

Transferable Semi-Supervised 3D Object Detection From RGB-D Data

Yew Siang Tang Gim Hee Lee

Department of Computer Science, National University of Singapore

{yewsiang.tang, gimhee.lee}@comp.nus.edu.sg

Abstract

We investigate the direction of training a 3D object detector for new object classes from only 2D bounding box labels of these new classes, while simultaneously transferring information from 3D bounding box labels of the existing classes. To this end, we propose a transferable semi-supervised 3D object detection model that learns a 3D object detector network from training data with two disjoint sets of object classes - a set of strong classes with both 2D and 3D box labels, and another set of weak classes with only 2D box labels. In particular, we suggest a relaxed re-projection loss, box prior loss and a Box-to-Point Cloud Fit network that allow us to effectively transfer useful 3D information from the strong classes to the weak classes during training, and consequently, enable the network to detect 3D objects in the weak classes during inference. Experimental results show that our proposed algorithm outperforms baseline approaches and achieves promising results compared to fully-supervised approaches on the SUN-RGBD and KITTI datasets. Furthermore, we show that our Box-to-Point Cloud Fit network improves performances of the fully-supervised approaches on both datasets.

1. Introduction

3D object detection refers to the problem of classification and estimation of a tight oriented 3D bounding box for each object in a 3D scene represented by a point cloud. It plays an essential role in many real-world applications such as home robotics, augmented reality and autonomous driving, where machines must actively interact and engage with its surroundings. In the recent years, significant advances in 3D object detection have been achieved with fully supervised deep learning-based approaches [4, 5, 6, 12, 16, 20, 31, 35, 40]. However, these strongly supervised approaches rely on 3D ground truth datasets [7, 25, 32, 33] that are tedious and time consuming to label even with good customised 3D annotation tools such as those in [7, 25]. In contrast, calculations based on [17] show that it can be ~ 3 -16 times faster (depending on the annotation tool used) to label 2D than 3D bounding boxes (details in supplementary). This

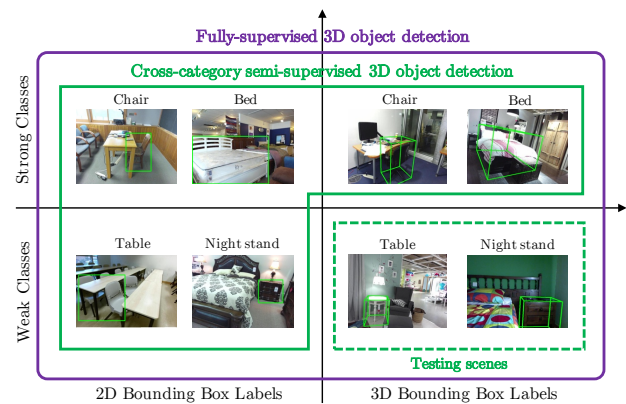


Figure 1. Illustration of fully-supervised 3D object detection where 3D box labels of all classes (*weak* and *strong* classes) are available, and cross-category semi-supervised 3D object detection where 3D box labels of inference classes (*weak* classes e.g. Table and Night stand) are not available. Conventional semi-supervised learning (which is in-category) requires strong annotations from the inference classes (3D box labels in the case of 3D object detection) while cross-category semi-supervised learning requires only the weak annotations from the inference classes (2D box labels in the case of 3D object detection).

observation inspires us to investigate the direction of training a 3D object detector for new object classes from only 2D bounding box labels of these new classes, while simultaneously transferring information from 3D bounding box labels of the existing classes. We hope that this direction of research leads to cost reductions, and improves the efficiency and practicality of learning 3D object detectors in new applications with new classes.

More specifically, the objective of this paper is to train a 3D object detector network with data from two disjoint sets of object classes - a set of classes with both 2D and 3D box labels (*strong* classes), and another set of classes with only 2D box labels (*weak* classes). In other words, our goal is to transfer useful 3D information from the *strong* classes to the *weak* classes (the new classes of a new 3D object detection application correspond to these *weak* classes). We refer to this as cross-category semi-supervised 3D object detection as illustrated in Fig. 1.

