

# Modeling and Measuring Redundancy in Multisource Multimodal Data for Autonomous Driving

Yuhan Zhou

Dept. of Information Science  
University of North Texas  
Denton, Texas, USA  
yuhanzhou@my.unt.edu

Mehri Sattari

Dept. of Information Science  
University of North Texas  
Denton, Texas, USA  
MehriSattari@my.unt.edu

Haihua Chen

Dept. of Data Science  
University of North Texas  
Denton, Texas, USA  
Haihua.Chen@unt.edu

Kewei Sha

Dept. of Data Science  
University of North Texas  
Denton, Texas, USA  
Kewei.Sha@unt.edu

**Abstract**—Next-generation autonomous vehicles (AVs) rely on large volumes of multisource and multimodal ( $M^2$ ) data to support real-time decision-making. In practice, data quality (DQ) varies across sources and modalities due to environmental conditions and sensor limitations, yet AV research has largely prioritized algorithm design over DQ analysis. This work focuses on redundancy as a fundamental but underexplored DQ issue in AV datasets. Using the nuScenes and Argoverse 2 (AV2) datasets, we model and measure redundancy in multisource camera data and multimodal image–LiDAR data, and evaluate how removing redundant labels affects the YOLOv8 object detection task. Experimental results show that selectively removing redundant multisource image object labels from cameras with shared fields of view improves detection. In nuScenes, mAP50 gains from 0.66 to 0.70, 0.64 to 0.67, and from 0.53 to 0.55, on three representative overlap regions, while detection on other overlapping camera pairs remains at the baseline even under stronger pruning. In AV2, 4.1–8.6% of labels are removed, and mAP50 stays near the 0.64 baseline. Multimodal analysis also reveals substantial redundancy between image and LiDAR data. These findings demonstrate that redundancy is a measurable and actionable DQ factor with direct implications for AV performance. This work highlights the role of redundancy as a data quality factor in AV perception and motivates a data-centric perspective for evaluating and improving AV datasets. Code, data, and implementation details are publicly available at: <https://github.com/yhZHOU515/RedundancyAD>.

**Index Terms**—Autonomous vehicles, data-centric AI, data quality, multimodal, multisource

## I. INTRODUCTION

Autonomous vehicles (AVs) influence human daily life in transportation mobility and driving safety [1], [2]. It also pushes the frontier of data engineering in data fusion, data simulation, and real-time data processing. The real-world challenge in AVs lies in modeling the highly dynamic environment, which contains long-tail risky cases and uncertainty [1], [3]–[6]. Current AV frameworks and data engineering approaches mainly emphasize benchmark development and evaluation, and model-centric architectural designs [7]. However, raw compute or larger neural networks cannot eliminate risk when the input data from various sensors is noisy, inconsistent, or biased. Despite the significant role of data quality (DQ) in

assuring the performance, efficiency, and reliability of AVs, these strategies not only overlook DQ evaluation but also treat all AV tasks with similar DQ requirements [8]. Consequently, the lack of DQ evaluation and task-centric awareness prevents AV systems from feeding quality diagnostics back into the loop and iterating on performance [4], [5], [9].

However, the data quality work in AV is still restrained due to the absence of a systematic evaluation of data quality dimensions under specific tasks [10], the general standards and definitions for “high-quality” data [5], and limited evidence and explanation on data quality impacts [11]. Prior data work in AV focused more on dataset characteristics [1], [5], [8], [12], such as weather conditions, geospatial coverage, sensor types, dataset sizes, and instance diversity [9], [13]–[15]. In the meantime, few papers have discussed data quality evaluation specifically for AV pipelines [16]. In real-world settings, the quality of different data sources and modalities usually varies due to unexpected environmental factors or sensor issues [17]. Therefore, it is crucial to research the dynamically varying quality of multisource and multimodal ( $M^2$ ) data and its impact on model performance.

Fig. 1 illustrates the  $M^2$  data and the redundancy of an example autonomous vehicle (AV). **Multisource** data are collected from heterogeneous onboard sensors (e.g., cameras, LiDAR, and RADAR), which offer complementary yet partially redundant representations of the same driving scene. These data are further integrated with external inputs such as GPS, traffic, infrastructure, and weather data, as well as data shared by connected autonomous vehicles (CAVs). The redundancy across multisource data improves situational awareness and system robustness by mitigating sensor uncertainty and failures, thereby supporting more reliable perception and decision making [18]. **Multimodal** data combines signals captured through different sensing modalities. For example, images from cameras and point clouds from LiDAR can contribute complementary information about the same scene, enhancing the accuracy of AV tasks such as object detection and localization.

Redundancy in autonomous driving object detection arises

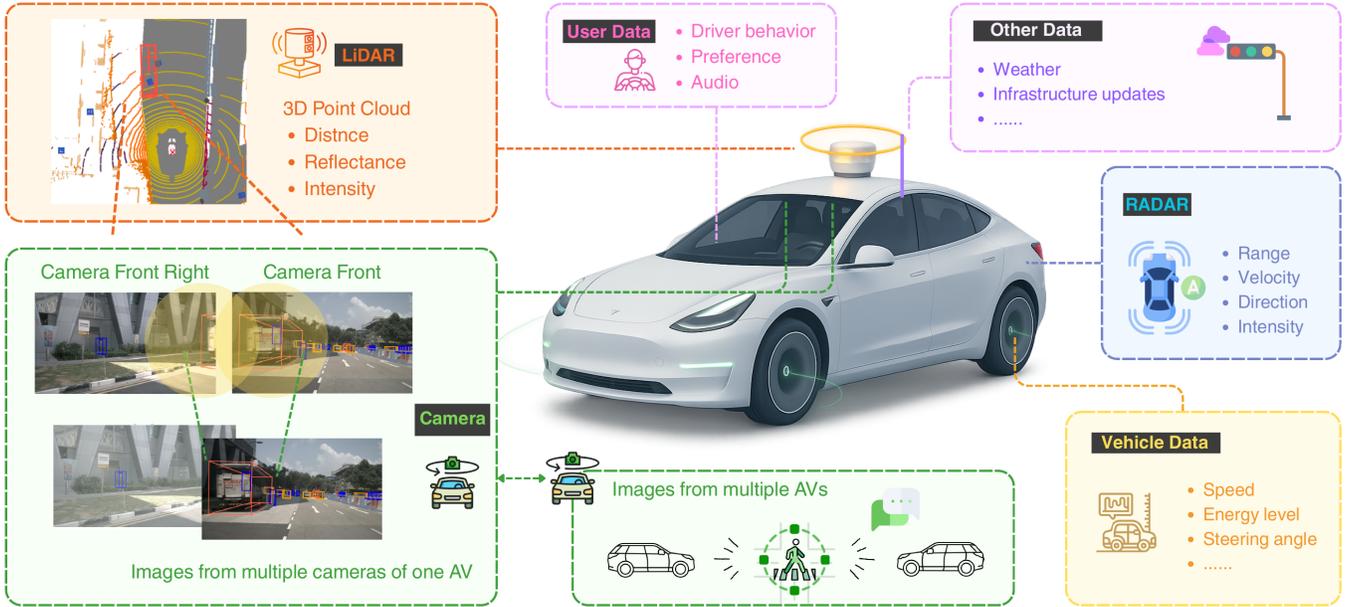


Fig. 1: Illustration of multisource and multimodal data in autonomous vehicles (AVs). LiDAR point clouds, multi-camera images, RADAR, vehicle dynamics, user-related data, and external contextual inputs feed into the AV system. Redundancy arises from overlapping fields of view among multiple onboard cameras (green), as well as from multimodal sensing in which cameras and LiDAR (orange) observe the same objects. The rest colored blocks represent additional data sources and modalities.

from multiple sensors on the same vehicle [19]. Although redundancy can improve reliability, its absence of explicit quantification poses major risks: duplicate detections increase computational costs and harm real-time efficiency [20]–[23]. Inconsistent redundant predictions can cause noise, degrading localization and confidence. Thus, modeling and measuring redundancy in a specific AV task is crucial, which can help spatial verification and fault tolerance without overwhelming the detection pipeline [24], [25]. This motivates us to model and measure  $M^2$  redundancy in the object detection task in AD. Based on the above, the research questions are formulated as follows.

