

Long-Tailed 3D Detection via Multi-Modal Late-Fusion

Yechi Ma*, Neehar Peri*, Shuoquan Wei, Achal Dave, Wei Hua, Yanan Li, Deva Ramanan, Shu Kong

Abstract—Contemporary autonomous vehicle (AV) benchmarks have advanced techniques for training 3D detectors, particularly on large-scale multi-modal (LiDAR + RGB) data. Surprisingly, although semantic class labels naturally follow a long-tailed distribution, existing benchmarks only focus on a few common classes (e.g., `pedestrian` and `car`) and neglect many rare but crucial classes (e.g., `emergency vehicle` and `stroller`). However, AVs must reliably detect both common and rare classes for safe operation in the open world. We address this challenge by formally studying the problem of *Long-Tailed 3D Detection* (LT3D), which evaluates *all* annotated classes, including those in-the-tail. We address LT3D with hierarchical losses that promote feature sharing across classes, and introduce diagnostic metrics that award partial credit to “reasonable” mistakes with respect to the semantic hierarchy (e.g., mistaking a `child` for an `adult`). Further, we point out that rare-class accuracy is particularly improved via *multi-modal late fusion* (MMLF) of independently trained *uni-modal* LiDAR and RGB detectors. Importantly, such an MMLF framework allows us to leverage large-scale uni-modal datasets (with more examples for rare classes) to train better uni-modal detectors, unlike prevailing end-to-end trained multi-modal detectors that require paired multi-modal data. Finally, we examine three critical components of our simple MMLF approach from first principles and investigate whether to train 2D or 3D RGB detectors for fusion, whether to match RGB and LiDAR detections in 3D or the projected 2D image plane, and how to fuse matched detections. Extensive experiments reveal that 2D RGB detectors achieve better recognition accuracy for rare classes than 3D RGB detectors and matching on the 2D image plane mitigates depth estimation errors for better matching. Our proposed MMLF significantly improves LT3D performance over prior work, particularly improving rare class performance from 12.8 to 20.0 mAP! Our code and models are available on our project page.

Index Terms—Long-Tailed Distribution, 3D Detection, Multi-Modal Late-Fusion, Autonomous Vehicles, Open World, LiDAR, RGB.

1 INTRODUCTION

Three-dimensional (3D) object detection is a key component in many robotics systems such as autonomous vehicles (AVs). To facilitate research in this space, the AV industry has released many large-scale 3D annotated multi-modal datasets [2], [3], [4], [5], [6]. However, these datasets often benchmark on only a few common classes (e.g., `car` and `pedestrian`) and ignore rare classes (although annotated by some datasets) like `stroller` and `emergency vehicle` (Fig. 1). In the real open world, safe navigation [7], [8] requires AVs to reliably detect rare-class objects such as `debris` and `stroller`. This motivates the study of *Long-Tailed 3D Detection* (LT3D), a problem that requires detecting objects from both common and rare classes.

Status Quo. Among contemporary AV datasets, nuScenes [6] exhaustively annotates objects of various classes crucial to AVs (Fig. 1) and organizes them into a semantic hierarchy (Fig. 4). We are motivated to study LT3D by re-purposing *all* annotated classes in nuScenes because detecting rare classes is useful for downstream tasks such as motion planning. Importantly, LT3D is not simply solved by retraining state-of-the-art methods on both common and rare classes [9]. For example, CMT [10], a multi-modal transformer-based detector, achieves only 4.8 mAP on rare categories despite achieving 79.9 mAP on common classes (Table 1).

Protocol. LT3D requires 3D localization and recognition of objects from each of the common (e.g., `adult` and `car`) and

rare classes (e.g., `child` and `stroller`). Moreover, for safety-critical robots such as AVs, detecting but mis-classifying rare objects (e.g., mis-classifying a `child` as an `adult`) is preferable to failing to detect them at all for safe operation. Therefore, we propose a new diagnostic metric to quantify the severity of classification mistakes in LT3D that exploits inter-class relationships w.r.t to the semantic hierarchy when awarding partial credit (Fig. 4). We use both the standard and proposed metrics to evaluate 3D detectors on all classes. Further, since prior work focuses on only a few common classes, they miss opportunities to exploit this semantic hierarchy during training. We propose hierarchical losses to promote feature sharing across both common and rare classes

Technical Insights. To address LT3D, we start by retraining state-of-the-art 3D detectors on *all* classes. Somewhat surprisingly, prior work performs rather poorly on rare classes, e.g., CenterPoint [12] achieves only 0.1 AP on `child` and 0.1 AP on `stroller` (Table 5). We propose several algorithmic innovations to improve these results. First, we allow feature sharing across common and rare classes by training a single feature trunk, adding in hierarchical losses that ensure features will be useful for all classes (Table 4). Second, noting that LiDAR data is simply too impoverished for even humans to recognize certain tail objects that tend to be small (e.g., `strollers`), we propose a *Multi-Modal Late-Fusion* (MMLF) framework (Fig. 2) that fuses detections from a LiDAR-only detector (for precise 3D localization) and an RGB-only detector (for better recognition). We introduce Multi-Modal Filtering (MMF), a simple late-fusion approach that post-processes uni-modal 3D predictions from 3D LiDAR-only (e.g., CenterPoint [12]) and 3D RGB-only (e.g., FCOS3D [13]) detectors, filtering away detections that are inconsistent across modalities. This significantly improves LT3D performance by 3.2

First two authors contributed equally. Neehar Peri and Deva Ramanan are with Carnegie Mellon University. Achal Dave is with Anthropic. Yechi Ma is with Zhejiang University and Zhejiang Lab. Shuoquan Wei, Wei Hua and Yanan Li are with Zhejiang Lab. Shu Kong is with the University of Macau, and Institute of Collaborative Innovation.
Corresponding Email: skong@um.edu.mo

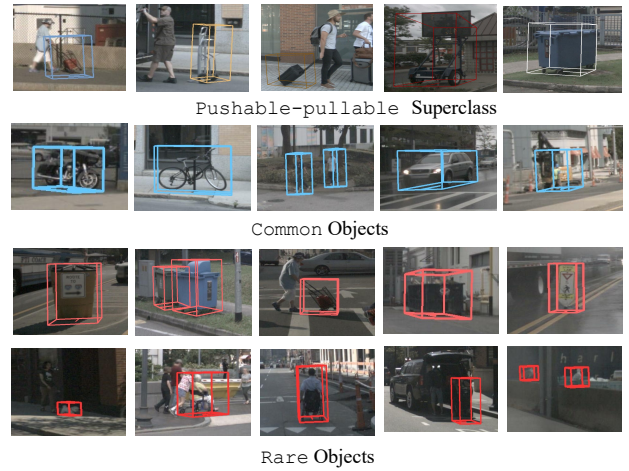
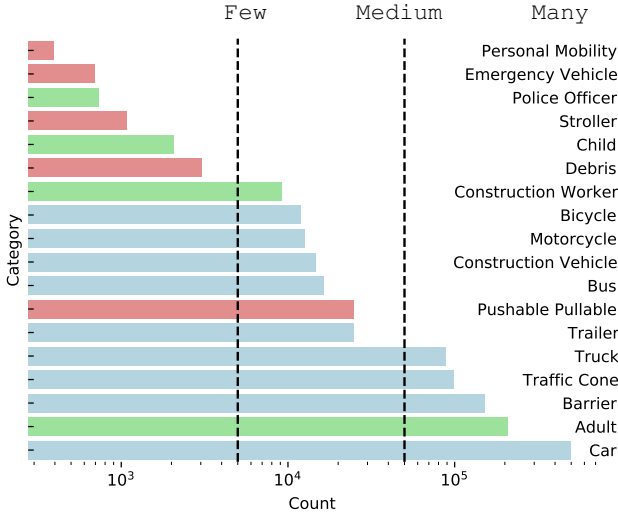


Fig. 1: According to the histogram of per-class object counts (on the left), the nuScenes benchmark focuses on the common classes in cyan (e.g., car and barrier) but ignores rare ones in red (e.g., stroller and debris). In fact, the benchmark creates a superclass pedestrian by grouping multiple classes in green, including the common class adult and several rare classes (e.g., child and police-officer); this complicates the analysis of detection performance as pedestrian performance is dominated by adult. Moreover, the ignored superclass pushable-pullable also contains diverse objects such as shopping-cart, dolly, luggage and trash-can as shown in the top row (on the right). We argue that AVs should also detect rare classes as they can affect AV behaviors. Following [1], we report performance for three groups of classes based on their cardinality (split by dotted lines): Many, Medium, and Few.

Late-fusion requires *matching* and *fusing* uni-modal detections.

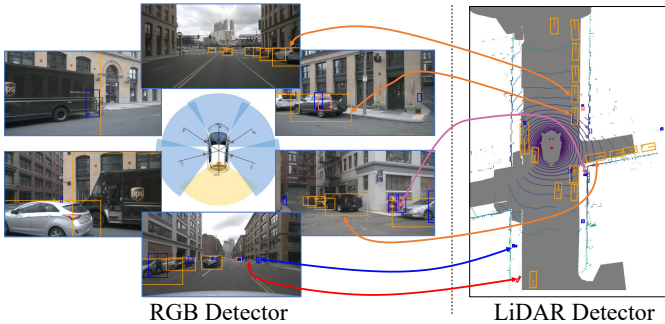


Fig. 2: We extensively explore the simple *multi-modal late-fusion* (MMLF) framework for LT3D by ensembling RGB and LiDAR uni-modal detectors [9]. We rigorously examine three critical components within this framework (Fig. 3) and propose a simple method that fuses detections produced by a 2D RGB-detector (e.g., DINO [11]) and a 3D LiDAR-detector (e.g., CenterPoint [12]). Our method achieves 51.4 mAP on the nuScenes [6] LT3D benchmark, significantly improving over state-of-the-art detectors by 5.9% (Table 1).

mAP on average (Table 2). Next, we delve into this MMLF framework and study three critical design choices (Fig. 3): (A) whether to train a 2D or a 3D monocular RGB detector for late-fusion, (B) whether to match detections in the 2D image plane or in 3D, and (C) how to optimally fuse detections. Our exploration reveals that using 2D RGB detectors, matching on the 2D image plane, and combining score-calibrated predictions with Bayesian fusion yields state-of-the-art LT3D performance, significantly outperforming end-to-end trained multi-modal 3D detectors (Table 1).

