

Multichannel Semantic Segmentation with Unsupervised Domain Adaptation

Kohei Watanabe¹, Kuniaki Saito¹, Yoshitaka Ushiku¹, and Tatsuya Harada^{1,2}

¹ The University of Tokyo

{watanabe, k-saito, ushiku, harada}@mi.t.u-tokyo.ac.jp

² RIKEN

Abstract. Most contemporary robots have depth sensors, and research on semantic segmentation with RGBD images has shown that depth images boost the accuracy of segmentation. Since it is time-consuming to annotate images with semantic labels per pixel, it would be ideal if we could avoid this laborious work by utilizing an existing dataset or a synthetic dataset which we can generate on our own. Robot motions are often tested in a synthetic environment, where multichannel (*e.g.*, RGB + depth + instance boundary) images plus their pixel-level semantic labels are available. However, models trained simply on synthetic images tend to demonstrate poor performance on real images. In order to address this, we propose two approaches that can efficiently exploit multichannel inputs combined with an unsupervised domain adaptation (UDA) algorithm. One is a fusion-based approach that uses depth images as inputs. The other is a multitask learning approach that uses depth images as outputs. We demonstrated that the segmentation results were improved by using a multitask learning approach with a post-process and created a benchmark for this task.

Keywords: semantic segmentation, domain adaptation, RGB-Depth, multi-task learning

1 Introduction

Semantic segmentation is a fundamental task for robots to understand their surroundings in detail. Most robots have depth sensors. In fact, research on semantic segmentation with RGBD images has been conducted and has demonstrated that depth images boost the accuracy of segmentation [13]. However, semantic pixel-level labels are necessary to train semantic segmentation models in general and it is time-consuming to annotate the image per pixel. For instance, the pixel labeling of one Cityscapes image takes 1.5 hours on average [6]. It would be ideal to avoid this laborious work by utilizing an existing dataset or a synthetic dataset which we could generate on our own.

Recently, the number of RGBD datasets taken in the real world has increased. In addition to the widely used 2.5D dataset such as NYUDv2 [31] and SUNRGBD [32], large-scale real 3D datasets such as Stanford2d3d [2, 1], ScanNet [7] have been generated due to the development of the 3D scanner and

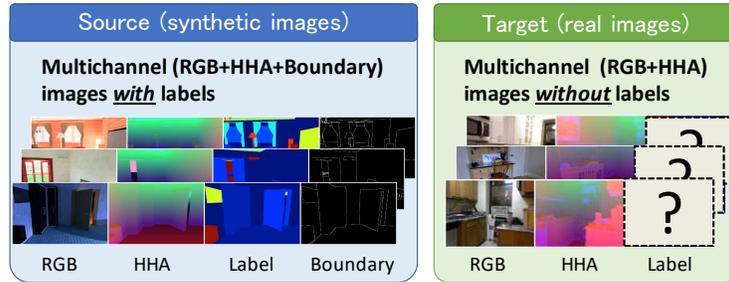


Fig. 1: Setting of this research. Left is samples of SUNCG [38] and right is samples of NYUDv2 [31]. HHA is the three dimensional encoding of depth [11].

scalable RGB-D capture system. However, the number of real datasets is small compared to the RGB dataset.

Conversely, computer graphics technology has been developed and large-scale synthetic datasets have also been generated. For example, SUNCG [38] contains 400K physically-based rendered images from 45K realistic 3D indoor scenes. SceneNet [24] contains 5 million images rendered of 16,895 indoor scenes. It is also possible to purchase 3D CAD models online and create customized synthetic datasets using UnrealCV [28]. The appearance of synthetic images is a bit different from that of real ones but the synthetic datasets still look real. In fact, it is ideal if a model trained on these dataset performs well on real datasets because robot motions are often tested in a synthetic environment before being tested in a real environment.

However, such a model is known not to generalize well because of the pixel-level distribution shift [15]. In order to solve this problem, a domain adaptation technique is necessary. Although several research studies on unsupervised domain adaptation for semantic segmentation have been conducted, they use only RGB input and do not consider the utilization of a multichannel (here we mean RGB + depth images, which are now easy to obtain in both synthetic and real environments).

We propose two approaches that can efficiently use multichannel inputs with an UDA algorithm. One is a fusion-based approach that uses different modal images as inputs and the other is a multitask learning approach that uses only RGB images as inputs but other modal images as outputs. Fusing different modalities (RGB, depth or boundary) efficiently is known to boost the segmentation accuracy compared to a simple concatenation of inputs known as *early fusion* in past research. Except for *early fusion*, many fusion methods [13, 5, 26, 12] exist and their efficacy is task-specific, which makes us rethink their ideas. Multitask learning is also a promising approach. Multitask learning that solves related tasks such as semantic segmentation and depth estimation tasks simultaneously is known to boost each task’s performance [17, 19]. In multitask learning it is easy to add another single task, such as a boundary detection task, which can be thought to render feature maps more aware of boundaries. Boundary de-



Fig. 2: Example of network architectures for domain adaptation [(a) domain adversarial neural network (DANN) [15] and (b) maximum classifier discrepancy (MCD) [30]]. When inputs belong to source samples which have labels, we train a feature generator and a classifier using the labels. When inputs belong to target samples which have **no** labels, we train a feature generator and (one discriminator in the case DANN is used or two classifiers in the case MCD is used) in an adversarial manner.

tection output can be utilized collaterally to refine the messy domain-adapted segmentation output.

In summary, the specific contribution of this paper includes:

- We combine a multichannel semantic segmentation task with an unsupervised domain adaptation (UDA, see Fig. 1) task and propose two approaches (fusion-based and multitask learning)
- We show that the multitask learning approach outperforms the simple *early fusion* approach according to all evaluation metrics.
- We propose adding a boundary detection task to the multitask learning approach and use the detection result to refine the segmentation output, which improves both the qualitative and quantitative results.

2 Related work

Here, we describe two related research themes, *domain adaptation for semantic segmentation* and *semantic segmentation with multichannel image*.