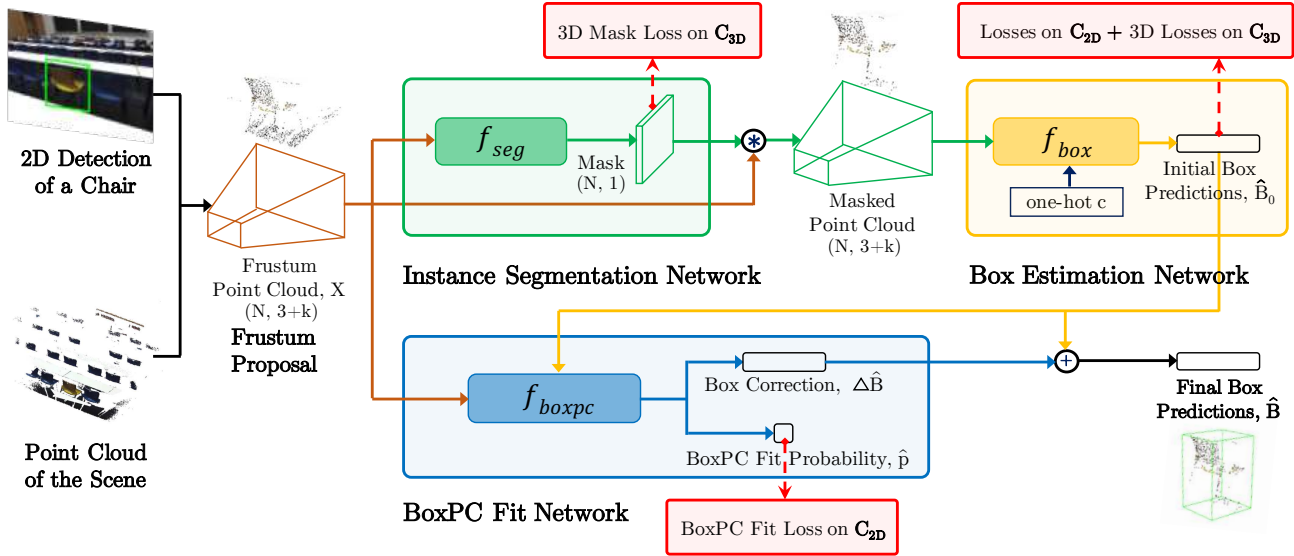


Figure 2. Overall architecture with a RGB-D (a pair of image and point cloud) input. (1) Frustum point clouds X are extracted from the input point cloud and 2D object detection boxes on the image. (2) f_{seg} takes X as input and outputs class-agnostic segmentations used to mask X . (3) f_{box} predicts an initial 3D box \hat{B}_0 with the masked point cloud. (4) The pretrained f_{boxpc} model refines \hat{B}_0 to \hat{B} according to X , and predicts the BoxPC fit probability used to supervise f_{box} .

To this end, we propose a novel transferable semi-supervised 3D object detection network. Our network leverages on the state-of-the-art Frustum PointNets [20] as the backbone. We train the backbone network to make class-agnostic segmentations and class-conditioned initial 3D box predictions on the *strong* classes in a fully-supervised manner. We also supervise the network with 2D box labels of the *weak* classes using a relaxed reprojection loss which corrects 3D box predictions that violate boundaries specified by the 2D box labels and through the prior knowledge of the object sizes. To transfer knowledge from 3D box labels from the *strong* classes, we first train a Box-to-Point Cloud (BoxPC) Fit network on 3D box and point cloud pairs to reason about 3D boxes and the fit with their corresponding point clouds. More specifically, we proposed an effective and differentiable method to encode a 3D box and point cloud pair. Finally, the 3D box predictions on the *weak* classes are supervised and refined by the BoxPC Fit network.

The contributions of our paper are as follows:

- We propose a network to perform 3D object detection on *weak* classes, where *only* 2D box labels are available. This is achieved using relaxed reprojection and box prior losses on the *weak* classes, and transferred knowledge from the *strong* classes.
- A differentiable BoxPC Fit network is designed to effectively combine BoxPC representation. It is able to supervise and improve the 3D object detector on the

weak classes after training on the *strong* classes.

- Our transferable semi-supervised 3D object detection model outperforms baseline approaches and achieves promising results compared to fully-supervised methods on both SUN-RGBD and KITTI datasets. Additionally, we show that our BoxPC Fit network can be used to improve the performance of fully-supervised 3D object detectors.

2. Related Work

3D Object Detection 3D object detection approaches have advanced significantly in the recent years [4, 5, 6, 12, 13, 16, 20, 23, 24, 31, 35, 40]. However, most approaches are fully-supervised and highly-dependent on 3D box labels that are difficult to obtain. To the best of our knowledge, there are no existing weakly- or semi-supervised 3D object detection approaches. Recently, [28] proposed a self-supervised Augmented Autoencoder that trains on CAD models, removing the need for pose-annotated training data. However, it is unclear if the method generalizes to general 3D object detection datasets with high intra-class shape variations and noisy depth data.

Weakly- and Semi-Supervised Learning There is growing interest in weakly- and semi-supervised learning in many problem areas [1, 2, 3, 27, 29, 37, 38, 39] because it is tedious and labor intensive to label large amounts of data for fully supervised deep learning. There is a wide literature of weakly-supervised [11, 19, 30, 41] and semi-

supervised [8, 15, 18, 36] learning approaches for semantic segmentation. Both strong and weak labels on the inference classes are required in conventional semi-supervised learning. Hence, the approach remains expensive for applications with new classes. [9, 10, 34] proposed the cross-category semi-supervised semantic segmentation which is a more general approach of semi-supervised segmentation, where the model is able to learn from strong labels provided for classes outside of the inference classes. They outperformed weakly- and semi-supervised methods, and showed on-par performances with fully-supervised methods. Inspired by the effectiveness of the transferred knowledge, we propose to tackle the same problem in the 3D object detection domain.

3. Problem Formulation

Let $C = C_{2D} \cup C_{3D}$, where C_{2D} is the set of classes with only 2D box labels (*weak* classes) and C_{3D} is the disjoint set of classes with 2D and 3D box labels (*strong* classes). We tackle the cross-category semi-supervised 3D object detection problem where 3D object detection is performed on the object classes in C_{2D} while training is done on strong labels from C_{3D} and weak labels from C_{2D} . In 3D object detection with RGB-D data, our goal is to classify and predict *amodal* 3D bounding boxes for objects in the 3D space. Depth data can be provided by various depth sensors e.g. RGB-D sensors and LiDAR etc. Each 3D bounding box is parameterized by its center (b_x, b_y, b_z) , size (b_h, b_w, b_l) and orientation b_θ along the vertical axis.

4. Method

4.1. Frustum PointNets

In this section, we briefly discuss Frustum PointNets (FPN), the 3D object detection framework which we adopt as the backbone of our framework. Refer to [20] and our supplementary for more details.