Contributions. We present three major contributions. First, we formulate the problem of LT3D, emphasizing the detection of both common and rare classes in safety-critical applications like AVs. Second, we design LT3D’s benchmarking protocol and develop a diagnostic metric that awards partial credit depending on the severity of misclassifications (e.g., misclassifying child-vs-adult is less problematic than misclassifying child-vs-car).

Third, we propose several architecture-agnostic approaches to address LT3D, including a simple *multi-modal late-fusion* (MMLF) strategy that generalizes across different RGB and LiDAR architectures. We conduct extensive experiments to ablate our design choices and demonstrate that our simple MMLF approach achieves state-of-the-art results on the nuScenes and Argoverse 2 LT3D benchmarks.

2 RELATED WORK

3D Object Detection for AVs. Contemporary approaches to 3D object detection can be broadly classified as LiDAR-only, RGB-only, and sensor-fusion methods. Recent work in 3D detection is heavily inspired by prior work in 2D detection [14], [15], [16]. LiDAR-based detectors like PointPillars [17], CBGS [18], and PVRCNN++ [19] adopt an SSD-like architecture [15] that regresses amodal bounding boxes from a bird’s-eye-view (BEV) feature map. More recently, CenterPoint [12] adopts a center-regression loss that is inspired by CenterNet [14]. Despite significant progress, LiDAR-based detectors often produce many false positives because it is difficult to distinguish foreground objects from the background given sparse LiDAR returns. Monocular RGB-based methods have gained popularity in recent years due to increased interest in camera-only perception. FCOS3D [13] extends FCOS [20] by additionally regressing the size, depth, and rotation for each object. More recently, methods such as BEVDet and BEVFormer [21], [22], [23] construct a BEV feature-map by estimating the per-pixel depth of each image feature [24]. PolarFormer [25] introduces a polar-coordinate transformation that improves near-field detection. Importantly, many of these state-of-the-art 3D RGB detectors are commonly pre-trained on large external datasets like DDAD [26]. Monocular RGB detectors accurately classify objects but struggle to estimate depth, particularly for far-field detections [27]. Despite recent advances in LiDAR and RGB 3D detectors, we find that multi-modal fusion is essential for LT3D (detailed next). Importantly, using both RGB (for better recognition) and LiDAR (for better 3D localization) helps detect rare classes. We study the multi-modal late-fusion




Fig. 3: We examine three key components in the multi-modal late-fusion (MMLF) of uni-modal RGB and LiDAR detectors from first principles: **A.** whether to train 2D or 3D RGB detectors, **B.** whether to match uni-modal detections on the 2D image plane or in the 3D bird’s-eye-view (BEV), and **C.** how to best fuse matched detections. Our exploration reveals that using 2D RGB detectors, matching on the 2D image plane, and combining calibrated scores with Bayesian fusion yields state-of-the-art LT3D performance (Table 3 and Table 4).

framework described in Fig. 2 to determine how to effectively fuse RGB and LiDAR uni-modal detectors for LT3D.

Multi-modal 3D Detection. Conventional wisdom suggests that fusing multi-modal cues, particularly using LiDAR and RGB, can improve 3D detection. Intuitively, LiDAR faithfully measures the 3D world (although it has notoriously sparse point returns), and RGB has high-resolution that captures semantic features for recognition (but lacks 3D information). Multi-modal fusion for 3D detection is an active field of exploration. Popular approaches can be categorized as input-fusion, feature-fusion, and late-fusion. Input-fusion methods typically augment LiDAR points using image-level features. For example, PointPainting [28] projects LiDAR points onto the output mask of a semantic segmentation model and appends corresponding class scores to each point. MVP [29] densifies regions of LiDAR sweeps that correspond with objects in semantic segmentation masks. In contrast, Frustum PointNets [30] leverage 2D RGB detections to localize objects within the box frustum using PointNets [31].

Recent works show that feature-fusion can be more effective than input-fusion. PointFusion [32] fuses global image and point-cloud features prior to detection and MSMD Fusion [33] fuses LiDAR and RGB features at multiple scales. TransFusion [34] and BEVFusion [35] fuse features in the BEV space using multi-headed attention. Despite the success of transformers for detecting common objects, [9] finds that TransFusion struggles to detect rare classes, and posits that the transformer architecture, as adopted in TransFusion and BEVFusion, suffers from limited training data (particularly for classes in the long tail). For transformers to work well in practice, they should be trained on diverse, large-scale datasets [36], [37]. It is also worth noting that end-to-end trained multi-modal detectors require paired multi-modal data for training, increasing the cost of data collection and modality alignment. In this work, we opt to study a multi-modal late-fusion strategy, a framework that ensembles uni-modal detectors, which do not require aligned RGB-LiDAR paired training data. CLOCs [38] is a late-fusion method that learns a separate network to fuse RGB and LiDAR detections, showing promising results for 3D detection. Importantly, prior late-fusion methods like CLOCS [38] only perform late-fusion on matched predictions with semantic agreement and do not fix misclassifications. We find that handling such misclassifications is critical for improving rare class performance. We delve into this simple multi-modal late-fusion framework, study three crucial design choices, and present a method that significantly outperforms the state-of-the-art for LT3D.

Long-Tailed Perception. AV datasets follow a long-tailed class distribution: a few classes like `car` and `pedestrian` are dominant, while others like `stroller` and `debris` are rarely seen. However, this problem is not unique to the AV domain [39] – Long-Tailed Perception (LTP) is a long-standing problem in the literature [1] and has been widely studied through the lens

of image classification, aiming for high accuracy averaged across imbalanced classes [1], [40], [41]. Existing LTP methods propose reweighting losses [42], [43], [44], [45], [46], [47], rebalancing data sampling [48], [49], [50], balancing gradients computed from imbalanced classes [51], and balancing network weights [41]. Others study LTP through the lens of 2D object detection with RGB images [52]. Compared to 2D image-based recognition, long-tailed 3D detection has unique opportunities and challenges because sensors such as LiDAR directly provide geometric and ego-motion cues that are difficult to extract from 2D images. Further, 2D detectors must detect objects of different scales due to perspective image projection, dramatically increasing the complexity of the output space (e.g., requiring more anchor boxes). In contrast, 3D objects do not exhibit as much scale variation, but far-away objects tend to have sparse LiDAR returns [27], [53], imposing different challenges. Finally, 3D detectors often use class-aware heads (i.e. each class has its own binary classifier) while 2D long-tail recognition approaches typically use shared softmax heads.

Recently, CBGS [18] explicitly addresses rare-class 3D detection by up-sampling LiDAR-sweeps with instances of rare classes, and pasting instances of rare objects copied from different scenes. Although this works well for improving detection of infrequently-seen classes (e.g. classes with medium number of examples like `bicycle` and `construction vehicle`), it does not provide significant improvement for classes with only a few examples like `debris` and `stroller` [9]. Additionally, rare classes, such as `child` and `stroller`, are typically small in size and have a limited number of LiDAR returns. As a result, LiDAR-only detectors struggle to accurately recognize these rare classes. In this work, we address LT3D by multi-modal late-fusion, which ensembles RGB and LiDAR uni-modal detectors.

3 LT3D BENCHMARKING PROTOCOL

Conceptually, LT3D extends the traditional 3D detection problem, which focuses on identifying objects from K common classes, by further requiring detection of N rare classes. As LT3D emphasizes detection performance on *all* classes, we report metrics for three groups of classes based on their cardinality (Fig. 1-left): *many* ($>50k$ instance/class), *medium* ($5k \sim 50k$ instance/class), and *few* ($<5k$ instance/class). To better analyze LT3D performance, we present two metrics below.

Mean average precision (mAP) is a well-established metric for object detection [2], [54], [55]. For 3D detection on LiDAR sweeps, a true positive is defined as a detection that has a center distance within a distance threshold to a ground-truth annotation [6]. mAP computes the mean of AP over classes, where per-class AP is the averaged area under the precision-recall curves with distance thresholds of [0.5, 1, 2, 4] meters.

Hierarchical Mean Average Precision (mAP_H). For safety-critical applications, it is more desirable to localize-but-misclassify an object than miss-detect this object, e.g., detecting but misclassifying a `child` as `adult` is better than not detecting this `child`. Therefore, we present hierarchical mean average precision (mAP_H), which considers such semantic relationships across classes to award partial credit. We use AP_H as a supplementary diagnostic tool to analyze how LT3D detectors make mistakes. To encode inter-class relationships, we leverage the semantic hierarchy (Fig. 4) such as the one defined by nuScenes [6]. We derive partial credit as a function of semantic similarity using the least common ancestor (LCA) distance metric. LCA has been proposed for image classification [56], [57] but not for object detection. Extending this metric to object detection is challenging as we must jointly evaluate semantic and spatial overlap. For clarity, we describe how to compute AP_H for a class C .

