

# NeRF-RPN: A general framework for object detection in NeRFs

Benran Hu<sup>1\*</sup> Junkai Huang<sup>1\*</sup> Yichen Liu<sup>1\*</sup> Yu-Wing Tai<sup>1,2</sup> Chi-Keung Tang<sup>1</sup>

<sup>1</sup>The Hong Kong University of Science and Technology <sup>2</sup>Kuaishou Technology

## Abstract

*This paper presents the first significant object detection framework, NeRF-RPN, which directly operates on NeRF. Given a pre-trained NeRF model, NeRF-RPN aims to detect all bounding boxes of objects in a scene. By exploiting a novel voxel representation that incorporates multi-scale 3D neural volumetric features, we demonstrate it is possible to regress the 3D bounding boxes of objects in NeRF directly without rendering the NeRF at any viewpoint. NeRF-RPN is a general framework and can be applied to detect objects without class labels. We experimented NeRF-RPN with various backbone architectures, RPN head designs and loss functions. All of them can be trained in an end-to-end manner to estimate high quality 3D bounding boxes. To facilitate future research in object detection for NeRF, we built a new benchmark dataset which consists of both synthetic and real-world data with careful labeling and clean up. Code and dataset are available at [https://github.com/lyc1yc52/NeRF\\_RPN](https://github.com/lyc1yc52/NeRF_RPN).*

## 1. Introduction

3D object detection is fundamental to important applications such as robotics and autonomous driving, which require scene understanding in 3D. Most existing relevant methods require 3D point clouds input or at least RGB-D images acquired from 3D sensors. Nevertheless, recent advances in Neural Radiance Fields (NeRF) [34] provide an effective alternative approach to extract highly semantic features of the underlying 3D scenes from 2D multi-view images. Inspired by Region Proposal Network (RPN) for 2D object detection, in this paper, we present the first 3D NeRF-RPN, which directly operates on the NeRF representation of a given 3D scene learned entirely from RGB images and camera poses. Specifically, given the radiance field and the density extracted from a NeRF model, our method produces bounding box proposals, which can be deployed in downstream tasks.

Recently, NeRF has provided very impressive results

\*Equal contribution. The order of authorship was determined alphabetically.

†This research is supported in part by the Research Grant Council of the Hong Kong SAR under grant no. 16201420.




Figure 1. **Region proposal results on a NeRF.** Top 12 proposals in eight orientations with highest confidence are visualized. The NeRF is trained from the *Living Room* scene from INRIA [38].

in novel view synthesis, while 3D object detection has become increasingly important in many real-world applications such as autonomous driving and augmented reality. Compared to 2D object detection, detection in 3D is more challenging due to the increased difficulty in data collection where various noises in 3D can be captured as well. Despite some good works, there is a lot of room for exploration in the field of 3D object detection. Image-based 3D object detectors either use a single image (e.g., [1, 4, 62]) or utilize multi-view consensus of multiple images (e.g., [29, 51, 63]). Although the latter use multi-view projective geometry to combine information in the 3D space, they still use 2D features to guide the pertinent 3D prediction. Some other 3D detectors based on point cloud representation (e.g., [31, 33, 41, 73]) heavily rely on accurate data captured by sensors. To our knowledge, there is still no representative work on direct 3D object detection in NeRF.

Thus, we propose NeRF-RPN to propose 3D ROIs in a given NeRF representation. Specifically, the network takes as input the 3D volumetric information extracted from NeRF, and directly outputs 3D bounding boxes of ROIs. NeRF-RPN will thus be a powerful tool for 3D object detection in NeRF by adopting the “3D-to-3D learning” paradigm, taking full advantages of 3D information inherent in NeRF and predicting 3D region proposals directly in 3D space.

As the first significant attempt to perform 3D object detection directly in NeRFs trained from multi-view images, this paper’s focus contributions consist of:

- First significant attempt on introducing RPN to NeRF

for 3D objection detection and related tasks.

- A large-scale public indoor NeRF dataset for 3D object detection, based on the existing synthetic indoor dataset Hypersim [46] and 3D-FRONT [11], and real indoor dataset ScanNet [6] and SceneNN [19], carefully curated for NeRF training.
- Implementation and comparisons of NeRF-RPNs on various backbone networks, detection heads and loss functions. Our model can be trained in 4 hrs using 2 NVIDIA RTX3090 GPUs. At runtime, it can process a given NeRF scene in 115 ms (excluding postprocessing) while achieving a 99% recall on our 3D-FRONT NeRF dataset.
- Demonstration of 3D object detection over NeRF and related applications based on our NeRF-RPN.

## 2. Related Work

### 2.1. NeRF

Neural radiance field (NeRF) [34] has become the mainstream approach for novel view reconstruction, which models the geometry and appearance of a given scene in a continuous and implicit radiance field parameterized by an MLP. Following this work, instant neural graphics primitive [36] applies hash encoding to reduce the training time dramatically. PlenOctrees [67] uses an octree-based radiance field and a grid of spherical basis functions to accelerate rendering and appearance decoding. TensoRF [3] projects a 3D point onto three 2D planes to encode the positional information. Although these works use different approaches to model structures, they achieve the same goal of taking as input xyz coordinates and 3D camera poses to generate view-dependent RGB color and volume density at each position to render the images from a given view point. NeRF not only provides structural details of a 3D scene but is also conducive to 3D training, where only posed RGB images are required, thus making this representation also suitable for 3D object detection.

### 2.2. Object Detection and Region Proposal Network

Subsequent to [20] and recent GPU advances, deep convolutional neural network (CNN) has become the mainstream approach for object detection given single images. Object detection based on deep learning can be divided into anchor-based methods and anchor-free methods. Anchor-based methods, including two-stage methods [12, 13, 16, 17, 45] and one-stage methods [10, 23, 25–27, 44, 68], first generate a large number of preset anchors with different sizes and aspect ratios on a given image, then predict the labels and box regression offsets of these anchors. For two-stage methods, relatively coarse region proposals are first generated from anchors, followed by refining such coarse proposals and inferring the labels. In contrast to anchor-based methods, anchor-free methods [8, 9, 22, 28, 57, 71, 72] predict on feature maps directly.