4.1.1 Frustum Proposal

As seen in Fig. 2, the inputs to the network is an image with a 2D box for each object and a 3D point cloud. During inference, we obtain the 2D boxes from a 2D object detector trained on C . We project the point cloud onto the image plane using the camera projection matrix and select only the points that lie in the 2D box for each object. This reduces the search space to only points within a 3D frustum which we refer to as frustum point cloud. Next, variance of the points in each frustum point cloud is reduced by rotating the frustum about the vertical axis of the camera to face the front. Formally, let us denote a frustum point cloud by $X = [x_1, x_2, \dots, x_N]^T \in \mathbb{R}^{N \times (3+k)}$, where N is the number of points. $x_n \in \mathbb{R}^{3+k}$ (i.e. x, y, z, RGB or reflectance) is the n -th point in the point cloud.




Figure 3. Illustration of 3D bounding box and point cloud pairs with a good BoxPC fit (center) and a bad BoxPC fit (right). The ground truth 3D bounding box of the chair is given on the left.

4.1.2 3D Instance Segmentation

The frustum point cloud contains foreground points of an object which we are interested in and background points from the surroundings. To isolate the foreground object and simplify the subsequent 3D box estimation task, the frustum point cloud is fed into a 3D instance segmentation network f_{seg} which predicts a class-agnostic foreground probability for each point. These probabilities are used to mask and retrieve the points with high values, thus giving a masked point cloud. f_{seg} is trained by minimizing

$$L_{seg}^{3D} = \sum_{c \in C_{3D}} H(f_{seg}^*(X^c), f_{seg}(X^c)) \quad (1)$$

on C_{3D} , where H is the point-wise binary cross-entropy loss and f_{seg}^* is the ground truth mask.

4.1.3 3D Box Estimation

The 3D box estimation network f_{box} predicts an initial 3D box of the object \hat{B}_0 given the masked point cloud and a one hot class vector from the 2D detector. Specifically, we compute the centroid of the masked point cloud and use it to translate the masked point cloud to the origin position to reduce translational variance. Next, the translated point cloud is given to two PointNet-based [21] networks that will predict the initial 3D box \hat{B}_0 . The center (b_x, b_y, b_z) is directly regressed by the network while an anchor-based approach is used to predict the size (b_h, b_w, b_l) and rotation b_θ . There are NS size and NH rotation anchors, and $3NS$ size and NH rotation residuals. Hence, the network has $3 + 4NS + 2NH$ outputs that are trained by minimizing

$$L_{box}^{3D} = \sum_{c \in C_{3D}} L_{c1-reg}^c + L_{c2-reg}^c + L_{r-cls}^c + L_{r-reg}^c + L_{s-cls}^c + L_{s-reg}^c + L_{corner}^c \quad (2)$$

on C_{3D} . L_{c1-reg}^c and L_{c2-reg}^c are regression losses for the box center predictions, L_{r-cls}^c and L_{r-reg}^c are the rotation anchor classification and regression losses, L_{s-cls}^c and L_{s-reg}^c are the size anchor classification and regression losses, and L_{corner}^c is the regression loss for the 8 vertices of the 3D box. Note that each loss term in L_{box}^{3D} is weighted




Figure 4. Different feature representations for a pair of 3D bounding box and point cloud. On the left, the inputs are processed independently and concatenated in the feature space. On the right, the inputs are combined in the input space and processed as a whole.

by w_s , which we omit in Eq. 2 and all subsequent losses for brevity. Until this point, the network is fully-supervised by 3D box labels available only for classes C_{3D} . In the subsequent sections, we describe our contributions to train f_{box} on the *weak* classes in C_{2D} by leveraging on the BoxPC Fit Network (Sec. 4.2) and weak losses (Sec. 4.3).

4.2. Box to Point Cloud (BoxPC) Fit Network

The initial 3D box predictions \hat{B}_0 on classes C_{2D} are not reliable because there are no 3D box labels for these classes. As a result, the 3D box surfaces of the initial predictions \hat{B}_0 for classes C_{2D} are likely to cut the frustum point cloud X at unnatural places, as illustrated on the right of Fig. 3. We utilize our BoxPC Fit Network f_{boxpc} with a novel training method to transfer the knowledge of a good BoxPC fit from the 3D box labels of classes C_{3D} to classes C_{2D} . The input to our f_{boxpc} is a pair of 3D box B and frustum point cloud X . The outputs of f_{boxpc} are the BoxPC Fit probability, i.e. goodness-of-fit

$$\hat{p} = f_{boxpc-cl_s}(X, B) \in [0, 1], \quad (3)$$

between B and the object in X , and the correction

$$\Delta \hat{B} = f_{boxpc-reg}(X, B) \in \mathbb{R}^7, \quad (4)$$

required on B to improve the fit between B and the object in X . $\hat{p} \rightarrow 1$ when B encloses the object in X tightly, i.e. there is a high overlap between B and the 3D box label B^* .

A pretrained f_{boxpc} (see Sec. 4.2.1 for the training procedure) is used to supervise and improve 3D box predictions \hat{B} for classes C_{2D} . Specifically, we train f_{box} to make better initial 3D box predictions \hat{B}_0 by maximizing \hat{p} . By minimizing the loss:

$$L_{fit}^{2D} = \sum_{c \in C_{2D}} -\log(f_{boxpc-cl_s}(X^c, \hat{B}_0^c)), \quad (5)$$

i.e. maximizing \hat{p} , our network f_{box} learns to predict \hat{B}_0 that fit objects in their respective point clouds well. Finally, we obtain the final 3D box prediction \hat{B} by correcting the initial 3D box prediction \hat{B}_0 with $\Delta \hat{B}$, i.e. $\hat{B} = \hat{B}_0 + \Delta \hat{B}$.