LCA=0: Consider the predictions and ground-truth boxes for C . Label the set of predictions that overlap with ground-truth boxes for C as true positives. Other predictions are false positives. *This is identical to the standard mAP metric.*

LCA=1: Consider the predictions for C , and ground-truth boxes for C and its all sibling classes with LCA distances to C of 1. Label the predictions that overlap a ground-truth box of C as a true positive. Label the predictions that overlap sibling classes as ignored [55]. All other predictions for C are false positives.

LCA=2: Consider the predictions for C and ground-truth boxes for C and all its sibling classes with LCA distances to C no greater than 2. For nuScenes, this includes all classes. Label the set of predictions that overlap ground-truth boxes for C as true positives. Label the set of predictions that overlap other classes as ignored. All other predictions for C are false positives.

4 LT3D METHODS: THE DEVIL IS IN THE DETAILS

To address LT3D, we first retrain state-of-the-art (SOTA) detectors on *all* long-tailed classes. We consider the following detectors:

- *LiDAR-based 3D Detectors:* CenterPoint [12], PointPillars [17], TransFusion-L [34], BEVFusion-L [35], CMT-L [10];
- *RGB-based 3D Detectors:* FCOS3D [13], PolarFormer [25], BEVFormer [23];
- *RGB-based 2D Detectors:* YOLOV7 [58], and DINO [11];
- *End-to-End Multi-modal 3D Detectors:* TransFusion [34], BEVFusion [35], CMT [10], and DeepInteraction [59].

In Section 4.1 and 4.2, we introduce two modifications that consistently improve LT3D performance, and describe our multi-modal late-fusion framework.

4.1 Grouping-Free Detector Head

Extending existing 3D detectors to train with more classes is surprisingly challenging. Many contemporary networks use a multi-head architecture that groups classes of similar size and shape to facilitate efficient feature sharing [12], [18]. For example, CenterPoint [12] groups `pedestrian` and `traffic-cone` since these objects are both tall and skinny. However, this multi-head grouping design may not work for super-classes like `pushable-pullable` and `debris` that contain diverse objects of different sizes and shapes. Moreover, in the multi-head design, each head is a group-specific detector that consists of several layers with lots of parameters. This makes multi-head grouping difficult to scale for a large number of classes. To

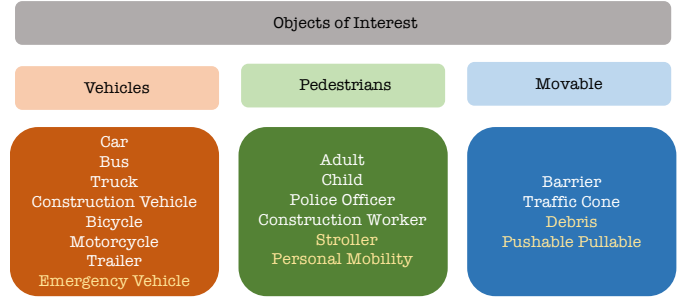


Fig. 4: The nuScenes dataset defines a semantic hierarchy for all annotated classes (Fig. 1). We highlight common classes in white and rare classes in gold. The standard nuScenes benchmark makes two choices for dealing with rare classes: (1) ignore them (e.g., `stroller` and `pushable-pullable`), or (2) group them into coarse-grained classes (e.g., `adult`, `child`, `construction-worker`, `police-officer` are grouped as `pedestrian`). Since the `pedestrian` class is dominated by `adult` (Fig. 1), the standard benchmarking protocol masks the challenge of detecting rare classes like `child` and `police-officer`. We explicitly incorporate this semantic hierarchy to evaluate LT3D performance via the proposed mAP_H metric described in Sec. 3.

address these issues, we treat each class as its own group to avoid hand-crafted grouping heuristics and design a group-free strategy, in which each class has only one linear layer as its detector (Fig. 7) and all classes share a single detector head. This design significantly reduces the number of parameters and allows learning the shared feature backbone collaboratively with all classes, effectively mitigating overfitting to rare classes. Adding a new class is as simple as adding a single linear layer to the detector head. Importantly, we demonstrate that our grouping-free detector head outperforms conventional grouping-based methods in the supplement.

4.2 Training with a Semantic Hierarchy

The nuScenes dataset defines a semantic hierarchy (Fig. 4) for all classes, grouping semantically similar classes under coarse-grained categories [6]. We leverage this hierarchy during training. Specifically, we train detectors to predict three labels for each object: its fine-grained class (e.g., `child`), its coarse class (e.g., `pedestrian`), and the root class `object`. We adopt a grouping-free detector head (Fig. 7) that outputs separate “multi-task” heatmaps for each class, and use a *per-class sigmoid focal loss* [60] rather than multi-class cross-entropy loss to avoid normalizing class probabilities across fine-grained and super-classes. It is worth noting that our approach does not explicitly enforce a tree-like hierarchy, and can be applied to more complex label relationships [61]. Crucially, adding a `vehicle` heatmap does not directly interfere with the `car` heatmap. However, this might produce repeated detections on the same test object (e.g. a single ground-truth `car` may be detected as a `car`, `vehicle`, and `object`). We address that by simply ignoring coarse detections at test time. We explore alternatives in the supplement and conclude that they achieve similar LT3D performance. Perhaps surprisingly, this training method improves detection performance not only for rare classes, but also for common classes (Table 4).

4.3 Multimodal Filtering for Detection Fusion

Rare class instances are often small and can be challenging to recognize from sparse (LiDAR) geometry alone – even humans




Fig. 5: Multi-Modal Filtering (MMF) effectively removes high-scoring false-positive LiDAR detections. The green boxes are ground-truth strollers, while the blue boxes are stroller detections from the LiDAR-based detector CenterPoint [12] (**left**) and RGB-based detector FCOS3D [13] (**mid**). The final filtered result removes LiDAR detections not within m meters of any RGB detection, shown in the red region, and keeps all other LiDAR detections, shown in the white region (**right**).

struggle to find strollers in LiDAR point clouds. This suggests that one can leverage multi-modal cues to improve LT3D. First, we find that, although LiDAR-based detectors are widely adopted for 3D detection, they produce many high-scoring false positives (FPs) for rare classes due to misclassification. We focus on removing such FPs. To this end, we use an RGB-based 3D detector (e.g., FCOS3D [13]) to filter out high-scoring false-positive LiDAR detections by leveraging two insights: (1) LiDAR-based 3D-detectors can achieve high recall and precise 3D localizations for TPs, and (2) RGB-based 3D-detections are accurate w.r.t recognition although their 3D localization is poor. Fig. 5 demonstrates our Multi-Modal Filtering (MMF) strategy. For each RGB detection, we keep LiDAR detections within a radius of m meters and remove all the other LiDAR detections. We denote this method as $\text{MMF}(D_L, D_R)$, where D_L and D_R are any 3D LiDAR detector and 3D RGB detector, respectively. Table 2 demonstrates that MMF greatly improves LT3D.

Limitations of Multi-Modal Filtering. Despite the effectiveness of multi-modal filtering, this approach is sensitive to the classification accuracy of FCOS3D when matching LiDAR and RGB-based detections, leading to many correctly localized but misclassified detections. We use a confusion matrix to further analyze the misclassifications within each superclass, as shown in Fig. 6. We explain how to compute a confusion matrix for the detection task: For each superclass, we construct a confusion matrix, in which the entry (i, j) indicates the misclassification




Fig. 6: We analyze our multi-modal filtering (MMF) approach. Rare classes are most often confused by the dominant class (in blue) in each superclass: (**left**) Vehicle is dominated by car, (**mid**) Pedestrian is dominated by adult, and (**right**) Movable is dominated by barrier. We find that class confusions are reasonable. Car is often mistaken for truck. Similarly, truck, construction-vehicle and emergency-vehicle are most often mistaken for car. Bicycle and motorcycle are sometimes misclassified as car, presumably because they are sometimes spatially close to cars. Adults have similar appearance to police-officer and construction-worker, and they are often co-localized with child and stroller.

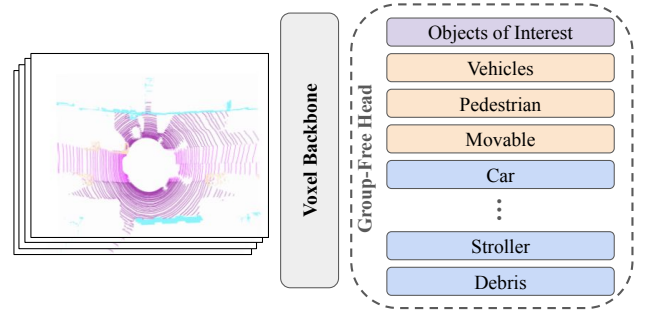


Fig. 7: We leverage the semantic hierarchy defined in the nuScenes dataset (Fig. 4) to train LT3D detectors by predicting class labels at *multiple* levels of the hierarchy for an object: its fine-grained label (e.g., child), its coarse class (e.g., pedestrian), and the root-level class object. This means that the final vocabulary of classes is no longer mutually exclusive, so we use a sigmoid focal loss that learns separate spatial heatmaps for each class.

rate of class- i objects as class- j . Specifically, given a class i , we find its predictions that match ground-truth boxes within a 2m center-distance of class- i and all its sibling classes (LCA=1, within the corresponding superclass); we ignore all unmatched detections. This allows us to count the mis-classifications of class- i objects into class- j . We normalize these counts to produce misclassification rates. We find that rare classes are most often confused by the dominant class in each superclass: Vehicle is dominated by car, Pedestrian is dominated by adult, and Movable is dominated by barrier. We posit that this can be addressed with more accurate RGB-based classifiers, motivating our study of *multi-modal late-fusion* (MMLF) from first principles.

