

Back to Reality: Weakly-supervised 3D Object Detection with Shape-guided Label Enhancement

Xiuwei Xu^{1,2}, Yifan Wang¹, Yu Zheng^{1,2}, Yongming Rao^{1,2}, Jie Zhou^{1,2}, Jiwen Lu^{1,2*}

¹Department of Automation, Tsinghua University, China

²Beijing National Research Center for Information Science and Technology, China

{xxw21, yifan-wa21, zhengyu19}@emails.tsinghua.edu.cn; raoyongming95@gmail.com;

{jzhou, lujiwen}@tsinghua.edu.cn

Abstract

In this paper, we propose a weakly-supervised approach for 3D object detection, which makes it possible to train a strong 3D detector with position-level annotations (i.e. annotations of object centers). In order to remedy the information loss from box annotations to centers, our method, namely Back to Reality (BR), makes use of synthetic 3D shapes to convert the weak labels into fully-annotated virtual scenes as stronger supervision, and in turn utilizes the perfect virtual labels to complement and refine the real labels. Specifically, we first assemble 3D shapes into physically reasonable virtual scenes according to the coarse scene layout extracted from position-level annotations. Then we go back to reality by applying a virtual-to-real domain adaptation method, which refine the weak labels and additionally supervise the training of detector with the virtual scenes. Furthermore, we propose a more challenging benchmark for indoor 3D object detection with more diversity in object sizes for better evaluation. With less than 5% of the labeling labor, we achieve comparable detection performance with some popular fully-supervised approaches on the widely used ScanNet dataset. Code is available at: <https://github.com/wyf-ACCEPT/BackToReality>.

1. Introduction

3D object detection is a fundamental scene understanding problem, which aims to detect 3D bounding boxes and semantic labels from a point cloud of 3D scene. Due to the irregular form of point clouds and complex contexts in 3D scenes, most existing 2D methods [31, 32, 50] cannot be directly applied to 3D object detection. Fortunately, with the development of deep learning techniques on point cloud understanding [27, 28], recent works [11, 20, 25, 35, 51] have

*Corresponding author.

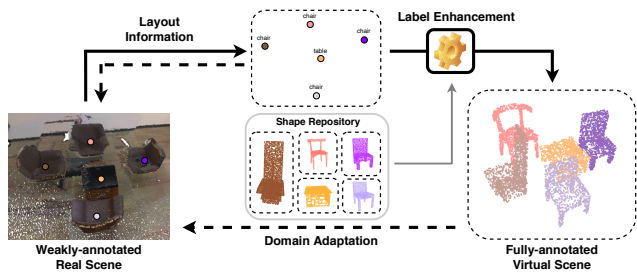


Figure 1. Demonstration of BR. We consider position-level annotations as the coarse layout of the scenes, which is utilized to generate virtual scenes from a 3D shape repository. Physical constraints are applied on the virtual scenes to remedy the information loss from box annotations to centers. Then a virtual-to-real domain adaptation method is presented to additionally supervise the real-scene 3D object detection with the virtual scenes. Dashed arrows indicate supervision for training.

employed deep neural networks to directly detect objects from point clouds and achieved favorable performance.

Despite the successes in deep learning based object detection on point clouds, massive amounts of labeled bounding boxes are required for training the detector. This issue significantly limits the applications of these methods, as labeling a precise 3D box takes more than 100s even by an experienced annotator [36]. Therefore, 3D object detection methods using cheap labels are desirable for practical applications. Motivated by this, increasing attention has been paid to weakly-supervised 3D object detection methods, which can be divided into two categories according to the form of annotation: scene-level [33] and position-level [21, 22] where only the class tag and both object center and class are annotated for each object respectively. The two types of annotation only require less than 1% and 5% time for one instance compared to labeling a bounding box, as shown in Table 1. While scene-level annotation is more time-saving, it is hard for the detector to learn how to precisely locate each object in a

Table 1. Annotating time and detection results of different methods based on various types of annotation. The benchmark is detailed in Section 4. (BBox refers to box annotation. S-L and P-L mean scene-level and position-level annotations respectively.)

| Annotation | BBox [20] | S-L [33] | P-L [21] | P-L(BR) |
|--------------------|-----------|----------|----------|---------|
| Time(s per object) | 110 | 1 | 5 | 5 |
| mAP@0.25(%) | 54.2 | <20 | 32.4 | 47.0 |

scene due to the lack of position information, and thus the performance is far from satisfactory [33]. Considering the time-accuracy tradeoff, position-level annotation is a more practical solution. However, previous position-level weakly-supervised 3D detection methods still require a number of precisely labeled boxes and can only cope with sparse outdoor scenes [21, 22]. Purely position-level weakly-supervised method for the complicated indoor detection task is still under exploration.

In this paper, we propose a shape-guided label enhancement approach called *Back to Reality* (BR) for weakly-supervised 3D object detection¹. To reduce the labor cost, we only label the center of each object in the 3D space and the labeling error of centers is allowed². While largely reducing the workload of labeling, the information loss is non-negligible from box annotations to centers. To address these, BR converts the weak labels into virtual scenes which contain much of the lost information, and in turn utilizes them to additionally supervise real-scene training, as shown in Figure 1. Our approach is based on two motivations: 1) in 3D vision, large-scale datasets of synthetic shapes are available. They contain rich geometry information, which can serve as strong prior to assist 3D object detection; 2) the position-level annotations are not only supervision for training, but they also provide coarse layout of the scene. Therefore, we assemble the 3D shapes into fully-annotated virtual scenes according to the coarse layout and apply physical constraints on them to remedy the information loss. Then a virtual-to-real domain adaptation method is presented to align the global features and object proposal features extracted by the detector between the real and virtual scenes. Moreover, our method can take advantage of the precise center labels in virtual scenes to correct the center error of position-level annotations. In this way the useful knowledge contained in virtual scenes is transferred back to reality. Experimental results on ScanNet [7] show the effectiveness of the proposed BR method.

2. Related Work

3D Shape to Scene: Since it is much easier to obtain a large scale synthetic 3D shape dataset than a real scene

¹Label enhancement (LE) is a technique to recover label distributions from logical labels, as defined in [46]. Here we extend the concept of LE to denote the process of recovering the lost information for weak labels.

