

Unsupervised Object Detection with LiDAR Clues

Hao Tian^{1*}, Yuntao Chen^{2,3*†}, Jifeng Dai¹, Zhaoxiang Zhang^{2,3,4}, Xizhou Zhu¹

¹SenseTime Research ²University of Chinese Academy of Sciences

³Center for Research on Intelligent Perception and Computing, CASIA

⁴Center for Excellence in Brain Science and Intelligence Technology, CAS

{zhuwalter, tianhao2, daijifeng}@sensetime.com

{chenyuntao2016, zhaoxiang.zhang}@ia.ac.cn

Abstract

Despite the importance of unsupervised object detection, to the best of our knowledge, there is no previous work addressing this problem. One main issue, widely known to the community, is that object boundaries derived only from 2D image appearance are ambiguous and unreliable. To address this, we exploit LiDAR clues to aid unsupervised object detection. By exploiting the 3D scene structure, the issue of localization can be considerably mitigated. We further identify another major issue, seldom noticed by the community, that the long-tailed and open-ended (sub-)category distribution should be accommodated. In this paper, we present the first practical method for unsupervised object detection with the aid of LiDAR clues. In our approach, candidate object segments based on 3D point clouds are firstly generated. Then, an iterative segment labeling process is conducted to assign segment labels and to train a segment labeling network, which is based on features from both 2D images and 3D point clouds. The labeling process is carefully designed so as to mitigate the issue of long-tailed and open-ended distribution. The final segment labels are set as pseudo annotations for object detection network training. Extensive experiments on the large-scale Waymo Open dataset suggest that the derived unsupervised object detection method achieves reasonable accuracy compared with that of strong supervision within the LiDAR visible range.

1. Introduction

Unsupervised object detection requires localization and classification of object instances without manual annotations in 2D images. Due to the importance of the problem, various relevant tasks have been studied. For example, some weakly-supervised object detection methods [18, 22, 5, 43, 74] seek to detect objects with image-level anno-

*Equal contribution. † This work is done when Hao Tian and Yuntao Chen are interns at SenseTime Research.

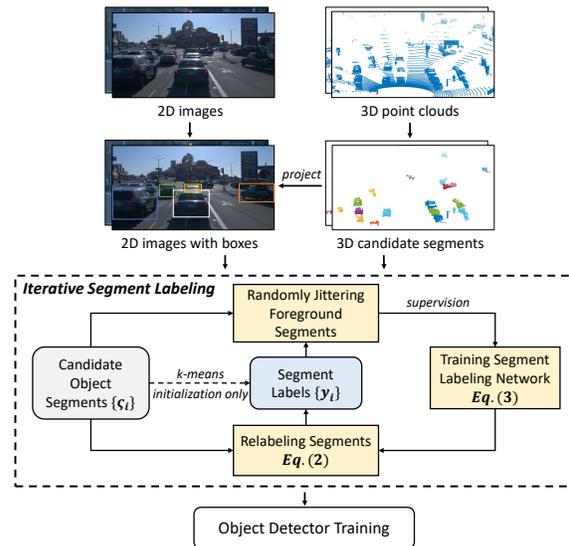


Figure 1: Illustration of the proposed approach.

tations only, while some semi-supervised object detection methods [62, 41, 73] are trained on both bounding box annotated data and additional massive unlabeled images. Unsupervised object proposal generation [12, 25, 77, 67, 1] has also been widely studied. However, to the best of our knowledge, there is no previous work addressing the unsupervised object detection problem.

Recently, considerable progress has been made in unsupervised feature learning [55]. The networks with the unsupervised learned features achieve accuracies on par with those of strong supervision when fine-tuned on downstream tasks. In this trend, some cluster discrimination based methods [85, 42, 35, 9, 91, 13, 79] have tried to address the *unsupervised image classification* problem. Competitive results compared with semi-supervised learning on ImageNet [21] are obtained in [79]. However, there is still a significant gap between unsupervised classification and unsupervised object detection, which involves both localizing and classifying multiple object instances in images.

One widely-known issue for unsupervised object detec-

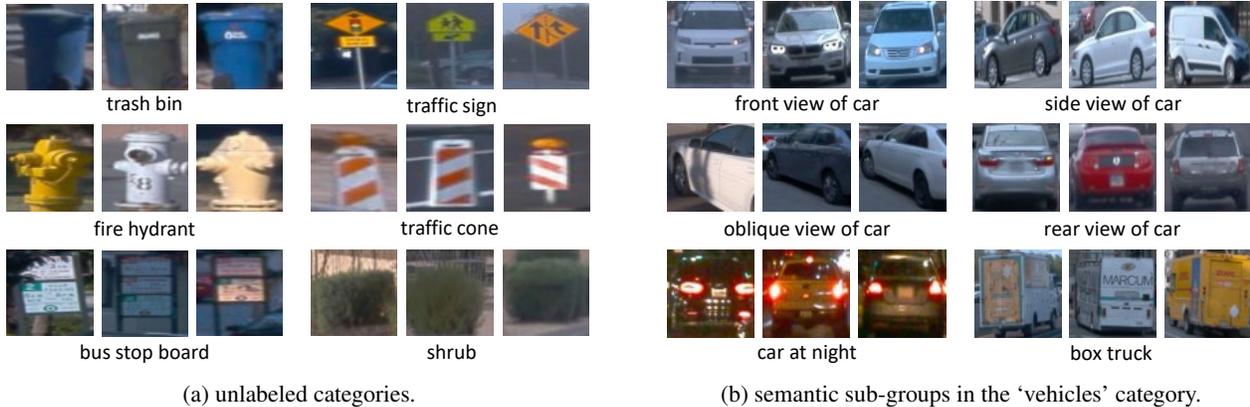


Figure 2: Examples of object categories and semantic sub-groups discovered in the training images, which are not annotated in the Waymo Open dataset [71].

tion is how to localize object instances precisely from the cluttered background without any human annotations. Unsupervised object localization from 2D images is extremely challenging because objects are of heterogeneous colors and textures with various shapes and occlusions. For example, unsupervised object proposals are usually generated by merging over-segmented regions according to color or texture clues [12, 25, 77, 67, 1], which suffer from ambiguous and unreliable object boundaries. Although the object recall rate can be satisfactory with large proposal numbers, the precision rate is very low due to the ambiguities.