- **RQ1:** How to define and model redundancy in multi-source and multimodal AV data for object detection?
- **RQ2:** How to measure redundancy and select the optimal data subset for object detection in AV?
- **RQ3:** How does redundancy removal affect object detection model performance?

Fig. 2 illustrates our overall research design, where each component directly addresses one of the three research questions. To answer RQ1, we evaluate redundancy across multisource (camera–camera) and multimodal (camera–LiDAR) observations in the object detection task. To answer RQ2, we investigate which redundant observations to remove and propose pruning strategies that retain the most complete object representations. To answer RQ3, we train and evaluate YOLOv8 using datasets curated with different redundancy levels, and analyze how redundancy removal influences detection performance and data efficiency.

The contributions of this paper can be summarized as follows.

- 1) For the first time, we model and measure redundancy in multisource and multimodal data in the object detection task. We design practical redundancy measurement strategies for camera–camera and camera–LiDAR data, highlighting the importance of data-quality research in autonomous driving.
- 2) We propose a task-driven data selection method based on bounding box completeness and spatial overlap constraints. We validate that our redundancy measurement and selection pipeline is not dataset-specific but generalizable to AV benchmarks.
- 3) We evaluate redundancy reduction on YOLOv8 object detection using nuScenes and Argoverse 2. The results show that reducing redundancy can maintain or even improve detection performance.
- 4) We reveal and analyze cross-modal redundancy between camera images and LiDAR point clouds, providing empirical evidence that multimodal redundancy is substantial and must be explicitly measured to balance robustness and efficiency.

The rest of the paper is organized as follows. Section II provides a literature review on the  $M^2$  data and data quality issues in AV, and specifically, the redundancy evaluation. Section III introduces the research design, object detection task, the model, and performance metrics. Section IV provides details of the datasets and experiments. Section V and Section VI present the experiment results and discussions, respectively.

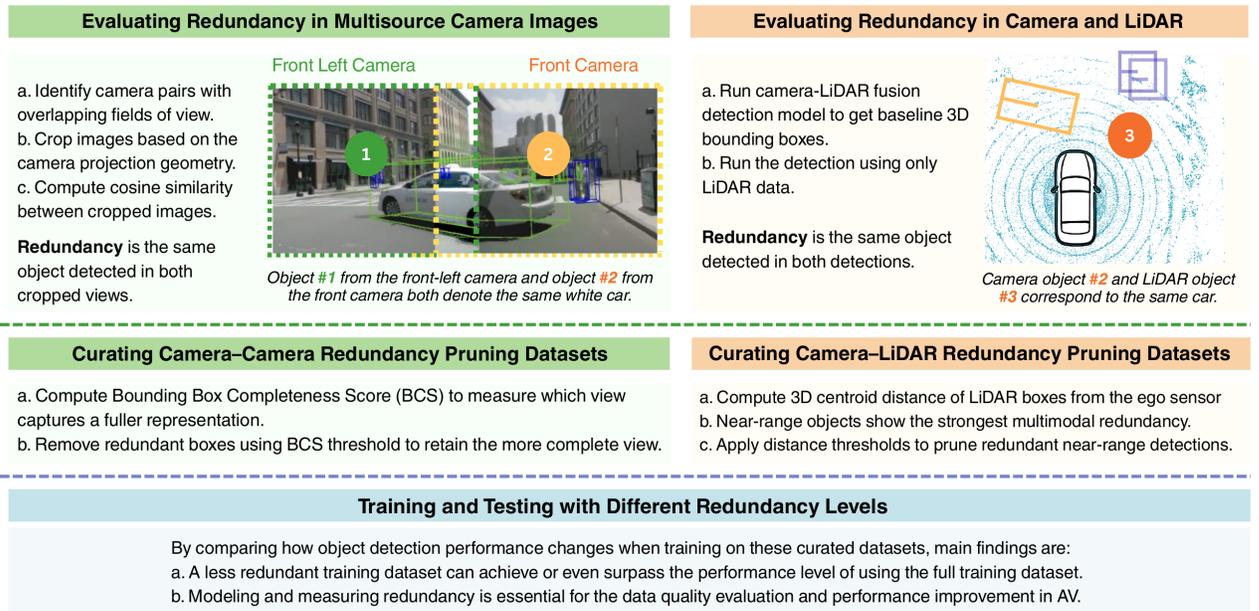


Fig. 2: Illustration of research design, using nuScenes dataset as examples. We first quantify redundancy in multisource and multimodal autonomous driving datasets (RQ1), then design data pruning to remove redundant observations (RQ2), and finally evaluate how redundancy removal affects object detection performance (RQ3).

Finally, Section VII concludes our work and puts forward future directions.

## II. RELATED WORK

### A. Multisource and Multimodal data in Autonomous Driving

At the **multisource** level, AVs gather information from various sensors, such as LiDAR, cameras, and RADAR, as well as other feeds including user behavior data, weather reports, and infrastructure updates. Since one AV can be equipped with more than one camera, the front, back, and side cameras are considered different data sources [26], [27]. The integration of multisource data can validate data and reduce the impact of errors or noises from individual sensors, which enables a more comprehensive view of the driving environment and a more accurate and robust performance [18], [28], [29], enhancing the system’s situational awareness, reliability, and robustness.

**Multimodal** data in AV refers to information obtained through different modes of data collection, mainly including various types of sensory data such as visual (camera), auditory (user), and spatial (LiDAR, RADAR) data, as shown in Fig. 1. Each mode provides unique and complementary insights that, when combined, offer a more holistic understanding of the environment. Multimodal data fusion enables these systems to perform complex tasks such as object detection, localization, and navigation with higher accuracy and reliability by leveraging the strengths of each data modality.

Combining  $M^2$  data can enhance performance and reliability in AV compared to a single data source or modality [30], [31]. Zhang et al. improved vehicle localization by integrating vehicle headlights and taillights detected from multiple

cameras [32]. Sanchez et al. found that single-source LiDAR detectors generalize poorly, and multisource LiDAR training can improve domain generalization by over 10% and accuracy by 5.3% over any individual dataset [33]. When LiDAR points are sparse for tiny objects, a camera could help recover them by visual cues [34]–[37].

### B. Data Quality in Autonomous Driving

Data quality (DQ) affects machine learning performance [38]. High-quality training data enhances accuracy, while low-quality data can lead to biased or erroneous outcomes [39], [40]. Recognizing this, scholars have developed frameworks [41]–[44] and DQ dimensions, such as accuracy [45], completeness [46], consistency [47], [48], redundancy [49] and fairness [50].