²We show the detailed labeling strategy in Section 3.1.

dataset, utilizing the shapes to assist scene understanding is a promising idea. Existing approaches can be divided into two categories: supervised [2, 3, 6, 40] and unsupervised [8, 19, 24, 30, 42]. For supervised methods, the synthetic shapes are usually used to complete the imperfect real scene scans. Given a set of CAD models and a real scan, a network is trained to predict how to place the CAD models in the scene and replace the partial and noisy real objects [2, 3, 6, 40]. Human-annotated pairs of raw scans and object-aligned scans are used in the training process. As supervised methods need extra human labor, that may limit the full utilization of 3D shape datasets. Unsupervised methods are usually used for data augmentation or dataset expansion. 3D CAD models are placed in a random manner following the basic physical constraints, in order to generate mixed reality scenes [8, 42] or virtual scenes [19, 24]. Recently, RandomRooms [30] proposes to use ShapeNet dataset for unsupervised pretraining of 3D detector. Our approach also utilizes 3D shapes to assist object detection in an unsupervised manner. Differently, we aim to make use of synthetic shapes to enhance the weak label and gain stronger supervision in position-level weakly-supervised detection task.

3D Object Detection: Early 3D object detection methods mainly include template-based methods [16, 18, 23] and sliding-window methods [37, 38]. Deep learning-based 3D detection methods for point clouds began to emerge thanks to PointNet/PointNet++ [27, 28]. However, methods in [4, 5, 15, 26] rely on generating 2D proposals and then project them into the 3D space, which is hard to handle scenes with heavy occlusion. More recently, networks that directly consume point clouds have been proposed [11, 20, 25, 35, 51]. While the development of 3D object detection methods is rapid, the application is still restricted partially due to the limited labeled data. To reduce the labor of human annotation, weakly-supervised methods [21, 22, 29, 33], semi-supervised methods [41, 49] and unsupervised pretraining methods [12, 30, 45, 47] have been proposed recently. However, pretraining methods rely on huge computing resources for training the networks in a contrastive learning manner. Semi-supervised methods follow the similar procedure as their 2D counterparts [39] and do not fully exploring the characteristics of 3d data. Therefore, we investigate weakly-supervised approach tailored for 3D object detection task.

3. Approach

Figure 2 illustrates the framework of our approach. Given real scenes with position-level annotations, we utilize 3D shapes to convert the weak labels into virtual scenes, which are utilized to provide additional supervision for the training of the detector. In this section, we first discuss our weakly-supervised setting and then demonstrate the key steps of BR.

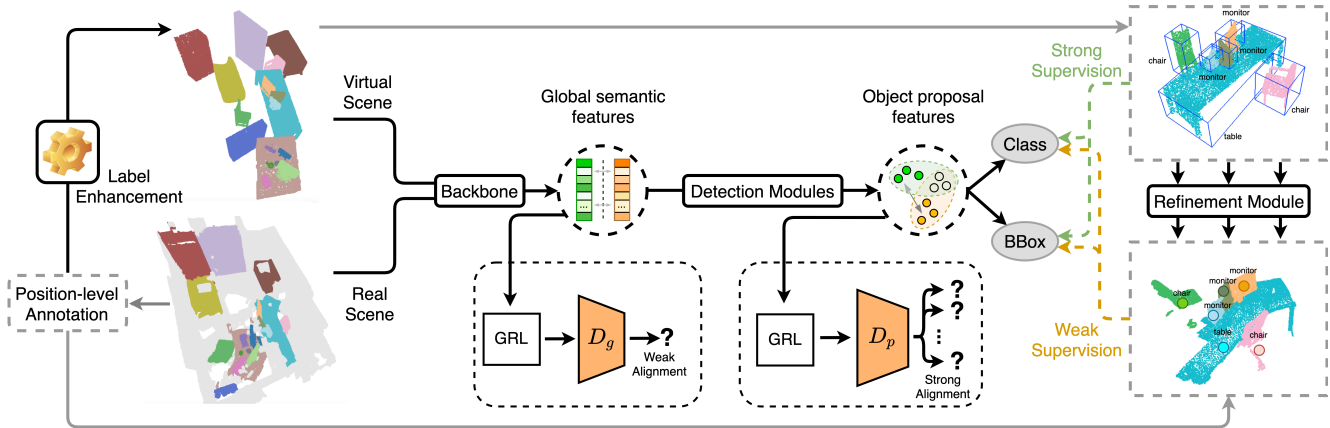


Figure 2. The framework of our *BR* approach. Given real scenes with position-level annotations, we first enhance the weak labels to get fully-annotated virtual scenes. Then the real scenes and virtual scenes are fed into the detector, trained with weakly-supervised and fully-supervised detection loss respectively. During training we use the precise object centers in virtual scenes to refine the imprecise centers in real scenes. Strong-weak adversarial domain adaptation method is utilized to align the distributions of features from both domains. The global discriminator outputs judgments for each scene, and the proposal discriminator outputs judgments for each object proposal. (Here GRL refers to gradient reversal layer; D_g and D_p stand for the global and proposal discriminators respectively.)

3.1. Position-level Annotation

As choosing a point in the 3D space is hard, we divide the labeling process into two steps: firstly we label the center of an object in a proper 2D view of the scene, and compute the line that goes through this center and the focus point of the camera according to the camera parameters of the 2D view. Secondly we choose a point on the line to determine the object’s center in the 3D space. This strategy requires less than 5s to label an instance, and the labeling error can be controlled within 10% of the instance size.

When the 3D scene is scanned, in many cases we can acquire **mesh** data. We assume the meshes are available in our input. Nevertheless, case where we only have point cloud data is also considered in our approach and experiments.

