

Understanding Real World Indoor Scenes With Synthetic Data

Ankur Handa, Viorica Pătrăucean, Vijay Badrinarayanan, Simon Stent and Roberto Cipolla

Department of Engineering,
University of Cambridge

handa.ankur@gmail.com, {vp344,vb292,sais2,rc10001}@cam.ac.uk

Abstract

Scene understanding is a prerequisite to many high level tasks for any automated intelligent machine operating in real world environments. Recent attempts with supervised learning have shown promise in this direction but also highlighted the need for enormous quantity of supervised data — performance increases in proportion to the amount of data used. However, this quickly becomes prohibitive when considering the manual labour needed to collect such data. In this work, we focus our attention on depth based semantic per-pixel labelling as a scene understanding problem and show the potential of computer graphics to generate virtually unlimited labelled data from synthetic 3D scenes. By carefully synthesizing training data with appropriate noise models we show comparable performance to state-of-the-art RGBD systems on NYUv2 dataset despite using only depth data as input and set a benchmark on depth-based segmentation on SUN RGB-D dataset.

1. Introduction

Many high level tasks in real world require some knowledge of objects present in the scene, their physical locations and the underlying physics involved as a means to understanding the scene. Autonomously navigating robots equipped with cameras can build upon the awareness and understanding of their environment to perform a range of simple to complex tasks in real world. In this work, we choose to focus on per-pixel semantic labelling as an intermediate step towards scene understanding. Indeed, not only can per-pixel labelling lead to object recognition, but it can also facilitate estimating volume and actual physical extent of the objects *e.g.* an indoor robot must understand how the floor bends or curves so as to adjust its path while navigating, a robot operating in a kitchen may need to have an idea of the volume of the objects to appropriately arrange them

in the cabinet — knowledge of the physical extent of the supporting surface can provide a rough estimate of where different objects can be placed. In 3D modelling, it is sometimes required to have precise knowledge of where different objects can be fitted or inserted into others. It is exactly in these high level tasks where the role of semantic segmentation is emphasized more than pure object detection. Furthermore, it is more likely to benefit from the context present in the scene to segment objects unlike any ‘blanket’ detection module that is generally run independently on a sliding window on the image.

Motivated by its recent success, we use deep learning as our computational framework for semantic segmentation. Inspired by the human brain [5], deep learning models have superseded many traditional approaches that relied on hand engineered features — past few years have seen a rapid proliferation of deep learning based approaches in many domains in AI. However, a major limitation of modern day deep learning models is the requirement of enormous supervised training data. Collecting big datasets can quickly become labour intensive and may not be a viable option in the end. In this work, we focus on the challenges of obtaining the desired training data for scene understanding.

Many existing datasets do not have the volume of data needed to make significant advances in scene understanding that we aim in this work. For instance, indoor scene datasets like, NYUv2 [25] contains only 795 training images for as many as 894 object classes. SUN RGB-D [27], on the other hand contains 5,285 training images for 37 classes. These are the only two indoor depth datasets with per-pixel labels and are limited in size considering the enormity of data needed to achieve good performance on unseen data. This limitation results from the difficulty of the labelling process. Indeed, per-pixel segmentation is a tedious, expensive, and even error-prone process, when done manually by human labellers. We believe scene understanding can greatly benefit from the computer graphics community that has long had the tradition of CAD model repos-

itories. Synthetic data is already used for many computer vision problems [11, 2, 3, 17, 16, 28] and [20, 18] in the context of robotics. We believe that the role of synthetic data and gaming environments [23] will continue to grow in providing training data with further advances in machine learning and data driven understanding.

Our main contribution in this work is to show the potential of synthesised ground truth depth data generated from annotated 3D scenes [19] in improving the performance of per-pixel labelling with a thorough and extensive evaluation on challenging real world indoor datasets. Consequently, we also highlight the usefulness of depth-only data, generally overlooked in favour of RGB, in segmenting functional categories of objects.

2. Related Work

Prior work on per-pixel indoor image labelling has been due to NYUv2 and SUN RGB-D datasets. The work of [8] was one of the first in the direction and built on top of a deep learning framework trained on NYUv2. However, it achieved only modest performance on the test data. Subsequently, [22] and [10] have improved the performance, again with deep learning inspired methods. We think that the potential of these new methods is yet to be explored fully and that the lack of training data is the primary hindrance. Both NYUv2 and SUN RGB-D are limited in their sizes and only provide per-pixel labels for low quality raw depth-maps and corresponding RGB frames — missing data and noise in raw sensor measurements exacerbate the problem even more. Also, since most of the labelling relies on human labour, missing labels and mislabelled data are very common, as shown in Fig. 1. This is inevitable as labelling for humans can be a tiring process and sometimes comes with a considerable monetary cost.