2.1 Domain adaptation for semantic segmentation

When we train a classifier in one (source) domain and apply it to classify samples in a new (target) domain, the classifier is known not to generalize well in the new domain due to the domain’s difference. Many methods tackle the problem by aligning distributions of features between the source and target domain [9, 23, 34, 3]. These methods are proposed to deal with classification problem. Recently, methods for semantic segmentation have been proposed too. Hoffman *et al.* [15] first tackled this problem. They adopted an adversarial training framework which has a feature extractor and a discriminator (see Fig. 2a). A discriminator tries

to detect whether the extracted feature correctly comes from source samples or target samples, while a feature extractor tries to generate features that deceive a discriminator in an adversarial manner. The other researches on this theme also leverage adversarial training. Zhang *et al.* [37] adopted curriculum learning that starts the easier task (global and super pixel label distribution of source samples matching those of target samples), then tries to solve difficult tasks (semantic segmentation). Cheng *et al.* [4] tried to tackle the cross-city adaptation problem via adversarial training and extract static-object priors that can be obtained from the Google Street View time-machine feature.

We utilize Saito *et al.* [30]’s method (MCD), which is shown to be effective in segmentation task. They proposed a method that uses two classifiers’ difference of output (called discrepancy) to align features between the source and target domain. They trained one feature extractor network and two different classifier networks for the same task (see Fig. 2b). Two classifiers are trained to increase the discrepancy for target samples whereas feature extractor is trained to decrease it. Details are in Sec. 3.1.

2.2 Semantic segmentation with multichannel (RGBD) image

Previously, RGBD segmentation was conducted based on handcrafted features specifically designed for capturing depth as well as color features [10]. Long *et al.* [22] proposed a fully convolutional neural network (FCN) for semantic segmentation. FCN not only replaced the fully connected layer in classification models such as AlexNet [18] with a convolutional layer but also proposed using two methods, *deconvolution* and *shortcut*, which are now widely used in many semantic segmentation models. Long *et al.* also reported the segmentation scores of the NYUDv2 dataset [31], where RGB + Depth were combined in the input (we call this *early fusion*) and RGB + HHA in output (we call this *score fusion*). HHA is the three dimensional encoding of depth (Horizontal disparity, Height above ground, and the Angle of the local surface normal with the inferred gravity direction) proposed by Gupta *et al.* [11]. FuseNet [13] prepared an RGB and depth encoder separately then fused the two encoders in certain middle layers (see Fig. 4d). The locality sensitive deconvolutional network with Gate Fusion [5] used an affinity matrix embedded with pairwise relations between neighboring RGB-D pixels to recover sharp boundaries of FCN maps. Gate Fusion (see Fig. 4c) learns to adjust to the contributions of RGB and depth that exist in the last layer of the network. RDFNet [26] fuses two networks with multi-modal feature fusion blocks and multi-level feature refinement blocks following RefineNet [21].

The above-mentioned approaches utilize all the different modals as input, but there is also an approach that utilizes only RGB as input and the other modals as output; this is the multitask learning approach. Multitask learning is a promising approach for efficiently and effectively addressing multiple mutually-related recognition tasks and its performance is known to outperform that of the single task methods. Kendall *et al.* worked on three tasks (semantic and instance

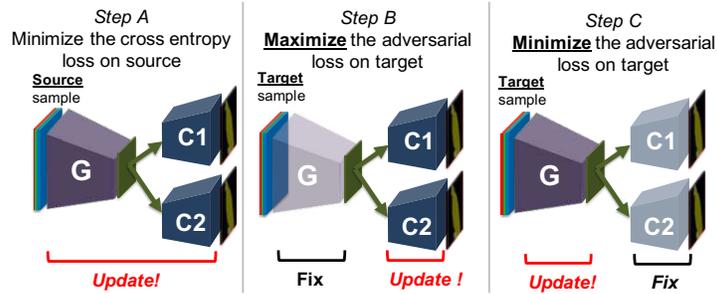


Fig. 3: Adversarial training steps in MCD [30].

segmentation, and depth estimation) [17] and Kuga *et al.* also worked on three tasks (RGB reconstruction, semantic segmentation, and depth estimation) [19].

There are other approaches using geometric cues obtained from depth images [27, 20], but in this research, we just focus on fusion-based and multitask learning approaches, which renders our model not only applicable to geometric applications but also other modal images, such as thermal images.

3 Proposed models

Our objective is to conduct unsupervised semantic segmentation with multichannel input. In order to realize that, the two required functionalities are:

- Annotation free (using no labels in a target dataset)
- Using different modalities [RGB, Depth (HHA), (Boundary)] efficiently

3.1 Annotation free (using no labels in a target dataset)

In order to satisfy the former function, a simple solution utilizes an existing dataset or synthetic dataset which we can generate on our own. However, if there is a domain shift (a difference of appearance or label distribution) between existing training data and test data, the performance can be poor. We tackle this case because a domain shift usually exists between synthetic and real datasets. Hence, we use a domain adaptation algorithm, which leverages adversarial training and enables the model to extract domain-robust features. In order to adopt the adversarial training algorithms, we separate an end-to-end segmentation model into a feature generator and a classifier as shown in Fig. 2. To utilize MCD [30], we prepare two classifiers (C_1 , C_2) and train them by three steps as shown in Fig. 3.

Formulation: We have access to a labeled source RGB image $\mathbf{x}_{\text{RGB}}^s$, HHA image $\mathbf{x}_{\text{HHA}}^s (= \mathbf{y}_2^s)$, instance boundary image \mathbf{y}_3^s and a corresponding semantic segmentation label \mathbf{y}_1^s drawn from a set of labeled source images $\{X_{\text{RGB}}^s, Y_1^s, Y_2^s (= X_{\text{HHA}}^s), Y_3^s\}$, as well as an unlabeled target image $\mathbf{x}_{\text{RGB}}^t, \mathbf{x}_{\text{HHA}}^t$ drawn from

unlabeled target images $\{X_{\text{RGB}}^t, Y_2^t (= X_{\text{HHA}}^t)\}$. We train a feature generator network G , which takes inputs \mathbf{x}^s or \mathbf{x}^t , and classifier networks C_1 and C_2 , which take features from G . C_1 and C_2 classify them into K classes per pixel, that is, they output a $(K \times |\mathbf{x}|)$ -dimensional vector of logits. Note that $|\mathbf{x}|$ denotes the number of pixels per image. We obtain class probabilities by applying the softmax function for the vector. We use the notation $p_1(\mathbf{y}_1|\mathbf{x})$, $p_2(\mathbf{y}_1|\mathbf{x})$ to denote the $(K \times |\mathbf{x}|)$ -dimensional probabilistic outputs for input \mathbf{x} obtained by C_1 and C_2 respectively.

