

3DIoUMatch: Leveraging IoU Prediction for Semi-Supervised 3D Object Detection

He Wang^{1*} Yezhen Cong^{2*} Or Litany³ Yue Gao² Leonidas J. Guibas¹
¹Stanford University ²Tsinghua University ³NVIDIA

Abstract

3D object detection is an important yet demanding task that heavily relies on difficult to obtain 3D annotations. To reduce the required amount of supervision, we propose 3DIoUMatch, a novel method for semi-supervised 3D object detection. We adopt VoteNet, a popular point cloud-based object detector, as our backbone and leverage a teacher-student mutual learning framework to propagate information from the labeled to the unlabeled train set in the form of pseudo-labels. However, due to the high task complexity, we observe that the pseudo-labels suffer from significant noise and are thus not directly usable. To that end, we introduce a confidence-based filtering mechanism. The key to our approach is a novel differentiable 3D IoU estimation module. This module is used for filtering poorly localized proposals as well as for IoU-guided bounding box deduplication. At inference time, this module is further utilized to improve localization through test-time optimization. Our method consistently improves state-of-the-art methods on both ScanNet and SUN-RGBD benchmarks by significant margins. For example, when training using only 10% labeled data on ScanNet, 3DIoUMatch achieves 7.7 absolute improvement on mAP@0.25 and 8.5 absolute improvement on mAP@0.5 upon the prior art.

1. Introduction

Object detection is a key task in 3D scene understanding. It provides a concise representation of raw sensor measurements in the form of semantically meaningful 3D bounding boxes. This low-dimensional representation can already serve numerous applications in AR/VR, as well as in robot navigation and manipulation. As a result, in recent years there has been a surge of interest in developing improved object detection pipelines and indeed current state-of-the-art methods show impressive performance. Yet, much of their success is attributed to the availability of large datasets of 3D scenes that are carefully annotated. While

rapid advances in sensor technology facilitate the collection of 3D scenes at scale, annotating them remains the main bottleneck. This calls for detection methods that can leverage both labeled and unlabeled data at train time.

In this work, we aim to address this requirement by proposing a novel semi-supervised 3D object detection method which we dub 3DIoUMatch. As our backbone object detection network we adopt VoteNet [18], a popular point-based object detector. To provide supervision to the unlabeled scenes, we leverage a teacher-student mutual learning framework [29] and use the bounding box predictions from the teacher network as pseudo-labels to supervise the student network on unlabeled data. However, unlike most pseudo-label techniques that were designed for classification, in the highly complex (joint regression and classification) task of object detection, we observe that the pseudo-labels suffer from significant noise, and using them directly is suboptimal.

Inspired by FixMatch [26], the state-of-the-art semi-supervised learning (SSL) method for 2D image classification that proposed confidence-based filtering to improve pseudo-label quality, we adopt a pseudo-label filtering mechanism for 3D object detection by setting thresholds on predicted class probabilities and objectness scores, so as to filter out teacher proposals with potentially erroneous semantic labels or ones not belong to foreground. While effective, these two criteria alone are not sufficient to capture localization quality, and the pseudo-labels may still have large errors in the bounding box parameters. To that end, we further propose to leverage estimated IoU (intersection over union) as a localization quality measure for pseudo-label filtering. IoU estimation was first proposed in the context of 2D object detection as a localization confidence in the pioneering work IoU-Net [12], where estimated IoU was proven successful in replacement of class confidence for test-time Non-Maximal Suppression (NMS). To the best of our knowledge, leveraging IoU estimation for pseudo-label filtering is a novel idea for SSL on both 2D and 3D object detection. With our newly devised differentiable 3D IoU estimation module, we are able to filter out poorly localized pseudo-labels and leverage estimated IoU

* indicates equal contributions.

Project page: <http://THU17cyz.github.io/3DIoUMatch>

for both train-time and test-time NMS.

A key challenge when filtering based on IoU estimation is how to properly set the threshold. Unlike objectness and class confidence for which high threshold values (e.g. 0.9) work well, 3D IoU is more sensitive to small errors. Setting the threshold too high would reduce the number of pseudo-labels to very few, from which little could be learned. To balance between quality and coverage, we propose a two-stage filtering process: first, using a relatively low IoU threshold; then, an IoU-guided class-aware *Lower-Half Suppression* (LHS) that removes only half of the highly-overlapping boxes with low predicted IoU. Our proposed LHS thus naturally sets a threshold that is both dynamic and class-aware. Our experiments show that LHS outperforms IoU-guided NMS, which suppresses all but the top one during semi-supervised training. Beyond train-time filtering, we also leverage our differentiable 3D IoU module for test-time IoU-guided NMS and optimization-based bounding box refinement, which is not possible with previous non-differentiable 3D IoU modules [32].

Our method consistently improves upon the previous state-of-the-art method, SESS [34], on both ScanNet and SUN-RGBD benchmarks by significant margins. When using only 10% labeled data on ScanNet, 3DIoUMatch outperforms SESS by 7.7 absolute improvement on mAP@0.25 and by 8.5 absolute improvement on mAP@0.5. When using 5% labeled data on SUN-RGBD, 3DIoUMatch outperforms SESS by 4.8 absolute improvement on mAP@0.25 and by 8.0 absolute improvement on mAP@0.5.

Our main contributions can be summarized as follows:

1. We propose a novel semi-supervised method for 3D object detection in point clouds based on pseudo-label propagation along with a carefully designed filtering mechanism.
2. For the first time, we leverage predicted 3D IoU as a localization confidence score for pseudo-label filtering, and further propose IoU-guided Lower-Half Suppression for robust pseudo-label deduplication.
3. We devise a 3D IoU module that enables our localization filtering, IoU-guided bounding box suppression, and IoU optimization for bounding box refinement.
4. We achieve markedly improved performance over prior art on the two major indoor object detection benchmarks, ScanNet and SUN-RGBD.

2. Related Works

Semi-Supervised Learning (SSL) Many of the recent SSL methods [2, 31, 1] leverage consistency regularization, first proposed in [23, 13], which enforces the model to predict consistently across label-preserving data augmentation of different intensity. Borrowing the concept from Mean Teacher [29], the model with frozen weight can be viewed

as the teacher model, otherwise student model. Some methods [2], following Mean Teacher, make the teacher model as the EMA of the student model for further regularization. Pseudo labeling [15] is another popular class of SSL method which can also be treated as a kind of consistency regularization, as one output of the unlabeled data is enforced to be consistent with the other (the pseudo-labels) by being supervised with the other. To improve the quality of pseudo-labels, FixMatch [26], a state-of-the-art SSL work on image classification, has shown that the student network can improve significantly by setting a classification confidence threshold τ_{cls} and filtering out low-confidence predictions from the teacher. With the filtered pseudo-labels, the student model only gets supervised on the unlabeled data whose pseudo-labels are kept. Another key factor to the success of these methods is strong data augmentation. It has been shown crucial to many SSL works [23, 13, 31]. Recent works [1, 26] proposed to adopt even more powerful augmentation such as RandAugment [3] and Cutout [5].