Ideally, one would like to have a fully labelled 3D model for every scene to generate annotations from arbitrary viewpoints but this is clearly missing in both the datasets. SUN3D [30] goes in the direction of providing annotated video sequences together with 3D point clouds obtained with SfM. However, they only provide 8 such sequences. The NYUv2 dataset provides a large number of videos, but only provide a couple of annotated frames per video. Additionally, the videos are captured without 3D reconstruction in mind and therefore not suitable for generating accurate 3D reconstructions or annotations, as observed from our own experiments. Furthermore, fusion of raw depth images within a temporal window can provide smooth depth-map to aid the segmentation as opposed to noisy raw depth-maps. SUN3D [30] also fuse the raw depth maps but again they are limited in the number of annotated sequences. Although, NYUv2 and SUN RGB-D do not provide fused depth measurements, fortunately, both datasets provide an in-painted version of the raw depth maps which



(a) Missing labels

(b) Mislabelled

Figure 1. Missing labels (a) and mislabelled frames (b) are very common in many real datasets. In (b) the toilet and sink have the same ground truth label. Both images are from SUN RGB-D[27].

has been used in [8, 22, 10] on per-pixel scene understanding. In this work, we build upon SceneNet [19], a repository of annotated synthetic 3D scenes, to collect potentially unlimited labelled training data with perfect ground truth for per-pixel labelling.

The idea of using synthetic scenes has existed in the past, in particular, [12], who released a number of scenes targetted towards the application of scene retrieval, which could potentially be used in the problem we are interested in. However, those scenes are small scale of the order of $4\text{m} \times 3\text{m} \times 3\text{m}$, and contain only one or two instances of characteristic objects that define the scene *e.g.* only one desk and monitor in the office room. On the other hand, object repositories have existed for a long time now particularly the popular Trimble Warehouse¹, ModelNet [29] and ShapeNet². Unfortunately, object repositories are not directly useful for the problem we are targetting. SceneNet contains sufficiently large and variegated scenes to provide the desired data we need in this work.

3. Synthesizing Training Data

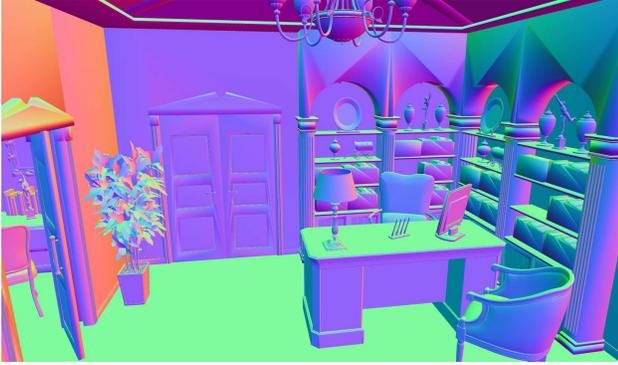
SceneNet³ [19] is inspired by the efforts developed in the graphics community to build large scale repositories of CAD models. It is an open-source repository of annotated synthetic indoor scenes — the SceneNet Basis Scenes (SNBS) — containing a significant number of manually labelled 3D models. Having a labelled scene gives a lot more flexibility in obtaining the desired training data — annotations from arbitrary view points are easily available, saving the expensive human labour required to label each image independently. Since our goal here is to show improvements on existing static-image datasets we do not render video trajectories and instead generate synthetic data from random poses.

SceneNet contains 3D models from five different scene categories: bedrooms, living rooms, offices, bathrooms, and kitchens, with at least 10 annotated scenes for each category. All the 3D models are metrically accurate and are available in standard *.obj* format. We use OpenGL to place virtual cameras in the synthetic scenes to generate ground

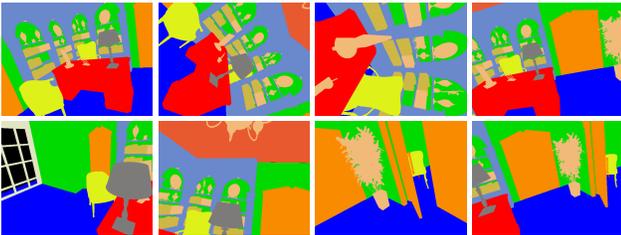
¹<http://www.sketchup.com>

²<http://shapenet.cs.stanford.edu/>

³robotvault.bitbucket.org



(a) Sample office scene from SceneNet.



(b) Examples of per-pixel semantically labelled views from the scene.

Figure 2. Annotated 3D models can conveniently allow for generation of potentially unlimited ground truth labelled data from different view-points. (a) shows one of the office scenes in SceneNet and (b) shows annotated images rendered from different view-points.

truth data from random viewpoints. Fig. 2(b) shows samples of rendered annotated views of a living room. In this work, we have used purely depth-based scene understanding. We also considered raytracing for RGB rendering but found it to be very time consuming when modelling complicated light transport. However, depth-only experiments allow us to study the effect of geometry in isolation.

Additionally, we generate new physically realistic scenes from object models downloaded from various online object repositories [29, 6] (see Table 1), using object co-occurrence statistics from SceneNet and optimising on the object positions within the scene with simulated annealing [31, 19]. Most importantly, scene generation is fully automated and sampling objects means that these scenes already come with annotations. We again render from random viewpoints to obtain the desired ground truth data and augment it to the SN-BS ground truth. In total, we use 10,030 rendered images.