Step A We train G , C_1 and C_2 to classify the source samples correctly. In order to make classifiers and generator obtain task-specific discriminative features, this step is crucial. We train the networks to minimize softmax cross entropy. The objective is as follows:

$$\min_{G, C_1, C_2} \mathcal{L}_{\text{seg}}(X_{\text{RGB}}^s, Y_1^s) \quad (1)$$

$$\mathcal{L}_{\text{seg}}(X_{\text{RGB}}^s, Y_1^s) = Y_1^s \log p(Y_1^s | X_{\text{RGB}}^s) \quad (2)$$

Step B We train C_1 and C_2 as a discriminator with fixing G . Let $\mathcal{L}_{\text{adv}}(X_t)$ be the adversarial loss that can be computed using target sample. This loss measures the discrepancy of C_1 and C_2 . A classification loss on the source samples is also added for better performance. The same number of source and target samples were randomly chosen to update the model. The objective is as follows:

$$\min_{C_1, C_2} \mathcal{L}_{\text{seg}}(X_{\text{RGB}}^s, Y_1^s) - \mathcal{L}_{\text{adv}}(X_{\text{RGB}}^t). \quad (3)$$

Step C We train G to minimize the adversarial loss with fixing C_1 and C_2 . The objective is as follows:

$$\min_G \mathcal{L}_{\text{adv}}(X_{\text{RGB}}^t). \quad (4)$$

The target and source images feed to the training randomly and these three steps are repeated until convergence of all the parts (classifiers and generator). The order of the three steps is not important but it is important to train the classifiers and generator in an adversarial manner under the condition that they can classify source samples correctly.

However, this still outputs messy segmentation results for the indoor scene recognition task (see Fig. 5). We propose to refine this by using boundary detection output that can be gained via a multitask learning approach (Details are in Sec. 3.2)

3.2 Using different modalities [RGB, Depth (HHA), (Boundary)] efficiently

In order to satisfy the latter function, we propose the two approaches below:

1. Fusion-based approach that uses all different modal images as input

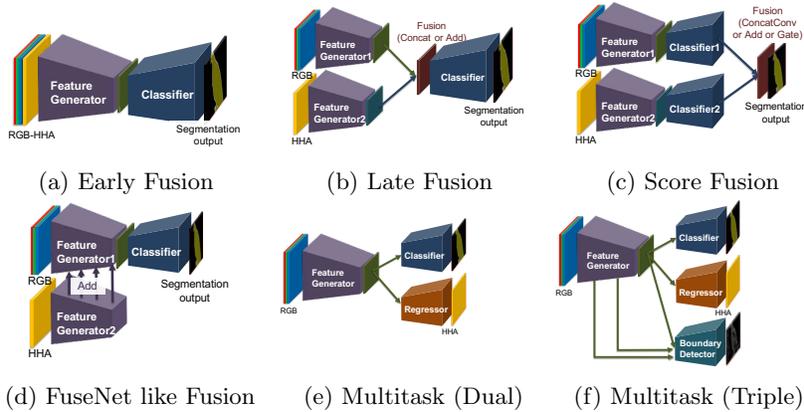


Fig. 4: Fusion-based [(a)-(d)] and multitask learning [(e), (f)] architectures. (e) model that solves the semantic segmentation and depth regression task and (f) model that solves the boundary detection task in addition to the other two tasks. In this figure, one classifier (two classifiers in the score fusion model) exists but actually two classifiers (four classifiers in the score fusion model) exist to utilize MCD as shown in Fig. 2b.

2. Multitask learning approach that uses only RGB as input and the other modals as output.

We will describe these two approaches in detail.

Fusion-based approach If multimodal images are inputs, an appropriate fusing method is known to boost segmentation accuracy. There are many fusion methods [13, 5, 26, 12] but in this research we focus on four comparatively simple fusions that are *early fusion*, *late fusion*, *score fusion*, *fusenet like fusion* (see Fig. 4a - Fig. 4d), because they are not specifically designed and widely used. *Early fusion* just concatenates the RGB and HHA (depth) in inputs. *Late fusion* fuses two encoders in the middle (in this research, the middle means the one layer before the final output). When fusing, we consider two ways of fusing, addition or concatenation. *Score fusion* fuses two encoders in the output. When fusing, we also consider three fusing methods, addition or concatenation + 1×1 convolution or gate fusion [5]. *Fusenet-like fusion* fuses two encoders in certain middle layers [13]. Past research [14, 33] showed that lower layers of a CNN are largely task and category agnostic but domain-specific, while higher layers are largely task and category specific but domain-agnostic. So *late fusion* is considered to be the best out of the four.

To incorporate these fusion-based models into MCD algorithm which we described in Sec. 3.1, we just replace X_{RGB} with X_{RGBHHA} .

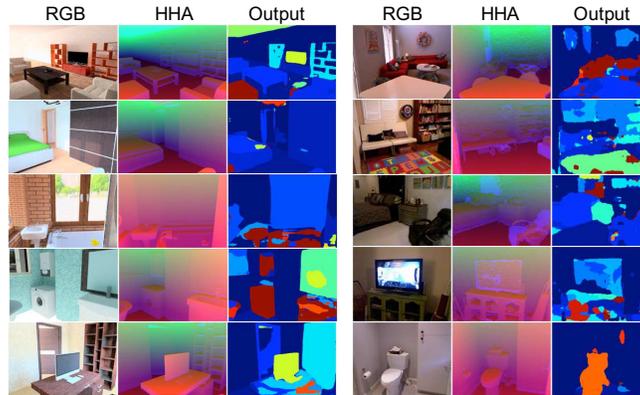


Fig. 5: Sample segmentation results. **Left:** Output of the model that was trained on SUNCG and tested on SUNCG. **Right:** Output of the domain-adapted model from SUNCG to the NYUDv2 train split, and tested on the NYUDv2 test split. The domain adapted model outputs messier results than usual.

