

RSN: Range Sparse Net for Efficient, Accurate LiDAR 3D Object Detection

Pei Sun¹ Weiyue Wang¹ Yuning Chai¹ Gamaleldin Elsayed²
Alex Bewley² Xiao Zhang¹ Cristian Sminchisescu² Dragomir Anguelov¹

¹Waymo LLC, ²Google
{peis}@waymo.com

Abstract

The detection of 3D objects from LiDAR data is a critical component in most autonomous driving systems. Safe, high speed driving needs larger detection ranges, which are enabled by new LiDARs. These larger detection ranges require more efficient and accurate detection models. Towards this goal, we propose Range Sparse Net (RSN) – a simple, efficient, and accurate 3D object detector – in order to tackle real time 3D object detection in this extended detection regime. RSN predicts foreground points from range images and applies sparse convolutions on the selected foreground points to detect objects. The lightweight 2D convolutions on dense range images results in significantly fewer selected foreground points, thus enabling the later sparse convolutions in RSN to efficiently operate. Combining features from the range image further enhance detection accuracy. RSN runs at more than 60 frames per second on a 150m × 150m detection region on Waymo Open Dataset (WOD) while being more accurate than previously published detectors. As of 11/2020, RSN is ranked first in the WOD leaderboard based on the APH/LEVEL_1 metrics for LiDAR-based pedestrian and vehicle detection, while being several times faster than alternatives.

1. Introduction

Concurrent with steady progress towards improving the accuracy and efficiency of 3D object detector algorithms [37, 24, 21, 16, 42, 32, 22, 30, 5, 35, 3], LiDAR sensor hardware has been improving in maximum range and fidelity, in order to meet the needs of safe, high speed driving. Some of the latest commercial LiDARs can sense up to 250m [12] and 300m [36] in all directions around the vehicle. This large volume coverage places strong demands for efficient and accurate 3D detection methods.

Grid based methods [43, 16, 42, 35, 8] divide the 3D space into voxels or pillars, each of these being optionally encoded using PointNet [25]. Dense convolutions are applied

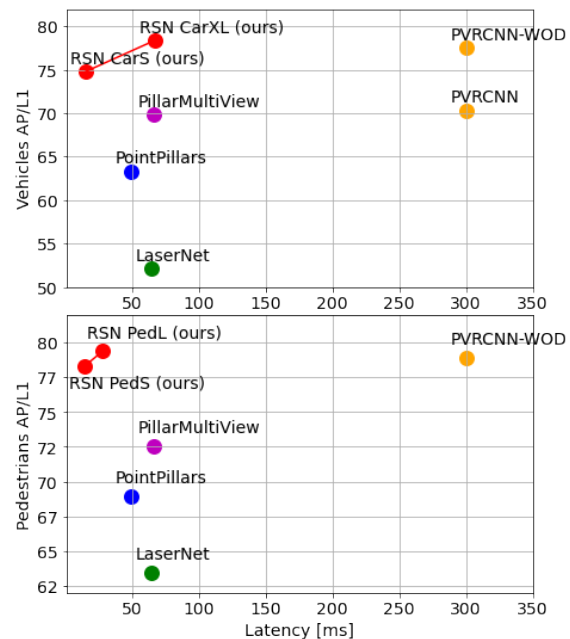


Figure 1. Accuracy (3D AP/L1 on WOD validation set) vs Latency (ms). RSN models significantly outperform others. See Table 1 and Table 2 for more details.

on the grid to extract features. This approach is inefficient for large grids which are needed for long range sensing or small object detection. Sparse convolutions [30] scale better to large detection ranges but are usually slow due to the inefficiencies of applying to all points. **Range images** are native, dense representations, suitable for processing point clouds captured by a single LiDAR. Range image based methods [21, 3] perform convolutions directly over the range in order to extract point cloud features. Such models scale well with distance, but tend to perform less well in occlusion handling, accurate object localization, and for size estimation. A second stage, refining a set of initial candidate detections, can help mitigate some of these quality issues, at the expense of significant computational cost.

To address the shortcomings of existing approaches, we

introduce a novel 3D object detection model – Range Sparse Net (RSN) – which boosts the 3D detection accuracy and efficiency by combining the advantages of methods based on both dense range images and grids. RSN first applies a lightweight 2D convolutional network to efficiently learn semantic features from the high-resolution range image. Unlike existing range image methods, which regress boxes directly from their underlying features, RSN is trained for high recall foreground segmentation. In a subsequent stage, sparse convolutions are applied only on the predicted foreground voxels and their learned range image features, in order to accurately regress 3D boxes. A configurable sparse convolution backbone and a customized CenterNet [41] head designed for processing sparse voxels is introduced in order to enable end-to-end, efficient, accurate object detection without non-maximum-suppression. Figure 1 summarizes the main gains obtained with RSN models compared to others on the WOD validation set to demonstrate RSN’s efficiency and accuracy.

RSN is a novel multi-view fusion method, as it transfers information from perspective view (range image) to the 3D view (sparse convolution on the foreground points). Its fusion approach differs from existing multi-view detection methods [42, 35] in that 1) RSN’s first stage directly operates on the high resolution range image while past approaches [42, 35] perform voxelization (in a cylindrical or spherical coordinate system) that may lose some resolution, especially for small objects at a distance. 2) RSN’s second stage processes only 3D points selected as foreground by the first stage, which yields improvements in both feature quality and efficiency.

RSN’s design combines several insights that make the model very efficient. The initial stage is optimized to rapidly discriminate foreground from background points, a task that is simpler than full 3D object detection and allows a lightweight 2D image backbone to be applied to the range image at full resolution. The downstream sparse convolution processing is only applied on points that are likely to belong to a foreground object, which leads to additional, significant savings in compute. Furthermore, expensive postprocessing such as non-maximum suppression are eliminated by gathering local maxima center-ness points on the output, similar to CenterNet [41].