Repository	Objects
ModelNet [29]	127,915
Archive3D	45,000
Stanford Database [12]	1,723

Table 1. Potential 3D object repositories for scene generation.

3.1. Adding noise to ground truth data

We realise the possible mismatch of the distribution of the noise characteristics in real world datasets and our synthetic depth-maps and therefore add noise to the perfect rendered depth-maps according to the simulated kinect noise model in [14, 20] (see Fig. 3 for visualisations). This is to ensure that any model trained on synthetic datasets can have a significant effect on real world depth data either directly or via fine-tuning [11].

4. Experimental Set-up

Our primary goal in this work is to show how synthetic data can enable improvements in per-pixel semantic segmentation on real world indoor scenes. For all our experiments, we choose the state-of-the-art segmentation algorithm [24, 4] with encoder-decoder architecture built on top of the popular VGG network [26]. Both algorithms have only been used for RGB image based segmentation, therefore, we adapted them to work on depth based three channel input, DHA,⁴ namely depth, height from ground plane and angle with gravity vector. Since we already know the camera poses when rendering, it is straightforward to obtain height and angle with gravity vector for synthetic data but we implemented a C++ GPU version of the otherwise slow MATLAB code of [15] to align the NYUv2 depth-maps to inertial frame to obtain the corresponding heights and angles. SUN RGB-D already provide the tilt angle and rotation matrices so we use them to obtain our DHA features. We initially report results on only 11 different categories (see Table 3). This is because, to generate new scenes, we sample objects from axis aligned ModelNet10 [29] which does not contain *painting* and *books* that in total add up to the 13 classes used in [10, 22]. However, we take special care in doing comparisons with [10] who made their per-pixel semantic predictions for the NYUv2 test data publicly available. We later, also report results on 13 classes by directly finetuning on the standard 13 class datasets.

Also, we inpainted the noisy depth-maps from our simulator with the MATLAB colorization code provided in the NYUv2 dataset. This was to ensure that our depth-maps qualitatively match with the inpainted depth-maps in NYUv2 (see Fig.3). We also ran the colorization with different kernel sizes and empirically found 3×3 to be good enough — our final results did not change much with 7×7 and bigger kernels.

5. Results

We collect a wide variety of random images sampled from SceneNet and convert each depth image into three

⁴A similar HHA encoding has been used in [16] but like [10] we did not observe much difference.

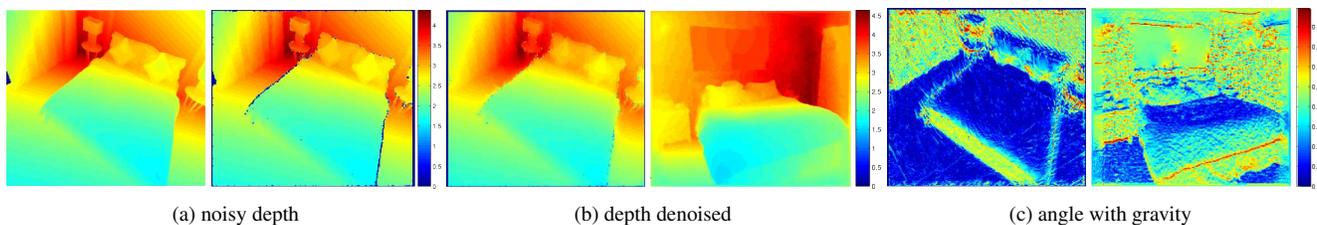


Figure 3. Left image in (a) shows the perfect rendered depth map and right image shows the corresponding noisy image provided by our simulator. In (b) we show a side by side comparison of the inpainted depth maps of one of the bedroom scenes in SceneNet with a similar looking bedroom in NYUv2 test data and (c) shows the comparison of angle with gravity vector images. Importantly, left image in (c) highlights the view point invariance of the angle with gravity.

channel DHA input. Our network is first trained on purely synthetic data and then fine tuned on the training images provided in the NYUv2 and SUN RGB-D datasets. Finally, NYUv2 and SUN RGB-D test data is used to compare the results with different variants of training data used in experiments.

We denote training performed on noise injected synthetic data from SceneNet, as SceneNet-DHA, and fine-tuning on training data in NYUv2 by SceneNet-FT-NYU-DHA. When using dropout we denote them by SceneNet-DO-DHA and SceneNet-FT-NYU-DO-DHA respectively. Similarly, the networks trained on NYU are denoted by NYU-DHA and NYU-DO-DHA. We also tried dropout at test time as mentioned in [13] but this comes at the expense of increased inference time — dropout at test time is shown to have similar effect of multiple model averaging and at the same time provides uncertainty in the final predictions. First, we quantify the performance of all these variants on the NYUv2 test data and later on a bigger real world dataset, SUN RGB-D, based on standard global and class accuracy metrics.

