EdgeNet: Semantic Scene Completion from a Single RGB-D Image

Aloisio Dourado, Teofilo Emidio de Campos
University of Brasilia
Brasilia, Brazil
aloisio.dourado.bh@gmail.com, t.decampos@st-annes.oxon.org

Hansung Kim, Adrian Hilton
University of Surrey
Surrey, UK
(h.kim, a.hilton)@surrey.ac.uk

Abstract—Semantic scene completion is the task of predicting a complete 3D representation of volumetric occupancy with corresponding semantic labels for a scene from a single point of view. In this paper, we present EdgeNet, a new end-to-end neural network architecture that fuses information from depth and RGB, explicitly representing RGB edges in 3D space. Previous works on this task used either depth-only or depth with colour by projecting 2D semantic labels generated by a 2D segmentation network into the 3D volume, requiring a two step training process. Our EdgeNet representation encodes colour information in 3D space using edge detection and flipped truncated signed distance, which improves semantic completion scores especially in hard to detect classes. We achieved state-of-the-art scores on both synthetic and real datasets with a simpler and a more computationally efficient training pipeline than competing approaches.

I. INTRODUCTION

The ability of reasoning about scenes in 3D is a natural task for humans, but remains a challenging problem in Computer Vision [1]. Knowing the complete 3D geometry of a scene and the semantic labels of each 3D voxel has many practical applications, like robotics and autonomous navigation in indoor environments, surveillance, assistive computing and augmented reality.

Currently available low cost RGB-D sensors generate data form a single viewing position and cannot handle occlusion among objects in the scene. For instance, in the scene depicted on the left part of Figure 1, parts of the wall, floor and furniture are occluded by the bed. There is also self-occlusion: the interior of the bed, its sides and its rear surfaces are hidden by the visible surface.

Given a partial 3D scene model acquired from a single RGB-D image, the goal of scene completion is to generate a complete 3D volumetric representation where each voxel is labelled as occupied by some object or free space. For occupied voxels, the goal of *semantic* scene completion is to assign a label that indicates to which class of object it belongs, as illustrated on the right part of Figure 1.

Before 2018, most of the work on scene reasoning only partially addressees this problem. A number of approaches only infer labels of the visible surfaces [2], [3], [4], while others only consider completing the occluded part of the scene, without semantic labelling [5]. Another line of work focuses on single objects, without the scene context [6].

The term semantic scene completion was introduced by Song *et al.* [7], who showed that scene completion and semantic labelling are intertwined and training a CNN to jointly deals with both tasks can lead to better results. Their approach only uses depth information, ignoring all information from RGB channels. Colour information is expected to be useful to distinguish objects that approximately share the same plane in the 3D space, and thus, are hard to be distinguished using only depth. Examples of such instances are flat objects attached to the wall, such as posters, paintings and flat TVs. Some types of closed doors and windows are also problematic for depth-only approaches.

Recent research also explored colour information from on RGB-D images to improve semantic scene completion scores. Some methods project colour information to 3D in a naive way, leading to a problem of data sparsity in the voxelised data that is fed to the 3D CNN [8], while others uses RGB information to train a 2D segmentation network and then project generated features to 3D, requiring a complex two step training process [9], [10].

Our work focuses on enhancing semantic scene segmentation scores using information from both depth and colour of RGB-D images in an end-to-end manner. In order to address the RGB data sparsity issue, we introduce a new strategy for encoding information extracted from RGB image in 3D space. We also present a new end-to-end 3D CNN architecture to combine and represent the features from colour and depth. Comprehensive experiments are conducted to evaluate the main aspects of the proposed solution. Results show that our fusion approach can enhance results of depth-only solutions and that EdgeNet achieves equivalent performance to current state-of-the-art fusion approach, with a much simpler training protocol.

To summarise, our main contributions are:

- EdgeNet, a new end-to-end CNN architecture that fuses depth, RGB edge information to achieve state-of-the-art performance in semantic scene completion with a much simpler approach;
- a new 3D volumetric edge representation using flipped signed-distance functions which improves performance and unifies data agregation for semantic scene completion from RGBD;
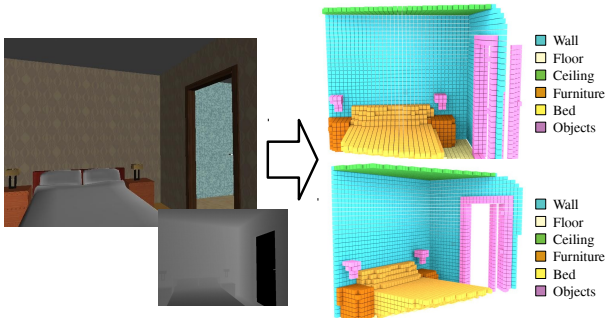
Fig. 1: Semantic scene completion. Given an RGB-D image, the goal is to infer a complete 3D occupancy grid with associated semantic labels.

- a more efficient end-to-end training pipeline for semantic scene completion with relation to previous approaches.

II. RELATED WORK

Scene semantic completion in 3D is a problem that was introduced relatively recently. Previous approaches to 3D SSC rely on Fully Convolutional Neural Network architectures (FCNs, introduced in [11]) and use SUNCG and NYUDv2 as training sources (these datasets are described in Section IV-A). We classify approaches into three main groups, based on the type of input of the semantic completion CNN: depth maps only, depth maps plus RGB and depth maps plus 2D segmentation maps.