4.4 Delving into Multi-Modal Late-Fusion for LT3D

Motivated by the effectiveness of our frustratingly simple multi-modal filtering (MMF) approach, we delve into this multi-modal late-fusion (MMLF) strategy. We explore three questions from first principles (Fig. 3): (1) how do we effectively incorporate RGB information, (2) how do we match RGB and LiDAR detections, and (3) how do we fuse them. We denote our final approach as $\text{MMLF}(D_L, D_R)$ in Table 1, where D_L is any 3D LiDAR detector and D_R is any 2D RGB detector.

4.4.1 How Do We Incorporate RGB Information?

Although LiDAR offers accurate localization, LiDAR-only detectors struggle to recognize objects using sparse LiDAR alone. RGB images provide complementary information that is essential for identifying objects and disambiguating those that are geometrically similar in point clouds but semantically different in images, e.g., construction worker vs. police officer. In the last subsection, we proposed multimodal filtering (MMF) (Fig. 5), a strategy that ensembles a 3D RGB detector and a 3D LiDAR detector, yielding remarkable improvements for LT3D (Table 2). However, we find that *ensembling a 2D RGB detector and a 3D LiDAR detector* yields more significantly better LT3D performance. We present insights on why using a 2D detector to incorporate RGB information is better and ablate the impact of using 2D vs. 3D RGB detectors for late-fusion in Table 3.

2D RGB detectors are more mature. 2D object detection is a fundamental problem in computer vision [55], [62], [63] that has matured in recent years and model trade-offs are well understood [15], [60], [63], [64]. In this work, we consider two

state-of-the-art 2D RGB detectors, YOLOV7 [58] and DINO [11]. YOLOV7 is a real-time detector that identifies a number of training techniques that nearly doubles the inference efficiency over prior work without sacrificing performance. DINO [11] is a recent transformer-based detector that improves upon DETR [16] using denoising anchor boxes. As 2D detectors do not make 3D predictions (e.g., depth and rotation), understanding how to best leverage them in the context of LT3D is a key challenge.

2D RGB detectors can be trained with more diverse data. Training 2D RGB detectors only requires *2D bounding box* annotations, which are significantly cheaper to collect than 3D cuboids used for training 3D RGB detectors [13], [25]. Since annotating 3D amodal cuboids is both expensive and non-trivial (compared to bounding-box annotations for 2D detection), datasets for monocular 3D RGB detection are considerably smaller and less diverse than their 2D detection counterparts. For example, nuScenes [6] (published in 2020) annotates 144K RGB images of 23 classes using 3D cuboids, while COCO [55] (an early 2D detection dataset published in 2014) annotates 330K images of 80 classes using 2D bounding boxes. This allows us to pretrain 2D RGB detectors on significantly larger, more diverse, publicly available datasets [65], [66], [67], [68], [69]. We demonstrate in Fig. 11 that leveraging existing 2D detection datasets helps train stronger 2D detectors, further improving LT3D performance.

4.4.2 How Do We Match Uni-Modal Detections?

Finding correspondence between two sets of RGB and LiDAR uni-modal detections is an essential step in the late-fusion framework (Fig. 3B). Our previous multi-modal filtering method [9] matched 3D RGB and 3D LiDAR detections using center distance (Fig. 5). However, precisely matching detections in 3D is difficult due to depth estimation errors from 3D RGB detectors. Instead, we opt to match *2D RGB* and 3D LiDAR detections. Prior work attempts to inflate 2D detections to 3D (using LiDAR points [70]), but we find that matching inflated 2D detections also yields noisy 3D boxes (due to outlier points from the background) and reduces overall match quality. In contrast, we match uni-modal detections by projecting 3D LiDAR detections onto the 2D image plane, avoiding additional noise due to imprecise depth estimates. We ablate the impact of matching in 3D versus on the 2D image plane in Table 3, and present our 2D matching algorithm below.

Spatial matching on the 2D image plane. Using the available sensor extrinsics, we project 3D LiDAR detections onto the 2D image plane. We then use the IoU metric to determine overlap between (projected) LiDAR and 2D RGB detections. A 2D RGB detection and a (projected) 3D LiDAR detection are considered

a match if their IoU is greater than a threshold. Although conceptually simple, this matching method works significantly better than using center distance to match detections in 3D (cf. bottom two rows in Table 3). Spatially matching uni-modal detections using 2D IoU yields three categories of detections: matched detections, unmatched RGB detections (that do not have corresponding LiDAR detections), and unmatched LiDAR detections (that do not have corresponding RGB detections). We handle unmatched detections below and present our fusion strategy for matched detections in the next subsection.

Handling unmatched detections. We remove unmatched 2D RGB detections, positing that any unmatched RGB detections are likely to be false positives given that LiDAR detectors tend to yield high recall [9]. Further, accurately inflating 2D RGB detections to 3D remains challenging. For unmatched 3D LiDAR detections, we down-weight their detection confidence scores by w (inspired by SoftNMS [71]). We optimize for w by grid search on a validation set and set $w = 0.4$.

4.4.3 How Do We Fuse Matched Uni-Modal Detections?

As illustrated by Fig. 3C and Fig. 9, detections may match spatially but not semantically. To address this, we propose a semantic matching heuristic. Given a pair of spatially matched RGB and LiDAR detections, we consider two cases: matched detections with semantic disagreement (e.g., RGB and LiDAR predict different classes), and matched detections with semantic agreement (e.g., RGB and LiDAR predict the same class).

Addressing semantic disagreement between modalities. If the two modalities predict different semantic classes, we use the confidence score (which is calibrated as explained below) and class label of the RGB-based detection and the 3D box extent from the LiDAR-based detection. Intuitively, RGB detectors can predict semantics more reliably from high resolution images than LiDAR-only detectors. This helps correct misclassifications of geometrically similar but semantically different objects produced by the 3D LiDAR detector, as shown in Fig. 9. Importantly, prior late-fusion methods like CLOCS [38] and our MMF only perform late-fusion on matched predictions with semantic agreement and do not fix misclassifications. We find that handling such misclassifications is critical for improving rare class performance (Table 1).

If both modalities predict the same semantic class, we perform score fusion and probabilistic ensembling [72] as described below. Note that the confidence scores of RGB and LiDAR detections are not directly comparable by default: LiDAR-based detectors are often under-confident as it is difficult to distinguish foreground-vs-background using sparse LiDAR alone. Therefore, score calibration is crucial for fusion. Below, we explore score calibration of RGB (x_{RGB}) and LiDAR (x_{LiDAR}) detections.

Score calibration. We calibrate detection confidences per model by tuning a temperature τ_c for the logit score of class c on a validation set before applying a sigmoid transform [72], [73], i.e., $\text{sigmoid}(\text{logit}_c/\tau_c)$. Optimally tuning per-class τ_c is computationally expensive as it requires tuning for all classes at the same time. Instead, we choose to greedily tune each τ_c , optimizing per-class AP on a val-set for each class progressively ordered by their cardinalities. It is worth noting that this score calibration is only performed once in training and tuned τ_c and $p(c)$ do not need further optimization during inference. Importantly, score calibration does not increase runtime or complexity.

Probabilistic ensembling. We assume independent class prior $p(c)$ and conditional independence given the class label c [72],




Fig. 8: Our multi-modal late-fusion (MMLF) strategy takes 3D LiDAR and 2D RGB detections as input (Fig. 3A), matches 2D RGB and (projected) 3D LiDAR detections on the image plane (Fig. 3B), and fuses matched predictions with score calibration and probabilistic ensembling to produce 3D detections (Fig. 3C).

i.e., $p(x_{\text{RGB}}, x_{\text{LiDAR}}|c) = p(x_{\text{RGB}}|c)p(x_{\text{LiDAR}}|c)$. We compute the final score as

$$\begin{aligned} p(c|x_{\text{RGB}}, x_{\text{LiDAR}}) &= p(x_{\text{RGB}}, x_{\text{LiDAR}}|c)p(c) / p(x_{\text{RGB}}, x_{\text{LiDAR}}) \\ &\propto p(x_{\text{RGB}}, x_{\text{LiDAR}}|c)p(c) \\ &\propto p(x_{\text{RGB}}|c)p(x_{\text{LiDAR}}|c)p(c) \\ &\propto p(c|x_{\text{RGB}})p(c|x_{\text{LiDAR}}) / p(c) \end{aligned}$$

where $p(c|x_{\text{RGB}})$ and $p(c|x_{\text{LiDAR}})$ are the posteriors after calibration. Unlike the balanced class distribution studied in [72], $p(c)$ can significantly impact the final long-tailed performance. To maximize detection accuracy, the optimal tuning practice is to jointly optimize all class priors $p(c)$ altogether, which is computationally expensive. Similar to score calibration, we greedily tune $p(c)$, one by one ordered by class cardinality. Probabilistic ensembling does not increase inference time or complexity.

5 EXPERIMENTS

We conduct extensive experiments to better understand the LT3D problem and gain insights by validating our techniques described in Sec. 4. Specifically, we aim to answer the following questions:¹

- 1) Are rare classes more difficult to detect than common classes?
- 2) Are objects from rare classes sufficiently localized but misclassified?
- 3) Does training with the semantic hierarchy improve detection performance for LT3D?
- 4) Does multi-modal late-fusion help detect rare classes?

We compare our multi-modal late-fusion (MMLF) approaches (Fig. 5 and 8) with prior works and present a detailed ablation study to further address the three motivating questions in Fig. 3. Our approach improves over prior works by 5.9 mAP, notably improving by 7.2 mAP on rare classes (Table 1).