We use DHA images of size 224×224 as input to the network and initialise the weights with pretrained VGG network. At first, it may seem that the network trained on images cannot be useful for training on depth data, however, since the first layer filters often look for edges and orientations [32], it is sensible to use them for relatively similar modalities like depth — edges in depth almost always align with the RGB edges — and deeper layers adapt accordingly. Therefore, in all our experiments the network always converged with accuracies in high nineties on the training set. Also, it is worth noting that FCN [22] quantise the HHA to $[0,255]$ for an equivalent RGB image and use pretrained VGG network to initialise. We do not quantise and maintain the floating point precision of the DHA images.

We render 10,030 depth images from random view points ensuring that a minimum number of objects is visible — avoiding camera looking at only walls or floor — and perturb the depth values according to Section 3.1 to generate depth maps qualitatively similar to NYUv2. Comparison of different proportions of objects in these rendered

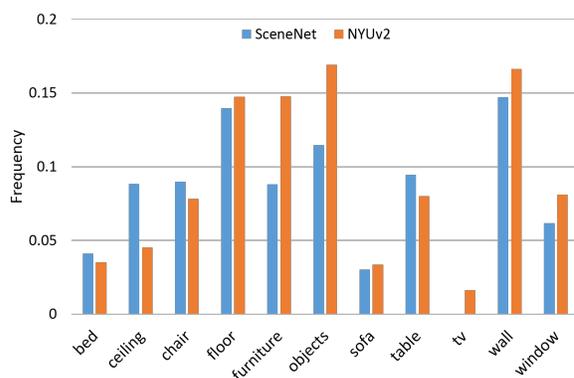


Figure 5. Side by side comparison of proportions of objects in the 9K images rendered with SceneNet and 795 training images in NYUv2.

images and NYUv2 training data is shown Fig. 5. All the models are trained with stochastic gradient descent with a starting learning rate of 0.01 which is multiplied by 0.9 every 3-4 epochs, a weight decay of 0.0005, and a momentum of 0.9. We characterise the experiments into comparisons with real training data in NYUv2 and the state-of-the-art, Eigen & Fergus [10], for both 11 and 13 class segmentations. We also perform experiments on SUN RGB-D for 13 classes and set a new benchmark for pure depth based segmentation.

Comparisons with NYUv2 training Training purely on synthetic data, SceneNet-DHA, results in only modest performance on the test data (see Table 2). Comparison with NYU-DHA reveals that fine-tuning is needed to obtain further improvements in the results. As a result, we see that the performance jump from NYU-DHA to SceneNet-FT-NYU-DHA is very clear — an increase of 5.4% in the class and 3.6% in the global accuracy showing the usefulness of synthetic data for real world scene segmentation. Importantly, convergence was twice as fast for SceneNet-FT-NYU-DHA compared to NYU-DHA. Specifically, prior to fine-tuning we trained SceneNet-DHA for 21K iterations (14 epochs). Fine-tuning on NYUv2 took another 4K iter-



Figure 4. Results for 11 classes on NYUv2 test data obtained with SceneNet-FT-NYU-DHA and NYU-DHA. First row shows the RGB images, predictions returned by SceneNet-FT-NYU-DHA and NYU-DHA are shown in second and third row respectively. Last row shows the ground truth.

ations (30 epochs) while NYU-DHA took about 7.5K iterations (56 epochs) to converge. Qualitative results of the segmentation are shown in Fig. 4 which highlights the impact of synthetic data in the final results.

Comparisons with Eigen & Fergus [10] We also compare our results to the state-of-the-art system of Eigen and Fergus [10] who use data augmentation and a combination of RGBD and normals. Since we use only depth, our system is not directly comparable to [10] but we obtain competitive performance. SceneNet-FT-NYU-DHA although performs better than NYU-DHA, it still falls short of the performance of the state-of-the-art method [10]. However, careful examination in class accuracies (see Table 3) reveals that we perform comparably only compromising on *tv* and *windows* — we expect RGB to play a bigger role here — emphasizing that for functional categories of objects geometry is a strong

cue for segmentation. We used the publicly available annotations of [10] with AlexNet⁵, to re-evaluate the accuracies for the 11 classes we use.

Comparisons with Dropout We also used dropout ratio of 0.5 in the fully connected layers to gauge the effect of regularisation in the parameters to prevent over-fitting. We see 3.2% improvement in class and 1.9% global accuracy in the results with real training data used in NYU-DO-DHA compared to NYU-DHA. However, we only observed minor improvements with both SceneNet-DO-DHA and SceneNet-FT-NYU-DO-DHA against SceneNet-DHA and SceneNet-FT-NYU-DHA respectively suggesting that increase in data acts as an implicit regulariser. Further improvements depend largely on the amount of training data