Region Proposal Network (RPN) was first introduced in [45] to propose regions in an image that may contain objects for subsequent refinement. RPN uses shared convolutional layers to slide through local regions on the feature maps from feature extraction layers and feeds the transformed features into a box-regression head and a box-classification head in parallel. In [45], RPN is applied on the feature map from the last shared convolution layer only, whereas more recent works such as Feature Pyramid Networks (FPN) [24] utilize multi-scale feature maps. Our proposed method adapts the idea of sliding window from 2D RPN and also utilizes FPN in a 3D fashion.

### 2.3. 3D Object Detection

Based on the input form, current 3D object detectors can be categorized as point cloud-based and RGB-based methods. Many point cloud-based methods first transform point clouds into voxel forms to subsequently operate on the 3D feature volume through convolution [15, 33, 49, 56, 73] or Transformers [33, 59]. However, the large memory footprint of the voxel representation constrains the resolution used. While sparse convolution [14] and 2D projection have been adopted to alleviate the issue, works directly operating on raw point clouds have been proposed recently [21, 31, 35, 40–42, 53, 60]. Most of them partition points into groups and apply classification and bounding box regression to each group. Criteria used for grouping include 3D frustums extruded from 2D detection [41], 3D region proposals [52, 53], and voting [39, 40, 60]. GroupFree3D [31] and Pointformer [37], on the other hand, use Transformers to attend over all points instead of grouping.

3D objection detection on single images or posed multi-view RGB images is more challenging and relatively less explored. Early attempts in monocular 3D objection detection first estimate the per-pixel depths [4, 65], pseudo-LiDAR signals [43, 62, 66], or voxel information [47] from an RGB image, performing detection on the reconstructed 3D features. Later works have extended 2D object detection methods to operate in 3D. For instance, M3D-RPN [1] and MonoDIS [54] use 2D anchors for detection and predict a 2D to 3D transformation. FCOS3D [61] extends FCOS [58] to predict 3D information. More research recently has been focused on the multi-view case. ImVoxelNet [51] projects 2D features back to a 3D grid and applies a voxel-based detector on it. DETR3D [63] and PETR [29] adopt similar designs as DETR [2], both trying to fuse 2D features and 3D position information. Although these image-based methods can assist the region proposal task in NeRF, they do not utilize the inherent 3D information from NeRF and are thus limited in their accuracy.

While it is possible to sample from NeRF and produce a voxel or point cloud representation on which previous 3D object detection methods can be applied, such conversion can be rather ad-hoc, depending on both the NeRF structure and the reconstruction quality. Noise and poor fine-level




Figure 2. **NeRF-RPN.** Our method first samples a grid of points in NeRF and extracts their RGB and density. The extracted volumetric features are then passed through a 3D backbone network to get deep multi-scale 3D features, which are fused with a 3D FPN and fed to the 3D RPN head to produce the region proposals.

geometry approximation in these converted representations also pose challenges to existing 3D object detectors. Note that unlike point cloud samples, which cover only the surface (crust) of objects, the density in NeRF distributes over the interior as well. Clearly, existing methods fail to utilize this important solid object information, which is adequately taken into account by our NeRF-RPN. Besides, there exists no 3D object detection dataset tailored for the NeRF representation, which also limits the advancement of 3D object detection in NeRF.

### 3. Method

Similar to the original RPN, our method has two major components, see Figure 2. The first consists of a feature extractor that takes a grid of radiance and density sampled from NeRF as input, and produces a feature pyramid as output. The second is an RPN head which operates on the feature pyramid and generates object proposals. The feature volumes corresponding to the proposals can subsequently be extracted and processed for any downstream tasks. Our method is flexible in the form of NeRF input features, the feature extractor architectures, and the RPN module, which can be adapted to multiple downstream tasks.

#### 3.1. Input Sampling from NeRF

Our method assumes a fully-trained NeRF model with reasonable quality model is provided. The first step is to uniformly sample its radiance and density field to construct a feature volume. Despite different NeRF variants exist with different radiance field representations or structures, they share the same property that radiance and density can be queried with view directions and spatial locations. As essentially the radiance and density go through a similar volumetric rendering process, our method uses them as the input so that it is agnostic to concrete NeRF structures.

We uniformly sample the radiance and density on a 3D grid that covers the full traceable volume of the NeRF model, which is determined by slightly enlarging the bounding box containing all the cameras and objects in the scene. The resolution of the grid in each dimension is proportional to the length of the traceable volume in

that dimension so that the aspect ratio of the objects are maintained. For NeRF models using plain RGB as radiance representation, we sample from the same set of viewing directions according to the camera poses used in NeRF training and average the results. If such camera poses are unknown, we uniformly sample directions from a sphere. Generally, each sample can be written as  $(r, g, b, \alpha)$ , where  $(r, g, b)$  is the averaged radiance and  $\alpha$  is converted from the density  $\sigma$ :

$$\alpha = \text{clip}(1 - \exp(-\sigma\delta), 0, 1), \quad (1)$$

where  $\delta = 0.01$  is a preset distance. For NeRF models adopting spherical harmonics or other basis functions as radiance representation, either the computed RGB values or the coefficients of the basis functions can be used as radiance information, depending on the downstream task.

#### 3.2. Feature Extractors

Given the sampled grid, the feature extractor will generate a feature pyramid. We adopt three backbones: VGG [55], ResNet [18] and Swin Transformer [30] in our experiments, but other backbone networks may also be applicable. Considering the large variation in object sizes for indoor NeRF scenes as well as the scale differences between different NeRF scenes, we incorporate an FPN [24] structure to generate multi-scale feature volumes and to infuse high-level semantics into higher resolution feature volumes. For VGG, ResNet, and the FPN layers, we replace all the 2D convolutions, poolings, and normalization layers with their 3D counterparts. For Swin Transformer, we correspondingly employ 3D position embedding and shifted windows.