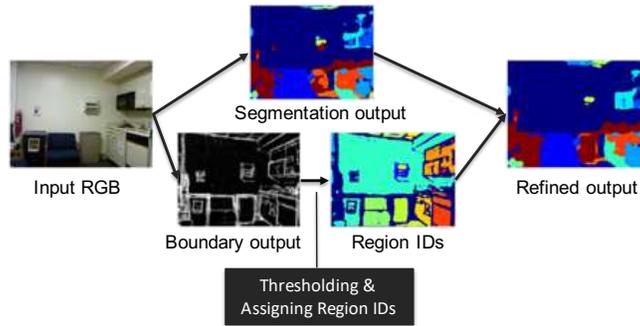


Fig. 6: Refinement of segmentation result using boundary detection output.

Multitask learning approach Multitask learning is a promising approach for efficiently and effectively addressing multiple mutually related recognition tasks and its performance is known to outperform that of single tasks [17]. We solve semantic segmentation and depth regression tasks simultaneously. Also, a general segmentation model is known to get not sharp boundaries [5, 26]. Kendall *et al.* [16] showed that the points on object boundaries have high aleatoric uncertainty. Hence, we add one extra task, a boundary detection task, which can be thought to render feature maps more aware of boundaries. One feature map is used as input for the semantic segmentation and depth estimation task. In addition, two lower feature maps were also used as inputs for the boundary detector following holistically-nested edge detection (HED) [35] as shown in Fig. 4f.

This boundary detection output can be utilized to refine the segmentation output. In fact, based on the segmentation results in Fig. 5, the outputs of the

domain-adapted model are messier than usual. In order to fix this, we propose post-processing the segmentation result based on boundary detection output as shown in Fig. 6. In detail, we first threshold the boundary detection output and then assign IDs to each separated region. Segmentation output (class label) in one region should be unique (one class label for one ID) and so the segmentation output is refined by voting in each separated region. However, boundary detection output is not always perfect. One region sometimes expands to the adjacent region and becomes too large. Therefore, we do not post-process a region whose area is bigger than the maximum-threshold (set to one-third of the image size). In addition, points exactly on the boundaries are not post-processed.

To incorporate these multitask learning models into MCD algorithm which we described in Sec. 3.1, we replace the semantic segmentation loss of *Step A* with a total multitask loss. When we compute the total loss, tuning the weight of each task is important. We adopted Kendall’s algorithm to automatically tune the weight [17] by introducing trainable homoscedastic uncertainty parameter σ_i , where i denotes the task index (1: semantic segmentation, 2: depth regression, 3: boundary detection). This total loss is computed as follows;

$$\begin{aligned} & \mathcal{L}_{\text{multitask}}(X_{\text{RGB}}^s, X_{\text{RGB}}^t, Y_1^s, Y_2^s, Y_2^t, Y_3^s) \\ &= \sum_{i \in \{1, 2, (3)\}} \left(\frac{1}{2\sigma_i^2} \mathcal{L}_i^s(X_{\text{RGB}}^s, Y_i^s) + \log \sigma_i^2 \right) + \mathcal{L}_2^t(X_{\text{RGB}}^t, Y_2^t) \end{aligned}$$

where $\mathcal{L}_1 (= \mathcal{L}_{\text{seg}})$, \mathcal{L}_2 , \mathcal{L}_3 denotes the cross entropy loss for semantic segmentation, the mean squared loss for depth (HHA) regression loss, the class-balanced cross entropy loss for boundary detection, respectively. When \mathcal{L}_3 is not used, the model corresponds to Fig. 4e and, otherwise it corresponds to Fig. 4f. Note that only depth regression loss (\mathcal{L}_2) is computed on both source and target samples but segmentation and boundary detection losses are computed only on source samples because we hypothesize that semantic labels and instance boundaries only exist in source samples. $\mathcal{L}_1 (= \mathcal{L}_{\text{seg}})$ is computed as Eq. 2 and \mathcal{L}_2 , \mathcal{L}_3 are computed as follows;

$$\mathcal{L}_2 = \|Y_2^t - f(X_{\text{RGB}})\|^2 \quad (5)$$

$$\mathcal{L}_3^s = - \frac{|Y_{3-}^s|}{|Y_3^s|} \sum_{j \in Y_{3+}^s} \log p(y_{3j}^s = 1 | X_{\text{RGB}}) - \frac{|Y_{3+}^s|}{|Y_3^s|} \sum_{j \in Y_{3-}^s} \log p(y_{3j}^s = 0 | X_{\text{RGB}}) \quad (6)$$

where f transforms input X_{RGB} to depth (HHA) regression output, and $|Y_{3-}|$ and $|Y_{3+}|$ denote the edge and non-edge ground truth label sets, respectively. (Note that $Y_3^s \in \{0, 1\}$ and $|Y_3| = |Y_{3+}| + |Y_{3-}|$.) $\log p(y_{3j}^s = 1 | X_3)$ and $\log p(y_{3j}^s = 0 | X_3)$ denotes the sigmoid output of predicted boundary detection on the edge and non-edge points, respectively.

4 Experiment

4.1 Setting

Implementation detail³: We use a dilated residual network (*drn_d_38*) [36] which is pre-trained on ImageNet [8], which was shown to perform well in [30]. We followed the public implementation⁴ and adopted MCD [30] as an unsupervised domain adaptation method because it had good performance on domain adaptation problems from synthetic GTA [29] to real CityScapes [6]. Then, we separated *drn_d_38* into a feature generator and a classifier (actually two classifiers) and trained them in an adversarial manner. In fusion-based models, the last transposed convolution layer was used as a classifier and all lower layers were used as feature generators. In multitask learning models, the feature generator is the same as fusion-based models but the classifier is composed of a bilinear up-sampling layer that enlarges the feature map eight times and three convolution layers following [17]. We used Momentum SGD to optimize our model and set the momentum rate to 0.9 and the learning rate to 1.0×10^{-3} in all experiments. The image size was resized to 640×480 and no data augmentation methods were used. We set one epoch to consist of 5000 iterations and chose test epoch numbers based on the entropy criteria following [25].