3.2. Shape-guided Label Enhancement

While position-level annotation requires far less labeling time, its information loss is severe, which is manifested in two aspects: 1) the information of objects’ sizes is lost; 2) the object centers are imprecise. In spite of this, position-level annotations can provide a coarse layout of the scenes. By assembling synthetic 3D shapes according to the layout, we are able to enhance the weak labels and generate accurately-annotated virtual scenes where sizes are available and centers are precise. Our label enhancement method is two-step: 1) first we calculate some basic properties of 3D shapes; 2) then we place these shapes to generate physically reasonable virtual scenes from the labels. We provide some implementation details in the supplementary³.

Definition of Shape Properties: Given a synthetic 3D shape, which is represented as $O \in R^{N \times 3}$, we assume it is

³we use * to indicate that the exact definition is in supplementary.

axis-aligned and normalized into a unit sphere. The length, width and height of O is defined as l , w and h . Then we divide the categories of shapes into three classes: supporter, stander and supportee. Supporters and standers are objects that can only be supported by ground, with the difference that standers are not likely to support other things. Other categories are supportees.

Then if a shape belongs to supporter, three properties are calculated: minimum-area enclosing rectangle (MER^*), supporting surface height (SSH^*) and compactness of the supporter surface (CSS^*). The MER is computed in XY plane, which is the minimum rectangle enclosing all the points of the shape. The SSH is the height of the highest surface on which other objects can stand. The CSS is a boolean value, indicating whether the supporting surface can be approximated by the MER .

Virtual Scene Generation: We utilize a three-stage approach to construct the virtual scenes, which is equivalent to generate the position of each shape stage by stage: 1) we first refine the coarse layout provided by position-level annotations and generate the initial positions; 2) then we generate gravity-aware positions by restoring the supporting relationships between objects; 3) lastly we generate collision-aware positions to make the virtual scenes physically reasonable. The pipeline is shown in Figure 3.

To generate *initial positions*, we need to recover a more precise layout from the geometric information of the scenes. Given a scene in mesh format, we first oversegment the meshes using a normal-based graph cut method [9, 13]. The result is a segment graph, where the nodes indicating segments and the edges denoting adjacency relations. Then for horizontal* segments whose area* is larger than A_{min} and height* is larger than H_{min} , we iteratively merge their

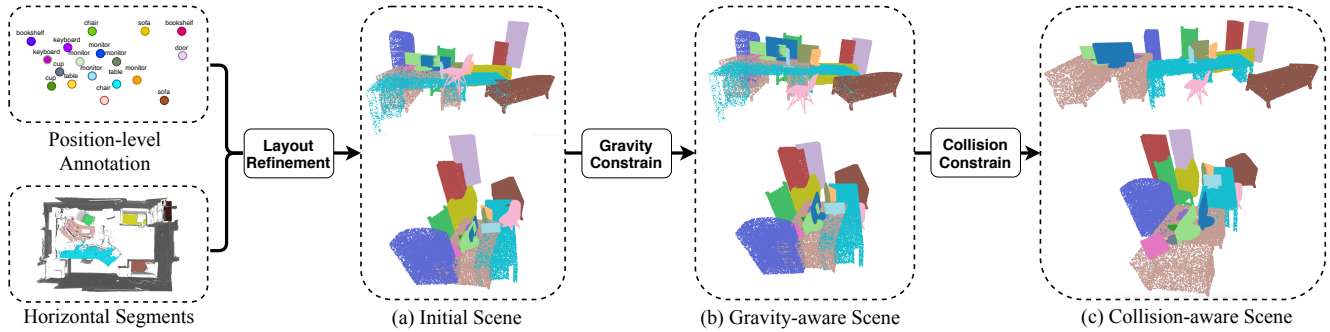


Figure 3. The pipeline of our three-stage virtual scene generation method. We first extract horizontal segments from the mesh data and use them to refine the coarse layout provided by position-level annotations. Then synthetic 3D shapes are placed in virtual scenes according to the new layout to construct initial virtual scenes. After that we apply gravity and collision constraints on the virtual scenes to restore the lost physical relationships between objects and make the scenes more realistic.

neighbors into them if the height difference between the horizontal segment and the neighbor segment is smaller than Δ_h . Once merged, the segments are considered as a whole and the height of the new merged segment is set to be same as the original horizontal segments. After merging, each horizontal segment is represented by its *MER*. If only one supporter’s center falls in a *MER*, we assign this *MER* to the supporter. When the centers of multiple supporters fall in the same *MER*, we perform K-means clustering of the horizontal segment according to these centers and calculate *MER* for each supporter respectively.

Then we place the 3D shapes of corresponding categories on the centers given by position-level annotations and utilize the horizontal segments to refine the layout. The initial positions of the shapes are represented by a dictionary, whose key is the instance index and value is a list:

$$[(x, y, z), (s_x, s_y, s_z), O, \theta, S, M, H] \quad (1)$$

where the instance index is a integer ranging from 1 to the number of objects in the scene. (x, y, z) denotes the center coordinates. (s_x, s_y, s_z) indicates the scales in three dimensions. θ is the rotation angle of the shape. S tells whether the shape is a supporter. M and H indicate the *MER* and *SSH* of supporter. They are set to None when S is false. If the shape has been assigned a horizontal segment, we use the *MER* of that segment to initialize the above parameters. That is, we choose a supporter whose *CSS* is True and make the *MER* of this supporter overlap with the horizontal segment. Otherwise we conduct random initialization. If only point cloud data is available, we simply perform random initialization and the following stages are the same.

Next we traverse the initial positions to generate *gravity-aware positions*. In this process we only need to change z and *SSH* in the position dictionary. For supporters and standers, we directly align their bottoms with the ground (i.e. the XY plane). For a supportee, if its (x, y) fall in any supporter’s *MER*, we assign it to the nearest supporter and align its bottoms with the supporting surface. Otherwise, it

is aligned to the ground.

After that we move the shapes to acquire *collision-aware positions*. This stage only x and y in the position dictionary will be changed. First we move the objects on the ground, the supported ones on which will move together if there are. Then for each supporter, we move its supportees until there is no overlap. Note that the three generation stages can not only make the virtual scenes more realistic, but also weaken the impact of imprecise center labels. Thus the virtual scene generation method is robust to labeling errors.

