

## OccAM’s Laser: Occlusion-based Attribution Maps for 3D Object Detectors on LiDAR Data

David Schinagl<sup>1,2</sup> Georg Krispel<sup>1</sup> Horst Possegger<sup>1</sup> Peter M. Roth<sup>3,4</sup> Horst Bischof<sup>1,2</sup>

{david.schinagl, georg.krispel, possegger, bischof}@icg.tugraz.at, peter.roth@tum.de

<sup>1</sup> Graz University of Technology <sup>2</sup> Christian Doppler Laboratory for Embedded Machine Learning

<sup>3</sup> Technical University of Munich <sup>4</sup> University of Veterinary Medicine, Vienna

### Abstract

While 3D object detection in LiDAR point clouds is well-established in academia and industry, the explainability of these models is a largely unexplored field. In this paper, we propose a method to generate attribution maps for the detected objects in order to better understand the behavior of such models. These maps indicate the importance of each 3D point in predicting the specific objects. Our method works with black-box models: We do not require any prior knowledge of the architecture nor access to the model’s internals, like parameters, activations or gradients. Our efficient perturbation-based approach empirically estimates the importance of each point by testing the model with randomly generated subsets of the input point cloud. Our sub-sampling strategy takes into account the special characteristics of LiDAR data, such as the depth-dependent point density. We show a detailed evaluation of the attribution maps and demonstrate that they are interpretable and highly informative. Furthermore, we compare the attribution maps of recent 3D object detection architectures to provide insights into their decision-making processes.

### 1. Introduction

Driven by the race to autonomous driving, the interest in 3D object detection is increasing in both, industry and the research community. Especially the detection of objects in point clouds acquired with LiDAR (Light Detection and Ranging) sensors is of great interest. On the one hand, this can be seen by the number and volume of published datasets that include LiDAR data, e.g. [3, 4, 8, 12, 42] and, on the other hand, by the number of recent publications that deal with this topic, e.g. [11, 14, 18, 33, 35, 36, 46, 49–54, 60]. In contrast, the explainability of these detectors based on deep neural networks is a rather unexplored area. For a safety-relevant area such as autonomous driving, however, it is of particular importance to make the decisions of these inher-

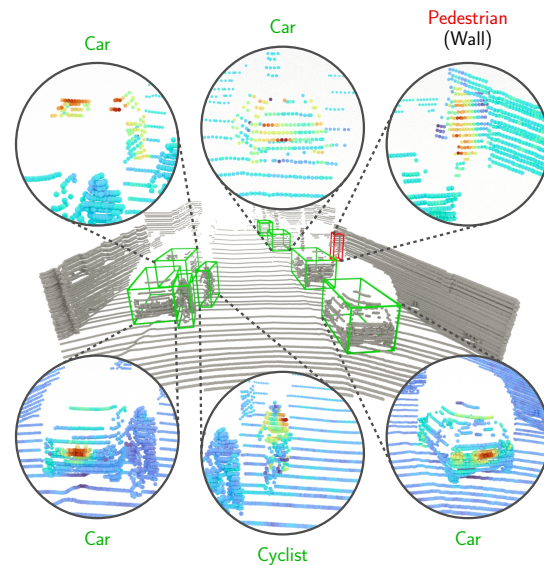


Figure 1. Our attribution maps show the importance of individual LiDAR points for the black-box detection model’s decisions.

ently opaque models more transparent.

An established technique for the analysis of models based on deep neural networks operating on image data is the generation of attribution or saliency maps, e.g. [30, 39, 41, 43, 55, 55]. The goal is to visualize the importance of input pixels on the model’s decision-making process. These have been proven useful for interpreting detections or analyzing the learned features. However, for models operating on point clouds there are only a few approaches that address the generation of point-level attribution maps, e.g. [10, 58]. These focus exclusively on the analysis of point cloud *classification* models. The analyzed models operate on artificially generated point clouds derived from CAD models, such as in the ModelNet dataset [48]. Moreover, these attribution map approaches assume prior knowledge of the architecture and require access to the gradients of the model. Thus, they are only applicable for the analysis of white-

box classification models. These properties prevent using them for the analysis of 3D object detectors on LiDAR point clouds. First, for 3D object detection, not only the *classification* of the object is important, but also its *localization*. These tasks usually go hand in hand and cannot be analyzed separately. Second, the detectors typically consist of complex architectures in which different parts contribute to the final detection, *e.g.* [33, 36], making gradient-based methods difficult or infeasible. In addition, a non-differentiable pre-processing step used in some models, *e.g.* [14, 50], prevents the backpropagation to the input points.

To the best of our knowledge, we present the first method to generate point-level attribution maps for 3D object detectors on LiDAR point clouds. For each detected object a separate map is created, as shown in Figure 1, which reflects the influence of each 3D point on the classification and localization of the object. Our method is applicable to black-box models and thus, in principle, to any detector, since neither prior knowledge of the model’s architecture nor access to its internals is required.

Inspired by perturbation approaches to generate saliency maps for image-based black-box models [23, 24, 55], we leverage the principle of *analysis by occlusion*. We propose *OccAM: Occlusion-based Attribution Maps* for 3D object detectors on LiDAR data. We estimate the importance of points by testing the model with randomly generated subsets of the input point cloud. The underlying assumption is that an object will be detected less accurately or not at all if areas of the input that are important for the detection have been removed or perturbed. However, the special characteristics of point clouds, like the unstructured nature and the depth-dependent point density in LiDAR data, poses new challenges compared to the analysis of image-based models. Thus, we propose a voxel-based sub-sampling strategy in which the sampling probability is adapted to the point density to challenge the detector appropriately. To evaluate the influence on the detections as precisely as possible, we further use a similarity metric that is optimized for the properties of 3D bounding boxes.

In a detailed analysis for PointPillars [14], we demonstrate that our attribution maps are interpretable and informative, *e.g.* allowing us to analyze the potential source of false detections. We show that by averaging multiple maps of a class, regions of particular importance can be derived. By comparing the average attribution maps of different state-of-the-art 3D object detectors, *i.e.* PointPillars, SECOND [49] and PV-RCNN [33], we further analyze whether differences can be detected.