4.2.1 Pretraining BoxPC Fit Network

We train the BoxPC Fit network f_{boxpc} on the classes in C_{3D} with the 3D box labels $B^* \in \mathbb{R}^7$ and their corresponding point clouds X . We sample varying degree of perturbations $\delta \in \mathbb{R}^7$ to the 3D box labels and get perturbed 3D boxes $B^* - \delta$. We define 2 sets of perturbations, $\mathbf{P}^+(B^*)$ and $\mathbf{P}^-(B^*)$, where \mathbf{P}^+ and \mathbf{P}^- are sets of small and large perturbations, respectively. Formally, $\mathbf{P}^+(B^*) = \{\delta : \alpha^+ \leq IOU(B^* - \delta, B^*) \leq \beta^+ \mid \{\alpha^+, \beta^+\} \in [0, 1]\}$, where IOU refers to the 3D Intersection-over-Union. \mathbf{P}^- is defined similarly but with α^- and β^- as the bounds. α^+ and β^+ are set such that the 3D boxes perturbed by $\delta \in \mathbf{P}^+$ have high overlaps with B^* or good fit with its point clouds and vice versa for \mathbf{P}^- . The loss function to train f_{boxpc} is given by

$$L_{boxpc}^{3D} = \sum_{c \in C_{3D}} L_{cls}^c + L_{reg}^c, \quad (6)$$

where L_{cls}^c is the classification loss

$$L_{cls} = \begin{cases} H(1, f_{boxpc-cl_s}(X, B^* - \delta)), & \text{if } \delta \in \mathbf{P}^+, \\ H(0, f_{boxpc-cl_s}(X, B^* - \delta)), & \text{if } \delta \in \mathbf{P}^-, \end{cases} \quad (7)$$

to predict the fit of a perturbed 3D box and its frustum point cloud, and L_{reg} is the regression loss

$$L_{reg} = l(\delta, f_{boxpc-reg}(X, B^* - \delta)), \quad (8)$$

to predict the perturbation to correct the perturbed 3D box. H is the binary cross-entropy and $l(y^*, \hat{y})$ is the Smooth L1 loss with y^* as the target for \hat{y} .

4.2.2 BoxPC Feature Representations

A possible way to encode the feature representations for the BoxPC network is illustrated on the left of Fig. 4. The features of the input 3D box and frustum point cloud are learned using several MLPs and PointNets [21], respectively. These features are then concatenated and fed into several layers of MLPs to get the BoxPC fit probability \hat{p} and Box Correction term $\Delta \hat{B}$. However, this neglects the intricate information of how the 3D box surfaces cut the

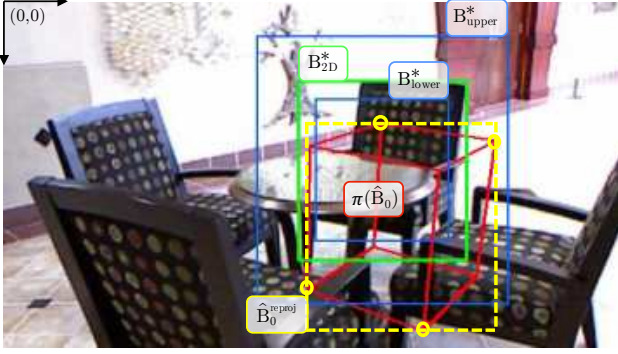


Figure 5. Illustration of the relaxed reprojection loss.

point cloud, leading to poor performance as shown in our ablation studies in Sec. 6.3. Hence, we propose another method to combine the 3D box and point cloud in the input space in a differentiable manner that also exploits the relationship between them. Specifically, we use the box center (b_x, b_y, b_z) to translate the point cloud to the origin to enhance translational invariance to the center positions. Next, we use the box sizes (b_h, b_w, b_l) and rotation b_θ to form six planes representing each of the six box surfaces. We compute a $N \times 6$ feature vector of the perpendicular distance (with direction) from each point to each of the six planes. This feature vector is concatenated with the original frustum point cloud $X \in \mathbb{R}^{N \times (3+k)}$ to give the box-combined frustum point cloud $X' \in \mathbb{R}^{N \times (9+k)}$ as illustrated on the right of Fig. 4. This representation allows the network to easily reason about regions where the box surfaces cut the point cloud since the perpendicular distance will be close to 0. Additionally, it is possible to perform feature learning on the 3D box and the point cloud *jointly* with a single PointNet-based [21] network, improving performance.

4.3. Weak Losses

We supervise the initial 3D box predictions \hat{B}_0 with 2D box labels B_{2D}^* and additional priors.

4.3.1 Relaxed Reprojection Loss

Despite the correlation, the 2D box that encloses all the points of a projected 3D box label does not coincide with the 2D box label. We show in Sec. 6.3 that performance deteriorates if we simply minimize the reprojection loss between the 2D box label and the 2D box enclosing the projection of the predicted 3D box onto the image plane. Instead, we propose a relaxed reprojection loss, where boundaries close to the 2D box labels are not penalized. Let $B_{2D}^* = [b_{left}^*, b_{top}^*, b_{right}^*, b_{bot}^*]$ (green box in Fig. 5) that consists of the left, top, right and bottom image coordinates be the 2D box label with the top-left hand corner as origin. We define an upper bound box $B_{upper}^* = B_{2D}^* + \mathcal{U}$ and a lower bound box $B_{lower}^* = B_{2D}^* + \mathcal{L}$ (outer and in-