Finally, we convert the collision-aware positions to point clouds with proper density. As larger surfaces are more likely to be captured by the sensor, we use the maximum of $(l_{s_x})(w_{s_y})$, $(w_{s_y})(h_{s_z})$ and $(l_{s_x})(h_{s_z})$ to approximate the surface area of shapes. Then the number of points for each object is set proportional to their surface areas using uniform sampling, the largest one remaining N points.

3.3. Virtual2Real Domain Adaptation

While the label enhancement approach is able to generate physically reasonable fully-annotated virtual scenes, there is still a huge domain gap between them and the real scenes (e.g. backgrounds like walls are missed in the virtual scenes). Therefore, we need to mining useful knowledge in the perfect virtual labels to make up for the information loss of position-level annotations, rather than just relying on the virtual scenes.

We refer to the virtual scenes and real scenes as source domain and target domain respectively. A virtual-to-real adversarial domain adaptation method is utilized to solve the above problem, whose overall objective is:

$$\begin{aligned} \max_D \min_O J &= L_{sup}(O) - L_{adv}(O, D) \\ &= (L_1 + L_2 + L_3) - (L_4 + L_5) \end{aligned} \quad (2)$$

where O refers to the object detection network (detector) and D indicates the discriminators used for adversarial feature alignment. L_{sup} aims to minimize the differences between the predicted bounding boxes and the annotations,

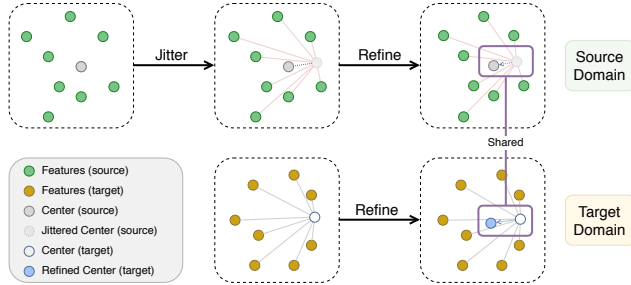


Figure 4. Demonstration of our center refinement method. We first jitter the center labels in source domain, and utilize a PointNet-like module to predict the center offset from the local graph of the jittered centers. This module can be directly utilized to predict the center error in target domain as the global semantic features from the two domains have been aligned.

which can be further divided into the loss for center refinement module (L_1), fully-supervised detection loss on source domain (L_2) and weakly-supervised detection loss on target domain (L_3). The objective of L_{adv} is to align the features from source domain and target domain, which aims to utilize the knowledge learned from source domain to assist object detection in target domain. L_{adv} can be divided into global feature alignment loss (L_4) and proposal feature alignment loss (L_5). Below we will explain these loss functions and our network in detail.

Firstly we elaborate on $L_{sup}(O)$. As shown in Figure 2, we divide the detector into three blocks: a backbone which extracts global semantic features from the scene, a detection module which generates object proposals from the semantic features, and a prediction head which predicts the semantic label and bounding box from each object proposal feature.

During training, we jointly refine the imprecise center labels in target domain and supervise the predictions of the detector. As shown in Figure 4, we jitter the center labels in source domain by adding noise within 10% of the objects' sizes to imitate the labeling error in target domain. Then for each jittered center, we query its k nearest neighbors in 3D euclidean space from the global semantic features to construct a local graph, and predict the center offset through a PointNet-like module:

$$p(c) = \text{MLP}_2 \left\{ \max_{i \in N(c)} \{ \text{MLP}_1 [f_i; c_i - c] \} \right\} \quad (3)$$

where p denotes the PointNet-like module, c indicates the jittered center label, $N(c)$ is the index set of the k nearest neighbors of c , f_i is the global semantic feature, whose coordinate is c_i , and \max refers to the channel-wise max-pooling. We set L_1 as the mean square error between the ground-truth center offset and $p(c)$. Then for fully-supervised training, the detection loss L_2 is the same as the loss utilized in the original method. For weakly-supervised training, we utilize p to predict the center error in target domain and acquire refined center labels. We set L_3 as a

simpler version of L_2 which ignores the supervision for box sizes. More details about L_3 can be found in supplementary.

Secondly we analyze $L_{adv}(O, D)$. We conduct feature alignment in an adversarial manner: the discriminator predicts which domain the features belong to, and the detector aims to generate features that are hard to discriminate. The sign of gradients is flipped by a gradient reversal layer [10].

As the virtual scenes and real scenes are processed by the same network, we hope L_3 helps the network learn how to locate each object in real scenes, and L_2 compensates for the information loss of centers and sizes. However, due to the domain gap, L_2 will introduce domain-specific knowledge of the virtual scenes, which impairs the influence of L_3 . Besides, the center refinement module is trained only on source domain, which may not perform well on target domain. Therefore, we align the global semantic features and object proposal features with L_4 and L_5 respectively. Inspired by [34], the features are aligned with different intensities at different stages. For global semantic features, we use a PointNet to predict the domain label. Focal loss [17, 34] is utilized to apply weak alignment:

$$L_4 = - \sum_{i=1}^B (1 - p_i)^\gamma \log(p_i), \quad \gamma > 1 \quad (4)$$

where B is the batch size, and p_i refers to the probability of the global discriminator's predictions on the corresponding domain. Features with high p is easy to judge, which means they are domain-specific features and forcing invariance to them can hurt performance. So a small weight is used to reduce their impact on training. For object proposal features, they will be directly taken to predict the properties for bounding boxes. As the properties are domain-invariant and have real physical meaning, we strongly align this stage of features using an objectness weighted L2 loss:

$$L_5 = \sum_{i=1}^B \sum_{j=1}^N s_{ij} (1 - p_{ij})^2 \quad (5)$$

where B is the batch size, N is the number of proposals, s_{ij} refers to the objectness label and p_{ij} is the probability of the proposal discriminator's predictions on the corresponding domain. We detail the architectures of center refinement module and discriminators in supplementary.