Semi-Supervised Object Detection Since the beginning of the deep learning era, tremendous progress has been made in 2D object detection, e.g region-based detectors [8, 7, 22] and single-stage detectors [16, 21, 30]. Similarly in 3D object detection, a number of deep learning methods have been proposed for different 3D data modalities, e.g. RGBD-based detectors [19, 17], point-based detectors [33, 25, 14, 18], voxel-based detectors [35], point-voxel-based detectors [24], etc.

Despite the great progress in both 2D and 3D object detection, most works focused on a fully-supervised setting. A few works [10, 6] have proposed to leverage unlabeled data or weakly-annotated data for 2D object detection. Under a standard SSL setting as we follow, CSD [11] proposed a consistency regularization method to enforce the consistency between predictions from an image and its flipped version. STAC [27] adopts a two-stage scheme for training Faster R-CNN [22]: in the first stage it pre-trains a detector with labeled data only and then predicts the pseudo labels for the unlabeled data; in the second stage, STAC leverages asymmetric data augmentation and the pseudo-label filtering mechanism to remove object proposals with low confidence. Note that the pseudo-labels are only generated once at the end of the first stage.

The only prior work on semi-supervised point-based 3D object detection, is SESS [34]. SESS is built upon VoteNet [18] and adopts a two-stage training scheme. It leverages a mutual learning framework composed of an EMA teacher and a student, uses asymmetric data augmentation, and enforces three kinds of consistency losses between the teacher and student outputs. Although SESS brings noticeable improvements upon a vanilla VoteNet when using only a small portion of labeled data, we find their consistency regularization suboptimal, as it is uni-




Figure 1. **3DIoUMatch pipeline at semi-supervised training stage.** We adopt as our backbone an extended version of VoteNet with an additional 3D IoU estimation module. For SSL, we utilize a teacher-student mutual learning framework, composed of a learnable student taking strongly augmented input data and an EMA teacher taking weakly augmented input samples. On labeled data, the student network is supervisedly trained. On unlabeled data, the student network takes pseudo-labels from its EMA teacher. To improve the quality of pseudo-label, we adopt a confidence-based filtering mechanism that filters out predictions that fail to pass all thresholds on class probability, objectness, and 3D IoU. We further use IoU-guided Lower-Half Suppression to remove the duplicated predictions. Using the filtered pseudo-labels, we selectively supervise the student predictions that are around the bounding boxes in the pseudo-labels.

formly enforced on all the student and teacher predictions. In this work, we instead propose to apply confidence-based filtering to improve the quality of pseudo-labels from the teacher predictions and we are the first (in both 2D and 3D object detection) to introduce IoU estimation for localization filtering.

IoU Estimation IoU estimation was first proposed in a 2D object detection work IoU-Net [12], which proposed an IoU head that runs in parallel to bounding box refinement and is differentiable w.r.t. bounding box parameters. IoU-Net adds an IoU estimation head to several off-the-shelf 2D detectors and uses IoU estimation instead of classification confidence to guide NMS, which improves the performance consistently over different backbones. Thanks to its differentiability, IoU-Net can perform IoU optimization on bounding box parameters for iterative refinement, which further brings noticeable performance improvement.

For 3D object detection, STD [32] follows IoU-Net to add a simple IoU estimation branch parallel with the box estimation branch and to guide NMS with IoU estimation. This IoU estimation module design, unfortunately, is not suitable for IoU optimization as the features fed to the IoU estimation branch are not differentiable w.r.t. the bounding box size. This is because STD concatenates the canonical coordinates w.r.t. box center with features of interior points to create the initial features for each bounding box proposal, which is unaware of the box size. Moreover, the bounding box location is regressed along with IoU estimation so the bounding box estimated and optimized is the one before regression, namely the feature offset problem [36], which is also a problem in IoU-Net. In contrast, our 3D IoU esti-

mation module is simple yet effective, can support the IoU optimization while avoiding the feature offset problem.

3. Method

In this section, we describe our solution in detail. After formulating the problem in 3.1, since our backbone detection module is based on VoteNet, we summarize it in 3.2. We then explain our newly devised 3D IoU estimation module in 3.3. These set the stage for a detailed description of our proposed 3DIoUMatch pipeline in 3.4. We further explain how we use the estimated 3D IoU for pseudo-label filtering and deduplication in 3.5. Finally, we illustrate how we leverage the pseudo-labels for supervision in 3.6.

3.1. Problem Definition

Given a 3D point cloud representation of a scene $x \in \mathbb{R}^{N \times 3}$ containing a set of objects $O = \{o^{(j)}\}$, we aim at detecting the amodal oriented 3D bounding boxes of all objects in O , along with their semantic class labels. In particular, we are interested in accomplishing this task under challenging conditions of limited supervision where we have access to a (small) set of labeled scenes $\{x_i^l, y_i^l\}_{i=1}^{N_l}$ and a set of unlabeled scenes $\{x_i^u\}_{i=1}^{N_u}$, where N_l and N_u are the number of labeled and unlabeled scenes, respectively. For a labeled scene x , the label y comprises bounding box parameters $\{b^{(j)}\}$ and semantic class labels $\{q^{(j)}\}$ of all ground truth objects $\{o^{(j)}\}$.

3.2. VoteNet for 3D Object Detection

VoteNet [18] is a single-stage object detector for 3D point clouds. Built upon PointNet++ [20] backbone,

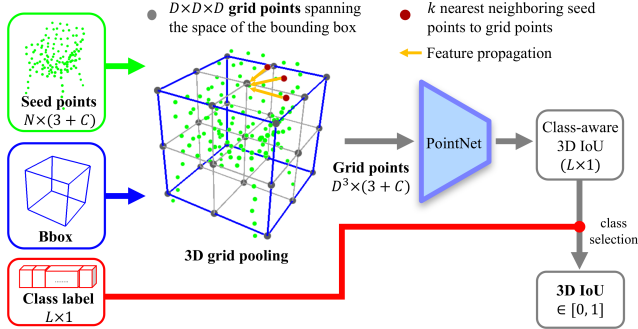


Figure 2. **3D IoU module** takes inputs seed feature points and a bounding box along with its predicted class label, and estimates the 3D IoU between the box and its maximum overlapping ground truth. The module constructs a 3D regular grid with D^3 virtual grid points spanning over the bounding box. We then perform a 3D grid feature pooling that applies a distance-weighted interpolation for feature propagation from the seed points to the grid points. Then the local coordinates of these grid points along with their features are pushed through a PointNet to regress class-aware 3D IoU. Finally, we use the input class label for output selection.