## 2. Related Work

**LiDAR-based 3D Object Detection:** These methods can be roughly categorized into point-based, grid-based, and hybrid methods depending on how they handle the unstruc-

tured format of point clouds. **Point-based methods** directly process the unordered points [15, 25, 26, 35, 37, 51]. PointNet [25] and its successor PointNet++ [26] showed in their pioneering work a way to directly extract features using shared MLPs along with a global pooling function. In PointRCNN [35], PointNet++ features are used to segment foreground points and to generate proposals. **Grid-based methods**, on the other hand, discretize the 3D space into regular grids that can then be processed more easily [14, 46, 49, 50, 60]. In the end-to-end trainable VoxelNet [60], the 3D space is divided into equally sized voxels where the features within each voxel are learned using PointNet. To increase the efficiency of VoxelNet, an improved 3D sparse convolution is used in SECOND [49]. Instead of using voxels, PointPillars [14] introduced a column (pillar) representation to further reduce the complexity. **Hybrid methods** such as PV-RCNN [33, 34] or STD [52] simultaneously process point and voxel information to obtain multi-scale features and fine-grained localization [11, 20, 21, 53]. For example, Part-A<sup>2</sup> [36] uses point-wise features to predict intra-object parts and then aggregates the part information for box refinement to improve the robustness. Recently, several approaches based on transformers [45] have been proposed using point-based and / or grid-based features [9, 19, 22, 32].

**Attribution Maps for Image-based Models:** To analyze **white-box** models, where access to the internals of the model is possible, several gradient-based approaches were proposed, which leverage the gradient of the output with respect to the input image [39]. Due to the gradient backpropagation, however, these approaches tend to create noisy maps. Hence, some methods average multiple gradients [1, 40, 43] or try to improve the backpropagation using gradient calculation rules [29, 38, 41, 55, 56]. Fong *et al.* [6, 7] create a minimal perturbation mask that minimizes the probability of the target class by backpropagating the error. Another type of methods for white-box analysis uses the activations of intermediate layers [2, 27, 30, 59]. For instance, Grad-CAM [30] uses the spatial information available in convolutional layers and weights the activation scores at every location with the gradients of the class score w.r.t. the layer. Bakken *et al.* [2] use PCA to decompose the features of the final convolutional layer.

On the contrary, to analyze **black-box** models neither prior knowledge about the architecture nor access to the internals is available. These methods generate the attribution map by perturbing the input and analyzing the effects on the model output [17, 61] and are closely related to the Shapley Value [31], a concept from game theory where the goal is to determine a player’s contribution to the final outcome. Zeiler and Fergus [55] use a sliding window to block out input patches while observing the drop in the output score of a classification model. The superpixel-based LIME [28]

uses a linear surrogate model that estimates the impact on the output if the superpixels are removed. RISE [23] averages random binary masks according to the model’s output class probability for the masked inputs. This is extended in D-RISE [24] by a similarity metric allowing its application to detection models as well.

Overall, black-box analysis is computationally expensive, but of course provides the advantage of more general applicability. Inspired by the insights of these explanation methods for image data, we address the creation of attribution maps for black-box 3D object detectors to better understand their decision-making processes.

**Attribution Maps for Point Cloud-based Models:** In contrast to attribution maps for image-based models, there are only few approaches that investigate point cloud-based models. The existing methods focus on the analysis of white-box classification models, mainly based on PointNet architectures that process synthetic point clouds, like in the ModelNet dataset [48]. Gupta *et al.* [10] adapt gradient-based attribution maps for image classification models to analyze point cloud classification models trained on ModelNet data. Zheng *et al.* [58] propose a gradient-based method based on the idea of approximating the non-differentiable point dropping by moving the points towards the point cloud center. The underlying assumption is that points in the center are uninformative for the classification model and thus, point dropping can be approximated by this differentiable operation. However, this is only true for point clouds derived from CAD models like in the ModelNet dataset. In real LiDAR data, there are no points on back-facing surfaces of an object, but potentially important points within the object. For example, a driver sensed through the side window of a car can be a cue for the detector.

Distantly related are adversarial attacks on point cloud classification models. More specifically, methods that try to identify the most important points for the model in order to target them [13, 16, 47, 57]. However, they also only target white-box classification models, trained almost exclusively on ModelNet data, and determine the saliency using gradient-based methods.

In contrast to existing approaches, we are able to generate attribution maps for 3D object detection models that process wide-range LiDAR point clouds. Since the goal of 3D object detection is not only the classification of a single object, but the localization and classification of multiple objects, these detectors often consist of complex architectures, *e.g.* [33, 36]. This makes gradient-based analysis difficult or infeasible at all. In addition, the models often use a non-differentiable pre-processing step, *e.g.* [14, 50], preventing the backpropagation to the 3D points. Therefore, inspired by model-agnostic approaches for image-based models [23, 24, 55], we propose a method to generate attribution maps for black-box 3D object detection models.

### 3. Occlusion Analysis for LiDAR Data

Our goal is to create an attribution map for each detected object in a point cloud. The map should indicate the importance of each 3D point for the detector’s prediction regarding the specific object. We do not require any knowledge about the architecture of the model, nor access to parameters, features or gradients. This allows universal applicability and thus, the comparison of different architectures independent of their design. OccAM provides insights into the decisions of a black-box detector by leveraging occlusion analysis. The basic idea is to estimate the importance of input points by systematically removing them while observing changes in the output of the model being analyzed, as illustrated in Figure 2. Compared to methods that have successfully applied this principle to analyze 2D image models [23, 24, 55], analyzing LiDAR-based 3D object detectors poses the following new challenges, which we address:

- The unstructured format of point clouds makes the systematic sub-sampling of the input data more complex. We propose an efficient voxel-based sampling approach that allows masking of adjacent regions while preserving point-level information (Section 3.1).
- An inherent characteristic of LiDAR data is the non-uniform depth-dependent point density which must be taken into account during sub-sampling. Hence, we propose a density-aware sampling strategy to allow simultaneously analyzing all objects regardless of their distance to the LiDAR sensor (Section 3.2).
- Intersection over Union (IoU) alone is not sufficient to evaluate subtle detection changes precisely. Therefore, we use a similarity metric that evaluates translation, rotation and scaling individually (Section 3.3).

**Definitions & Notations:** Let  $\mathcal{X} = \{\mathbf{x}_j\}_{j=1\dots M}$  be a LiDAR point cloud, where each of the  $M$  unordered points  $\mathbf{x}_j \in \mathbb{R}^4$  consists of the 3D coordinates and the intensity/reflectance value. Furthermore, let  $F$  be a black-box object detection pipeline and  $F(\mathcal{X})$  be a set of detections  $\mathcal{D} = \{\mathbf{d}_k\}$  to be analyzed. Our goal is to generate an attribution map  $\Psi_k = \{\psi_{k,j}\}_{j=1\dots M}$  for each  $\mathbf{d}_k \in \mathcal{D}$  that reflects the contribution  $\psi_{k,j}$  of each point  $\mathbf{x}_j \in \mathcal{X}$  to the decision of the model  $F$  w.r.t.  $\mathbf{d}_k$ .