A. Depth maps only

Song *et al*. [7] used depth maps from the SUNCG synthetic dataset to train a typical contracting fully convolutional CNN with 3D dilated convolutions, called SSCNet. They showed that jointly training for segmentation and completion leads to better results, as both tasks are inherently intertwined. To deal with data sparsity after projecting depth maps from 2D to 3D, the authors used a variation of Truncated Signed Distance Function (TSDF) that they called Flipped TSDF (F-TSDF). Zhang *et al*. [12] used dense conditional random field to enhance SSCNet results. Guo and Tong [13] applied a sequence of 2D convolutions to the depth maps, used a projection layer to projected the features to 3D and feed the output to a 3D CNN.

All solutions in this category are end-to-end approaches, in other words, the network is trained as a whole, with no need for extra training stages for specific parts. EdgeNet is an end-to-end network as well. RGB edges are aggregated in the same training pipeline of the depth information.

B. Depth maps plus RGB

Guedes *et al*. [8] reported preliminary results obtained by adding colour to an SSCNet-like architecture. In addition to the F-TSDF encoded depth volume, they used three extra projected volumes, corresponding to the channels of the RGB image, with no encoding, resulting in 3 sparse volumetric representation of the partially observed surfaces. The authors reported no significant improvement using the colour information in this sparse manner.

C. Depth maps plus 2D segmentation

Models in this category use a two step training protocol, where a 2D segmentation CNN is first trained and then it is used to generate input to a 3D semantic scene completion CNN. Current models differ in the way the generated 2D information is fed into the 3D CNN.

Garbade *et al*. [9] used a pre-trained 2D segmentation CNN with a fully connected CRF [14] to generate a segmentation map, which, after post-processing, was projected to 3D. Liu *et al*. [10] used depth maps and RGB information as input to an encoder-decoder 2D segmentation CNN. The encoder branch of the 2D CNN is a ResNet-101 [15] and the decoder branch contains a series of dense upsampling convolutions. The generated features from the 2D CNN are then reprojected to 3D using camera parameters, before being fed into a 3D CNN. The authors showed results using 2 different strategies to fuse depth and RGB: SNetFusion performs fusion just after the 2D segmentation network, while TNetFusion only performs fusion after the 3D convolutional network. TNetFusion achieves higher performance, with a much higher computational cost. The 2D CNN is also pre-trained offline.

Using 2D segmentation maps on 3D SSC brings an additional complexity to the training phase which is trains and evaluates the 2D segmentation network prior to the 3D CNN training. In this work, we propose an end-to-end approach to fuse information from depth and colour, where the network can be trained and evaluated as a whole, and still achieves state-of-the-art performance.

III. OUR SOLUTION: EDGENET

Our proposed solution is the first end-to-end approach that successfully uses information from RGB to improve semantic scene completion performance over depth only. It consists in a novel approach to encode information from RGB edges and depth maps and a new 3D CNN architecture to fuse both modalities that we call EdgeNet.

A. Encoding edges in 3D

As discussed earlier, colour information should complement depth maps for 3D semantic scene completion. However, combination of these modalities in a meaningful representation for learning is not trivial. Guedes *et al*. [8] naively added 3 channels to each voxel to insert R, G and B colour information into the representation, with no encoding. In this way, the vast majority of voxels have no colour data while only those at the visible surface have a colour value. This explains why they do not improve on the previous approach using depth only. Song *et al*. [7] demonstrate that F-TSDF encoding plays an important role in feeding a projected depth map to a 3D CNN and produces better results than TSDF and other encoding techniques.
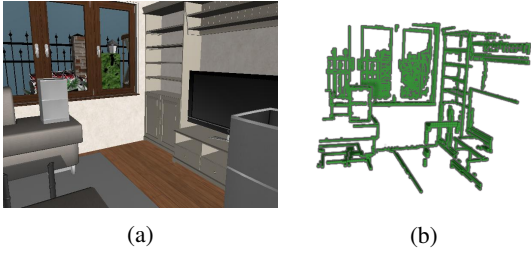
Fig. 2: Projection of Edges to 3D: (a) original RGB image, (b) voxelized edges after projection.
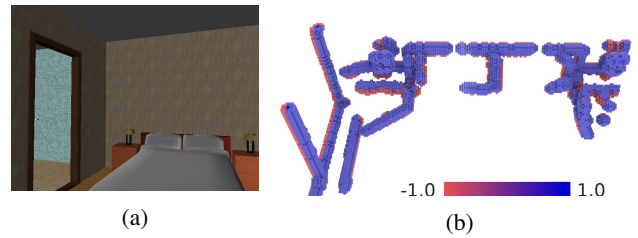

Fig. 3: (a) original scene. (b) F-TSDF of edges in 3D. The edge image is a horizontal cut of the scene, taken just above the bed. Only F-TSDF values with absolute value greater than 0.8 are shown (best viewed in colour).

Given a sparse 3D voxel volume, the Truncated Signed Distance Function (TSDF) consists in computing the Euclidean distance of each empty voxel to the nearest occupied voxel. The signal of occluded regions is set to be negative, while visible regions are given positive values. Near the occupied surface, TSDF produces a value that tends to zero on both sides (and its first derivative tends to zero as well). TSDF values are normalised to [-1,1]. Flipped TSDF (F-TSDF) follows the same principle, but the absolute values of both visible and occluded regions are flipped:

$$\text{F-TSDF} = \text{sign}(\text{TSDF}) \cdot (1 - |\,\text{TSDF}\,|). \qquad (1)$$

A discontinuity near the occupied surface (from -1 to 1) occurs and the first derivative tends to infinity.

F-TSDF encoding of volumetric data can be easily applied to depth maps after 3D projection because each voxel carries binary information: occupied or free. On the other hand, F-TSDF can not straightforwardly be applied to RGB or semantic segmentation maps, because they are not binary.

