Multimodal 3D Object Detection on Unseen Domains

Hegde, Deepti; Lohit, Suhas; Peng, Kuan-Chuan; Jones, Michael J.; Patel, Vishal M.

TR2025-078 June 07, 2025

Abstract

LiDAR datasets for autonomous driving exhibit biases in properties such as point cloud density, range, and object dimensions. As a result, object detection networks trained and evaluated in different environments often experience performance degrada- tion. Domain adaptation approaches assume access to unannotated samples from the test distribution to address this problem. However, in the real world, the exact conditions of deployment and access to samples representative of the test dataset may be unavailable while training. We argue that the more realistic and challenging formulation is to require robustness in performance to unseen target domains. We propose to address this problem in a two-pronged manner. First, we leverage paired LiDAR-image data present in most autonomous driving datasets to perform multimodal object detection. We suggest that working with multimodal features by leveraging both images and LiDAR point clouds for scene understanding tasks results in object detectors more robust to unseen domain shifts. Second, we train a 3D object detector to learn multimodal object features across different distributions and promote feature invariance across these source domains to improve generalizability to unseen target domains. To this end, we propose CLIX3D, a multimodal fusion and supervised contrastive learning framework for 3D object detection that performs alignment of object features from sameclass samples of different domains while pushing the features from different classes apart. We show that CLIX3D yields state- of-the-art domain generalization performance under multiple dataset shifts.

IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop 2025

© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Mitsubishi Electric Research Laboratories, Inc. 201 Broadway, Cambridge, Massachusetts 02139

Multimodal 3D Object Detection on Unseen Domains

Deepti Hegde¹ Suhas Lohit² Kuan-Chuan Peng² Michael J. Jones² Vishal M. Patel¹

¹Johns Hopkins University ²Mitsubishi Electric Research Laboratories

{dhegde1, vpatel36}@jhu.edu {slohit, kpeng, mjones}@merl.com

Abstract

LiDAR datasets for autonomous driving exhibit biases in properties such as point cloud density, range, and object dimensions. As a result, object detection networks trained and evaluated in different environments often experience performance degradation. Domain adaptation approaches assume access to unannotated samples from the test distribution to address this problem. However, in the real world, the exact conditions of deployment and access to samples representative of the test dataset may be unavailable while training. We argue that the more realistic and challenging formulation is to require robustness in performance to unseen target domains. We propose to address this problem in a two-pronged manner. First, we leverage paired LiDAR-image data present in most autonomous driving datasets to perform multimodal object detection. We suggest that working with multimodal features by leveraging both images and LiDAR point clouds for scene understanding tasks results in object detectors more robust to unseen domain shifts. Second, we train a 3D object detector to learn multimodal object features across different distributions and promote feature invariance across these source domains to improve generalizability to unseen target domains. To this end, we propose CLIX^{3D}, a multimodal fusion and supervised contrastive learning framework for 3D object detection that performs alignment of object features from sameclass samples of different domains while pushing the features from different classes apart. We show that CLIX^{3D} yields stateof-the-art domain generalization performance under multiple dataset shifts.

1. Introduction

LiDAR point clouds provide direct, albeit sparse, 3D geometric information of a scene through accurate depth estimates, while RGB images can provide high resolution 2D color and texture information. Recent years have seen the release of several large scale multimodal datasets containing registered LiDAR point clouds and RGB images [7, 40, 13, 9, 19], which have aided the development of numerous deep neural networks for solving perception tasks for autonomous navigation such as segmen-



Figure 1. (Left) Overview of the contrastive learning framework of **CLIX^{3D}**. We use multi-source training and supervised contrastive learning between region level features to improve robustness of 3D LiDAR and LiDAR-image based object detection networks. (Middleleft) Comparison of 3D detection precision performance of **CLIX^{3D}** against LiDAR-only detection and SOTA fusion methods. (Far-ight) Comparison of **CLIX^{3D}**'s robustness to unseen distributions with SOTA domain generalization work.

tation [33, 58, 2] and object detection[48, 37, 39, 38, 26]. Until recently, the best performing 3D object detectors operate only on LiDAR point clouds [38, 37, 26, 39], despite the availability of paired image-point cloud data in most autonomous driving datasets [7, 13, 40, 19, 9]. The image and LiDAR modalities offer complementary information while describing the same underlying scene, making the multimodal fusion approach a natural choice for training a detection network. Several recent works have demonstrated the effectiveness of LiDAR-image fusion for 3D scene understanding [28, 11, 31, 21], and outperformed LiDAR-only methods. In this work, we contribute to this direction and propose a multi-stage LiDAR-image fusion method for 3D object detection. In particular, we focus on how multimodal object detection networks can be trained to be robust when evaluated on scenes "in-the-wild" (see Fig. 1).

LiDAR point clouds collected from different environments vary widely between one another in terms of point cloud density, scene properties and object dimensions due to different modes of capture, locations, weather conditions, *etc.* LiDAR scenes from different datasets have large differences that are easily visible.

While the severity of the cross-dataset distribution gap is not

as large when dealing with images, changes in the time-of-day and adverse weather during capture can result in datasets that are biased to specific conditions. This becomes a problem when performing scene understanding tasks in real-world scenarios, where the conditions of capture at test time may differ from the training dataset. The performance of neural network models trained on data from a particular distribution drops when shown samples from a different distribution [45]. Including image information helps not only the baseline performance, but also in training networks robust to distribution shifts. Images provide dense color and texture information, while LiDAR point clouds provide sparse but accurate depth measurements. Data from both these modalities is prone to varying highly depending on the environmental conditions. Images are particularly prone to illumination changes, while LiDAR scenes are not. The density and range of point cloud scenes are dependent on the sensor specification, and tend to differ largely between datasets, while the corresponding images are unaffected. This can be observed in Fig. 2, where LiDAR-image pairs from the Waymo Open Dataset [40], nuScenes [7] and KITTI [13] are shown side-by-side to demonstrate the differing properties of each modality in differing conditions. The complementary nature of both the information and sources of domain shift makes LiDAR-RGB fusion the natural choice to improve the robustness of 3D object detection networks to distribution shift. DeepFusion [28] briefly explores this capability by evaluating their proposed LiDAR-RGB fusion pipeline on location-specific domain shift within the same dataset. In contrast, we address more challenging cross-dataset domain shifts and propose object-level contrastive learning to build more robust detectors.