ner blue boxes in Fig. 5), where $\mathcal{U} = [u_1, u_2, u_3, u_4]$ and $\mathcal{L} = [v_1, v_2, v_3, v_4]$ are vectors used to adjust the size of B_{2D}^* . Given an initial 3D box prediction \hat{B}_0 , we project it onto the image plane $\pi(\hat{B}_0)$ (red box in Fig. 5) and obtain an enclosing 2D box around the projected points $\hat{B}_0^{reproj} = g(\pi(\hat{B}_0))$ (yellow box in Fig. 5), where $g(\cdot)$ is the function that returns the enclosing 2D box. We penalize \hat{B}_0^{reproj} for violating the bounds specified by B_{upper}^* and B_{lower}^* . An example is shown in Fig. 5, the bounds on the left and right are not violated because the left and right sides of the 2D box \hat{B}_0^{reproj} stays within the bounds specified by B_{upper}^* and B_{lower}^* (blue boxes). However, the top and bottom of \hat{B}_0^{reproj} beyond the bounds are penalized. More formally, the relaxed reprojection loss is given by

$$L_{reproj}^{2D} = \sum_{c \in C_{2D}} \sum_{i=1}^4 \tilde{l}(v_i, u_i, \hat{B}_0^{c, reproj}(i)), \quad (9)$$

$$\tilde{l}(y_{lower}^*, y_{upper}^*, \hat{y}) = \begin{cases} l(y_{upper}^*, \hat{y}), & \text{if } \hat{y} > y_{upper}^*, \\ l(y_{lower}^*, \hat{y}), & \text{if } \hat{y} < y_{lower}^*, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

$\tilde{l}(y_{lower}^*, y_{upper}^*, \hat{y})$ is a relaxed Smooth L1 loss such that there is no penalty on \hat{y} if $y_{lower}^* \leq \hat{y} \leq y_{upper}^*$, and $\hat{B}_0^{c, reproj}(i)$ retrieves the i -th component of $\hat{B}_0^{c, reproj}$.

4.3.2 Box Prior Loss

We use prior knowledge on the object volume and size to regularize the training loss. More specifically, we add the prior loss

$$L_{prior}^{2D} = \sum_{c \in C_{2D}} L_{vol}^c + L_{s-var}^c, \quad (11)$$

to train our network, where

$$L_{vol}^c = \max(0, V^c - \hat{b}_{0,h}^c \hat{b}_{0,w}^c \hat{b}_{0,l}^c), \quad (12)$$

is to penalize predictions with volumes that are below class-specific thresholds from prior knowledge about the scale of objects, and

$$L_{s-var}^c = l(\bar{b}_{0,h}^c, \hat{b}_{0,h}^c) + l(\bar{b}_{0,w}^c, \hat{b}_{0,w}^c) + l(\bar{b}_{0,l}^c, \hat{b}_{0,l}^c), \quad (13)$$

is to penalize the size variance of the predictions within each class in each minibatch since there should not be excessive size variations within a class. $(\hat{b}_{0,h}^c, \hat{b}_{0,w}^c, \hat{b}_{0,l}^c)$ are the size predictions from \hat{B}_0 for class c , and $(\bar{b}_{0,h}^c, \bar{b}_{0,w}^c, \bar{b}_{0,l}^c)$ are the average size predictions per minibatch for class c and V^c is the volume threshold for class c .

5. Training

We train a Faster RCNN-based [22] 2D object detector with the 2D box labels from classes C and train the BoxPC Fit Network by minimizing L_{boxpc}^{3D} from Eq. 6 with the 3D

	bathtub	bed	toilet	chair	desk	dresser	nstand	sofa	table	bkshef	mAP
Fully-Supervised:											
DSS [26]	44.2	78.8	78.9	61.2	20.5	6.4	15.4	53.5	50.3	11.9	42.1
COG [23]	58.3	63.7	70.1	62.2	45.2	15.5	27.4	51.0	51.3	31.8	47.6
2D-driven [13]	43.5	64.5	80.4	48.3	27.9	25.9	41.9	50.4	37.0	31.4	45.1
LSS [24]	76.2	73.2	73.7	60.5	34.5	13.5	30.4	60.4	55.4	32.9	51.0
FPN [20]	43.3	81.1	90.9	64.2	24.7	32.0	58.1	61.1	51.1	33.3	54.0
FPN*	46.0	80.3	82.8	58.5	24.2	36.9	47.6	54.9	42.0	41.6	51.5
FPN* + BoxPC Refine	51.5	81.0	85.0	59.0	24.5	38.6	52.2	55.3	43.8	40.9	53.2
CC Semi-Supervised:											
FPN*	6.1	69.5	28.1	12.4	18.1	17.5	37.4	25.8	23.5	6.9	24.5
FPN* w/o OneHot	24.4	69.5	30.5	15.9	22.6	19.4	39.0	37.3	29.0	12.8	30.1
Ours + R	29.5	60.9	65.3	36.0	20.2	27.3	50.9	46.4	28.4	6.7	37.2
Ours + BoxPC	28.4	67.9	73.3	32.3	23.3	31.0	50.9	48.9	33.7	16.4	40.6
Ours + BoxPC + R	28.4	68.1	77.9	32.9	23.3	30.6	51.1	49.9	34.7	13.7	41.1
Ours + BoxPC + R + P	30.2	70.7	76.3	33.6	24.0	32.5	52.0	49.8	34.2	14.8	41.8

Table 1. 3D object detection AP on SUN-RGBD *val* set. Fully-supervised methods are trained on 2D and 3D box labels of all classes while Cross-category (CC) Semi-supervised methods are trained on 3D box labels of classes in C_{3D} and 2D box labels of all classes. BoxPC, R, P refer to the BoxPC Fit network, relaxed reprojection loss and box prior loss respectively. * refers to our implementation.