#### 3.3. 3D Region Proposal Networks

Our 3D Region Proposal Network takes the feature pyramid from the feature extractor and outputs a set of oriented bounding boxes (OBB) with corresponding objectness scores. As in most 3D object detection works, we constrain the bounding box rotation to  $z$ -axis only (yaw angle), which is aligned with the gravity vector and

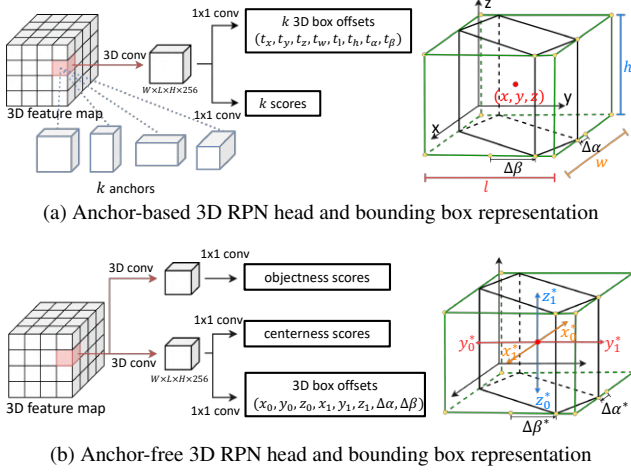


Figure 3. **3D RPN Head.** These two figures illustrate the architectures of anchor-based and anchor-free 3D RPN heads along with their 3D midpoint offset bounding box representations.

perpendicular to the ground. We experiment two types of region proposal methods: anchor-based and anchor-free method, see Figure 3.

**Anchor-Based RPNs** Conventional RPNs as originally proposed in Faster R-CNN [45] place anchors of different sizes and aspect-ratios at each pixel location and predict objectness scores and bounding box regression offsets for each anchor. We extend this approach to 3D by placing 3D anchors of different aspect-ratios and scales in voxels on different levels of the feature pyramid. We add  $k$  levels of 3D convolutional layers after the feature pyramids (typically  $k = 2$  or  $4$ ), on top of which two separate  $1 \times 1 \times 1$  3D convolutional layers are used to predict the probability  $p$  that an object exists, and the bounding box offsets  $\mathbf{t}$  for each anchor, see Figure 3(a). These layers are shared between different levels of the feature pyramid to reduce the number of parameters and improve the robustness to scale variation. The bounding box offsets  $\mathbf{t} = (t_x, t_y, t_z, t_w, t_l, t_h, t_\alpha, t_\beta)$  are parametrized similarly as in [64] but extended with a new dimension:

$$\begin{aligned} t_x &= (x - x_a)/w_a, & t_y &= (y - y_a)/l_a, \\ t_z &= (z - z_a)/h_a, & t_w &= \log(w/w_a), \\ t_l &= \log(l/l_a), & t_h &= \log(h/h_a), \\ t_\alpha &= \Delta\alpha/w, & t_\beta &= \Delta\beta/l, \end{aligned} \quad (2)$$

where  $x, y, w, l, \Delta\alpha, \Delta\beta$  describe the OBB projected onto the  $xy$ -plane, and  $z, h$  are for the additional dimension in height.  $x_a, y_a, z_a, w_a, l_a, h_a$  give the position and size of the reference anchor, see Figure 3(a). Note that this encoding does not guarantee the decoded OBBs are cuboids. We follow [64] to transform the projections into rectangles before using them as proposals.

To determine the label of each anchor, we follow the process in Faster R-CNN but with parameters adapted under the 3D setting: we assign a positive label to an anchor

if it has an Intersection-over-Union (IoU) overlap greater than 0.35 with any of the ground-truth boxes, or if it has the highest IoU overlap among all anchors with one of the ground-truth box. Non-positive anchors with IoU below 0.2 for all ground-truth boxes are regarded negative. Anchors that are neither positive nor negative are ignored in the loss computation. The loss is similar to that in Faster R-CNN:

$$\begin{aligned} L(\{p_i\}, \{\mathbf{t}_i\}) &= \frac{1}{N_{cls}} \sum_i L_{cls}(p_i, p_i^*) \\ &+ \frac{\lambda}{N_{reg}} \sum_i p_i^* L_{reg}(\mathbf{t}_i, \mathbf{t}_i^*), \end{aligned} \quad (3)$$

where  $p_i, \mathbf{t}_i$  are predicted objectness and box offsets,  $p_i^*, \mathbf{t}_i^*$  are ground-truth targets,  $N_{cls}, N_{reg}$  are the number of anchors involved in loss computation, and  $\lambda$  is a balancing factor between the two losses.  $L_{cls}$  is the binary cross entropy loss and  $L_{reg}$  is the smooth  $L_1$  loss in [12]. The regression loss is only computed for positive anchors.

**Anchor-Free RPNs** Anchor-free object detectors discard the expensive IoU computation between anchors and ground-truth boxes and can be used for region proposal in specific problem scopes (e.g., figure-ground segmentation). We choose FCOS which is a representative anchor-free method and extend it to 3D.