To deal with this problem, we introduce a new strategy to fuse colour appearance and depth information for 3D semantic scene completion. Our approach exploits edge detection in the image, which gives a 2D binary representation of the scene that can highlight objects that are hard to detect in depth maps. For instance, a poster on a wall is expected to be invisible in a depth map, especially after down-sampling. On the other hand, RGB edges highlight the presence of that object.

The main advantage of extracting edges and projecting them to 3D is the possibility to apply F-TSDF on both edges and surface volumes, as they are both binary, providing two meaningful input signals to the 3D CNNs. Another advantage is that due to their simplicity, edges are more transferable, removing the need for the application of a domain adaptation method when learning from synthetic images and applying on real images.

We apply F-TSDF to 3D edges, similarly to F-TSDF applied to 3D surfaces: for each voxel in the edge volume, we look for the nearest edge to calculate the Euclidean distance. Visible and occluded voxels are related only to edges, not to surfaces. We use the standard Canny edge detector [16] to perform edge detection. Each edge location is projected to a point in the 3D space using its depth information and the camera calibration matrix. The resulting point cloud is voxelised in the same way as the depth point cloud, resulting in a sparse volume

of $240 \times 144 \times 240$ voxels. Figure 2 shows a scene from the SUNCG dataset and its corresponding edges projected to 3D. Figure 3b shows in detail a region of the projected edges of 3a after F-TSDF encoding. Note that the greatest gradients occur along the edges.

B. EdgeNet architecture

In order to combine depth and edge modalities, we propose a new 3D semantic segmentation CNN architecture that we call **EdgeNet**. Our proposed solution is a 3D CNN inspired by the U-Net design [17] which has been successfully used in many 2D semantic segmentation problems, and is presented in Figure 4. We address the degradation problem of deeper networks [18], by replacing simple convolutional blocks of U-Net by ResNet modules [15]. On lower resolutions, the ResNet modules uses dilated convolutions to improve the receptive field. To match the resolution of the output, the input branch reduces the resolution to 1/4 of the input. Next blocks follows encoder-decoder design and the last stage of the decoding branch is responsible for reducing the number of channels to match the desired number of output classes and loss calculations.

Depth and Edges Fusion Schemes. The encoder-decoder structure of EdgeNet allows us to evaluate three fusion schemes: Early Fusion (EdgeNet-EF), Middle Fusion (EdgeNet-MF) and Late Fusion (EdgeNet-LF). In EdgeNet-EF, just after F-TSDF encoding, both input volumes are concatenated and fed into the main network. In EdgeNet-MF, the input branch is divided into two parts while in EdgeNet-LF, both input and encoding branches are divided. To keep the same memory requirement in all fusion schemes, the total quantity of channels in all scheme is always the same.

Data balancing and loss function. In volumetric data, occluded and occupied voxels are highly unbalanced, so we use a weighted version of categorical cross entropy as the loss function to train our models. To obtain the weights, for each training batch, we randomly initialize a tensor $rand_{occl}$ of the same shape as the batch with ones and zeroes using the ratio $r = (2 \sum occu / \sum occl)$, where $occl$ and $occu$ are two tensors obtained from the previously calculated occupancy grid relative to occluded and occupied voxels. The final weight tensor is $w = occu + occl \odot rand_{occl}$, where $\odot$ denotes the
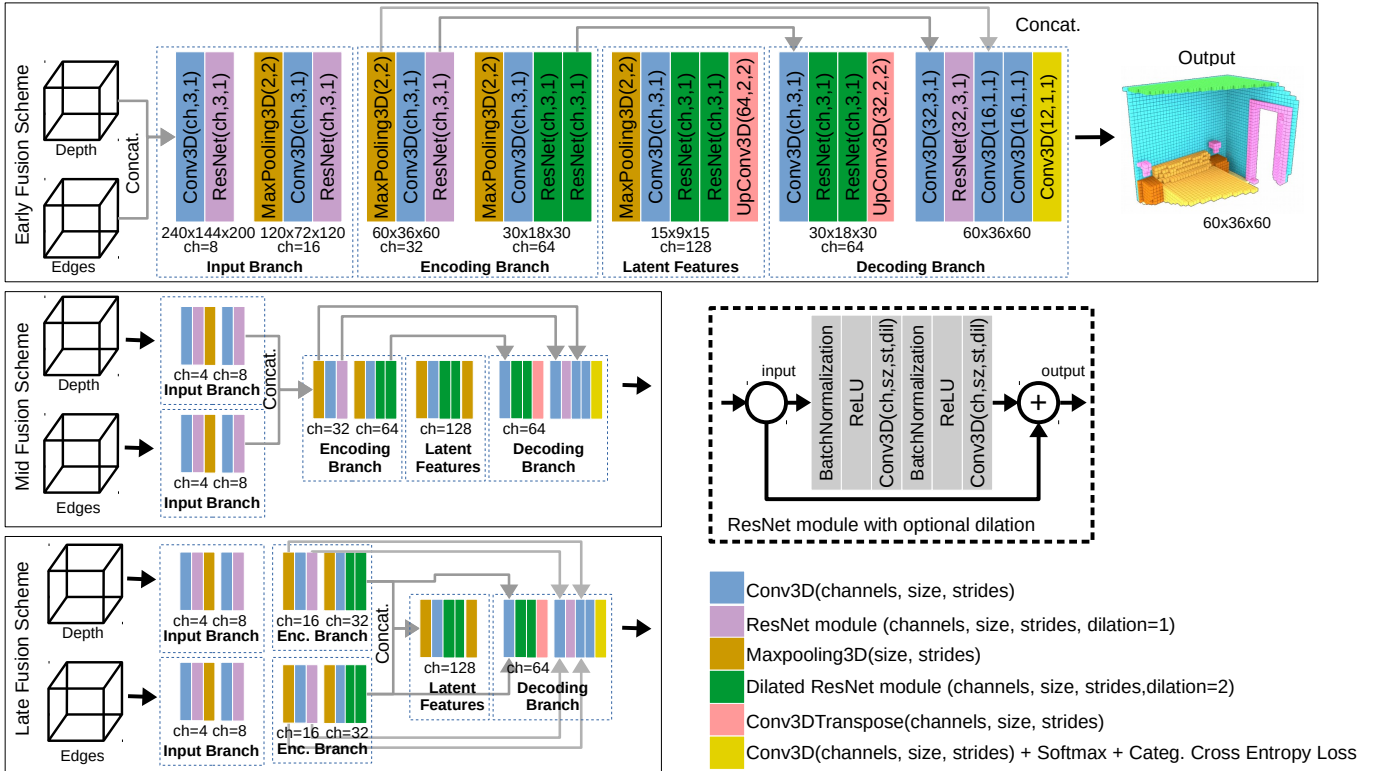
Fig. 4: EdgeNet architecture and fusion schemes (best viewed in colour).