**Attribution Map Generation:** Let  $\Omega = \{\omega_j\}_{j=1\dots M}$ ,  $\omega_j \in \{0, 1\}$ , be a random binary mask. Then, we sub-sample  $\tilde{\mathcal{X}} = \{\mathbf{x}_j \in \mathcal{X} \mid \omega_j = 1\}$  according to  $\Omega$ . The detected objects  $F(\tilde{\mathcal{X}})$  in this sub-sampled point cloud are given by  $\tilde{\mathcal{D}} = \{\tilde{\mathbf{d}}_l\}$ . For a single detection  $\mathbf{d}_k \in \mathcal{D}$  and a single 3D point  $\mathbf{x}_j \in \mathcal{X}$ , we then define the contribution  $\psi_{k,j} \in \Psi_k$  as the expected value of the similarity metric  $S$

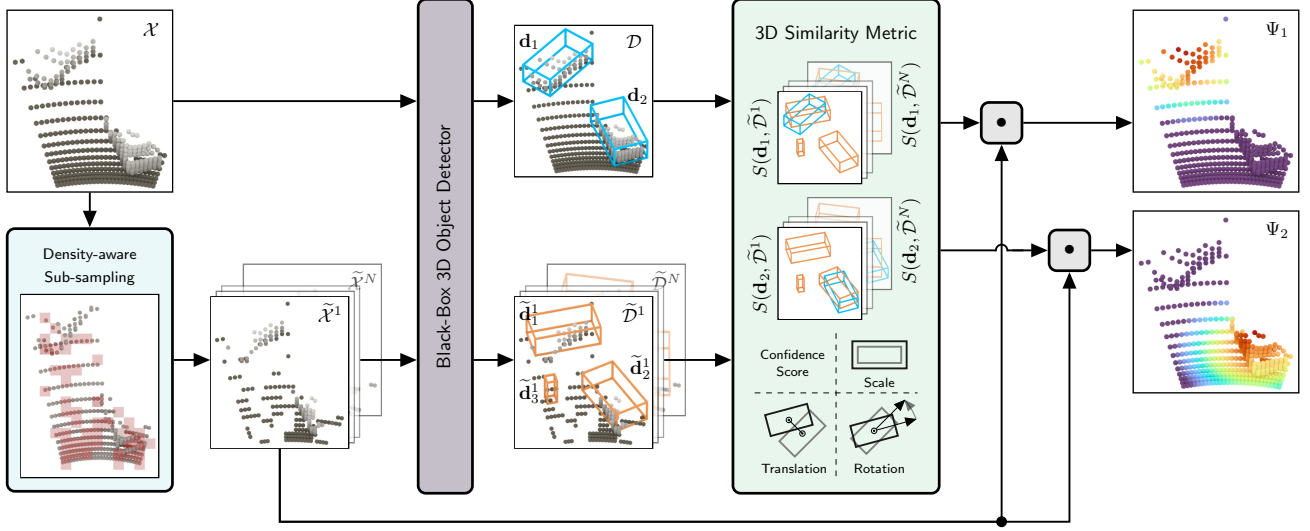


Figure 2. OccAM overview. We consider the specific point cloud characteristics during the sub-sampling of the input point cloud, such that the black-box 3D object detectors are challenged appropriately for all sensing ranges. Using a similarity metric tailored for this task, we can precisely analyze the changes in the detection output to obtain highly expressive attribution maps.

over all possible masks  $\Omega$ , where  $\mathbf{x}_j$  was part of the input, *i.e.*  $\omega_j = 1$ :

$$\psi_{k,j} = \mathbb{E}_{\forall \Omega} \left[ S(\mathbf{d}_k, \tilde{\mathcal{D}}) \mid \omega_j = 1 \right]. \quad (1)$$

$S(\mathbf{d}_k, \tilde{\mathcal{D}})$  evaluates how well the object described by  $\mathbf{d}_k$  was detected despite removing individual points, see Section 3.3. For this purpose, we compare  $\mathbf{d}_k$  with all detections in  $\tilde{\mathcal{D}} = \{\tilde{\mathbf{d}}_l\}$ , similar to the image-based analysis in [24]. The intuition behind this idea is that an object can be detected more accurately if important 3D points for the detection are visible. The expected value in Eq. (1) can be approximated using Monte Carlo sampling, as shown in [23]. Therefore, a set of masks  $\{\Omega^i\}_{i=1\dots N}$  is randomly generated from a probability distribution over the domain of possible masks and the corresponding sub-sampled point clouds  $\{\tilde{\mathcal{X}}^i\}_{i=1\dots N}$  are used as input to the detector. The resulting contribution scores are then aggregated to obtain

$$\psi_{k,j} \approx \frac{1}{\mathbb{E}[\omega_j] \cdot N} \sum_{i=1}^N S(\mathbf{d}_k, \tilde{\mathcal{D}}^i) \cdot \omega_j^i, \quad (2)$$

where  $\omega_j^i \in \Omega^i$  indicates the presence of the 3D point  $\mathbf{x}_j$  in point cloud  $\tilde{\mathcal{X}}^i$ . The contribution of each point is normalized by the expected number of samples in which the point was visible, *i.e.*  $\mathbb{E}[\omega_j] \cdot N$ . We can replace this normalization value by the empirical observation  $\sum_{i=1}^N \omega_j^i$ .

### 3.1. Voxel-based Sampling

Sub-sampling influences both, the accuracy of the resulting attribution maps and the number of necessary iterations  $N$  to get a good approximation of Eq. (2). If we remove too

few points, many iterations are necessary to gain information for the whole input. Removing large input areas, on the other hand, leads to coarse attribution maps. Thus, properly sub-sampling the point clouds is more complex in contrast to images, where the adjacency is well-defined and patches can easily be sampled in a regular grid.

Removing the 3D points individually would lead to a huge space of possible masks  $\Omega$  and  $N$  would be impractically high. In addition, the ability to evaluate the impact of hiding entire areas, such as the side of a vehicle, would be limited: The probability that all points of a specific area are removed while other areas remain untouched within a single mask would be very low.

