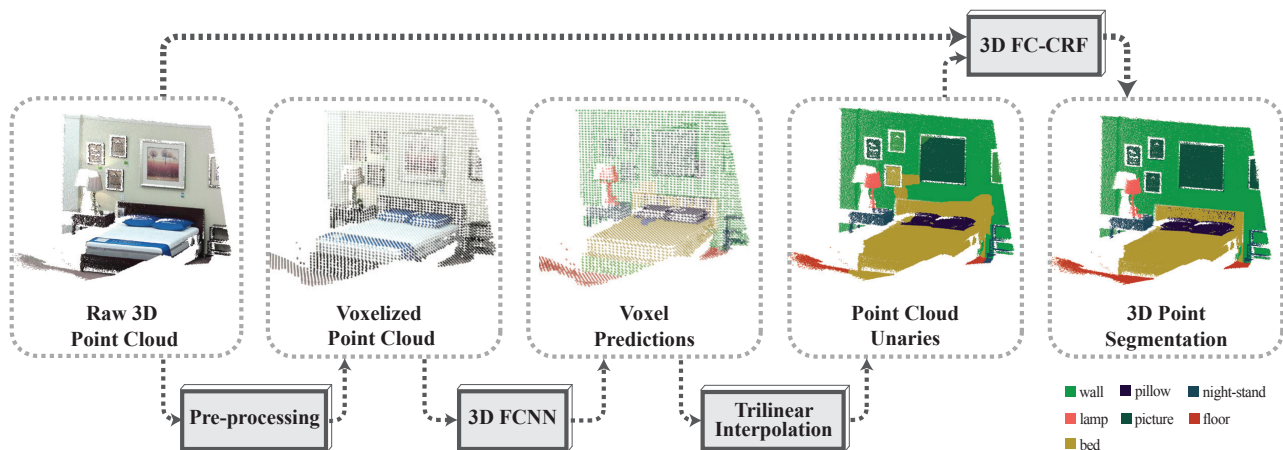


Figure 1: **SEGCloud**: A 3D point cloud is voxelized and fed through a 3D fully convolutional neural network to produce coarse down-sampled voxel labels. A trilinear interpolation layer transfers this coarse output from voxels back to the original 3D Points representation. The obtained 3D point scores are used for inference in the 3D fully connected CRF to produce the final results. Our framework is trained end-to-end.

Abstract

3D semantic scene labeling is fundamental to agents operating in the real world. In particular, labeling raw 3D point sets from sensors provides fine-grained semantics. Recent works leverage the capabilities of Neural Networks (NNs), but are limited to coarse voxel predictions and do not explicitly enforce global consistency. We present SEGCloud, an end-to-end framework to obtain 3D point-level segmentation that combines the advantages of NNs, trilinear interpolation (TI) and fully connected Conditional Random Fields (FC-CRF). Coarse voxel predictions from a 3D Fully Convolutional NN are transferred back to the raw 3D points via trilinear interpolation. Then the FC-CRF enforces global consistency and provides fine-grained semantics on the points. We implement the latter as a differentiable Recurrent NN to allow joint optimization. We evaluate the framework on two indoor and two outdoor 3D datasets (NYU V2, S3DIS, KITTI, Semantic3D.net), and show performance comparable or superior to the state-of-the-art on all datasets.

1. Introduction

Scene understanding is a core problem in Computer Vision and is fundamental to applications such as robotics, autonomous driving, augmented reality, and the construction industry. Among various scene understanding problems, 3D semantic segmentation allows finding accurate object boundaries along with their labels in 3D space, which is useful for fine-grained tasks such as object manipulation, detailed scene modeling, etc.

Semantic segmentation of 3D point sets or point clouds has been addressed through a variety of methods leveraging the representational power of graphical models [36, 44, 3, 48, 30, 35]. A common paradigm is to combine a classifier and a Conditional Random Field (CRF) [39] to predict spatially consistent labels for each data point [68, 69, 45, 66, 69]. Random Forests classifiers [7, 15] have shown great performance on this task, however the Random Forests classifier and CRF stage are often optimized independently and put together as separate modules, which limits the information flow between them.

3D Fully Convolutional Neural Networks (3D-

FCNN) [42] are a strong candidate for the classifier stage in 3D Point Cloud Segmentation. However, since they require a regular grid as input, their predictions are limited to a coarse output at the voxel (grid unit) level. The final segmentation is coarse since all 3D points within a voxel are assigned the same semantic label, making the voxel size a factor limiting the overall accuracy. To obtain a fine-grained segmentation from 3D-FCNN, an additional processing of the coarse 3D-FCNN output is needed. We tackle this issue in our framework which is able to leverage the coarse output of a 3D-FCNN and still provide a fine-grained labeling of 3D points using trilinear interpolation (TI) and CRF.

We propose an end-to-end framework that leverages the advantages of 3D-FCNN, trilinear interpolation [47], and fully connected Conditional Random Fields(FC-CRF) [39, 37] to obtain fine-grained 3D Segmentation. In detail, the 3D-FCNN provides class probabilities at the voxel level, which are transferred back to the raw 3D points using trilinear interpolation. We then use a Fully Connected Conditional Random Field (FC-CRF) to infer 3D point labels while ensuring spatial consistency. Transferring class probabilities to points before the CRF step, allows the CRF to use point level modalities (color, intensity, etc.) to learn a fine-grained labeling over the points, which can improve the initial coarse 3D-FCNN predictions. We use an efficient CRF implementation to perform the final inference. Given that each stage of our pipeline is differentiable, we are able to train the framework end-to-end using standard stochastic gradient descent.

The contributions of this work are:

- We propose to combine the inference capabilities of Fully Convolutional Neural Networks with the fine-grained representation of 3D Point Clouds using TI and CRF.
- We train the voxel-level 3D-FCNN and point-level CRF jointly and end-to-end by connecting them via Trilinear interpolation enabling segmentation in the original 3D points space.

Our framework can handle 3D point clouds from various sources (laser scanners, RGB-D sensors, etc.), and we demonstrate state-of-the art performance on indoor and outdoor, partial and fully reconstructed 3D scenes, namely on NYU V2[52], Stanford Large-Scale 3D Indoor Spaces Dataset (S3DIS)[5], KITTI[23, 22], and the Semantic3D.net benchmark for outdoor scenes[26].

2. Related Work

In this section, we present related works with respect to three main aspects of our framework: neural networks for 3D data, graphical models for 3D Segmentation and works that explore the combination of Convolutional Neural Net-

works (CNN) and CRF. Other techniques have been employed for 3D Scene Segmentation [13, 2, 40] but we focus mainly on the ones related to the above topics.