Most current works that address the problem of domain gaps assume the existence of an unlabeled dataset that is representative of the domain used for evaluation (target domain). These methods further adapt a source network using target samples and at times labelled source samples to perform unsupervised domain adaptation (UDA) [49, 35, 8, 46, 32, 50, 18, 17]. This formulation can be unrealistic, as the target domain characteristics are often unknown and can change dynamically. Additionally, the adapted model is suitable only for the target domain it is trained for, and must be re-trained for every new target distribution. In contrast, we formulate and propose a method to address the domain generalization (DG) problem, which is a more practical and challenging setting for the 3D object detection task [6, 57]. In the DG setting, no information about the target domain(s) is available during training. The guiding principle of DG is that training a model able to generalize over diverse source domains can help generalize to unseen target distributions [4].

Gulrajani and Lopez-Paz [14] have shown that performing empirical risk minimization across multiple diverse datasets results in generalizable models given enough data and an effective model selection strategy. DG through domain-invariance has also received significant attention for 2D images, especially for image recognition [53], including methods based on contrastive



Figure 2. LiDAR and image examples from the KITTI [13], Waymo [40], nuScenes [7] datasets in different environmental conditions (listed at the beginning of each row). Images are particularly prone to illumination conditions, while LiDAR scenes differ in density between datasets. The first column shows the image scene, the second column is a bird's-eye-view (BEV) of the LiDAR scene and the last row shows the LiDAR scene projected to the front image.

learning. However, extending these ideas to 3D object detection is not straightforward due to relative complexity of the scene compared to samples from image classification datasets, as well as the more difficult task of performing denser predictions. We suggest that this problem setting requires operating on more local regions in the scene. We focus on individual objects in a scene by **utilizing region features provided by existing object detection networks to enforce domain invariance** in the feature extraction backbones of 3D object detectors.

We thus propose to tackle performance degradation of object detectors evaluated on unseen domains in a two-pronged manner. First, we propose a multimodal fusion method for RGB and LiDAR data to leverage complementary information across modalities for increased robustness. Second, we design a novel framework called Contrastive LiDAR+Image Cross-Domain 3D Object Detection – **CLIX**^{3D} for training LiDAR and LiDAR-RGB 3D object detectors to improve performance on unseen domains by learning domain invariant representations of objects in a scene. Fig. 1 shows an overview of the contrastive training method. The features from object regions of each scene are aligned in a common embedding space according to

category to train the network to learn properties of the object that are invariant to domain specific conditions. We are the first to propose a multi-source multimodal setting to address robustness to unseen domains. **CLIX^{3D}** outperforms multi-source training and single-source DG on most domain shift scenarios.

- 1. We propose **CLIX^{3D}**, a supervised contrastive learning framework for learning domain-invariant features that are suited for 3D object detection by operating on small regions rather than at the global level.
- By leveraging various autonomous driving datasets acquired using different sensors under different conditions and environments, we show that CLIX^{3D} results in improved domain generalization to unseen target domains.
- We propose a new LiDAR-RGB fusion approach MS-Fusion – as a part of CLIX^{3D} that improves same-domain and cross-domain 3D detection performance, compared to state-of-the-art (SOTA) fusion methods.

2. Related works

2.1. 3D object detection

There are fundamentally three main approaches to performing purely LiDAR-based object detection with deep networks in literature – (a) operating directly on the irregular point clouds [51, 15], (b) first voxelizing or pillarizing the point clouds and operating on the voxels [26, 47, 39, 12], and (c) projecting the point clouds to 2D, e.g., bird's-eye view (BEV) [30, 20]. Each network may also be characterized as a single-stage or a twostage network. Single stage object detectors directly regress and classify bounding box predictions from features whereas two-stage object detectors use an additional refinement head operating on region proposals. The SOTA single stage object detector SE-SSD [56] uses IoU-based matching to align a studentteacher network to perform soft filtering. VoxelRCNN [12] proposes a novel region-of-interest pooling and box refinement and outperforms previous point and voxel based methods. In this work, we use VoxelRCNN as the base object detector due to its superior performance as a two-stage object detector. (See Sec. 4.2) VoxelRCNN also takes less time to train than more recent methods. PVRCNN, FocalsConv, and CaDNN take 5, 4, and 15 hours respectively to train on KITTI while VoxelRCNN takes around 2 hours under the same settings. This is an important factor when training on a large number of source samples.

LiDAR-image fusion: Several works seek to incorporate image data to boost 3D object detection performance. Early methods such as PointPainting [43] and PointAugmenting [44] concatenate the input LiDAR point cloud with semantic scores or deep features before being passed into the 3D network. Current single-view, convolutional approaches improve on these by performing fusion at the feature level mid-way through the network. DeepFusion [28] performs fusion by aligning the LiDAR and image scenes through inverting the LiDAR augmentations before fusing at the feature level using a cross-attention mechanism for feature selection. FocalsConv [11] proposes a method for sparse convolutions with learned sparsity which they leverage for LiDAR-image fusion. TransFusion [3] and LIFT [54] are transformer based fusion approaches. BEVFusion [31] leverages multi-view images to create bird's-eye-view image feature maps for fusion. In this work, we focus on single view convolutional fusion approaches due to their compatibility with existing LiDAR detectors and datasets.

2.2. Cross-domain transfer

A significant number of works in perception literature address dealing with domain shift, by either adapting to specific targets, or training robust networks.

Domain adaptation for 3D object detection: Recent works that address unsuerpvised domain adaptation (UDA) for 3D object detection include ST3D [49], ST3D++ [50], MLCNet [32], and scalable pseudo-labeling [8]. However, these works use a **single** source domain and assume access to annotated or unannotated samples from the target domain that are used for model adaptation. While Wang's method [45] does not train on target data, it uses the bounding box statistics of the target dataset to resize predicted bounding boxes to perform adaptation. Some works train with multiple source domains [52] or models [42], but still operate in the UDA formulation. Zhang *et al.* [55] show that multi-source training can aid in boosting detection performance, but do not perform cross-dataset transfer.