4. Experiment

In this section, we conduct experiments to show the effectiveness of our BR approach. We first describe the datasets and experimental settings. Then we evaluate the generated virtual scenes and present the detection results of our method. We also design experiments to show the robustness of our virtual scene generation method and demonstrate the practicality of our approach. Finally we design

Table 2. Number of objects in each category in the training set and validation set of ScanNet, and average number of points of objects in each category in the real scenes and the virtual scenes.

| | Property | Bath-tub | Bed | Bench | Book-shelf | Bottle | Chair | Cup | Cur-tain | Desk | Door | Dresser | Key-board | Lamp | Laptop | Monitor | Night-stand | Plant | Sofa | Stool | Table | Toilet | Ward-robe |
|------------|---------------|----------|------|-------|------------|--------|-------|-----|----------|------|------|---------|-----------|------|--------|---------|-------------|-------|------|-------|-------|--------|-----------|
| # train | Object Number | 113 | 308 | 58 | 786 | 234 | 4357 | 132 | 408 | 551 | 2028 | 174 | 193 | 376 | 86 | 574 | 190 | 293 | 406 | 315 | 1526 | 201 | 98 |
| # validate | Object Number | 31 | 81 | 21 | 234 | 41 | 1368 | 34 | 95 | 127 | 467 | 43 | 53 | 83 | 25 | 191 | 34 | 50 | 97 | 51 | 407 | 58 | 19 |
| # real | Point Number | 2941 | 3905 | 1015 | 2679 | 101 | 726 | 66 | 2919 | 1525 | 1110 | 1274 | 74 | 272 | 173 | 370 | 700 | 792 | 2718 | 525 | 1282 | 1445 | 2762 |
| # virtual | Point Number | 6891 | 8683 | 4097 | 6258 | 162 | 2135 | 91 | 5495 | 5004 | 6048 | 2703 | 480 | 609 | 343 | 939 | 1088 | 1249 | 7250 | 1391 | 5421 | 3716 | 6105 |

Table 3. The class-specific detection results (mAP@0.25) of different weakly-supervised methods on ScanNet validation set. (FSB is the fully-supervised baseline. [†] indicates the method requires a small proportion of bounding boxes to refine the prediction. Other methods only use position-level annotations as supervision. We set best scores in bold, runner-ups underlined.)

| | Setting | batht. | bed | bench | bsf. | bot. | chair | cup | curt. | desk | door | dres. | keyb. | lamp | lapt. | monit. | n.s. | plant | sofa | stool | table | toil. | ward. | mAP@0.25 |
|------------------------|------------------------|-------------|-------------|-------------|-------------|------------|-------------|------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| VoteNet | FSB [25] | 66.8 | 86.2 | <u>24.4</u> | 55.6 | 0.0 | <u>88.3</u> | 0.0 | <u>48.5</u> | <u>62.8</u> | <u>45.8</u> | 24.1 | 0.1 | 47.2 | 5.2 | 62.1 | 73.2 | 13.4 | <u>88.7</u> | 35.1 | <u>62.6</u> | 94.6 | 7.8 | 45.1 |
| | WSB | 21.9 | 46.9 | 0.3 | 2.3 | 0.0 | 53.7 | 0.0 | 0.9 | 32.1 | 1.0 | 6.6 | 0.1 | 0.2 | 0.1 | 1.8 | 53.6 | 0.1 | 57.0 | 4.6 | 6.4 | 19.7 | 0.0 | 14.1 |
| | WS3D [†] [21] | 22.0 | 58.5 | 10.3 | 5.8 | 0.0 | 60.4 | 0.0 | 4.1 | 26.7 | 3.2 | 1.6 | 0.0 | 14.0 | 0.6 | 18.6 | 46.3 | 0.4 | 32.7 | 11.8 | 23.5 | 65.0 | 0.0 | 18.4 |
| | WSBP _P | 43.2 | 58.0 | 2.4 | 16.1 | 0.0 | 75.1 | 0.7 | 7.9 | 54.2 | 6.4 | 7.1 | 2.3 | 35.2 | 18.4 | 12.8 | 64.0 | 4.4 | 68.5 | 20.2 | 22.0 | 71.6 | 5.2 | 27.1 |
| | WSBP _M | 45.0 | 49.6 | 5.5 | 18.5 | 0.0 | 62.7 | 2.9 | 11.4 | 49.6 | 6.9 | 2.5 | 1.0 | 30.0 | 7.6 | 21.4 | 64.8 | 7.3 | 79.6 | 23.1 | 35.2 | 80.9 | 2.2 | 27.6 |
| | BR _P (Ours) | 51.2 | 73.0 | 16.4 | 27.1 | 0.1 | 70.3 | 0.0 | 8.3 | 44.5 | 7.3 | 16.0 | 1.5 | 40.2 | 7.7 | 42.1 | 50.8 | 7.4 | 67.1 | 10.7 | 39.0 | 88.4 | <u>18.1</u> | 31.2 |
| BR _M (Ours) | 57.1 | 80.4 | 14.3 | 31.7 | 0.0 | 77.4 | 0.0 | 13.2 | 49.7 | 11.3 | 14.8 | 1.0 | 43.5 | 6.0 | 56.5 | 65.0 | 10.6 | 80.2 | 26.9 | 44.2 | 91.4 | 6.5 | 35.5 | |
| GroupFree3D | FSB [20] | 86.2 | <u>87.5</u> | 16.3 | <u>49.6</u> | 0.6 | 92.5 | 0.0 | 70.9 | 78.5 | 53.5 | 56.0 | 6.4 | 68.2 | 11.5 | 81.5 | 88.5 | 15.2 | 88.2 | 45.6 | 65.0 | 99.7 | 31.2 | 54.2 |
| | WSB | 75.0 | 75.7 | 4.3 | 17.2 | 0.0 | 81.4 | 0.0 | 3.5 | 34.0 | 4.7 | 3.2 | 2.1 | 46.6 | 3.3 | 45.8 | 52.8 | 8.3 | 71.0 | 15.7 | 18.1 | 90.8 | 0.7 | 29.7 |
| | WS3D [†] [21] | 71.9 | 78.3 | 0.9 | 20.2 | 0.8 | 79.2 | 1.0 | 2.9 | 47.6 | 7.7 | 10.6 | 19.2 | 41.6 | 13.5 | 65.6 | 41.2 | 0.8 | 74.6 | 17.7 | 26.3 | 88.9 | 1.7 | 32.4 |
| | WSBP _P | 71.9 | 77.1 | 7.7 | 25.2 | 3.0 | 80.6 | 0.4 | 3.2 | 50.1 | 10.5 | 36.3 | 17.0 | 52.9 | 30.3 | 59.9 | 63.8 | 9.6 | 78.2 | 28.4 | 25.3 | 93.3 | 14.4 | 38.2 |
| | WSBP _M | 81.8 | 82.6 | 0.0 | 35.0 | 0.0 | 77.5 | 0.4 | 27.1 | 38.4 | 7.6 | 22.3 | 9.7 | 44.3 | 24.4 | 65.4 | 76.5 | 5.5 | 62.4 | 34.7 | 28.7 | 99.7 | 5.4 | 37.7 |
| | BR _P (Ours) | 72.3 | 73.5 | 45.8 | 27.7 | 0.0 | 77.2 | 8.2 | 30.8 | 35.0 | 17.8 | <u>51.7</u> | 0.3 | 64.2 | 25.0 | 63.5 | 66.6 | <u>23.8</u> | 86.7 | 33.9 | 37.6 | 98.3 | 5.2 | 43.0 |
| BR _M (Ours) | <u>85.3</u> | 90.9 | 8.8 | 34.3 | <u>1.9</u> | 80.0 | <u>7.7</u> | 24.7 | 58.0 | 20.8 | 45.4 | 31.3 | <u>64.4</u> | <u>25.8</u> | <u>67.5</u> | <u>76.7</u> | 27.3 | 91.4 | <u>43.3</u> | 46.7 | 94.8 | 8.3 | 47.1 | |