VoteNet first processes the input point cloud $\{x_i\}_{i=1}^N$ to generate a sub-sampled set of $M < N$ seed points enriched with high-dimensional features $\{[x_i; f_i] \in \mathbb{R}^{3+C}\}_{i=1}^M$. Next, each seed point votes for the center of the object it belongs to, and the votes are grouped into K clusters. Finally, each of the K vote clusters is aggregated to make a prediction of a 3D bounding box parameters $b^{(k)}$, a corresponding objectness score $s_k = P(b^{(k)} \text{ is an object})$, and a probability distribution $\{p_{cls}\}$ over L possible semantic classes. The bounding box parameters b are its center location $c \in \mathbb{R}^3$, scale $d \in \mathbb{R}^3$, and orientation θ around the upright axis.

At train time, VoteNet jointly minimizes a weighted combination of the following target losses: vote coordinate regression, objectness score binary classification, box center regression, bin classification and residual regression for heading angle, scale regression, and category classification. We refer the readers to [18] for a detailed description of each loss term.

At test time, VoteNet applies Non-Maximum Suppression (NMS) based on objectness score to remove duplicated bounding boxes. Here, we instead rely on a 3D IoU estimation module which we describe next.

3.3. 3D IoU Estimation Module

As will be made clear in the full pipeline description, a key contribution in this work is an IoU-guided filtering scheme. To facilitate the rejection of poorly localized proposals, as well as guiding deduplication and test-time refinement, we devise a new 3D IoU estimation module differentiable w.r.t bounding box parameters.

In detail, for each predicted bounding box $b^{(k)}$, we wish to estimate its 3D IoU $v^{(k)} \in [0, 1]$ with respect to its corresponding ground-truth box $\{o^k \in \{o^{(j)}\} | k = \text{argmax}_j (\text{IoU}(b^{(k)}, o^{(j)}))\}$. VoteNet does not have intermediate region proposals and only output bounding box parameters at the end stage. Features used for bounding box parameter regression are gathered from vote points in a fixed-radius ball vicinity around each vote cluster, which are unaware of the final bounding box prediction. So, different from implementation in IoU-Net [12] that parallel the bounding refinement and IoU estimation, we need to do it serially by pooling features again specifically for 3D IoU estimation using the final predicted bounding box.

This feature pooling layer takes a bounding box as input and should generate continuous features with respect to the change in bounding box parameters. Existing RoI pooling methods proposed in GSPN [33] and PointRCNN [25] and 3D IoU module proposed in STD [32] simply set a hard cropping boundary at the bounding box surface, taking the point features inside the proposal and ignoring any points outside. These designs have poor differentiability and cause discontinuities whenever a change in the box parameters modifies the point population inside the box, thus are not suitable for 3D IoU optimization (see Table 3).

Here, for the first time, we devise a 3D pooling layer, 3D Grid Pooling, that is differentiable with respect to the change in all bounding box parameters. Inspired by RoIAlign [9] in 2D object detection, we propose to construct virtual grid points spanning the space of the bounding box and their features are obtained by distance-weighted interpolation from real points not restricted inside the box.

Network architecture Taking as inputs the seed points $\{z_i\}_{i=1}^M$, predicted bounding box $b = \{c, d, \theta\}$, and a predicted label l , our 3D IoU module estimates the largest 3D IoU between B and all ground truth bounding boxes. Following IoU-Net [12], the IoU estimation is class-aware.

To build a differentiable 3D IoU module, we first construct a $D \times D \times D$ grid $\{g_m \in \mathbb{R}^3 | m \in [0, D^3 - 1]\}$ that exactly span over the space of b and evenly divide its width, length, and height. For each grid point g_m , we find its k nearest neighbours among all seed points and interpolate their features to get $f_m = \frac{\sum_{i=1}^k w_i f_i}{\sum_{i=1}^k w_i}$, where $w_i = \frac{1}{d(g_m, g_i)^2}$ and d is the L2 distance. Ideally, if k is equal to the number of all seed points, then the IoU module is continuously differentiable. Due to the computational cost, we empirically find $k = 3$ is sufficient for accurate 3D IoU estimation with smooth gradients. We then concatenate g_m and f_m for each grid point and form a grid feature set $\{[g_m; f_m]\}$. The feature set will be pushed towards a PointNet to predict class-aware 3D IoU. A final 3D IoU selection will be performed using the input class label.

Training IoU Estimation Module To train the 3D IoU estimation branch in both stages, we generate on-the-fly training data via jittering the bounding box predictions, i.e. adding Gaussian noise to the box center and size. As a way of data augmentation, this jittering is essential for the generalization of IoU estimation to unlabeled data. We use an L1 loss to supervise the IoU estimation module.

3.4. 3DIoUMatch for SSL on 3D object detection

With the incorporation of 3D IoU module into VoteNet, we construct an IoU-aware VoteNet for SSL on 3D object detection. Our proposed solution is comprised of two training stages: a pre-training stage, where we train our IoU-aware VoteNet on the labeled data, followed by an SSL stage where the entire data is utilized by pseudo-labeling the unlabeled scenes.

Pre-training. We start by training our IoU-aware VoteNet in a supervised manner, using the labeled set $\{\mathbf{x}_i^l, \mathbf{y}_i^l\}_{i=1}^{N_l}$. The training loss is a sum over the original VoteNet losses L_{votenet} and 3D IoU loss L_{IoU} . Once converged, we clone the network to create a pair of student and teacher networks.

Semi-supervised training through a teacher-student framework. We follow a teacher-student mutual learning framework [29] and train our networks on both labeled $\{\mathbf{x}_i^l, \mathbf{y}_i^l\}_{i=1}^{N_l}$ and unlabeled data $\{\mathbf{x}_i^u\}_{i=1}^{N_u}$. Each training batch contains a mixture of $\{\mathbf{x}_i^l\}_{i=1}^{B_l}$ labeled samples and $\{\mathbf{x}_i^u\}_{i=1}^{B_u}$ unlabeled samples.

For labeled samples, we supervise the student network using ground truth supervisions (as done in the pre-training stage) whereas for unlabeled samples, the student networks is supervised using pseudo-labels $\{\tilde{\mathbf{y}}_i^u\}_{i=1}^{N_u}$ generated from the teacher network. The final loss is formed as:

$$L = L_l(\{\mathbf{x}_i^l\}_{i=1}^{N_l}, \{\mathbf{y}_i^l\}_{i=1}^{N_l}) + \lambda_u L_u(\{\mathbf{x}_i^u\}_{i=1}^{N_u}, \{\tilde{\mathbf{y}}_i^u\}_{i=1}^{N_u})$$

where λ_u is the unsupervised loss weight.