Datasets. To explore LT3D, we use the well established nuScenes [6] and Argoverse 2.0 (AV2) [4] datasets. Other AV datasets such as KITTI [2] and Waymo [5] do not support the study of LT3D as they only annotate three common classes. nuScenes and AV2 have 18 and 26 fine-grained classes, respectively, which follow long-tailed distributions. To quantify the long-tail, we compute the imbalance factor (IF), defined as the ratio between the numbers of annotations of the most and the least common classes [44]: nuScenes and AV2 have an IF=1,670 and 2,500 respectively, which are significantly more imbalanced than existing long-tail image recognition benchmarks, e.g., iNaturalist (IF=500) [74] and ImageNet-LT (IF=1,000) [75]. nuScenes arranges classes in a semantic hierarchy (Fig. 4); AV2 does not provide a semantic hierarchy but we construct one based on the nuScenes’ hierarchy (please refer to the supplement for details). Following prior work, we use official training and validation sets. We focus on nuScenes in the main paper and AV2 in the supplement. Our primary conclusions hold for both datasets.

5.1 State-of-the-Art Comparison on nuScenes

Table 1 compares our MMF and MMLF approaches with prior work on nuScenes. Fig. 9 presents qualitative results. We adapt existing methods (which were previously trained on 10 classes in nuScenes) for LT3D by retraining them on all 18 classes.

1. Answers: yes, yes, yes, yes.

TABLE 1: Benchmarking results on nuScenes. We denote the RGB and LiDAR modalities by C and L, respectively. Our MMLF fuses 3D LiDAR and 2D RGB detections (from CenterPoint and DINO, respectively) with score calibration and probabilistic ensembling. MMLF performs the best averaged across *all* common and rare classes, notably outperforming end-to-end multi-modal methods such as BEVFusion [35], DeepInteraction [59], and CMT [10]. Importantly, our MMLF approach nearly doubles the detection performance achieved by prior work on classes with few examples!

Method	Modality	All	Many	Medium	Few
FCOS3D [13]	C	20.9	39.0	23.3	2.9
BEVFormer [23]	C	27.3	52.3	31.6	1.4
PolarFormer [25]	C	28.0	54.0	31.6	2.2
CenterPoint [12]	L	40.4	77.1	45.1	4.3
TransFusion-L [34]	L	38.5	68.5	42.8	8.4
BEVFusion-L [35]	L	42.5	72.5	48.0	10.6
CMT-L [10]	L	34.7	73.4	35.9	1.1
CLOCS [38]	C+L	40.0	68.2	45.7	10.0
TransFusion [34]	C+L	39.8	73.9	41.2	9.8
BEVFusion [35]	C+L	45.5	75.5	52.0	12.8
DeepInteraction [59]	C+L	43.7	76.2	51.1	7.9
CMT [10]	C+L	44.4	79.9	53.0	4.8
CenterPoint [12] + RCNN [76]	C+L	34.0	64.8	37.5	4.3
MMF(CenterPoint, FCOS3D) (Ours)	C+L	43.6	77.1	49.0	9.4
MMLF(CenterPoint, DINO) (Ours)	C+L	51.4	77.9	59.4	20.0

CenterPoint [12], a popular LiDAR-only 3D detector, is unable to detect rare objects, achieving just 4.3 mAP on classes with few examples. This is expected as it is difficult to identify rare objects from sparse LiDAR points alone. Perhaps surprisingly, the transformer-based 3D LiDAR detector BEVFusion-L performs considerably better on few classes, achieving 10.6 mAP, but it underperforms CenterPoint by 4.6 mAP on classes with many examples. We posit that limited training data in-the-tail and class imbalance make it difficult to learn generalizable features, preventing robust LT3D performance. In contrast, BEVFusion [23], which is an end-to-end trained multi-modal method, performs 3.0 mAP better than the LiDAR-only variant (BEVFusion-L), confirming the benefit of using both RGB and LiDAR for LT3D.

Next, we implement another simple baseline, termed “CenterPoint + RCNN”, that trains a region-based CNN (RCNN) [76] classifier on cropped regions corresponding to projected 3D detections. Notably, it underperforms CenterPoint by 6.4 mAP, presumably because learning classifiers on cropped regions does not exploit contextual information and leads to worse classification accuracy. This suggests that late-fusion cannot be simply solved with a 3D LiDAR detector and a strong 2D classifier. Our multimodal filtering method MMF(CenterPoint, FCOS3D) keeps CenterPoint detections that are close to monocular 3D RGB detections produced by FCOS3D in 3D and discards all other LiDAR predictions. This simple baseline achieves 9.4 mAP on classes with few examples. Lastly, we re-implement CLOCS [38] by fusing detections from DINO and CenterPoint respectively. Note that CLOCS only fuses predictions from the same class, which prevents re-labeling misclassified LiDAR detections. However, CLOCS performs worse than our MMF baseline. By carefully considering design choices outlined in Fig. 3, MMLF(CenterPoint, DINO) improves over MMF(CenterPoint, FCOS3D) by 7.8 mAP!

5.2 Ablation Studies and Empirical Analysis

We conduct extensive ablations to validate our proposed MMLF approach. We provide additional results in the supplement.

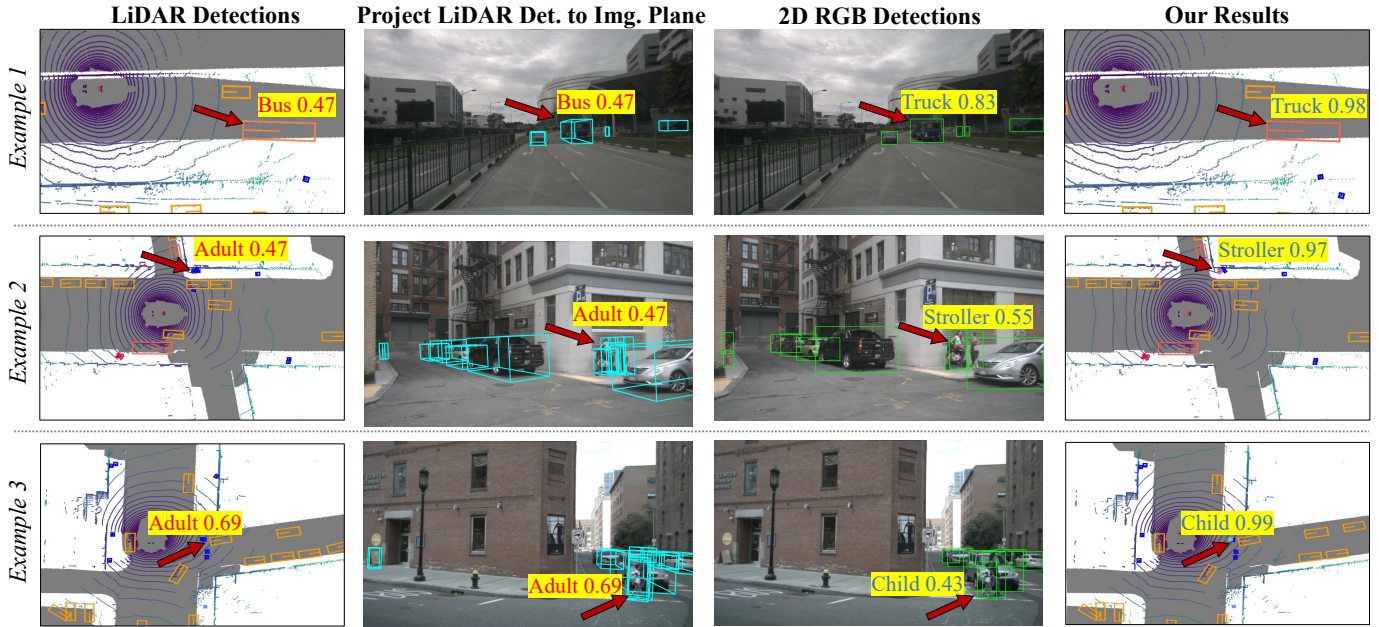


Fig. 9: Three examples demonstrate how our multi-modal late-fusion (MMLF) approach improves LT3D by ensembling 2D RGB detections (from DINO [11]) and 3D LiDAR detections (from CenterPoint [12]). In all examples, MMLF correctly relabels detections which are geometrically similar (w.r.t size and shape) in LiDAR but are visually distinct in RGB, such as bus-vs-truck, adult-vs-stroller, and adult-vs-child.

TABLE 2: **Ablation on Hierarchy Loss and Multi-Modal Filtering (MMF)**. Training with the semantic hierarchy improves both PointPillars [17] and CenterPoint [12] for LT3D by >1% mAP over All classes. Moreover, MMF yields 4~11 mAP improvement on Medium and Few classes, showing the benefit of late fusion!

Method	Modality	All	Many	Medium	Few
FCOS3D [13]	C	20.9	39.0	23.3	2.9
PointPillars [17]	L	30.0	64.2	28.4	3.4
+ Hierarchy Loss	L	31.2	66.4	30.4	2.9
+ MMF(PointPillars, FCOS3D)	C+L	35.8	66.2	41.0	4.4
CenterPoint [12]	L	39.2	76.4	43.1	3.5
+ Hierarchy Loss	L	40.4	77.1	45.1	4.3
+ MMF(CenterPoint, FCOS3D)	C+L	43.6	77.1	49.0	9.4

5.2.1 Analysis on Training with Semantic Hierarchy

We validate the benefits of training with semantic hierarchy by modifying LiDAR-based detectors (PointPillars [17] and CenterPoint [12]) to jointly predict class labels at different levels of the semantic hierarchy. For example, we modify the detector to additionally classify *stroller* as *pedestrian* and *object*. The semantic hierarchy naturally groups classes based on shared attributes and may have complementary features, helping train a detector that allows feature sharing across common and rare classes and leads to improved performance. Table 2 highlights that this approach (“+ Hierarchy Loss”) effectively improves LT3D accuracy by >1 mAP over all classes.