For the localization challenge, we argue that the key missing piece is the 3D scene structure, which is essential for human vision. Object boundaries can be better distinguished in the 3D point cloud because different objects cannot occupy the same 3D location. On the other hand, 3D shape information can also be used to better classify objects. Thanks to the popularization of LiDAR sensors, such synchronized 2D images and 3D point clouds have become much easier to obtain. Here we choose to generate candidate object segments based on 3D topology and to learn to label these segments into different categories / clusters¹. The labeling predictions are also based on features of both 2D images and 3D point clouds.

With the localization issue mitigated, we further identify another major issue, seldom noticed by the community: accommodating the long-tailed and open-ended (sub-)category distribution in unsupervised object detection. Research works in unsupervised classification [85, 42, 35, 9, 91, 13, 79] are mostly experimented on balanced and closed-world datasets (e.g., ImageNet [21] and CIFAR [45]), which consist of known number of categories with balanced number of images. In unsupervised object proposal generation [12, 25, 77, 67, 1], all the object cat-

egories merge into one single class of foreground objects. Thus, the long-tail and open-ended distribution problem does not bother. However, for object detection, the datasets (e.g., LVIS [34] and Open Images [47]) often have a long-tailed distribution, due to the nature of natural scenes. In the situation of unsupervised object detection, the difficulty is even further magnified, where the actual number of appeared object categories is unknown. Besides, the long-tailed and open-ended distribution not only exists in object categories but also exists in different semantic sub-groups (e.g., different views, poses) of the same object category. There is no hint which sub-groups belong to the same category till final evaluation or human examination. These semantic sub-groups need to be accurately discovered and detected so as to capture the whole semantic categories. The challenge of long-tailed and open-ended distribution brings difficulty for the labeling of candidate object segments. The head categories contain many object segments, while the tail categories contain very few. Proper labeling mechanism is necessary so as to avoid the tail categories being buried with the cluttered background.

This paper presents the first practical method for unsupervised object detection with the aid of LiDAR clues. The input is a training set composed of synchronized 2D image and 3D point cloud pairs without any type of human annotations, while the output is an object detection network applicable to 2D images. Our approach is illustrated in Figure 1. For each training pair, candidate object segments are first extracted from the 3D point cloud, based on the 3D topology instead of 2D image appearance. Iterative segment labeling is then conducted to assign segment labels and to train a segment labeling network, assuming segments with similar 2D image appearances and 3D shapes are of the same category. Such iterative optimization makes the predicted categories fit the long-tailed and open-ended distributions. The final segment labels are set as pseudo annotations for object detection network training. During testing,

¹The predictions made by our approach are actually of clusters because no annotations are utilized. But for the coherence of terminology, we use the term categories without confusion in the paper.

the trained detection network is applied to 2D images.

In iterative labeling, the segment labeling network and its training mechanism are carefully designed to accommodate the nature of the long-tailed and open-ended distribution. In the segment labeling network, motivated by [72], we cancel the competition among foreground categories to prevent misclassifying objects in tail categories as background. During network training, only segments labeled as foreground are utilized to provide training losses, to avoid the impact of those foreground segments wrongly labeled as background. The negative examples are generated by jittering the segments labeled as foreground. Starting from a large number of allowed object categories (10,000 by default), our approach will automatically discover the effective number and distribution of appeared categories.

Extensive experiments on the large-scale Waymo Open dataset [71] suggest that the derived unsupervised object detection method achieves reasonable accuracy compared with that of strong supervision within the LiDAR visible range. Besides, our approach can detect object categories appear in the training images but are not annotated in the dataset, such as ‘trash bin’, ‘traffic sign’, and ‘fire hydrant’. Figure 2 and Figure 4 show some example results of our proposed approach. Code shall be released.

2. Related Work

Discriminative Unsupervised Feature Learning Recently, learning visual features by discriminative tasks without human supervision has shown great promise [55]. The networks with the learned features achieve accuracy on par with those of strong supervision when fine-tuned on downstream tasks. Most approaches can be categorized into two classes: instance discrimination or cluster discrimination.

Instance discrimination methods treat each image in a dataset as an individual instance and learn discrimination among instances. [23] firstly proposed to learn a classifier with each image as a category. InstDisc [83] replaces the classifier with a contrastive loss over the memory bank that stores previously-computed representations for other images. CMC [76] further extends it by taking multiview of the same image as positive samples. MoCo [37, 16] improves the contrastive learning methods by storing representations from a momentum encoder. SimCLR [14, 15] shows that the memory bank is unnecessary if the batch size is large enough. Very recently, BYOL [32] even discards negative sampling in self-supervised learning.

Cluster discrimination methods [9, 10, 101, 2, 40, 29, 11] learn discrimination among groups of images with similar appearance instead of among individual images. In DeepCluster [9], given the encoded features, k-means is applied to generate pseudo labels. The encoded features are further refined by learning to classify according to these pseudo labels. [10] scales this method to massive non-

curated data. [101] improves the performance by replacing the mutual-exclusive clustering with a local soft-clustering. SeLa [2] and SwAV [11] further formulate the problem as simultaneous clustering and representation learning by maximizing the information between pseudo labels and input data indices.

Most of these methods focus on the quality of the learned features, where the performance is usually evaluated via linear classifier with fixed features, few-shot image classification, and transfer learning [31]. While some cluster discrimination methods [85, 42, 35, 9, 91, 13, 79] would also directly evaluate the quality of clustering. By viewing the task as *unsupervised classification*, they evaluate how well the learned clusters align with the semantic categories. Typically, Normalized Mutual Information [9, 91, 13, 79], Adjusted Rand Index [13, 79], or Accuracy under the best mapping between learned clusters and ground-truth categories [85, 42, 35, 13, 79] are used as the evaluation metrics.

The cluster discrimination methods are most relevant to the proposed method. In our approach, we also iteratively refine the cluster assignment and conduct feature learning. The key difference is we leap forward from image classification to object detection, which involves both localizing and classifying multiple object instances. Besides, the datasets in image classification are usually balanced and of closed-world. While we conduct unsupervised object detection on long-tailed and open-world distributions.

Unsupervised-, Weakly- and Semi-Supervised Object Detection To reduce the hunger of human annotations, various training settings have been studied for object detection. [18, 22, 5, 43, 74] seek to detect objects with image-level annotations only. While [39, 75, 65] train object detectors on data with bounding box annotations for some categories and image-level annotations for other categories. Another prevalent setting is training on bounding box annotated data with additional massive unlabeled data [62, 41, 73]. Recently, there are also some works focusing on training with partial bounding box annotated images [82, 87, 97, 99].