Unlike anchor-based methods, our FCOS-based RPN predicts a single objectness  $p$ , a set of bounding box offsets  $\mathbf{t} = (x_0, y_0, z_0, x_1, y_1, z_1, \Delta\alpha, \Delta\beta)$ , and a centerness score  $c$  for each voxel, see Figure 3(b). We extend the encoding of box offsets in FCOS and define the regression target  $\mathbf{t}_i^* = (x_0^*, y_0^*, z_0^*, x_1^*, y_1^*, z_1^*, \Delta\alpha^*, \Delta\beta^*)$  as following:

$$\begin{aligned} x_0^* &= x - x_0^{(i)}, & x_1^* &= x_1^{(i)} - x, \\ y_0^* &= y - y_0^{(i)}, & y_1^* &= y_1^{(i)} - y, \\ z_0^* &= z - z_0^{(i)}, & z_1^* &= z_1^{(i)} - z, \\ \Delta\alpha^* &= v_x^{(i)} - x, & \Delta\beta^* &= v_y^{(i)} - y, \end{aligned} \quad (4)$$

where  $x, y, z$  are the voxel position,  $x_0^{(i)} < x_1^{(i)}$  are the left and right boundary of the **axis-aligned bounding box** (AABB) of  $i$ -th ground-truth OBB, and likewise for  $y_0^{(i)}, y_1^{(i)}, z_0^{(i)}, z_1^{(i)}$ .  $v_x^{(i)}$  denotes the  $x$  coordinate of the upmost vertex in the  $xy$ -plane projection of the OBB, and  $v_y^{(i)}$  is the  $y$  coordinate of the rightmost vertex, see Figure 3(b). The ground-truth centerness is given by:

$$c^* = \sqrt{\frac{\min(x_0^*, x_1^*)}{\max(x_0^*, x_1^*)} \times \frac{\min(y_0^*, y_1^*)}{\max(y_0^*, y_1^*)} \times \frac{\min(z_0^*, z_1^*)}{\max(z_0^*, z_1^*)}}. \quad (5)$$

The overall loss is then given by

$$\begin{aligned} L(\{p_i\}, \{\mathbf{t}_i\}, \{c_i\}) &= \frac{1}{N_{pos}} L_{cls}(p_i, p_i^*) \\ &+ \frac{\lambda}{N_{pos}} p_i^* L_{reg}(\mathbf{t}_i, \mathbf{t}_i^*) + \frac{1}{N_{pos}} p_i^* L_{ctr}(c_i, c_i^*), \end{aligned} \quad (6)$$






Figure 4. **Binary Classification Network.** The binary classification network architecture with rotated 3D ROI pooling along with the bounding box representation used in this network.

where  $L_{cls}$  is the focal loss in [25] and  $L_{reg}$  is the IoU loss for rotated boxes in [70].  $L_{ctr}$  is the binary cross entropy loss.  $p_i^* \in \{0, 1\}$  is the ground-truth label of each voxel in the feature pyramid, which is determined using the same center sampling and multi-level prediction process as in [58];  $\lambda$  is the balancing factor and  $N_{pos}$  is the number of voxels with  $p_i^* = 1$ . The regression and centerness loss only account for positive voxels.

To learn  $p, t, c$ , we adapt the network in FCOS by adding  $k = 2$  or  $4$  3D convolutional layers independently for the classification and regression branch after the feature pyramid. We append a convolutional layer on top of the classification and regression branch, respectively, to output  $p$  and  $t$ , and a parallel convolutional layer on the regression branch for predicting  $c$ . Like in our anchor-based method, we transform the possibly skewed box predictions into cuboids before further post-processing.

### 3.4. Additional Loss Functions

**Objectness Classification** Although NeRF-RPN mainly targets on high recalls, some downstream tasks may prefer a low false-positive rate as well. To improve the precision of ROIs, we add a binary classification network as a sub-component to achieve foreground/background classification. More specifically, the network takes 1) the ROIs from RPN, and 2) the feature pyramid from the feature extractor as input, and outputs an objectness score and bounding box refinement offsets for each ROI, see Figure 4. We extract rotation invariant features for each proposal via rotated ROI pooling. Each proposal is parameterized by  $(x_r, y_r, z_r, w_r, l_r, h_r, \theta_r)$ , where  $(x_r, y_r, z_r)$  describes the center coordinate,  $w_r, l_r, h_r$  are the three dimensions, and  $\theta_r \in [-\frac{\pi}{2}, \frac{\pi}{2}]$  is the yaw angle. Referring to [64], we first enlarge the box and locate it in the corresponding feature volume, then apply trilinear interpolation to calculate the value on each feature point, and pad the ROI feature volume with zero before forwarding it to a pooling layer. The feature volume is pooled to  $N \times 3 \times 3 \times 3$  and used for further regression and classification. Referring to [7], the bounding box offset  $\mathbf{g} = (g_x, g_y, g_z, g_w, g_l, g_h, g_\theta)$  is defined as

$$\begin{aligned} g_x &= ((x - x_r) \cos \theta_r + (y - y_r) \sin \theta_r) / w_r \\ g_y &= ((y - y_r) \cos \theta_r - (x - x_r) \sin \theta_r) / l_r \\ g_z &= (z - z_r) / h_r, & g_w &= \log(w / w_r), \\ g_l &= \log(l / l_r), & g_h &= \log(h / h_r), \\ g_\theta &= (\theta - \theta_r) / 2\pi \end{aligned} \quad (7)$$

The classification layer estimates the probability over 2 classes (namely, *non-object* and *object* class). ROIs with IoU overlap greater than 0.25 with any of the ground-truth boxes are labeled as *object*, while all the others are labeled *non-object*. The loss function is similar to Equation 3, where the box offsets are replaced by  $\mathbf{g}, \mathbf{g}^*$ .

**2D Projection Loss** We project 3D bounding box coordinates  $b_i = (x_i, y_i, z_i)$  into 2D  $b'_i = (x'_i, y'_i)$  and construct a 2D projection loss as following:

$$L_{2d \text{ proj}}(\{b'_i\}) = \frac{1}{N_{cam} N_{box}} L_{reg}(b'_i, b_i^*), \quad (8)$$

where  $N_{cam}, N_{box}$  are the number of cameras and the number of proposals. We set 4 cameras at 4 top corners of the room, pointing to the room center. Refer to the supplementary material for more discussion.

## 4. NeRF Dataset for 3D Object Detection

There has been no representative NeRF dataset constructed for 3D object detection. Thus, we build the first NeRF dataset for 3D object detection utilizing Hypersim [46] and 3D-FRONT [11] datasets. In addition to these synthetic datasets, we incorporate a subset of the real-world datasets from SceneNN [19] and ScanNet [6] to demonstrate that our method is robust to real-world data. Figure 5 shows some selected examples of the 3D groundtruth boxes we carefully labeled from 3D-FRONT. Table 1 summarizes our dataset.