In this work, we make the following main contributions:

- We propose a simple, efficient and accurate 3D LiDAR detection model RSN, which utilizes LiDAR range images to perform foreground object segmentation, followed by sparse convolutions to efficiently process the segmented foreground points to detect objects.
- We propose a simple yet effective temporal fusion strategy in RSN with little additional inference cost.
- In experiments on the Waymo Open Dataset [34] (WOD), we demonstrate the state of art accuracy and

efficiency for vehicle and pedestrian detection. Experiments on an internal dataset further demonstrate RSN’s scalability for long-range object detection.

- We conduct ablation studies to examine the effectiveness of range image features and the impact of aspects like foreground point selection thresholds, or end-to-end model training, on both latency and accuracy.

2. Related Work

2.1. LiDAR Data Representation

There are four major LiDAR data representations for 3D object detection including voxel grids, point sets, range images, and hybrids.

Voxel grid based methods. 3D points are divided into a grid of voxels. Each voxel is encoded with hand-crafted metrics such as voxel feature means and covariances. Vote3Deep [7] was the first to apply a deep network composed of sparse 3D convolutions to 3D detection. They also proposed an L_1 penalty to favour sparsity in deeper layers. The voxels can be scattered to a pseudo-image which can be processed by standard image detection architectures. MV3D [4], PIXOR [38] and Complex YOLO [33] are notable models based on this approach. VoxelNet [43] applied PointNet [25] in each voxel to avoid handcrafted voxel features. PointPillars [16] introduced 2D pillar to replace 3D voxel to boost model efficiency. For small enough 3D voxel sizes, the PointNet can be removed if 3D sparse convolutions are used. Notable examples based on this approach include Second [37] and PVRCNN [30].

There are three major drawbacks to voxel based methods. 1) Voxel size is constant at all ranges which limits the model’s capability at distance and usually needs larger receptive fields. 2) The requirement of a full 3D grid poses a limitation for long-range, since both complexity and memory consumption scale quadratically or cubically with the range. Sparse convolutions can be applied to improve scalability but is usually still limited by the large number of voxels. 3) The voxel representation has a limited resolution due to the scalability issue mentioned above.

Point set based methods. This line of methods treats point clouds as unordered sets. Most approaches are based on the seminal PointNet and variants [25, 26]. FPointNet[24] detects objects from a cropped point cloud given by 2D proposals obtained from images; PointRCNN[32] proposes objects directly from each point; STD [39] relies on a sparse to dense strategy for better proposal refinement; DeepHough [23] explores deep hough voting to better group points before generating box proposals. Although these methods have the potential to scale better with range, they lag behind the quality of voxel methods. Moreover, they require nearest neighbor search for the input, scaling with the number of points, which can be costly.

Range image based methods. Despite being a native and dense representation for 3D points captured from a single view-point e.g. LiDAR, prior work on using 2D range images is not extensive. LaserNet [21] applied a traditional 2D convolution network to range image to regress boxes directly. RCD-RCNN [3] pursued range conditioned dilation to augment traditional 2D convolutions, followed by a second stage to refine the proposed range-image boxes which is also used by Range-RCNN [18]. Features learned from range images alone are very efficient when performing 2D convolutions on 2D images but aren't that good at handling occlusions, for accurate object localization, and for size regression, which usually requires more expressive 3D features.

Hybrid methods. MultiView [42] fuses features learned from voxels in both spherical and Cartesian coordinates to mitigate the limited long-range receptive fields resulting from the fixed-voxel discretization in grid based methods. Pillar-MultiView [35] improves [42] by further projecting fused spherical and cartesian features to bird-eye views followed by additional convolution processing to produce stronger features. These methods face similar scalability issues as voxel approaches.

2.2. Object Detection Architectures

Typical two-stage detectors [10, 9, 28, 6] generate a sparse set of regions of interest (RoIs) and classify each of them by a network. PointRCNN [32], PVRCNN [30], RCD-RCNN [3] share similar architectures with Faster-RCNN but rely on different region proposal networks designed for different point cloud representations. Single-stage detectors were popularized by the introduction of YOLO [27], SSD [20] and RetinaNet [19]. Similar architectures are used to design single stage 3D point cloud methods [43, 16, 37, 42, 35]. These achieve competitive accuracy compared to two stage methods such as PVRCNN [30] but have much lower latency. Keypoint-based architectures such as CornerNet [17] and CenterNet [41] enable end to end training without non-maximum-suppression. AFDet [8] applies a CenterNet-style detection head to a PointPillars-like detector for 3D point clouds. Our proposed RSN method also relies on two stages. However the first stage performs segmentation rather than box proposal estimation, and the second stage detects objects from segmented foreground points rather than performing RoI refinement. RSN adapts the CenterNet detection head to sparse voxels.

3. Range Sparse Net

The main contribution of this work is the Range Sparse Net (RSN) architecture (Fig. 2). RSN accepts raw LiDAR range images [34] as input to an efficient 2D convolution backbone that extracts range image features. A segmentation head is added to process range image features. This segments background and foreground points, with the foreground be-