several ablation studies to verify our scene generation and domain adaptation method.

4.1. Experiments Setup

Datasets: We choose ModelNet40 [43] as the dataset of synthetic 3D shapes. ModelNet40 contains 12,311 synthetic CAD models from 40 categories, split into 9,843 for training and 2,468 for testing. We perform experiments on the ScanNet [7] dataset. ScanNet is a richly annotated dataset of indoor scenes with 1201 training scenes and 312 validation scenes. For each object appeared in the scenes, ScanNet officially provides its corresponding class in ModelNet40. Therefore we choose 22 categories of ModelNet40 which have more than 50 objects in ScanNet training set and 20 in the validation set, and report detection performance on them. Since ScanNet does not provide human-labeled bounding boxes, we predict axis-aligned bounding boxes and evaluate the prediction on validation set as in [20, 25, 44, 48]. We name this benchmark ScanNet-md40.

Compared to the 18-category setting in previous works [20, 25, 44], our ScanNet-md40 benchmark is more challenging. Apart from the categories of big objects (e.g. desk and bathtub), we also aim to detect relatively small objects, such as laptop, keyboard and monitor. Hence our benchmark can better evaluate the performance of both detectors and weakly-supervised learning methods.

Compared Methods: To illustrate the effect of our BR approach, the popular VoteNet [25] and state-of-the-art GroupFree3D [20] are selected as our detectors. We

compare BR with the following settings: 1) FSB: fully-supervised baseline, which serves as the upper bound of weakly-supervised methods; 2) WSB: weakly-supervised baseline, which trains the detector on real scenes by using L_3 only; 3) WS3D: another position-level weakly-supervised approach proposed in [21], which makes use of a number of precisely annotated bounding boxes; 4) WSBP: WSB pretrained on the virtual scenes. For settings which require the virtual scenes, we conduct experiments on two versions of virtual scenes (from points/meshes), which are distinguished by subscripts M and P respectively.

Implementation Details: We set $N = 10000$, $A_{min} = 0.1m^2$, $H_{min} = 0.1m$, $\Delta_h = 0.02m$, $k = 16$ and $\gamma = 3$. During training, as real scenes are more complicated, the converging of L_3 is much slower than L_2 . Therefore we multiple L_2 by 0.1 to slow down the training on virtual scenes and stabilize the process of feature alignment. To better train our center refinement module, the global semantic features should not change rapidly. Therefore we first train BR without L_1 until convergence, and then use the whole loss function to fine-tune the network. For GroupFree3D which has several decoders and each one outputs a stage of proposal features, we conduct feature alignment only for the last stage.

Different from previous works [20, 25], in our setting we need to detect small objects, such as bottle, cup and keyboard. As it is difficult for the network to extract high-quality features of these objects, we utilize an augmentation strategy to alleviate the problem, which is similar to [14].



Figure 5. The qualitative visualization results of our virtual scene generation. In (b), (c) and (d), the same color indicates the same object. Gray points are floors, walls and objects that we do not care. It can be seen that the virtual scenes preserve the coarse scene context and the supporting relationships between objects.

Please refer to the supplementary for more details.

4.2. Results and Analysis

Virtual Scene Evaluation: We first evaluate the statistics of the generated virtual scenes by computing the average number of points of objects in each category in real scenes and virtual scenes. As the input point clouds are downsampled to a given number before fed into the network, we only care about the ratio of average point numbers of objects in each category as the numbers can be controlled by the downsampling scale. We demonstrate the results in Table 2. It shows that the ratio in our virtual scenes is similar with that in the real scenes, which indicates the statistics of the virtual scenes are reasonable.

We also show qualitative visualizations to demonstrate our scene generation method in Figure 5. The virtual scenes generated with mesh information are named as mesh-version virtual scenes. Otherwise they are named as point-version virtual scenes. It is shown that the mesh-version virtual scenes can largely preserve the layout of the real scenes, and the point-version ones successfully combine the individual 3D shapes in a meaningful way.