There are also some relevant tasks focusing on unsupervised object localization, where the classification of object instances is not involved. *Unsupervised object proposal generation* has been widely studied on static 2D images (e.g., by grouping similar super-pixels) [12, 25, 77, 67, 1]. There are also research works on unsupervised proposal generation on videos [68, 78, 57, 28, 84], 3D point clouds [44, 19, 24, 6, 98, 7, 36], etc. *Motion segmentation* is a binary labeling task of identifying the individual moving pixels in videos w.r.t. the background motion [8], which is conducted on 2D image frames via background modeling [58, 92] or structure from motion [56, 96, 64, 63].

To the best of our knowledge, there is no previous work addressing the unsupervised object detection problem.

Compared with weakly- and semi-supervised object detection methods, we do not use any type of human annotation. Our approach is relevant to unsupervised object localization. Here we utilize the unsupervised object proposal generation approach as a component in our method.

Supervised 3D Object Detection A typical 3D object detector takes the point cloud of a scene as its input and produces oriented 3D bounding boxes localizing each detected object [33]. These methods can be divided into two categories: region-proposal-based and single-shot-based methods. Region-proposal-based methods generate and classify the region proposals, which typically use multi-view [17, 46, 52, 51], point cloud segmentation [95, 69, 94] or frustum [80, 86, 60, 70] as representations. Single-shot-based methods directly predict class probabilities and regress 3D bounding boxes via single-stage networks. The network is applied on Bird’s Eye View (BEV) [4, 89, 90], discretized voxel [50, 26, 49, 100, 88, 48], or point clouds [93] as representations.

Here we also use point cloud segmentation to generate region proposals for object detection. Aside from being unsupervised, the key difference is our approach utilizes the point clouds in train-time only, providing 3D information to aid unsupervised object detector training. Our trained object detectors are applied to 2D images.

3. Algorithm

3.1. Overview

In our approach, the input is a training set composed of 2D image and 3D point cloud pairs, capturing natural scenes at synchronized times. Note that no annotations are available in the training set. The output is an object detection network applicable to 2D images, which is capable of detecting object categories seen in training².

Figure 1 provides an overview of our approach. For each training pair, candidate object segments are extracted from the 3D point clouds, based on the 3D topology. Each object segment is represented by its corresponding 3D points. A segment labeling network is trained to assign labels for the candidate segments, indicating whether the segments highly overlap with the object instances (foreground / background), and which categories they belong to. Iterative optimization is conducted to assign labels and to train the segment labeling network. Assuming segments with similar 2D image appearances and 3D shapes are of the same category, we may expect such iterative optimization can correct

²The trained object detection network can localize and classify the objects in the test input image. However, it is unaware of the semantic naming of the categories. For example, the network can localize frontal and side-view cars in the image, and classify them to be of “1-st category” and “2-nd category” respectively. But it has no information to associate the “1-st category” and the “2-nd category” to “car”. Further category naming association is necessary for numerical evaluation, as described in Section 3.4.

inconsistent segment labels. The final segment labels are set as pseudo annotations for object detection network training. The trained network is applied to 2D images.

3.2. Exploiting 2D Images and 3D Point Clouds

Object Segment Generation on Point Cloud We generate a set of candidate segments based on some object proposal approaches. Foreground objects and their object categories are discovered by labeling the generated object segments. We found it more reliable to differentiate objects on 3D point clouds than on 2D images because different objects cannot occupy the same 3D location. Here we utilize the 3D point cloud segmentation algorithm in [7] to generate candidate segments. For each image x and its corresponding 3D point cloud P , the generated candidate segments are represented by $S = \{\varsigma_i\}_{i=1}^n$, where n is the number of segments, and ς_i denotes the i -th segment. A segment $\varsigma_i \subseteq P$ is a collection of 3D points, as a subset of point cloud P . Each segment ς_i corresponds to a 2D bounding box b_i , by projecting it on the image plane. Some segments may well localize the objects, while others are on the cluttered background.

Exploiting 2D Images and 3D Point Clouds in Segment Labeling Network

The segment labeling network N takes both the 2D image and 3D point cloud as input. There are two sub-network extracting 2D image features and 3D point cloud features, respectively. Given the 3D segment ς_i , PointNet [61] is applied to extract a 1024-d 3D shape feature. Given the image x and the 2D bounding box b_i , Fast R-CNN [30] with ResNet-50 [38] is applied to extract a 1024-d 2D image appearance feature for the i -th segment. The concatenation of these two features is further fed into a linear classifier, producing a C -d vector (C denotes the number of foreground categories) followed by a sigmoid function to predict the foreground category probabilities s_i , which shall be further discussed in network design & training in Section 3.3.

3.3. Iterative Segment Labeling

Assuming the maximum number of allowed object categories is C , for each image x , labels $Y = \{y_i\}_{i=1}^n$ are assigned for all the candidate segments $S = \{\varsigma_i\}_{i=1}^n$, where $y_i \in \{0, 1, \dots, C\}$. $y_i = 0$ indicates the segment ς_i is on cluttered background. $y_i = c > 0$ indicates the segment ς_i highly overlaps with an object belonging to the c -th category. In unsupervised object detection, the (sub-)category distribution is long-tail and open-ended. Because there is no hint whether a tail category corresponds to / is a subgroup for an annotated category for evaluation, we need to ensure the labeling quality for both the head and tail categories. Here the segment labeling is conducted in an iterative manner, where the assigned labels gradually conform

to the underlying long-tail and open-ended distribution.

Segment Label Initialization First, we need to derive an initial guess of the labels $\{y_i\}_{i=1}^n$. Here each object segment is assigned as one of the C foreground categories. We conduct clustering based on the concatenation of the PointNet features extracted on 3D segment ς_i , and the ResNet-50 features extracted from 2D image patch within bounding box b_i . The PointNet sub-network is randomly initialized. While the ResNet-50 backbone is initialized with ImageNet [21] MoCo v2 [16] pre-trained weights. Note that no annotated labels are involved here. The 3072-d concatenation of PointNet and ResNet-50 features form the representation of the object segments. The features are tuned by applying MoCo v2 self-supervised training, where each object segment is treated as an individual sample. The PointNet and ResNet-50 network weights are updated accordingly. After MoCo v2 training, k-means clustering [3] is performed on the trained 3072-d object segment features (with C clusters). Each object segment is assigned with the corresponding cluster index as its initial category label, $y_i \in \{1, 2, \dots, C\}$.