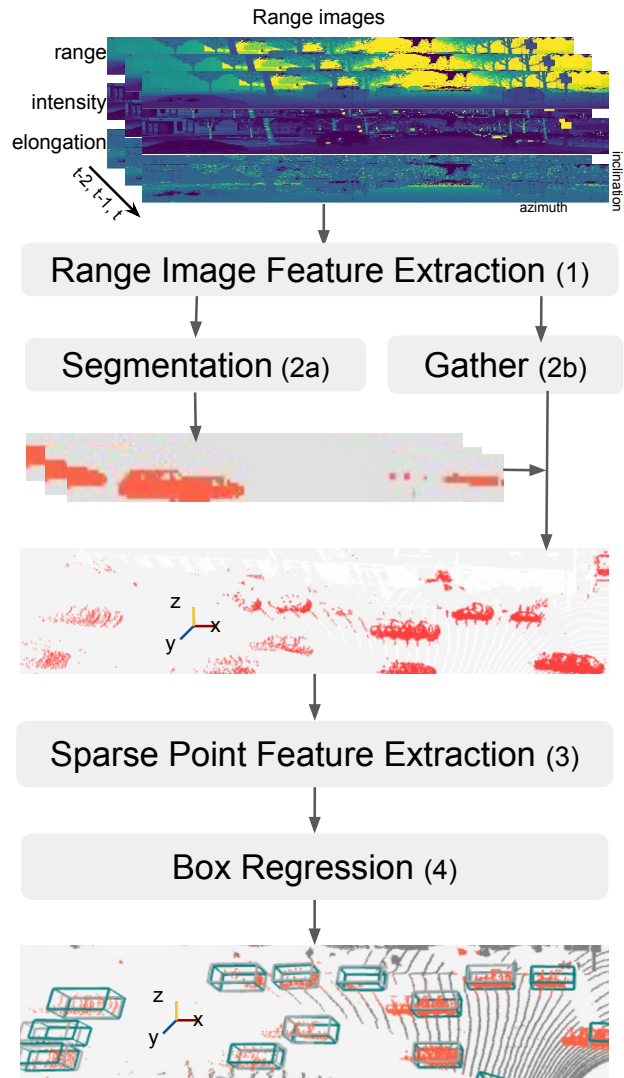


Figure 2. (Best viewed in color) Range Sparse Net object detection architecture. The net consists of five components: 1) Range image feature extraction: a 2D convolution net on range images to extract associated image features. 2) Foreground point selection: foreground points are segmented on range images in 2a); together with the learned range image features, they are gathered to sparse points in 2b). 3) Sparse point feature extraction: per-point features are extracted on the selected foreground points by applying sparse convolutions. 4) A sparse CenterNet head to regress boxes. Red points are selected foreground points. Light gray boxes are ground truths. Teal boxes are detection results.

ing points inside ground truth objects. Unlike traditional semantic segmentation, recall is emphasized over high precision in this network. We select foreground points based on the segmentation result. The selected foreground points are further voxelized and fed into a sparse convolution network. These sparse convolutions are very efficient because we only need to operate on a small number of foreground points. At the end, we apply a modified CenterNet [41] head to regress 3D boxes efficiently without non-maximum-suppression.

3.1. Range Image Feature Extraction (RIFE)

Range images are a native dense representation of the data captured by LiDAR sensors. Our input range images contain range, intensity and elongation channels, where range is the distance from LiDAR to the point at the time the point is collected, while intensity and elongation are LiDAR return properties which can be replaced or augmented with other LiDAR specific signals. The channel values of the input range images are normalized by clipping and rescaling to $[0, 1]$.

A 2D convolution net is applied on the range image to simultaneously learn range image features and for foreground segmentation. We adopt a lightweight U-Net [29] with its structure shown in Fig. 3. Each $D(L, C)$ downsampling block contains L resnet [13] blocks each with C output channels. Within each block the first has stride 2. Each $U(L, C)$ block contains 1 upsampling layer and L resnet blocks. All resnet blocks have stride 1. The upsampling layer consists of a 1×1 convolution followed by a bilinear interpolation.

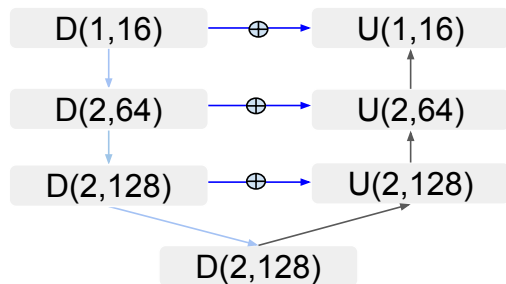


Figure 3. Range image U-Net feature extractor to compute high level semantic range features. See section 3.1 for details.

3.2. Foreground Point Selection

To maximize efficiency through sparsity in the downstream processing, the output of this 2D convolutional network is an ideal place to reduce the input data cloud to only points most likely to belong to an object. Here, a 1×1 convolutional layer performs pixelwise foreground classification on the learned range image features from 3.1. This layer is trained using the focal loss [19] with ground truth labels derived from 3d bounding boxes by checking whether the corresponding pixel point is in any box.

$$L_{\text{seg}} = \frac{1}{P} \sum_i L_i, \quad (1)$$

P is the total number of valid range image pixels. L_i is the focal loss for point i . Points with foreground score s_i greater than a threshold γ are selected. As false positives can be removed in the later sparse point feature extraction phase (§3.3) but false negatives cannot be recovered, the foreground threshold is selected to achieve high recall and acceptable precision.

3.3. Sparse Point Feature Extraction (SPFE)

We apply dynamic voxelization [42] on the selected foreground points. Similar to PointPillars [16], we append each point p with $p - m$, var , $p - c$ where m , var is the arithmetic mean and covariance of each voxel, c is the voxel center. Voxel sizes are denoted as $\Delta_{x,y,z}$ along each dimension. When using a pillar style voxelization where 2D sparse convolution is applied, Δ_z is set to $+\infty$. The selected foreground points are encoded into sparse voxel features which can optionally be further processed by a PointNet [25].

A 2D or 3D sparse convolution network (for pillar style, or 3D type voxelization, respectively) is applied on the sparse voxels. Fig. 4 illustrates the net building blocks and example net architectures used for vehicle and pedestrian detection. More network architecture details can be found in the Appendix A.