Dataset: We used the publicly available synthetic dataset SUNCG [38] as the source domain dataset and the real dataset NYUDv2 [31] as the target domain dataset. SUNCG contains two types of RGB images, an OpenGL-based and physically-based color image. We would like to use more realistic data and therefore used the latter type. 568,793 RGB + HHA + instance boundary (only for multitask:triple) images of SUNCG and 795 RGB + HHA images of NYUDv2 train set were used for training and the NYUDv2 test set that contains 654 images was used for evaluation. During training, we randomly sampled just a single sample (setting the batch size to 1 because of the GPU memory limit) from both the images (and their labels) of the source dataset and the remaining images of the target dataset yet with no labels. Removing 6 classes (*books, paper, towel, box, person, bag*) that do not exist in SUNCG from the NYUDv2 40 class, 34 common classes were considered. According to the author of SUNCG, they removed *person* and *plant* in the rendered data because these two types of objects can be hardly rendered photo-realistic. Fig. 7 shows the class label distribution of SUNCG and NYUDv2. The imbalanced distribution demands the application of the four evaluation metrics.

Evaluation Metrics: We report on four metrics from common semantic segmentation and scene parsing evaluations. They are pixel accuracy (pixAcc), mean accuracy (mAcc), mean intersection over union (mIoU), and frequency weighted intersection over union (fwIoU). Let k be the number of classes, n_{ij} be the number of pixels of class i predicted to belong to class j , t_i be the total number of pixels of class i in ground truth segmentation. We compute:

³ our code: <https://github.com/LittleWat/multichannel-semseg-with-uda>

⁴ <https://github.com/fyu/drn>

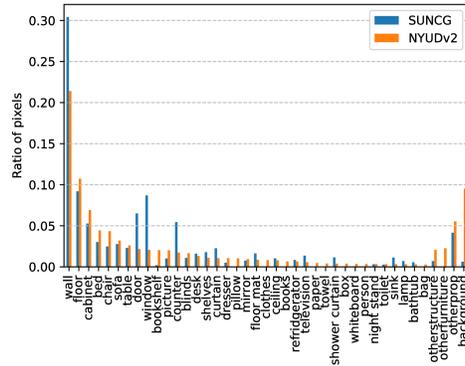


Fig. 7: Distribution of SUNCG [38] and NYUDv2 [31].

$$\text{pixAcc} = \frac{\sum_i n_{ii}}{\sum_i t_i} \left(= \frac{\sum_i n_{ii}}{\sum_j \sum_i n_{ij}} \right), \text{mAcc} = \frac{1}{k} \sum_i \frac{n_{ii}}{t_i}, \text{mIoU} = \frac{1}{k} \sum_i \frac{n_{ii}}{\sum_j (n_{ij} + n_{ji}) - n_{ii}},$$

$$\text{fwIoU} = \frac{1}{\sum_k t_k} \sum_i \frac{t_i n_{ii}}{\sum_j (n_{ij} + n_{ji}) - n_{ii}}.$$

4.2 Results

Fig. 8 and Table 1 show the qualitative and quantitative results, respectively. We can confirm the effect of domain adaptation. *Adapt (Multitask:Triple+Refine)* was the best according to three evaluation metrics and outperformed *Adapt (RGB)* and *Adapt (EarlyFusion)* in all the evaluation metrics. The post-process which refines the segmentation results using the boundary detection outputs could improve all the metrics and lead better qualitative results. In the fusion-based models, *Adapt (LateFusion:Add)* could outperform *Adapt (RGB)* in three evaluation metrics but the performance of most of the fusion-based models is not that different nor worse than that of *Adapt (RGB)*. This is due to the fact that the visibility of RGB images was better than that of HHA images for almost all the classes in the dataset in addition to the fact that objects which exist far away from the camera cannot be seen in HHA images due to the depth sensor range limitation, which could have a negative effect in fusion-based approaches.

IoUs (see Table 3) RGB v.s. HHA The classes whose IoU of RGB outperformed those of HHA were *Ceiling* and *Floor*, whose visibility is better in HHA than in RGB. Conversely, the classes whose IoU of HHA outperformed those of RGB were *Window*, *Blinds* or *Television*. This result is reasonable from the perspective of such classes as shown in: Fig. 9.

Fusion-based v.s. Multitask learning Multitask learning approaches, especially the *Adapt (Triple+Refine)* model, outperformed fusion-based approaches significantly in classes such as *Bed*, *Picture*, *Toilet* except for *Ceiling*. This indicates that multitask learning approaches work well for *object* classes while fusion-based approaches work well for *region* classes.

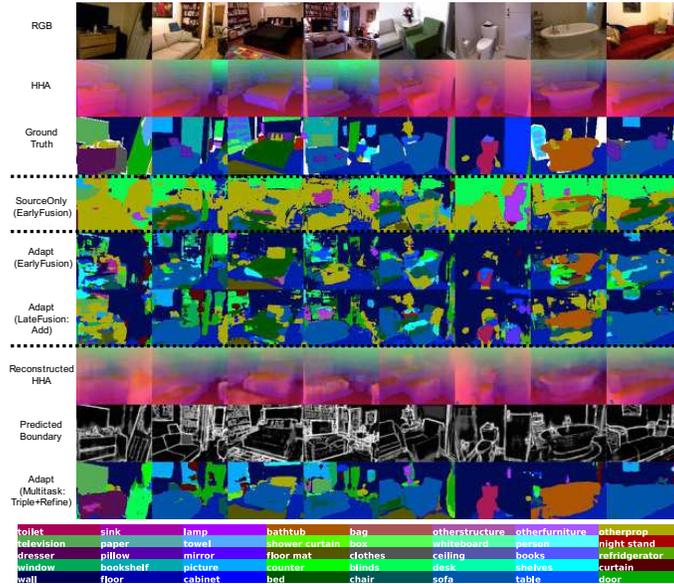


Fig. 8: Qualitative results. The three bottom rows show the results of *Adapt* (*Multitask: Triple+Refine*).