Segment Labeling Network Design and Training At each round in the iterative segment labeling, given the segment labels $\{y_i\}_{i=1}^n$ generated from the previous round, the segment labeling network is re-trained from initialization. Note that no annotated data is utilized in initializing the segment labeling network. The PointNet sub-network is initialized from the weights in segment label initialization. While the ResNet-50 backbone in Fast R-CNN is initialized with ImageNet MoCo v2 pre-trained weights. The remaining weights in the network are randomly initialized.

Due to the long-tail distribution, a straight-forward implementation of letting segment labeling network directly predict y_i does not work well. The segments corresponding to tail categories are usually low scored, which are hard to be differentiated from the background. Careful sample reweighing in training the segment labeling network may mitigate the problem. But in our approach, the sample labels are actually the predictions by the segment labeling network in the previous iteration. The errors can get magnified if the sample reweighing is not well-tuned.

Here, we find a simple design of the segment labeling network can make it much more robust to long-tail distribution. Motivated by [72], in the segment labeling network, the probabilities of each foreground category are independently predicted using the sigmoid function. In the design, no competition is involved between foreground categories. Each foreground category only needs to differentiate itself from the background clutters, as

$$\mathbf{s}_i = \text{sigmoid}(N(x, b_i, \varsigma_i | \theta)), \quad (1)$$

where N denotes the object segment classification network,

θ is the network parameter, $\mathbf{s}_i \in [0, 1]^C$ are the probabilities that segment ς_i belongs to different foreground object categories, which is trained with a simplified version of Equalization Loss [72], which shall be further elaborated.

Given the scores \mathbf{s}_i of ς_i , the label y_i is derived as

$$y_i = \begin{cases} 0, & \text{if } \max \mathbf{s}_i < \eta, \\ \arg \max \mathbf{s}_i, & \text{otherwise,} \end{cases} \quad (2)$$

where η is the foreground probability threshold. If the segment ς_i has the estimated foreground probabilities \mathbf{s}_i less than η for all categories ($\eta = 0.95$ by default), it is assumed to be on background ($y_i = 0$). Otherwise, the category with the largest probability is set as the label ($y_i = \arg \max \mathbf{s}_i$).

During training, object segments labeled as foreground are utilized to provide training losses. While the object segments labeled as background are discarded. This is because we find some actual foreground segments maybe wrongly labeled as background. Such labeling error deteriorates performance. Here we generate background training examples by randomly jittering the foreground segments. Given a foreground segment ς_i ($y_i > 0$) and its 2D bounding box b_i , random jittering is applied to produce jittered segment $\hat{\varsigma}_i$ and bounding box \hat{b}_i . Specifically, for each bounding box b_i , a target box IoU value $\text{IoU}_{\text{target}}$ is randomly sampled from a uniform distribution between 0.1 and 1.0. Then the jittered box \hat{b}_i is generated by randomly sampling the top-left and bottom-right corners until the box IoU between b_i and \hat{b}_i is within $[\text{IoU}_{\text{target}} - 0.005, \text{IoU}_{\text{target}} + 0.005]$. The jittered segment $\hat{\varsigma}_i$ is derived from modifying ς_i according to \hat{b}_i . In it, the 3D points whose projected 2D coordinates locate outside of the jittered bounding box \hat{b}_i are removed from ς_i . The foreground / background label $\hat{z}_i \in \{0, 1\}$ is defined as $\hat{z}_i = 1$ if the jittered bounding box \hat{b}_i overlaps with b_i large than 0.5 in terms of box IoU, and $\hat{z}_i = 0$ otherwise.

These jittered segments serve as training samples. The training loss for each jittered segment is defined as

$$L(\hat{\mathbf{s}}_i; \hat{z}_i, y_i) = -\hat{z}_i \log(\hat{\mathbf{s}}_{i, y_i}) - \sum_{c=1}^C (1 - \hat{z}_i) \log(1 - \hat{\mathbf{s}}_{i, c}) \quad (3)$$

where $\hat{\mathbf{s}}_i = N(x, \hat{b}_i, \hat{\varsigma}_i | \theta)$ are the predictions made by the segment labeling network, $\hat{\mathbf{s}}_{i, c}$ denotes the c -th value in $\hat{\mathbf{s}}_i$, which is the predicted probability that segment $\hat{\varsigma}_i$ belonging to the c -th category, y_i is the segment label assigned in the previous iteration. This loss function is a simplified version of Equalization Loss [72]. In the Equalization Loss, to prevent misclassifying objects in tail categories as background, the classifier of each tail category will ignore the discouraging gradients from foreground samples of other categories. However, how to define tail categories needs careful tuning. We experimentally found that treating all categories as tail in the Equalization Loss works well. In our simplified

Equalization Loss, jittered segments labeled as foreground ($\hat{z}_i = 1$) with category y_i will encourage the probability prediction \hat{s}_{i,y_i} . Background segments ($\hat{z}_i = 0$) will discourage the probability predictions for all foreground categories. During training, the jittered segments with foreground and background labels are randomly sampled with a 1:3 ratio.

Summary Figure 1 illustrates the iterative segment labeling process. First, we get the initial guess of the labels $\{y_i\}_{i=1}^n$ via k-means clustering [3] based on features learned with self-supervised training. Then, the network training and segment labeling process are applied iteratively for several rounds. By such iterative labeling, the predicted categories fit the long-tailed and open-ended distribution.

Here, as the actual number of appeared object categories is unknown, the maximum number of allowed object categories C is set as a large number in this paper (by default, $C = 10,000$). We empirically observed that as the iteration runs, the produced segment labels naturally follow the long-tail distribution. The head categories occupy most candidate segments, while the tail categories are of the minority. Those non-existing categories within $\{1, \dots, C\}$ finally do not claim any segments. More ablation studies are provided in Section 4.3.

3.4. Object Detector Training and Evaluation

In our implementation, Faster R-CNN [66] with FPN [53] is chosen as the object detection network, where ImageNet MoCo v2 pre-trained ResNet-50 is utilized as the backbone. In training, the 2D bounding boxes and labeled categories of foreground segments (c_i with $y_i \neq 0$) are set as pseudo annotations. During testing, the trained network is applied to 2D images.

A key point during training is to avoid feeding missed foreground objects as negative training examples. In Faster R-CNN, for the RPN head, thanks to excessive background anchor boxes, by uniform random sampling (which is the default choice), the chances of sampling a missed foreground object as negative is low. But for the Fast R-CNN head, the negative region proposals produced by RPN have a good chance to be an actually missed foreground object. Therefore, we only collect region proposals with box IoU between 0.1 and 0.5 with the pseudo annotated boxes as negative examples for Fast R-CNN training. By forcing the sampled region proposals to overlap with the pseudo-annotated boxes, we considerably reduce the chance of sampling a missed foreground object as a negative example.

