This CVPR Workshop paper is the Open Access version, provided by the Computer Vision Foundation.

Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.



Two Stream 3D Semantic Scene Completion

Martin Garbade¹, Yueh-Tung Chen², Johann Sawatzky¹, and Juergen Gall¹

¹University of Bonn, {garbade, sawatzky, gall}@iai.uni-bonn.de ²b-it, RWTH Aachen University, yuehtung.chen@gmail.com

Abstract

Inferring the 3D geometry and the semantic meaning of surfaces, which are occluded, is a very challenging task. Recently, a first end-to-end learning approach has been proposed that completes a scene from a single depth image. The approach voxelizes the scene and predicts for each voxel if it is occupied and, if it is occupied, the semantic class label. In this work, we propose a two stream approach that leverages depth information and semantic information, which is inferred from the RGB image, for this task. The approach constructs an incomplete 3D semantic tensor, which uses a compact three-channel encoding for the inferred semantic information, and uses a 3D CNN to infer the complete 3D semantic tensor. In our experimental evaluation, we show that the proposed two stream approach substantially outperforms the state-of-the-art for semantic scene completion.

1. Introduction

Humans quickly infer the 3D semantics of a scene, i.e., an estimate of the 3D geometry and the semantic meaning of the surfaces. While RGB-D sensors in combination with CNNs provide geometry and semantic information, the resulting representation is very sparse since large parts of the 3D scene are occluded and not visible. The perception, however, is not limited to the visible part of the scene. When looking at a mug on a table, a human can estimate the full geometry of both objects including parts which are invisible since they are occluded by the objects themselves. This information is obtained from semantic understanding of the scene which allows to estimate the spatial extent of the objects from experience. Such an ability is highly desirable for autonomous agents, e.g., to navigate or interact with objects. A robot that has an intuition about the geometry behind the surface it sees, for example, could plan ahead given a single view instead of exhaustively explore the occluded parts of a scene first.

In this work, we aim to estimate the semantics not only of the visible part, but of the entire scene including the occluded space. To this end, we build on the work of Song et al. [36]. They show that semantic scene understanding and 3D scene completion benefit from each other. On one hand, recognizing a part of the object helps to estimate its location in the 3D space and the voxels it occupies. On the other hand, knowing the occupancy in the 3D space gives information on form and size of the object and thus facilitates semantic recognition. For estimating for each voxel in the scene the occupancy and semantic label, they proposed an end-to-end trainable 3D convolutional neural network (3D CNN) which incorporates context from a large field of view via dilated convolutions. The approach, however, only uses depth as input and neglects the RGB image. This means that the semantic label has to be inferred from the geometry alone and properties such as color, texture, or reflectance are not taken into account.

We therefore extend the approach [36] by keeping its beneficial context incorporation and end-to-end trainability while modifying it to leverage semantic information inferred from the RGB image at the input stage as well as at the loss. Given a single RGB-D image, we first use a 2D CNN to infer the semantic labels from the RGB data and construct an incomplete 3D semantic tensor. To this end, we map the inferred semantic labels to the 3D space and label each visible surface voxel by the inferred class label. The 3D semantic tensor is incomplete since it only contains the labels of the visible voxels but not of the occluded voxels. The 3D projection is performed using the depth image. The tensor is then used as input for a 3D CNN that infers a complete 3D semantic tensor, which includes the occupancy and semantic labels for all voxels.

Using the RGB images as input leads to a significant performance gain in scene completion and semantic scene completion as our experiments show. We outperform [36] by a substantial margin of up to 9.4 % on NYU. This implies that RGB images provide a rich discriminative signal.

2. Related Work

Several works address the problem of semantic segmentation of RGB-D images, e.g., [30, 21, 11, 34], but they infer semantic labels only for the visible pixels of the image, which means that occluded voxels are not reconstructed.

A possible strategy for semantic scene completion is the generation of 3D object proposals and subsequent 3D shape completion of the respective object. The problem of completing the 3D shape has been addressed in several works [31, 37, 38, 40, 39, 41, 13, 7]. In the context of inferring a voxel-wise segmentation, holes between objects have to be filled. As long as these missing parts are small, they can be filled using plane fitting [28] or object symmetry [19, 26]. Objects that are not detected, however, heavily disturb the 3D scene completion. Completing the scene geometry without predicting the semantics has been addressed by [9]. Their model assumes that objects of semantically dissimilar classes can still be represented by similar 3D shapes, i.e., it is possible to predict the unobserved voxels from the frontal geometry. However, this approach fails for complex scenes where these geometric constraints are violated.

An alternative is to use instances of 3D mesh models and fit them to the scene geometry [10, 12, 19, 22, 27, 29, 35, 32, 23, 33]. The mesh models, however, do not model shape variations of objects of the same category and increasing the number of mesh models per category is not practical since the number of available CAD models is limited and the shape retrieval becomes more expensive. The approaches [17, 24] even neglect fine-grained details and simply fit 3D primitives to the scene.