Therefore, we propose to subdivide the 3D space along the  $x$ ,  $y$  and  $z$  axis into uniformly distributed voxels of equal edge length. We then group the points according to the voxel in which they are located. In each iteration, all points of the  $v^{\text{th}}$  voxel are kept jointly with probability  $P_v$ , by setting the corresponding mask entries to  $\omega_j = 1$ , or removed jointly,  $\omega_j = 0$ , with the counter probability  $1 - P_v$ . We argue that the pattern of several neighboring points being removed simultaneously is also very common in real LiDAR data, *e.g.* when an object is partially occluded by another object. To avoid always grouping the same points into the same voxel in each iteration, the voxel grid is randomly rotated and translated in all directions.

This sub-sampling approach has two advantages: First, by grouping the points into voxels we drastically reduce the number of necessary iterations  $N$ . Second, by randomly changing the voxel grid in each iteration it is still possible to obtain point-level insights.

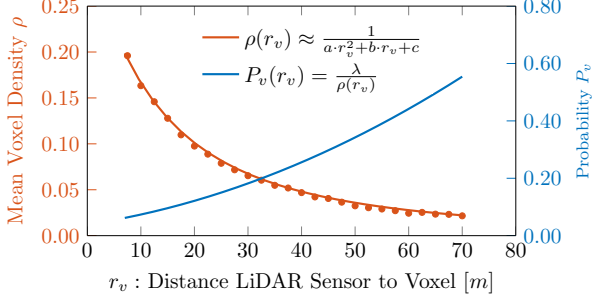


Figure 3. Visualization of the mean voxel density,  $\rho \in [0, 1]$ , for the KITTI dataset [8] (voxel size 20 cm), defined as the mean percentage of non-empty voxels within 1 m, depending on the distance  $r_v$ . We approximate this density by a second order polynomial  $\rho(r_v)$  and use its inverse, scaled by the hyperparameter  $\lambda$ , to determine the sampling probability  $P_v$  for each voxel.

### 3.2. Density-aware Sampling Strategy

As shown in Eq. (2), each mask is weighted by how accurately the object of interest was detected in the sub-sampled point cloud. Therefore, the quality of the final attribution map depends on how well the detector was challenged. If the probability  $P_v$  to keep the relevant voxels and corresponding points is too low and the object of interest is never detected, no information can be obtained. The same is true if  $P_v$  is too high and the object is always detected accurately. We must also take into account that the density of the points is not uniformly distributed, in contrast to the pixels in an image. The point density depends strongly on the distance to the LiDAR sensor. This means that for close objects, more voxels have to be removed than for distant objects to challenge the detector appropriately. From Figure 3, we see that a car placed at a distance of 10 m consists, on average, of about 4 times as many voxels as a car placed 50 m away. Thus, the contribution of a single voxel to the model’s decision increases with the distance from the LiDAR sensor. More formally, the distance between two neighboring laser rays that hit a surface orthogonal to the LiDAR sensor is given by  $r_{\text{obj}} \cdot \tan(\alpha)$ , where  $r_{\text{obj}}$  is the distance to the object and  $\alpha$  the angular resolution of the LiDAR sensor. Thus, assuming a square area of a certain size, the number of data points hitting that area decreases quadratically with the distance from the LiDAR sensor.

Therefore, we approximate the mean voxel density once for a dataset with a second order polynomial  $\rho(r_v) = (a \cdot r_v^2 + b \cdot r_v + c)^{-1}$ . When generating the attribution maps, we use the scaled inverse of this mean density to determine  $P_v$  for each voxel,

$$P_v(r_v) = \lambda \frac{1}{\rho(r_v)}, \quad (3)$$

where  $r_v$  is the distance from the LiDAR sensor to the voxel

center and  $\lambda$  is a hyperparameter. The probability  $P_v$  for each voxel thus depends on the mean voxel density  $\rho$ , scaled by  $\lambda$  to adjust the desired density in the sub-sampled point cloud. The relationship between the mean voxel density  $\rho$  and  $P_v$  is also shown in Figure 3. This sub-sampling strategy enables us to challenge the detector appropriately, regardless of the distance from the LiDAR sensor.

### 3.3. Similarity Metric

The similarity metric  $S$  measures how well the (presumed) object  $\mathbf{d}_k$  was still detected after sub-sampling the input. For this purpose,  $\mathbf{d}_k$  is compared with all detections  $\tilde{\mathcal{D}} = \{\tilde{\mathbf{d}}_l\}$  from the sub-sampled point cloud. Inspired by the image-based approach [24], we define  $S$  as maximum pairwise similarity between  $\mathbf{d}_k$  and the detections  $\tilde{\mathbf{d}}_l \in \tilde{\mathcal{D}}$ ,

$$S(\mathbf{d}_k, \tilde{\mathcal{D}}) = \max_{\tilde{\mathbf{d}}_l \in \tilde{\mathcal{D}}} s(\mathbf{d}_k, \tilde{\mathbf{d}}_l), \quad (4)$$

and adapt it to the requirements of 3D object detection on LiDAR data. Each detection  $\mathbf{d}_k$  is defined by a class label  $y_k$ , a confidence score  $c_k$  and a 3D bounding box  $\mathbf{b}_k$ . To compute the pairwise similarity between two detections,

$$s(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \prod_{A \in \mathbb{A}} s_A(\mathbf{d}_k, \tilde{\mathbf{d}}_l), \quad (5)$$

we define a set of sub-metrics  $\mathbb{A}$  to determine the individual differences in the label, the confidence score, and the consensus of the bounding boxes. The first two sub-metrics  $s_{\text{class}}$  and  $s_{\text{overlap}}$  ensure that only detections of the same class and overlapping detections,  $\text{IoU} > 0$ , can lead to a non-zero similarity:

$$s_{\text{class}}(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \begin{cases} 1 & \text{if } y_k = \tilde{y}_l \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

$$s_{\text{overlap}}(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \begin{cases} 1 & \text{if } \text{IoU}(\mathbf{b}_k, \tilde{\mathbf{b}}_l) > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

To ensure that a potential increase in the confidence score compared to the original detection, *i.e.*  $\tilde{c}_l > c_k$ , does not lead to a penalization, we directly use  $\tilde{c}_l$  for

$$s_{\text{conf}}(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \tilde{c}_l. \quad (8)$$