Data Quality in autonomous driving is the backbone for assuring performance, efficiency, reliability, and robustness. Particularly,  $M^2$  data in CAVs are complex to handle, making the DQ evaluation more significant and challenging. In real-world settings, the quality of different modalities and data sources usually varies due to unexpected environmental factors or sensor issues [17]. This brings major challenges (but not limited to): (1) mitigate the underlying influence of arbitrary noise, bias, and discrepancy, (2) enhance data quality (DQ) in terms of different dimensions, (3) adapt the quality of dynamically varying nature of  $M^2$  data.

However, there is a lack of systematic evaluation of data quality dimensions under specific tasks [10]. Most existing studies primarily focus on dataset characteristics [8], [12], such as geospatial coverage, sensor types, and instance diversity

[10]. Few studies have defined specific DQ dimensions of  $M^2$  data in AVs from a task-centric perspective and answered the question of how to select useful data.

### C. Redundancy in Autonomous Driving Datasets

By having multiple sources of similar information, the system can validate data, reducing the impact of errors or noise from individual sensors. When one sensor fails or generates deviated data, the other sensors can still perform [24], [25], [51]. This enables a more accurate and robust performance [28], [29]. In object detection, multiple sensors are set up to obtain comprehensive and reliable environmental information [35], [36].

Redundancy can negatively impact the AD system. Firstly, it will decrease the efficiency in computational time and storage [20]–[23]. For example, when the weather conditions are good and clear, there is less need to keep high-overlapping information from LiDAR and cameras. The models trained on such data would have little information gain and performance improvement [52], [53]. The redundant information includes the temporal correlation between point cloud scans, similar urban environments, and the symmetries in the driving environment caused by driving in opposite directions at the same location [21], [52]. As an example, after randomly removing 30% of the redundant background cloud points, Zhao et al. found only a subtle drop in performance, which proves the necessity and efficiency of researching the redundancy data quality issues in the datasets [20].

Evaluating cross-modal redundancy poses challenges. Each modality has a different informative level and contributes to different tasks [54]. Reducing redundancy is not simply removing duplicate records, as in previous single-modality data analysis tasks. It requires investigations on the sensor characteristics, task specificity, performance experiments, trade-off evaluation, and so on. Li et al. view redundancy as a content data quality issue and divide redundancy into element redundancy and data redundancy [55], quantifying how much one or several elements co-occur and repeat the same behavior for longer than the threshold. Duong et al. define redundancy as the similarity between pairs of point clouds resulting from geometric transformations of ego-vehicle movement and environmental changes [52]. Rao et al. also calculated the similarity between the images and text pairs [56]. From an information science perspective, calculating mutual information is also an approach [57].

## III. METHODOLOGY

### A. Research Design

This section introduces the research design on redundancy evaluation across the multisource camera views and the multimodal camera and LiDAR data for the object detection task. Fig. 3 introduces the workflow of experiments. (i) For multisource redundancy evaluation, when 2D labels are present, we directly estimate overlapping fields of view (FoVs) between camera pairs and crop shared regions. For datasets that provide 3D annotations, we project 3D cuboids and identify

object views based on projection overlap. We then design a Bounding-Box Completeness Score (BCS) to measure how fully each object instance is captured per view, and remove redundant camera observations by BCS thresholding. (ii) To measure multimodal redundancy, we compute the 3D centroid of the predicted box using LiDAR data and measure its ego-centric distance from the vehicle sensor origin. Rather than pruning by confidence alone, we remove LiDAR boxes within a foreground distance threshold where visual coverage is already strong, preserving observations for distant or visually incomplete objects.

The research design is guided by three considerations, consistent with our objective to model and measure redundancy in AD. (i) Overlapping FoVs in multisource images produce duplicate 2D labels for the same object, requiring explicit spatial redundancy modeling. (ii) In camera–LiDAR fusion, object importance is strongly conditioned on sensor range, motivating distance-aware redundancy measurement and pruning beyond IoU-based criteria. (iii) Redundancy removal must be tailored to the specific task. To preserve the most informative instance per object in object detection, we adopt BCS to retain the most complete 2D camera observation in overlapping views and apply LiDAR box centroid-to-ego distance thresholds to filter temporally correlated 3D duplicates.

### B. Task Formulation

Our task objective is to evaluate whether removing redundant observations of the same object across sources and modalities can improve object detection performance. The task is defined as learning a mapping

$$f_{\theta} : \mathcal{X} \rightarrow \mathcal{Y}, \quad (1)$$

where  $\mathcal{X}$  denotes multisource or multimodal inputs (overlapping view pairs) and  $\mathcal{Y}$  denotes object detection outputs.

### C. Redundancy Modeling and Evaluation

1) *Modeling Redundancy in Multisource Data:* We consider overlapping camera views that can produce multiple observations of the same physical object. For each object instance  $o_i$  at time step  $t$ , let  $\mathcal{P}_t$  denote the set of overlapping camera-view pairs at time  $t$  (obtained either from dataset-defined overlap pairs or from calibration-driven FoV overlap estimation). For a selected pair  $(c_a, c_b) \in \mathcal{P}_t$ , the input is a pair of image regions extracted from the overlapping fields of view (FoVs) of the two cameras,

$$\mathcal{X}_{t,i}^{\text{src}} = (\tilde{I}_{t,i}^{(a)}, \tilde{I}_{t,i}^{(b)}), \quad (2)$$

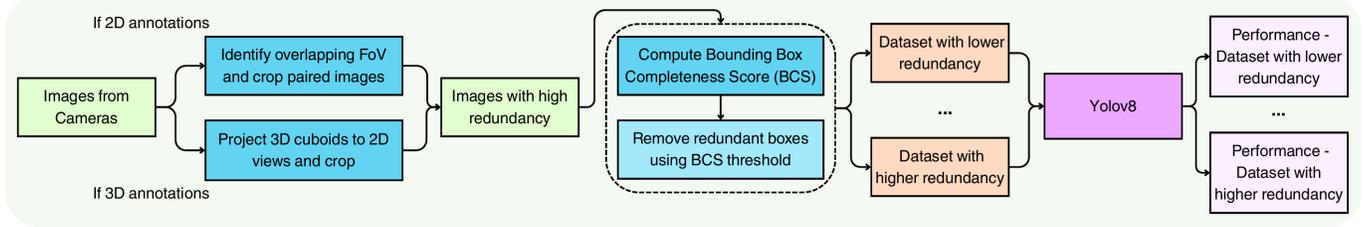
where  $\tilde{I}_{t,i}^{(a)}$  and  $\tilde{I}_{t,i}^{(b)}$  are RGB image regions from cameras  $c_a$  and  $c_b$  corresponding to their overlapping FoV and potentially observing the same physical object.

Given an input  $\mathcal{X}_{t,i}$ , the detector predicts

$$\hat{\mathcal{Y}}_{t,i} = \{(\hat{b}_i, \hat{c}_i, \hat{s}_i)\}, \quad (3)$$

where  $\hat{b}_i$  denotes the bounding box,  $\hat{c}_i$  the class label, and  $\hat{s}_i$  the confidence score.