To succeed in semi-supervised learning, it is crucial for the teacher network to generate high-quality pseudo-labels and maintain a reliable performance margin over the student network throughout the training. As commonly used in SSL literature, e.g. Mean Teacher [29] and SESS [34], we adopt an EMA teacher. We further leverage asymmetric data augmentation and pseudo-label filtering (see Sec.3.5).

To be in a position of advantage, the teacher network takes input data with weak augmentation only while the student network uses stronger data augmentation. We share the same data augmentation strategy with SESS. The input point clouds to our teacher network are augmented only by random sub-sampling while the inputs to the student network further undergo a set of stochastic transformation \mathcal{T} , including random flip, random rotation around the upright axis, and a random uniform scaling.

3.5. Pseudo-Label Filtering and Deduplication

In the teacher-student framework, the performance gap between the teacher and the student is usually quite marginal given that these two models are just different by EMA on weight and data augmentation strength. Hence, it is not always true that the teacher prediction is more accurate than the student’s on a specific training sample. On unlabeled data, the student model will only benefit from the pseudo-labels that are more accurate than its predictions. Therefore we should filter out low-quality predictions from the teacher model and only supervise the student model with the rest of the teacher model predictions.

Jointly filtering based on class, objectness, localization confidences For VoteNet, we propose to set an objectness threshold τ_{obj} and filter out bounding box predictions with objectness score $s < \tau_{obj}$. We further propose to set a classification confidence threshold τ_{cls} for filtering out predictions that are likely to contain a wrong class label.

Note that none of these two confidence measures capture the accuracy of bounding box parameter predictions. We propose to predict a 3D IoU for each predicted bounding box, use the 3D IoU estimation as a localization confidence, and set a localization threshold τ_{IoU} to filter out poorly localized predictions. Formally, we remove all the predictions that fail to satisfy all three confidence thresholds, i.e. $s > \tau_{obj}$, $\max(p_{cls}) > \tau_{cls}$, and $v > \tau_{\text{IoU}}$.

IoU-guided lower-half suppression for deduplication

After the confidence-based filtering, there is still a lot of duplicated bounding box predictions that may introduce harmful noise to our pseudo-labels. NMS is a standard process in object detection for duplicate removal before evaluation, which takes a set of highly overlapped bounding box predictions that share the same class prediction, ranks them according to a confidence score and removes all but the top-1 prediction. STAC [27] applies class confidence based NMS to teacher predictions during pseudo-label generation.

The default NMS used in VoteNet is based on objectness confidence. Given that objectness score doesn’t capture the localization quality, a train-time IoU-guided NMS will naturally perform better (see Table.2), where we use the product of predicted IoU and predicted objectness as the ranking metric. However, using the top one selected by IoU-guided NMS can still be suboptimal, since the predicted IoU will inevitably carry some errors. We argue that different from the test time scenario, pseudo-labels do not need to be fully deduplicated. Imagine this situation: if a bounding box predicted by the student is $0.2m$ to the left of its corresponding ground truth, it is a foreground object and will get bounding box supervision in VoteNet. However, if unfortunately the pseudo-label survives after non-maximal suppression is to the right of the ground truth more than $0.1m$, this predicted bounding box may lose supervision

Dataset	Model	5%		10%		20%		100%	
		mAP @0.25	mAP @0.5	mAP @0.25	mAP @0.5	mAP @0.25	mAP @0.5	mAP @0.25	mAP @0.5
ScanNet	VoteNet	27.9±0.5	10.8±0.6	36.9±1.6	18.2±1.0	46.9±1.9	27.5±1.2	57.8	36.0
	SESS reported	\	\	39.7±0.9	18.6	47.9±0.4	26.9	62.1	38.8
	SESS	32.0±0.7	14.4±0.7	39.5±1.8	19.8±1.3	49.6±1.1	29.0±1.0	61.3	39.0
	Ours	40.0±0.9	22.5±0.5	47.2±0.4	28.3±1.5	52.8±1.2	35.2±1.1	62.9	42.1
	Abs. improve.	+8.0	+8.1	+7.7	+8.5	+3.2	+6.2	+1.6	+3.1
SUN-RGBD	VoteNet	29.9±1.5	10.5±0.5	38.9±0.8	17.2±1.3	45.7±0.6	22.5±0.8	58.0	33.4
	SESS reported	\	\	42.9±1.0	14.4	47.9±0.5	20.6	61.1	37.3
	SESS	34.2±2.0	13.1±1.0	42.1±1.1	20.9±0.3	47.1±0.7	24.5±1.2	60.5	38.1
	Ours	39.0±1.9	21.1±1.7	45.5±1.5	28.8±0.7	49.7±0.4	30.9±0.2	61.5	41.3
	Abs. improve.	+4.8	+8.0	+3.4	+7.9	+2.6	+6.4	+1.0	+3.2

Table 1. Comparison with VoteNet and SESS on ScanNet val set and SUN RGB-D val set under different ratios of labeled data. We report the mAP@0.25 and mAP@0.5 as mean±standard deviation across 3 runs under different random data splits. Due to the randomness of the data splits and our better pre-training protocol, SESS results provided by us are higher than those reported in the paper on mAP@0.5, and the mAP@0.25 results differ a little (the only difference is the pre-trained weights and data splits). The final improvement is the absolute improvement of our method over SESS results provided by us. Following SESS, we also report the results with 100% labeled data, where we simply make a copy of the full dataset as unlabeled data and train our method.

and be treated as a background box. This example shows that strict non-maximal suppression can lead to a smaller number of student model predictions that can receive supervision. Since we cannot know the best pseudo label among a bunch of highly-overlapped ones, it’s fine to be less strict. To this end, we propose a novel Lower-Half Suppression, or in short, LHS, that only discards half of the proposals with lower predicted IoU. We argue that since LHS suppresses bounding boxes sharing the same class label, this suppression can be seen as a second-step class-aware self-adjusted filtering, which sets dynamic thresholds among the overlapping bounding boxes to keep the ones with higher confidence and hence find a better balance between pseudo-label quality and the amount of supervision. We also use the product of predicted IoU and predicted objectness as the confidence metric.