5.2.2 Analysis on Multi-Modal Late-Fusion

We study the trade-off between using 2D and 3D RGB detectors, and matching in the 2D image plane and in 3D in Table 3. We further examine the impact of using additional data and study different fusion strategies in Table 4.

How do we incorporate RGB information? Table 2 demonstrates that incorporating RGB information (by matching and filtering 3D LiDAR detections) using a 3D RGB detector greatly improves LT3D. We now evaluate the impact of using 2D RGB-based detectors (e.g., YOLOV7 and DINO) versus 3D RGB-

TABLE 3: **Fusing uni-modal detections in 3D vs. on the 2D image plane**. We evaluate the impact of fusing 3D LiDAR detections (from CenterPoint trained with hierarchy loss) with 2D RGB and 3D RGB detections in 3D versus on the image plane. We match and filter detections in 3D using center distance (MMF), and match and filter detections in the 2D image plane using IoU (MMLF). Following [70], we inflate 2D detections to 3D using LiDAR points within the box frustum. We project 3D detections to the image plane using provided sensor extrinsics. Results show that matching 3D RGB detections in 3D and on the image plane yields similar results. Unsurprisingly, inflating 2D RGB detections for matching in 3D performs worse than matching 3D RGB detections in 3D. In contrast, filtering LiDAR-based detections using 2D detections in the image plane (bottom right panel) significantly improves performance for classes with medium and few examples by >10 mAP. This suggests that 2D detectors achieve better detection performance compared to 3D RGB detectors. We further verify this in the supplement.

Method	Fusion in 3D (MMF)				Fusion in 2D (MMLF)			
	All	Many	Medium	Few	All	Many	Medium	Few
CenterPoint	40.4	77.1	45.1	4.3	40.4	77.1	45.1	4.3
+ FCOS3D [13]	42.9	76.6	48.7	8.1	42.6	75.0	49.4	7.7
+ BEVFormer [23]	43.2	76.9	50.8	6.3	42.8	75.2	51.4	5.7
+ PolarFormer [25]	42.8	76.8	50.0	6.1	42.6	75.1	51.1	5.6
+ YOLOV7 [58]	40.1	76.1	43.8	5.8	45.7	77.1	52.8	11.2
+ DINO [11]	40.3	76.2	44.1	5.9	49.5	77.4	57.7	16.7

based detectors (e.g., FCOS3D, BEVFormer, PolarFormer), and matching LiDAR detections with 2D RGB detections in 3D versus on the 2D image plane in Table 3. Results show that matching LiDAR detections with 2D RGB detections on the 2D image plane (bottom right) performs best.

How do we match detections from uni-modal detectors?

MMF keeps LiDAR detections within a radius of m meters for each RGB detection and removes all the others that are not close to any RGB-based detections. This works well when using a 3D RGB detector, e.g., MMF with FCOS3D, BEVFormer, and PolarFormer improves over the LiDAR-only detector CenterPoint by 2 mAP, as shown in Table 3. Moreover, matching inflated 2D RGB detections in 3D, performs worse than matching 3D RGB detections, notably achieving marginally lower accuracy than the




Fig. 10: Both our MMLF method (columns 1-3) and TransFusion [34] (columns 4 -5) share the same failure cases. In the first and second row, the 2D RGB detector DINO detects the heavily occluded cars but 3D LiDAR detector fails to detect them. As a result, our MMLF misses these cars because it throws away unmatched RGB detections which do not have accurate 3D information. In the third row, we see that, although both the LiDAR and RGB detectors fire on the object (whose ground-truth label is `police-officer`), the LiDAR detector classifies it as `adult` and RGB detector classifies it as `construction-worker`. As a result, the final detection is incorrect w.r.t the predicted label. TransFusion also misclassifies this object as an `adult`.

TABLE 4: Progressively adding training components improves LT3D performance, confirming that 2D RGB detectors are better suited for multi-modal late-fusion, matching projected 3D LiDAR detections on the 2D image plane outperforms matching 2D RGB detections inflated to 3D, and score calibration prior and probabilistic fusion improves performance. Importantly, a majority of our performance improvement can be attributed to matching and filtering on the 2D image plane with DINO. In addition, training with external data further improves rare class accuracy by 4.2 mAP.

Method	Δ	All	Many	Medium	Few
CenterPoint [12]		39.2	76.4	43.1	3.5
+ Hierarchy Loss	+1.2	40.4	77.1	45.1	4.3
+ MMF w/ DINO	+7.5	47.9	77.1	55.8	14.4
+ External Data	+1.9	49.8	77.1	57.1	18.6
+ Score Calibration	+0.7	50.5	77.8	58.2	18.7
+ Probabilistic Ensembling	+0.9	51.4	77.9	59.4	20.0

LiDAR-only baseline.

Table 3 shows that projecting LiDAR detections on the 2D image plane and fusing them with 2D RGB detections significantly improves performance for classes with medium and few examples by more than 10 mAP. In contrast, projecting 3D RGB detections for matching on the 2D image plane performs worse than matching 2D RGB detections on the 2D image plane, suggesting that MMLF performs better with 2D RGB detectors than 3D RGB detectors.

How do we fuse matched detections? Prior to fusion, we first calibrate the scores of LiDAR and RGB detections to ensure that they are comparable. This improves accuracy by 0.7 mAP averaged over all classes (Table 4). Once calibrated, we probabilistically fuse matched detections to generate the final set of detections. Table 4 shows that probabilistic fusion improves performance by 0.9 mAP.

5.2.3 Analysis of Misclassification

We evaluate CenterPoint [12] and CMT [10] using the hierarchical mAP metric (mAP_H), which considers semantic relationships across classes to award partial credit. In safety-critical applications, detecting but misclassifying objects (as a semantically related category) is more desirable than a missed detection (e.g., detecting but misclassifying a `child` as an `adult` is preferable to not detecting this `child`). Table 5 lists the results of applying

TABLE 5: **Diagnosis using the mAP_H metric on selected classes.** We analyze CenterPoint [12] and CMT [10], with our proposed group-free detector head, hierarchy loss, and MMLF (using the 2D RGB detector DINO [11]). Comparing the rows of LCA=0 for CenterPoint and CMT, we see that MMLF significantly improves on rare classes (in blue), e.g., `construction-vehicle` (CV), `bicycle`, `motorcycle` (MC), `construction-worker` (CW), `stroller`, and `pushable-pullable` (PP). Moreover, performance increases significantly from LCA=0 to LCA=1 compared against LCA=1 to LCA=2, suggesting that objects from rare classes are often detected but misclassified as sibling classes.

Method	mAP_H	Car	Adult	Truck	CV	Bicycle	MC	Child	CW	Stroller	PP
CenterPoint [12] (original)	LCA=0	82.4	81.2	49.4	19.7	33.6	48.9	0.1	14.2	0.1	21.7
	LCA=1	83.9	82.0	58.7	20.5	35.2	50.5	0.1	18.3	0.1	22.0
	LCA=2	84.0	82.4	58.8	20.7	36.4	51.0	0.1	19.5	0.1	22.6
CenterPoint (Group-Free)	LCA=0	88.1	86.3	62.7	24.5	48.5	62.8	0.1	22.2	4.3	32.7
	LCA=1	89.0	87.1	71.6	26.7	50.2	64.7	0.1	29.4	4.5	32.9
	LCA=2	89.1	87.5	71.7	26.8	51.1	65.2	0.1	30.5	4.8	33.4
CenterPoint (Group-Free) + Hierarchy Loss	LCA=0	88.6	86.9	63.4	25.7	50.2	63.2	0.1	25.3	8.7	36.8
	LCA=1	89.5	87.6	72.4	27.5	52.2	65.2	0.1	32.4	9.4	37.0
	LCA=2	89.6	88.0	72.5	27.7	53.2	65.7	0.1	34.0	9.8	37.6
CenterPoint (Group-Free) + Hier. + MMLF	LCA=0	86.3	87.7	60.6	35.3	70.0	75.9	8.8	55.9	37.7	58.1
	LCA=1	86.8	88.3	68.5	37.3	70.4	77.1	16.2	66.0	51.5	58.2
	LCA=2	86.9	88.6	68.6	37.7	70.9	77.4	16.3	69.0	52.4	58.9
CMT [10] (RGB + LiDAR)	LCA=0	88.6	87.7	65.2	36.9	66.7	76.3	4.7	34.4	0.9	34.1
	LCA=1	89.1	88.3	73.3	38.8	67.6	77.2	7.3	50.6	1.2	34.6
	LCA=2	89.1	88.6	73.4	39.0	68.3	77.7	7.7	52.6	1.3	35.6
CMT-L (LiDAR-only) + MMLF	LCA=0	87.1	87.3	62.7	35.3	72.0	76.1	11.9	53.9	29.6	56.4
	LCA=1	87.6	87.8	69.6	36.4	72.3	76.8	16.7	63.4	50.6	56.6
	LCA=2	87.6	88.1	69.7	36.7	72.7	77.3	16.9	66.3	53.2	57.3

our developed techniques including the group-free detector head, the hierarchy loss, and MMLF using the 2D detector DINO. Results reveal that (1) classes are most often misclassified as their LCA=1 siblings within coarse-grained superclasses, (2) our developed techniques significantly improve mAP_H (especially on rare classes), reducing misclassification between sibling classes.