⁵at the time of publication only AlexNet annotations were available

11 Class Semantic Segmentation: NYUv2

Training	bed	ceiling	chair	floor	furn	objs.	sofa	table	tv	wall	window
NYU-DHA	64.5	68.2	51.0	95.0	51.0	48.2	49.7	41.9	12.8	84.2	24.5
SceneNet-DHA	60.8	43.4	68.5	90.0	26.5	24.3	21.2	42.1	0	92.1	0.3
SceneNet-FT-NYU-DHA	70.3	75.9	59.8	96.0	60.7	49.7	59.9	49.7	24.3	84.8	27.9
NYU-DO-DHA	69.0	74.6	54.0	95.6	57.1	48.7	55.7	42.5	18.5	84.7	25.5
SceneNet-DO-DHA	67.7	40.9	67.5	87.8	38.6	22.6	15.8	44.2	0	89.0	0.8
SceneNet-FT-NYU-DO-DHA	72.5	74.1	61.0	96.2	60.4	50.0	62.8	43.8	19.4	85.3	30.0
Eigen <i>et al.</i> (rgbd+normals) [10]	61.1	78.3	72.1	96.5	55.1	52.1	45.8	45.0	41.9	88.7	57.7

Table 3. Different training variants on the NYUv2 test set. The benefits of adding synthetic data are clear. For fair comparison, it should be noted that [10] use RGBD+Normals than just depth as used in this work. We would like to stress here that our network is trained end-to-end as compared to multi-stage training done in Eigen *et al.* [10] and does not use any RGB data or data augmentation. We outperform [10] on *sofas* and *beds* but fall behind on *chairs*, *tv* and *windows*. We expect *tv* and *windows* are likely to be segmented better with RGB data. However, we obtain comparable performance on rest of the classes further emphasising that for functional categories of objects shape is a strong cue. Poor performance of SceneNet-DHA and SceneNet-DO-DHA on *tv* and *windows* is mainly due to limited training data for these classes in SceneNet. Also note that results of Hermans *et al.* [21] are not directly comparable for 11 classes because they do not make their predicted annotations publicly available.

11 Class Semantic Segmentation: NYUv2

Training	global acc.	class acc.
SceneNet-DHA	54.4	42.6
NYU-DHA	63.8	53.7
SceneNet-FT-NYU-DHA	67.4	59.1
SceneNet-DO-DHA	54.6	43.2
NYU-DO-DHA	65.7	56.9
SceneNet-FT-NYU-DO-DHA	68.0	59.9
Eigen <i>et al.</i> (rgbd+normals) [10]	69.5	63.0

Table 2. Different variants of training data that we use in our experiments. The performance jump from NYU-DHA to SceneNet-FT-NYU-DHA is clear. Adding dropout helps most in the NYU-DO-DHA but SceneNet-FT-NYU-DO-DHA shows only a marginal improvement. Overall increase in performance from NYU-DO-DHA to SceneNet-FT-NYU-DO-DHA is by 2.3% and 3% in global and class accuracy respectively. Note that we recomputed the accuracies of [10] using their publicly available annotations of 320×240 and resizing them to 224×224 . Note that Hermans *et al.* [21] predictions are not publicly available. As a result, we cannot evaluate their performance on these 11 classes.

used [1]. Although there is no limit to the amount of training data we can render we are only limited by GPU speed and memory — training time for a batch size of 6 images takes about 0.8s for forward pass and 1.2s for backward pass, a combined total of 2s per iteration on NVidia Tesla K80. We also anticipate that training can be made efficient by either completely forgoing the fully connected layers [4] or reducing their number of features [7] — fully connected layers in VGG contain nearly $\frac{204800}{246912} \sim 83\%$ of the parameters. Training on data of the order of ImageNet [9] remains an exciting opportunity for the future.

We also used dropout at test time [13] but observed very similar performance gain without it. However, dropout at

test time [13] makes the network robust to out-of-domain data. We leave it as an interesting future direction to explore. Overall, we have SceneNet-FT-NYU-DO-DHA as a clear winner against NYU-DO-DHA as shown in Table 2.

Results on 13 Class Semantic Segmentation: NYUv2

We also performed our experiments on the 13 class semantic segmentation task. It is worth remembering that the two extra classes we add to the 11 class experiment are *painting* and *books*. Although, we are limited by the number of labels for *painting* and *books* in SceneNet, we fine tune directly on 13 class semantic segmentation task on NYUv2. The performance gain from NYU-DHA to SceneNet-FT-NYU-DHA is evident, highlighting the role of synthetic data. As seen in Table 4, SceneNet-FT-NYU-DHA performs consistently better than NYU-DHA for all the classes. We observe similar trend for the comparison between NYU-DO-DHA and SceneNet-FT-NYU-DO-DHA. It is worth remembering that Eigen & Fergus [10] use RGB and normals together with the depth channel and hence maintain superior overall performance over our methods that use only depth data. Again, careful examination reveals that we only compromise on *books*, *painting*, *tv* and *windows* — we expect RGB to play a bigger role here in segmenting these classes. This is also reflected in the overall mean global and class accuracies as shown in Table 5. Figures 6 and 7 compare predictions returned by SceneNet-FT-NYU-DHA, NYU-DHA and Eigen & Fergus [10] on a variety of test images in NYUv2 dataset.