To mitigate the issue of long-tailed distribution, the simplified version of Equalization Loss [72] in Eq. (3) is also applied for training the classifier in the Fast R-CNN head. Because the competition between foreground categories is canceled, each regional proposal may have multiple high-confidence category predictions. Class-agnostic NMS is

utilized following the Fast R-CNN head to force each detected object with only one category label.

During the evaluation, we test the accuracy of the predictions made by the object detection network w.r.t. the annotated ground-truths. However, we have no idea which discovered cluster³ in training corresponds to which semantic category annotated, without annotations or human examination. Besides, it is also possible that a discovered cluster captures just a subgroup of an annotated category. Thus, for quantitative evaluation, a many-to-one mapping between the discovered clusters and the ground-truth semantic categories is indispensable. Here we set up the mapping on the training set. Assuming there are K ground-truth categories labeled in the dataset, each cluster in our approach is mapped to one of the K labeled categories or a special “others” category. The “others” category is introduced because our approach will discover unlabeled categories. Specifically, for each generated candidate object segment during training, if the segment has overlapped with some ground-truth bounding boxes larger than 0.5, it will be assigned with the ground-truth category label of the bounding box with the largest overlap. Otherwise, the segment will be assigned as “others”. Then, each cluster will be assigned to either a ground-truth category or the “others” category to achieve the minimum error rate. The category mapping and the object detector are finally applied on the test set, where the traditional AP metric can be utilized for evaluation.

4. Experiments

4.1. Dataset and Implementation Details

Evaluation is conducted on Waymo Open [71], which is a recently released large-scale dataset for autonomous driving. The dataset collects 2D videos and 3D point clouds at synchronized time steps. The training and validation sets contain 798 and 202 videos with around 158k and 40k image frames, respectively. In experiments, 2D images from the ‘front’ camera and 3D point clouds from the ‘top’ LiDAR are utilized. The temporal information from videos is not used. There are 3 object categories annotated with 2D bounding boxes in the dataset, i.e., ‘vehicles’, ‘pedestrians’, and ‘cyclists’. We test the feature transferability by fine-tuning on Cityscapes [20] and PASCAL VOC [27].

We also establish a manual-annotated detector baseline. In our approach, the discovered objects for training are all LiDAR visible. However, for the manual-annotated 2D bounding boxes, the LiDAR invisible distant boxes (with distance more than 75 meters) are also included. For a fair comparison, the manual-annotated baseline is trained with LiDAR-visible 2D bounding box annotations only.

Evaluation is measured by the traditional AP metric at

³Here we call the predictions by our approach *cluster* to avoid confusion with the annotated categories.

annotation setting	#images	#bboxes	network weights initialized from	vehicles				pedestrians				cyclists			
				AP	AP _s	AP _M	AP _L	AP	AP _s	AP _M	AP _L	AP	AP _s	AP _M	AP _L
(a) our approach	158k	0	ImageNet MoCo v2	28.7	1.0	34.0	77.2	28.5	6.8	46.0	60.2	1.0	0.0	0.0	1.7
(b) manual annotations	15k	91k	ImageNet MoCo v2	29.2	1.0	33.8	83.2	22.7	1.9	38.9	57.0	3.9	0.0	3.3	17.9
(c) manual annotations	15k	91k	ImageNet supervised	30.1	1.0	36.3	84.1	22.2	1.8	36.6	63.3	4.6	0.0	3.3	23.4
(d) manual annotations	158k	1087k	ImageNet MoCo v2	34.7	1.0	43.8	89.0	37.7	9.6	61.9	79.1	23.8	3.8	28.6	80.2

Table 1: Comparison of different annotation settings on Waymo Open validation.

pre-training task and dataset	Cityscapes						PASCAL VOC		
	AP ^{bbox}	AP ₅₀ ^{bbox}	AP ₇₅ ^{bbox}	AP ^{mask}	AP ₅₀ ^{mask}	AP ₇₅ ^{mask}	AP ^{bbox}	AP ₅₀ ^{bbox}	AP ₇₅ ^{bbox}
random initialization	30.6	55.0	29.9	24.6	48.8	19.8	45.3	71.8	48.9
image classification on ImageNet	37.5	62.1	38.5	31.8	56.9	29.7	54.6	80.9	60.8
our unsupervised object detection on Waymo	36.1	60.8	36.2	30.6	57.1	27.0	53.8	79.3	59.3

Table 2: Transfer results of different pre-trained models on Cityscapes and PASCAL VOC.

the box IoU threshold of 0.5. For better analysis, following COCO evaluation [54], results are also reported on small (area < 32² pixels), medium (32² pixels < area < 96² pixels) and large objects (area > 96² pixels), denoted as AP_s, AP_M and AP_L, respectively.

In our approach, the iteration number for iterative segment labeling is set as 10 rounds when compared with manual annotations, and 1 round in ablation study by default, for experimental efficiency. Note that no manual annotation is involved until the final evaluation. For all experiments, the hyper-parameter setting for Faster R-CNN [66] follows the open source Detectron2 [81] code base. For more dataset and implementation details please refer to Appendix.

4.2. Comparison on Annotation Settings

Table 1 compares the results of our unsupervised object detection approach and those with manual annotations. Comparing Table 1 (a) and (b), our purposed approach achieves on par accuracy on the ‘vehicles’ and ‘cyclists’ categories with that of using 10% manual annotations, where a considerable higher AP is achieved on the ‘pedestrians’ category. Table 1 (d) also presents the result of training with 100% manual annotations, which achieves relatively high APs. Our unsupervised approach delivers reasonable accuracy compared to those of strongly supervised.

Because traditional supervised object detection methods use ImageNet supervised pre-training features, we also ablate the effect of unsupervised pre-training (i.e., MoCo v2). Comparing Table 1 (b) and (c), the difference of unsupervised and supervised pre-training is negligible.

Table 2 presents the feature transferability performance by fine-tuning on Cityscapes [20] and PASCAL VOC [27]. Pre-training by our unsupervised object detection approach demonstrates good feature transferability, which achieves on par performance with that of ImageNet pre-training, which is much better than that of random initialization.

Note that our approach only uses around 10% number of unlabeled images for pre-training compared with that of ImageNet, which consists of 1.28M images with manual classification annotations.