Hadamard product. Let p be the predicted probabilities of the 12 classes for each voxel and y be the one hot encoded ground truth tensor. The categorical cross entropy loss function is then given by

$$L_{cce}(p, y) = - \sum (w \odot y \odot \log p). \qquad (2)$$

C. Training pipeline with offline data preparation

As F-TSDF calculation is computationally intensive, to reduce overall training time, the F-TSDF volumes that feed the models are preprocessed off-line once. The preprocessed dataset is then stored, and may be used as many times as needed, including by different models. Following previous works, we rotate the 3D Scene to align it with gravity and have room orientation based on the Manhattan assumption. We fixed the dimensions of the 3D space to 4.8 m horizontally, 2.88 m vertically and 4.8 m in depth. Voxel grid size is 0.02 m, resulting in a $240 \times 144 \times 240$ 3D volume. The TSDF truncation value is 0.24 m. Surface and edge projection as well as F-TSDF encoding of all volumes are done in this stage. During preprocessing, we also calculate an occupancy grid where we distinguish occupied voxels inside the room and FOV; non-occupied occluded voxels inside the room and FOV; and all other voxels. This occupancy grid will be further used to balance the dataset during training time.

IV. EXPERIMENTS

In this section we describe the datasets and the evaluation protocol used in this paper.

A. Datasets

We train and validate our proposed approach on SUNCG [7] and NYUDv2 [19] datasets. The SUNCG dataset consists of about 45K synthetic scenes from which more than 130K 3D scenes were rendered with corresponding depth maps and ground truth, divided in train and test datasets. As the original training and test sets did not include RGB images, we extracted the camera poses from the provided ground truth and rendered a new set of depth and RGB images from the SUNCG synthetic scenes. To avoid misalignments, the ground truth volumes were regenerated from the scene meshes.

NYUDv2 is a widely used dataset of indoor scenes that includes depth and RGB images captured by the Kinect depth sensor, divided in 795 samples for training and 654 for test. Following the majority of works in semantic segmentation we used ground truth obtained by voxelizing the 3D mesh annotations from [20] and and mapped object categories based on [21].

B. Training protocols

Our experiments consists in training our models from scratch on SUNCG and NYUDv2, and also fine-tuning models trained from SUNCG to NYUDv2. For experiments in which we trained our models from scratch, we use the technique known as One Cycle Learning [22], which is a combination of Curriculum Learning [23] and Simulated Annealing [24]. After some preliminary tests, we found 0.01 to be a good base learning rate. We use a maximum of 30 epochs, in order to

maintain total training time in a acceptable limit. Following [22], we start with the base learning rate and linearly increase the effective learning until 0.1 in the 10th epoch, then linearly decrease the learning rate until reach the start-up level in the 20th epoch. During the annealing phase, we linearly go from 0.01 to 0.0005 in a further 10 epochs. Due to GPU memory size constraints, we use a batches of 3 samples. We also use SGD optimizer with a momentum of 0.9 and decay of 0.0005 in all experiments, as used in most previous works. For SUNCG, each epoch consists of 30,540 scenes randomly selected from the whole training set. For NYUDv2, each epoch comprises the whole training set. For fine tuning, we initialize the network with parameters trained on SUNCG and use the standard training policy with SGD with fixed learning rate of 0.01 and 0.0005 of weight decay.

Thanks to our lightweight training pipeline with offline F-TSDF preprocessing, our training time is only 4 days on SUNCG and 6 hours on NYUDv2, using a GTX 1080 TI.

C. Evaluation

For the semantic scene completion task, we report the Intersection over Union (IoU) of each object class on both the observed and occluded voxels. For the scene completion task, all non-empty object classes are considered as one category, and we report Precision, Recall and IoU of the binary predictions on occluded voxels. Voxels outside the view or the room are not considered.

D. Experimental results

We compare our results to semantic scene completion approaches that use depth-only [13], [7], [12], depth plus RGB [8] and depth plus 2D segmentation maps [9], [10]. We also investigate the effects of the main aspects of our proposed solution on SUNCG. Comparative results were extracted from the original papers.

1) Ablation Studies and results on SUNCG: In Table I, investigate the effects of the main aspects of our proposed solution. At first, we analyse the effect of our training pipeline. We took SSCNet as a baseline and retrain it, using our lightweight training framework, that allows a batch size of 3 samples in comparison to the 1 sample batch size of original SSCNet. Results of that experiment are shown as SSNet*. We observed a large improvement on SSC scores just using our pipeline.