Various contextual cues proved to be helpful for semantic scene completion. While physical reasoning is employed in [44], [18] predict voxel labels with a conditional random field (CRF) whose unary potentials are determined by floor plans. The CRF, however, only models contextual information within a short distance. In [1] and [15], semantic scene completion and multi-view reconstruction are jointly performed. The approaches do not rely on end-to-end learning approaches, but they use predefined features and heuristics to integrate context information.

To facilitate learning of scene completion in an end-toend manner, [6, 43, 2] collected large scale datasets with real world data. Previously, synthetic datasets were employed to provide ground truth data for object completion [3, 40] or for entire scenes [14]. Due to these datasets, training of an end-to-end approach for semantic scene completion became feasible [36]. Parallel to our work, [42] have proposed a new network architecture which leverages sparse feature map encodings and allows for much deeper network architectures. While we address the problem of scene completion from a single viewpoint as in [36, 42], semantic scene completion from multiple RGB-D images is addressed in [8].

3. Two Stream Semantic Scene Completion

3.1. Semantic Scene Completion

The goal of 3D semantic scene completion is to classify every voxel in the view frustum into one of K + 1 labels $c = c_0, ..., c_K$ where c_0 represents an empty voxel and $c_1, ..., c_K$ represents one of K = 11 class labels like ceiling, floor, wall, window, chair, bed, sofa, table, tv, furniture and object. As illustrated in Figure 1, the camera observes only a part of the scene while other voxels are occluded. The occluded voxels can either be empty (c_0) or belong to one of the K classes.

To address the task of 3D semantic scene completion, we propose an approach that leverages two input streams, namely RGB and depth. An overview of the approach is given in Figure 2 a). While the depth data is converted into a volumetric representation, the RGB image is first processed in a separate branch to infer 2D semantic segmentation maps and then transformed into a volumetric representation referred to as color-volume (Section 3.3). The volumetric representation is then fed to a 3D convolutional neural network (3D-CNN). The 3D-CNN infers a 3D semantic tensor where every voxel is classified as either being empty or belonging to one of the K semantic classes. In the following, each step will be discussed in detail.

3.2. Depth Input Stream

To obtain the volumetric input encoding, the depth map is projected into a regular voxel grid using the camera pose, which is provided along with each image. The voxel grid is of size $240 \times 144 \times 240$ voxels and encodes a scene of 4.80m horizontally, 2.88m vertically, and 4.80m in depth with a resolution of 0.02m. For every pixel in the depth map, its corresponding voxel in the 3D input volume is computed using the camera pose. The obtained binary voxel mask encodes the location of surface points that are visible to the camera, see Figure 2 b).

As pre-processing, all 3D scenes are rotated such that the room orientations are aligned. For indoor room scenes, one can assume that most of the observed surface normals are oriented either like the normals of the walls, floor or ceiling, which are usually planar. Therefore a principal component analysis of the surface normals is used to infer the room orientation, which is used to align the scene.

3.3. Color Input Stream

The input RGB image is first processed by a 2D-CNN [4], which is an adaptation of the Resnet101 architecture [16] for semantic segmentation. While all but one pooling layer are omitted, dilated convolutions are used to keep the output resolution high while simultaneously increasing the receptive field. The 2D-CNN predicts the softmax probabilities for every class and pixel at a resolution which is



Figure 1. Using the protocol of [36], ground truth labels are provided for all voxels of a 3D volume. Voxels that are outside the intersection of the camera frustum and ground truth volume are outside the room or outside the view and not taken into account. Within the intersection, there are observed surface voxels (green) and observed non-occupied voxels (light gray), but other voxels are not observed by the camera. These voxels are either non-occupied (blue) or belong to an object (black).



Figure 2. a) The proposed two stream approach for semantic scene completion transforms first the depth data and RGB image into a volumetric representation, which represents the geometry and semantic of the visible scene and then uses a 3D-CNN to infer a 3D semantic tensor for the entire scene. b) Given 2D depth map and camera pose, a binary voxel mask is created by setting each voxel that belongs to a depth pixel to one and all other voxels to zero (blue). c) Visualization of TSDF vs. flipped TSDF. One can see the long 'shadow' caused by the observed surface which produces high gradients at the occlusion boundary (between -1 and 1). In the flipped TSDF, this effect is suppressed. The gradient is highest at the surface.

four times smaller than the input image. The output is then upsampled to the original resolution of the image using bilinear interpolation. A densely connected CRF [20] is then used in combination with the inferred class probabilities and the RGB image to refine the semantic segmentation map. For training, we use the same setting as in [4]. As initialization, we use a model that is pre-trained on MSCOCO [25] and fine-tune it on the dataset for 3D semantic scene completion. Furthermore, we present results using the more recent model Deeplab v3+[5] which is pretrained on ADE-20k and finetune it on NYUv2 using an initial learning rate of 0.001. We also apply a CRF [20] on the resulting outputs.