### Research Design 1 — Multisource (Camera–Camera)



### Research Design 2 — Multimodal (Camera–LiDAR)

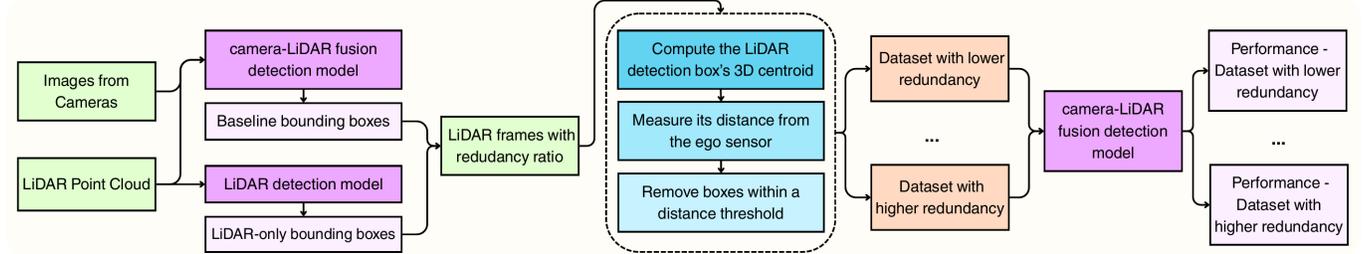


Fig. 3: The workflow of the research design on redundancy evaluation. nuScenes emphasizes 2D labeling and AV2 adopts 3D annotations. Beyond dataset-specific configurations, this workflow can be generalized to other AD datasets.

2) *Modeling Redundancy in Multimodal Data*: The input consists of paired frames from image and LiDAR modalities,

$$\mathcal{X}_{t,i}^{\text{mod}} = (I_{t,i}, L_{t,i}), \quad (4)$$

where  $I_{t,i}$  and  $L_{t,i}$  denote synchronized image and LiDAR observations. The output is the same as the multisource data redundancy evaluation.

#### D. Model for Object Detection

Our redundancy evaluation is conducted under the object detection (OD) task, where duplicated inputs can directly influence localization quality, model bias, and training efficiency. OD aims to jointly localize and classify objects of interest, enabling autonomous driving systems to parse dynamic road scenes and support safety-critical decision-making [36]. OD pipelines typically rely on heterogeneous sensors, including cameras, LiDAR, and RADAR. Given  $M^2$  input, the model outputs 2D and/or 3D bounding boxes for detected objects. Representative OD architectures include CNN-based detectors (e.g., YOLO [58], VoxelNet [59], PointNet [60]), transformer models [61], and multimodal fusion methods such as BEVFusion [62] and BEV-space camera–LiDAR frameworks.

To evaluate the impact of redundancy pruning, we require a model with a strong and stable baseline accuracy to observe performance shifts from redundant supervision. We therefore adopt YOLOv8, a high-performance, real-time CNN detector on AD benchmarks. YOLOv8 was trained with a batch size of 16 and 50 epochs for the object detection task for evaluation.

#### E. Performance Evaluation Metrics

Object detection performance is commonly measured using mean Average Precision (mAP) and recall. In this work,

mAP50 is used to assess bounding box alignment with ground-truth labels at  $\geq 50\%$  IoU, while recall measures the fraction of true objects successfully detected.

Given multiple redundancy levels  $\{r_1, r_2, \dots, r_m\}$  constructed under both multisource and multimodal settings, we compare

$$\text{Performance}(\mathcal{X}^{(r_1)}) \text{ vs. } \text{Performance}(\mathcal{X}^{(r_2)}), \quad (5)$$

across these settings to evaluate the effect of redundancy on object detection performance.

## IV. EXPERIMENTS

This section introduces the datasets and experiments on redundancy evaluation to validate the proposed methodology.

#### A. Dataset

The datasets used are nuScenes-mini, nuScenes-in-KITTI, and Argoverse 2. We select nuScenes and Argoverse 2 due to their suitability for this redundancy evaluation study. First of all, they are well-known, large-scale, multisensor AD benchmarks. Secondly, both datasets include public calibration files and dense annotations, which enable a robust evaluation of redundancy across spatial and temporal axes. nuScenes provides a 360-degree multi-camera rig with substantial inter-view FoV overlap and densely annotated 3D boxes, making it well-suited for studying camera–camera redundancy and cross-modal pruning in fusion models. The details are introduced in Table I and as follows.

a) *nuScenes-mini*: The nuScenes dataset is a large-scale, multimodal dataset designed for autonomous driving research [13]. It includes 360-degree sensor coverage with data collected from six cameras, five RADARs, and one LiDAR sensor, along with high-definition maps and detailed object

TABLE I: Summary of datasets used for redundancy evaluation

Dataset	Year	Sensors	Annotations	Size	Modalities
nuScenes-mini [13]	2020	6 cameras (12 Hz), 1 LiDAR, RADAR, GPS/IMU	3D bounding boxes, object attributes, tracking, HD maps	10 scenes, 404 frames	Multisource camera images
nuScenes-in-KITTI [63]	2021	Cameras, 1 LiDAR, GPS/IMU	2D/3D bounding boxes, calibration, tracking	85 scenes in KITTI format	Image and LiDAR
Argoverse 2 [64]	2023	9 cameras, 2 LiDARs (10 Hz), GPS/IMU	3D bounding boxes, tracking, map annotations, motion labels	1,000 scenes, 20,000 LiDAR sequences	Multisource camera images

annotations. The cameras are of 12Hz capture frequency. The full dataset consists of 1,000 driving scenes, providing a rich resource for perception tasks such as object detection, tracking, and scene understanding. The nuScenes-mini dataset is a smaller subset containing 10 scenes and 404 frames with the same sensor setting and annotation structure.

*b) nuScenes-in-KITTI:* nuScenes-in-KITTI [63] is a nuScenes version that is compatible with the KITTI format [65]. KITTI is a collection of synchronized sensor recordings from a car equipped with cameras, LiDAR, and GPS/IMU, driving through city, residential, campus, rural, and highway scenes in Germany.

*c) Argoverse 2:* Argoverse 2 is designed to support robust 3D perception, tracking, and motion forecasting research [64]. It provides multimodal sensor data collected across diverse urban environments. It consists of 1,000 annotated driving sequences with dense 3D bounding box annotations and tracking labels. The dataset’s scale, sensor redundancy, and diverse scenarios make it well-suited for evaluating redundancy in both image and LiDAR modalities.