box labels from classes C_{3D} . Finally, when we train the 3D object detector, we alternate between optimizing the losses for classes in C_{2D} and classes in C_{3D} . Specifically, when optimizing for C_{2D} , we train f_{box} to minimize

$$L_{box}^{2D} = L_{fit}^{2D} + L_{reproj}^{2D} + L_{prior}^{2D}, \quad (14)$$

where L_{fit}^{2D} , L_{reproj}^{2D} and L_{prior}^{2D} are from Eq. 5, 9 and 11 respectively. When optimizing for C_{3D} , we train f_{seg} and f_{box} to minimize L_{seg}^{3D} and L_{box}^{3D} from Eq. 1 and 2 respectively. Hence, the loss functions for f_{seg} and f_{box} are respectively given by

$$L_{seg} = L_{seg}^{3D}, \quad (15)$$

$$L_{box} = L_{box}^{3D} + L_{box}^{2D}. \quad (16)$$

6. Experiments

6.1. Datasets

We evaluate our model on SUN-RGBD and KITTI benchmarks for 3D object detection. For the SUN-RGBD benchmark, we evaluate on 10 classes as in [20]. To test the performance of cross-category semi-supervised 3D object detection (CS3D), we randomly split the 10 classes into 2 subsets (C_A and C_B) of 5 classes each (the 5 classes on the left and on the right of Tab. 1). First, we let C_A be the *strong* classes and C_B be the *weak* classes by setting $C_{3D} = C_A, C_{2D} = C_B$ to obtain the evaluation results for C_B . Next, we reversed C_A and C_B to obtain the evaluation results for C_A . We follow the same *train/val* split and performance measure of average precision (AP) with 3D Intersection-over-Union (IoU) of 0.25 as in [20].

The KITTI benchmark evaluates on *Cars*, *Pedestrians* and *Cyclists* classes. To test the CS3D setting, we set

any 2 classes to C_{3D} and the remaining class to C_{2D} to obtain the evaluation results on C_{2D} . For example, to obtain the evaluation results on *Cyclists*, we set $C_{3D} = \{Cars, Pedestrians\}$. We follow the same *train/val* split as in [20] and measure performance using AP with 3D IoU of 0.25 for all classes. For fully-supervised setting, we use 3D IoU of 0.7 for *Cars* and 0.5 for *Pedestrians* and *Cyclists*.

6.2. Comparison with Baselines

To the best of our knowledge, there are no other methods that demonstrate weakly- or semi-supervised 3D object detection, therefore we design baselines (details in supplementary) with the state-of-the-art FPN [20] where it trains on strong labels from C_{3D} and tests on C_{2D} . Specifically, we use the original network (“FPN*”) and a network without the one hot class vector (“FPN* w/o OneHot”). The latter performs better since there are no 3D box labels for classes C_{2D} . Finally, we train fully-supervised BoxPC Fit networks to further improve the fully-supervised FPN [20].

SUN-RGBD In Tab. 1, BoxPC, R, P refer to the usage of the proposed BoxPC Fit network, relaxed reprojection loss and box prior loss. The fully-supervised “FPN*” serves as an upper-bound performance for the semi-supervised methods. The semi-supervised baselines “FPN*” and “FPN* w/o OneHot” perform poorly because they are not able to reason about the *weak* classes and predict boxes that are only roughly correct. By adding R and BoxPC, it allows the network to reason about 3D box predictions on *weak* classes, improving performance significantly from 30.1% to 41.1%. Finally, the addition of the prior knowledge P allows our model to achieve a good performance of 41.8% mAP, which is 81.2% of the fully-supervised “FPN*” (vs

Method	Cars			Pedestrians			Cyclists		
	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
Fully-Supervised:									
FPN [20]	96.69	96.29	88.77	85.40	79.20	71.61	86.59	67.60	63.95
CC Semi-Supervised:									
FPN	0.00	0.00	0.00	4.31	4.37	3.74	1.05	0.71	0.71
FPN w/o OneHot	0.08	0.07	0.07	55.66	47.33	41.10	54.35	36.46	34.85
Ours + BoxPC + R	14.71	12.19	11.21	75.13	62.46	57.18	61.25	41.96	40.07
Ours + BoxPC + R + P	69.78	58.66	51.40	76.85	66.92	58.71	63.29	47.54	44.66

Table 2. 3D object detection AP on KITTI *val* set, evaluated at 3D IoU threshold of 0.25 for all classes.

Method	Cars			Pedestrians			Cyclists		
	Easy	Moderate	Hard	Easy	Moderate	Hard	Easy	Moderate	Hard
Fully-Supervised:									
VeloFCN [14]	15.20	13.66	15.98	-	-	-	-	-	-
MV3D [6]	71.29	62.68	56.56	-	-	-	-	-	-
Voxelnet [40]	89.60	84.81	78.57	65.95	61.05	56.98	74.41	52.18	50.49
FPN [20]	84.40	71.37	63.38	65.80	56.40	49.74	76.11	56.88	53.17
FPN + BoxPC Refine	85.63	72.10	64.25	68.83	59.41	52.08	78.77	58.41	54.52

Table 3. 3D object detection AP on KITTI *val* set, evaluated at 3D IoU threshold of 0.7 for *Cars* and 0.5 for *Pedestrians* and *Cyclists*.