As in Section 3.2, we convert the 2D segmentation map into a volumetric representation. Since each pixel in the depth map corresponds to a pixel in the 2D semantic segmentation map, every class pixel can be projected into the 3D volume at the location of its corresponding depth value. This yields an incomplete 3D semantic tensor that assigns to every surface voxel its corresponding class label. The class labels can be encoded by one-hot encoding, i.e., a channel for each class, or by a single channel for the class label. In our experiments, however, we show that none of them is optimal. Encoding semantic classes with only one channel implies a semantic proximity of classes by the numerical proximity of their class values, which introduces undesirable artifacts based on the class values. The one-hot encoding has the disadvantage that it is insufficient in terms of memory consumption since it requires to store a K dimensional vector per voxel. We therefore represent the semantic information by a lower dimensional vector. We use

a three-dimensional vector and encode the classes linearly from (0,0,1) over (0,1,1), (0,1,0), (1,1,0) to (1,0,0).

3.4. 3D-CNN

For the 3D-CNN, we adapt the architecture of [36] by increasing the number of input channels of the first convolutional layer such that it fits to our input. The architecture is illustrated in Figure 3. It is inspired by the 2D-CNN for semantic segmentation. The major difference apart from using 3D instead of 2D convolutions is that the network only has a depth of 14 convolutional layers. The network has therefore significantly less parameters than its two dimensional counter part. Moreover, batch-normalization layers are omitted due to the small size of the batches.

We adapt the training protocol of [36] as follows. We train for 150,000 steps with a learning rate of 0.01 that is reduced by a factor of 0.1 after 100,000 iterations. As optimizer, stochastic gradient (SGD) with momentum is applied. As initialization, we chose a random initialization with a Gaussian distribution with mean $\mu = 0$ and a standard deviation of $\sigma = 0.01$.

The output of the 3D-CNN is a semantic tensor of size 60 x 36 x 60 x (K + 1), where K is the number of object classes and an additional class is added for empty voxels. We compute a softmax cross entropy loss on the unnormalized network outputs y:

$$\mathcal{L} = -\sum_{i,c} w_{ic} \hat{y}_{ic} \log\left(\frac{e^{y_{ic}}}{\sum\limits_{c' \in \mathcal{C}} e^{y_{ic'}}}\right) \tag{1}$$

where \hat{y}_{ic} are the binary ground truth vectors, i.e., $\hat{y}_{ic} = 1$ if voxel *i* is labeled by class *c*, and w_{ic} are the loss weights. Since the ratio of empty vs. occupied voxels is 9:1, the empty space is randomly subsampled. Therefore w_{ic} is chosen as binary mask such that only 2N empty voxels are selected for loss calculation where N is the number of occupied voxels in the scene.

4. Experimental Evaluation

4.1. Evaluation Metric

For evaluation, we follow the evaluation protocol of [36], which evaluates the accuracy on a subset of voxels. The evaluation considers only voxels that are part of the occluded space and within both the room and the field-of-view as shown in Figure 1. While generating the 3D semantic labels from the annotated CAD models, every voxel in the input volume is marked as being on surface, free space, occluded space, outside field of view, outside room or outside ceiling. For semantic scene completion, a binary evaluation mask is computed such that the evaluation metric is only computed for voxels which are either occluded, on surface or close to the surface (within the range of the TSDF function defined by [36]). For scene completion another mask is computed which comprises all voxels in the occluded space. To assess the quality of 3D scene completion, several metrics are computed. First we compute the Jaccard index, which measures the intersection over union (IoU) between ground truth and predicted voxel for every object category $c_1, ..., c_K$. As an overall segmentation performance, we compute the average across all classes. For scene completion all voxels are considered to belong to one of the two classes empty vs. non-empty. All object categories $c_1, ..., c_K$ are counted as 'non-empty'. For completion, IoU as well as precision and recall are computed.

4.2. Datasets

We evaluate our method on the NYUv2 dataset, which is in the following denoted as NYU. NYU consists of indoor scenes that are captured via a Kinect sensor. For 3D semantic scene completion labels, we use the annotated 3D labels provided by [31]. They provide 1449 scenes, annotated with 11 classes, 795 of which are used for training and 654 for testing. These annotations consist of CAD models that are fitted into the scene. Since the CAD models do not exactly fit the shape of the annotated objects and neglect small objects such as clutter, there is a significant mismatch between the Kinect input data and the output labels. To address this problem, depth maps generated from the projections of the 3D annotations as in [31] are used for training. For evaluation, we consider two test sets. The first test set, which is denoted by NYU Kinect, consists of the depth maps from the Kinect sensor and the second test set, which is denoted as NYU CAD, uses the depth maps generated by projection.

4.3. Ablation Study

We conduct an ablation study on Kinect to analyze the design choices of our model.

4.3.1 Effect of Semantic Input

As mentioned in Section 3.3, we compare two network architectures for the 2D-CNN, namely Deeplab v2 (DLv2) [4] and Deeplab v3+ (DLv3+) [5]. Table 1 shows that DLv3+ increases the accuracy for semantic scene completion from 31.3 % to 33.8 %. This is expected since DLv3+ provides a better 2D segmentation accuracy compared to DLv2 as shown in Table 2. For scene completion, IoU is slightly higher for DLv2 than for DLv3+.