Neural Networks for 3D Data: 3D Neural Networks have been extensively used for 3D object and parts recognition [60, 54, 46, 25, 53, 21], understanding object shape priors, as well as generating and reconstructing objects [73, 71, 19, 70, 12]. Recent works have started exploring the use of Neural Networks for 3D Semantic Segmentation [53, 16, 32]. Qi *et al.* [53] propose a Multilayer Perceptron (MLP) architecture that extracts a global feature vector from a 3D point cloud of $1m^3$ physical size and processes each point using the extracted feature vector and additional point level transformations. Their method operates at the point level and thus inherently provides a fine-grained segmentation. It works well for indoor semantic scene understanding, although there is no evidence that it scales to larger input dimensions without additional training or adaptation required. Huang *et al.* [32] present a 3D-FCNN for 3D semantic segmentation which produces coarse voxel-level segmentation. Dai *et al.* [16] also propose a fully convolutional architecture, but they make a single prediction for all voxels in the same voxel grid column. This makes the wrong assumption that a voxel grid column contains 3D points with the same object label. All the aforementioned methods are limited by the fact that they do not explicitly enforce spatial consistency between neighboring points predictions and/or provide a coarse labeling of the 3D data. In contrast, our method makes fine-grained predictions for each point in the 3D input, explicitly enforces spatial consistency and models class interactions through a CRF. Also, in contrast to [53], we readily scale to larger and arbitrarily sized inputs, since our classifier stage is fully convolutional.

Graphical Models for 3D Segmentation: Our framework builds on top of a long line of works combining graphical models([61, 62, 39, 20, 38]) and highly engineered classifiers. Early works on 3D Semantic Segmentation formulate the problem as a graphical model built on top of a set of features. Such models have been used in several works to capture contextual relationships based on various features and cues such as appearance, shape, and geometry. These models are shown to work well for this task [50, 49, 36, 58, 44, 3, 48].

A common paradigm in 3D semantic segmentation combines a classifier stage and a Conditional Random Field to impose smoothness and consistency [68, 69, 45, 66, 69]. Random Forests [7, 15] are a popular choice of classifier in this paradigm and in 3D Segmentation in general [75, 17, 9, 8, 51, 67]; they use hand-crafted features to robustly provide class scores for voxels, oversegments or 3D Points. In [45], the spin image descriptor is used as a feature, while [68] uses a 14-dimensional feature vector based on geometry and appearance. Hackel *et al.* [27] also

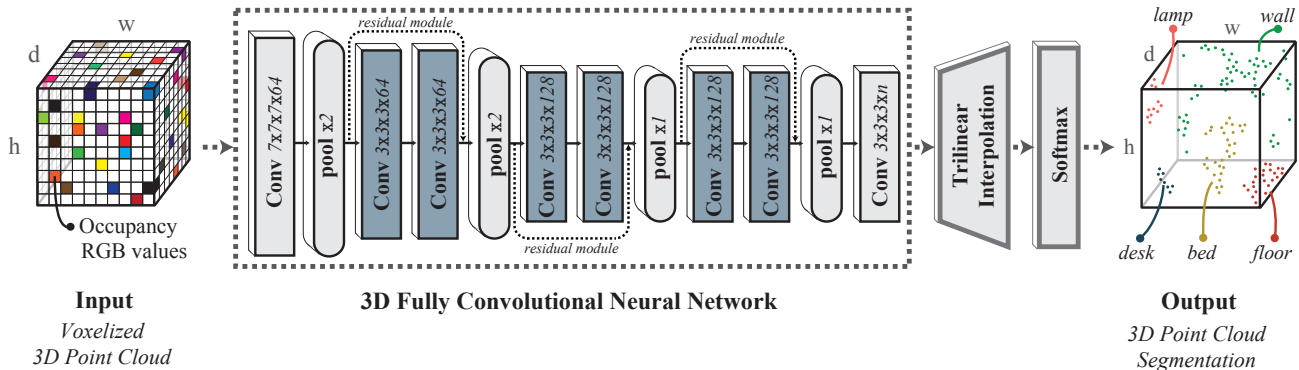


Figure 2: **Network architecture:** The 3D-FCNN is made of 3 residual layers sandwiched between 2 convolutional layers. Max Pooling in the early stages of the network yields a 4X downsampling.

define a custom set of features aimed at capturing geometry, appearance and location. In these works, the Random Forests output is used as unary potentials (class scores) for a CRF whose parameters are learned independently. The CRF then leverages the confidence provided by the classifier, as well as similarity between an additional set of features, to perform the final inference. In contrast to these methods, our framework uses a 3D-FCNN which can learn higher dimensional features and provide strong unaries for each data point. Moreover, our CRF is implemented as a fully differentiable Recurrent Neural Network, similar to [76]. This allows the 3D-FCNN and CRF to be trained end-to-end, and enables information flow from the CRF to the CNN classification stage.

Joint CNN + CRF: Combining 3D CNN and 3D CRF has been previously proposed for the task of lesion segmentation in 3D medical scans. Kamnitsas *et al.* [34] propose a multi-scale 3D CNN with a CRF to classify 4 types of lesions from healthy brain tissues. The method consists of two modules that are not trained end-to-end: a 2-stream architecture operating at 2 different scan resolutions and a CRF. In the CRF training stage, the authors reduce the problem to a 2-class segmentation task in order to find parameters for the CRF that can improve segmentation accuracy.

Joint end-to-end training of CNN and CRF was first demonstrated by [76] in the context of image semantic segmentation, where the CRF is implemented as a differentiable Recurrent Neural Network (RNN). The combination of CNN and CRF trained in an end-to-end fashion demonstrated state-of-the-art accuracy for semantic segmentation in images. In [76] and other related works [42, 10], the CNN has a final upsampling layer with learned weights which allows to obtain pixel level unaries before the CRF stage. Our work follows a similar thrust by defining the CRF as an RNN and using a trilinear interpolation layer to transfer the coarse output of the 3D-FCNN to individual 3D points before the CRF stage. In contrast to [34], our framework is a single stream architecture which jointly optimizes the 3D CNN and CRF, targets the domain of 3D Scene Point Clouds, and is able to handle a large number of classes both at the CNN and CRF stage. Unlike [76, 42, 10], we choose

to use deterministic interpolation weights that take into account the metric distance between a 3D point and its neighboring voxel centers (Section 3.2). Our approach reduces the number of parameters to be learned, and we find it to work well in practice. We show that the combination of jointly trained 3D-FCNN and CRF with TI consistently performs better than a stand alone 3D-FCNN.