Datasets	# Scenes	# Boxes					
		Total #	Average # (per scene)	# Boxes in size (# voxels)			
				$< 16^3$	$16^3 \sim 32^3$	$32^3 \sim 64^3$	$> 64^3$
Hypersim	250	4798	19.2	3836	770	184	8
3D-FRONT	159	1191	7.5	129	703	324	35
ScanNet	90	1086	12.1	508	488	88	2
SceneNN	16	367	22.9	182	112	54	19

Table 1. **Statistics of our NeRF dataset for 3D Object Detection.**

**Hypersim** Hypersim is a very realistic synthetic dataset for indoor scene understanding containing a wide variety of rendered objects with 3D semantics. However, the dataset is not specifically designed for NeRF training, where the object annotations provided are noisy for direct use in region proposal tasks. Thus, we perform extensive cleaning based on both the NeRF reconstruction quality and the usability of object annotations (see supp mtrl). Finally we keep around 250 scenes after cleanup. The original 3D object bounding boxes in Hypersim are not carefully pruned, as some objects are invisible in all images. Furthermore, many instances are too fine in scale, while some are of less or little interest, e.g., floors and windows. We remove ambiguous objects that may interfere our training. Then, we filter out tiny or thin objects by checking if the smallest dimension of their AABB is below a certain threshold. After these automatic pre-processing, we manually examine each remaining object. Objects that are visible in less than three images, or with over half of their AABBs invisible in all images, are removed.



Figure 5. **3D-FRONT NeRF Dataset Samples.** Rows 1–2 show the NeRF reconstruction quality and ground-truth bounding box quality of our 3D-FRONT NeRF dataset. Rows 3–4 show groundtruth boxes with diverse object appearance in the dataset.

**3D-FRONT** 3D-FRONT [11] is a large-scale synthetic indoor scene dataset with room layouts and textured furniture models. Due to its size, effort has been spent on splitting complex scenes into individual rooms and cleaning up bounding boxes (supp mtrl). A total of 159 usable rooms are manually selected, cleaned, and rendered in our dataset. More rooms can be generated for NeRF training using our code and 3D-FRONT dataset, which will be released when the paper is accepted for publication. We perform extensive manual cleaning on the bounding boxes in each room. Similar to Hypersim, bounding boxes for construction objects such as ceilings and floors are removed automatically based on their labels. Moreover, we manually merge the relevant parts bounding boxes to label the entire semantic object (e.g., seat, back panel and legs are merged into a chair box). Refer to Figure 5 for examples.

**Real-World Dataset** We construct our real-world NeRF dataset leveraging ScanNet [6], SceneNN [19], and a dataset from INRIA [38]. ScanNet is a commonly used real-world dataset for indoor 3D object detection which contains over 1,500 scans. We randomly select 90 scenes and for each scene, we uniformly divide the video frames into 100 bins and select the sharpest frame in each bin based on the variance of Laplacian. We use the provided depth and a depth-guided NeRF [48] to train the models. For object annotations, we compute the minimum bounding boxes based on the annotated meshes and discard objects of

certain classes and sizes as similarly done for Hypersim.

## 5. Experiments

### 5.1. Training & Testing

**Training** During training, input scenes are randomly flipped along  $x, y$  axes and rotated along  $z$ -axis by  $\frac{\pi}{2}$  with probability 0.5 for each augmentation operation. Additionally, the scenes are slightly rotated along  $z$ -axis by  $\alpha \in [-\frac{\pi}{18}, \frac{\pi}{18}]$  with a probability of 0.5, which we find can significantly improve the average precision (AP) in RPN outputs. We optimize our network with AdamW [32] with an initial learning rate of 0.0003 and a weight decay of 0.001. In our training, we set  $\lambda = 5.0$  for Eq. (3) and  $\lambda = 1.0$  for Eq. (6). For the anchor-based approach, we adopt a 4-level FPN and anchors of 13 different aspect ratios, which are 1:1:1, 1:1:2, 1:1:3, 2:2:1, 3:3:1, and their permutations. All anchors on the same level of feature volume share the same size for their shortest side, which is in  $\{8, 16, 32, 64\}$ , from fine to coarse scale. Following the RPN training strategy in [45], we randomly sample 256 anchors from each scene in each iteration to compute the loss, where the ratio of positive and negative anchors is 1:1. For anchor-free approach, all output proposals are used to compute the loss.

**Testing** After obtaining the ROIs with objectness scores, we first discard the boxes whose geometry centers are beyond the scene boundary. Then, we select the top 2,500 proposals on each level of the feature volumes independently. To remove redundant proposals, we apply Non-Maximum Suppression (NMS) to the aggregated boxes based on rotated-IoU with threshold 0.1, after which we select the 2,500 boxes with the highest objectness scores.

### 5.2. Ablation Study

**Backbones and Heads** Table 2 tabulates the recall and average precision of different combinations of feature extraction backbones and RPN heads. When fixing the backbones and comparing the RPN heads only, we observe that anchor-free models achieve a higher AP on all three datasets. The two RPN methods attain similar recalls on 3D-FRONT and ScanNet, while on Hypersim anchor-free models are generally higher in recalls. We believe the better performance of anchor-free models results are twofold: 1) The centerness prediction of anchor-free models helps suppress proposals that are off from the centers, which is particularly helpful when the bounding box center is misaligned with the mass center, or when the NeRF input is noisy; 2) The limited number of aspect ratios and scales for anchors limits the performance of anchor-based models as 3D objects vary greatly in sizes.

Furthermore, when comparing the performance between different backbones, we notice that models with VGG19 generally achieve better recall and AP compared to others.