3D Object Detection Results: As shown in Table 3, with position-level annotations only, WSB reduces the detection accuracy by a large margin in terms of mAP@0.25

compared to FSB. That’s mainly because WSB fails to learn the ability of predicting precise centers and sizes of bounding boxes according to the scene context. WS3D makes use of some box annotations and achieve better performance. However, as it is specially designed for outdoor 3D object detection, WS3D is still far from satisfactory when coping with the complicated indoor scenes. With pretraining on the virtual scenes, WSBP has more than 8% improvement over the WSB. That shows the ability of predicting precise bounding boxes learned in the source domain has been successfully transferred to the target domain. With our domain adaptation method to conduct better transferring, the improvement over the WSB is boosted to a higher level. The above results shows each step in BR is necessary: the virtual scenes are helpful to boost the detection performance, and the domain adaptation method can further explore the potential of the virtual scenes. Interestingly, as the virtual scenes become more realistic (from point-version to mesh-version), the performance of BR improves a lot while WSBP has little change, which indicates that layout may not be that important in pretraining as in domain adaptation.

In terms of class-specific results, on some categories the mAP@0.25 of the BR_M (for GroupFree3D) is even the highest among all the methods including the FSB. However, all methods fail to precisely detect cup and bottle, which

Table 4. The detection results (mAP@0.25) of BR under different error rate for center labeling on ScanNet. We adopt GroupFree3D as the detector and utilize mesh-version virtual scenes for BR.

| Method | Error Rate | | | | |
|------------------------|-------------|-------------|-------------|-------------|-------------|
| | 10% | 20% | 30% | 40% | 50% |
| WSB | 29.7 | 26.8 | 25.0 | 22.3 | 19.7 |
| BR _M (Ours) | 47.1 | 46.0 | 43.9 | 43.1 | 41.2 |

shows current 3D detectors still face huge challenges in small object detection. More detection results (mAP@0.5) can be found in supplementary.

Robustness for Labeling Error: In our labeling strategy, the center error is within 10%, which we define as the error rate, of the object’s size. To show the robustness of our approach, we gradually increase this rate from 10% to 50% by randomly jittering the centers according to the box sizes, and report the detection results of WSB and BR_M (for GroupFree3D) in terms of mAP@0.25. As shown in Table 4, with the increasing of error rate, the performance of BR degrades more slowly than WSB. Even if the error rate is 50%, which allows us to label the centers in a more time-saving strategy, BR can still achieve satisfactory results (higher than 0.41 in terms of mAP@0.25).

4.3. Ablation Study

We further design ablation experiments to study the influences of each scene generation step and each domain adaptation loss to the performance of our BR approach. In this section, we adopt VoteNet as the detector and use point-version virtual scenes for universality.

In Table 5, we illustrate that in our virtual scene generation pipeline, the physical constraints and density control are effective. As the virtual scenes become more realistic, the performance of our BR approach is getting better.

As shown in Table 6, we show the effect of each domain adaptation module and the center refine module. It can be seen that with global alignment or object proposal alignment, the detection performance can be boosted by 3.5% and 2.2% respectively. By combining the two kinds of feature alignments, we achieve higher detection accuracy. Having applied the center refinement method, the performance is further boosted by 1.0%.

4.4. Limitation

Due to the limited number of categories in ModelNet40, we selectively evaluate the performance of BR on 22 classes. However, as online repositories of user-generated 3D shapes, such as the 3D Warehouse repository [1], contain 3D shapes in almost any category, BR can be easily extended to 3D object detection on more classes once these online synthetic shapes are organized into a standard dataset. Therefore, ideally we can leverage a larger syn-

Table 5. The detection results (mAP@0.25) of BR with virtual scenes at different generation stages on ScanNet. Here the detector is VoteNet and the virtual scenes are point-version.

| Gravity Constrain | Collision Constrain | Density Control | mAP@0.25 |
|-------------------|---------------------|-----------------|-------------|
| | | | 26.3 |
| ✓ | | | 27.2 |
| ✓ | ✓ | | 28.5 |
| ✓ | ✓ | ✓ | 31.2 |

Table 6. The detection results (mAP@0.25) of BR with different domain adaptation modules on ScanNet. Here the detector is VoteNet and the virtual scenes are point-version.

| Global Alignment | Proposal Alignment | Center Refinement | mAP@0.25 |
|------------------|--------------------|-------------------|-------------|
| | | | 24.2 |
| ✓ | | | 28.7 |
| | ✓ | | 27.4 |
| ✓ | ✓ | | 30.2 |
| ✓ | ✓ | ✓ | 31.2 |

thetic 3D shape dataset, which covers all objects that may appear in indoor scenes. This dataset can promote more researches on 3D scene understanding with synthetic shapes, which we leave for future work.

5. Conclusion

In this paper, we have proposed a new label enhancement approach, namely Back to Reality (BR), for 3D object detection trained using only object centers and class tags as supervision. To fully exploit the information contained in the position-level annotations, we consider them as the coarse layout of scenes, which is utilized to assemble 3D shapes into fully-annotated virtual scenes. We apply physical constraints on the generated virtual scenes to make sure the relationship between objects is reasonable. In order to make use of the virtual scenes to remedy the information loss from box annotations to centers, we present a virtual-to-real domain adaptation method, which transfers the useful knowledge learned from the virtual scenes to real-scene 3D object detection. Experimental results on ScanNet dataset show the effectiveness of our BR approach.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62125603, and Grant U1813218, and in part by a grant from the Beijing Academy of Artificial Intelligence (BAAI).