In summary, our work differs from previous works in the design of an end-to-end deep learning framework for fine-grained 3D semantic segmentation, the use of deterministic trilinear interpolation to obtain point-level segmentation, and the use of a jointly trained CRF to enforce spatial consistency. The rest of the paper is organized as follows. Sections 3 and 4 present the components of our end-to-end framework and Section 5 provides implementation details. Section 6 presents our experiments including datasets (6.1), benchmark results (6.2), and system analysis (6.3). Section 7 concludes with a summary of the presented results.

3. SEGCloud Framework

An overview of the SEGCloud pipeline is shown in Figure 1. In the first stage of our pipeline, the 3D data is voxelized and the resulting 3D grid is processed by a 3D fully convolutional neural network (3D-FCNN)¹. The 3D-FCNN down-samples the input volume and produces probability distributions over the set of classes for each down-sampled voxel (Section 3.1). The next stage is a trilinear interpolation layer which interpolates class scores from down-sampled voxels to 3D points (Section 3.2). Finally, inference is performed using a CRF which combines the original 3D points features with interpolated scores to produce fine-grained class distributions over the point set (Section 3.3). Our entire pipeline is jointly optimized and the CRF inference and joint optimization processes are presented in Section 4.

¹Depending on the type of 3D data a pre-processing step of converting it to a 3D point cloud representation might be necessary.

3.1. 3D Fully Convolutional Neural Network

Our framework uses a 3D-FCNN to learn a representation suitable for semantic segmentation. Moreover, the fully convolutional network reduces the computational overhead needed to generate predictions for each voxel by sharing computations [43]. In the next section, we describe how we represent 3D point clouds as an input to the 3D-FCNN.

3D-FCNN data representation: Given that the 3D-FCNN input should be in the form of a voxel grid, we convert 3D point clouds as follows. Each data point is a 3D observation o_i , that consists of the 3D position p_i and other available modalities, such as the color intensity I_i and sensor intensity S_i . We place the 3D observations $\mathbf{O} = \{o_i\}$ in a metric space so that the convolution kernels can learn the scale of objects. This process is usually handled in most 3D sensors. Then we define a regular 3D grid that encompasses the 3D observations. We denote each cell in the 3D grid as a voxel v_i and for simplicity, each cell is a cube with length $V = 5cm$. Most of the space in the 3D input is empty and has no associated features. To characterize this, we use a channel to denote the occupancy as a binary value (zero or one). We use additional channels to represent other modalities. For instance, three channels are used for RGB color, and one channel is used for sensor intensity when available. **Architecture:** Our 3D-FCNN architecture is illustrated in Figure 2. We use 3 residual modules [28] sandwiched between 2 convolutional layers, as well as 2 destructive pooling layers in the early stages of the architecture to down-sample the grid, and 2 non-destructive ones towards the end. The early down-sampling gives us less memory footprint. The entire framework is fully convolutional and can handle arbitrarily sized inputs. For each voxel v_i , the 3D-FCNN outputs scores(logits) L_i associated with a probability distribution q_i over labels. The resulting scores are transferred to the raw 3D points via trilinear interpolation.

3.2. 3D Trilinear Interpolation

The process of voxelization and subsequent down-sampling in the 3D-FCNN converts our data representation to a coarse 3D grid which limits the resolution of semantic labeling at the CRF stage (to 20 cm in our case). Running the CRF on such coarse voxels results in a coarse segmentation. One option to avoid this information loss is to increase the resolution of the voxel grid (*i.e.* decrease the voxel size) and/or remove the destructive pooling layers, and run the CRF directly on the fine-grained voxels. However, this quickly runs into computational and memory constraints, since for given 3D data dimensions, the memory requirement of the 3D-FCNN grows cubically with the resolution of the grid. Also, for a given 3D-FCNN architecture, the receptive field decreases as the resolution of the grid increases, which can reduce performance due to having less context available during inference(see [63]).

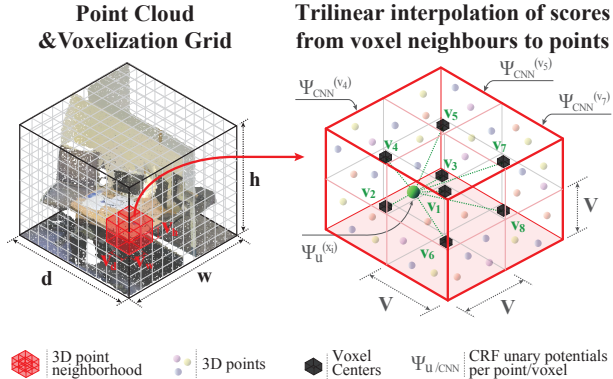


Figure 3: **Trilinear interpolation of class scores from voxels to points:** Each point’s score is computed as the weighted sum of the scores from its 8 spatially closest voxel centers.

We therefore dismiss a voxel-based CRF approach and resort to running CRF inference using the raw 3D points as nodes. In this way, the CRF can leverage both the 3D-FCNN output and the fine-grained modalities of the input 3D points to generate accurate predictions that capture scene and object boundaries in detail. We achieve this using trilinear interpolation to transfer the voxel-level predictions from the 3D-FCNN to the raw 3D points as illustrated in Figure 3. Specifically, for each point, $o_i = \{p_i, I_i, S_i\}$, we define a random variable x_i that denotes the semantic class, and the scores(logits) L_i associated with the distribution of x_i are defined as a weighted sum of scores $L_{i,n}(x_{i,n})$ of its 8 spatially closest voxels $v_{i,n}$, $n \in \{1, \dots, 8\}$ whose centers are $(p_{i,n}^x, p_{i,n}^y, p_{i,n}^z)$ as follows:

$$\psi_u(x_i = l) = L_i(x_i = l) = \sum_{n=1}^8 w_{i,n} L_{i,n}(x_{i,n} = l) \quad (1)$$

$$w_{i,n} = \prod_{s \in \{x,y,z\}} (1 - |p_i^s - p_{i,n}^s|/V)$$

where V is the voxel size. During back propagation, we use the same trilinear interpolation weights $w_{i,n}$ to splat the gradients from the CRF to the 3D-FCNN. The obtained point level scores are then used as unaries in the CRF.