Methods	Backbones	Hypersim				3D-FRONT				ScanNet			
		Recall <sub>25</sub>	Recall <sub>50</sub>	AP <sub>25</sub>	AP <sub>50</sub>	Recall <sub>25</sub>	Recall <sub>50</sub>	AP <sub>25</sub>	AP <sub>50</sub>	Recall <sub>25</sub>	Recall <sub>50</sub>	AP <sub>25</sub>	AP <sub>50</sub>
Anchor-based	VGG19	57.1	14.9	11.2	1.3	97.8	<b>76.5</b>	65.9	43.2	88.7	42.4	40.7	14.4
	ResNet-50	49.8	13.0	9.7	1.3	96.3	70.6	65.7	45.1	86.2	32.0	34.4	9.0
	Swin-S	69.8	<b>28.3</b>	24.6	6.2	<b>98.5</b>	63.2	51.8	26.6	<b>93.6</b>	<b>44.3</b>	38.7	12.9
Anchor-free	VGG19	66.7	27.3	<b>30.9</b>	<b>11.5</b>	96.3	69.9	<b>85.2</b>	<b>59.9</b>	89.2	42.9	55.5	18.4
	ResNet-50	63.2	17.5	23.2	6.0	95.6	67.7	83.9	55.6	91.6	35.5	55.7	16.1
	Swin-S	<b>70.8</b>	21.0	27.7	7.7	96.3	62.5	78.7	41.0	90.6	39.9	<b>57.5</b>	<b>20.5</b>

Table 2. **Ablation on different backbones and heads.** Recall<sub>25</sub> and Recall<sub>50</sub> denote the recall scores at an IoU threshold of 0.25 and 0.5, respectively.

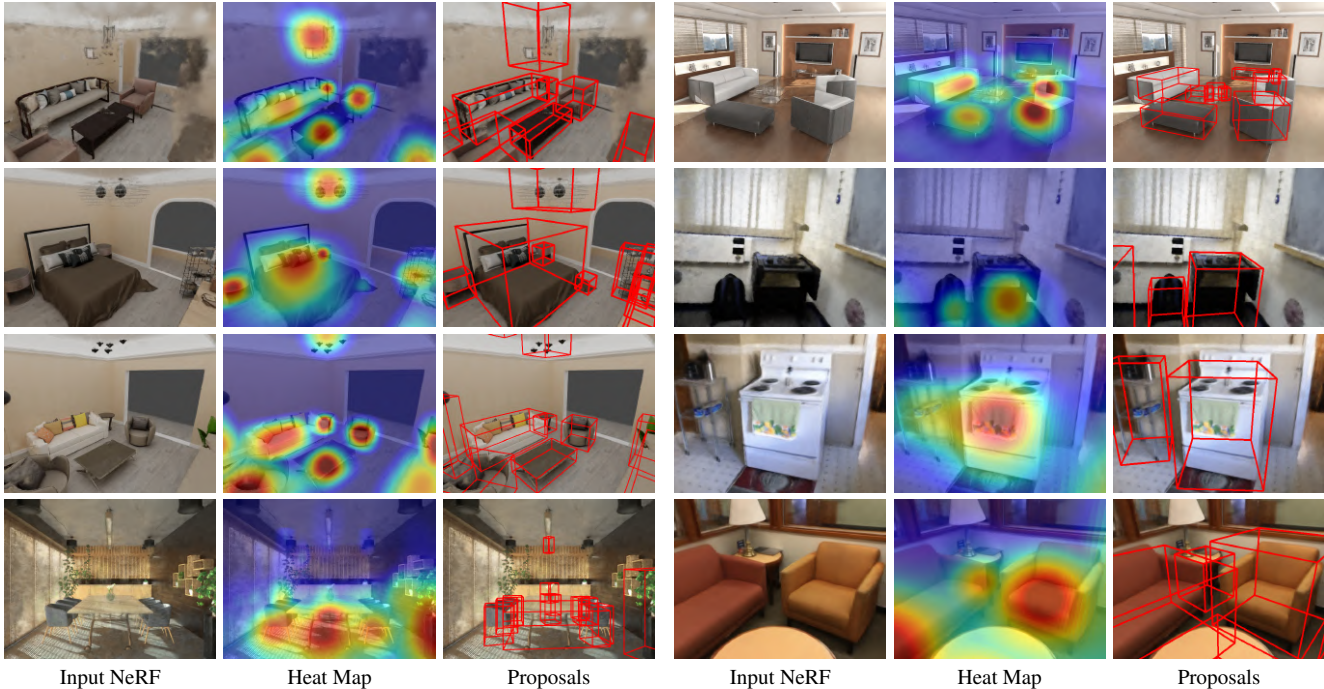


Figure 6. **Qualitative Results.** The “Heat Map” columns show the distribution of proposal confidence scores where red means higher confidence. The “Proposals” columns show a few top bounding boxes after NMS. From top to bottom, left to right, scene 1-3 are from 3D-FRONT, 4-5 from Hypersim, and 6-8 from ScanNet.




Figure 7. **Failure Cases.** (a)(b) Missing and merging proposals, (c) wrong rotation, (d) no proposal for tiny/second-level objects.

The major exception concerns the performance of anchor-based models on Hypersim, where Swin-S demonstrates superior recall and AP. Given that the NeRF results on Hypersim are significantly noisier and the scenes are more complex, we suspect that the larger receptive fields and the richer semantics enabled by the shifted windows, and attention of Swin Transformers are crucial to our anchor-based method in this case.

**NeRF Sampling Strategies** While the density field from NeRF is view-independent, the radiance depends on the viewing direction and can be encoded with different schemes. In the supplemental material we investigate the

effect of this view-dependent information, and conclude that using density alone is the best strategy.

**Regression Loss** We test three common loss functions for bounding box regression on the 3D-FRONT dataset using Swin-S as the backbone in Table 4. IoU loss directly optimizes the IoU between the predicted and ground-truth bounding boxes while DIOU loss [69] penalizes the normalized distance between the two for faster convergence. We use the variants for oriented boxes of these two losses as proposed in [70]. Our results illustrate that IoU loss consistently outperforms the other two for the anchor-based approach, while for anchor-free models IoU and DIOU loss produce similar performance.