References

- [1] Trimble 3d warehouse. [EB/OL]. <http://3dwarehouse.sketchup.com/>. 8
- [2] Armen Avetisyan, Manuel Dahnert, Angela Dai, Manolis Savva, Angel X Chang, and Matthias Nießner. Scan2cad: Learning cad model alignment in rgb-d scans. In *CVPR*, pages 2614–2623, 2019. 2
- [3] Armen Avetisyan, Angela Dai, and Matthias Nießner. End-to-end cad model retrieval and 9dof alignment in 3d scans. In *ICCV*, pages 2551–2560, 2019. 2
- [4] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *CVPR*, pages 2147–2156, 2016. 2
- [5] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *CVPR*, pages 1907–1915, 2017. 2
- [6] Manuel Dahnert, Angela Dai, Leonidas J Guibas, and Matthias Niessner. Joint embedding of 3d scan and cad objects. In *ICCV*, pages 8749–8758, 2019. 2
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 2, 6
- [8] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Hausser, Caner Hazirbas, Vladimir Golkov, Patrick Van Der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. In *ICCV*, pages 2758–2766, 2015. 2
- [9] Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient graph-based image segmentation. *IJCV*, 59(2):167–181, 2004. 3
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. 5
- [11] Ji Hou, Angela Dai, and Matthias Nießner. 3d-sis: 3d semantic instance segmentation of rgb-d scans. In *CVPR*, pages 4421–4430, 2019. 1, 2
- [12] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *CVPR*, pages 15587–15597, 2021. 2
- [13] Andrej Karpathy, Stephen Miller, and Li Fei-Fei. Object discovery in 3d scenes via shape analysis. In *ICRA*, pages 2088–2095, 2013. 3
- [14] Mate Kisantal, Zbigniew Wojna, Jakub Murawski, Jacek Naruniec, and Kyunghyun Cho. Augmentation for small object detection. *arXiv preprint arXiv:1902.07296*, 2019. 6
- [15] Jean Lahoud and Bernard Ghanem. 2d-driven 3d object detection in rgb-d images. In *ICCV*, pages 4622–4630, 2017. 2
- [16] Yangyan Li, Angela Dai, Leonidas Guibas, and Matthias Nießner. Database-assisted object retrieval for real-time 3d reconstruction. In *CGF*, volume 34, pages 435–446, 2015. 2
- [17] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. 5
- [18] Or Litany, Tal Remez, Daniel Freedman, Lior Shapira, Alex Bronstein, and Ran Gal. Asist: automatic semantically invariant scene transformation. *CVIU*, 157:284–299, 2017. 2
- [19] Xingyu Liu, Charles R. Qi, and Leonidas J. Guibas. FlowNet3d: Learning scene flow in 3d point clouds. In *CVPR*, pages 529–537, 2019. 2
- [20] Ze Liu, Zheng Zhang, Yue Cao, Han Hu, and Xin Tong. Group-free 3d object detection via transformers. *arXiv preprint arXiv:2104.00678*, 2021. 1, 2, 6
- [21] Qinghao Meng, Wenguan Wang, Tianfei Zhou, Jianbing Shen, Luc Van Gool, and Dengxin Dai. Weakly supervised 3d object detection from lidar point cloud. In *ECCV*, pages 515–531, 2020. 1, 2, 6
- [22] Qinghao Meng, Wenguan Wang, Tianfei Zhou, Jianbing Shen, Yunde Jia, and Luc Van Gool. Towards a weakly supervised framework for 3d point cloud object detection and annotation. *TPAMI*, 2021. 1, 2
- [23] Liangliang Nan, Ke Xie, and Andrei Sharf. A search-classify approach for cluttered indoor scene understanding. *TOG*, 31(6):1–10, 2012. 2
- [24] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, pages 523–540, 2020. 2
- [25] Charles R. Qi, Or Litany, Kaiming He, and Leonidas J. Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, pages 9277–9286, 2019. 1, 2, 6
- [26] 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*, pages 918–927, 2018. 2
- [27] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. 1, 2
- [28] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017. 1, 2
- [29] Zengyi Qin, Jinglu Wang, and Yan Lu. Weakly supervised 3d object detection from point clouds. In *ACM MM*, pages 4144–4152, 2020. 2
- [30] Yongming Rao, Benlin Liu, Yi Wei, Jiwen Lu, Cho-Jui Hsieh, and Jie Zhou. Randomrooms: Unsupervised pre-training from synthetic shapes and randomized layouts for 3d object detection. In *ICCV*, pages 3283–3292, 2021. 2
- [31] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016. 1
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: towards real-time object detection with region proposal networks. *TPAMI*, 39(6):1137–1149, 2016. 1
- [33] Zhongzheng Ren, Ishan Misra, Alexander G Schwing, and Rohit Girdhar. 3d spatial recognition without spatially labeled 3d. In *CVPR*, pages 13204–13213, 2021. 1, 2
- [34] Kuniaki Saito, Yoshitaka Ushiku, Tatsuya Harada, and Kate Saenko. Strong-weak distribution alignment for adaptive object detection. In *CVPR*, pages 6956–6965, 2019. 5
- [35] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, pages 770–779, 2019. 1, 2

- [36] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, pages 567–576, 2015. 1
- [37] Shuran Song and Jianxiong Xiao. Sliding shapes for 3d object detection in depth images. In *ECCV*, pages 634–651, 2014. 2
- [38] Shuran Song and Jianxiong Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *CVPR*, pages 808–816, 2016. 2
- [39] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017. 2
- [40] Mikaela Angelina Uy, Jingwei Huang, Minhyuk Sung, Tolga Birdal, and Leonidas Guibas. Deformation-aware 3d model embedding and retrieval. In *ECCV*, pages 397–413, 2020. 2
- [41] He Wang, Yezhen Cong, Or Litany, Yue Gao, and Leonidas J Guibas. 3dioumatch: Leveraging iou prediction for semi-supervised 3d object detection. In *CVPR*, pages 14615–14624, 2021. 2
- [42] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *CVPR*, pages 2642–2651, 2019. 2
- [43] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, pages 1912–1920, 2015. 6
- [44] Qian Xie, Yu-Kun Lai, Jing Wu, Zhoutao Wang, Yiming Zhang, Kai Xu, and Jun Wang. Mlcvnet: Multi-level context votenet for 3d object detection. In *CVPR*, pages 10447–10456, 2020. 6
- [45] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *ECCV*, pages 574–591, 2020. 2
- [46] Ning Xu, Yun-Peng Liu, and Xin Geng. Label enhancement for label distribution learning. *TKDE*, 2019. 2
- [47] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. *arXiv preprint arXiv:2101.02691*, 2021. 2
- [48] Zaiwei Zhang, Bo Sun, Haitao Yang, and Qixing Huang. H3dnet: 3d object detection using hybrid geometric primitives. In *ECCV*, pages 311–329, 2020. 6
- [49] Na Zhao, Tat-Seng Chua, and Gim Hee Lee. Sess: Self-ensembling semi-supervised 3d object detection. In *CVPR*, pages 11079–11087, 2020. 2
- [50] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 1
- [51] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, pages 4490–4499, 2018. 1, 2