Results on 13 Class Semantic Segmentation: SUN RGB-D

We perform similar experiments on 13 classes on SUN RGB-D. The dataset provides 5,825 training images and 5,050 images for testing in total. This is one order of

13 Class Semantic Segmentation: NYUv2

Training	bed	books	ceil.	chair	floor	furn	objs.	paint.	sofa	table	tv	wall	window
NYU-DHA	67.7	6.5	69.9	47.9	96.2	53.8	46.5	11.3	50.7	41.6	10.8	85.0	25.8
SceneNet-DHA	60.8	2.0	44.2	68.3	90.2	26.4	27.6	6.3	21.1	42.2	0	92.0	0.0
SceneNet-FT-NYU-DHA	70.8	5.3	75.0	58.9	95.9	63.3	48.4	15.2	58.0	43.6	22.3	85.1	29.9
NYU-DO-DHA	69.6	3.1	69.3	53.2	95.9	60.0	49.0	11.6	52.7	40.2	17.3	85.0	27.1
SceneNet-DO-DHA	67.9	4.7	41.2	67.7	87.9	38.4	25.6	6.3	16.3	43.8	0	88.6	1.0
SceneNet-FT-NYU-DO-DHA	70.8	5.5	76.2	59.6	95.9	62.3	50.0	18.0	61.3	42.2	22.2	86.1	32.1
Eigen <i>et al.</i> (rgbd+normals) [10]	61.1	49.7	78.3	72.1	96.0	55.1	40.7	58.7	45.8	44.9	41.9	88.7	57.7
Hermans <i>et al.</i> (rgbd+crf)[21]	68.4	N/A	83.4	41.9	91.5	37.1	8.6	N/A	28.5	27.7	38.4	71.8	46.1

Table 4. Results on NYUv2 test data for 13 semantic classes. We see a similar pattern here — adding synthetic data helps immensely in improving the performance of nearly all functional categories of objects using DHA as input channels. As expected, accuracy on *books*, *painting*, *tv*, and *windows*, is compromised highlighting that the role of depth as a modality to segment these objects is limited. Note that we recomputed the accuracies of [10] using their publicly available annotations of 320×240 and resizing them to 224×224. Hermans *et al.* [21] use “*Decoration*” and “*Bookshelf*” instead of *painting* and *books* as the other two classes. Therefore, they are not directly comparable. Also, their annotations are not publicly available but we have still added their results in the table. Note that they use 640×480. Poor performance of SceneNet-DHA and SceneNet-DO-DHA on *tv* and *windows* is mainly due to limited training data for these classes in SceneNet.

13 Class Semantic Segmentation: SUNRGBD

Training	bed	books	ceil.	chair	floor	furn	objs.	paint.	sofa	table	tv	wall	window
SceneNet-DHA	33.2	2.5	40.6	54.0	71.1	26.2	22.1	9.5	15.0	29.2	0	89.2	0.0
SceneNet-DO-DHA	46.1	5.2	43.6	54.8	63.1	37.4	23.2	10.7	12.2	29.8	0	83.6	1.0
SUNRGBD-DHA	70.4	11.2	64.7	69.2	94.0	48.4	35.3	13.7	48.2	63.0	3.5	89.7	27.9
SUNRGBD-DO-DHA	73.6	16.6	71.6	70.1	93.5	47.9	38.7	17.2	58.5	61.8	6.8	88.7	33.9
SceneNet-FT-SUNRGBD-DHA	69.0	20.0	70.3	70.7	93.7	49.7	35.5	15.7	57.8	65.9	14.1	89.0	33.8
SceneNet-FT-SUNRGBD-DO-DHA	75.6	13.5	69.2	73.6	93.8	52.0	37.1	16.8	57.2	62.7	9.5	88.8	36.5

Table 6. Results on SUNRGBD test data for 13 semantic classes. We see a similar pattern here — adding synthetic data helps immensely in improving the performance of nearly all functional categories of objects using DHA as input channels. As expected, accuracy on *books*, *painting*, *tv*, and *windows*, is compromised highlighting that the role of depth as a modality to segment these objects is limited.

magnitude bigger in size than NYUv2. As shown in Table 7, we observe similar trends *i.e.* training on synthetic data and finetuning on real dataset helps in improving the accuracies by 1% and 3.5% in global and class accuracy respectively when comparing SUNRGBD-DHA and SceneNet-FT-SUNRGBD-DHA. However, we quickly see diminishing returns when using dropout *i.e.* SceneNet-FT-SUNRGBD-DHA and SceneNet-FT-SUNRGBD-DO-DHA perform nearly the same. Furthermore, SceneNet-FT-SUNRGBD-DO-DHA performs only 0.8% and 0.9% better in global and class accuracy respectively, compared to SUNRGBD-DO-DHA. It is worth remembering that when experimenting with NYUv2, the proportion of synthetic data was 10 times the real training data while a bigger SUN RGB-D dataset means that this proportion is only 2 times ($\frac{10,030}{5,825} \sim 2$) the real training data, suggesting that further improvements can be possible either through another order

of magnitude increase in data or a possible change in the architecture. Nonetheless, we set a benchmark on pure depth based segmentation on SUN RGB-D dataset. Breakdown of class accuracies for different training variants is shown in Table 6.