After isolating the effect of our training protocol, we investigate the effect of our encoder-decoder architecture, with dilated ResNet modules. To accomplish this, we used EdgeNet-D, that is the Ednet architecture fed only with depth, without edges. Once again we observed a high level of improvement, comparing to SSCNet*. EdgeNet-D also got the best overall scores amongst the depth-only approaches. Next experiment evaluates the effect of adding edges to an existing depth-only architecture. We took SSCNet and fed it with both depth and edges after F-TSDF encoding (SSCNet-E). We observed improvements compared to SSCNet* on overall scores and especially on hard-to-detect classes like TVs and objects.

Finally, we evaluate the benefits of adding Edges to our architeture in three fusion schemes: EdgeNet-EF, EdgeNet-MF and EdgeNet-LF. Performance gains from EdgeNet-D show, once again, that adding edges is useful. A discussion about fusion schemes is provided on Section V.

We also compare EdgeNet results to previous approaches. Overall, our proposed solutions achieve the best performance by a large margin. EdgeNet-EF achieves best average scores, while EdgeNet-MF achieves the best score in some classes. EdgeNet-EF surpassed VVNetR120, the best previous approach on average SSC, by 3.3%. As expected, the highest improvements are observed on hard to detect classes, like objects and TVs. Although SUNCG is synthetic, evaluation on this dataset is quite important because of the poor quality of the ground truth in NYU, which impacts negatively accurate models like EdgeNet.

E. Results on NYUDv2

Table II shows the results of EdgeNet on NYUDv2 dataset and compares it with previous approaches. We compare results for models trained only on synthetic data, only on NYUDv2 and on both synthetic and NYUDv2 using fine tuning.

On SUNCG-only and on NYUDv2-only training scenarios, EdgeNet-MF achieved the best overall scores on Scene Completion and Semantic Scene Completion. On SUNCG+NYU training scenario, however, TNetFuse presented the best result. EdgeNet-MF achieved best scores on structural elements and chair. It is worth mentioning that the NYUDv2 dataset has severe ground truth errors and misalignment, so results are not precise, and small differences in results may be questioned (see Section IV-F).

Despite these problems on NYU ground truth, EdgeNet achieves state-of-the-art level results with a much simpler and more computationally efficient training pipeline. EdgeNet is an end-to-end approach, and its memory consumption allows a batch size of 3 samples in a GTX 1080TI GPU, while TNetFuse requires a complex two step training procedure and uses a batch size of only 1 sample, in the same GPU.

F. Qualitative Results

Qualitative results on NYUDv2 are shown in Figure 5. Models used to generate the inferences were trained on SUNCG and fine tuned on NYUDv2. We compare results of SSCNet* to our three models. It is visually perceptible that EdgeNet presents more accurate results.

In the first row of images of Figure 5, note the presence of a picture and a window, and observe that the ground truth misses the window. SCCNet* did not detect the picture and the window while EdgeNet-MF detects the window and some parts of the picture. This ground truth mislabelling affects negatively the performance of EdgeNet.

The second row of Figure 5 also depicts some problems related to Ground Truth annotations on NYUDv2 dataset. Note that neither the papers fixed on the wall nor the shelf appear in the Ground Truth. All models captured the shelf, but only EdgetNet inferred the presence of objects fixed on the wall.

input	model	scene completion			semantic scene completion (IoU, in percentages)											
		prec.	rec.	IoU	ceil.	floor	wall	win.	chair	bed	sofa	table	tvs	furn.	objs.	avg.
d	SSCNet[7]	76.3	**95.2**	73.5	96.3	84.9	56.8	28.2	21.3	56.0	52.7	33.7	10.9	44.3	25.4	46.4
	SSCNet*	92.7	89.7	83.8	97.0	94.6	74.3	51.1	43.7	78.2	70.9	49.5	45.2	61.0	51.3	65.2
	DCRF [12]				95.4	84.3	57.7	24.5	28.2	63.4	55.3	34.5	19.6	45.8	28.7	48.8
	VVNetR-120 [13]	90.8	91.7	84.0	**98.4**	87.0	61.0	54.8	49.3	83.0	**75.5**	55.1	43.5	68.8	57.7	66.7
	EdgeNet-D	93.1	90.4	84.8	97.2	94.4	78.4	56.1	50.4	80.5	73.8	54.5	49.8	69.5	59.2	69.5
d+s	SNetFuse[10]	56.7	91.7	53.9	65.5	60.7	50.3	56.4	26.1	47.3	43.7	30.6	37.2	44.9	30.0	44.8
	TNetFuse[10]	53.9	**95.2**	52.6	60.6	57.3	53.2	52.7	27.4	46.8	53.3	28.6	41.1	44.1	29.0	44.9
d+e	SSCNet-E	92.8	89.6	83.8	97.0	94.5	74.6	51.8	43.9	77.0	70.8	49.3	49.2	62.1	52.0	65.7
	EdgeNet-EF(Ours)	**93.7**	90.3	**85.1**	97.2	94.9	**78.6**	57.4	49.5	80.5	74.4	**55.8**	51.9	70.1	**62.5**	**70.3**
	EdgeNet-MF(Ours)	93.3	90.6	**85.1**	97.2	**95.3**	78.2	**57.5**	**51.4**	**80.7**	74.1	54.5	**52.6**	**70.3**	60.1	70.2
	EdgeNet-LF(Ours)	93.0	89.6	83.9	97.0	94.6	76.4	52.0	44.6	79.8	71.5	48.9	48.3	66.1	55.9	66.8

TABLE I: **Results and ablation studies on SUNCG test set**. We took SSCNet as a baseline and show the effect of each one of the main aspects of our proposed approach. Column 'input' indicates the type of input: d = depth only; d+e = depth + edges. SSCNet* is our implementation of the original SSCNet, with our training pipeline. EdgeNet-D has the same architecture of the other versions of EdgeNet, but the edge volume is not fed into the network. EdgeNet-EF achieves the best overall scores and surpassed VVNetR-120 by 3.3% on average IoU for semantic scene completion.