*B. Evaluating Redundancy in Multisource data in nuScenes*

This experiment focuses on the multisource camera images. As shown in Fig. 4, the nuScenes camera setup has six pairs of overlapping fields of view (FoV). These overlapping parts indicate areas where cross-camera redundancy may occur. In other words, the front right camera and the back left camera can not capture the same instance simultaneously. Therefore, these six pairs are our research focus. To illustrate our framework by redundancy in this kind of multisource data, we experiment in the following steps:

(a) Identify the overlapping FoV based on the nuScenes dataset sensor setting:

- **Pair 1:** CAM\_FRONT and CAM\_FRONT\_RIGHT overlap by 15°
- **Pair 2:** CAM\_FRONT and CAM\_FRONT\_LEFT overlap by 15°
- **Pair 3:** CAM\_FRONT\_RIGHT and CAM\_BACK\_RIGHT overlap by 15°
- **Pair 4:** CAM\_FRONT\_LEFT and CAM\_BACK\_LEFT overlap by 15°
- **Pair 5:** The rear camera CAM\_BACK overlaps with CAM\_BACK\_RIGHT by 20°
- **Pair 6:** The rear camera CAM\_BACK overlaps with CAM\_BACK\_LEFT by 20°

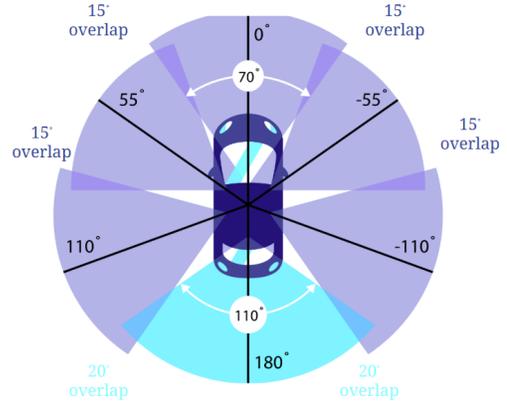


Fig. 4: Camera settings and overlapping fields of view (FoVs). Note that the degree of FoV overlap and camera configuration shown here are specific to nuScenes. For datasets with different sensor layouts, the same redundancy modeling and measurement approach remains applicable by adjusting the geometric parameters.

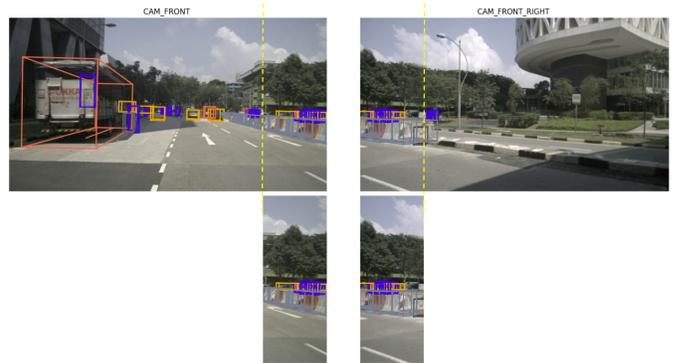


Fig. 5: Illustration of cropping based on overlapping angles of the cameras

(b) Crop images based on overlapped angles: Fig. 5 illustrates the process of cropping each pair of camera images based on the overlapping FoV. Redundancy occurs in the bottom two cropped images.

(c) Calculate cosine similarity: for the cropped images in the six pairs, we calculate the cosine similarity of each sample and exclude the ones without redundant instances. This provides us with a preliminary understanding of redundancy.

(d) Create different levels of redundant training datasets to investigate how redundancy affects the inference performance: for each pair of overlapping detections, we compute a **Bounding Box Completeness Score (BCS)**, which indicates how completely the bounding box (BB<sub>ox</sub>) presents the instance.

Let  $\text{BB}_{\text{ox}}^{\text{full}}$  be the original (uncropped) 2D bounding box area in the image, and let  $\text{BB}_{\text{ox}}^{\text{clipped}}$  be the visible portion after clipping to the image boundaries. We define:

$$\text{BCS}(b) = \frac{\text{Area}(\text{BB}_{\text{ox}}^{\text{clipped}}(b))}{\text{Area}(\text{BB}_{\text{ox}}^{\text{full}}(b))}, \quad (6)$$

where  $b$  indexes a candidate box. Within each redundant group, if

$$\max_b \text{BCS}(b) - \min_b \text{BCS}(b) > \tau_{\text{BCS}}, \quad (7)$$

We retain only the box with the higher BCS and discard the lower one; otherwise, we preserve both boxes. As the threshold  $\tau_{\text{BCS}}$  increases, fewer boxes are removed, thus retaining more redundancy while still preferring more complete annotations.

(e) Train YOLOv8 on each pair of overlapping camera images: from the previous step, we obtain training sets with different levels of redundancy. Next, we train the model using these training datasets and evaluate how removing redundancy affects the inference performance.

### C. Evaluating Redundancy in Multisource Data in Argoverse 2

For Argoverse 2 (AV2), the multisource redundancy experiment follows the same overall pipeline as the nuScenes multi-source study, but is adapted to AV2’s log-based organization, surround-view sensor layout, and 3D-first annotation format. Table II summarizes the methodological alignment between the nuScenes and AV2 multisource experiments, highlighting the shared instance-level BCS-guided pruning framework while clarifying dataset-specific differences in sensor layout, annotation space, and calibration-driven overlap pairing.

Conceptually, the two datasets share the same core design: object-centric supervision with track-level identifiers, a BCS to compare alternative views, and a consistent YOLOv8 configuration so that observed performance changes are attributable to redundancy removal rather than model hyperparameters. In nuScenes-mini, redundancy candidates may be pre-screened using overlap-region similarity cues before BCS-guided pruning. In AV2, redundancy is defined from 3D ego-frame cuboids that are projected into each camera view to obtain 2D boxes and their BCS scores. In both cases, pruning is applied at the instance level: the lower-BCS observation is removed when the BCS gap between two overlapping views exceeds the pruning threshold  $\tau_{\text{BCS}}$ .

On the AV2 side, multiple logs are converted into a unified multi-camera YOLO dataset, with a shared train/validation split defined over  $(\text{log\_id}, \text{timestamp})$  keys and a global class map across categories. Camera overlap pairs are not hard-coded; instead, they are derived from calibration by estimating each camera’s yaw center and horizontal field of view (FoV) from extrinsics and intrinsics, then computing angular overlap

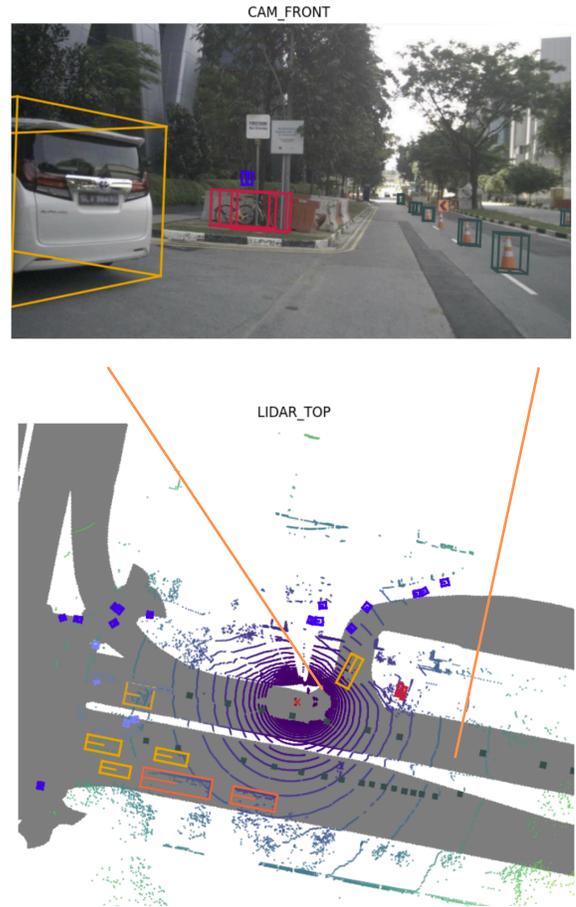


Fig. 6: Illustration of the redundancy between the Front Camera and LiDAR in nuScenes

on the viewing circle to select overlapping camera pairs for pruning. Using this overlap graph, the experiment constructs a baseline (unpruned) dataset and a family of BCS-pruned datasets that differ only in the pruning threshold  $\tau_{\text{BCS}}$ . Among these,  $\tau_{\text{BCS}} = 0.5$  provides a balanced operating point in our sweep, removing redundant labels while keeping performance close to the unpruned baseline.