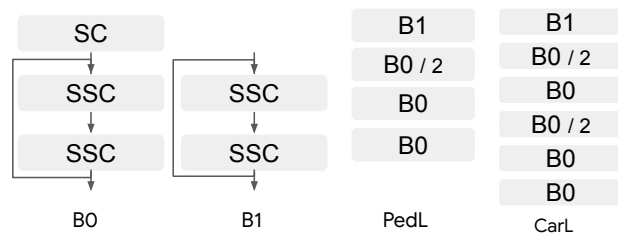


Figure 4. SPFE building blocks and example net architectures. See section 3.3 for usage details. SC denotes 3×3 or $3 \times 3 \times 3$ sparse convolution [11] with stride 1 or 2. SSC denotes 3×3 or $3 \times 3 \times 3$ submanifold sparse convolution. PedL (2D) and CarL (2D) are the large pedestrian and vehicle SPFEs. /2 denotes stride 2.

3.4. Box Regression

We use a modified CenterNet [41, 8] head to regress boxes from point features efficiently. The feature map consists of voxelized coordinates $V = \{v | v \in \mathbb{N}_0^d\}$, where $d \in \{2, 3\}$ depending on whether 2D or 3D SPFE is used. We scale and shift it back to the raw point Cartesian coordinate as $\tilde{V} = \{\tilde{v} | \tilde{v} \in R^d\}$. The ground truth heatmap for any $\tilde{v} \in \tilde{V}$ is computed as $h = \min\{\exp(-\frac{\|\tilde{v} - b_c\| - \|\tilde{V} - b_c\|}{\sigma^2}) | b_c \in B_c(\tilde{v})\}$ where $B_c(\tilde{v})$ is the set of centers of the boxes that contain \tilde{v} . $h = 0$ if $|B_c(\tilde{v})| = 0$. This is illustrated in Fig. 5. σ is a per class constant. We use a single fully connected layer to predict heatmap and box parameters. The heatmap is regressed with a penalty-reduced focal loss [41, 19].

$$L_{\text{hm}} = -\frac{1}{N} \sum_{\tilde{p}} \{(1 - \tilde{h})^\alpha \log(\tilde{h}) I_{h > 1 - \epsilon} + (1 - h)^\beta \tilde{h}^\alpha \log(1 - \tilde{h}) I_{h \leq 1 - \epsilon}\}, \quad (2)$$

where \tilde{h} and h are the predicted and ground truth heatmap values respectively. ϵ , added for numerical stability, is set to $1e - 3$. We use $\alpha = 2$ and $\beta = 4$ in all experiments, following [41, 17].

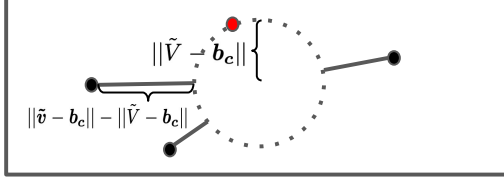


Figure 5. RSN centerness heatmap computation. The heatmap value is computed by the distance between the point and the circle placed at the box center, with radius being the distance from the box center to the closest point (red point).

The 3D boxes are parameterized as $\mathbf{b} = \{d_x, d_y, d_z, l, w, h, \theta\}$ where d_x, d_y, d_z are the box center offsets relative to the voxel centers. Note that d_z is the same as the absolute box z center if 2D point feature extraction backbone is used (see 3.3). l, w, h, θ are box length, width, height and box heading. A bin loss [32] is applied to regress heading θ . The other box parameters are directly regressed under smooth L1 losses. IoU loss [40] is added to further boost box regression accuracy. Box regression losses are only active on the feature map pixels that have ground truth heatmap values greater than a threshold δ_1 .

$$L_{\theta_i} = L_{bin}(\theta_i, \tilde{\theta}_i), \quad (3)$$

$$L_{\mathbf{b}_i \setminus \theta_i} = \text{SmoothL1}(\mathbf{b}_i \setminus \theta_i - \tilde{\mathbf{b}}_i \setminus \tilde{\theta}_i), \quad (4)$$

$$L_{\text{box}} = \frac{1}{N} \sum_i (L_{\theta_i} + L_{\mathbf{b}_i \setminus \theta_i} + L_{\text{iou}_i}) I_{h_i > \delta_1}, \quad (5)$$

where \tilde{b}_i, b_i are the predicted and ground truth box parameters respectively, $\tilde{\theta}_i, \theta_i$ are the predicted and ground truth box heading respectively. h_i is the ground truth heatmap value computed at feature map pixel i .

The net is trained end to end with the total loss defined as

$$L = \lambda_1 L_{\text{seg}} + \lambda_2 L_{\text{hm}} + L_{\text{box}} \quad (6)$$

We run a sparse submanifold max pooling operation on the sparse feature map voxels that have heatmap prediction greater than a threshold δ_2 . Boxes corresponding to local maximum heatmap predictions are selected.

3.5. Temporal Fusion

Existing range image based detection methods [21] [3] are not temporal fusion friendly as range images are constructed while the self-driving-car (SDC) moves. Stacking range images directly gives little benefit for detection performance due to ego-motion. Removing ego-motion from the range images is not optimal because range reconstructions at a different frame results in non-trivial quantization errors.

Temporal RSN takes a sequence of temporal invariant range images as input as shown in Fig. 2. RIFE is applied on each range image to segment foreground points and extract range image features. Then we transform all the selected

points to the latest frame to remove ego-motion. During the SPFE phase, we append to each point voxel features computed from its own frame instead of all frames. This works better because it avoids mixing points from different frames together during voxelization. In addition, we append the time difference in seconds w.r.t. the latest frame to each point to differentiate points from different frames. The selected foreground points from all frames are processed by the SPFE backbone same as single frame models.

4. Experiments

We introduce the RSN implementation details and illustrate its efficiency and accuracy in multiple experiments. Ablation studies are conducted to understand the importance of various RSN components.