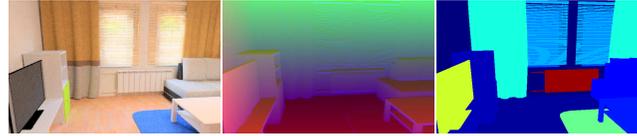


Fig. 9: Sample (RGB, HHA, GT from left to right) of SUNCG [38]. *Window*, *Blinds* or *Television* looks clear in the RGB image. However, *Floor* looks clear in the HHA image.

Boundary detection result Following HED [35], we computed three evaluation metrics, fixed contour threshold (ODS), per-image best threshold (OIS) and average precision (AP) using public code⁵. We compared with handcrafted edge detection methods (Sobel, Canny, and Laplacian, whose hyper-parameter was set to default of OpenCV) that do not use ground truth. As shown in Table 2, the boundary detection output outperformed these handcrafted methods, but the adaptation was not so effective. Based on this, boundary detection is considered to be a domain agnostic task compared to semantic segmentation.

Failures From Fig. 7 and Table 3, the IoUs of the rare classes such as *BookShelf*, *Pillow*, *Mirror*, *Clothes*, *BookShel*, *Shower-curtain*, *Whiteboard*, *Other-furniture* were zero or almost zero. Rare classes in source samples seem to be difficult to

⁵ <https://github.com/pdollar/edges>

Table 1: Four evaluation metrics [%] of the the domain adaptation results from SUNCG [38] to NYUDv2 [31]. (Oracle is the result of the model trained on the train split of NYUDv2)

	pixAcc	mAcc	fwIoU	mIoU
Oracle (Target Only)	60.7	38.7	45.7	28.0
Source Only (RGB)	13.0	6.7	9.9	3.2
Source Only (HHA)	15.6	9.7	9.0	3.8
Source Only (EarlyFusion)	17.9	9.9	10.0	4.2
Adapt (RGB)	42.3	19.1	27.2	11.4
Adapt (HHA)	40.5	13.4	22.7	8.6
Adapt (EarlyFusion)	43.6	17.1	27.6	10.7
Adapt (LateFusion:Add)	41.8	19.7	28.6	12.2
Adapt (LateFusion:Concat)	40.9	17.1	25.8	10.6
Adapt (ScoreFusion:Add)	39.7	19.0	27.6	10.7
Adapt (ScoreFusion:ConcatConv)	38.8	17.0	26.2	10.7
Adapt (ScoreFusion:Gate)	37.3	14.8	24.1	9.2
Adapt (FusenetFusion)	29.5	14.0	19.5	6.9
Adapt (Multitask:Dual)	44.0	20.2	30.1	12.8
Adapt (Multitask:Triple)	42.6	22.6	30.0	13.1
Adapt (Multitask:Triple+Refine)	43.7	22.8	30.6	13.2

recognize. Fig. 10 shows the confusion matrix of the best model, *Adapt (Multitask:Triple+Refine)*. *Floormat* is mispredicted as *Floor*, *Whiteboard* is mispredicted as *Wall*, *Picture* and *Blinds* are mispredicted as *Window*, *Pillow* is mispredicted as *Sofa*. If we consider the source label distribution shown in Fig. 7, a rare class is often mispredicted as a common class whose position is the same as the rare one. The ratio of pixels that are mispredicted as *Othrprop* was high. This is probably because *Othrprop* contains various kinds of classes such as *car*, *motorcycle*, *soccer goal post*, *gas stove* as shown in Fig. 11.

5 Conclusion

We combined a multichannel semantic segmentation task with an unsupervised domain adaptation task and proposed two architectures (fusion-based and multitask learning). We demonstrated that the multitask learning approach outperforms the simple *early fusion* approach in all the evaluation metrics. In addition, we propose adding a boundary detection task in the multitask learning approach and using the detection result to refine the segmentation output. We qualitatively and quantitatively show this post-process is effective especially in the classes whose boundaries look clear. However, the scores of the adaptation result are still poor when compared to *oracle*. In future work, we would like to use a few labeled and many unlabeled target samples (semi-supervised setting) and improve the results.

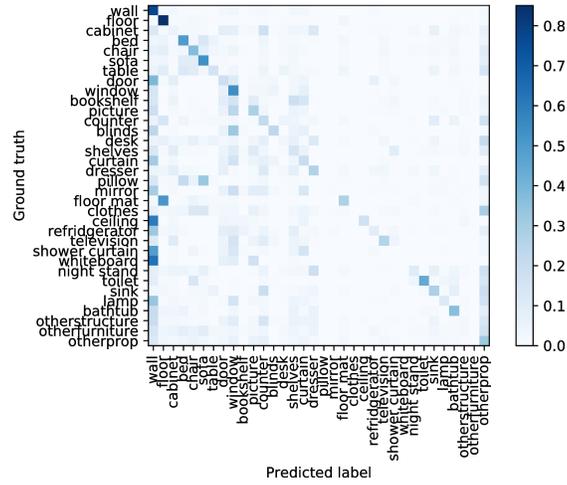


Fig. 10: Confusion matrix of the result of *Adapt (Multitask:Triple+Refined)*. Total value of each row is 1.0.



Fig. 11: Sample of *Otherprop*.

Method	ODS	OIS	AP
Sobel	0.446	0.473	0.190
Canny	0.322	0.322	0.000
Laplacian	0.450	0.467	0.262
Source Only (Multitask:Triple)	0.542	0.557	0.451
Adapt (Multitask:Triple)	0.567	0.591	0.445

Table 2: Result of boundary detection

6 Acknowledgements

The work was partially funded by the ImPACT Program of the Council for Science, Technology, and Innovation (Cabinet Office, Government of Japan).