To determine the correspondence of the pairwise 3D bounding boxes as precisely as possible, we evaluate the rotation, translation and scaling separately, inspired by the true positive metrics of nuScenes [3]. Each bounding box  $\mathbf{d}$  is given by its center coordinates  $\mathbf{C}$ , its box dimensions and its yaw angle  $\theta$ . The translation similarity is computed as the Euclidean distance between the bounding box centers,

$$s_{\text{translation}}(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \max\left(1 - \|\mathbf{C}_k - \tilde{\mathbf{C}}_l\|, 0\right). \quad (9)$$

For the scale similarity, we align the bounding boxes w.r.t. their center and yaw angle and then compute their IoU,

$$s_{\text{scale}}(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \text{IoU}(\mathbf{b}_k^{\text{aligned}}, \tilde{\mathbf{b}}_l^{\text{aligned}}). \quad (10)$$

Finally, the orientation similarity is determined by the smallest yaw angle difference between  $\theta_k$  and  $\tilde{\theta}_l$  in radians,

$$s_{\text{orientation}}(\mathbf{d}_k, \tilde{\mathbf{d}}_l) = \max(1 - |\theta_k - \tilde{\theta}_l|, 0). \quad (11)$$

The advantage of these separate scores is that the localization of the object can be evaluated more accurately than by using only the IoU. Additionally, this allows the creation of attribution maps for individual properties.

## 4. Experiments

We demonstrate OccAM’s insights into the behavior of a PointPillars [14] model in detail and compare the behaviors across several recent detectors, including SECOND [49] and PV-RCNN [33].

### 4.1. Implementation Details

Our experimental setup<sup>1</sup> is based on the open-source toolbox OpenPCDet [44]. Unless otherwise stated, all detectors are trained on the KITTI dataset [8] with their default configurations and using the standard training policies. We use the standard split [5] into training and validation samples. All shown attribution maps are exclusively from the 3769 samples of the validation set, which we also use to estimate the voxel density  $\rho$ . We use a voxel size of 20 cm (cube-shaped) in all our experiments.

We determine the values of the hyperparameters  $N$  and  $\lambda$  empirically. We found that  $N = 3000$  is a good compromise between the runtime and the quality of the attribution maps (details in supplemental material). The parameter  $\lambda$  was chosen such that at a distance of  $r_v = 25$  m the sampling probability  $P_v$  is equal to 0.15, unless explicitly stated. This configuration gives stable results regardless of the detector’s architecture.

Since all  $N$  probing steps are independent of each other, the entire process is fully parallelizable. Furthermore, the perturbed point clouds can be efficiently generated on the CPU. The runtime is thus dominated by the detector runtime on the GPU. For PointPillars, the analysis of a KITTI point cloud takes on average 50 s (for  $N = 3000$ ) on a desktop PC using a single NVIDIA® GeForce® RTX 3090 GPU.

### 4.2. Qualitative Results

We demonstrate that our attribution maps are interpretable and informative. Figure 4 shows examples for both, true and false positive detections. Considering the *car* class we have found that the license plates, due to their high

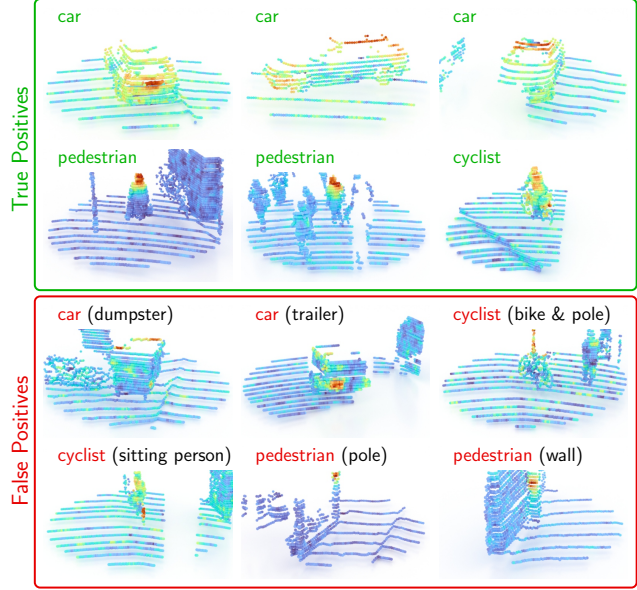


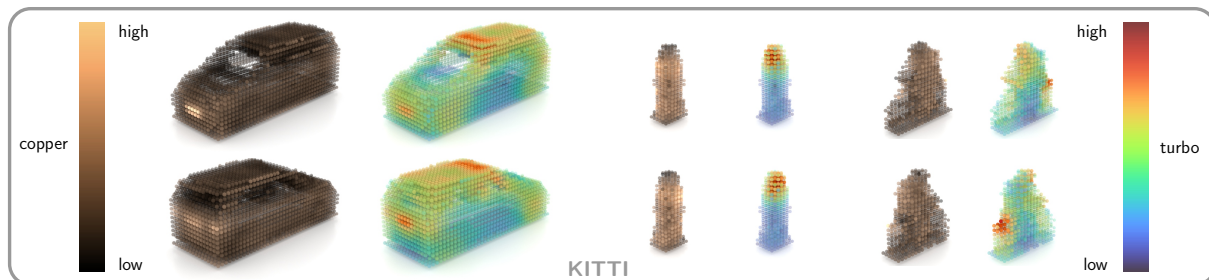
Figure 4. Attribution map examples for PointPillars detections on KITTI. Warmer colors (turbo colormap) denote higher contribution of a point to this detection. Crops are for visualization only.

reflectance, represent a very discriminative feature. Furthermore, the roof railing, the A-pillar and the B-pillar are very distinct features of this class. We argue that these areas have a high identification value due to their shape and the transparency of the windows. Especially, if a car is only seen from the side or the license plate is occluded by other objects, these areas are very important for detection. For *pedestrians*, as expected, the head-shoulder silhouette proved to be a highly discriminatory feature and was consistently shown as the most important area. This area is also very relevant for the detection of *cyclists* in addition to the reflectors on the front and rear of the bike. The high interpretability is also evident for false detections. For example, in the case of the mis-detection of the dumpster as a vehicle, the top-rear area was obviously incorrectly interpreted as roof railing. In the case of the false detection of the trailer, the detector was obviously misled by the license plate. Note also how well our attribution maps allow us to analyze single objects even in crowded scenes, *e.g.* see the pedestrian in the group. These examples show that our attribution maps can provide useful insights into the model behavior and are also well suited to analyze false detections.