4.1. Waymo Open Dataset

We primarily benchmark on the challenging Waymo Open Dataset (WOD) [34]. WOD released its raw data in high quality range image format directly, which makes it a better fit for building range image models. The dataset contains 1150 sequences in total, split into 798 training, 202 validation, and 150 test. Each sequence contains about 200 frames, where each frame captures the full 360 degrees around the ego-vehicle that results in a range image of a dimension 64×2650 pixels. The dataset has one long range LiDAR with range capped at 75 meters and four near range LiDARs. We only used data from the long range LiDAR but still evaluated our results on full range. In practice, RSN can be adapted to accept multiple LiDAR images as inputs.

4.2. Implementation Details

RSN is implemented in the Tensorflow framework [1] with sparse convolution implementation similar as [37]. Pedestrians and vehicles are trained separately with different SPFEs (§3.3). They share the same RIFE (§3.1). We show results from 3 vehicle models CarS, CarL, CarXL and 2 pedestrian models PedS, PedL with network details described in §3.3 and Appendix A. Each model can be trained with single frame input (e.g. CarS_1f) or 3 frame input (e.g. CarS_3f). The input images are normalized by $\min(v, m)/m$ where v is range, intensity and elongation, m is 79.5, 2.0, 2.0 respectively. The last return is picked if there are multiple laser returns.

The foreground score cutoff γ in §3.2 is set to 0.15 for vehicle and 0.1 for pedestrian. The segmentation loss weight λ_1 in Eq.6 is set to 400. The voxelization region is $[-79.5m, 79.5m] \times [-79.5m, 79.5m] \times [-5m, 5m]$. The voxel sizes are set to 0.2 meter and 0.1 meter for vehicle model and pedestrian model respectively. Per object σ in the heatmap computation is set to 1.0 for vehicle and 0.5 for pedestrian. The heatmap loss weight λ_2 is set to 4 in Eq. 6. The heatmap thresholds δ_1, δ_2 in §3.4 are both set to 0.2. We

use 12 and 4 bins in the heading bin loss in Eq. 3 for heading regression for vehicle and pedestrian, respectively.

4.3. Training and Inference

RSN is trained from scratch end-to-end using an ADAM optimizer [15] on Tesla V100 GPUs. Different SPFE backbones are trained with the maximum possible batch sizes to fit the net in GPU memory. Single frame models are trained on 8 GPUs. 3-frame temporal models are trained on 16 GPUs. We adopted the cosine learning rate decay, with initial learning rate set to 0.006, 5k warmup steps starting at 0.003, 120k steps in total. We observed that accuracy metrics such as AP fluctuate during training because the points selected to SPFE keep changing, although networks always stabilize at the end. This input diversity to SPFE adds regularization to RSN. Layer normalization [2] instead of batch normalization [14] is used in the PointNet within each voxel because the number of foreground points varies a lot among input frames.

We rely on two widely adopted data augmentation strategies, including random flipping along the X axis and global rotation around the Z axis with a random angle sampled from $[-\pi/4, \pi/4]$ on the selected foreground points.

During inference, we reuse past learned range features and segmentation results (outputs of foreground point selection) such that the inference cost for temporal models remains similar as the single frame models.

4.4. Results

All detection results are measured using the official WOD evaluation detection metrics which are BEV and 3D average precision (AP), heading error weighted BEV and 3D average precision (APH) for L1 (easy) and L2 (hard) difficulty levels [34]. The IoU threshold is set as 0.7 for vehicle, 0.5 for pedestrian. We show results on the validation set for all our models in Table 1 and Table 2, results on the official test set in Table 3. The latency numbers are obtained on Tesla V100 GPUs with float32 without TensorRT except PVRCNN which is obtained on Titan RTX from PVRCNN authors. In order to better show the latency improvement from our RSN model, NMS timing is not included in all of the baselines because our efficient detection head can be adapted to most of other baselines. We do not show timing for our single frame models as their latency is bounded by their multi-frame correspondences.

Table 1 shows that our single frame model CarS_1f is at least 3x more efficient than the baselines while still being more accurate than all single stage methods. Its temporal version boosts the accuracy further at negligible additional inference costs. CarXL_3f significantly outperforms all published methods. It also outperforms PVRCNN-WOD [31], the most accurate LiDAR only model submitted in the Waymo Open Dataset Challenge.

Table 2 shows more significant improvements on both efficiency and accuracy on pedestrian detection. The efficient single frame model PedS_1f is significantly more accurate and efficient than all published single-stage baseline models. Its temporal version further improves accuracy. The less efficient model PedL_3f, outperforms PVRCNN-WOD [31], while still being significantly more efficient than all baselines. We see additional efficiency gains on pedestrian detection compared with vehicle detection because there are much fewer people-foreground points. Given the high resolution range image and the high recall foreground segmentation, our model is a great fit for real time small object detection.

Table 3 shows that RSN ensemble outperforms the PVRCNN WOD challenge submission [31] which is an ensemble of many models.

Fig. 6 shows a few examples picked from Waymo Open Dataset validation set to demonstrate the model quality in dealing with various hard cases such as a crowd of pedestrians, small objects with few points, large objects, and moving objects in temporal model.

4.5. Foreground Point Selection Experiments

Foreground point selection is one of the major contributions in the RSN model that supports better efficiency. We conduct experiments by scanning the foreground selection threshold γ described in §3.3. As shown in Fig. 7, there exists a score threshold γ that reduces model latency significantly with negligible impact on accuracy.

In practice, γ and λ_1 in Eq 6 need to be set to values so that selected foreground points have high recall and enough accuracy to achieve good speedup. In our experiments, foreground precision/recall is 77.5%/99.6% for CarS_3f and 15.3%/97.6% for PedS_3f. We can start with low γ and scanning some possible values of λ_1 to pick one λ_1 . Then we grid search a few γ .