**Additional Loss** We discuss the impact of aforementioned losses in the supplemental material.

### 5.3. Results

We performed experiments with different model configurations on various NeRF datasets constructed from

Methods	Hypersim				3D-FRONT				ScanNet			
	AR <sub>25</sub>	AR <sub>50</sub>	AP <sub>25</sub>	AP <sub>50</sub>	AR <sub>25</sub>	AR <sub>50</sub>	AP <sub>25</sub>	AP <sub>50</sub>	AR <sub>25</sub>	AR <sub>50</sub>	AP <sub>25</sub>	AP <sub>50</sub>
ImVoxelNet	19.7	5.7	9.7	2.3	88.3	71.5	<b>86.1</b>	<b>66.4</b>	51.7	20.2	37.3	9.8
FCAF3D	47.6	19.4	30.7	8.8	89.1	56.9	73.1	35.2	<b>90.2</b>	42.4	<b>63.7</b>	<b>18.5</b>
Ours (anchor-based)	57.1	14.9	11.2	1.3	<b>97.8</b>	<b>76.5</b>	65.9	43.2	88.7	42.4	40.7	14.4
Ours (anchor-free)	<b>66.7</b>	<b>27.3</b>	<b>30.9</b>	<b>11.5</b>	96.3	69.9	85.2	59.9	89.2	<b>42.9</b>	55.5	18.4

Table 3. **Quantitative Comparison.** Our results are reported on the VGG19 backbone. AR refers to the recall score at the specified IoU threshold, instead of the average recall.

Methods	Loss	Recall		AP	
		0.25	0.50	0.25	0.50
Anchor-based	Smooth $L_1$	<b>98.5</b>	63.2	51.8	26.6
	IoU	<b>98.5</b>	<b>71.3</b>	<b>61.6</b>	<b>36.7</b>
	DIoU	97.1	<b>71.3</b>	59.5	32.8
Anchor-free	Smooth $L_1$	96.3	56.6	76.5	39.9
	IoU	96.3	62.5	<b>78.7</b>	<b>41.0</b>
	DIoU	<b>97.1</b>	<b>64.0</b>	77.4	40.2

Table 4. **Ablation results of the bounding box regression loss.**



Figure 8. **Application: Scene Editing.** Removing an object in a bounding box proposed by our NeRF-RPN.

Hypersim [46], 3D-FRONT [11], ScanNet [6], SceneNN [19] and INRIA [38], where [19] and [38] are only used in test time due to their relatively small numbers of usable scenes. Detailed quantitative results are shown in Table 2. Figure 6 shows the qualitative results produced by the model with VGG19 backbone and anchor-free RPN head. Figure 7 shows typical failure cases. During our experiments, we found that bad NeRF reconstruction can severely hamper the prediction. As aforementioned, the region proposal task largely depends on 3D geometry in NeRF. Similar to 2D RPN for images, our method also has missing/merging proposals or wrong rotation after NMS. Presently, our dataset handles first-level objects; tiny or second-level objects are future work.

**Application: Scene Editing** We can edit the scene in NeRF given the proposals produced by our NeRF-RPN. See Figure 8 for a demonstration which sets the density inside the proposal to zero when rendering.

#### 5.4. Comparison

To demonstrate the effectiveness of our method, we compare NeRF-RPN to recent 3D object detection methods like ImVoxelNet [51] and FCAF3D [50], with results shown in Table 3 and Figure 9. ImVoxelNet takes multi-view RGB as input, hence NeRF input images are used for training. FCAF3D is a point cloud-based 3D detector, so we use the ground-truth depth from Hypersim and ScanNet, NeRF

rendered depth for 3D-FRONT, and the corresponding RGB images to build point clouds. We adapt the implementation of these two methods in [5] and train them from scratch on the three datasets we use. Our method outperforms ImVoxelNet by a large margin on all datasets except 3D-FRONT, although ImVoxelNet takes advantage from our object-centric camera trajectory in 3D-FRONT. NeRF-RPN also outperforms FCAF3D on 3D-FRONT and Hypersim, despite ground-truth depth is used for Hypersim, giving FCAF3D extra advantages.




Figure 9. **Qualitative Comparison.** The first two rows are from Hypersim and the rest are from 3D-FRONT.

## 6. Conclusion

We propose the first significant 3D object detection framework for NeRF, NeRF-RPN, which operates on the voxel representation extracted from NeRF. By performing comprehensive experiments on different backbone networks, namely, VGG, ResNet, Swin Transformer along with anchor-based, anchor-free RPN heads, and multiple loss functions, we validate our NeRF-RPN can regress high-quality boxes directly from NeRF, without rendering images from NeRF in any view. To facilitate future work on 3D object detection in NeRF, we built a new benchmark dataset consisting of both synthetic and real-world data, with high NeRF reconstruction quality and careful bounding box labeling and cleaning. We hope NeRF-RPN will become a good baseline that can inspire and enable future work on 3D object detection in NeRFs.