58.4% for baseline “FPN* w/o OneHot”). We consider this a promising result for semi-supervised methods, and it demonstrates the effectiveness of our proposed model in transferring knowledge to unseen classes. In addition, we show the effectiveness of the BoxPC Fit network in fully-supervised settings by training the BoxPC Fit network on all classes C and using the network to refine the 3D box predictions of the state-of-the-art FPN [20]. As seen in Tab. 1, “FPN* + BoxPC Refine” outperforms the vanilla fully-supervised “FPN*” in every single class except for *bookshelf*, showing the usefulness of the BoxPC Fit network even in fully-supervised settings.

KITTI In Tab. 2, the baselines are unable to achieve any performance on *Cars* when trained on *Pedestrians* and *Cyclists* because of huge differences in sizes. The network simply assumes a *Car* instance is small and makes small predictions. With the addition of a prior volume, it becomes possible to make predictions on *Cars*. We observe that adding “BoxPC + R + P” significantly improves performance over the baseline’s mean AP of 0.07% to 59.95% for *Cars*. It also improved the baseline’s performance from 48.03% to 67.49% for *Pedestrians* and from 41.89% to 51.83% for *Cyclists*. Similarly, we show that a fully-supervised BoxPC Fit network is able to refine and improve the 3D box predictions made by “FPN” in Tab. 3. In this case, “FPN + BoxPC Refine” improves AP for all classes at all difficulty levels over the original “FPN”.

6.3. Ablation Studies

Relaxed Reprojection Loss To illustrate the usefulness of using a relaxed version of the reprojection loss, we vary

the lower and upper bounds B_{lower}^* and B_{upper}^* . We set B_{lower}^*, B_{upper}^* to different scales of B_{2D}^* , i.e. $1.5 B_{2D}^*$ refers to the 2D box centered at the same point as B_{2D}^* but with $1.5\times$ size. By setting $B_{lower}^*, B_{upper}^* = B_{2D}^*$, it reverts from the relaxed version to the direct version of re-projection loss. We train with similar settings as “Ours + R” in Tab. 1, and show in Tab. 4 that mean AP improves from 25.6% ($B_{upper}^* = B_{2D}^*$) to 37.2% ($B_{upper}^* = 1.5 B_{2D}^*$).

BoxPC Fit Network Objectives The BoxPC Fit network is trained on classification and regression objectives derived from the 3D box labels. To show the effectiveness of having both objectives, we adopt different settings in which classification or regression objectives are active. We train with similar settings as “Ours + BoxPC” in Tab. 1. Without the classification objective for the BoxPC Fit network, the 3D object detector is unable to maximize $\hat{p} = f_{boxpc-cls}(X, \hat{B}_0)$ during training. Omission of the regression objective prevents the initial box predictions \hat{B}_0 from being refined further by $\Delta\hat{B} = f_{boxpc-reg}(X, \hat{B}_0)$ to the final box predictions \hat{B} . From Tab. 5, we see that training without either classification or regression objective causes performance to drop from 40.6% to 32.1% or 39.8%.

BoxPC Representation To demonstrate the importance of a joint Box to Point Cloud input representation as seen on the right of Fig. 4, we train on the “Ours + BoxPC” setting in Tab. 1 with different BoxPC representations. As shown in Tab. 5, the performance of combined learning of BoxPC features exceeds the performance of independent learning of BoxPC features by 6.2%.

B_{lower}^*	B_{upper}^*	bathtub	bed	toilet	chair	desk	dresser	nstand	sofa	table	bkshef	mAP
B_{2D}^*	1.0 B_{2D}^*	17.0	11.9	60.1	24.3	11.2	24.0	46.8	30.3	19.2	11.3	25.6
B_{2D}^*	1.5 B_{2D}^*	29.5	60.9	65.3	36.0	20.2	27.3	50.9	46.4	28.4	6.7	37.2
B_{2D}^*	2.0 B_{2D}^*	16.9	23.4	55.8	35.0	19.6	25.1	49.2	48.3	31.3	3.5	30.8

Table 4. 3D object detection AP on SUN-RGBD *val* set with varying B_{upper}^* for the relaxed reprojection loss function.

Features	Cls	Reg	bathtub	bed	toilet	chair	desk	dresser	nstand	sofa	table	bkshef	mAP
Combined	✓		25.9	67.3	71.8	30.6	23.6	31.8	48.8	48.1	32.6	17.3	39.8
Combined		✓	27.4	69.4	39.8	17.7	23.1	18.5	39.9	41.7	30.2	12.9	32.1
Independent	✓	✓	39.6	68.1	24.5	26.8	22.3	34.4	38.8	46.7	30.3	12.3	34.4
Combined	✓	✓	28.4	67.9	73.3	32.3	23.3	31.0	50.9	48.9	33.7	16.4	40.6

Table 5. 3D object detection AP on SUN-RGBD *val* set when BoxPC Fit Network is trained on different representations and objectives.




Figure 6. Qualitative comparisons between the baseline, fully-supervised and proposed models on the SUN-RGBD *val* set.

6.4. Qualitative Results

We visualize some of the predictions by the baseline, fully-supervised and proposed models in Fig. 6. We used 2D box labels for frustum proposals to reduce noise due to 2D detections. We observe that the proposed model’s predictions are closer to the fully-supervised model than the baseline’s. The baseline’s predictions tend to be inaccurate since it is not trained to fit the objects from *weak* classes.

7. Conclusion

We propose a transferable semi-supervised model that is able to perform 3D object detection on *weak* classes with only 2D box labels. Our method achieves strong performance over the baselines and improves the practicality and usefulness of 3D object detectors in new applications.

Acknowledgement This work is supported in part by a Singapore MOE Tier 1 grant R-252-000-A65-114.