4.6. Ablation study

In this section, we show additional ablation studies in order to gain insight into model design choices. All experiments in this section are conducted on our efficient models CarS_3f and on PedS_3f.

Table 4 shows that features learnt from range image not only help segment foreground points, thus supporting model efficiency, but also improve model accuracy as shown in row -RI. Accuracy improvement is higher for pedestrians because of the high resolution semantic feature learned especially impacting the long range. Gradients passed from SPFE to RIFE help detection accuracy as shown in row -E2E. Temporal variant features (x, y, z) with ego-motion removed hurt detection accuracy for pedestrian detection as shown in row +xyz. Detection accuracy drops if the heatmap normalization described in §3.4 is disabled as shown in row -Norm.

Method	Latency (ms)	AP/APH L1		AP/APH L2		AP/APH L1 3D by distance		
		BEV	3D	BEV	3D	<30m	30-50m	> 50m
LaserNet CVPR'19 [21] *	64.3	71.2/67.7	52.1/50.1	-	-	70.9/68.7	52.9/51.4	29.6/28.6
P.Pillars CVPR'19 [16] †	49.0	82.5/81.5	63.3/62.7	73.9/72.9	55.2/54.7	84.9/84.4	59.2/58.6	35.8/35.2
PillarMultiView ECCV'20 [35]	66.7 ‡	87.1/-	69.8/-	-	-	88.5/-	66.5/-	42.9/-
PVRCNN CVPR'20 [30]	-	83.0/82.1	70.3/69.7	77.5/76.6	65.4/64.8	91.9/91.3	69.2/68.5	42.2/41.3
PVRCNN WOD'20 [31]	300 ¶	-	77.5/76.9	-	68.7/68.2	-	-	-
RCD CORL'20 [3]	-	82.1/83.4	69.0/68.5	-	-	87.2/86.8	66.5/66.1	44.5/44.0
RSN CarS_1f (Ours)	-	86.7/86.0	70.5/70.0	77.5/76.8	63.0/62.6	90.8/90.4	67.8/67.3	45.4/44.9
RSN CarS_3f (Ours)	15.5	88.1/87.4	74.8/74.4	80.8/80.2	65.8/65.4	92.0/91.6	73.0/72.5	51.8/51.2
RSN CarL_1f (Ours)	-	88.5/87.9	75.1/74.6	81.2/80.6	66.0/65.5	91.8/91.4	73.5/73.1	53.1/52.5
RSN CarL_3f (Ours)	25.4	91.0/90.3	75.7/75.4	82.1/81.6	68.6/68.1	92.1/91.7	74.6/74.1	56.1/55.4
RSN CarXL_3f (Ours)	67.5	91.3/90.8	78.4/78.1	82.6/82.2	69.5/69.1	92.1/91.7	77.0/76.6	57.5/57.1

Table 1. Performance comparisons on the Waymo Open Dataset *validation set* for vehicle detection. (*) is re-implemented by [3]. (†) is our re-implementation with flip and rotation data augmentation following PointPillar setting in [34] which is better than other PointPillars re-implementations such as [42]. (‡) is from [35]. ¶ is obtained privately from PVRCNN authors who benchmarked on Titan RTX. All the other latency numbers are obtained based on our own implementations on Tesla V100 GPUs. They are averaged on 10 scenes, each has more than 100 vehicles.

Method	Latency (ms)	AP/APH L1		AP/APH L2		AP/APH L1 3D by distance		
		BEV	3D	BEV	3D	< 30m	30-50m	> 50m
LaserNet CVPR'19 [21] *	64.3	70.0/-	63.4/-	-	-	73.5/-	61.6/-	42.7/-
P.Pillars CVPR'19 [16] †	49.0	76.0/62.0	68.9/56.6	67.2/54.6	60.0/49.1	76.7/64.3	66.9/54.3	52.9/40.5
PillarMultiView ECCV'20 [35]	66.7 ‡	78.5/-	72.5/-	-	-	79.3/—	72.1/—	56.8/—
PVRCNN WOD'20 [31]	300 ¶	-	78.9/75.1	-	69.8/66.4	-	-	-
RSN PedS_1f (Ours)	-	80.7/74.9	74.8/69.6	71.2/65.9	65.4/60.7	81.4/77.4	72.8/66.8	59.0/50.6
RSN PedS_3f (Ours)	14.4	84.2/80.7	78.3/75.2	74.8/71.6	68.9/66.1	81.7/78.8	74.4/71.3	64.9/61.5
RSN PedL_1f (Ours)	-	83.4/77.6	77.8/72.7	73.9/68.6	68.3/63.7	83.9/79.7	74.1/68.2	62.1/54.1
RSN PedL_3f (Ours)	28.2	85.0/81.4	79.4/76.2	75.5/72.2	69.9/67.0	84.5/81.5	78.1/74.7	68.5/65.0

Table 2. Performance comparison on the Waymo Open Dataset *validation set* for pedestrian detection. See Table 1 for details on ways to obtaining latency numbers. All the latency of our models are averaged on 10 scenes, each has more than 50 pedestrians.