Domain generalization: In contrast, in this paper, we study a related but novel setting for 3D object detection - multi-source DG where the target domain is completely unseen during training. We show that multi-source and multimodal training from diverse domains can lead to improved generalizability in 3D object detection models. We design a supervised contrastive learning method CLIX^{3D} that can be used with LiDAR and LiDAR-RGB detector architectures to lead to higher generalization performance. Lehner et al. [27] recently propose 3D-VField, an adversarial data augmentation strategy, for DG with point clouds. However, they only show single-source DG and do not discuss domain invariance. Several other works perform single source DG [24] and DG for semantic segmentation [23, 36]Although the idea of contrastive learning has been employed for DG in 2D image recognition [53], it is not straightforward to apply such ideas for 3D object detection from LiDAR point clouds and we discuss the challenges in design and implementation in Sec. 3.

3. CLIX^{3D}

We first describe the multi-stage LiDAR-image feature fusion approach, followed by the supervised contrastive learning framework for multi-source object detection.

3.1. LiDAR-image fusion

We design a multi-stage fusion module, called **MSFusion**, for LiDAR and image data that are incorporated into existing 3D object detection networks to consider image information during the feature extraction stage. Consider a LiDAR scene and image pair $\{p,q\}$ such that $p \in \mathbb{R}^{g \times 3}$ and $q \in \mathbb{R}^{h \times w \times 3}$, where q denotes the number of points in the LiDAR scene and $h \times w$ denotes the spatial dimensions of the image. Let ϕ be a convolution-based encoder for the image modality decomposed into s stages $\{\phi_i\}_{i=1}^s$ and ψ be a 3D feature extraction network that processes point clouds decomposed into $\{\psi_i\}_{i=1}^s$. The image is passed to ϕ to obtain a set of feature maps F_{ϕ_i} of successively smaller feature dimensions from each stage. The 3D encoder consists of a series of 3D convolution blocks and set abstraction layers that process a voxelized point cloud to give 3D feature maps consisting of stacked voxel features. The voxelized LiDAR scene is passed to ψ to obtain sets of voxels that undergo successive set abstractions to give sets of features with reducing spatial dimensions F_{ψ_i} . For each stage pair $\{F_{\phi_i}, F_{\psi_i}\}$ we find the voxel feature to pixel feature correspondence before performing deep feature fusion.

Finding voxel-pixel correspondence. Given the camera matrix, each voxel may be projected to the image feature plane to obtain pixel-voxel level correspondences. Since the number of points that represent the scene are fewer than the number of pixels ($m < h \times w$), we assign a neighborhood of pixels of size k to each projected point, with the pixel location of the point as the center of the neighborhood. Each LiDAR point now corresponds to a [$k \times k$] patch of the image. With reducing spatial dimensions of each 2D feature map, the camera intrinsics are scaled accordingly.

LiDAR augmentation reversal. The use of data augmentations are standard practice in 3D LiDAR-only object detection networks, as they aid in performance [39, 12, 47]. However, these augmentations are unique to the LiDAR modality and include transformations such as rotation and scaling, which affect the correspondence of points with pixels. To ensure that the 3D encoder benefits from these augmentations and that accurate voxel-pixel mapping takes place, we keep a record of each augmentation type and degree of augmentation and perform reversal in the voxel feature space before the fusion step.

Deep feature fusion. Consider the set of LiDAR-image feature maps $\{F_{\phi_i}, F_{\psi_i}\}_{i=1}^s$. After applying the reverse augmentation on the voxel feature maps and obtaining voxel-pixel correspondences for each stage, each feature vector path is average pooled, concatenated along the feature dimension and passed to a projection layer that embeds the feature into the 3D embedding space. The result is passed to the next stage of the object detection network.

3.2. Multi-source 3D object detection

We now focus on the later stages of the object detection network. The methodology for contrastive multi-source training applies to both single and multimodal object detectors. For the purpose of clarity, we consider LiDAR object detectors here. The task of 3D object detection is to localize and classify salient objects in the scene. This is done by estimating location and dimensions of a 3D bounding box that describe the position of each object, and classifying it into one of the predefined classes. Let a point cloud dataset of n samples be represented as $\mathcal{D} = \{p_i, l_i\}_{i=1}^n$, where p_i is the LiDAR scene consisting of 3D points that describe the scene, and l_i consists of object labels such that each object in the scene is described with the label $\{x, y, z, dx, dy, dz, \theta, c\}$ where $\{x, y, z\}$ denote the center of the bounding box, and $\{dx, dy, dz\}$ denote the displacement of the bounding box edges from the center to represent bounding box dimension, θ denotes the orientation, and c denotes the class label.

Borrowing the notation from [1], we consider a set of LiDAR point cloud datasets \mathcal{D}_e consisting of n_e labeled samples of cobject categories such that $\mathcal{D}_e = \{p_i^e, l_i^e\}_{i=1}^{n_e}$, where $e \in \mathcal{E}_{tr}$ denotes the data acquisition environments of the training datasets. We assume the occurrence of the same object categories across training and evaluation datasets. We call the labeled datasets seen during training as source datasets which are samples from source domain distributions, and denote by $n_s = |\mathcal{E}_{tr}|$, the number of source domains available for training. The goal of DG is to train a network on these samples to perform well on unseen environments from unknown distributions (*i.e.*, target domains).

Region-level supervised contrastive learning. A LiDAR scene contains some regions that provide little information, due to the sparse nature of the point clouds. This makes global scene-level comparisons of samples from different distributions difficult beyond simplistic attributes such as point cloud density. We propose to focus on meaningful regions in the LiDAR scene, *i.e.*, regions with objects of interest such as cars, bicycles, and pedestrians. Point cloud objects from the scene belonging to the same category tend to share geometric properties. A point cloud of a car is identifiable as a car even in the absence of visual attributes such as color and texture. Thus, we hope to learn a universal representation of a LiDAR "Car" object that is consistent across domains. We wish to train a network to learn region features that capture this domain invariant representation of each object category. From our multi-source dataset with samples from different environments, we have access to different data representations of the same object category. By pushing together objects from the same category and pulling apart objects from different categories, features across domains are aligned in a common embedding space.