3.3. 3D Fully Connected Conditional Random Field

The energy function of a CRF consists of a set of unary and pairwise potential energy terms. The unary potentials are a proxy for the initial probability distribution across semantic classes and the pairwise potentials enforce smoothness and consistency between predictions. The energy of the CRF is defined as,

$$E(x) = \sum_i \psi_u(x_i) + \sum_{i < j} \psi_p(x_i, x_j) \quad (2)$$

where ψ_u denotes the unary potential which is defined in Equation 3.2 and ψ_p denotes the pairwise potential. Note that all nodes in the CRF are connected with each other through the pairwise potentials. We use the Gaussian kernels from [37] for the pairwise potentials,

$$\begin{aligned} \psi_p(x_i, x_j) = & \mu(x_i, x_j) \left[w_s \exp\left(-\frac{|p_i - p_j|^2}{2\theta_\gamma^2}\right) \right. \\ & \left. + w_b \exp\left(-\frac{|p_i - p_j|^2}{2\theta_\alpha^2} - \frac{|I_i - I_j|^2}{2\theta_\beta^2}\right) \right] \end{aligned} \quad (3)$$

where w_b and w_s are the weights of the bilateral and spatial kernel respectively, μ is the label compatibility score, and $\theta_\alpha, \theta_\beta, \theta_\gamma$ are the kernels' bandwidth parameters. When RGB information is not available, we only use the spatial kernel. Using Gaussian kernels enables fast variational inference and learning through a series of convolutions on a permutohedral lattice [1] (Section 4).

4. CRF Inference and Joint Optimization

Exact energy minimization in CRF is intractable, therefore we rely on a variational inference method which allows us to jointly optimize both the CRF and 3D-FCNN [76, 37]. The output after the CRF energy minimization gives us fine-grained predictions for each 3D point that takes smoothness and consistency into account. Given the final output of the CRF, we follow the convention and use the distance between the prediction and ground truth semantic labels as a loss function and minimize it.

CRF Inference: The CRF with Gaussian potential has a special structure that allows fast and efficient inference. Krähenbühl *et al.* [37] presented an approximate inference method which assumes independence between semantic label distributions $Q(\mathbf{X}) = \prod_i Q_i(x_i)$, and derived the update equation:

$$\begin{aligned} Q_i^+(x_i = l) = & \frac{1}{Z_i} \exp\left\{-\psi_u(x_i)\right. \\ & \left.- \sum_{l' \in \mathcal{L}} \mu(l, l') \sum_{m=1}^K w^{(m)} \sum_{j \neq i} k^{(m)}(f_i, f_j) Q_j(l')\right\} \end{aligned} \quad (4)$$

The above update equation can be implemented using simple convolutions, sums and softmax as shown by Zheng *et al.* [76], who implemented CRF inference and learning as a Recurrent Neural Network (RNN), named CRF-RNN. CRF-RNN can be trained within a standard CNN framework, so we follow the same procedure to define our 3D CRF as an RNN for inference and learning. This formulation allows us to integrate the CRF within our 3D-FCNN framework for joint training.

Loss: Once we minimize the energy of the CRF in Equation 2, we obtain the final prediction distribution of the semantic class x_i on each 3D observation o_i . Denoting

the ground truth discrete label of the observation o_i as y_i , we follow the convention and define our loss function as the distance between a final prediction distribution and the ground truth distribution using KL divergence:

$$L(\mathbf{x}, \mathbf{y}) = \frac{1}{N} \sum_{i=1}^N E_{y_i}[-\log p(x_i)] \quad (5)$$

where N is the number of observations. Since the entropy of y_i is a constant with respect to all parameters, we do not include it in the loss function equation.

5. Implementation Details

We implemented the SEGCloud framework using the Caffe neural network library [33]². Within the Caffe framework, we adapted the bilinear interpolation of [11] and implemented trilinear interpolation as a neural network layer. All computations within the 3D-FCNN, trilinear interpolation layer, and CRF are done on a Graphical Processing Unit (GPU). For CRF inference, we adapt the RNN implementation of Zheng *et al.* [76] to 3D point clouds.

To address the lack of data in some datasets and make the network robust, we applied various data augmentation techniques such as random color augmentation, rotation along the upright direction, and points sub-sampling. The above random transformations and sub-sampling allow to increase the effective size of each dataset by at least an order of magnitude, and can help the network build invariance to rotation/viewpoint changes, as well as reduced and varying context (see [63]).

Training is performed in a 2-step process, similar to [76] (see Figure 4). In the first stage, we train the 3D-FCNN in isolation via trilinear interpolation for 200 epochs.

In the second stage, we jointly train the 3D-FCNN and the CRF end-to-end (both modules connected through the trilinear interpolation layer). The approximate variational inference method we used for the CRF [37] approximates convolution in a permutohedral grid whose size depends on the bandwidth parameters $\theta_\alpha, \theta_\beta, \theta_\gamma$. We fixed θ_γ at 5cm, θ_β at 11 and used a grid search with small perturbation on a validation set to find the optimal θ_α (see [63]).

6. Experiments

In this section, we evaluate our framework on various 3D datasets and analyze the performance of key components.

6.1. Datasets

Several 3D Scene datasets have been made available to the research community [56, 4, 5, 31, 59, 72, 52, 16, 24, 74]. We chose four of them so that they cover indoor and

²We use [64] that supports 3D convolution.

Table 1: Results on the Semantic3D.net Benchmark (*reduced-8 challenge*)

Method	man-made terrain	natural terrain	high vegetation	low vegetation	buildings	hard scape	scanning artefacts	cars	mIOU	mAcc ³
TMLC-MSR [27]	89.80	74.50	53.70	26.80	88.80	18.90	36.40	44.70	54.20	68.95
DeePr3SS [41]	85.60	83.20	74.20	32.40	89.70	18.50	25.10	59.20	58.50	88.90
SnapNet [6]	82.00	77.30	79.70	22.90	91.10	18.40	37.30	64.40	59.10	70.80
3D-FCNN-TI(Ours)	84.00	71.10	77.00	31.80	89.90	27.70	25.20	59.00	58.20	69.86
SEGCloud (Ours)	83.90	66.00	86.00	40.50	91.10	30.90	27.50	64.30	61.30	73.08

Table 2: Results on the Large-Scale 3D Indoor Spaces Dataset (S3DIS)

Method	ceiling	floor	wall	beam	column	window	door	chair	table	bookcase	sofa	board	clutter	mIOU	mAcc
PointNet [53]	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	41.09	48.98
3D-FCNN-TI(Ours)	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	47.46	54.91
SEGCloud (Ours)	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	48.92	57.35

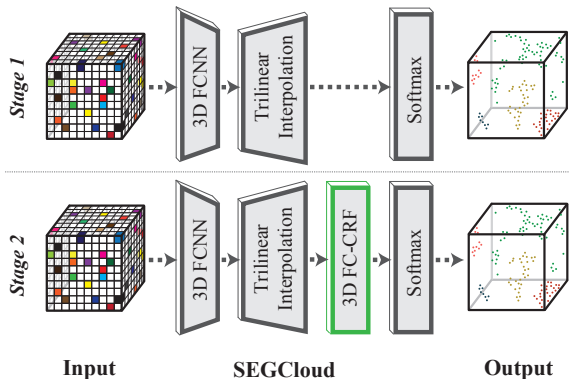


Figure 4: We follow a 2-stage training by first optimizing over the point-level unary potentials (no CRF) and then over the joint framework for point-level fine-grained labeling.