train	input	model	scene completion			semantic scene completion (IoU, in percentages)											
			prec.	rec.	IoU	ceil.	floor	wall	win.	chair	bed	sofa	table	tvs	furn.	objs.	avg.
SUNCG	d	SSCNet[7]	55.6	91.9	53.2	5.8	81.8	19.6	5.4	12.9	34.4	26	13.6	6.1	9.4	7.4	20.2
	d+e	EdgeNet-EF(Ours)	**61.9**	80.0	**53.6**	9.1	**92.9**	18.3	5.7	15.8	40.4	30.7	9.2	3.3	13.7	11.6	22.8
		EdgeNet-MF(Ours)	60.7	80.3	52.8	**11.0**	92.3	**20.5**	7.2	**16.3**	42.8	**32.8**	**10.5**	6.0	**15.7**	**11.8**	**24.3**
		EdgeNet-LF(Ours)	59.9	**80.5**	52.3	3.2	87.1	19.9	**8.6**	15.4	**43.5**	32.3	8.8	4.3	13.7	10.0	22.4
NYU	d	SSCNet[7]	57.0	**94.5**	55.1	15.1	94.7	24.4	0.0	**12.6**	32.1	35.0	**13.0**	7.8	27.1	10.1	24.7
	d+e	EdgeNet-EF(Ours)	**78.1**	65.1	55.1	**21.8**	**95.0**	27.3	**8.4**	6.8	**53.1**	38.6	7.5	0.0	30.4	**13.3**	27.5
		EdgeNet-MF(Ours)	76.0	68.3	**56.1**	17.9	94.0	**27.8**	2.1	9.5	51.8	**44.3**	9.4	3.6	**32.5**	12.7	**27.8**
		EdgeNet-LF(Ours)	75.5	67.5	55.4	19.8	94.9	24.4	5.7	7.2	50.3	38.8	10.0	0.0	33.2	12.2	27.0
SUNCG + NYU	d	SSCNet[7]	59.3	92.9	56.6	15.1	94.6	24.7	10.8	17.3	53.2	45.9	15.9	13.9	31.1	12.6	30.5
		DCRF[12]	-	-	-	18.1	92.6	27.1	10.8	18.8	54.3	47.9	17.1	15.1	34.7	13.0	31.8
		VVNetR-120[13]	69.8	83.1	61.1	19.3	94.8	28.0	12.2	19.6	57.0	50.5	17.6	11.9	35.6	15.3	32.9
	d+c	Guedes *et al.* [8]	-	-	56.6	-	-	-	-	-	-	-	-	-	-	-	30.5
	d+s	Garbade *et al.* *[9]	69.5	82.7	**60.7**	12.9	92.5	25.3	20.1	16.1	56.3	43.4	17.2	10.4	33.0	14.3	31.0
		SNetFuse[10]	67.6	**85.9**	**60.7**	22.2	91.0	28.6	**18.2**	19.2	56.2	51.2	16.2	12.2	37.0	17.4	33.6
		TNetFuse[10]	67.3	85.8	**60.7**	17.3	92.1	28.0	16.6	19.3	**57.5**	**53.8**	**17.7**	**18.5**	**38.4**	**18.9**	**34.4**
	d+e	EdgeNet-EF(Ours)	77.0	70.0	57.9	16.3	**95.0**	27.9	14.2	17.9	55.4	50.8	16.5	6.8	37.3	15.3	32.1
		EdgeNet-MF(Ours)	**79.1**	66.6	56.7	**22.4**	**95.0**	**29.7**	15.5	**20.9**	54.1	53.0	15.6	14.9	35.0	14.8	33.7
		EdgeNet-LF(Ours)	77.6	69.5	57.9	20.6	94.9	29.5	9.8	18.1	56.2	50.5	11.4	5.2	35.9	15.3	31.6

TABLE II: **Semantic scene completion results on NYUDv2 test set**. Column input indicates the type of input: d=depth only; d+s=depth and segmentation maps; d+e=depth and edges. Column train indicates dataset used for training the models. SUNCG + NYU means trained on SUNCG and fine tuned on NYUDv2.

When quantitative results are computed, these ground truth annotation flaws unfairly benefit the less precise models and harm more precise models like ours.

V. DISCUSSION

In this section we discuss key aspects and contributions of our proposed approach.

Has the new training pipeline any influence over results? We compared the results originally achieved by SSCNet to results of the version of it trained with our pipeline (SSCNet*). On SUNCG we observed an improvement of almost 20% on semantic scene completion and more than 10% on scene completion. Besides the improvements on model performance, the more computationally efficient pipeline also contributed to reduce training time from 7 days to 4 days when training on SUNCG and from 30 hours to 6 hours when training on NYUDv2, with a batch of size 3, whereas the original framework only allowed a batch size of 1 sample on a GTX 1080Ti (which has 11GB of memory). Besides reducing training time, larger batch sizes enhance training stability, acting as a regularizer [25].

Is a deeper U-shaped CNN with dilated ResNet modules helpful? We investigated the effects of our architecture with and without aggregating edges. On both scenarios, our proposed architecture outperformed the shallower network, confirming that our network architecture is helpful.