Feature selection. Anchor-based networks provide region proposals for a given scene. These proposals consist of features corresponding to foreground as well as background regions. Given the ground-truth bounding box information, we can identify the most accurate proposals and use them for contrastive training. Consider a batch of input point clouds from different source domains $\{p_i^{e_j}, l_i^{e_j}\}_{i=1}^{b}$ where j indexes the source domains and b is the batch size. Let the RPN head be denoted by ϕ_{rpn} . After feature extraction and region pooling, ϕ_{rpn} generates a set of ROIs represented by features f and predicted bounding boxes B which we denote by $\{\{f_k^{e_j}, B_k^{e_j}\}_{k=1}^{n_r}\}_{j=1}^{n_s}$ for a given set of samples in a batch, where e denotes the environment/domain of the sample, and n_r and n_s are the numbers of ROIs and domains,



Figure 3. Description of the proposed **CLIX^{3D}** for generalizing 3D object detectors to unseen target domains. Samples from multiple diverse source domains are used to train the object detectors. Following multi-stage deep feature fusion in the feature extraction backbone, supervised contrastive learning is applied on ROI features obtained after the pooling step which encourages domain invariance. As illustrated on the right side, region features that belong to the same class, but from different domains are encouraged to be closer together in feature space, while those that belong to different classes are pushed apart.

respectively. Many region proposals are redundant or belong to background areas, so we select the foreground ROIs closest to the nearest ground truth label, and sample *m* background proposals. We now have a subset of ROI features, ground truth class labels *c*, and domain labels from samples across different source distributions denoted by $\{\{f_k^{e_j}, c_k^{e_j}\}_{k=1}^{n_r+m}\}_{j=1}^{n_s}$.

Supervised contrastive learning. We now have the region features of objects across samples in the batch, which originate from different domains. We enforce domain invariance in the learned features by minimizing the distance between region features of objects from the same category and maximizing the distance between objects of different categories on ROIs prior to being passed to further layers. Consider the set of region features and corresponding category labels from scenes across different source distributions $\{\{f_k^{e_j}, c_k^{e_j}\}_{k=1}^{n_r}\}_{j=1}^{n_s}$. To draw an analogy to contrastive learning frameworks like [22], each region from each LiDAR scene is treated as a training sample, with the input batch consisting of N regions (samples). Each sample has many positive pairs, decided by the number of samples sharing the same object category, and an arbitrary number of negative pairs. Unlike [22], we do not have augmented views, since we obtain samples directly in the features space. We now show the way we perform contrastive learning which can handle multiple positives. To enforce domain invariance, distances between region features across LiDAR scenes are optimized by comparing region features from two source domains.

Optimizing feature similarity: We use the cosine similarity for the unit-normalized feature vectors computed using $\sin = \langle f^{e_1}, f^{e_2} \rangle$. Inspired by the multimodal self-supervised

contrastive learning method CLIP [34], we utilize binary cross entropy (BCE) loss to optimize feature similarity. We modify the formulation to handle multiple positive pairs and compute the sum of BCE loss for sample pairs that are constructed for a particular class across all object categories, such that:

$$\mathcal{L}_{con} = \sum_{i=1}^{N_c} \frac{1}{|R(i)|} \sum_{\substack{j \in R(i); h \in \mathbf{R} \\ k, l \in D(i)}} -\omega_i \{q_i y_j \log(s) + (1 - y_j) \log(1 - s)\},$$

where $s = \sigma(\sin(f_j^{e_k}, f_h^{e_l})), \sigma(\cdot)$ denotes the softmax operation, q_i is the weight value for positive samples computed by the ratio of number of negative and positive samples, **R** denotes the indices of all foreground regions and sampled background regions, and $y_j = \mathbb{1}(j = h)$ with $\mathbb{1}(\cdot)$ as the indicator function, N_c denotes the number of classes, R(i)and D(i) denote the regions of interest and their respective domain labels. Optimizing these loss functions trains the feature extraction and region proposal branches to produce aligned feature representations for all regions in the batch that belong to the same object category, including those from different domains. These region features are then passed to a bounding box regression branch that performs localization, and a classification branch that categorizes each proposal.

Losses for object detection: In addition to the contrastive loss, the network is supervised by the localization (\mathcal{L}_{loc}) and classification (\mathcal{L}_{cls}) losses present in their original training frameworks. The contrastive loss \mathcal{L}_{con} is added to the existing RPN loss \mathcal{L}_{rpn} and made to share a similar scale in value. The

final loss \mathcal{L} below is used to train the network:

$$\mathcal{L} = \mathcal{L}_{loc} + \mathcal{L}_{cls} + \mathcal{L}_{rpn} + \mathcal{L}_{con}.$$
 (1)

4. Experiments

4.1. Datasets

In autonomous driving, the environment of data capture is characterized by a variety of variables, including geographic location, type of LiDAR sensor, position of the sensor, and weather conditions. These vary between datasets, providing a diverse set of training data with which to train a learning algorithm. For the purpose of DG, we seek diverse datasets that express different forms of distribution shift, e.g., from average size of vehicles to point cloud density. For our experiments, we choose four popular autonomous driving LiDAR-image datasets for 3D object detection: Lyft [19], KITTI [13], Waymo Open Dataset [40], and nuScenes [7]. In the supplement, we compare various properties of these datasets. We construct our domain shift scenarios to cover shifts from larger annotation-rich source datasets to smaller annotation-poor datasets as well as the reverse case. Due to their varying data formats, to train on multiple datasets at a time, we convert all the datasets to the format of KITTI for both training and evaluation.