Method	AP/APH L1 3D				AP/APH L2 3D			
	Overall	< 30m	30-50m	> 50m	Overall	< 30m	30-50m	> 50m
VEHICLE								
P.Pillars CVPR'19 [16] †	68.6/68.1	87.2/86.7	65.5/64.9	40.9/40.2	60.5/60.1	85.9/85.4	58.9/58.3	31.3/30.8
PVRCNN Ensem WOD'20 [31]	81.1/80.6	93.4/93.0	80.1/79.6	61.2/60.5	73.7/73.2	92.5/92.0	74.0/73.5	49.3/48.6
RSN CarXL_3f (Ours)	80.7/80.3	92.2/91.9	79.1/78.7	63.0/62.5	71.9/71.6	91.5/91.1	71.4/71.1	49.9/49.5
RSN CarXL_3f Ensem (Ours)	81.4/81.0	92.4/92.0	80.2/79.8	64.7/64.1	72.8/72.4	91.5/91.1	74.2/73.8	51.3/50.8
PEDESTRIAN								
P.Pillars CVPR'19 [16] †	68.0/55.5	76.0/63.5	66.8/54.1	54.3/42.1	61.4/50.1	73.4/61.2	61.5/49.8	43.9/34.0
PVRCNN Ensem WOD'20 [31]	80.3/76.3	86.7/82.9	78.9/74.8	70.5/66.4	74.0/70.2	84.8/80.9	73.6/69.6	59.2/55.5
RSN PedL_3f (Ours)	78.9/75.6	85.5/82.4	77.5/74.2	67.3/64.1	70.7/67.8	81.9/78.9	70.3/67.3	55.8/53.0
RSN PedL_3f Ensem (Ours)	82.4/78.0	89.1/85.0	81.1/76.8	70.7/66.3	74.7/70.7	86.0/82.0	74.6/70.6	58.7/54.8

Table 3. Performance comparison on the Waymo Open Dataset *test set*. (†) is our re-implementation as described in Table 1. 'Ensem' is short for ensemble. See Appendix C for details.

4.7. Scalability

To further demonstrate RSN's model scalability, we conducted experiments on an internal dataset collected from higher quality longer range LiDARs. Here, the detection

range is a square of size $[-250m, 250m] \times [-250m, 250m]$ and that is centered at the SDC. This is beyond the memory capacity of PointPillars [16] running on a Tesla v100 GPU. We have trained RSN CarS_3f and a variant with RIFE

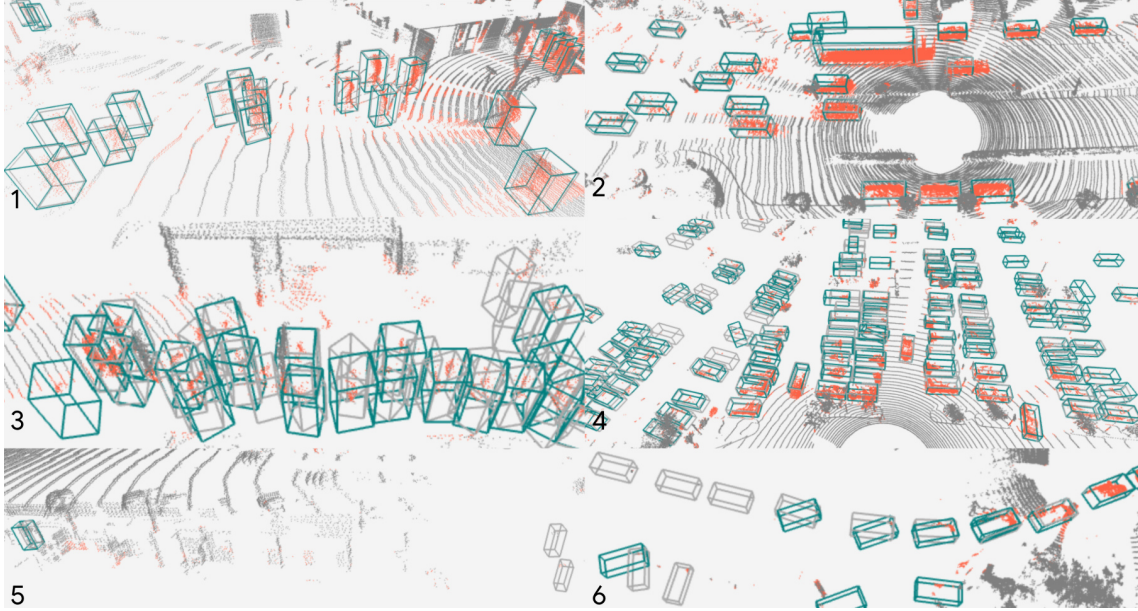


Figure 6. (Best viewed in color) Example pedestrian and vehicle detection results of CarS_3f and PedS_3f on the Waymo Open Dataset validation set. Light gray boxes are ground-truth and teal boxes are our prediction results. Red points are selected foreground points. Ex 1, 2: RSN performs well when objects are close and mostly visible. Both vehicles and pedestrians are predicted with high accuracy, including dynamic vehicles, large vehicles. Ex 3, 4: RSN handles large crowds with severe occlusion with few false positives and false negatives. Many of the false-negatives in Ex 4 have very few points in the ground-truth boxes. Ex 5, 6: Typical failures of RSN are for distant or heavily occluded objects and having very few points observed.

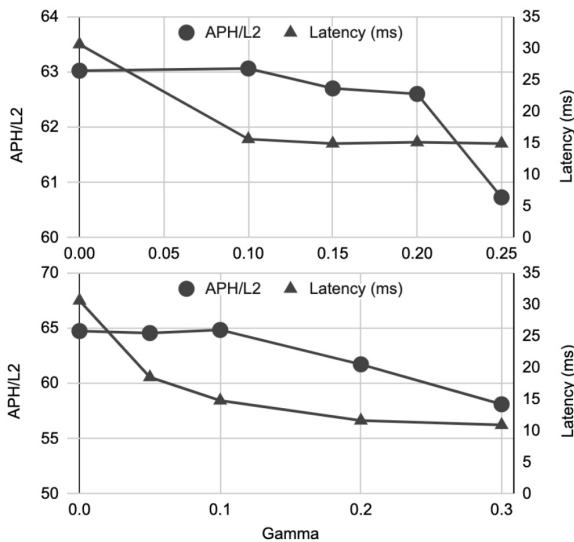


Figure 7. Model performance for different foreground point selection thresholds γ as defined in §3.2. Top: vehicle result for RSN CarS_3f. Bottom: pedestrian result for model PedS_3f. The model accuracy (by APH/L2) does not decrease much but latency drops rapidly when γ is less than a certain threshold.

and foreground point selection removed on this dataset. As shown in table 5, RSN can scale to a significantly larger detection range with good accuracy and efficiency. This demonstrates that foreground sampling and range image features remain effective in the larger detection range.