Final-step pseudo-label processing After the filtering and IoU-guided LHS, we now have high-quality predictions $\{\hat{y}_T^u\}_{k=1}^{K'}$ from the teacher network, where K' is the number of bounding boxes remains. Given that the student model inputs go through a stronger augmentation including an additional geometric transformation \mathcal{T} , in synchronize with the student model inputs, the bounding box parameters of the pseudo-labels need to go through the same transformation \mathcal{T} , namely $\tilde{\mathbf{b}}^u = \mathcal{T}(\hat{\mathbf{b}}_T^u)$. We further take convert the predicted class probability distribution \hat{p}_T^u into semantic class label via $\tilde{q}^u = \max(\hat{p}_T^u)$. Now we obtain the filtered pseudo-labels $\tilde{y} = \{\tilde{\mathbf{b}}^u, \tilde{q}^u\}_{k=1}^{K'}$.

3.6. Selective Supervision using Pseudo-Labels

For our generated pseudo-labels, there is no guarantee that the labels can cover all the ground truth objects from

O due to the filtering and potentially inaccurate teacher predictions. Given the incompleteness of our filtered pseudo-labels, we are relatively confident about the bounding boxes in this set but student predictions far away from all of our pseudo-labels are not necessarily negative. Our experiments show that supervising objectness on unlabeled data using the pseudo-labels seriously hurts the performance. For similar reasons, we do not supervise vote loss, which is a unique element in VoteNet and not shown in other detectors. For more analysis and experimental proof for this, we refer the readers to the supplementary materials. In this case, we will only supervise the bounding boxes in the vicinity of the pseudo bounding boxes and aim to improve their bounding box quality. More specifically, we stick to the way how VoteNet select foreground objects for bounding box parameter supervision: we supervise bounding box parameters and class for a prediction only if the vote that generates this prediction is within 0.3m of any bounding box in the pseudo-labels. For this set of pseudo-foreground predictions, we adopt the same way that VoteNet establishes association and enforce original VoteNet losses except for objectness loss and vote loss.

4. Experiments

4.1. Datasets and Evaluation Metrics

We evaluate our 3DIoUMatch on ScanNet [4] and SUN RGB-D [28]. ScanNet is an indoor dataset consisting of 1,513 reconstructed meshes, among which 1,201 are training samples and the rest are validation samples. SUN RGB-D contains 10,335 RGB-D images of indoor scenes which are split into 5,285 training samples and 5,050 validation samples. For both datasets, we follow[18, 34] for prepro-

Obj&Cls Filter	IoU Filter	Train-time Suppression	Test-time Suppression	Test-time IoU opt.	ScanNet 10%		SUN-RGBD 5%	
					mAP@0.25	mAP@0.5	mAP@0.25	mAP@0.5
			Obj-NMS		38.4	19.8	32.9	12.5
✓			Obj-NMS		44.5	24.7	36.9	17.5
✓		Obj-NMS	Obj-NMS		44.2	25.2	37.1	17.4
✓	✓	IoU-NMS	Obj-NMS		45.9	26.8	37.4	18.7
✓	✓	IoU-LHS	Obj-NMS		46.5	26.9	37.9	18.5
✓	✓	IoU-LHS	IoU-NMS		47.0	28.2	38.8	20.8
✓	✓	IoU-LHS	IoU-NMS	✓	47.2	28.3	39.0	21.1

Table 2. Effects of the different components, including train-time filtering and deduplication, and test-time improvements.

cessing data and labels to train our method and we report mAP@0.25 (mean average precision with 3D IoU threshold 0.25) and mAP@0.5 in the following experiments.

4.2. Implementation Details

Training For the pre-training stage, we train with a batch size of 8 and follow the same data augmentation of SESS [34]. We then use those pre-trained weights to initialize the student and teacher networks. For the SSL stage, we construct each batch by taking 4 labeled samples and 8 unlabeled samples, with the same data augmentation. The weights of different loss terms (e.g. center regression loss, size regression loss, etc.) are the same as VoteNet and we set $\lambda_u = 2$. The student network is trained for 1000 epochs (the labeled data is traversed in one epoch), optimized by an ADAM optimizer with an initial learning rate of 0.002, and the learning rate is decayed by 0.3, 0.3, 0.1, 0.1 at the 400th, 600th, 800th and 900th epoch, respectively. The number of generated 3D proposals is 128. We use $k = 3, D = 4$ for the IoU module. The three thresholds are set to be $\tau_{obj} = 0.9, \tau_{cls} = 0.9, \tau_{IoU} = 0.25$. For more details, we refer the readers to the supplementary materials.

Inference We forward the input to the student network to generate proposals. We first apply IoU optimization to refine the bounding box parameters following IoU-Net [12], followed by an IoU-guided NMS with a 3D IoU threshold of 0.25.

4.3. Result Comparison

Table 1 shows the results of our method compared to SESS and VoteNet under different ratios of labeled data on ScanNet and SUN RGB-D, respectively. The results illustrate that, with our effective train-time filtering and test-time improvement leveraging IoU estimation, we are able to significantly outperform current state-of-the-art, SESS, under all labeled ratio settings. With 5% labeled data, our method outperforms SESS by **8.1** and **8.0** on mAP@0.5 on ScanNet and SUN RGB-D, respectively. Note that our method gains more improvement on mAP@0.5, thanks to the high quality of pseudo labels and the IoU guidance for test-time NMS.

4.4. Ablation Study

Filtering and Deduplication Mechanism. We study the effect of each component of the filtering and deduplication mechanism. In Table 2, the second row shows the results of naive pseudo labeling, which takes all predictions from the teacher model for supervision. Expectedly the results are not satisfying, only a little higher than VoteNet. Simply applying the dual filtering of classification and objectness confidence gives significant improvement, as the filtering picks out the teacher model proposals that are very likely to be close to true objects and have the correct class. The conventional objectness-based NMS in VoteNet, however, fails to improve further, since the remaining proposals already have high objectness scores and the objectness-based NMS is not capable of picking the ones with higher localization accuracy.

As shown in the fifth and sixth row, after we introduce IoU during train time, IoU filtering and train-time IoU-guided NMS contribute to better performance under both settings. Our proposed IoU-guided LHS improves over IoU-guided NMS on mAP@0.25, since LHS finds a better balance between quality and coverage. With better filtering and deduplication leveraging IoU estimation during train time, we gain 2.3 and 1.7 absolute improvement over the without-IoU version on mAP@0.25 and mAP@0.5 respectively on ScanNet 10%. This verifies that considering localization confidence is important for getting high-quality pseudo labels. With test-time improvements, our method gains in total 3.0 and 3.1 absolute improvement respectively.