4.2. Experimental setup

Domain shift scenarios. To demonstrate our DG method, we explore cross-dataset distribution shift that covers a change in location, time of day, weather conditions, and rates of LiDAR return. We conduct transfer from multiple source datasets to a single target dataset at a time in the form $\{S_1, S_2\} \rightarrow T$ where S denotes the source dataset and T denotes the target dataset. In particular, we conduct experiments with the settings:

1. $S_1 =$ Waymo, $S_2 =$ nuScenes; $T \in \{$ Lyft, KITTI $\}$

2. $S_1 = Lyft, S_2 = KITTI; T \in \{Waymo, nuScenes\}$

This covers a broad variety of types of cross-dataset shift, including transferring to and from the smaller annotation-poor datasets such as KITTI and between datasets with dense and sparse point clouds. For baseline experiments, we also conduct single-source domain transfer experiments. A network trained on a particular source dataset setting may be evaluated on any target distribution.

3D object detection networks: We demonstrate our multimodal domain generalization framework primarily on VoxelRCNN [12], which we choose due to its superior baseline object detection performance, the symmetry of the 3D feature extraction network to that of ResNet [16], and the existence of a region proposal network, which we require to leverage region features. We also demonstrate the pure-LiDAR approach on Part- A^2 for consistent comparison.

LiDAR network architecture: The 3D feature extraction backbone consists of 4 stages of sequential sparse 3D convolutional blocks that produce feature maps of dimensions 16, 32, 64, and 128 with successively decreasing spatial dimensions. In terms of voxel features, set abstraction at each stage results in sparser voxels in scaled down ranges. This architecture may be observed in several 3D object detector networks such as [38, 26, 5, 12]. **Image network architecture:** We leverage a ResNet-50 backbone [16] pre-trained for image segmentation on COCO [29] under DeepLabv3 [10] as the 2D image extraction backbone. This branch is trained end-to-end.

Baselines and oracle: We compare our multi-source contrastive learning-based DG framework against single-source and multi-source direct transfer baselines, in which a source-trained model is directly evaluated on the target datasets. That is, in "direct transfer," the object detection networks are trained only with the standard classification and bounding box regression loss functions, without contrastive loss. We also compare our method against the single-source DG work 3D-VField for the backbone network Part- A^2 which uses adversarial data augmentation. As the implementation of 3D-VField has not been made public, we compare our results against those reported in the paper. We also provide the "oracle" results for both LiDAR and LiDAR-image detectors, *i.e.*, when both source and target domains are the same, and trained using only classification and bounding box regression losses in Table 1.

Evaluation metric: We evaluate the object detection networks in average precision (AP) in the KITTI format. We report performance on the object categories of "Car," "Pedestrian," and "Cyclist" with respective 3D IoU thresholds of 0.7, 0.5, 0.5 for the KITTI, Lyft, and Waymo datasets and 0.5, 0.25, 0.25 for the nuScenes dataset, respectively. For a fair comparison with the reported numbers of [27], we include evaluation results on Waymo with the 3D IoU threshold 0.5 for the "moderate" category in Table 3. Each object is further categorised in terms of difficulty as "easy," "moderate," or "hard" based on the level of occlusion, truncation, and distance from the camera. We follow [13] for this evaluation convention.

Implementation details: We train the object detection networks in our supervised contrastive learning framework on an equal number of samples from each source dataset, and ensure that an equal number of scenes from each dataset make up an input batch. As the quality of the selected region features depends on the proposal network, we pretrain the network for 30 epochs on the source datasets without contrastive loss before training for a further 30 epochs with the contrastive loss. We use standard data augmentation methods for object detectors such as global scaling, rotation, and ground-truth sampling. We train the networks with a batch size of 32, Adam optimizer [25] and cyclic learning rate scheduling. The initial learning rate is 0.01 with a weight decay of 0.01. For implementing the object detectors, we follow the codebase OpenPCDet [41]. We train each model on four 48GB NVIDIA A40 GPUs.

5. Results and discussion

We present and discuss the results of the MSFusion method on each LiDAR-image dataset as well as the full

Dataset	Modality	Method	Car			Pedestrian			Cyclist		
			easy	moderate	hard	easy	moderate	hard	easy	moderate	hard
KITTI	L	VoxelRCNN	92.36	82.74	80.04	65.85	58.67	52.35	88.05	69.53	64.92
	L + I	VoxelRCNN - DeepFusion*	92.19	82.59	80.06	60.56	52.82	46.63	84.79	62.70	58.61
	L + I	VoxelRCNN - FocalsConv	92.55	82.92	80.34	61.01	53.28	48.02	89.17	70.26	65.67
	L+I	VoxelRCNN - MSFusion (ours)	92.22	83.32	82.45	66.19	58.41	52.63	89.10	70.23	66.34
Waymo	L	VoxelRCNN	76.19	74.59	72.17	57.89	58.25	57.88	65.79	66.36	66.26
	L + I	VoxelRCNN - DeepFusion*	76.41	74.82	72.51	56.67	56.40	56.05	64.97	65.80	63.29
	L + I	VoxelRCNN - FocalsConv	76.40	76.07	72.81	57.18	57.62	57.93	68.39	66.60	66.06
	L+I	VoxelRCNN - MSFusion (ours)	76.17	74.74	72.38	58.49	58.31	57.66	73.89	71.97	68.04
nuScenes	L	VoxelRCNN	22.05	18.35	17.30	8.99	7.98	7.52	8.48	7.42	6.69
	L + I	VoxelRCNN - DeepFusion*	24.44	20.52	19.14	13.99	12.59	11.95	9.62	8.97	8.01
	L + I	VoxelRCNN - FocalsConv	20.74	17.46	16.28	8.59	7.80	7.49	11.55	10.69	9.47
	L+I	VoxelRCNN - MSFusion (ours)	25.69	21.50	20.41	9.45	8.58	8.55	12.28	11.44	10.01
Lyft	L	VoxelRCNN	85.91	75.44	71.57	45.77	33.74	32.75	73.95	54.93	51.23
	L + I	VoxelRCNN - DeepFusion*	88.67	78.64	76.68	53.53	40.97	39.88	77.55	60.67	56.77
	L+I	VoxelRCNN - FocalsConv	88.99	79.26	77.29	56.59	40.86	40.10	78.91	62.87	58.84
	L+I	VoxelRCNN - MSFusion (ours)	89.11	79.86	78.00	55.53	43.19	42.27	78.56	62.51	59.15

Table 1. 3D average precision (AP) results of VoxelRCNN [12] trained on KITTI [13], Waymo [40], nuScenes [7], and Lyft [19]. Compares performance of training the network with just LiDAR (L) or LiDAR and image (L+I) using DeepFusion [28] and FocalsConv [11] with the proposed MSFusion approach. * indicates the method is re-implemented by us. This table shows that (a) fusion of LiDAR and RGB inputs leads to improved performance, and (b) the proposed fusion method outperforms the SOTA methods.