We compare the results to a setting when we use ground truth semantic segmentation masks for the RGB images as input to the 3D-CNN, which is denoted by GT in Table 1. This also serves as an upper bound for our method when the used 2D-CNN provides perfect segmentation masks. As expected, using ground truth segmentation masks improves the semantic scene completion compared to DLv3+



Figure 3. Architecture of the 3D-CNN. The parameters of the convolution kernels are denoted as (number of filters, kernel size, stride, dilation). All but the last convolution layer have a ReLU activation function assigned to it. Arrows indicate skip connections [16] where the output of one convolution layer is added to another output at a later stage. Pool denotes max pooling. The output is a volume that is 4 fold downsampled with respect to the input of the 3D CNN and encodes for every voxel the probability of it being empty (label 0) or to belong to one of K semantic classes.

	Scene Completion		Semantic Scene Completion											
semantic	IoU	ceil.	floor	wall	win.	chair	bed	sofa	table	tv	furn.	objs	avg	
DLv2	60.1	7.0	93.1	25.9	16.8	14.7	53.3	46.0	16.8	22.7	34.1	13.8	31.3	
DLv3+	60.0	6.7	93.2	26.2	20.6	17.8	56.9	49.2	16.5	29.4	37.6	18.3	33.8	
GT	61.0	6.6	93.4	27.5	24.6	23.0	61.6	57.9	24.3	33.5	46.4	24.5	38.5	
RGB image	58.2	4.2	93.4	19.3	4.4	10.8	34.9	20.2	11.8	4.9	17.2	10.3	21.0	
T-11-1 I	- f the second the second		mand D		(D	CD :		ID CNIN	T ::4		41 1	1	- £ 41	

Table 1. Impact of the quality of the semantic input. For the version 'RGB image', the 2D-CNN is omitted and the color values of the pixels instead of the semantic information is stored in the semantic volume.

by +4.7 %. For scene completion, the improvements compared to DLv3+ are +1.0 %. This shows that the quality of the 2D-CNN has a strong impact on the accuracy of semantic scene completion but only a minor impact on scene completion, which is also not the focus of this work.

As it is illustrated in Figure 2 a), the RGB images are processed by a 2D-CNN and the inferred pixel-wise labels are used to construct the semantic volume. We also evaluate in Table 1 what happens if the three-channel encoding is not based on the semantic labels but if the RGB values of the pixels are directly used for the encoding, i.e., without inferring semantic information from the visible part of the scene. This setting is denoted by 'RGB image' and performs as expected poorly. Besides of the class 'floor', all categories are poorly estimated.

4.3.2 Input Encoding

As we discussed in Section 3.3, the encoding of the semantic information in the semantic volume should provide a numerical equidistance between classes, which can be achieved by using one-hot encoding. However, this approach has a high memory footprint. As an alternative, we evaluate a one-channel and a three-channel input encoding. In the one-channel setup, the numeric class values are normalized to the range from 0 to 1. For the proposed 3-channel input encoding, every label is mapped to a 3 dimensional vector as described in Section 3.3. Table 3 shows that using only one channel performs poorly since it introduces undesirable artifacts based on the class values. While some classes like 'floor' and 'bed' are well recognized, the accuracy for 'window' and 'tv' is very low. Using one-hot encoding (12 channels) performs much better than 1 channel but it is expensive in terms of memory consumption. The proposed three-channel encoding requires less memory while it only slightly decreases the accuracy. Also the training time of the 3 channel setup is by a factor of 1.7 faster which reduces the training time from 4 to 2.5 days. Therefore we adopt the 3 channel setup as it provides an efficient alternative to the one-hot encoding.

4.3.3 Fusion with flipped TSDF

Furthermore, we have conducted an experiment where we combine our input with the flipped truncated signed distance function (fTSDF) proposed by [36] and evaluate different fusing schemes.

The fTSDF is computed as follows: The previously computed binary voxel mask (Figure 2 b) is used to first compute a truncated signed distance function (TSDF) encoding as illustrated in Figure 2 c). In the TSDF, every voxel contains as value the distance d to the next surface point. The sign of the distance value indicates whether a voxel lies in the empty (1) or occluded space (-1). The TSDF has the disadvantage of having high gradients at the occlusion boundary, i.e., the boundary between observed and unobserved

	ceil.	floor	wall	win.	chair	bed	sofa	table	tvs	furn.	objs.	avg.
Deeplab v2	58.1	85.7	76.6	62.9	58.5	65.8	62.8	37.9	56.8	56.5	54.7	61.5
Deeplab v3+	71.1	89.8	82.8	72.8	65.8	72.4	66.1	50.7	63.0	64.7	62.9	69.3
Table 2. 2D	semantic	e segmen	tation ac	curacies	s on the N	VYUv2	lataset (% IoU).	In both o	cases, a C	CRF is us	ed.

	Scene Completion		Semantic Scene Completion											
channels	IoU	ceil.	floor	wall	win.	chair	bed	sofa	table	tv	furn.	objs	avg	
1	59.3	8.3	93.3	25.0	13.4	11.5	43.0	31.8	11.2	2.4	26.5	16.8	25.8	
3	60.0	6.7	93.2	26.2	20.6	17.8	56.9	49.2	16.5	29.4	37.6	18.3	33.8	
12 (one-hot)	60.0	9.7	93.4	25.5	21.0	17.4	55.9	49.2	17.0	27.5	39.4	19.3	34.1	
Ta	ble 3. Impact of the nun	ber of d	channels	for the s	emantic	volume.	12 chan	nels refe	ers to one	-hot enc	oding.			

space behind a surface. Therefore in the TSDF encoding every surface yields a shadow into the unobserved space as shown in Figure 2 c).