6. Conclusions

We presented an effective solution to the problem of indoor scene understanding — system trained with large number of rendered synthetic depth frames is able to achieve near state-of-the-art performance on per-pixel image labelling despite using only depth data. We specifically show that adding synthetic data improves the performance on NYUv2 and SUN RGB-D for depth-only experiments and offers a promising route in further improvements in the state-of-the-art. We hope to continue indoor scene segmentation on real world scenes with a reconstruction system

13 Class Semantic Segmentation: NYUv2

Training	global acc.	class acc.
Hermans <i>et al.</i> (rgbd+crf) [21]	54.2	48.0
SceneNet-DHA	54.4	37.1
NYU-DHA	63.6	47.2
SceneNet-FT-NYU-DHA	66.5	51.7
SceneNet-DO-DHA	55.3	37.6
NYU-DO-DHA	65.0	48.6
SceneNet-FT-NYU-DO-DHA	67.2	52.5
Eigen <i>et al.</i> (rgbd+normals) [10]	68.0	60.8

Table 5. The performance jump from NYU-DHA to SceneNet-FT-NYU-DHA follows similar trend — the role of synthetic data is evident. Adding dropout shows only a marginal improvement. Overall increase in performance from NYU-DO-DHA to SceneNet-FT-NYU-DO-DHA is by 2.2% and 3.9% in global and class accuracy respectively. Note that we recomputed the accuracies of [10] using their publicly available annotations of 320×240 and resizing them to 224×224 . Hermans *et al.*[21] annotations are not publicly available but we have still added their results in the table. Moreover, they use 640×480 .



Figure 6. Results for 13 class labels: First row is the RGB image. Second row shows the results obtained with SceneNet-FT-NYU-DHA while the third row shows the results with NYU-DHA. Fourth row has the predictions returned by Eigen and Fergus [10]. Last row shows the ground truth. Notice that when the camera is viewing at cluttered scenes or far away objects, our results have more blob-like segments in the output. Also, note that monitors tend to get confused with chairs for all algorithms.

running in the loop to bring real-time semantic segmentation using fused depth-maps.

13 Class Semantic Segmentation: SUNRGBD

Training	global acc.	class acc.
SceneNet-DHA	56.9	30.2
SUNRGBD-DHA	73.7	49.2
SceneNet-FT-SUNRGBD-DHA	74.7	52.7
SceneNet-DO-DHA	54.7	31.6
SUNRGBD-DO-DHA	74.2	52.2
SceneNet-FT-SUNRGBD-DO-DHA	75.0	53.1

Table 7. Global and class accuracies for 13 class experiments on SUN RGB-D. We see improvements of 1% and 3.5% in global and class accuracy comparing SUNRGB-DHA and SceneNet-FT-SUNRGBD-DHA. However, when using dropout, SUNRGBD-DO-DHA, SceneNet-FT-SUNRGBD-DHA and SceneNet-FT-SUNRGBD-DO-DHA perform nearly the same suggesting that increase in data is only helpful up to a point and that further improvements can be possible either through another order of magnitude of data as seen in NYUv2 experiments or a possible change in the architecture.

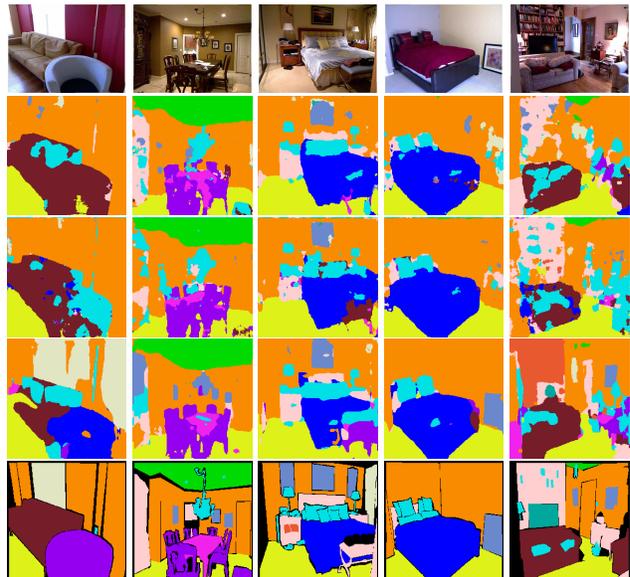


Figure 7. Results for 13 class labels: First row is the RGB image. Second row shows the results obtained with SceneNet-FT-NYU-DHA while the third row shows the results with NYU-DHA. Fourth row has the predictions returned by Eigen and Fergus [10]. Last row shows the ground truth. Looking at the results, it is evident that the predictions by [10] tend to be smoother and contained unlike NYU-DHA as well as SceneNet-FT-NYU-DHA. We believe this difference is a result of rich input channels in RGB, normals and depth used by [10] while we use only depth in our experiments.