CLIX^{3D}ramework evaluated on various domain shifts. We point to the supplementary material for qualitative results.

LiDAR-RGB fusion. In Table 1 we compare the proposed MSFusion method for fusing LiDAR and RGB with the SOTA convolution-based multimodal object detectors that use single view images, DeepFusion [28] and FocalsConv [11], as well as the pure-LiDAR baseline for VoxelRCNN. In this setting, the training and testing domains are the same, and we convert each dataset to KITTI FOV format with standard splits. We observe consistent improvement from the pure-LiDAR baseline, particularly in the "hard" category of objects, which are characterized by high occlusion, truncation, and larger distances from the camera. We show that multi-stage fusion of LiDAR-RGB features outperforms or matches the SOTA in most cases. In cases where the proposed method does not perform the best, it comes second with only a small margin.

Detection on unseen domains. We demonstrate our multimodal framework for four domain shift scenarios. We compare our method against direct transfer baselines and that of 3D-Vfield [27]. By direct transfer, we mean that the network is trained on the source dataset and evaluated on the target dataset with no adaptation or changes to training. In Table 2 we show the performance of VoxelRCNN [12] trained under the domain shift listed in Sec. 4.2. We point out observations that are two-fold. Firstly, introducing image information improves domain transfer performance in most cases. Particularly for the "Cyclist" category, we observe a large performance improvement when evaluating on KITTI and Lyft. This supports the suggestion that multimodal detectors are more robust to unseen distributions. Secondly, training the object detector under the multi-source

contrastive framework also aids in improved robustness to out-of-distribution samples. Compared to single source direct transfer cases, we observe a significant performance improvement in most cases when the network is trained under **CLIX**^{3D}, and this trend continues when applied to the multimodal detector. This indicates that encouraging domain invariance in a diverse training set can aid generalization. In Table 3, we compare our LiDAR-only contrastive learning framework on the detector Part-A² with the domain generalization work 3D-Vfield [27]. Since the code is not publicly available, we compare against the reported domain scenario and single category. We outperform [27], but point out that their learnable data augmentations may be incorporated into our framework.

Ablation study. We examine each component of the multimodal multi-source contrastive learning framework and demonstrate empirically the role they play in training a generalizable object detection network. In Table 4, we compare the transfer performance to the Lyft dataset under different modules of the proposed framework. We find that adding the contrastive objective aids in multi-source training.

Limitations. There are certain limitations with our approach. We observe some inconsistent performance when training with multiple sources. In the case of transferring to the Waymo dataset, we observe that the multi-source trained model performs worse than Lyft/Waymo direct in the "Car" category, but outperforms all single-source networks in the "Pedestrian" and "Cyclist" categories. This could be due to the fact that the average size of cars in the Lyft dataset is closer than KITTI to that of the Waymo dataset. We believe this problem can be mitigated by training with more diverse source data, and point

Turin / Trat	M- 1-1:4-	Method	Car		Pedestrian			Cyclist			
Train / Test	wodanty		easy	moderate	hard	easy	moderate	hard	easy	moderate	hard
Waymo / KITTI	L	DT	16.12	14.95	14.43	60.39	53.01	47.89	67.51	57.40	51.96
	L+I	DT	6.33	5.98	_ 5.48	_57.26	50.69	45.93	_67.75	62.91	58.11
nuScenes / KITTI	L	DT	3.05	2.89	2.58	34.03	28.92	25.93	7.05	4.60	4.50
huseeles/ KITTI	L+I	DT	12.51	8.92	7.41	31.31	25.87	23.52	28.56	16.45	15.84
	L	CLIX ^{3D}	37.90	35.94	36.77	58.92	51.81	45.94	70.43	57.91	54.35
waymo + nuscenes / KITTI	L+I	CLIX ^{3D}	37.72	35.22	36.37	65.07	57.94	51.84	68.71	55.33	52.35
	L	DT	60.08	41.99	41.62	38.43	23.80	23.55	41.95	22.17	20.08
Waymo / Lyft	L+I	DT	59.33	42.31	41.95	37.25	21.88	21.86	46.71	25.99	23.39
		DT	22.67	13.77	14.12	7.55	3.98	4.23	7.15	3.84	3.21
nuScenes / Lyft	L+I	DT	36.15	23.74	23.96	7.07	4.67	4.51	11.07	5.80	5.08
	 L	CLIX ^{3D}	56.85	39.66	38.79	35.57	22.74	22.47	45.86	26.96	26.11
Waymo + nuScenes / Lyft	L+I	CLIX ^{3D}	63.87	44.18	43.02	39.38	25.22	24.87	37.82	21.64	20.46
	L	DT	14.10	14.53	13.89	32.56	33.16	31.32	30.46	27.46	24.51
KIIII/ waymo	L+I	DT	18.62	18.63	17.67	22.60	23.48	23.66	36.35	35.44	32.46
T C / XX/	L	DT	46.78	45.21	41.53	27.99	28.69	28.85	21.34	22.81	22.47
Lyft / waymo	L+I	DT	54.98	54.93	51.73	37.06	38.54	38.24	31.29	31.41	32.16
	L	CLIX ^{3D}	28.13	28.06	25.18	43.12	43.43	42.51	47.17	43.53	41.33
KIIII + Lyft / Waymo	L+I	CLIX ^{3D}	48.13	46.57	43.65	47.62	46.93	45.59	48.06	49.03	46.05
	L	DT	24.75	20.70	19.12	9.26	8.09	8.09	3.45	3.67	3.55
KITTI/ nuscenes	L+I	DT	28.23	23.15	21.56	9.63	8.21	8.17	4.30	4.62	3.90
	L	DT	21.36	17.35	16.21	6.33	6.42	6.53	4.36	5.25	4.74
Lyft / nuScenes	L+I	DT	24.25	19.95	18.63	9.03	7.57	7.44	6.55	5.89	5.30
	 L	CLIX ^{3D}	25.02	20.60	19.41	5.63	5.49	4.32	3.44	3.62	2.96
KIIII + Lyft / nuScenes	L + I	CLIX ^{3D}	31.06	25.41	24.09	4.82	4.79	4.82	7.10	6.80	6.07