To provide a more meaningful input signal, the signed distance function is transformed into a flipped TSDF [36], where every signed distance value d is converted into a distance d_f which is 1 or -1 at a surface and linearly falls to 0 at a distance d_{max} from the surface:

$$d_f = sign(d)\mathcal{H}(d_{max} - |d|)\frac{d_{max} - |d|}{d_{max}}$$
(2)

where d_{max} is the maximum distance of 24 cm and \mathcal{H} is the Heaviside function:

$$\mathcal{H}(x) = \begin{cases} 1 & \text{if } x \ge 0\\ 0 & \text{if } x < 0. \end{cases}$$
(3)

We perform different fusion experiments to evaluate whether the proposed fTSDF encoding can give us a meaningful signal for semantic scene completion. As one can see from Figure 3 the 3D-CNN consists of several blocks. We concatenate the fTSDF before block 1 (early fusion) and also after block 1, 2 and 5 (fusion 1, 2 and 5). Before concatenation, both the color and the fTSDF input stream are processed separately. In the case of "late fusion" we take the maximum of the softmax probabilities of both streams.

As can be seen from Table 4, all fusion schemes perform slightly worse than our approach. This indicates that fTSDF provides a superfluous signal for our approach. This is interesting since computing the flipped TSDF volume is the most time-consuming part for inference and our approach provides a substantial faster alternative while also increasing the accuracy.

4.4. Comparison to State-of-the-Art

We evaluate our approach on the two test sets NYU CAD and NYU Kinect, which are in the following denoted as CAD and Kinect, and we compare our approach to the stateof-the-art. The results for scene completion and semantic scene completion are reported in Table 5. Both of our approaches with 3-channel input encoding and one-hot encoding perform comparably. Since one-hot encoding yields a slightly higher accuracy, we only discuss the difference between the latter and other approaches from the literature.

Our approach sets the new state-of-the-art for semantic scene completion. We achieve 46.2 % on CAD and 34.1 %on Kinect and outperform the approach by Song et al. [36] by +6.2 % on CAD and +3.6 % on Kinect, although they use SUNCG as additional training data. If the same training data, i.e. only NYU, is used, our approach outperforms [36] by +9.4 % on Kinect. For scene completion, we outperform [36] by +5.8 % on CAD and +4.9 % on Kinect if NYU is used as training data. However, even if [36] uses additional training data from SUNCG, our approach still outperforms it by +2.9 % on CAD and +3.4 % on Kinect. Compared to the recent ESSCN approach [42], we perform better in both scene completion (+3.8 %) and semantic scene completion (+7.4 %). Note also that pretraining on SUNCG [36] and using a stronger 3D-CNN architecture [42] are orthogonal to our proposed method. One can assume that our performance would further increase by incorporating both ideas.

Table 5 also includes the results of other approaches that do not rely on end-to-end learning [44, 9, 24, 10]. Furthermore, the approaches [44] and [9] only address scene completion but not semantic scene completion. These methods perform substantially worse than the end-to-end learning approaches.

5. Conclusions

In this work, we have proposed a two stream approach for 3D semantic scene completion. In contrast to previous works, the proposed approach leverages depth and semantic information of the visible part of the scene for this task. In our experiments, we have shown that the proposed three-channel encoding for the semantic volume is not only memory efficient but it also results in higher accuracies compared to a single-channel encoding and is competitive to a memory expensive one-hot encoding. The proposed approach achieves state-of-the-art results for semantic scene completion on the NYUv2 dataset while also providing much faster inference times than approaches based of TSDF input features.

	Scene Completion					Seman	tic Scei	ne Com	pletion				
input	IoU	ceil.	floor	wall	win.	chair	bed	sofa	table	tv	furn.	objs	avg
proposed, no fTSDF	60.0	6.7	93.2	26.2	20.6	17.8	56.9	49.2	16.5	29.4	37.6	18.3	33.8
with fTSDF, early fusion	60.2	8.1	94.4	25.6	17.1	17.8	53.6	48.0	17.0	28.0	36.0	18.4	33.1
with fTSDF, fusion 1	60.0	4.8	94.1	25.5	21.5	16.6	56.9	47.2	16.7	27.5	37.3	18.1	33.3
with fTSDF, fusion 2	54.4	5.9	93.6	22.0	11.0	16.5	50.4	41.0	12.4	23.1	31.9	12.4	29.1
with fTSDF, fusion 5	59.1	5.1	92.9	23.0	19.4	15.1	53.9	46.7	16.3	28.2	34.6	15.0	31.8
with fTSDF, late fusion	60.4	5.7	93.9	25.7	20.3	15.9	55.7	44.8	17.0	28.1	34.9	16.0	32.5
RGB image	58.2	4.2	93.4	19.3	4.4	10.8	34.9	20.2	11.8	4.9	17.2	10.3	21.0
Table 4. Impact of the input for	r the 3D-CNN. The pro	posed a	rchitectu	re is sho	wn in F	igure 2 a). The v	versions	'with fT	SDF' re	fers to a		

version where not only the semantic volume but also the flipped TSDF volume [36] are used.