### D. Evaluating Redundancy in Multimodal Data

Fig. 6 shows an example of multimodal data redundancy in the nuScenes dataset. The LiDAR on top and the front camera capture the same objects. This experiment investigates the redundancy between these two sensors:

(a) Run the full camera-LiDAR fusion detection model [66] to get baseline 3D bounding boxes. Check the overlapping detections between this and the detection using only LiDAR data as redundancy. Let  $\mathcal{B}_{\text{LiDAR}}$  be the set of boxes detected using LiDAR only, and  $\mathcal{B}_{\text{base}}$  the set detected by the image-LiDAR fusion. The *redundancy ratio* for one frame is

$$\text{RR} = \frac{|\{b \in \mathcal{B}_{\text{base}} \mid \exists b' \in \mathcal{B}_{\text{LiDAR}} : \text{IoU}(b, b') \geq \theta\}|}{|\mathcal{B}_{\text{base}}|}. \quad (8)$$

TABLE II: Comparison between the redundancy experiments on nuScenes and AV2 datasets. Both apply instance-level, BCS-guided pruning; AV2 adds calibration-driven overlap pairing and 3D-to-2D projection. All the experiments were conducted on a PC equipped with one NVIDIA GeForce 4090 RTX GPU and 24 GB of memory. Model configurations are fixed across experiments.

Aspect	nuScenes multisource	AV2 multisource
Sensor layout	6 calibrated surround-view cameras with predefined overlapping FoVs.	9 calibrated cameras (7 ring + 2 front stereo); surround-view configuration with denser overlap.
Source of supervision	Object identity inferred implicitly from dataset annotations and temporal alignment.	Explicit object identity via <code>track_uuid</code> from 3D ego-frame cuboids.
Annotation space	2D image-space object-centric observations derived from overlap regions.	3D-first cuboids in a unified ego frame, projected into each camera view to obtain 2D observations.
Redundancy unit	Multiple 2D observations of the same object instance across overlapping cameras.	Multiple projected 2D observations originating from the same 3D object instance across cameras.
Overlap pairing	Fixed camera overlap pairs defined by dataset specification.	Camera overlap pairs computed from calibration (yaw center and horizontal FoV) using angular overlap.
BCS computation	Boundary-aware 2D coverage score computed on valid image regions.	Boundary-aware 2D coverage score computed after projecting 3D cuboids and clipping to image bounds.
Pruning rule	Drop the lower-BCS observation if the BCS gap exceeds $\tau_{\text{BCS}}$ .	Same instance-level rule applied to projected views.
Training protocol	YOLOv8 with matched hyperparameters to isolate redundancy effects.	Same YOLOv8 configuration; baseline (unpruned) and BCS-pruned datasets differ only by $\tau_{\text{BCS}}$ .

(b) Compute each box’s 3D centroid in the LiDAR frame, and measure its distance from the ego sensor:

$$\mathbf{c}(b) = \frac{1}{8} \sum_{i=1}^8 \mathbf{v}_i \quad \text{and} \quad d(b) = \|\mathbf{c}(b)\|_2, \quad (9)$$

where  $\{\mathbf{v}_i\}$  are the eight corner vertices of  $b$ . We then set a single pruning threshold  $T_{\text{dist}}$  (meters). All boxes whose centroids lie within this distance are removed:

$$\mathcal{B}_{\text{pruned}} = \{\mathcal{B}_{\text{LiDAR}} \mid d(\mathcal{B}_{\text{LiDAR}}) \geq T_{\text{dist}}\}. \quad (10)$$

By sweeping  $T_{\text{dist}}$ , we could trace how many boxes are pruned and what fraction of true detections is lost, thus choosing a trade-off point that maximally reduces redundancy without excessively harming detection performance. This distance threshold is chosen based on statistical results explained in the next section.

(c) Run LiDAR detection again using the partially removed LiDAR data and observe how removing redundancy affects the performance. Let  $\mathcal{B}_{\text{pruned}}$  be the set after pruning. We define the *lost ratio*  $\ell$  as the fraction of baseline boxes removed due to removing the redundancy:

$$\ell = \frac{|\mathcal{B}_{\text{base}} \setminus \mathcal{B}_{\text{pruned}}|}{|\mathcal{B}_{\text{base}}|} = 1 - \frac{|\mathcal{B}_{\text{base}} \cap \mathcal{B}_{\text{pruned}}|}{|\mathcal{B}_{\text{base}}|}. \quad (11)$$

## V. RESULTS

### A. Overall Results

We conduct the multisource redundancy experiments on two datasets under a unified protocol. We first establish the workflow on nuScenes and then repeat the same redundancy pruning procedure on Argoverse 2 (AV2) under a controlled

evaluation setup. Table III summarizes the 2D detection performance. For each dataset, we report results for the unpruned baseline (trained and evaluated using all valid camera views) and the results using pruned training datasets.

- **nuScenes.** The baseline detector was evaluated on 1,401 images containing 12,286 labeled object instances across 14 categories. Overall, the detector achieved a box precision of 0.780, a recall of 0.630, and an mAP50 of 0.720. These results serve as the nuScenes baseline for subsequent multisource redundancy analysis. Using a threshold of  $\tau_{\text{BCS}} = 0.8$  on pair 4, which removes approximately 7% of the baseline, detection performances improved and became the best among all pairs.
- **Argoverse 2 (AV2).** For AV2, the baseline YOLOv8 was evaluated on the full training set, which contains 187,265 projected 2D labels corresponding to 95,266 3D object tracks across 19 classes. Using a threshold of  $\tau_{\text{BCS}} = 0.5$ , which removes 9,442 labels (about 5.0% of the baseline), precision performance remained close, with only modest reductions in recall and mAP50.

Compared to the AV2 baseline, pruning at  $\tau_{\text{BCS}} = 0.5$  yields a modest reduction in recall and mAP, while precision remains slightly higher than the unpruned baseline. Overall, detection performance remains close to the baseline despite removing redundant camera-level supervision. A more detailed analysis of the performance–redundancy trade-off across pruning thresholds is presented in the following sections.

### B. Impact of Redundancy in Multisource Data in nuScenes

For multisource redundancy evaluation, Fig. 7 shows an example of high-similarity (left) and low-similarity (right) views in the front and front left cameras in nuScenes. In high-similarity examples (left), the detected objects and background

TABLE III: Multisource detection performance with and without BCS-guided pruning

Dataset	Experiment	Prec.	Rec.	mAP50	mAP50 : 95
nuScenes	Train unpruned, Eval unpruned	0.780	0.630	0.720	0.450
nuScenes	Train pruned ( $\tau_{BCS}=0.8$ on pair 4), Eval unpruned	0.867	0.800	0.770	0.638
AV2	Train unpruned, Eval unpruned	0.815	0.565	0.640	0.447
AV2	Train pruned ( $\tau_{BCS}=0.5$ ), Eval unpruned	0.818	0.557	0.622	0.433

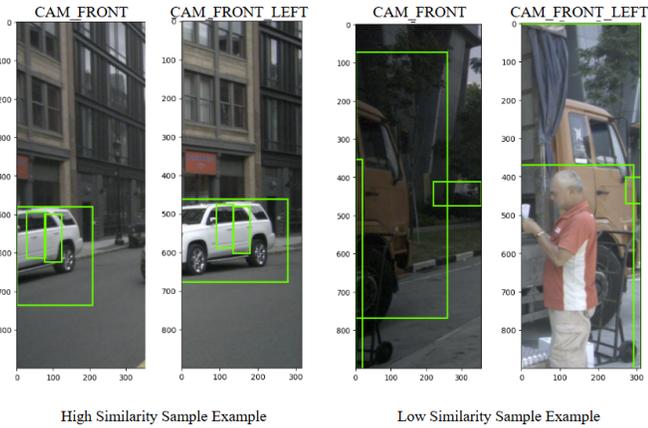


Fig. 7: Examples of high and low similarity samples in nuScenes

almost provide the same information. By contrast, the low-similarity ones (right) present more complementary detections.