	Vehicle		Pedestrian	
	AP/APH L2	Latency	AP/APH L2	Latency
Baseline	64.2/63.9	15.4	68.88/66.07	14.4
-RI	60.9/60.3	27.0	63.5/60.5	30.0
-E2E	60.1/59.7	15.6	64.8/61.7	14.6
+xyz	64.1/63.7	-	64.2/61.3	-
-Norm	60.6/60.2	-	64.7/61.9	-

Table 4. The 3D AP/APH at LEVEL_2 difficulty and latency in milliseconds on the validation set for several ablations (§4.6).

RIFE stage	BEV APH	3D APH	Latency (ms)
✓	83.6	61.2	22
✗	79.4	53.4	44

Table 5. The APH and latency in milliseconds on the test set of an internal long range LiDAR dataset for vehicle detection with model CarS_3f.

5. Conclusions

We have introduced RSN, a novel range image based 3D object detection method that can be trained end-to-end using LiDAR data. The network operates in the large detection range required for safe, fast-speed driving. In the Waymo Open Dataset, we show that RSN outperforms all existing LiDAR-only methods by offering higher detection performance (AP/APH on both BEV and 3D) as well as faster running times. For future work, we plan to explore alternative detection heads and optimized SPFE in order to better take advantage of the sparsity of the foreground points.

References

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, 2016. 5
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 6
- [3] Alex Bewley, Pei Sun, Thomas Mensink, Dragomir Anguelov, and Cristian Sminchisescu. Range conditioned dilated convolutions for scale invariant 3d object detection. In *Conference on Robot Learning*, 2020. 1, 3, 5, 7
- [4] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 2
- [5] Shuyang Cheng, Zhaoqi Leng, Ekin Dogus Cubuk, Barret Zoph, Chunyan Bai, Jiquan Ngiam, Yang Song, Benjamin Caine, Vijay Vasudevan, Congcong Li, et al. Improving 3d object detection through progressive population based augmentation. In *ECCV*, 2020. 1
- [6] Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*, pages 379–387, 2016. 3
- [7] Martin Engelcke, Dushyant Rao, Dominic Zeng Wang, Chi Hay Tong, and Ingmar Posner. Vote3deep: Fast object detection in 3d point clouds using efficient convolutional neural networks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1355–1361. IEEE, 2017. 2
- [8] Runzhou Ge, Zhuangzhuang Ding, Yihan Hu, Yu Wang, Sijia Chen, Li Huang, and Yuan Li. Afdet: Anchor free one stage 3d object detection. *arXiv preprint arXiv:2006.12671*, 2020. 1, 3, 4
- [9] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 3
- [10] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 3
- [11] Benjamin Graham and Laurens van der Maaten. Sub-manifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307*, 2017. 4
- [12] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 4
- [14] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 6
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [16] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, 2019. 1, 2, 3, 4, 7
- [17] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 734–750, 2018. 3, 4
- [18] Zhidong Liang, Ming Zhang, Zehan Zhang, Xian Zhao, and Shiliang Pu. Rangercnn: Towards fast and accurate 3d object detection with range image representation, 2021. 3
- [19] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 3, 4
- [20] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, 2016. 3
- [21] Gregory P Meyer, Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, and Carl K Wellington. Lasernet: An efficient probabilistic 3d object detector for autonomous driving. In *CVPR*, 2019. 1, 3, 5, 7
- [22] Jiquan Ngiam, Benjamin Caine, Wei Han, Brandon Yang, Yuning Chai, Pei Sun, Yin Zhou, Xi Yi, Ouais Alsharif, Patrick Nguyen, et al. Starnet: Targeted computation for object detection in point clouds. *arXiv preprint arXiv:1908.11069*, 2019. 1
- [23] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9277–9286, 2019. 2
- [24] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018. 1, 2
- [25] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 1, 2, 4
- [26] 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*, 2017. 2
- [27] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 3
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016. 3

- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 4
- [30] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, 2020. 1, 2, 3, 7
- [31] Shaoshuai Shi, Chaoxu Guo, Jihan Yang, and Hongsheng Li. Pv-rcnn: The top-performing lidar-only solutions for 3d detection/3d tracking/domain adaptation of waymo open dataset challenges. *arXiv preprint arXiv:2008.12599*, 2020. 6, 7
- [32] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Point-rcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, 2019. 1, 2, 3, 5
- [33] Martin Simon, Stefan Milz, Karl Amende, and Horst-Michael Gross. Complex-yolo: Real-time 3d object detection on point clouds. *arXiv preprint arXiv:1803.06199*, 2018. 2
- [34] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020. 2, 3, 5, 6, 7
- [35] Yue Wang, Alireza Fathi, Abhijit Kundu, David Ross, Caroline Pantofaru, Tom Funkhouser, and Justin Solomon. Pillar-based object detection for autonomous driving. In *ECCV*, 2020. 1, 2, 3, 7
- [36] Waymo. Waymo’s 5th generation driver. <https://blog.waymo.com/2020/03/introducing-5th-generation-waymo-driver.html>. 1
- [37] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 2018. 1, 2, 3, 5
- [38] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7652–7660, 2018. 2
- [39] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1951–1960, 2019. 2
- [40] Dingfu Zhou, Jin Fang, Xibin Song, Chenye Guan, Junbo Yin, Yuchao Dai, and Ruigang Yang. Iou loss for 2d/3d object detection, 2019. 5
- [41] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 2, 3, 4
- [42] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *CORL*, 2019. 1, 2, 3, 4, 7
- [43] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, 2018. 1, 2, 3