Table 2. 3D average precision (AP) of unseen domain transfer to the Waymo, nuScenes, KITTI and Lyft datasets. DT = direct transfer, L=LiDAR, I=Image. We show that multi-source training of multimodal detectors under **CLIX**^{3D}'s contrastive learning framework results in improved robustness to unseen domains.

Turin (Trat	Madaad		AP				
Train / Test	Method	Car	Pedestrian	Cyclist			
KITTI / Waymo	DT 3D-Vfield [27]	47.64 56.80	26.45	53.83			
Lyft / Waymo	DT	56.97	22.44	48.42			
KITTI + Lyft / Waymo	DT CLIX ^{3D}	65.34 66.59	32.41 34.68	67.78 66.28			

Table 3. 3D detection results for IoU thresholds $\{0.5, 0.5, 0.25\}$ on Li-DAR cross-domain experiments on Part-A² [39]. DT = direct transfer.

Multi-source	Multimodal	Contractive loss	Car				
With source	Withhodai	Contrastive 1055	easy	moderate	hard		
\checkmark	×	×	59.11	40.53	39.53		
×	\checkmark	×	59.95	41.39	40.55		
\checkmark	\checkmark	×	62.07	43.08	41.79		
\checkmark	\checkmark	\checkmark	63.87	44.18	43.02		

Table 4. Ablation study for Waymo+nuScenes transfer to the Lyft dataset with 3D precision values for the "Car" category. Single source experiments indicate source is Waymo.

out that in the DG setting, *it is not possible to choose a single source that is closer in distribution to the target data*, in which case the multi-source training that performs better on average

is preferred. We also observe that the "Pedestrian" and "Cyclist" categories of the nuScenes dataset are difficult to transfer to, resulting in relatively lower precision values. This is due to the sparse nature of objects from these categories in the considered datasets, which we hope to explicitly address in future work.

6. Conclusion

For the problem of unseen domain shift in 3D object detection, we show that MSFusion, our proposed LiDAR+RGB fusion method, outperforms prior methods and that multimodal fusion improves robustness when the target domains are unseen during training. We further improve generalization by proposing a multi-source training framework **CLIX**^{3D} with simple yet effective region-level contrastive learning, which promotes invariance among features belonging to the same class across domains and pushes features from different classes apart. **CLIX**^{3D} beats state-of-the-art methods in most cases under many distribution shifts for multiple object classes and datasets.

Acknowledgments This work was done when Deepti Hegde was an intern at MERL. For this work, Deepti Hegde, Suhas Lohit, Kuan-Chuan Peng, and Michael J. Jones were supported by

Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA. Vishal M. Patel was partially supported by the NSF CAREER Award 2045489.

References

- Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. arXiv preprint arXiv:1907.02893, 2019. 4
- [2] Mehmet Aygun, Aljosa Osep, Mark Weber, Maxim Maximov, Cyrill Stachniss, Jens Behley, and Laura Leal-Taixé. 4d panoptic lidar segmentation. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 5527–5537, 2021. 1
- [3] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 1090–1099, 2022. 3
- [4] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. Advances in neural information processing systems, 19, 2006. 2
- [5] Prarthana Bhattacharyya and Krzysztof Czarnecki. Deformable pv-rcnn: Improving 3d object detection with learned deformations. arXiv preprint arXiv:2008.08766, 2020. 6
- [6] Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. In Advances in Neural Information Processing Systems. Curran Associates, Inc., 2011. 2
- [7] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Q. Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11618–11628, 2020. 1, 2, 6, 7
- [8] Benjamin Caine, R. Roelofs, Vijay Vasudevan, Jiquan Ngiam, Yuning Chai, Z. Chen, and Jonathon Shlens. Pseudo-labeling for scalable 3D object detection. *ArXiv*, abs/2103.02093, 2021. 2, 3
- [9] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8748–8757, 2019. 1
- [10] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis* and machine intelligence, 40(4):834–848, 2017. 6
- [11] Yukang Chen, Yanwei Li, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Focal sparse convolutional networks for 3d object detection.

In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 5428–5437, 2022. 1, 3, 7

- [12] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1201–1209, 2021. 3, 4, 6, 7
- [13] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition* (*CVPR*), 2012. 1, 2, 6, 7
- [14] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. arXiv preprint arXiv:2007.01434, 2020. 2
- [15] Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11873–11882, 2020. 3
- [16] 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*, pages 770–778, 2016. 6
- [17] Deepti Hegde and Vishal M Patel. Attentive prototypes for source-free unsupervised domain adaptive 3d object detection. arXiv preprint arXiv:2111.15656, 2021. 2
- [18] Deepti Hegde, Velat Kilic, Vishwanath Sindagi, A Brinton Cooper, Mark Foster, and Vishal M Patel. Source-free unsupervised domain adaptation for 3d object detection in adverse weather. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pages 6973–6980. IEEE, 2023. 2
- [19] John Houston, Guido Zuidhof, Luca Bergamini, Yawei Ye, Long Chen, Ashesh Jain, Sammy Omari, Vladimir Iglovikov, and Peter Ondruska. One thousand and one hours: Self-driving motion prediction dataset. arXiv preprint arXiv:2006.14480, 2020. 1, 6, 7
- [20] Peixiang Huang, Li Liu, Renrui Zhang, Song Zhang, Xinli Xu, Baichao Wang, and Guoyi Liu. Tig-bev: Multi-view bev 3d object detection via target inner-geometry learning. arXiv preprint arXiv:2212.13979, 2022. 3
- [21] Philip Jacobson, Yiyang Zhou, Wei Zhan, Masayoshi Tomizuka, and Ming C Wu. Center feature fusion: Selective multi-sensor fusion of center-based objects. In 2023 IEEE International Conference on Robotics and Automation (ICRA), pages 8312–8318. IEEE, 2023. 1
- [22] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33:18661–18673, 2020. 5
- [23] Hyeonseong Kim, Yoonsu Kang, Changgyoon Oh, and Kuk-Jin Yoon. Single domain generalization for lidar semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer* vision and pattern recognition, pages 17587–17598, 2023. 3