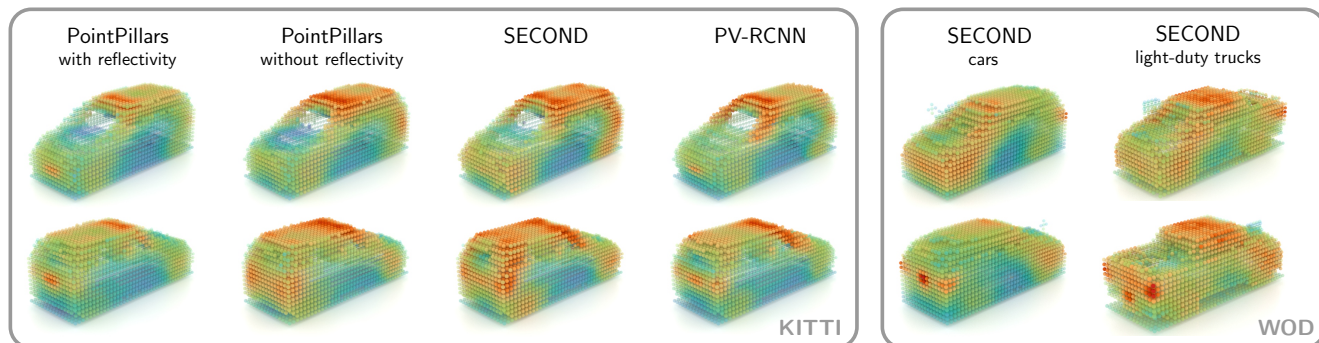
### 4.3. Average Attribution Maps

To verify these findings from the individual attribution maps, we also average the maps of the specific classes. To do so, we first scale all boxes and the associated points to a uniform size and then align them w.r.t. their center and yaw angle. We then voxelize the resulting point cloud and aver-

<sup>1</sup><https://github.com/dschinagl/occam>



(a) Average attribution maps (turbo-colored) for PointPillars [14] trained and evaluated on KITTI [8]. We also show the average LiDAR reflectivity / intensity values (copper-colored). From left to right: cars, pedestrians and cyclists.



(b) Comparison of PointPillars [14] (trained with and without reflectivity values), SECOND [49] and PV-RCNN [33] trained and evaluated on KITTI [8]. Note how the average attribution maps reflect the granularity of the underlying detector architecture.

(c) SECOND [49] trained and evaluated on the Waymo Open Dataset (WOD) [42]. For visualization purposes, we clustered the vehicles by size.

Figure 5. Average attribution maps (turbo-colored) on (a) (b) KITTI [8] and, for comparison, on (c) Waymo Open Dataset (WOD) [42]. We adjust the opacity of the voxels according to their attribution score. Best viewed on screen.

age the attribution values of the individual points within a voxel. The resulting maps are shown in Figure 5a. Our previous findings from the individual attribution maps, such as the importance of license plates or the head-shoulder silhouette, are also clearly visible in these averaged maps.

#### 4.4. Metrics

Evaluating the explainability of attribution maps is still an open question, even for image-based analysis methods. We adopt the *point dropping* approach [58] for point cloud classification models. The basic idea is to iteratively remove the most important points and re-evaluate the model w.r.t. IoU and confidence. If the attribution maps are expressive, then the model’s performance should degrade faster than by just removing points randomly. On the contrary, if the least important points are removed first, then the performance should degrade slower than during random removal.

For 3D object detection, however, this idea unfairly penalizes the random removal baseline. Thus, we apply point dropping only to points within actual detections. For this we use the original black-box detector output as pseudo ground truth. Iteratively, we then remove the most/least important points, re-run the detector and compare its output to the pseudo ground truth. Since the random baseline also

removes points only within detections, this is a fair comparison: these points are obviously more important than unrelated points outside, far-away from any detection.

As shown in Figure 6, removing the most/least important points first leads to a significantly faster/slower model degradation in contrast to random removal. Thus, our attribution maps are expressive. Note that both IoU and the confidence score do not reach zero, even if all points within the bounding boxes are removed. This is caused by the interesting finding that surrounding features also notably contribute to the detector’s decision, *e.g.* see Figure 7.

#### 4.5. Importance of Density-aware Sampling

To demonstrate the importance of our density-aware sampling strategy, we compare our point density-derived distance-dependent choice of  $P_v(r_v)$  with a fixed choice for  $P_v$ . Figure 8 shows the mean similarity score during the attribution map computation as a function of the distance between the object of interest and the LiDAR sensor. With a fixed choice of  $P_v$  the mean similarity score strongly depends on the distance to the object. Intuitively, close objects can be detected more accurately than distant objects. However, if the objects are recognized too precisely or only very rarely during this probing phase, these samples do not

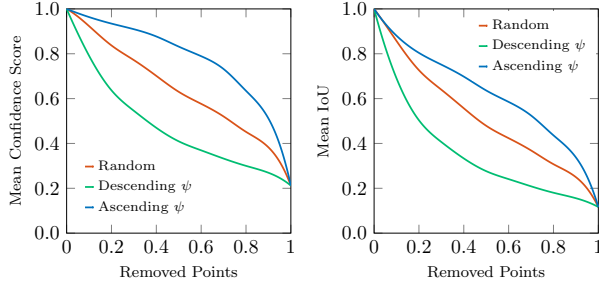


Figure 6. Point dropping evaluation. Points within each bounding box are removed in descending or ascending order of importance, and the resulting detections are compared with the original detections using confidence score and IoU.

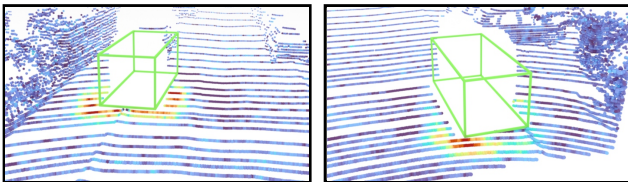


Figure 7. Detections despite empty bounding boxes: The attribution maps show that distinctive cutouts on the ground are responsible for these PointPillars detections. Since detections with no points are obviously more sensitive to perturbation, we set  $\lambda = 0.3$  at  $r_v = 25$  m for these visualizations.

lead to any information gain. This means that for fixed  $P_v$  there is only a small distance range in which meaningful attribution maps would be created.

In contrast, our density-aware sampling strategy leads to an almost constant similarity score independent of the distance to the object. This allows the computation of attribution maps for all objects regardless of their distance to the sensor with only a single probing process.