7. Acknowledgements

We thank Andrew Davison, Daniel Cremers, and Zoubin Ghahramani, and colleagues at CUED, Cambridge for their encouragement and valuable suggestions.

References

- [1] <http://cs229.stanford.edu/materials/CS229-DeepLearning.pdf>.
- [2] M. Aubry, D. Maturana, A. Efros, B. Russell, and J. Sivic. Seeing 3D chairs: exemplar part-based 2D-3D alignment using a large dataset of CAD models. In *CVPR*, 2014.
- [3] M. Aubry and B. Russell. Understanding deep features with computer-generated imagery. In *ICCV*, 2015.
- [4] V. Badrinarayanan, A. Handa, and R. Cipolla. SegNet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling. *arXiv preprint arXiv:1505.07293*, 2015.
- [5] C. F. Cadieu, H. Hong, D. L. K. Yamins, N. Pinto, D. Ardila, E. A. Solomon, N. J. Majaj, and J. J. DiCarlo. Deep Neural Networks Rival the Representation of Primate IT Cortex for Core Visual Object Recognition. *PLoS Computational Biology*, 2014.
- [6] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu. ShapeNet: An Information-Rich 3D Model Repository. *CoRR*, abs/1512.03012.
- [7] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. In *ICLR*, 2015.
- [8] C. Couprie, C. Farabet, L. Najman, and Y. LeCun. Indoor semantic segmentation using depth information. In *ICLR*, 2013.
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009.
- [10] D. Eigen and R. Fergus. Predicting Depth, Surface Normals and Semantic Labels with a Common Multi-Scale Convolutional Architecture. In *ICCV*, 2015.
- [11] P. Fischer, A. Dosovitskiy, E. Ilg, P. Häusser, C. Hazrbaş, V. Golkov, P. van der Smagt, D. Cremers, and T. Brox. FlowNet: Learning Optical Flow with Convolutional Networks. In *ICCV*, 2015.
- [12] M. Fisher, D. Ritchie, M. Savva, T. Funkhouser, and P. Hanrahan. Example-based synthesis of 3d object arrangements. In *ACM SIGGRAPH Asia*, SIGGRAPH Asia, 2012.
- [13] Y. Gal and Z. Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. *arXiv preprint arXiv:1506.02142*, 2015.
- [14] M. Gschwandtner, R. Kwitt, A. Uhl, and W. Pree. BlenSor: blender sensor simulation toolbox. In *Advances in Visual Computing*, pages 199–208. Springer, 2011.
- [15] S. Gupta, P. Arbelaez, and J. Malik. Perceptual organization and recognition of indoor scenes from RGB-D images. In *CVPR*, 2013.
- [16] S. Gupta, P. A. Arbeláez, R. B. Girshick, and J. Malik. Aligning 3D models to RGB-D images of cluttered scenes. In *CVPR*, 2015.
- [17] S. Gupta, R. Girshick, P. Arbelaez, and J. Malik. Learning Rich Features from RGB-D Images for Object Detection and Segmentation. In *ECCV*, 2014.
- [18] A. Handa, R. A. Newcombe, A. Angeli, and A. J. Davison. Real-Time Camera Tracking: When is High Frame-Rate Best? In *ECCV*, 2012.
- [19] A. Handa, V. Pătrăucean, S. Stent, and R. Cipolla. SceneNet: an Annotated Model Generator for Indoor Scene Understanding. In *ICRA*, 2016.
- [20] A. Handa, T. Whelan, J. B. McDonald, and A. J. Davison. A Benchmark for RGB-D Visual Odometry, 3D Reconstruction and SLAM. In *ICRA*, 2014.
- [21] A. Hermans, G. Floros, and B. Leibe. Dense 3d semantic mapping of indoor scenes from rgb-d images. In *ICRA*, 2014.
- [22] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [23] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 2015.
- [24] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. *arXiv preprint arXiv:1505.04366*, 2015.
- [25] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from RGBD images. In *ECCV*, 2012.
- [26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [27] S. Song, S. P. Lichtenberg, and J. Xiao. SUN RGB-D: A RGB-D scene understanding benchmark suite. In *CVPR*, pages 567–576, 2015.
- [28] J. Wu, I. Yildirim, W. Freeman, and J. Tenenbaum. Perceiving physical object properties by integrating a physics engine with deep learning. In *NIPS*, 2015.
- [29] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. Xiao. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *CVPR*, 2015.
- [30] J. Xiao, A. Owens, and A. Torralba. Sun3d: A database of big spaces reconstructed using sfm and object labels. In *ICCV*, 2013.
- [31] L.-F. Yu, S. K. Yeung, C.-K. Tang, D. Terzopoulos, T. F. Chan, and S. Osher. Make It Home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics*, 30(4):86, 2011.
- [32] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*. 2014.