We set 0.9 for both classification and objectness confidence threshold following STAC [27] and investigate the effect of different IoU thresholds on ScanNet 10%, as shown in Figure 3. The performance (with test-time improvements) is higher than the without-IoU baseline by large margins when $\tau_{IoU} \leq 0.5$. Note that the performance peaks at $\tau_{IoU} = 0.25$ for mAP@0.25 while peaking at 0.5 for mAP@0.5, simply because mAP@0.5 prefers a stronger filtering on localization quality. When $\tau_{IoU} > 0.5$, further increasing τ_{IoU} may lead to a drastic drop in pseudo-label coverage and hence is detrimental to the performance.

Method	ScanNet		SUN RGB-D	
	mAP @0.25	mAP @0.5	mAP @0.25	mAP @0.5
Obj-NMS [18]	57.84	35.99	58.01	33.44
Box query IoU-NMS only	57.56	37.07	58.16	34.81
Box query IoU-NMS + Optim.	57.62	37.17	58.19	34.90
Ours IoU-NMS only	57.92	37.01	58.82	36.22
Ours IoU-NMS + Optim.	58.46	37.43	59.11	36.71

Table 3. Comparison of our IoU module with box-query on ScanNet 100% and SUN RGB-D 100% .

Test-time IoU-guided NMS and IoU optimization. We then evaluate the improvement brought by using IoU estimations at test time. The last two rows in Table 2 shows that IoU-guided NMS and IoU optimization improves the performance further.

As mentioned in 3.3, some region-based 3D detectors, e.g. STD [32], crop the features inside a predicted bounding box and regress the offset. To capture their core characteristics under IoU optimization, we build a simple IoU estimation module which only queries points inside the predicted box and passes the queried feature points through a PointNet to predict the 3D IoU, namely box query. In principle, the differentiability of this module is the same as that in STD, which doesn’t release their code and misses the IoU optimization step in their paper. For a fair comparison, we train another IoU-aware VoteNet with box query as the IoU estimation module and show the comparison between it and our proposed method on the full set of ScanNet and SUN RGB-D. From the results in Table 3, we prove that our method is more effective on both IoU-guided NMS and IoU optimization than box query.

4.5. Result Analysis

In this section, we examine how our 3DIoUMatch works during training on ScanNet 10%. The upper two curves in Figure 4 show that as the training goes, the performance on unlabeled data and test data increases conformably. The increasing performance on unlabeled data indicates the increasing quality of pseudo-labels. We also show how the coverage of the pseudo-labels changes on the unlabeled data over the training. Here coverage at a certain threshold simply means the class-agnostic recall, measuring the percentage of ground truth objects that can find a pseudo-label with an IoU larger than the threshold. As we can see from the lower two curves in Figure 4: at the beginning, the coverage of the pseudo-labels is relatively low due to the strict filtering mechanism; as the semi-supervised learning goes on, the improving detection performance leads to a higher




Figure 3. 3DIoUMatch results with different IoU thresholds on ScanNet 10%.

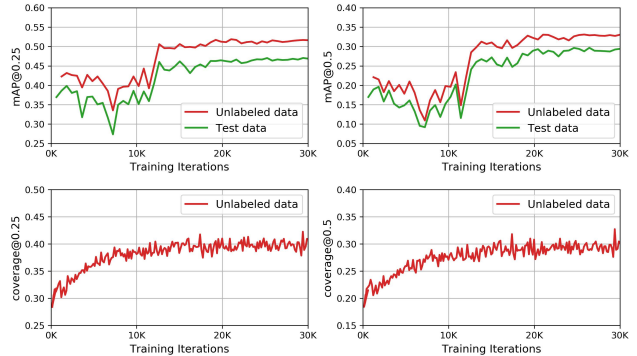


Figure 4. The performance improvements and pseudo-label coverage over the semi-supervised learning stage on ScanNet 10%.

passing rate of the filter and hence a higher coverage of the pseudo-labels, which in return fuels the semi-supervised learning; by the end of training, the coverage at 0.25 and at 0.5 both increase by about 10%.

4.6. Limitations

As discussed in 3.6, we do not supervise objectness loss and vote loss on unlabeled data due to the uncertainty about negative samples. Also, we do not supervise the 3D IoU module on unlabeled data. We expect further performance improvements if the unlabeled data can be utilized to improve their training.

5. Conclusion

In this paper, we propose 3DIoUMatch, a novel semi-supervised 3D object detection method leveraging IoU estimation. Built upon a teacher-student mutual learning framework, we use EMA teacher, asymmetric data augmentation and pseudo-label filtering and deduplication to make the student effectively learn from the teacher. With our IoU estimation module, we make filtering and deduplication aware of localization confidence and apply test-time IoU-guided NMS and IoU optimization, leading to further improvement. Experiment results on two real-world datasets verify

the effectiveness of our method and we achieve significant gain over the prior art under each setting. We believe our idea of leveraging IoU estimation is generally helpful and can be coupled with different kinds of 2D and 3D object detectors to improve their semi-supervised learning.

Acknowledgement: This research is supported by a grant from the SAIL-Toyota Center for AI Research, NSF grant CHS-1528025, a Vannevar Bush Faculty fellowship, and gifts from the Adobe, Amazon AWS, and Snap corporations.

A. Implementation Details

Network Design Our IoU-aware VoteNet shares the same structure with VoteNet[18] except for the IoU estimation module. We provide a more detailed description of the IoU estimation module here. The IoU estimation module is appended after the proposal generation module of VoteNet and takes the bounding box proposals as input. For each bounding box proposal, we create $4 \times 4 \times 4$ virtual grid points. We obtain the relative coordinates of the grid points by subtracting the coordinates of the bounding box center. For every grid point we find its k nearest neighbours among all seed points and interpolate their features to get $f_m = \frac{\sum_{i=1}^k w_i f_i}{\sum_{i=1}^k w_i}$, where $w_i = \frac{1}{d(g_m, g_i)^2}$ and d is the L2 distance. The interpolated features of every grid point is then concatenated with the relative coordinates and forwarded into an MLP with channel dimensions of [256+3, 128, 128, 128] to learn a new feature. Then the features of all grid points go through a global max pooling, after which go through another MLP with channels [128, 128, 128, C], where C is the number of classes, to make the IoU prediction class-aware. Finally, we select the per box IoU estimation by using the class label (during training) or class prediction (during inference).

Training For the pre-training stage, we find that the network does not converge using the same protocol as fully-supervised VoteNet. We instead use a new protocol, where the network is trained for 900 epochs, optimized by an ADAM optimizer with an initial learning rate of 0.001, and the learning rate is decayed by 0.1, 0.1, 0.1 at the 400th, 600th and 800th, respectively. We observe convergence using this protocol on all ratios of labeled data.