#### 4.6. Comparison of Detector Architectures

The attribution maps generated by OccAM are well suited to compare the behavior of different object detectors. In Figure 5b, we show the average car attribution maps for PointPillars [14], SECOND [49] and PV-RCNN [33]. For pedestrians and cyclists we found out that the head-shoulder silhouette is the most important feature, regardless of the analyzed architecture. These additional maps and further detectors are included in the supplemental material.

For all detectors, the points on the roof are an important feature to distinguish a vehicle from other classes while the door areas are of minor relevance. Another distinct characteristic of cars are the well-reflecting license plates and lights. Here, however, the detectors differ: while PointPillars focuses on the license plates, these are less important for SECOND. We also trained PointPillars without reflectivity values, which then focuses more on the edges of the

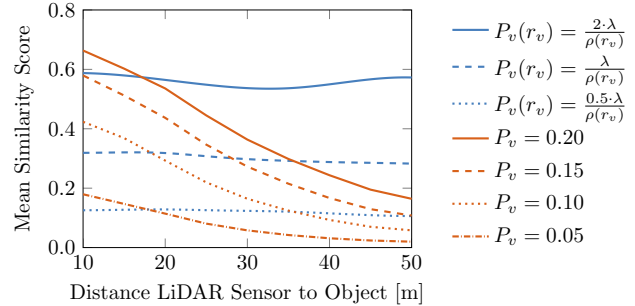


Figure 8. Different sampling schemes for a PointPillar model. We plot the mean similarity score over the distance between the object of interest and the LiDAR sensor. With a fixed  $P_v$ , attribution maps would only be expressive in a narrow distance range.

object, similar to SECOND. For PV-RCNN, the vehicle’s license plates and lights appear to be of balanced importance.

The influence of the underlying architecture and the point cloud processing method can best be seen at the A-pillars. Due to the transparent windows of the car, the A-pillar represents a very distinct feature. While this feature is of little importance for PointPillars, due to its coarse (pillar) processing, it is notably more important for the voxel-based SECOND model. PV-RCNN leverages these fine structures best, since it processes point and voxel information jointly.

To further demonstrate OccAM’s general applicability, we provide the average attribution maps of SECOND on the Waymo Open Dataset (WOD) [42] in Figure 5c. Since, in contrast to KITTI, WOD has only a single *vehicle* class (which also includes trucks, etc.), we cluster the individual maps by vehicle size before averaging. These maps show both, similarities (focus on roof and A-pillars) and differences (varying importance license plates, lights and fender regions) between KITTI (Figure 5b) and WOD (Figure 5c).

## 5. Conclusion

We proposed OccAM, the first method to generate attribution maps for black-box 3D object detectors on LiDAR data. Our attribution maps are well interpretable, informative and provide useful insights into the behavior of the detectors. While some behavior was already anticipated, e.g. head-shoulder silhouette, we also found unexpected insights, e.g. the importance of the A-pillars and the local surroundings. With OccAM, we offer the community a way to analyze the behavior of any 3D object detector on any LiDAR dataset.

**Acknowledgements** This work was partially funded by the Austrian FFG project iLIDS4SAM (878713) and by the Christian Doppler Laboratory for Embedded Machine Learning. Peter M. Roth received support from the German Federal Ministry of Education and Research (BMBF) within the international future AI lab AI4EO (01DD20001).



## References

- [1] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, 10(7), 2015. [2](#)
- [2] Marianne Bakken, Johannes Kvam, Alexey A. Stepanov, and Asbjørn Berge. Principal Feature Visualisation in Convolutional Neural Networks. In *Proc. ECCV*, 2020. [2](#)
- [3] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multi-modal dataset for autonomous driving. In *Proc. CVPR*, 2020. [1](#), [5](#)
- [4] Ming-Fang Chang, John Lambert, Patsorn Sangkloy, Jagjeet Singh, Slawomir Bak, Andrew Hartnett, De Wang, Peter Carr, Simon Lucey, Deva Ramanan, et al. Argoverse: 3D Tracking and Forecasting with Rich Maps. In *Proc. CVPR*, 2019. [1](#)
- [5] 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 *NeurIPS*, 2015. [6](#)
- [6] Ruth Fong, Mandela Patrick, and Andrea Vedaldi. Understanding Deep Networks via Extremal Perturbations and Smooth Masks. In *Proc. ICCV*, 2019. [2](#)
- [7] Ruth Fong and Andrea Vedaldi. Interpretable Explanations of Black Boxes by Meaningful Perturbation. In *Proc. ICCV*, 2017. [2](#)
- [8] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *Proc. CVPR*, 2012. [1](#), [5](#), [6](#), [7](#)
- [9] Tianrui Guan, Jun Wang, Shiyi Lan, Rohan Chandra, Zuxuan Wu, Larry Davis, and Dinesh Manocha. M3DeTR: Multi-representation, Multi-scale, Mutual-relation 3D Object Detection with Transformers. *arXiv:2104.11896*, 2021. [2](#)
- [10] Ananya Gupta, Simon Watson, and Hujun Yin. 3D Point Cloud Feature Explanations Using Gradient-Based Methods. In *Proc. IJCNN*, 2020. [1](#), [3](#)
- [11] Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure Aware Single-stage 3D Object Detection from Point Cloud. In *Proc. CVPR*, 2020. [1](#), [2](#)
- [12] Xinyu Huang, Peng Wang, Xinjing Cheng, Dingfu Zhou, Qichuan Geng, and Ruigang Yang. The ApolloScape Open Dataset for Autonomous Driving and its Application. *PAMI*, 42:2702–2719, 2020. [1](#)
- [13] Jaeyeon Kim, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Minimal Adversarial Examples for Deep Learning on 3D Point Clouds. In *Proc. ICCV*, 2021. [3](#)
- [14] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast Encoders for Object Detection from Point Clouds. In *Proc. CVPR*, 2019. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [15] Zhichao Li, Feng Wang, and Naiyan Wang. LiDAR R-CNN: An Efficient and Universal 3D Object Detector. In *Proc. CVPR*, 2021. [2](#)
- [16] Daniel Liu, Ronald Yu, and Hao Su. Extending Adversarial Attacks and Defenses to Deep 3D Point Cloud Classifiers. In *Proc. ICIP*, 2019. [3](#)
- [17] Scott M. Lundberg and Su-In Lee. A Unified Approach to Interpreting Model Predictions. In *NeurIPS*, 2017. [2](#)
- [18] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid R-CNN: Towards Better Performance and Adaptability for 3D Object Detection. In *Proc. ICCV*, 2021. [1](#)
- [19] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel Transformer for 3D Object Detection. In *Proc. ICCV*, 2021. [2](#)
- [20] Zhenwei Miao, Jikai Chen, Hongyu Pan, Ruiwen Zhang, Kaixuan Liu, Peihan Hao, Jun Zhu, Yang Wang, and Xin Zhan. PVGNet: A Bottom-Up One-Stage 3D Object Detector With Integrated Multi-Level Features. In *Proc. CVPR*, 2021. [2](#)
- [21] Jongyoun Noh, Sanghoon Lee, and Bumsub Ham. HVPR: Hybrid Voxel-Point Representation for Single-stage 3D Object Detection. In *Proc. CVPR*, 2021. [2](#)
- [22] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3D Object Detection with Pointformer. In *Proc. CVPR*, 2021. [2](#)
- [23] Vitali Petsiuk, Abir Das, and Kate Saenko. RISE: Randomized Input Sampling for Explanation of Black-box Models. In *Proc. BMVC*, 2018. [2](#), [3](#), [4](#)
- [24] Vitali Petsiuk, Rajiv Jain, Varun Manjunatha, Vlad I. Morariu, Ashutosh Mehra, Vicente Ordonez, and Kate Saenko. Black-box Explanation of Object Detectors via Saliency Maps. In *Proc. CVPR*, 2021. [2](#), [3](#), [4](#), [5](#)
- [25] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proc. CVPR*, 2017. [2](#)
- [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] Sylvestre-Alvise Rebuffi, Ruth Fong, Xu Ji, and Andrea Vedaldi. There and Back Again: Revisiting Backpropagation Saliency Methods. In *Proc. CVPR*, 2020. [2](#)
- [28] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proc. KDD*, 2016. [2](#)
- [29] Adria Ruiz, Antonio Agudo, and Francesc Moreno-Noguer. Generating Attribution Maps with Disentangled Masked Backpropagation. In *Proc. ICCV*, 2021. [2](#)
- [30] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization. In *Proc. ICCV*, 2017. [1](#), [2](#)
- [31] L. S. Shapley. A Value for n-Person Games. In *Contributions to the Theory of Games II*, pages 307–317. Princeton University Press, 1953. [2](#)
- [32] Hualian Sheng, Sijia Cai, Yuan Liu, Bing Deng, Jianqiang Huang, Xian-Sheng Hua, and Min-Jian Zhao. Improving 3D Object Detection with Channel-wise Transformer. In *Proc. ICCV*, 2021. [2](#)