4.3. Ablation Study

Due to limited space, here we only ablate some key factors. More detailed ablations are provided in Appendix. In the ablation studies, for experimental efficiency, only 1 round is applied in iterative segment labeling by default.

Exploiting 2D Images and 3D Point Clouds During training, our unsupervised object detection approach utilizes 3D point clouds for object segment generation and in the segment labeling network. Here, we ablate the necessity of 3D information. Table 3a presents the results of not using 3D point clouds. Without point clouds in the segment labeling network, the result is considerably worse, which indicates the classification can benefit from 3D shape information. We further tried generating object segments on 2D images. Here we adopt MCG [59] to generate candidate segments from 2D images, which incurs large drop in AP. This indicates 3D point clouds are important in localizing objects.

Iterative Segment Labeling Table 4 and Figure 3 present the results of applying different iteration rounds in the iterative segment labeling. At initialization, the number of object segments in different categories are close to uniform (see Figure 3). As the algorithm iterates, the predicted categories gradually conform the long-tail distribution. The redundant categories in initialization disappear. The final detector accuracy also gets improved as the iteration runs, where the generated pseudo annotations fit better with the underlying distributions.

Loss Function for Foreground Classification In this work, we use a simplified version of Equalization Loss [72] (see Eq. (3)) to mitigate the classification challenge of long-

setting	AP		
	veh	ped	cyc
our approach	26.9	21.1	5.2
– segment labeling network N without point clouds	26.3	18.7	3.3
– – candidate segments S estimated by MCG [59] on 2D images	11.0	0.0	0.0

(a) exploiting 2D images and 3D point clouds (1-st iter)

loss function	AP		
	veh	ped	cyc
softmax loss	25.4	20.3	2.0
sigmoid loss	23.8	20.7	0.6
our approach	28.7	28.5	1.0

(b) loss function for classification (10-th iter)

video sequences	AP		
	veh	ped	cyc
10%	23.1	15.5	1.6
25%	25.8	19.2	2.6
50%	26.0	21.6	4.9
100%	26.9	21.1	5.2

(c) quantity of training data (1-st iter)

Table 3: Ablations on Waymo Open validation, where ‘veh’, ‘ped’ and ‘cyc’ are abbreviations for ‘vehicles’, ‘pedestrians’ and ‘cyclists’, respectively. Table (a) gradually removes the usage of 3D point clouds. Table (b) verifies our simplified version of Equalization Loss [72]. Table (c) explores the effectiveness of training data.

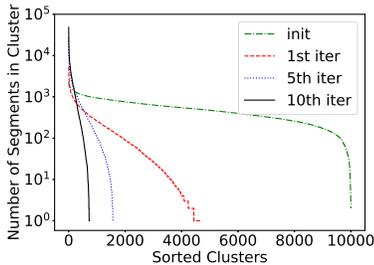


Figure 3: Cluster distribution during the iterative segment labeling.

iterations		1	2	3	4	5	6	7	8	9	10
AP	vehicles	26.9	27.7	28.5	28.2	27.9	28.7	28.7	28.9	28.8	28.7
	pedestrians	21.1	23.8	25.6	26.8	27.0	27.9	28.0	28.3	28.5	28.5
	cyclists	5.2	5.5	6.0	6.3	4.9	2.6	3.2	2.5	1.9	1.0
#cluster	100% segments	4699	2921	2278	1889	1586	1353	1152	986	834	729
	90% segments	1813	1282	982	772	610	489	402	332	281	243
	80% segments	1240	888	673	521	403	320	260	216	185	161

Table 4: Ablation on iterations in the iterative segment labeling on Waymo Open validation.

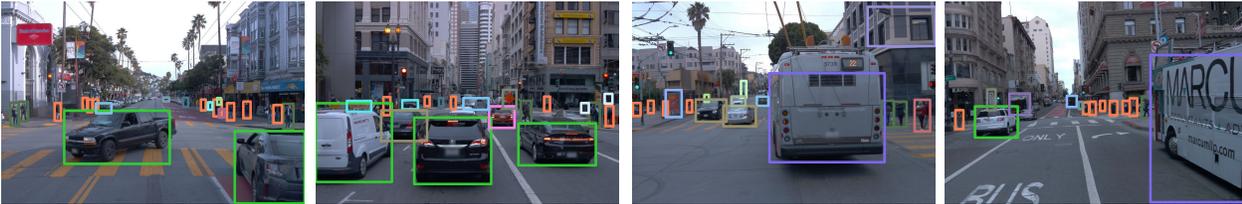


Figure 4: Object detection results on Waymo Open validation by Faster R-CNN trained with our unsupervised object detection method. Different colors of bounding boxes indicate the corresponding cluster id.

tailed distribution. Here we also tried using the vanilla softmax loss and sigmoid loss. Table 3b presents the results after 10 iterations. Using the vanilla softmax loss and sigmoid loss achieve relatively lower performance, which indicates the effectiveness of the proposed loss function.

Quantity of Training Data We further explore the impact of data quantity. Different portions of video sequences are randomly sampled from Waymo for training. As showed in Table 3c, our approach can effectively exploit vast unlabeled images for unsupervised object detection.

4.4. Visualization

Example detection results produced by our unsupervised approach are presented in Figure 4. Our approach can discover object categories and semantic subgroups in the training images but are not annotated. Figure 2a shows some discovered unlabeled categories. Figure 2b shows some dis-

covered semantic subgroups in the ‘vehicles’ category.

5. Conclusion

In this paper, we present the first practical method for unsupervised object detection. During training, 3D point clouds are utilized to mitigate the difficulty of differentiating and localizing objects. We further identify another major issue, seldom noticed by the community: the accommodation of the long-tailed and open-ended distribution in object (sub-)categories. A carefully designed iterative segment labeling process is conducted to generate pseudo annotations for object detection network training. Extensive experiments on the large-scale Waymo Open dataset demonstrate the effectiveness of our approach.

Acknowledgments The work is supported by the National Key R&D Program of China (2020AAA0105200), Beijing Academy of Artificial Intelligence.