5.2.4 Analysis of Failure Cases

We visualize common failure cases of MMLF and compare them with the failure cases of TransFusion [34], an end-to-end trained multi-modal detector. We find that our MMLF fails in cases of occlusions (where there is no 3D information) and in cases where the 2D RGB detector misclassifies objects. See Fig. 10 for detailed analysis.




Fig. 11: Training 2D detectors with more data (e.g., training with nuScenes + nuImages vs. nuScenes only) and using better 2D detectors (e.g., DINO vs. YOLOV7) improve performance on the proxy task of 2D detection and the downstream 3D detection achieved by our multi-modal late-fusion algorithm.

5.3 Limitations and Future Work

Our work does not directly study how addressing LT3D affects downstream perception tasks. Future work should address this limitation. Moreover, as shown in Fig. 11, simply training better 2D RGB detectors with more data provides a natural pathway for improving LT3D performance. We find that 2D detection accuracy is a strong proxy for final 3D LT3D performance. Recent works in large-scale vision language models [77], [78], [79], [80], [81] show promising zero-shot results in detecting rare classes. Identifying ways of incorporating foundation models into our late-fusion framework might improve LT3D further. In addition, our proposed late-fusion approach is unlikely to work well for objects that are visible in LiDAR but occluded in RGB images (or visible in RGB images but occluded in LiDAR, but this is unlikely because LiDARs are typically mounted higher than RGB cameras). Notably, due to LiDAR and RGB sensors not being co-located on the autonomous vehicle, some objects may be visible in one modality but occluded in another. If an object is visible in LiDAR but not in RGB, we would down-weight the confidence scores of LiDAR-based detections that not supported by RGB-based detections. In contrast, if an object is visible in RGB but not in LiDAR, our 2D late-fusion pipeline would not be able to detect this object in 3D.

6 CONCLUSION

We explore the problem of long-tailed 3D detection (LT3D), detecting objects not only from common classes but also from many rare classes. This problem is motivated by the operational safety of autonomous vehicles (AVs) but has broad robotic applications, e.g., elder-assistive robots [82] that fetch diverse items [83] should address LT3D. To study LT3D, we establish rigorous evaluation protocols that allow for partial credit to better diagnose 3D detectors. We propose several algorithmic innovations to improve LT3D, including a group-free detector head, a hierarchical loss that promotes feature sharing across long-tailed classes, and present a detailed exploration of multi-modal late-fusion for LT3D. We find that 2D RGB detectors are better suited for late-fusion, matching projected 3D LiDAR detections on the 2D image-plane outperforms matching 2D RGB detections inflated to 3D, and score calibration and probabilistic fusion notably improve performance. Our simple multi-modal late-fusion achieves state-of-the-art LT3D performance on the nuScenes benchmark, improving over prior works by 5.9 mAP over all long-tailed distributed classes!

ACKNOWLEDGMENTS


This work was supported by the University of Macau (SRG2023-00044-FST), NSFC (No.62206256), CMU Center for Autonomous Vehicle Research, and NSF GRFP (DGE2140739).

REFERENCES

- [1] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu, "Large-scale long-tailed recognition in an open world," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [2] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3354–3361.
- [3] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan *et al.*, "Argoverse: 3d tracking and forecasting with rich maps," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [4] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, and J. Hays, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," in *Neural Information Processing Systems Datasets and Benchmarks Track*, 2021.
- [5] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, "Scalability in perception for autonomous driving: Waymo open dataset," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [6] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [7] A. Taeihagh and H. S. M. Lim, "Governing autonomous vehicles: emerging responses for safety, liability, privacy, cybersecurity, and industry risks," *Transport Reviews*, vol. 39, no. 1, pp. 103–128, 2019.
- [8] K. Wong, S. Wang, M. Ren, M. Liang, and R. Urtasun, "Identifying unknown instances for autonomous driving," in *CoRL*, 2020.
- [9] N. Peri, A. Dave, D. Ramanan, and S. Kong, "Towards long-tailed 3d detection," in *Conference on Robot Learning (CoRL)*, 2022.
- [10] J. Yan, Y. Liu, J. Sun, F. Jia, S. Li, T. Wang, and X. Zhang, "Cross modal transformer: Towards fast and robust 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 18 268–18 278.
- [11] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. Ni, and H. Shum, "Dino: Detr with improved denoising anchor boxes for end-to-end object detection," in *International Conference on Learning Representations*, 2022.
- [12] T. Yin, X. Zhou, and P. Krahenbuhl, "Center-based 3d object detection and tracking," in *CVPR*, 2021.
- [13] T. Wang, X. Zhu, J. Pang, and D. Lin, "FCOS3D: fully convolutional one-stage monocular 3d object detection," in *ICCV*, 2021.
- [14] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking objects as points," *CoRR*, vol. abs/2004.01177, 2020.
- [15] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *ECCV*, 2016.
- [16] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoryyko, "End-to-end object detection with transformers," in *European Conference on Computer Vision (ECCV)*, 2020.
- [17] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [18] B. Zhu, Z. Jiang, X. Zhou, Z. Li, and G. Yu, "Class-balanced grouping and sampling for point cloud 3d object detection," *arXiv preprint arXiv:1908.09492*, 2019.
- [19] S. Shi, L. Jiang, J. Deng, Z. Wang, C. Guo, J. Shi, X. Wang, and H. Li, "Pv-rcnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection," *International Journal of Computer Vision*, pp. 1–21, 2022.
- [20] Z. Tian, C. Shen, H. Chen, and T. He, "Fcos: Fully convolutional one-stage object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9627–9636.
- [21] J. Huang, G. Huang, Z. Zhu, and D. Du, "Bevdet: High-performance multi-camera 3d object detection in bird-eye-view," *arXiv preprint arXiv:2112.11790*, 2021.