NYU CAD		Scene Completion					Seman	tic Scer	ne Com	pletion				
method	trained on	IoU	ceil.	floor	wall	win.	chair	bed	sofa	table	tv	furn.	objs	avg
Zheng et al. [44]	NYU	34.6												
Firman et al. [9]	NYU	50.8												
SSCNet [36]	NYU	70.3												
SSCNet [36]	SUNCG+NYU	73.2	32.5	92.6	49.2	8.9	33.9	57.0	59.5	28.3	8.1	44.8	25.1	40.0
Ours, 3ch	NYU	76.1	28.3	94.0	48.6	33.0	33.4	67.9	54.7	31.1	33.8	50.8	30.6	46.0
Ours, one-hot	NYU	76.1	25.9	93.8	48.9	33.4	31.2	66.1	56.4	31.6	38.5	51.4	30.8	46.2
		1												
NYU Kinect		Scene Completion					Seman	tic Scer	ne Comp	pletion				
NYU Kinect Lin et al. [24]	NYU	Scene Completion 36.4	0.0	11.7	13.3	14.1	Seman 9.4	tic Scer 29.0	ne Comj 24.0	pletion 6.0	7.0	16.2	1.1	12.0
NYU Kinect Lin et al. [24] Geiger et al.[10]	NYU NYU	Scene Completion 36.4 44.4	0.0 10.2	11.7 62.5	13.3 19.1	14.1 5.8	Seman 9.4 8.5	tic Scer 29.0 40.6	ne Comp 24.0 27.7	pletion 6.0 7.0	7.0 6.0	16.2 22.6	1.1 5.9	12.0 19.6
NYU Kinect Lin et al. [24] Geiger et al.[10] SSCNet[36]	NYU NYU NYU	Scene Completion 36.4 44.4 55.1	0.0 10.2 15.1	11.7 62.5 94.7	13.3 19.1 24.4	14.1 5.8 0.0	Seman 9.4 8.5 12.6	tic Scer 29.0 40.6 32.1	ne Comp 24.0 27.7 35.0	0.0 6.0 7.0 13.0	7.0 6.0 7.8	16.2 22.6 27.1	1.1 5.9 10.1	12.0 19.6 24.7
NYU Kinect Lin et al. [24] Geiger et al.[10] SSCNet[36] SSCNet[36]	NYU NYU NYU SUNCG+NYU	Scene Completion 36.4 44.4 55.1 56.6	0.0 10.2 15.1 15.1	11.7 62.5 94.7 94.6	13.3 19.1 24.4 24.7	14.1 5.8 0.0 10.8	Seman 9.4 8.5 12.6 17.3	tic Scer 29.0 40.6 32.1 53.2	ne Comp 24.0 27.7 35.0 45.9	6.0 7.0 13.0 15.9	7.0 6.0 7.8 13.9	16.2 22.6 27.1 31.1	1.1 5.9 10.1 12.6	12.0 19.6 24.7 30.5
NYU Kinect Lin et al. [24] Geiger et al.[10] SSCNet[36] SSCNet[36] ESSCN [42]	NYU NYU NYU SUNCG+NYU NYU	Scene Completion 36.4 44.4 55.1 56.6 56.2	0.0 10.2 15.1 15.1 17.5	11.7 62.5 94.7 94.6 75.4	13.3 19.1 24.4 24.7 25.8	14.1 5.8 0.0 10.8 6.7	Seman 9.4 8.5 12.6 17.3 15.3	tic Scer 29.0 40.6 32.1 53.2 53.8	ne Comp 24.0 27.7 35.0 45.9 42.4	6.0 7.0 13.0 15.9 11.2	7.0 6.0 7.8 13.9 0.0	16.2 22.6 27.1 31.1 33.4	1.1 5.9 10.1 12.6 11.8	12.0 19.6 24.7 30.5 26.7
NYU Kinect Lin et al. [24] Geiger et al.[10] SSCNet[36] SSCNet[36] ESSCN [42] Ours, 3ch	NYU NYU NYU SUNCG+NYU NYU NYU	Scene Completion 36.4 44.4 55.1 56.6 56.2 60.0	0.0 10.2 15.1 15.1 17.5 6.7	11.7 62.5 94.7 94.6 75.4 93.2	13.3 19.1 24.4 24.7 25.8 26.2	14.1 5.8 0.0 10.8 6.7 20.6	Seman 9.4 8.5 12.6 17.3 15.3 17.8	tic Scer 29.0 40.6 32.1 53.2 53.8 56.9	ne Comp 24.0 27.7 35.0 45.9 42.4 49.2	6.0 7.0 13.0 15.9 11.2 16.5	7.0 6.0 7.8 13.9 0.0 29.4	16.2 22.6 27.1 31.1 33.4 37.6	1.1 5.9 10.1 12.6 11.8 18.3	12.0 19.6 24.7 30.5 26.7 33.8
NYU Kinect Lin et al. [24] Geiger et al.[10] SSCNet[36] SSCNet[36] ESSCN [42] Ours, 3ch Ours, one-hot	NYU NYU SUNCG+NYU NYU NYU NYU	Scene Completion 36.4 44.4 55.1 56.6 56.2 60.0 60.0	0.0 10.2 15.1 15.1 17.5 6.7 9.7	11.7 62.5 94.7 94.6 75.4 93.2 93.4	13.3 19.1 24.4 24.7 25.8 26.2 25.5	14.1 5.8 0.0 10.8 6.7 20.6 21.0	Seman 9.4 8.5 12.6 17.3 15.3 17.8 17.8	tic Scer 29.0 40.6 32.1 53.2 53.8 56.9 55.9	ne Comp 24.0 27.7 35.0 45.9 42.4 49.2 49.2	6.0 7.0 13.0 15.9 11.2 16.5 17.0	7.0 6.0 7.8 13.9 0.0 29.4 27.5	16.2 22.6 27.1 31.1 33.4 37.6 39.4	1.1 5.9 10.1 12.6 11.8 18.3 19.3	12.0 19.6 24.7 30.5 26.7 33.8 34.1