outdoor, partial and fully reconstructed, as well as small, medium and large scale point clouds. For our evaluation, we favor those for which previous 3D Semantic Segmentation works exist, with replicable experimental setups for comparison. We choose baselines so that they are representative of the main research thrusts and topics related to our method (*i.e.*, Neural Networks, Random Forests, and CRFs). The datasets we chose for evaluation are the Semantic3D.net Benchmark [26], the Stanford Large-Scale 3D Indoor Spaces Dataset (S3DIS) [5], KITTI [23, 22], and NYU V2 [52]. The datasets showcase a wide range of sizes from the smallest KITTI dataset with 12 million training points, to the largest Semantic3D.net with 1.9 billion training points⁴(details in [63]). We evaluate our method on each dataset and provide a comparison against the state-of-the-art.

6.2. Results

We present quantitative and qualitative results for each of the datasets introduced above. We compare against the state-of-the-art, and perform an ablation study to showcase the benefits of the CRF. The metrics reported are mean

IOU and mean Accuracy across classes unless otherwise stated.

Semantic3D.net benchmark: We evaluate our architecture on the recent Semantic3D.net benchmark [26], which is currently the largest labeled 3D point cloud dataset for outdoor scenes. It contains over 3 billion points and covers a range of urban scenes. We provide results on the *reduced-8* challenge of the benchmark in Table 1. Our method outperforms [6] by 2.2 mIOU points and 2.28% accuracy and sets a new state-of-the-art on that challenge. When compared against the best method that does not leverage extra data through ImageNet [57] pretrained networks, our method outperforms [27] by 7.1 mIOU points, 4.1% accuracy. Note that we also do not leverage extra data or ImageNet [57] pretrained networks. Our base 3D-FCNN trained with Trilinear Interpolation (3D-FCNN-TI) already achieves state-of-the-art performance, and an additional improvement of 3.1 mIOU points and 3.22% can be attributed to the CRF. An example segmentation of our method is shown in Figure 5. The 3D-FCNN-TI produces a segmentation which contains some noise on the cars highlighted in the figure. However, the combination with the CRF in the SEGCloud is able to remove the noise and provide a cleaner segmentation of the point cloud.

Stanford Large-Scale 3D Indoor Spaces Dataset

(S3DIS): The S3DIS dataset [5] provides 3D point clouds for six fully reconstructed large-scale areas, originating from three different buildings. We train our architecture on two of the buildings and test on the third. We compare our method against the MLP architecture of Qi *et al.*, (PointNet) [53]. Qi *et al.* [53] perform a six-fold cross validation across areas rather than buildings. However, with this experimental setup, areas from the same building end up in both the training and test set resulting in increased performance and do not measure generalizability. For a more principled evaluation, we choose our test set to match their fifth fold (*ie.* we test on Area 5 and train on the rest). We obtain the results from the authors for

⁴This excludes the validation set in our data split

Table 3: Results on the NYUV2 dataset

Method	Bed	Objects	Chair	Furniture	Ceiling	Floor	Deco.	Sofa	Table	Wall	Window	Booksh.	TV	mIOU	mAcc	glob Acc
Coupric et al. [14]	38.1	8.7	34.1	42.4	62.6	87.3	40.4	24.6	10.2	86.1	15.9	13.7	6.0	-	36.2	52.4
Wang et al. [65]	47.6	12.4	23.5	16.7	68.1	84.1	26.4	39.1	35.4	65.9	52.2	45.0	32.4	-	42.2	-
Hermans et al. [29]	68.4	8.6	41.9	37.1	83.4	91.5	35.8	28.5	27.7	71.8	46.1	45.4	38.4	-	48.0	54.2
Wolf et al. [69]	74.56	17.62	62.16	47.85	82.42	98.72	26.36	69.38	48.57	83.65	25.56	54.92	31.05	39.51	55.6±0.2	64.9±0.3
3D-FCNN-TI(Ours)	69.3	40.26	64.34	64.41	73.05	95.55	21.15	55.51	45.09	84.96	20.76	42.24	23.95	42.13	53.9	67.38
SEGCloud (Ours)	75.06	39.28	62.92	61.8	69.16	95.21	34.38	62.78	45.78	78.89	26.35	53.46	28.5	43.45	56.43	66.82

Table 4: Results on the KITTI dataset.

Method	building	sky	road	vegetation	sidewalk	car	pedestrian	cyclist	signage	fence	mIOU	mAcc
Zhang et al. [75]	86.90	-	89.20	55.00	26.20	50.0	49.00	19.3	51.7	21.1	-	49.80
3D-FCNN-TI(Ours)	85.83	-	90.57	70.50	25.56	65.68	46.35	7.78	28.40	4.51	35.65	47.24
SEGCloud (Ours)	85.86	-	88.84	68.73	29.74	67.51	53.52	7.27	39.62	4.05	36.78	49.46

comparison shown in Table 2. We outperform the MLP architecture of [53] by 7.83 mIOU points and 8.37% in mean accuracy. Our base 3D-FCNN-TI also outperforms their architecture and the effect of our system’s design choices on the performance of the 3D-FCNN and 3D-FCNN-TI are analyzed in Section 6.3. Qualitative results on this dataset (Figure 5) show an example of how detailed boundaries are captured and refined by our method.

NYU V2: The NYU V2 dataset [52] contains 1149 labeled RGB-D images. Camera parameters are available and are used to obtain a 3D point cloud for each RGB-D frame. In robotics and navigation applications, agents do not have access to fully reconstructed scenes and labeling single frame 3D point clouds becomes invaluable. We compare against 2D and 3D-based methods except for those that leverage additional large scale image datasets (e.g. [35], [18]), or do not use the official split or the 13-class labeling defined in [14] (e.g. [35], [68]). We obtain a confusion matrix for the highest performing method of [69] to compute mean IOU in addition to the mean accuracy numbers they report. Wolf et al. [69] evaluate their method by aggregating results of 10 random forests. Similarly, we use 10 different random initializations of network weights, and use a validation set to select our final trained model for evaluation. Results are shown in Table 3. We outperform the 3D Entangled Forests method of [69] by 3.94 mIOU points and 0.83% mean accuracy.