- [22] J. Huang and G. Huang, "Bevdet4d: Exploit temporal cues in multi-camera 3d object detection," *arXiv preprint arXiv:2203.17054*, 2022.
- [23] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, "Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers," in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*. Springer, 2022, pp. 1–18.
- [24] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *ECCV*, 2020.
- [25] Y. Jiang, L. Zhang, Z. Miao, X. Zhu, J. Gao, W. Hu, and Y.-G. Jiang, "Polarformer: Multi-camera 3d object detection with polar transformers," *arXiv preprint arXiv:2206.15398*, 2022.
- [26] V. Guizilini, R. Ambrus, S. Pillai, A. Raventos, and A. Gaidon, "3d packing for self-supervised monocular depth estimation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2485–2494.
- [27] S. Gupta, J. Kanjani, M. Li, F. Ferroni, J. Hays, D. Ramanan, and S. Kong, "Far3det: Towards far-field 3d detection," in *NeurIPS*, 2022.
- [28] S. Vora, A. H. Lang, B. Helou, and O. Beijbom, "Pointpainting: Sequential fusion for 3d object detection," in *IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- [29] T. Yin, X. Zhou, and P. Krähenbühl, "Multimodal virtual point 3d detection," *NeurIPS*, 2021.
- [30] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas, "Frustum pointnets for 3d object detection from rgb-d data," in *IEEE conference on computer vision and pattern recognition*, 2018.
- [31] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.
- [32] D. Xu, D. Anguelov, and A. Jain, "Pointfusion: Deep sensor fusion for 3d bounding box estimation," in *IEEE conference on computer vision and pattern recognition*, 2018.
- [33] Y. Jiao, Z. Jie, S. Chen, J. Chen, X. Wei, L. Ma, and Y.-G. Jiang, "Msmdfusion: Fusing lidar and camera at multiple scales with multi-depth seeds for 3d object detection," *arXiv preprint arXiv:2209.03102*, 2022.
- [34] X. Bai, Z. Hu, X. Zhu, Q. Huang, Y. Chen, H. Fu, and C.-L. Tai, "Transfusion: Robust lidar-camera fusion for 3d object detection with transformers," in *CVPR*, 2022.
- [35] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," *arXiv preprint arXiv:2205.13542*, 2022.
- [36] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [37] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [38] S. Pang, D. Morris, and H. Radha, "Cloccs: Camera-lidar object candidates fusion for 3d object detection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [39] W. J. Reed, "The pareto, zipf and other power laws," *Economics letters*, vol. 74, no. 1, pp. 15–19, 2001.
- [40] Y. Zhang, B. Kang, B. Hooi, S. Yan, and J. Feng, "Deep long-tailed learning: A survey," *arXiv:2110.04596*, 2021.
- [41] S. Alshammari, Y.-X. Wang, D. Ramanan, and S. Kong, "Long-tailed recognition via weight balancing," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [42] Y. Cui, M. Jia, T.-Y. Lin, Y. Song, and S. Belongie, "Class-balanced loss based on effective number of samples," in *CVPR*, 2019.
- [43] S. H. Khan, M. Hayat, M. Bennamoun, F. A. Sohel, and R. Togneri, "Cost-sensitive learning of deep feature representations from imbalanced data," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 8, pp. 3573–3587, 2017.
- [44] K. Cao, C. Wei, A. Gaidon, N. Arechiga, and T. Ma, "Learning imbalanced datasets with label-distribution-aware margin loss," in *NeurIPS*, 2019.
- [45] S. Khan, M. Hayat, S. W. Zamir, J. Shen, and L. Shao, "Striking the right balance with uncertainty," in *CVPR*, 2019.
- [46] C. Huang, Y. Li, C. C. Loy, and X. Tang, "Deep imbalanced learning for face recognition and attribute prediction," *PAMI*, vol. 42, no. 11, pp. 2781–2794, 2019.
- [47] S. Zhang, Z. Li, S. Yan, X. He, and J. Sun, "Distribution alignment: A unified framework for long-tail visual recognition," in *CVPR*, 2021.
- [48] C. Drummond, R. C. Holte *et al.*, "C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling," in *Workshop on learning from imbalanced datasets II*, 2003.
- [49] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: synthetic minority over-sampling technique," *Journal of Artificial Intelligence Research*, vol. 16, pp. 321–357, 2002.
- [50] H. Han, W.-Y. Wang, and B.-H. Mao, "Borderline-smote: a new over-sampling method in imbalanced data sets learning," in *International Conference on Intelligent Computing*. Springer, 2005, pp. 878–887.
- [51] K. Tang, J. Huang, and H. Zhang, "Long-tailed classification by keeping the good and removing the bad momentum causal effect," in *NeurIPS*, 2020.
- [52] A. Gupta, P. Dollar, and R. Girshick, "Lvis: A dataset for large vocabulary instance segmentation," in *CVPR*, 2019.
- [53] N. Peri, M. Li, B. Wilson, Y.-X. Wang, J. Hays, and D. Ramanan, "An empirical analysis of range for 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4074–4083.
- [54] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes challenge: A retrospective," *International Journal of Computer Vision (IJCV)*, 2015.
- [55] T. Lin, M. Maire, S. J. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: common objects in context," in *European Conference on Computer Vision (ECCV)*, 2014.
- [56] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *International journal of computer vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [57] J. Shi, G. Gare, J. Tian, S. Chai, Z. Lin, A. Vasudevan, D. Feng, F. Ferroni, and S. Kong, "Lca-on-the-line: Benchmarking out-of-distribution generalization with class taxonomies," in *International Conference on Machine Learning (ICML)*, 2024.
- [58] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," *arXiv preprint arXiv:2207.02696*, 2022.
- [59] Z. Yang, J. Chen, Z. Miao, W. Li, X. Zhu, and L. Zhang, "Deepinteraction: 3d object detection via modality interaction," in *NeurIPS*, 2022.
- [60] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *ICCV*, 2017.
- [61] Z. Lin, D. Pathak, Y.-X. Wang, D. Ramanan, and S. Kong, "Continual learning with evolving class ontologies," *Advances in Neural Information Processing Systems*, vol. 35, pp. 7671–7684, 2022.
- [62] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan, "Object detection with discriminatively trained part-based models," *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, no. 9, pp. 1627–1645, 2009.
- [63] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems*, 2015.
- [64] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [65] L. H. Li, P. Zhang, H. Zhang, J. Yang, C. Li, Y. Zhong, L. Wang, L. Yuan, L. Zhang, J.-N. Hwang *et al.*, "Grounded language-image pre-training," in *CVPR*, 2022.
- [66] X. Zhou, V. Koltun, and P. Krähenbühl, "Simple multi-dataset detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 7571–7580.
- [67] X. Wang, Z. Cai, D. Gao, and N. Vasconcelos, "Towards universal object detection by domain attention," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 7289–7298.
- [68] H. Xu, L. Fang, X. Liang, W. Kang, and Z. Li, "Universal-rnnc: Universal object detector via transferable graph r-cnn," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 12492–12499.
- [69] J. Redmon and A. Farhadi, "Yolo9000: better, faster, stronger," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 7263–7271.
- [70] B. Wilson, Z. Kira, and J. Hays, "3d for free: Crossmodal transfer learning using hd maps," *arXiv preprint arXiv:2008.10592*, 2020.
- [71] N. Bodla, B. Singh, R. Chellappa, and L. S. Davis, "Soft-nms—improving object detection with one line of code," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 5561–5569.

- [72] Y.-T. Chen, J. Shi, Z. Ye, C. Mertz, D. Ramanan, and S. Kong, "Multimodal object detection via probabilistic ensembling," in *European Conference on Computer Vision (ECCV)*, 2022.
- [73] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *ICML*, 2017.
- [74] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, and S. Belongie, "The inaturalist species classification and detection dataset," in *CVPR*, 2018.
- [75] Z. Liu, Z. Miao, X. Zhan, J. Wang, B. Gong, and S. X. Yu, "Large-scale long-tailed recognition in an open world," in *CVPR*, 2019, pp. 2537–2546.
- [76] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014.
- [77] X. Zhou, R. Girdhar, A. Joulin, P. Krähenbühl, and I. Misra, "Detecting twenty-thousand classes using image-level supervision," in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*. Springer, 2022, pp. 350–368.
- [78] H. Zhang, P. Zhang, X. Hu, Y.-C. Chen, L. Li, X. Dai, L. Wang, L. Yuan, J.-N. Hwang, and J. Gao, "Glipv2: Unifying localization and vision-language understanding," *NeurIPS*, 2022.
- [79] M. Minderer, A. Gritsenko, and N. Houlsby, "Scaling open-vocabulary object detection," *arXiv:2306.09683*, 2023.
- [80] A. Madan, N. Peri, S. Kong, and D. Ramanan, "Revisiting few-shot object detection with vision-language models," *arXiv preprint arXiv:2312.14494*, 2023.
- [81] S. Parashar, Z. Lin, T. Liu, X. Dong, Y. Li, D. Ramanan, J. Caverlee, and S. Kong, "The neglected tails in vision-language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- [82] N. Savage, "Robots rise to meet the challenge of caring for old people," *Nature*, 2022.
- [83] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, and et al., "Ego4d: Around the world in 3, 000 hours of egocentric video," in *Computer Vision and Pattern Recognition*, 2022.



Achal Dave is a Member of Technical Staff at Anthropic. Prior to this, he was a Research Scientist at the Toyota Research Institute and Amazon. He completed his Ph.D. at Carnegie Mellon University, advised by Prof. Deva Ramanan. His research interest centers around efficient scaling of vision and language models. He has co-organized workshops on these topics at leading conferences, including ECCV and CVPR. He received his B.Sc. in EECS from UC Berkeley.



Wei Hua received a PhD degree in Applied Mathematics from Zhejiang University in 2002. He is currently a Senior Research Fellow at Zhejiang Lab. His current research interests include autonomous driving, intelligent simulation, digital twin, reinforcement learning, and their applications to intelligent transportation, smart city, etc.



Yanan Li is currently a researcher in the Research Center for Frontier Fundamental Studies at Zhejiang Lab. She earned her PhD in the Department of Computer Science and Technology from Zhejiang University in 2018. Her research interests include zero-shot learning, few-shot learning, continual learning and their applications in object/instance detection and segmentation.



Deva Ramanan is a Professor in the Robotics Institute at Carnegie-Mellon University and the director of the CMU Argo AI Center for Autonomous Vehicle Research. His research interests span computer vision and machine learning, with a focus on visual recognition. He was awarded the David Marr Prize in 2009, the PASCAL VOC Lifetime Achievement Prize in 2010, the IEEE PAMI Young Researcher Award in 2012, named one of Popular Science's Brilliant 10 researchers in 2012, named a National Academy of Sciences Kavli Fellow in 2013, won the Longuet-Higgins Prize in 2018 for fundamental contributions in computer vision, and was recognized for best paper awards in CVPR 2019, ECCV 2020, and ICCV 2021. His work is supported by NSF, ONR, DARPA, as well as industrial collaborations with Intel, Google, and Microsoft. He served as the program chair of the IEEE Computer Vision and Pattern Recognition (CVPR) 2018. He is on the editorial board of the International Journal of Computer Vision (IJCV) and is an associate editor for the IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI). He regularly serves as a senior program committee member for CVPR, the International Conference on Computer Vision (ICCV), and the European Conference on Computer Vision (ECCV). He also regularly serves on NSF panels for computer vision and machine learning.



Shu Kong is an Assistant Professor in the Faculty of Science and Technology at University of Macau. He was an Assistant Professor in the Department of Computer Science and Engineering at Texas A&M University. He was a postdoc fellow and Project Scientist in the Robotics Institute at Carnegie Mellon University. He earned a PhD degree in Computer Science at the University of California, Irvine. His research lies in Computer Vision and its interactions with other fields (e.g., ML, NLP, HCI, robotics, and graphics), broad applications (e.g., AR/VR, autonomous driving, etc.), and diverse disciplines (e.g., biology, paleoecology, psychology, special education, etc.). His current research focuses on Visual Perception via Learning in the Open World (VPLOW). He has organized workshops on this topic at CVPR from 2021 through 2024. His previous paper on this topic was recognized for Best Paper / Marr Prize at the International Conference on Computer Vision (ICCV) 2021.



Yechi Ma is a PhD student at the Computer Science Department, Zhejiang University. He obtained his MS degree in Electrical Engineering from Zhejiang University in 2018. Since 2019, he has been a Senior Engineer working on visual perception in autonomous driving at Zhejiang Lab. His research interests lie in computer vision and machine learning, specializing in 2D and 3D detection, and their applications in autonomous driving and smart city.



Neehar Peri is a Ph.D. student in Robotics at Carnegie Mellon University. He earned his M.S. in Computer Science at the same institution. He completed his undergraduate degree in Computer Engineering at the University of Maryland, College Park. His research interests include computer vision, machine learning, and robotics.



Shuquan Wei earned his BS and MS degrees in Electrical Engineering from Zhejiang University in 2017 and 2020, respectively. He is currently an engineer at Zhejiang Lab, working on autonomous driving and object detection. His research interests include computer vision and its applications to autonomous vehicles.