Acknowledgement

A special thanks goes to Christian Grund for visualizing the qualitative results in Blender. The work has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) GA 1927/2-2 (FOR 1505 Mapping on Demand).

References

- M Blaha, C Vogel, A Richard, D Wegner, T Pock, and K Schindler. Large-scale semantic 3d reconstruction: An adaptive multi-resolution model for multi-class volumetric labeling. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [2] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from RGB-D Data in Indoor Environments. In *International Conference* on 3D Vision, 2017. 2
- [3] Angel X. Chang, Thomas A. Funkhouser, Leonidas J. Guibas, Pat Hanrahan, Qi-Xing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. In *CoRR*, *abs*/1512.03012, 2015. 2
- [4] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs. In *International Conference on Learning Representations*, 2015. 2, 3, 4

- [5] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *CoRR*, *abs*/1802.02611, 2018. 3, 4
- [6] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [7] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape Completion using 3D-Encoder-Predictor CNNs and Shape Synthesis. In *IEEE Conference on Computer Vision* and Pattern Recognition, 2017. 2
- [8] A Dai, D Rithie, M Bokeloh, S Reed, J Sturm, and M Nießner. ScanComplete: Large-Scale Scene Completion and Semantic Segmentation for 3D Scans. In *IEEE Conference* on Computer Vision and Pattern Recognition, 2018. 2
- [9] Michael Firman, Oisin Mac Aodha, Simon Julier, and Gabriel J. Brostow. Structured Prediction of Unobserved Voxels From a Single Depth Image. In *IEEE Conference* on Computer Vision and Pattern Recognition, 2016. 2, 6, 7
- [10] Andreas Geiger and Chaohui Wang. Joint 3D Object and Layout Inference from a single RGB-D Image. In *German Conference on Pattern Recognition*, volume 9358, pages 183–195, 2015. 2, 6, 7
- [11] Saurabh Gupta, Pablo Arbelaez, and Jitendra Malik. Perceptual organization and recognition of indoor scenes from RGB-D images. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 564–571, 2013. 2



Figure 4. Qualitative results on NYUv2 Kinect. From left to right: Input RGB-D image, ground truth, result obtained by [36], and result obtained by our approach. Overall, our completed semantic 3D scenes are less cluttered and show a higher voxel class accuracy compared to [36].

- [12] Saurabh Gupta, Pablo Andrés Arbeláez, Ross B. Girshick, and Jitendra Malik. Aligning 3D models to RGB-D images of cluttered scenes. In *IEEE Conference on Computer Vision* and Pattern Recognition, pages 4731–4740, 2015. 2
- [13] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-Resolution Shape Completion Using Deep Neural Networks for Global Structure and Local Geometry Inference. In *IEEE International Conference on Computer Vision*, 2017. 2
- [14] Ankur Handa, Viorica Patraucean, Vijay Badrinarayanan, Simon Stent, and Roberto Cipolla. SceneNet: Understanding Real World Indoor Scenes With Synthetic Data. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2
- [15] Christian Häne, Christopher Zach, Andrea Cohen, Roland Angst, and Marc Pollefeys. Joint 3D Scene Reconstruction and Class Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 97–104, 2013. 2
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016. 2, 5
- [17] H Jiang and J Xiao. A linear approach to matchig cuboids in RGBD images. In *IEEE Conference on Computer Vision* and Pattern Recognition, 2013. 2
- [18] Byung-soo Kim, Pushmeet Kohli, and Silvio Savarese. 3D scene understanding by Voxel-CRF. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2013. 2
- [19] Young Min Kim, Niloy J. Mitra, Dong-Ming Yan, and Leonidas Guibas. Acquiring 3D Indoor Environments with Variability and Repetition. In ACM Transactions on Graphics, volume 31, pages 138:1–138:11, 2012. 2
- [20] Philipp Krähenbühl and Vladlen Koltun. Efficient Inference in Fully Connected CRFs with Gaussian Edge Potentials. In *Neural Information Processing Systems*, 2011. 3
- [21] Kevin Lai, Liefeng Bo, and Dieter Fox. Unsupervised feature learning for 3d scene labeling. In *IEEE International Conference on Robotics and Automation*, pages 3050–3057, 2014. 2
- [22] Kevin Lai and Dieter Fox. Object Recognition in 3D Point Clouds Using Web Data and Domain Adaptation. In *The International Journal of Robotics Research*, volume 29, pages 1019–1037, 2010. 2
- [23] Yangyan Li, Angela Dai, Leonidas Guibas, and Matthias Niessner. Database-Assisted Object Retrieval for Real-Time 3D Reconstruction. In *Computer Graphics Forum*, volume 34, pages 435–446, 2015. 2
- [24] D Lin, S Fidler, and R Urtasun. holistic scene understanding for 3D object detection with RGBD cameras. In *International Conference on Computer Vision*, 2013. 2, 6, 7
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 3
- [26] Oliver Mattausch, Daniele Panozzo, Claudio Mura, Olga Sorkine-Hornung, and Renato Pajarola. Object detection