Inspired by IoU-Net[12], for both stages, we generate on-the-fly training data via jittering the bounding box predictions for the IoU estimation module. Specifically, we add $\epsilon_{size} \sim N(0, (0.3\mathbf{d})^2)$ to each bounding box size prediction \mathbf{d} and add $\epsilon_{center} \sim N(0, (0.3\mathbf{d})^2)$ to each bounding box center prediction \mathbf{c} to obtain $N_{proposal}$ more training samples. The final IoU estimation loss is the L1 loss averaged over all IoU training samples, original predictions or jitters. The IoU estimation loss weight is 1.

Inference As IoU-Net[12] did not release code, we implemented a simple version of test-time IoU optimization.

1. We obtain the original bounding box proposals.
2. We calculate the gradients of the IoU estimation w.r.t. to bounding box size and center, $grad_{size}$, $grad_{center}$, and update the bounding box size and center by adding $grad_{size} * \lambda$, $grad_{center} * \lambda$ to the box size and center, respectively, where λ is the optimization step size.
3. We repeat the second step for T times.

We find setting T to 10 yields noticeable improvement while not slowing inference speed too much. Choosing λ from the range of $[1e^{-4}, 5e^{-4}]$ has similar performance.

B. Overhead of the IoU module

Our light-weighted IoU estimation module brings moderate overhead to the network, as shown in Table 4. The memory reported in the second column refers to the memory consumed by training with batch size 8 on a single GTX 1080Ti GPU. The last two columns mean the time consumed by a full pass (forwarding and backwarding) of a batch of 8 on a single GTX 1080Ti GPU, training ScanNet and SUNRGB-D respectively. Note that regardless of the network design, there is overhead introduced by calculating the ground truth IoU for supervision.

Method	Mem. (GB)	ScanNet (s)	SUNRGB-D (s)
VoteNet	6.56	0.282	0.316
Ours	6.60	0.325	0.377

Table 4. Memory and time overhead of the IoU module.

C. More on IoU Module Comparison

We provide more explanation on why an IoU estimation network design like that in STD[32] is less effective in IoU estimation and is not differentiable. Given a bounding box proposal, STD concatenates the canonized coordinates and features of the points inside the bounding box to form new features of the points. Therefore, the new feature of a point f' can be denoted as a function of the point coordinates \mathbf{p} , the original point feature f , the bounding box center \mathbf{c} and the bounding box heading angle θ . Then STD voxelizes the bounding box and sample points in each voxel to produce the voxel feature f_v . The process of producing the voxel feature from points in voxels consist of no other parameters except from the point features $\{f'_i\}$ and point coordinates \mathbf{p}_i , so f_v is still a function of $\{f_i\}, \{\mathbf{p}_i\}, \mathbf{c}, \theta$. As all voxel features are flattened and fed to an MLP, which outputs the final IoU, we can conclude that the IoU estimation is not differentiable w.r.t. bounding box size.

	cab	bed	chair	sofa	table	door	wind	bkshf	pic	cntr	desk	curt	fridge	showr	toil	sink	bath	ofurn
VoteNet mAP@0.25	17.9	74.7	74.5	75.3	45.6	18.3	11.7	21.7	0.7	28.4	49.4	21.5	23.2	18.5	79.6	25.7	66.3	11.7
SESS mAP@0.25	20.5	75.1	76.2	76.4	48.1	20.0	14.4	19.4	1.2	30.0	51.8	25.0	30.0	26.4	82.2	29.2	72.3	14.1
Without IoU mAP@0.25	22.6	79.5	77.8	77.8	49.6	25.4	18.6	27.7	3.3	41.4	56.2	27.4	30.4	53.6	81.3	28.5	74.5	18.8
3DIoUMatch mAP@0.25	26.6	82.6	80.9	83.3	52.1	28.0	19.9	29.4	3.7	45.0	61.9	29.2	34.1	51.2	85.7	32.3	82.8	21.5
VoteNet mAP@0.5	3.2	64.6	43.4	49.3	25.1	2.8	1.1	8.7	0.0	2.4	14.7	3.9	7.6	1.1	46.8	11.9	39.4	1.5
SESS mAP@0.5	3.7	61.2	48.0	44.8	29.5	3.2	2.8	8.4	0.2	7.5	19.2	5.0	12.2	1.8	48.7	15.3	40.8	2.2
Without IoU mAP@0.5	3.9	66.1	52.7	50.7	35.1	7.9	5.0	13.1	0.9	14.5	26.1	10.3	17.5	7.0	63.9	11.7	62.1	4.9
3DIoUMatch mAP@0.5	5.9	72.0	60.5	56.6	39.7	10.3	5.2	18.1	0.7	16.0	35.3	8.3	21.4	6.2	67.5	13.2	67.6	5.2

Table 5. Per class mAP@0.25 and mAP@0.5 on ScanNet val set, with 10% labeled data.

	bathtub	bed	bookshelf	chair	desk	dresser	nightstand	sofa	table	toilet
VoteNet mAP@0.25	67.8	32.2	39.4	58.5	53.5	8.0	1.9	14.7	3.2	20.3
SESS mAP@0.25	70.8	34.7	41.9	60.4	63.0	9.8	3.7	25.2	4.0	28.0
Without IoU mAP@0.25	75.1	33.5	43.0	59.5	76.9	6.8	5.1	33.0	3.5	34.8
3DIoUMatch mAP@0.25	75.4	37.7	45.2	64.2	77.0	6.0	5.7	34.6	4.5	39.4
VoteNet mAP@0.5	31.2	6.2	15.5	29.6	14.6	0.5	0.2	2.0	0.3	5.2
SESS mAP@0.5	36.7	7.2	19.2	31.8	20.4	0.7	0.5	7.0	0.4	7.1
Without IoU mAP@0.5	41.5	9.7	25.7	34.5	40.8	0.8	0.8	8.3	0.8	11.4
3DIoUMatch mAP@0.5	45.2	14.4	27.8	43.6	47.2	0.8	1.9	15.7	0.6	13.4

Table 6. Per class mAP@0.25 and mAP@0.5 on SUNRGB-D val set, with 5% labeled data.

We also argue that for VoteNet, since the number of seed points with features are small (1024), box query methods may have difficulty querying points inside a bounding box, especially if a bounding box is too small. Our method, instead won’t suffer from this as we are not confined to points inside the bounding box.

Although STD didn’t release code, we still implemented an IoU estimation module according to the paper for better comparison. However, some issues need to be stated. First, since the backbone of STD is very different from VoteNet, the comparison between IoU estimation module alone is inherently problematic. Second, STD aims at outdoor object detection, where the task is slightly different. Third, we adopted most of the parameters of STD in the paper, but changed number of voxels (to 27) and number of points sampled per voxel (to 6) due to memory concerns and the small number of seed points in VoteNet. The results in Table 7 show the better performance of our IoU module. We also observe serious overfitting using the STD IoU module, suggesting that it may not be suitable for our problem.