## References

- [1] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9287–9296, 2019. 1, 2
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision (ECCV)*, pages 213–229. Springer, 2020. 2
- [3] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [4] Xiaozhi Chen, Kaustav Kundu, Yukun Zhu, Andrew G Berneshawi, Huimin Ma, Sanja Fidler, and Raquel Urtasun. 3d object proposals for accurate object class detection. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, 2015. 1, 2
- [5] MMDetection3D Contributors. MMDetection3D: Open-MMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020. 8
- [6] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 5, 6, 8
- [7] Jian Ding, Nan Xue, Yang Long, Gui-Song Xia, and Qikai Lu. Learning roi transformer for oriented object detection in aerial images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2849–2858, 2019. 5
- [8] Zhiwei Dong, Guoxuan Li, Yue Liao, Fei Wang, Pengju Ren, and Chen Qian. Centripetalnet: Pursuing high-quality keypoint pairs for object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [9] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [10] Cheng-Yang Fu, Wei Liu, Ananth Ranga, Amrith Tyagi, and Alexander C. Berg. Dssd : Deconvolutional single shot detector. *arXiv preprint arXiv:1701.06659*, 2017. 2
- [11] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10933–10942, 2021. 2, 5, 6, 8
- [12] Ross Girshick. Fast r-cnn. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, December 2015. 2, 4
- [13] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [14] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9224–9232, 2018. 2
- [15] JunYoung Gwak, Christopher Choy, and Silvio Savarese. Generative sparse detection networks for 3d single-shot object detection. In *ECCV*, 2020. 2
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2017. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(9):1904–1916, 2015. 2
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3
- [19] Binh-Son Hua, Quang-Hieu Pham, Duc Thanh Nguyen, Minh-Khoi Tran, Lap-Fai Yu, and Sai-Kit Yeung. Scenenn: A scene meshes dataset with annotations. In *International Conference on 3D Vision (3DV)*, pages 92–101. Ieee, 2016. 2, 5, 6, 8
- [20] Alex Krizhevsky. One weird trick for parallelizing convolutional neural networks. *CoRR*, abs/1404.5997, 2014. 2
- [21] Jean Lahoud and Bernard Ghanem. 2d-driven 3d object detection in rgb-d images. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4622–4630, 2017. 2
- [22] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *European Conference on Computer Vision (ECCV)*, September 2018. 2
- [23] Zuoxin Li and Fuqiang Zhou. Fssd: Feature fusion single shot multibox detector. *arXiv preprint arXiv:1712.00960*, 2017. 2
- [24] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 2, 3
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollar. Focal loss for dense object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct 2017. 2, 5
- [26] Songtao Liu, Di Huang, and andYunhong Wang. Receptive field block net for accurate and fast object detection. In *European Conference on Computer Vision (ECCV)*, September 2018. 2
- [27] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: Single shot MultiBox detector. In *European Conference on Computer Vision (ECCV)*, pages 21–37. Springer International Publishing, 2016. 2
- [28] Wei Liu, Irtiza Hasan, and Shengcai Liao. Center and scale prediction: Anchor-free approach for pedestrian and face detection. *arXiv preprint arXiv:1904.02948*, 2019. 2
- [29] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. *arXiv preprint arXiv:2203.05625*, 2022. 1, 2

- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 3
- [31] Ze Liu, Zheng Zhang, Yue Cao, Han Hu, and Xin Tong. Group-free 3d object detection via transformers. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2949–2958, 2021. 1, 2
- [32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*. OpenReview.net, 2019. 6
- [33] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3164–3173, 2021. 1, 2
- [34] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [35] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2906–2917, 2021. 2
- [36] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 41(4):102:1–102:15, July 2022. 2
- [37] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3d object detection with pointformer. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7463–7472, June 2021. 2
- [38] Julien Philip, Sébastien Morgenthaler, Michaël Gharbi, and George Drettakis. Free-viewpoint indoor neural relighting from multi-view stereo. *CoRR*, abs/2106.13299, 2021. 1, 6, 8
- [39] Charles R Qi, Xinlei Chen, Or Litany, and Leonidas J Guibas. Imvotenet: Boosting 3d object detection in point clouds with image votes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4404–4413, 2020. 2
- [40] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9277–9286, 2019. 2
- [41] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 918–927, 2018. 1, 2
- [42] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 652–660, 2017. 2
- [43] Rui Qian, Divyansh Garg, Yan Wang, Yurong You, Serge Belongie, Bharath Hariharan, Mark Campbell, Kilian Q Weinberger, and Wei-Lun Chao. End-to-end pseudo-lidar for image-based 3d object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5881–5890, 2020. 2
- [44] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [45] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 2, 4, 6
- [46] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M. Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2, 5, 8
- [47] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. *arXiv preprint arXiv:1811.08188*, 2018. 2
- [48] Barbara Roessle, Jonathan T. Barron, Ben Mildenhall, Pratul P. Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022. 6
- [49] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Fcaf3d: Fully convolutional anchor-free 3d object detection. *arXiv preprint arXiv:2112.00322*, 2021. 2
- [50] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Fcaf3d: fully convolutional anchor-free 3d object detection. In *ECCV*, 2022. 8
- [51] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *WACV*, 2022. 1, 2, 8
- [52] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. *arXiv preprint arXiv:2102.00463*, 2021. 2
- [53] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [54] Andrea Simonelli, Samuel Rota Bulo, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disentangling monocular 3d object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1991–1999, 2019. 2
- [55] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [56] Shuran Song and Jianxiong Xiao. Deep sliding shapes for amodal 3d object detection in rgb-d images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 808–816, 2016. 2
- [57] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2

- [58] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9627–9636, 2019. 2, 5
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017. 2
- [60] Haiyang Wang, Shaoshuai Shi, Ze Yang, Rongyao Fang, Qi Qian, Hongsheng Li, Bernt Schiele, and Liwei Wang. Rbgnet: Ray-based grouping for 3d object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1110–1119, 2022. 2
- [61] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 913–922, 2021. 2
- [62] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8445–8453, 2019. 1, 2
- [63] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, pages 180–191. PMLR, 2022. 1, 2
- [64] Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han. Oriented r-cnn for object detection. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3520–3529, 2021. 4, 5
- [65] Bin Xu and Zhenzhong Chen. Multi-level fusion based 3d object detection from monocular images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2345–2353, 2018. 2
- [66] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019. 2
- [67] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. 2
- [68] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z. Li. Single-shot refinement neural network for object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2
- [69] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *AAAI*, volume 34, pages 12993–13000, 2020. 7
- [70] Dingfu Zhou, Jin Fang, Xibin Song, Chenye Guan, Junbo Yin, Yuchao Dai, and Ruigang Yang. Iou loss for 2d/3d object detection. In *International Conference on 3D Vision (3DV)*, pages 85–94. IEEE, 2019. 5, 7
- [71] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 2
- [72] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krahenbuhl. Bottom-up object detection by grouping extreme and center points. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [73] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4490–4499, 2018. 1, 2