Is aggregating edges helpful? May Other 3D CNN architectures benefit from aggregating edges? We compared the original SSCNet architecture trained with our pipeline to a modified version of it that aggregates edges encoded with F-TSDF (SSCNet-E). SSCNet-E presented better results on SUNCG, demonstrating that the aggregation of edge information is helpful. We also observed improvements using a deeper depth-only network (EdgeNet-D). This experiments demonstrates that the proposed 3D volumetric representation
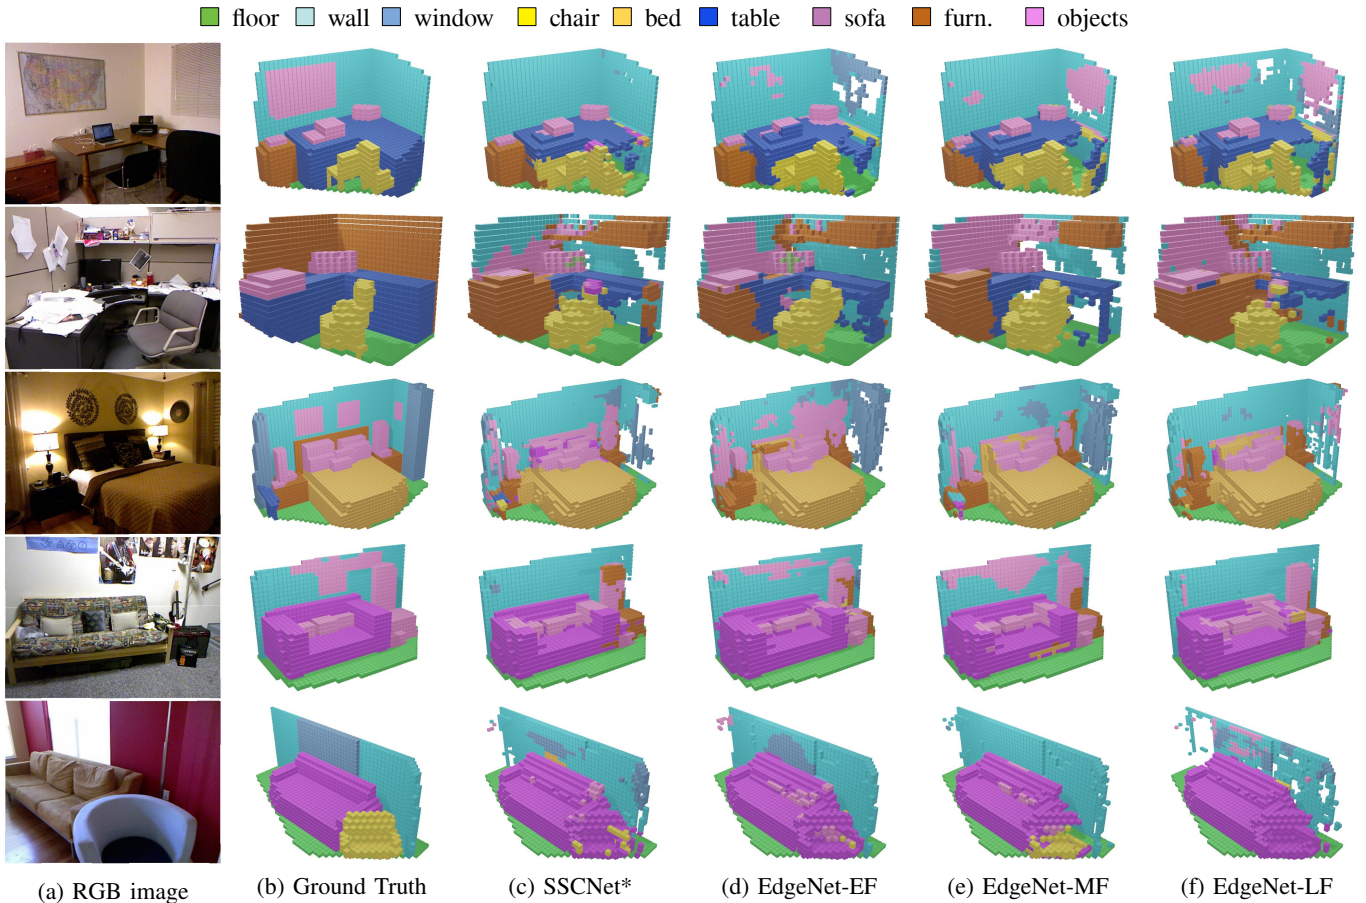
| (a) RGB image | (b) Ground Truth | (c) SSCNet* | (d) EdgeNet-EF | (e) EdgeNet-MF | (f) EdgeNet-LF |

Fig. 5: **Qualitative Results**. We compare EdgeNet results using SSCNet* as a baseline on SUNCG and NYUDv2. Overall, EdgeNet gives more accurate voxel predictions, especially for hard to detect clases (best viewed in colour).

of color edges can improve the performance of other previous depth only approaches.

What is the best fusion strategy? The later the fusion, the higher is the memory requirement, due to the duplication of convolutional branches. Higher memory may imply in smaller batch sizes which may negatively impact learning. Liu *et al.* [10] observed better results using late fusion, but they faced the problem of higher memory consumption. Our choice was to fix the memory footprint, reducing the number of channels of duplicated branches without compromising the training time and stability. However, very late fusion schemes may suffer from accuracy degradation due to reduced number of parameters in deeper layers. Taking those aspects into account, we found that a mid-level fusion strategy works and generalizes better for EdgeNet considering both synthetic and real datasets.

How does EdgeNet compare to other RGB + depth approaches? We have compared EdgeNet with other RGB + depth approachs on SUNCG (Table I) and NYUDv2 (Table II. On SUNCG, EdgeNet versions surpassed previous approaches by a large margin. On NYU, EdgeNet got similar results as the solutions from TNetFuse [10], with less than 1% difference. It is important to observe NYU ground truth annotations are

not precise, which impacts negatively more accurate models. Another aspect that is worth mentioning is that TNetFuse needs a complex and less computationally efficient two-step training protocol, while EdgeNet and the previous depth-only solutions cited in this paper are end-to-end networks, with a much simpler and efficient training pipeline.