Fig. 8 shows how removing redundancy affects detection performance on six pairs in Fig. 4. For each pair, the thresholds are set from 0.0 to 1.0, with a 0.2 interval. As the threshold increases, more instances in the training dataset are retained, meaning that more redundancy is preserved until the 1.0 threshold retains all instances for training. The trends reveal that a less redundant training dataset can achieve or even surpass the performance level of using the full training dataset.

### C. Impact of Redundancy in Multisource data in Argoverse 2

To examine whether multisource redundancy generalizes beyond nuScenes, we repeat the instance-level BCS-guided pruning experiment on Argoverse 2 (AV2). With different  $\tau_{BCS}$  thresholds, Table IV summarizes the redundancy–performance trade-off on AV2.

Across the sweep, BCS-based pruning removes a non-trivial fraction of camera-level labels (approximately 4.1–8.6%, depending on  $\tau_{BCS}$ ), while the number of unique 3D object tracks observed during training remains constant (95,266 tracks). This confirms that redundancy is primarily introduced by overlapping redundant multi-camera observations rather than by removing distinct object instances.

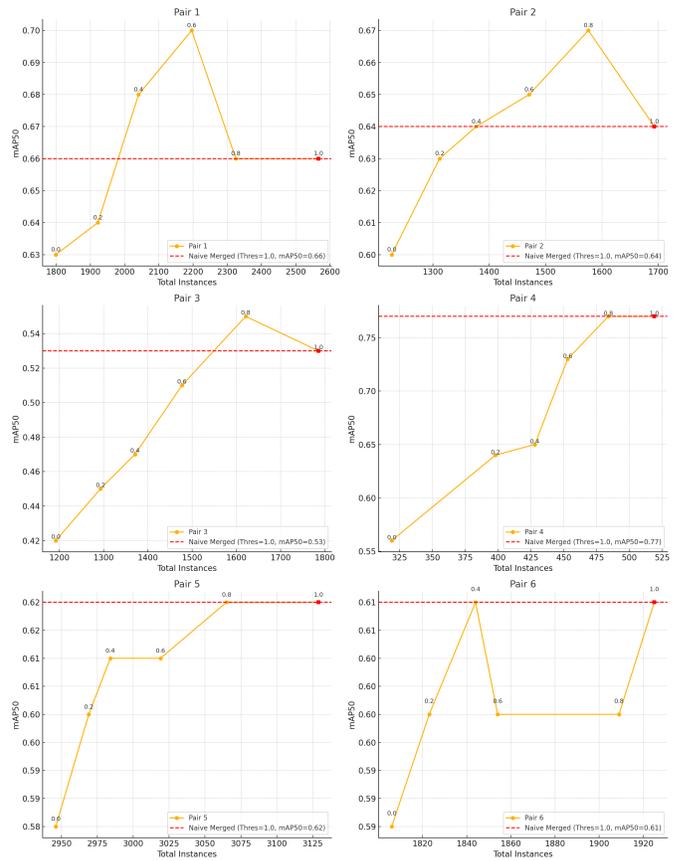


Fig. 8: Effects of removing redundant instances in overlapping FoV images on mAP50 in nuScenes

TABLE IV: AV2 multisource BCS pruning sweep. Training uses different  $\tau_{BCS}$  values while evaluation is fixed to the same unpruned validation split. “Del.” and “Rem.” denote deleted and remaining camera-level 2D labels in the training set. “Tracks” is the number of unique 3D object tracks observed during training.

Train	$\tau_{BCS}$	Del.	Rem.	Tracks	Prec.	Rec.	mAP50
Baseline	–	0	187,265	95,266	0.815	0.565	0.640
Pruned	0.1	16,056	171,209	95,266	0.811	0.531	0.604
Pruned	0.2	14,345	172,920	95,266	0.788	0.544	0.605
Pruned	0.3	12,810	174,455	95,266	0.826	0.538	0.617
Pruned	0.4	11,018	176,247	95,266	0.823	0.545	0.618
Pruned	0.5	9,442	177,823	95,266	0.818	0.557	0.622
Pruned	0.6	7,651	179,614	95,266	0.812	0.554	0.630

Detection performance remains within a narrow margin of the unpruned baseline across a broad range of thresholds. In particular,  $\tau_{BCS} = 0.5$  discards 9,442 labels (about 5.0%) while maintaining comparable precision (0.818 vs. 0.815), with only modest reductions in recall and mAP50. Overall, these trends mirror the nuScenes results, indicating that redundant multisource supervision can be reduced substantially with minimal accuracy cost under a controlled, fixed evaluation protocol.

Compared to nuScenes, AV2 introduces additional im-

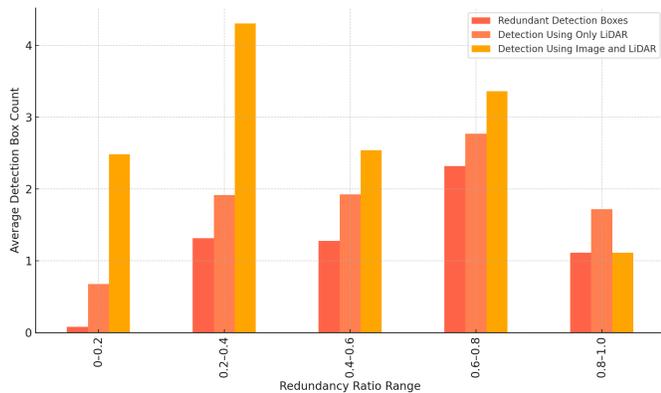


Fig. 9: Redundancy in LiDAR and image data in nuScenes

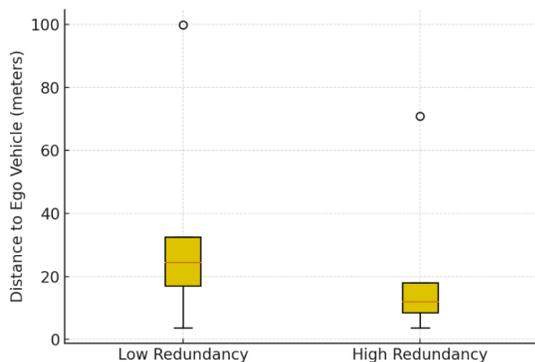


Fig. 10: T-test of distance threshold and cross-modal redundancy in nuScenes

plementation complexity due to its log-based organization, calibration-driven overlap estimation, and 3D-first annotation pipeline. Despite these differences, the consistent performance–redundancy behavior suggests that multisource redundancy is an intrinsic property of multi-camera supervision rather than a dataset-specific artifact.