- [33] 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 *Proc. CVPR*, 2020. 1, 2, 3, 6, 7, 8
- [34] 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:2102.00463*, 2021. 2
- [35] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud. In *Proc. CVPR*, 2019. 1, 2
- [36] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From Points to Parts: 3D Object Detection from Point Cloud with Part-aware and Part-aggregation Network. *PAMI*, 43:2647–2664, 2021. 1, 2, 3
- [37] Weijing Shi and Ragunathan Rajkumar. Point-GNN: Graph Neural Network for 3D Object Detection in a Point Cloud. In *Proc. CVPR*, 2020. 2
- [38] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning Important Features Through Propagating Activation Differences. In *Proc. ICML*, 2017. 2
- [39] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. *arXiv:1312.6034*, 2013. 1, 2
- [40] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. SmoothGrad: Removing noise by adding noise. *arXiv:1706.03825*, 2017. 2
- [41] Jost Tobias Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for Simplicity: The All Convolutional Net. In *Proc. ICLR Workshops*, 2015. 1, 2
- [42] Pei Sun, Henrik Kretschmar, 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 *Proc. CVPR*, 2020. 1, 7, 8
- [43] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic Attribution for Deep Networks. In *Proc. ICML*, 2017. 1, 2
- [44] OpenPCDet Development Team. OpenPCDet: An open-source toolbox for 3D object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020. 6
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 2
- [46] Yue Wang, Alireza Fathi, Abhijit Kundu, David A. Ross, Caroline Pantofaru, Tom Funkhouser, and Justin Solomon. Pillar-based Object Detection for Autonomous Driving. In *Proc. ECCV*. Springer, 2020. 1, 2
- [47] Matthew Wicker and Marta Kwiatkowska. Robustness of 3D Deep Learning in an Adversarial Setting. In *Proc. CVPR*, 2019. 3
- [48] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *Proc. CVPR*, 2015. 1, 3
- [49] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely Embedded Convolutional Detection. *Sensors*, 18(10):3337, 2018. 1, 2, 6, 7, 8
- [50] Bin Yang, Wenjie Luo, and Raquel Urtasun. PIXOR: Real-time 3D Object Detection from Point Clouds. In *Proc. CVPR*, 2018. 1, 2, 3
- [51] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3DSSD: Point-based 3D Single Stage Object Detector. In *Proc. CVPR*, 2020. 1, 2
- [52] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. STD: Sparse-to-Dense 3D Object Detector for Point Cloud. In *Proc. ICCV*, 2019. 1, 2
- [53] Maosheng Ye, Shuangjie Xu, and Tongyi Cao. HVNet: Hybrid Voxel Network for LiDAR Based 3D Object Detection. In *Proc. CVPR*, 2020. 1, 2
- [54] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3D Object Detection and Tracking. In *Proc. CVPR*, 2021. 1
- [55] Matthew D. Zeiler and Rob Fergus. Visualizing and Understanding Convolutional Networks. In *Proc. ECCV*, 2014. 1, 2, 3
- [56] Jianming Zhang, Sarah Adel Bargal, Zhe Lin, Jonathan Brandt, Xiaohui Shen, and Stan Sclaroff. Top-Down Neural Attention by Excitation Backprop. *IJCV*, 126(10):1084–1102, 2018. 2
- [57] Qiang Zhang, Jiancheng Yang, Rongyao Fang, Bingbing Ni, Jinxian Liu, and Qi Tian. Adversarial Attack and Defense on Point Sets. *arXiv:1902.10899*, 2019. 3
- [58] Tianhang Zheng, Changyou Chen, Junsong Yuan, Bo Li, and Kui Ren. Pointcloud Saliency Maps. In *Proc. ICCV*, 2019. 1, 3, 7
- [59] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning Deep Features for Discriminative Localization. In *Proc. CVPR*, 2016. 2
- [60] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In *Proc. CVPR*, 2018. 1, 2
- [61] Luisa M. Zintgraf, Taco S. Cohen, Tameem Adel, and Max Welling. Visualizing Deep Neural Network Decisions: Prediction Difference Analysis. In *Proc. ICLR*, 2017. 2