Table 3: IoUs [%] of the the domain adaptation results from SUNCG [38] to NYUDv2 [31]. (Oracle is the result of the model using train split of NYUDv2.)

	Wall	Floor	Cabinet	Bed	Chair	Sofa	Table	Door	Window	BookShelf	Picture	Counter	Blinds	Desks	Shelves	Curtain	Dresser
Oracle (Target Only)	66.9	78.8	43.3	53.6	39.7	38.6	25.8	16.3	33.9	32.1	43.6	38.9	42.5	9.2	7.0	24.0	19.0
Source Only (RGB)	27.0	19.5	1.4	0.2	1.7	2.3	3.9	8.5	2.8	0.1	2.9	0.6	2.2	0.2	0.2	0.6	2.2
Source Only (HHA)	9.0	30.7	1.7	11.5	13.3	8.7	4.2	0.2	0.7	0.0	0.0	0.6	3.9	0.0	1.2	0.4	0.9
Source Only (EarlyFusion)	7.5	46.8	0.2	11.0	3.0	8.1	6.1	3.6	1.5	0.0	3.6	2.4	5.9	0.4	0.6	1.0	0.7
Adapt (RGB)	54.4	56.8	5.8	12.7	27.8	28.8	10.6	8.7	13.3	0.5	9.0	4.1	26.3	4.1	2.0	9.4	8.8
Adapt (HHA)	36.6	75.9	1.9	16.6	20.0	21.3	17.7	3.9	0.6	0.3	0.1	4.3	1.0	4.6	2.3	1.3	1.7
Adapt (EarlyFusion)	50.1	75.8	2.6	19.4	24.7	19.8	21.3	6.1	8.3	0.4	6.0	3.5	8.9	2.9	4.0	5.1	2.7
Adapt (LateFusion:Add)	50.4	74.3	6.7	15.1	30.6	31.0	17.7	7.9	8.1	0.5	5.5	7.3	14.9	3.4	1.8	5.4	3.9
Adapt (LateFusion:Concat)	51.5	57.1	4.7	10.8	28.0	29.2	12.6	9.2	7.9	0.2	9.3	2.2	17.4	3.2	2.0	3.1	5.6
Adapt (ScoreFusion:Add)	50.9	74.3	6.7	8.1	25.3	25.5	20.7	8.8	7.5	0.6	9.9	4.8	18.7	4.6	2.6	4.9	1.9
Adapt (ScoreFusion:ConcatConv)	46.5	69.5	5.5	24.1	20.2	18.6	12.4	9.1	13.5	0.1	7.6	6.8	13.7	3.6	1.3	11.4	2.2
Adapt (ScoreFusion:Gate)	41.8	75.3	2.7	5.6	18.9	22.1	17.0	4.3	18.3	0.1	2.2	4.8	2.1	2.9	2.1	4.9	3.0
Adapt (FusenetFusion)	28.9	70.5	5.0	6.9	17.8	19.9	12.9	6.7	6.0	0.0	0.0	2.5	0.0	4.2	3.3	1.5	0.0
Adapt (Multitask:Dual)	55.5	67.6	7.8	27.0	27.0	30.6	16.6	9.2	18.2	0.0	16.0	7.4	15.1	2.7	3.0	10.6	2.6
Adapt (Multitask:Triple)	55.3	67.6	4.2	33.0	25.9	31.0	12.0	8.1	15.9	0.4	16.1	7.7	17.3	2.1	3.5	6.5	10.3
Adapt (Multitask:Triple+Refine)	56.7	68.2	4.2	35.0	26.0	32.6	12.2	8.5	16.3	0.4	16.3	7.7	18.4	1.9	3.4	6.8	10.8

	Pillow	Mirror	Floor-mat	Clothes	Ceiling	Refrigerator	Television	Shower-curtain	Whiteboard	NightStand	Toilet	Sink	Lamp	Bathtub	Other-structure	Other-furniture	Other-prop
Oracle (Target Only)	27.2	10.9	17.9	14.2	58.0	9.2	30.1	6.5	0.1	2.1	51.1	28.3	23.5	18.5	9.0	6.1	27.4
Source Only (RGB)	0.0	0.0	1.7	0.0	4.9	0.0	12.1	0.0	0.0	1.3	0.3	0.0	10.5	0.0	0.0	0.0	0.4
Source Only (HHA)	4.5	0.0	0.0	0.0	14.7	0.1	0.0	0.0	0.0	1.5	0.0	2.8	3.8	2.6	0.0	0.0	13.8
Source Only (EarlyFusion)	0.0	0.0	7.6	0.0	6.1	0.0	2.8	0.0	0.0	0.3	0.8	2.8	3.0	5.0	0.3	0.0	10.4
Adapt (RGB)	0.0	0.5	7.6	0.0	18.0	2.7	13.7	0.7	0.0	1.7	9.6	9.1	17.2	5.8	2.0	0.0	17.1
Adapt (HHA)	0.0	0.0	0.4	0.0	42.6	0.1	1.1	0.9	0.0	3.2	4.1	5.6	10.6	5.6	0.1	0.0	6.8
Adapt (EarlyFusion)	0.0	0.1	12.3	0.0	29.4	1.3	6.0	1.3	0.0	1.0	9.1	5.9	10.9	6.8	1.5	0.0	15.3
Adapt (LateFusion:Add)	0.0	2.0	6.0	0.0	45.8	4.3	11.5	2.3	0.0	3.2	16.6	6.0	8.9	10.1	0.6	0.0	14.5
Adapt (LateFusion:Concat)	0.0	1.0	8.1	0.0	44.2	1.4	1.3	0.5	0.0	4.6	8.5	9.4	6.7	7.1	0.8	0.0	11.7
Adapt (ScoreFusion:Add)	0.0	0.6	3.0	0.0	27.3	2.4	7.6	0.9	0.0	4.3	10.0	4.0	5.1	9.9	2.0	0.0	12.0
Adapt (ScoreFusion:ConcatConv)	0.0	1.0	2.7	0.0	30.4	0.9	18.2	0.1	0.0	0.4	5.3	10.0	8.1	5.6	0.3	0.0	13.9
Adapt (ScoreFusion:Gate)	0.0	0.4	2.0	0.0	39.0	0.6	10.9	1.2	0.0	2.0	4.1	1.5	8.3	2.6	0.2	0.0	11.4
Adapt (FusenetFusion)	0.0	0.0	1.4	0.0	38.1	0.0	4.6	0.0	0.0	0.0	0.0	0.0	0.2	0.0	0.0	0.0	4.2
Adapt (Multitask:Dual)	0.0	0.8	7.8	0.0	36.5	2.7	12.7	1.6	0.0	2.6	11.5	3.1	12.0	10.2	1.7	0.0	16.6
Adapt (Multitask:Triple)	0.0	1.4	13.9	0.0	20.4	4.7	10.7	0.0	0.0	4.5	30.0	6.2	6.7	10.0	1.7	0.0	17.6
Adapt (Multitask:Triple+Refine)	0.0	0.4	14.7	0.0	15.8	4.7	11.2	0.0	0.0	5.5	30.2	6.2	5.7	11.2	1.6	0.0	17.3