KITTI: The KITTI dataset [23, 22] provides 6 hours of traffic recording using various sensors including a 3D laser scanner. Zhang et al. [75] annotated a subset of the KITTI tracking dataset with 3D point cloud and corresponding 2D image annotations for use in sensor fusion for 2D semantic segmentation. As part of their sensor fusion process, they train a unimodal 3D point cloud classifier using Random

Forests. We use this classifier as a baseline for evaluating our framework’s performance. The comparison on the labeled KITTI subset is reported in Table 4. We demonstrate performance on par with [75] where a Random Forests classifier is used for segmentation. Note that for this dataset, we train on the laser point cloud with no RGB information.

Analysis of results: In all datasets presented, our performance is on par with or better than previous methods. As expected, we also observe that the addition of a CRF improves the 3D-FCNN-TI output and the qualitative results showcase its ability to recover clear object boundaries by smoothing out incorrect regions in the bilateral space (e.g. cars in Semantic3D.net or chairs in S3DIS). Quantitatively, it offers a relative improvement of 3.0-5.3% mIOU and 4.4-4.7% mAcc for all datasets. Specifically, we see the largest relative improvement on Semantic3D.net - 5.3% mIOU. Since Semantic3D.net is by far the largest dataset (at least 8X times larger), we believe that such characteristic might be representative for large scale datasets as the base networks are less prone to overfitting. We notice however that several classes in the S3DIS dataset, such as *board*, *column* and *beam* are often incorrectly classified as *walls*. These elements are often found in close proximity to walls and have similar colors, which can present a challenge to both the 3D-FCNN-TI and the CRF.

6.3. System Analysis

We analyze two additional components of our framework: geometric data augmentation and trilinear interpolation. The experiments presented in this section are performed on the S3DIS dataset. We also analyzed the effect of joint training versus separate CRF initialization (details and results in supplementary material [63]).

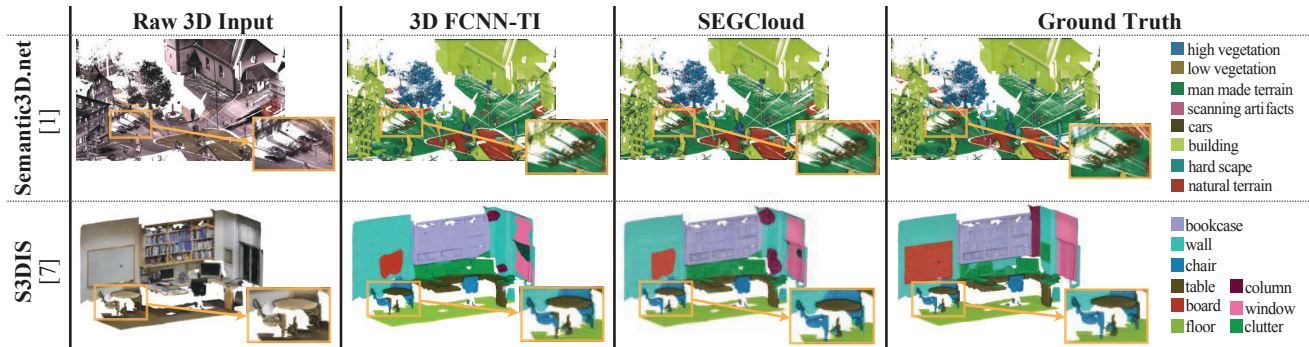


Figure 5: Qualitative results of our framework on Semantic3D.net and S3DIS. Additional results provided in suppl. [63].

Table 5: Effect of Geometric Augmentation

	Method	mIOU
	PointNet [53]	41.09
	Ours- no augm. (3D-FCNN-TI)	43.67
	Ours (3D-FCNN-TI)	47.46

Effect of Geometric Data Augmentation: Our framework uses several types of data augmentation. Our geometric data augmentation methods in particular (random 360° rotation along the z-axis and scaling) are non-standard. Qi *et al.* [53] use different augmentation, including random rotation along the z-axis, and jittering of x, y, z coordinates to augment object 3D point clouds, but it is not specified whether the same augmentation is used on 3D scenes. We want to determine the role of our proposed geometric augmentation methods on the performance of our base 3D-FCNN-TI architecture. We therefore train the 3D-FCNN-TI without any geometric augmentation and report the performance in Table 5. We observe that the geometric augmentation does play a significant role in the final performance and is responsible for an improvement of 3.79 mIOU points. However, even without any geometric augmentation, our base 3D-FCNN-TI outperforms the MLP architecture of [53] by 2.58 mIOU points.

Trilinear interpolation analysis: We now present a study on the effect of trilinear interpolation on our framework. For simplicity, we perform this analysis on the combination of 3D-FCNN and interpolation layer only (no CRF module). We want to study the advantage of our proposed 8-neighbours trilinear interpolation scheme (Section 3.2) over simply assigning labels of points according to the voxel they belong to (see Figure 6 for a schematic explanation of the two methods). The results of the two interpolation schemes are shown in Table 6. We observe that trilinear interpolation helps improve the 3D-FCNN performance by 2.62 mIOU points over simply transferring the voxel label to the points within the voxel. This shows that considering the metric distance between points and voxels, as well a larger neighborhood of voxels can help improve accuracy in predictions.

Table 6: Effect of trilinear interpolation

	Method	mIOU
	PointNet [53]	41.09
	Ours-NN (3D-FCNN-NN)	44.84
	Ours (3D-FCNN-TI)	47.46

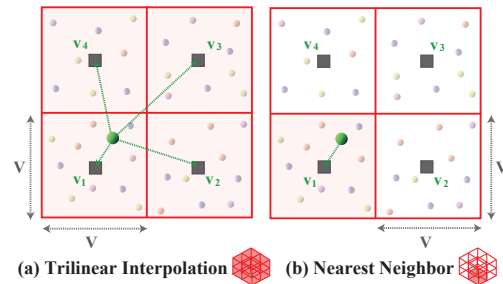


Figure 6: Assigning voxel labels to 3D points (top view): Trilinear interpolation (a) versus the conventional approach of the nearest voxel center (b).

7. Conclusion

We presented an end-to-end 3D Semantic Segmentation framework that combines 3D-FCNN, trilinear interpolation and CRF to provide class labels for 3D point clouds. Our approach achieves performance on par or better than state-of-the-art methods based on neural networks, random forests and graphical models. We show that several of its components such as geometric 3D data augmentation and trilinear interpolation play a key role in the final performance. Although we demonstrate a clear advantage over some Random Forests methods and a point-based MLP method, our implementation uses a standard voxel-based 3D-FCNN and could still adapt to the sparsity of the voxel grid using sparse convolutions (*e.g.* [55]) which could add an extra boost in performance, and set a new state-of-the-art in 3D Semantic Segmentation.