- [24] Jaeyeul Kim, Jungwan Woo, Jeonghoon Kim, and Sunghoon Im. Rethinking lidar domain generalization: Single source as multiple density domains. In *European Conference on Computer Vision*, pages 310–327. Springer, 2024. 3
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6
- [26] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. 1, 3, 6
- [27] Alexander Lehner, Stefano Gasperini, Alvaro Marcos-Ramiro, Michael Schmidt, Mohammad-Ali Nikouei Mahani, Nassir Navab, Benjamin Busam, and Federico Tombari. 3D-VField: Adversarial augmentation of point clouds for domain generalization in 3D object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17295–17304, 2022. 3, 6, 7, 8
- [28] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc V Le, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17182–17191, 2022. 1, 2, 3, 7
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part* V 13, pages 740–755. Springer, 2014. 6
- [30] Meng Liu and Jianwei Niu. Bev-net: A bird's eye view object detection network for lidar point cloud. In 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 5973–5980. IEEE, 2021. 3
- [31] Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In *IEEE International Conference on Robotics and Automation* (ICRA), 2023. 1, 3
- [32] Zhipeng Luo, Zhongang Cai, Changqing Zhou, Gong-Duo Zhang, Haiyu Zhao, Shuai Yi, Shijian Lu, Hongsheng Li, Shanghang Zhang, and Ziwei Liu. Unsupervised domain adaptive 3D detection with multi-level consistency. *ArXiv*, abs/2107.11355, 2021. 2, 3
- [33] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In 2019 IEEE/RSJ international conference on intelligent robots and systems (IROS), pages 4213–4220. IEEE, 2019. 1
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya

Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 2021. 5

- [35] Cristiano Saltori, St'ephane Lathuili'ere, N. Sebe, E. Ricci, and Fabio Galasso. SF-UDA^{3D}: Source-free unsupervised domain adaptation for LiDAR-based 3D object detection. 2020 International Conference on 3D Vision (3DV), pages 771–780, 2020. 2
- [36] Jules Sanchez, Jean-Emmanuel Deschaud, and François Goulette. Domain generalization of 3d semantic segmentation in autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18077–18087, 2023. 3
- [37] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN:
 3D object proposal generation and detection from point cloud. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 770–779, 2019.
- [38] 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 *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10529–10538, 2020. 1, 6
- [39] 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. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43 (08):2647–2664, 2021. 1, 3, 4, 8
- [40] 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 *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2446–2454, 2020. 1, 2, 6, 7
- [41] OpenPCDet Development Team. OpenPCDet: An open-source toolbox for 3D object detection from point clouds. https: //github.com/open-mmlab/OpenPCDet, 2020. 6
- [42] Darren Tsai, Julie Stephany Berrio, Mao Shan, Eduardo Nebot, and Stewart Worrall. Ms3d: Leveraging multiple detectors for unsupervised domain adaptation in 3d object detection. arXiv preprint arXiv:2304.02431, 2023. 3
- [43] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 4604–4612, 2020. 3
- [44] Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. Pointaugmenting: Cross-modal augmentation for 3d object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11794–11803, 2021. 3
- [45] Yan Wang, X. Chen, Yurong You, Li Erran, Bharath Hariharan, M. Campbell, Kilian Q. Weinberger, and Wei-Lun Chao. Train in Germany, test in the USA: Making 3D object detectors generalize. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 11710–11720, 2020. 2, 3

- [46] Qiangeng Xu, Yin Zhou, Weiyue Wang, Charles R. Qi, and Dragomir Anguelov. Grid-GCN for fast and scalable point cloud learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [47] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely embedded convolutional detection. *Sensors*, 18(10), 2018. 3, 4
- [48] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 1
- [49] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. ST3D: Self-training for unsupervised domain adaptation on 3D object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 2, 3
- [50] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. St3d++: Denoised self-training for unsupervised domain adaptation on 3d object detection. *IEEE transactions on pattern analysis and machine intelligence*, 45(5):6354–6371, 2022. 2, 3
- [51] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3DSSD: Point-based 3D single stage object detector. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11040–11048, 2020. 3
- [52] Xingxu Yao, Sicheng Zhao, Pengfei Xu, and Jufeng Yang. Multi-source domain adaptation for object detection. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 3273–3282, 2021. 3
- [53] Xufeng Yao, Yang Bai, Xinyun Zhang, Yuechen Zhang, Qi Sun, Ran Chen, Ruiyu Li, and Bei Yu. PCL: Proxy-based contrastive learning for domain generalization. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 7097–7107, 2022. 2, 3
- [54] Yihan Zeng, Da Zhang, Chunwei Wang, Zhenwei Miao, Ting Liu, Xin Zhan, Dayang Hao, and Chao Ma. Lift: Learning 4d lidar image fusion transformer for 3d object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 17172–17181, 2022. 3
- [55] Bo Zhang, Jiakang Yuan, Botian Shi, Tao Chen, Yikang Li, and Yu Qiao. Uni3d: A unified baseline for multi-dataset 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9253–9262, 2023. 3
- [56] Wu Zheng, Weiliang Tang, Li Jiang, and Chi-Wing Fu. Se-ssd: Self-ensembling single-stage object detector from point cloud. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14494–14503, 2021. 3
- [57] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 2
- [58] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9939–9948, 2021. 1