References

- [1] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5297–5307, 2016.
- [2] Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *European Conference on Computer Vision (ECCV)*, 2018.
- [3] Liang-Chieh Chen, Sanja Fidler, Alan L Yuille, and Raquel Urtasun. Beat the mturkers: Automatic image labeling from weak 3d supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3198–3205, 2014.
- [4] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2147–2156, 2016.
- [5] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In *Advances in Neural Information Processing Systems*, pages 424–432, 2015.
- [6] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *IEEE CVPR*, volume 1, page 3, 2017.
- [7] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3354–3361. IEEE, 2012.
- [8] Seunghoon Hong, Hyeonwoo Noh, and Bohyung Han. Decoupled deep neural network for semi-supervised semantic segmentation. In *Advances in neural information processing systems*, pages 1495–1503, 2015.
- [9] Seunghoon Hong, Junhyuk Oh, Honglak Lee, and Bohyung Han. Learning transferrable knowledge for semantic segmentation with deep convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3204–3212, 2016.
- [10] Ronghang Hu, Piotr Dollár, Kaiming He, Trevor Darrell, and Ross Girshick. Learning to segment every thing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4233–4241, 2018.
- [11] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *European Conference on Computer Vision (ECCV)*, pages 695–711. Springer, 2016.
- [12] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven Waslander. Joint 3d proposal generation and object detection from view aggregation. *arXiv preprint arXiv:1712.02294*, 2017.
- [13] Jean Lahoud and Bernard Ghanem. 2d-driven 3d object detection in rgb-d images. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 4632–4640. IEEE, 2017.
- [14] Bo Li. 3d fully convolutional network for vehicle detection in point cloud. In *Intelligent Robots and Systems (IROS), 2017 IEEE/RSJ International Conference on*, pages 1513–1518. IEEE, 2017.
- [15] Qizhu Li, Anurag Arnab, and Philip H.S. Torr. Weakly- and semi-supervised panoptic segmentation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [16] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Košecká. 3d bounding box estimation using deep learning and geometry. In *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*, pages 5632–5640. IEEE, 2017.
- [17] Dim P Papadopoulos, Jasper RR Uijlings, Frank Keller, and Vittorio Ferrari. Extreme clicking for efficient object annotation. In *Proceedings of the ICCV*, pages 4940–4949. IEEE, 2017.
- [18] George Papandreou, Liang-Chieh Chen, Kevin P Murphy, and Alan L Yuille. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1742–1750, 2015.
- [19] Deepak Pathak, Philipp Krahenbuhl, and Trevor Darrell. Constrained convolutional neural networks for weakly supervised segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 1796–1804, 2015.
- [20] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 918–927, 2018.
- [21] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. *Proc. Computer Vision and Pattern Recognition (CVPR), IEEE*, 1(2):4, 2017.
- [22] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [23] Zhile Ren and Erik B Sudderth. Three-dimensional object detection and layout prediction using clouds of oriented gradients. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1525–1533, 2016.
- [24] Zhile Ren and Erik B Sudderth. 3d object detection with latent support surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 937–946, 2018.
- [25] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- [26] Shuran Song and Jianxiong Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 808–816, 2016.
- [27] David Stutz and Andreas Geiger. Learning 3d shape completion from laser scan data with weak supervision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1955–1964, 2018.

- [28] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *European Conference on Computer Vision (ECCV)*, pages 712–729. Springer, 2018.
- [29] Peng Tang, Xinggang Wang, Angtian Wang, Yongluan Yan, Wenyu Liu, Junzhou Huang, and Alan Yuille. Weakly supervised region proposal network and object detection. In *European Conference on Computer Vision (ECCV)*, pages 352–368, 2018.
- [30] Paul Vernaza and Manmohan Chandraker. Learning random-walk label propagation for weakly-supervised semantic segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 3, 2017.
- [31] Dominic Zeng Wang and Ingmar Posner. Voting for voting in online point cloud object detection. In *Robotics: Science and Systems*, volume 1, page 5, 2015.
- [32] Yu Xiang, Wonhui Kim, Wei Chen, Jingwei Ji, Christopher Choy, Hao Su, Roozbeh Mottaghi, Leonidas Guibas, and Silvio Savarese. Objectnet3d: A large scale database for 3d object recognition. In *European Conference Computer Vision (ECCV)*, 2016.
- [33] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *Applications of Computer Vision (WACV), 2014 IEEE Winter Conference on*, pages 75–82. IEEE, 2014.
- [34] Huaxin Xiao, Yunchao Wei, Yu Liu, Maojun Zhang, and Jia-ashi Feng. Transferable semi-supervised semantic segmentation. In *Association for the Advancement of Artificial Intelligence*, 2017.
- [35] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. *arXiv preprint arXiv:1711.10871*, 2017.
- [36] Jia Xu, Alexander G Schwing, and Raquel Urtasun. Learning to segment under various forms of weak supervision. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pages 3781–3790. IEEE, 2015.
- [37] Zi Jian Yew and Gim Hee Lee. 3dfeat-net: Weakly supervised local 3d features for point cloud registration. In *European Conference on Computer Vision (ECCV)*, 2018.
- [38] Dingwen Zhang, Junwei Han, Yang Yang, and Dong Huang. Learning category-specific 3d shape models from weakly labeled 2d images. In *Proc. CVPR*, pages 4573–4581, 2017.
- [39] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *IEEE International Conference on Computer Vision*, 2017.
- [40] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. *arXiv preprint arXiv:1711.06396*, 2017.
- [41] Yanzhao Zhou, Yi Zhu, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Weakly supervised instance segmentation using class peak response. In *Proceedings of the IEEE international conference on computer vision*, pages 3791–3800, 2018.