Acknowledgments

We acknowledge the support of Facebook and MURI (1186514-1-TBCJE) for this research.

References

- [1] A. Adams, J. Baek, and M. A. Davis. Fast High-Dimensional Filtering Using the Permutohedral Lattice. *Computer Graphics Forum*, 2010. 5
- [2] A. K. Aijazi, P. Checchin, and L. Trassoudaine. Segmentation based classification of 3d urban point clouds: A super-voxel based approach with evaluation. *Remote Sensing*, 5(4):1624–1650, 2013. 2
- [3] A. Anand, H. S. Koppula, T. Joachims, and A. Saxena. Contextually Guided Semantic Labeling and Search for 3D Point Clouds. *International Journal of Robotics Research*, 32(1):19–34, 2013. 1, 2
- [4] I. Armeni, S. Sax, A. R. Zamir, and S. Savarese. Joint 2D-3D-Semantic Data for Indoor Scene Understanding. *ArXiv e-prints*, Feb. 2017. 5
- [5] I. Armeni, O. Sener, A. Zamir, H. Jiang, and S. Savarese. 3D Semantic Parsing of Large-Scale Indoor Spaces. *CVPR*, pages 1534–1543, 2016. 2, 5, 6
- [6] A. Boulch, B. L. Saux, and N. Audebert. Unstructured Point Cloud Semantic Labeling Using Deep Segmentation Networks. In I. Pratikakis, F. Dupont, and M. Ovsjanikov, editors, *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2017. 6
- [7] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. 1, 2
- [8] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla. Segmentation and recognition using structure from motion point clouds. In *Proceedings of the 10th European Conference on Computer Vision: Part I, ECCV '08*, pages 44–57, Berlin, Heidelberg, 2008. Springer-Verlag. 2
- [9] N. Chehata, L. Guo, and C. Mallet. Airborne lidar feature selection for urban classification using random forests. In *Proceedings of the ISPRS Workshop: Laserscanning09*, pages 207–212, 2009. 2
- [10] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *CoRR*, abs/1606.00915, 2016. 3
- [11] C. B. Choy, J. Gwak, S. Savarese, and M. Chandraker. Universal correspondence network. In *Advances in Neural Information Processing Systems 29*. 2016. 5
- [12] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016. 2
- [13] A. Cohen, A. G. Schwing, and M. Pollefeys. Efficient structured parsing of facades using dynamic programming. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3206–3213, June 2014. 2
- [14] C. Couprie, C. Farabet, L. Najman, and Y. Lecun. *Indoor semantic segmentation using depth information*. 2013. 7
- [15] A. Criminisi and J. Shotton. *Decision Forests for Computer Vision and Medical Image Analysis*. Springer Publishing Company, Incorporated, 2013. 1, 2
- [16] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. *arXiv preprint arXiv:1702.04405*, 2017. 2, 5
- [17] D. Dohan, B. Matejek, and T. Funkhouser. Learning hierarchical semantic segmentations of LIDAR data. In *International Conference on 3D Vision (3DV)*, Oct. 2015. 2
- [18] D. Eigen and R. Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2650–2658, Dec 2015. 7
- [19] H. Fan, H. Su, and L. Guibas. A Point Set Generation Network for 3D Object Reconstruction from a Single Image. *ArXiv e-prints*, Dec. 2016. 2
- [20] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient graph-based image segmentation. *Int. J. Comput. Vision*, 59(2):167–181, Sept. 2004. 2
- [21] A. Garcia-Garcia, F. Gomez-Donoso, J. Garcia-Rodriguez, S. Orts-Escolano, M. Cazorla, and J. Azorin-Lopez. Pointnet: A 3d convolutional neural network for real-time object class recognition. In *2016 International Joint Conference on Neural Networks (IJCNN)*, pages 1578–1584, July 2016. 2
- [22] A. Geiger. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '12, pages 3354–3361, Washington, DC, USA, 2012. IEEE Computer Society. 2, 6, 7
- [23] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The KITTI dataset. *International Journal of Robotics Research*, 32(11):1231 – 1237, Sept. 2013. 2, 6, 7
- [24] G. Georgakis, M. A. Reza, A. Mousavian, P. Le, and J. Kosecka. Multiview RGB-D dataset for object instance detection. *CoRR*, abs/1609.07826, 2016. 5
- [25] K. Guo, D. Zou, and X. Chen. 3d mesh labeling via deep convolutional neural networks. *ACM Trans. Graph.*, 35(1):3:1–3:12, Dec. 2015. 2
- [26] T. Hackel, N. Savinov, L. Ladicky, J. D. Wegner, K. Schindler, and M. Pollefeys. SEMANTIC3D.NET: A NEW LARGE-SCALE POINT CLOUD CLASSIFICATION BENCHMARK. *to appear in ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci 2017*, 2017. 2, 6
- [27] T. Hackel, J. D. Wegner, and K. Schindler. Fast Semantic Segmentation of 3D Point Clouds With Strongly Varying Density. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, III-3(July):177–184, 2016. 2, 6
- [28] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 4
- [29] A. Hermans, G. Floros, and B. 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, May 2014. 7
- [30] H. Hu, D. Munoz, J. A. Bagnell, and M. Hebert. Efficient 3-d scene analysis from streaming data. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2013. 1
- [31] B.-S. Hua, Q.-H. Pham, D. T. Nguyen, M.-K. Tran, L.-F. Yu, and S.-K. Yeung. Scenenn: A scene meshes dataset with annotations. In *International Conference on 3D Vision (3DV)*, 2016. 5