and classification from large-scale cluttered indoor scans. In *Computer Graphics Forum*, volume 33, pages 11–21, 2014. 2

- [27] Oliver Mattausch, Daniele Panozzo, Claudio Mura, Olga Sorkine-Hornung, and Renato Pajarola. Object Detection and Classification from Large-Scale Cluttered Indoor Scans. In *Computer Graphics Forum*, 2014. 2
- [28] Aron Monszpart, Nicolas Mellado, Gabriel J. Brostow, and Niloy J. Mitra. RAPter: Rebuilding Man-made Scenes with Regular Arrangements of Planes. In ACM Transactions on Graphics, volume 34, pages 103:1–103:12, 2015. 2
- [29] Liangliang Nan, Ke Xie, and Andrei Sharf. A Searchclassify Approach for Cluttered Indoor Scene Understanding. In ACM Transactions on Graphics, volume 31, pages 137:1–137:10, New York, NY, USA, 2012. 2
- [30] X Ren, L Bo, and Dieter Fox. RGB-(D) scene labeling: Features and Algorithms. In *IEEE Conference on Computer Vi*sion and Pattern Recognition, 2012. 2
- [31] Jason Rock, Tanmay Gupta, Justin Thorsen, JunYoung Gwak, Daeyun Shin, and Derek Hoiem. Completing 3D object shape from one depth image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2, 4
- [32] Tianjia Shao, Weiwei Xu, Kun Zhou, Jingdong Wang, Dongping Li, and Baining Guo. An Interactive Approach to Semantic Modeling of Indoor Scenes with an RGBD Camera. In ACM Transactions on Graphics, 2012. 2
- [33] Yifei Shi, Pinxin Long, Kai Xu, Hui Huang, and Yueshan Xiong. Data-driven Contextual Modeling for 3D Scene Understanding. In *Computers & Graphics*, volume 55, pages 55–67, 2016. 2
- [34] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor Segmentation and Support Inference from RGBD Images. In *European Conference on Computer Vi*sion, 2012. 2
- [35] Shuran Song and Jianxiong Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
 2
- [36] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic Scene Completion from a Single Depth Image. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2, 3, 4, 5, 6, 7, 8
- [37] Thanh Nguyen, Duc and Hua, Binh-Son and Tran, Khoi and Pham, Quang-Hieu and Yeung, Sai-Kit. A field model for repairing 3D shapes. In *IEEE Conference on Computer Vision* and Pattern Recognition, 2016. 2
- [38] Jacob Varley, Chad DeChant, Adam Richardson, Joaquín Ruales, and Peter Allen. Shape completion enabled robotic grasping. In *IEEE Conference on Intelligent Robots and Systems*, 2017. 2
- [39] Weiyue Wang, Qiangui Huang, Suya You, Chao Yang, and Ulrich Neumann. Shape Inpainting using 3D Generative Adversarial Network and Recurrent Convolutional Networks. In *CoRR*, abs/1711.06375, 2017. 2
- [40] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D

ShapeNets: A deep representation for volumetric shapes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2

- [41] Bo Yang, Hongkai Wen, Sen Wang, Ronald Clark, Andrew Markham, and Niki Trigoni. 3D Object Reconstruction from a Single Depth View with Adversarial Learning. In *CoRR*, *abs/1708.07969*, 2017. 2
- [42] Jiahui Zhang, Hao Zhao, Anbang Yao, Yurong Chen, Li Zhang, and Hongen Liao. Efficient Semantic Scene Completion Network with Spatial Group Convolution. In *European Conference on Computer Vision*, 2018. 2, 6, 7
- [43] Yinda Zhang, Shuran Song, Ersin Yumer, Manolis Savva, Joon-Young Lee, Hailin Jin, and Thomas Funkhouser. Physically-Based Rendering for Indoor Scene Understanding Using Convolutional Neural Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 2
- [44] B Zheng, Y Zhao, J C Yu, K Ikeuchi, and Zhu Sx-Cx. Beyond point clouds: Scene understanding by reasoning geometry and physics. In *IEEE Conference on Computer Vision* and Pattern Recognition, 2016. 2, 6, 7