VI. CONCLUSION

This paper presented a new approach to fuse depth and colour into a CNN for semantic scene completion. We introduced the use of F-TSDF encoded 3D projected edges extracted from RGB images. We also presented a new end-to-end network architecture capable of properly aggregating edges and depth, extracting useful information from both sources, without requiring previous 2D semantic segmentation training as is the case of previous approaches that combine depth and colour. Experiments with alternate models, showed that both aggregating edges and the new proposed architecture have positive impact on semantic scene completion, especially for hard to detect objects. Qualitative results show significant improvement for objects such as pictures, which cannot be differentiated by depth only. On SUNCG, we have achieved the best overall result, and on NYU, we have achieved the

state-of-the-art results of other approaches that use a more complex training protocol.

Experiments showed that our proposed approach of aggregating Edges may be applied to other existing solutions, opening room for further improvements.

We also developed a lightweight training pipeline for the task, which reduced the memory footprint in comparison to other solutions and reduced the training time on SUNCG from 7 to 4 days and on NYUDv2 from 30 to 6 hours, that we intend to make public upon acceptance of this paper.

REFERENCES

[1] D. Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. MIT Press, 1982.

[2] S. Gupta, P. Arbelez, and J. Malik, "Perceptual organization and recognition of indoor scenes from RGB-D images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2013, pp. 564–571.

[3] X. Qi, R. Liao, J. Jia, S. Fidler, and R. Urtasun, "3D graph neural networks for RGBD semantic segmentation," in *IEEE International Conference on Computer Vision (ICCV)*, 10 2017, pp. 5209–5218.

[4] X. Ren, L. Bo, and D. Fox, "Rgb-(d) scene labeling: Features and algorithms," in *IEEE Conference on Computer Vision and Pattern Recognition*, June 2012, pp. 2759–2766.

[5] M. Firman, O. M. Aodha, S. Julier, and G. J. Brostow, "Structured prediction of unobserved voxels from a single depth image," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 5431–5440.

[6] D. T. Nguyen, B. Hua, M. Tran, Q. Pham, and S. Yeung, "A field model for repairing 3D shapes," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 5676–5684.

[7] S. Song, F. Yu, A. Zeng, A. X. Chang, M. Savva, and T. Funkhouser, "Semantic Scene Completion from a Single Depth Image," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[8] A. B. S. Guedes, T. E. de Campos, and A. Hilton, "Semantic scene completion combining colour and depth: preliminary experiments," *CoRR arXiv*, vol. 1802.04735, 2018. [Online]. Available: http://arxiv.org/abs/1802.04735

[9] M. Garbade, J. Sawatzky, A. Richard, and J. Gall, "Two stream 3D semantic scene completion," *CoRR arXiv*, vol. 1804.03550, 2018. [Online]. Available: http://arxiv.org/abs/1804.03550

[10] S. Liu, Y. HU, Y. Zeng, Q. Tang, B. Jin, Y. Han, and X. Li, "See and think: Disentangling semantic scene completion," in *Conference on Neural Information Processing Systems (NeurIPS)*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 263–274. [Online]. Available: http://papers.nips.cc/paper/7310-see-and-think-disentangling-semantic-scene-completion.pdf

[11] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 3431–3440.

[12] L. Zhang, L. Wang, X. Zhang, P. Shen, M. Bennamoun, G. Zhu, S. A. A. Shah, and J. Song, "Semantic scene completion with dense CRF from a single depth image," *Neurocomputing*, vol. 318, pp. 182–195, Nov. 2018. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0925231218310142

[13] Y. Guo and X. Tong, "View-Volume Network for Semantic Scene Completion from a Single Depth Image," in *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence ((IJCAI))*, Stockholm, Sweden, Jul. 2018, pp. 726–732. [Online]. Available: https://www.ijcai.org/proceedings/2018/101

[14] 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," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, April 2018.

[15] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 770–778.

[16] J. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 8, no. 6, pp. 679–698, Nov 1986.

[17] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.

[18] K. He and J. Sun, "Convolutional neural networks at constrained time cost," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.

[19] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgbd images," in *European Conference in Computer Vision (ECCV)*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 746–760.

[20] R. Guo, C. Zou, and D. Hoiem, "Predicting complete 3D models of indoor scenes," *CoRR arXiv*, vol. 1504.02437, 2015. [Online]. Available: http://arxiv.org/abs/1504.02437

[21] A. Handa, V. Patraucean, V. Badrinarayanan, S. Stent, and R. Cipolla, "SceneNet: Understanding real world indoor scenes with synthetic data," *CoRR arXiv*, vol. 1511.07041, 2015. [Online]. Available: http://arxiv.org/abs/1511.07041

[22] L. N. Smith, "A disciplined approach to neural network hyperparameters: Part 1 - learning rate, batch size, momentum, and weight decay," *CoRR arXiv*, vol. 1803.09820, 2018. [Online]. Available: http://arxiv.org/abs/1803.09820

[23] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning (ICML)*. New York, NY, USA: ACM, 2009, pp. 41–48. [Online]. Available: http://doi.acm.org/10.1145/1553374.1553380

[24] E. Aarts and J. Korst, *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. New York, NY, USA: John Wiley & Sons, Inc., 1989.

[25] S. Smith, P. jan Kindermans, C. Ying, and Q. V. Le, "Don't decay the learning rate, increase the batch size," in *Sixth International Conference on Learning Representations (ICLR)*, 2018.