- [32] S. Y. J. Huang. Point Cloud Labeling using 3D Convolutional Neural Network. In *International Conference on Pattern Recognition*, pages 1–6, 2016. 2
- [33] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014. 5
- [34] K. Kamnitsas, C. Ledig, V. F. J. Newcombe, J. P. Simpson, A. D. Kane, D. K. Menon, D. Rueckert, and B. Glocker. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Medical Image Analysis*, 36:61–78, 2017. 3
- [35] B.-S. Kim, P. Kohli, and S. Savarese. 3D Scene Understanding by Voxel-CRF. *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1425–1432, 2013. 1, 7
- [36] H. S. Koppula, A. Anand, T. Joachims, and A. Saxena. Semantic Labeling of 3D Point Clouds for Indoor Scenes. *Neural Information Processing Systems*, pages 1–9, 2011. 1, 2
- [37] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*. 2011. 2, 5
- [38] L. Ladicky, C. Russell, P. Kohli, and P. H. S. Torr. Associative hierarchical random fields. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(6):1056–1077, June 2014. 2
- [39] J. D. Lafferty, A. McCallum, and F. C. N. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning, ICML '01*, pages 282–289, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. 1, 2
- [40] J.-F. Lalonde, N. Vandapel, D. Huber, and M. Hebert. Natural terrain classification using three-dimensional lidar data for ground robot mobility. *Journal of Field Robotics*, 23(10):839 – 861, November 2006. 2
- [41] F. J. Lawin, M. Danelljan, P. Tosteborg, G. Bhat, F. S. Khan, and M. Felsberg. Deep projective 3d semantic segmentation, 2017. 6
- [42] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. *Computer Vision and Pattern Recognition (CVPR)*, 2015. 2, 3
- [43] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June, pages 3431–3440, 2015. 4
- [44] Y. Lu and C. Rasmussen. Simplified markov random fields for efficient semantic labeling of 3D point clouds. In *IEEE International Conference on Intelligent Robots and Systems*, pages 2690–2697, 2012. 1, 2
- [45] A. Martinović, J. Knopp, H. Riemenschneider, and L. Van Gool. 3D all the way: Semantic segmentation of urban scenes from start to end in 3D. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June, pages 4456–4465, 2015. 1, 2
- [46] D. Maturana and S. Scherer. VoxNet: A 3D Convolutional Neural Network for Real-Time Object Recognition. In *IROS*, 2015. 2
- [47] E. Meijering. A chronology of interpolation: From ancient astronomy to modern signal and image processing. *Proceedings of the IEEE*, 90(3):319–342, March 2002. 2
- [48] D. Munoz, J. A. Bagnell, and M. Hebert. Co-inference machines for multi-modal scene analysis. In *European Conference on Computer Vision (ECCV)*, 2012. 1, 2
- [49] D. Munoz, N. Vandapel, and M. Hebert. Directional associative markov network for 3-d point cloud classification. *Fourth international symposium on 3D data processing, visualization and transmission*, pages 1–8, 2008. 2
- [50] D. Munoz, N. Vandapel, and M. Hebert. Onboard contextual classification of 3-d point clouds with learned high-order markov random fields. In *2009 IEEE International Conference on Robotics and Automation*, pages 2009–2016, May 2009. 2
- [51] L. Nan, K. Xie, and A. Sharf. A search-classify approach for cluttered indoor scene understanding. *ACM Trans. Graph.*, 31(6):137:1–137:10, Nov. 2012. 2
- [52] P. K. Nathan Silberman, Derek Hoiem and R. Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. 2, 5, 6, 7
- [53] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *CoRR*, abs/1612.00593, 2016. 2, 6, 7, 8
- [54] 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. 2
- [55] G. Riegler, A. O. Ulusoy, and A. Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 8
- [56] H. Riemenschneider, A. Bódis-Szomorú, J. Weissenberg, and L. Van Gool. *Learning Where to Classify in Multi-view Semantic Segmentation*, pages 516–532. Springer International Publishing, Cham, 2014. 5
- [57] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 6
- [58] R. Shapovalov, A. Velizhev, and O. Barinova. Non-Associative Markov Networks for 3D Point Cloud Classification. *Isprs, XXXVIII-3A*:103–108, 2010. 2
- [59] S. Song, S. P. Lichtenberg, and J. Xiao. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 07-12-June, pages 567–576, 2015. 5
- [60] S. Song and J. Xiao. Deep Sliding Shapes for amodal 3D object detection in RGB-D images. In *CVPR*, 2016. 2
- [61] B. Taskar, V. Chatalbashev, and D. Koller. Learning Associative Markov Networks. *Proc. of the International Conference on Machine Learning*, pages 102–110, 2004. 2

- [62] B. Taskar, C. Guestrin, and D. Koller. Max margin Markov networks. *Neural Information Processing Systems*, 2003. 2
- [63] L. P. Tchapmi, C. B. Choy, I. Armeni, J. Gwak, and S. Savarese. Supplementary Material for SEGCloud: Semantic Segmentation of 3D Point Clouds. <http://segcloud.stanford.edu/supplementary.pdf>. 4, 5, 6, 7, 8
- [64] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4489–4497. IEEE, 2015. 5
- [65] A. Wang, J. Lu, G. Wang, J. Cai, and T.-J. Cham. *Multi-modal Unsupervised Feature Learning for RGB-D Scene Labeling*, pages 453–467. Springer International Publishing, Cham, 2014. 7
- [66] T. Wang, J. Li, and X. An. An efficient scene semantic labeling approach for 3d point cloud. In *ITSC*, pages 2115–2120. IEEE, 2015. 1, 2
- [67] M. Weinmann, B. Jutzi, and C. Mallet. Semantic 3D scene interpretation: A framework combining optimal neighborhood size selection with relevant features. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, pages 181–188, Aug. 2014. 2
- [68] D. Wolf, J. Prankl, and M. Vincze. Fast Semantic Segmentation of 3D Point Clouds using a Dense CRF with Learned Parameters. *Icra*, 2015. 1, 2, 7
- [69] D. Wolf, J. Prankl, and M. Vincze. Enhancing semantic segmentation for robotics: The power of 3-d entangled forests. *IEEE Robotics and Automation Letters*, 1(1):49–56, Jan 2016. 1, 2, 7
- [70] J. Wu, T. Xue, J. J. Lim, Y. Tian, J. B. Tenenbaum, A. Torralba, and W. T. Freeman. Single image 3d interpreter network. In *European Conference on Computer Vision (ECCV)*, 2016. 2
- [71] J. Wu, C. Zhang, T. Xue, W. T. Freeman, and J. B. Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in Neural Information Processing Systems*, pages 82–90, 2016. 2
- [72] J. Xiao, A. Owens, and A. Torralba. Sun3d: A database of big spaces reconstructed using sfm and object labels. *2013 IEEE International Conference on Computer Vision (ICCV)*, 00:1625–1632, 2013. 5
- [73] X. Yan, J. Yang, E. Yumer, Y. Guo, and H. Lee. Perspective transformer nets: Learning single-view 3d object reconstruction without 3d supervision. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 1696–1704. Curran Associates, Inc., 2016. 2
- [74] Q. Zhang, X. Song, X. Shao, R. Shibasaki, and H. Zhao. Category modeling from just a single labeling: Use depth information to guide the learning of 2d models. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pages 193–200, June 2013. 5
- [75] R. Zhang, S. A. Candra, K. Vetter, and A. Zakhor. Sensor fusion for semantic segmentation of urban scenes. *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1850–1857, 2015. 2, 7
- [76] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. S. Torr. Conditional Random Fields as Recurrent Neural Networks. *International Conference on Computer Vision*, pages 1529–1537, 2015. 3, 5