References

1. Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2d-3d-semantic data for indoor scene understanding. arXiv:1702.01105 (2017)
2. Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S.: 3d semantic parsing of large-scale indoor spaces. In: CVPR (2016)
3. Bousmalis, K., Silberman, N., Dohan, D., Erhan, D., Krishnan, D.: Unsupervised pixel-level domain adaptation with generative adversarial networks. In: CVPR (2017)
4. Chen, Y.H., Chen, W.Y., Chen, Y.T., Tsai, B.C., Wang, Y.C.F., Sun, M.: No more discrimination: Cross city adaptation of road scene segmenters. In: ICCV (2017)
5. Cheng, Y., Cai, R., Li, Z., Zhao, X., Huang, K.: Locality-sensitive deconvolution networks with gated fusion for rgb-d indoor semantic segmentation. In: CVPR (2017)
6. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
7. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: CVPR (2017)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
9. Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. In: ICML (2014)
10. Gupta, S., Arbelaez, P., Malik, J.: Perceptual organization and recognition of indoor scenes from rgb-d images. In: CVPR (2013)
11. Gupta, S., Girshick, R., Arbeláez, P., Malik, J.: Learning rich features from rgb-d images for object detection and segmentation. In: ECCV (2014)
12. Ha, Q., Watanabe, K., Karasawa, T., Ushiku, Y., Harada, T.: Mfnet: Towards real-time semantic segmentation for autonomous vehicles with multi-spectral scenes. In: IROS (2017)
13. Hazirbas, C., Ma, L., Domokos, C., Cremers, D.: Fusetnet: Incorporating depth into semantic segmentation via fusion-based cnn architecture. In: ACCV (2016)
14. Hoffman, J., Gupta, S., Leong, J., Guadarrama, S., Darrell, T.: Cross-modal adaptation for rgb-d detection. In: ICRA (2016)
15. Hoffman, J., Wang, D., Yu, F., Darrell, T.: Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. arXiv:1612.02649 (2016)
16. Kendall, A., Gal, Y.: What uncertainties do we need in bayesian deep learning for computer vision? In: NIPS (2017)
17. Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. CVPR (2018)
18. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS (2012)
19. Kuga, R., Kanezaki, A., Samejima, M., Sugano, Y., Matsushita, Y.: Multi-task learning using multi-modal encoder-decoder networks with shared skip connections. In: ICCV Workshop (2017)
20. Lin, D., Chen, G., Cohen-Or, D., Heng, P.A., Huang, H.: Cascaded feature network for semantic segmentation of rgb-d images. In: ICCV (2017)
21. Lin, G., Milan, A., Shen, C., Reid, I.: Refinenet: Multi-path refinement networks with identity mappings for high-resolution semantic segmentation. CVPR (2017)

22. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: CVPR (2015)
23. Long, M., Cao, Y., Wang, J., Jordan, M.I.: Learning transferable features with deep adaptation networks. In: ICML (2015)
24. McCormac, J., Handa, A., Leutenegger, S., J.Davison, A.: Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? In: ICCV (2017)
25. Morerio, P., Cavazza, J., Murino, V.: Minimal-entropy correlation alignment for unsupervised deep domain adaptation. ICLR (2018)
26. Park, S.J., Hong, K.S., Lee, S.: Rdfnet: Rgb-d multi-level residual feature fusion for indoor semantic segmentation. In: ICCV (2017)
27. Qi, X., Liao, R., Jia, J., Fidler, S., Urtasun, R.: 3d graph neural networks for rgb-d semantic segmentation. In: ICCV (2017)
28. Qiu, W., Zhong, F., Zhang, Y., Qiao, S., Xiao, Z., Kim, T.S., Wang, Y., Yuille, A.: Unrealcv: Virtual worlds for computer vision. ACM MM Open Source Software Competition (2017)
29. Richter, S.R., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: ECCV (2016)
30. Saito, K., Watanabe, K., Ushiku, Y., Harada, T.: Maximum classifier discrepancy for unsupervised domain adaptation. CVPR (2018)
31. Silberman, N., Hoiem, D., Kohli, P., Fergus, R.: Nyu depth dataset v2. In: ECCV (2012)
32. Song, S., Lichtenberg, S.P., Xiao, J.: Sun rgb-d: A rgb-d scene understanding benchmark suite. In: CVPR (2015)
33. Song, X., Herranz, L., Jiang, S.: Depth cnns for rgb-d scene recognition: Learning from scratch better than transferring from rgb-cnns. In: AAAI (2017)
34. Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial discriminative domain adaptation. In: CVPR (2017)
35. Xie, S., Tu, Z.: Holistically-nested edge detection. In: ICCV (2015)
36. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: CVPR (2017)
37. Zhang, Y., David, P., Gong, B.: Curriculum domain adaptation for semantic segmentation of urban scenes. In: ICCV (2017)
38. Zhang, Y., Song, S., Yumer, E., Savva, M., Lee, J.Y., Jin, H., Funkhouser, T.: Physically-based rendering for indoor scene understanding using convolutional neural networks. In: CVPR (2017)