D. Why not Supervise Votes and Objectness?

As we mentioned, we supervise all VoteNet loss terms on unlabeled data except for vote regression loss and objectness binary classification loss. As we observe, supervising votes or objectness with pseudo labels leads to degrading performance. The main reason is that by rigorous filtering and deduplication we can only be highly confident of a true object being close to a pseudo bounding box, but we are not sure whether or not there is a true object where there are no pseudo bounding boxes nearby. If we supervise object-

Method	ScanNet 100%		SUNRGB-D 100%	
	mAP @0.25	mAP @0.5	mAP @0.25	mAP @0.5
	Improv.	Improv.	Improv.	Improv.
Ours IoU-NMS	+0.08	+1.02	+0.81	+2.78
Ours IoU Opt.	+0.54	+0.42	+0.29	+0.49
STD IoU-NMS	-0.03	+0.22	+0.20	+1.32
STD IoU Opt.	+0.04	+0.00	+0.00	-0.01

Table 7. Effectiveness of the IoU module in STD compared with our method.

ness on unlabeled data with the pseudo labels the same way as VoteNet, it’s not difficult to imagine the network would be more and more biased on detecting objects. In Table 8, our experiments on ScanNet 10% and SUNRGB-D 5% show that the performance suffers a drop after supervising objectness on unlabeled data.

Vote prediction is an unique component of VoteNet. For a point, the label for its vote is the center of the object it belongs to. To generate pseudo vote labels, the straightforward way is to count every point inside a pseudo bounding box as a vote. However, since this pseudo vote label set is also far from complete, we face a similar problem supervising with it. In Table 8, our experiments on ScanNet 10% and SUNRGB-D 5% also show that the performance drops after supervising vote prediction on unlabeled data.




Figure 5. **Qualitative results on ScanNet, with 10% labeled data.** Here green bounding boxes have an $\text{IoU} \geq 0.25$ while red bounding boxes are with an $\text{IoU} < 0.25$.




Figure 6. Qualitative results on SUNRGB-D, with 5% labeled data.

Method	ScanNet 10%		SUNRGB-D 5%	
	mAP @0.25	mAP @0.5	mAP @0.25	mAP @0.5
3DIoUMatch	47.2	28.3	39.0	21.1
+vote sup. on unlabeled	45.4	28.3	37.9	20.9
+obj. sup. on unlabeled	40.1	26.0	38.2	20.4

Table 8. Objectness & vote supervision on unlabeled data using pseudo-labels.

E. Per-class Evaluation

We report per-class average precision on ScanNet with 10% labeled data and SUNRGB-D with 5% labeled data, respectively. The bold numbers are the highest per class. The results in Table 5, 6 show that our method improves the average precision on nearly all classes over SESS. Our 3DIoUMatch also has better performance on most classes than the without-IoU version.

F. Qualitative Results

We show the qualitative results on ScanNet val set with 10% labeled training data, Figure 5 and on SUNRGB-D val set with 5% labeled training data, Figure 6. For the results of our method, SESS and VoteNet, green bounding boxes are the predicted bounding boxes whose IoU ≥ 0.25 , and the red bounding boxes are those with an IoU < 0.25 . As can be seen in both figures, our method give more accurate predictions and significantly reduces the number of false positives.

References

- [1] David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and Colin Raffel. Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. *arXiv preprint arXiv:1911.09785*, 2019. 2
- [2] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 5049–5059, 2019. 2
- [3] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020. 2
- [4] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. 6
- [5] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 2
- [6] Jiyang Gao, Jiang Wang, Shengyang Dai, Li-Jia Li, and Ram Nevatia. Note-rcnn: Noise tolerant ensemble rcnn for semi-supervised object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9508–9517, 2019. 2
- [7] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 2
- [8] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 2
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 4
- [10] Judy Hoffman, Sergio Guadarrama, Eric S Tzeng, Ronghang Hu, Jeff Donahue, Ross Girshick, Trevor Darrell, and Kate Saenko. Lsda: Large scale detection through adaptation. In *Advances in Neural Information Processing Systems*, pages 3536–3544, 2014. 2
- [11] Jisoo Jeong, Seungeui Lee, Jeeseo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *Advances in neural information processing systems*, pages 10759–10768, 2019. 2
- [12] Borui Jiang, Ruixuan Luo, Jiayuan Mao, Tete Xiao, and Yun-jing Jiang. Acquisition of localization confidence for accurate object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784–799, 2018. 1, 3, 4, 7, 9
- [13] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. 2
- [14] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. 2
- [15] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, 2013. 2
- [16] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 2
- [17] Charles R. Qi, Xinlei Chen, Or Litany, and Leonidas J. Guibas. Invotenet: Boosting 3d object detection in point clouds with image votes, 2020. 2
- [18] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9277–9286, 2019. 1, 2, 3, 4, 6, 8, 9

- [19] 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. 2
- [20] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in neural information processing systems*, pages 5099–5108, 2017. 3
- [21] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [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. 2
- [23] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *Advances in neural information processing systems*, pages 1163–1171, 2016. 2
- [24] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10529–10538, 2020. 2
- [25] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointrcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–779, 2019. 2, 4
- [26] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv preprint arXiv:2001.07685*, 2020. 1, 2
- [27] Kihyuk Sohn, Zizhao Zhang, Chun-Liang Li, Han Zhang, Chen-Yu Lee, and Tomas Pfister. A simple semi-supervised learning framework for object detection. *arXiv preprint arXiv:2005.04757*, 2020. 2, 5, 7
- [28] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, 2015. 6
- [29] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*, pages 1195–1204, 2017. 1, 2, 5
- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019. 2
- [31] Qizhe Xie, Zihang Dai, Eduard Hovy, Minh-Thang Luong, and Quoc V Le. Unsupervised data augmentation for consistency training. *arXiv preprint arXiv:1904.12848*, 2019. 2
- [32] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Ji-aya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1951–1960, 2019. 2, 3, 4, 8, 9
- [33] Li Yi, Wang Zhao, He Wang, Minhyuk Sung, and Leonidas J Guibas. Gspn: Generative shape proposal network for 3d instance segmentation in point cloud. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3947–3956, 2019. 2, 4
- [34] Na Zhao, Tat-Seng Chua, and Gim Hee Lee. Sess: Self-ensembling semi-supervised 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11079–11087, 2020. 2, 5, 6, 7
- [35] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4490–4499, 2018. 2
- [36] Li Zhu, Zihao Xie, Liman Liu, Bo Tao, and Wenbing Tao. Iou-uniform r-cnn: Breaking through the limitations of rpn. *arXiv preprint arXiv:1912.05190*, 2019. 3