References

- [1] Pablo Arbeláez, Jordi Pont-Tuset, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale combinatorial grouping. In *CVPR*, 2014. 1, 2, 3
- [2] Yuki Markus Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. *arXiv preprint arXiv:1911.05371*, 2019. 3
- [3] Bahman Bahmani, Benjamin Moseley, Andrea Vattani, Ravi Kumar, and Sergei Vassilvitskii. Scalable k-means++. *arXiv preprint arXiv:1203.6402*, 2012. 5, 6
- [4] Jorge Beltrán, Carlos Guindel, Francisco Miguel Moreno, Daniel Cruzado, Fernando Garcia, and Arturo De La Escalera. Birdnet: a 3d object detection framework from lidar information. In *ITSC*, 2018. 4
- [5] Hakan Bilen and Andrea Vedaldi. Weakly Supervised Deep Detection Networks. In *CVPR*, 2016. 1, 3
- [6] Igor Bogoslavskiy and Cyrill Stachniss. Fast range image-based segmentation of sparse 3d laser scans for online operation. In *IROS*, 2016. 3
- [7] Igor Bogoslavskiy and Cyrill Stachniss. Efficient Online Segmentation for Sparse 3D Laser Scans. *PFG—Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, 2017. 3, 4
- [8] Alan C Bovik. *The essential guide to video processing*. Academic Press, 2009. 3
- [9] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *ECCV*, 2018. 1, 2, 3
- [10] Mathilde Caron, Piotr Bojanowski, Julien Mairal, and Armand Joulin. Unsupervised pre-training of image features on non-curved data. In *CVPR*, 2019. 3
- [11] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 2020. 3
- [12] Joao Carreira and Cristian Sminchisescu. Constrained parametric min-cuts for automatic object segmentation. In *CVPR*, 2010. 1, 2, 3
- [13] Jianlong Chang, Lingfeng Wang, Gaofeng Meng, Shiming Xiang, and Chunhong Pan. Deep adaptive image clustering. In *ICCV*, 2017. 1, 2, 3
- [14] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. *arXiv:2002.05709*, 2020. 3
- [15] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *NeurIPS*, 2020. 3
- [16] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3, 5
- [17] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *CVPR*, 2017. 4
- [18] Ramazan Gokberk Cinbis, Jakob Verbeek, and Cordelia Schmid. Weakly supervised object localization with multi-fold multiple instance learning. *TPAMI*, 2016. 1, 3
- [19] Alvaro Collet, Bo Xiong, Corina Gurau, Martial Hebert, and Siddhartha S Srinivasa. Exploiting domain knowledge for object discovery. In *ICRA*, 2013. 3
- [20] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *CVPR*, 2016. 6, 7
- [21] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 1, 2, 5
- [22] Ali Diba, Vivek Sharma, Ali Pazandeh, Hamed Pirsiavash, and Luc Van Gool. Weakly Supervised Cascaded Convolutional Networks. In *CVPR*, 2017. 1, 3
- [23] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *TPAMI*, 2015. 3
- [24] Bertrand Douillard, James Underwood, Noah Kuntz, Vsevolod Vlaskine, Alastair Quadros, Peter Morton, and Alon Frenkel. On the segmentation of 3d lidar point clouds. In *ICRA*, 2011. 3
- [25] Ian Endres and Derek Hoiem. Category independent object proposals. In *ECCV*, 2010. 1, 2, 3
- [26] Martin Engelcke, Dushyant Rao, Dominic Zeng Wang, Chi Hay Tong, and Ingmar Posner. Vote3deep: Fast object detection in 3d point clouds using efficient convolutional neural networks. In *ICRA*, 2017. 4
- [27] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The PASCAL Visual Object Classes Challenge: A Retrospective. *IJCV*, 2015. 6, 7
- [28] Katerina Fragkiadaki, Pablo Arbeláez, Panna Felsen, and Jitendra Malik. Learning to segment moving objects in videos. In *CVPR*, 2015. 3
- [29] Spyros Gidaris, Andrei Bursuc, Nikos Komodakis, Patrick Pérez, and Matthieu Cord. Learning representations by predicting bags of visual words. In *CVPR*, 2020. 3
- [30] Ross Girshick. Fast R-CNN. In *ICCV*, 2015. 4
- [31] Priya Goyal, Dhruv Mahajan, Abhinav Gupta, and Ishan Misra. Scaling and benchmarking self-supervised visual representation learning. In *ICCV*, 2019. 3
- [32] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *NeurIPS*, 2020. 3
- [33] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. Deep learning for 3d point clouds: A survey. *TPAMI*, 2020. 4
- [34] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 2
- [35] Philip Haeusser, Johannes Plapp, Vladimir Golkov, Elie Aljalbout, and Daniel Cremers. Associative deep clustering: Training a classification network with no labels. In *GCPR*, 2018. 1, 2, 3

- [36] Frederik Hasecke, Lukas Hahn, and Anton Kummert. Fast lidar clustering by density and connectivity. *arXiv preprint arXiv:2003.00575*, 2020. 3
- [37] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 3
- [38] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016. 4
- [39] 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 *NeurIPS*, 2014. 3
- [40] Jiabo Huang, Qi Dong, Shaogang Gong, and Xiatian Zhu. Unsupervised deep learning by neighbourhood discovery. *arXiv preprint arXiv:1904.11567*, 2019. 3
- [41] Jisoo Jeong, Seungeui Lee, Jeesoo Kim, and Nojun Kwak. Consistency-based semi-supervised learning for object detection. In *NeurIPS*, 2019. 1, 3
- [42] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, 2019. 1, 2, 3
- [43] Vadim Kantorov, Maxime Oquab, Minsu Cho, and Ivan Laptev. ContextLocNet: Context-aware Deep Network Models for Weakly Supervised Localization. In *ECCV*, 2016. 1, 3
- [44] Andrej Karpathy, Stephen Miller, and Li Fei-Fei. Object discovery in 3d scenes via shape analysis. In *ICRA*, 2013. 3
- [45] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 2
- [46] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *IROS*, 2018. 4
- [47] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Tom Duerig, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*, 2018. 2
- [48] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, 2019. 4
- [49] Bo Li. 3d fully convolutional network for vehicle detection in point cloud. In *IROS*, 2017. 4
- [50] Bo Li, Tianlei Zhang, and Tian Xia. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*, 2016. 4
- [51] Ming Liang, Bin Yang, Yun Chen, Rui Hu, and Raquel Urtasun. Multi-task multi-sensor fusion for 3d object detection. In *CVPR*, 2019. 4
- [52] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *ECCV*, 2018. 4
- [53] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. In *CVPR*, 2017. 6
- [54] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 7
- [55] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Zhaoyu Wang, Li Mian, Jing Zhang, and Jie Tang. Self-supervised learning: Generative or contrastive. *arXiv preprint arXiv:2006.08218*, 2020. 1, 3
- [56] Rahul Kumar Namdev, Abhijit Kundu, K Madhava Krishna, and CV Jawahar. Motion segmentation of multiple objects from a freely moving monocular camera. In *ICRA*, 2012. 3
- [57] Dan Oneata, Jérôme Revaud, Jakob Verbeek, and Cordelia Schmid. Spatio-temporal object detection proposals. In *ECCV*, 2014. 3
- [58] Anestis Papazoglou and Vittorio Ferrari. Fast Object Segmentation in Unconstrained Video. In *ICCV*, 2013. 3
- [59] Jordi Pont-Tuset, Pablo Arbelaez, Jonathan T Barron, Ferran Marques, and Jitendra Malik. Multiscale Combinatorial Grouping for Image Segmentation and Object Proposal Generation. *TPAMI*, 2016. 7, 8
- [60] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, 2018. 4
- [61] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 4
- [62] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omni-supervised learning. In *CVPR*, 2018. 1, 3
- [63] Rene Ranftl, Vibhav Vineet, Qifeng Chen, and Vladlen Koltun. Dense monocular depth estimation in complex dynamic scenes. In *CVPR*, 2016. 3
- [64] Anurag Ranjan, Varun Jampani, Lukas Balles, Kihwan Kim, Deqing Sun, Jonas Wulff, and Michael J Black. Competitive collaboration: Joint unsupervised learning of depth, camera motion, optical flow and motion segmentation. In *CVPR*, 2019. 3
- [65] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *CVPR*, 2017. 3
- [66] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-time Object Detection with Region Proposal Networks. In *NIPS*, 2015. 6, 7
- [67] Zhile Ren and Gregory Shakhnarovich. Image segmentation by cascaded region agglomeration. In *CVPR*, 2013. 1, 2, 3
- [68] Gilad Sharir and Tinne Tuytelaars. Video object proposals. In *CVPR Workshops*, 2012. 3
- [69] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, 2019. 4
- [70] Kiwoo Shin, Youngwook Paul Kwon, and Masayoshi Tomizuka. Roarnet: A robust 3d object detection based on region approximation refinement. In *IV*, 2019. 4
- [71] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan,

- Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. *arXiv:1912.04838*, 2019. 2, 3, 6
- [72] Jingru Tan, Changbao Wang, Buyu Li, Quanquan Li, Wanli Ouyang, Changqing Yin, and Junjie Yan. Equalization loss for long-tailed object recognition. In *CVPR*, 2020. 3, 5, 6, 7, 8
- [73] Peng Tang, Chetan Ramaiah, Ran Xu, and Caiming Xiong. Proposal learning for semi-supervised object detection. *arXiv preprint arXiv:2001.05086*, 2020. 1, 3
- [74] Peng Tang, Xinggang Wang, Xiang Bai, and Wenyu Liu. Multiple Instance Detection Network with Online Instance Classifier Refinement. In *CVPR*, 2017. 1, 3
- [75] Yuxing Tang, Josiah Wang, Boyang Gao, Emmanuel Delandréa, Robert Gaizauskas, and Liming Chen. Large scale semi-supervised object detection using visual and semantic knowledge transfer. In *CVPR*, 2016. 3
- [76] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint arXiv:1906.05849*, 2019. 3
- [77] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *IJCV*, 2013. 1, 2, 3
- [78] Michael Van den Bergh, Gemma Roig, Xavier Boix, Santiago Manen, and Luc Van Gool. Online video seeds for temporal window objectness. In *ICCV*, 2013. 3
- [79] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *ECCV*, 2020. 1, 2, 3
- [80] Zhixin Wang and Kui Jia. Frustum convnet: Sliding frustums to aggregate local point-wise features for amodal 3d object detection. *arXiv preprint arXiv:1903.01864*, 2019. 4
- [81] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 7
- [82] Zhe Wu, Navaneeth Bodla, Bharat Singh, Mahyar Najibi, Rama Chellappa, and Larry S Davis. Soft sampling for robust object detection. *arXiv preprint arXiv:1806.06986*, 2018. 3
- [83] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised Feature Learning via Non-Parametric Instance Discrimination. In *CVPR*, 2018. 3
- [84] Fanyi Xiao and Yong Jae Lee. Track and segment: An iterative unsupervised approach for video object proposals. In *CVPR*, 2016. 3
- [85] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *ICML*, 2016. 1, 2, 3
- [86] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *CVPR*, 2018. 4
- [87] Mengmeng Xu, Yancheng Bai, Bernard Ghanem, Boxiao Liu, Yan Gao, Nan Guo, Xiaochun Ye, Fang Wan, Haihang You, Dongrui Fan, et al. Missing labels in object detection. In *CVPR Workshops*, 2019. 3
- [88] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 2018. 4
- [89] Bin Yang, Ming Liang, and Raquel Urtasun. Hdnet: Exploiting hd maps for 3d object detection. In *Robot Learning*, 2018. 4
- [90] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *CVPR*, 2018. 4
- [91] Jianwei Yang, Devi Parikh, and Dhruv Batra. Joint unsupervised learning of deep representations and image clusters. In *CVPR*, 2016. 1, 2, 3
- [92] Yanchao Yang, Antonio Loquercio, Davide Scaramuzza, and Stefano Soatto. Unsupervised moving object detection via contextual information separation. In *CVPR*, 2019. 3
- [93] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *CVPR*, 2020. 4
- [94] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Ipod: Intensive point-based object detector for point cloud. *arXiv preprint arXiv:1812.05276*, 2018. 4
- [95] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *ICCV*, 2019. 4
- [96] Zhichao Yin and Jianping Shi. GeoNet: Unsupervised Learning of Dense Depth, Optical Flow and Camera Pose. In *CVPR*, 2018. 3
- [97] Jihun Yoon, Seungbum Hong, Sanha Jeong, and Min-Kook Choi. Semi-supervised object detection with sparsely annotated dataset. *arXiv preprint arXiv:2006.11692*, 2020. 3
- [98] Dimitris Zermas, Izzat Izzat, and Nikolaos Papanikolopoulos. Fast segmentation of 3d point clouds: A paradigm on lidar data for autonomous vehicle applications. In *ICRA*, 2017. 3
- [99] Han Zhang, Fangyi Chen, Zhiqiang Shen, Qiqi Hao, Chenchen Zhu, and Marios Savvides. Solving missing-annotation object detection with background recalibration loss. In *ICASSP*, 2020. 3
- [100] Yin Zhou and Oncel Tuzel. Voxnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, 2018. 4
- [101] Chengxu Zhuang, Alex Lin Zhai, and Daniel Yamins. Local aggregation for unsupervised learning of visual embeddings. In *ICCV*, 2019. 3