#### D. Impact of Redundancy in Multimodal Data

For multimodal redundancy evaluation, Fig. 9 shows the distribution of redundancy across the Lidar and image detection in nuScenes. We further group and compare the high-redundancy and low-redundancy samples. T-test results of distance threshold  $T_{\text{dist}}$  and cross-modal redundancy ( $p\text{-value}=1.17e\text{-}76$ ) in Fig. 10 suggest that the objects with high cross-modal redundancy remain very close to the ego-vehicle. This supports the removal of close-range LiDAR data as redundancy. Fig. 11 shows that the detections can be almost unaffected by choosing a reasonable threshold and removing near-range LiDAR points. In conclusion, for multimodal data in object detection, closer-range LiDAR data is more redundant. While removing them could have little impact, the efficiency can be improved due to the decrease in data points to be processed.

## VI. DISCUSSION

We demonstrate that redundancy exists in (1) multisource and single modality data, which is the multi-view camera

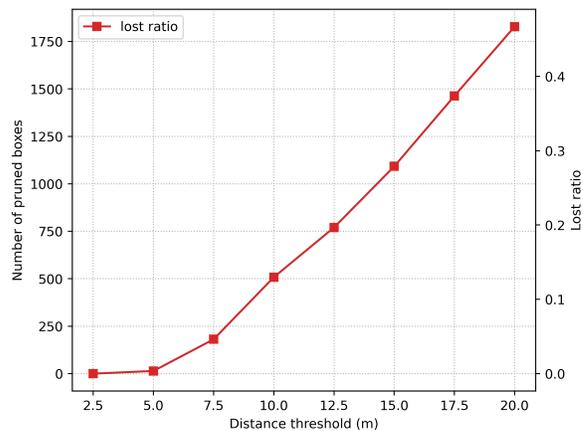


Fig. 11: The effects of removing close-range redundant LiDAR data in nuScenes on detection performance

images; and (2) multisource and multimodal data, which is the image and LiDAR data, using qualitative illustrations and quantitative statistics in three experiments.

#### A. Redundancy in Multisource Data

This existing multisource redundancy in two datasets is not surprising, as the autonomous driving system requires a certain level of redundancy to cross validate and ensure safety [24], [25], [28], [29], [51]. From a data quality perspective, however, simply accumulating more overlapping views of the same physical object does not necessarily provide more informative supervision. Instead, it can introduce low-informative redundancy that increases dataset size and computational cost without improving model behavior. Accordingly, our multisource experiments focus on (1) modeling and measuring redundancy at the instance level for object detection, and (2) removing low-informative yet computation-consuming redundancy while preserving complementary observations that contribute to performance.

From the nuScenes multisource experiment (Fig. 8), introducing the Bounding Box Completeness Score (BCS) effectively reduces redundancy by preferentially removing lower-quality redundant boxes while preserving complementary views when overlapping observations are of similar quality. In several camera pairs, models trained on a less redundant subset achieve comparable, and in some cases slightly higher, detection performance than models trained on the full unpruned dataset. This observation highlights that redundant multisource supervision may provide diminishing returns once at least one high-quality view of an object is available, reinforcing the importance of data quality over sheer quantity.

The Argoverse 2 (AV2) results further support this conclusion. Across a sweep of pruning thresholds  $\tau_{\text{BCS}} \in [0.1, 0.6]$  (Table IV), the pruning procedure removes between 7,651 and 16,056 camera-level labels out of 187,265 candidates (approximately 4.1%–8.6%), while the number of unique 3D object tracks observed during training remains constant at

95,266. This indicates that multisource redundancy in AV2 primarily arises from overlapping camera projections of the same physical objects, rather than from the presence of distinct instances.

Despite this substantial reduction in supervision, detection performance remains close to the unpruned baseline across the entire sweep. For example, at  $\tau_{BCS} = 0.5$ , the model discards 9,442 labels (about 5.0%) while maintaining comparable precision (0.818 vs. 0.815) and only modest reductions in recall and mAP50. These trends closely mirror the nuScenes results and demonstrate that instance-level redundancy pruning yields a more compact training signal without substantially degrading downstream detection performance.

Finally, BCS-guided pruning preserves the relative representation of instances across overlapping camera pairs. The proportion of instances contributed by key camera overlaps remains similar before and after pruning, suggesting that redundancy is reduced without introducing strong viewpoint imbalance. Overall, these findings suggest that multisource redundancy should be treated as a controllable data quality dimension: by explicitly managing redundant supervision at the instance level, it is possible to improve dataset efficiency while maintaining robust perception performance.

### B. Redundancy in Multimodal Data

The evaluation of image-LiDAR redundancy reveals that high redundancy ratios tend to occur for objects located close to the ego-vehicle, where dense LiDAR returns. This suggests that these instances are being detected equally well by both modalities and therefore contribute little new signal when both are included. This finding calls back to the multisource redundancy that selectively removing redundancy could maintain the performances, while offering other benefits, such as reducing the training dataset size.

The BCS in the first experiment also guides us in looking into the relationship between the distance to the ego vehicle and the 3D bounding box size. Zhang et al. find that the distance captured in the point cloud is a crucial factor for achieving a high point cloud quality [67]. Large, well-structured objects are easy for sensors to detect. Ideally, the further the object is located, the smaller the 3D bounding box size is. However, the bounding box quality remains a concern. From qualitative checks, the boxes could not accurately present the object size, despite correctly providing the centroid of the location. This was also observed in Kim et al.’s study on label quality [7]. They pointed out that many boxes are labeled as either too small or too big, resulting in low-quality datasets.

## VII. CONCLUSION AND FUTURE WORK

In this work, we first put forward a research gap: while AVs collect increasingly large volumes of  $M^2$  data, redundancy as a crucial DQ dimension has been overlooked. To address this, we model and measure redundancy across  $M^2$  data using nuScenes and Argoverse 2 datasets in object detection tasks using YOLOv8.

The results prove that selectively removing redundant multisource image data leads to measurable improvements in detection performance. In nuScenes, removing redundant labels preserves accuracy and frequently improves mAP50, with representative gains including Pair 1 increasing from 0.66 to 0.70, Pair 2 from 0.64 to 0.67, and Pair 3 from 0.53 to 0.55, while other pairs (4, 5, 6) match their respective baselines even under more aggressive pruning. In AV2, approximately 4.1–8.6% labels are removed, and detection performance remains near the baseline 0.64 mAP50, with the best redundancy–accuracy trade-off achieved at  $\tau_{BCS}=0.6$ .

Additionally, our analysis of multimodal data in image and LiDAR point clouds highlights existing redundancy issues that traditional pipelines do not consider, revealing opportunities to make perception systems more efficient without compromising reliability. Finally, further exploration of redundancy in AVs will focus on the following perspectives:

- 1) This paper has initially explored data redundancy in AD using the nuScenes and AV2 datasets, demonstrating that partially removing redundant data could improve model performance. We will extend this evaluation across diverse, unlabeled, and larger-scale datasets, suitable and SOTA models.
- 2) We aim to thoroughly evaluate the impacts of redundancy under varied conditions. For example, we will investigate how redundancy patterns shift across driving environments (e.g., dense urban vs. highway) and other factors such as lighting, weather, and sensor visibility.
- 3) Additional modalities, such as RADAR or user-generated data, introduce complexity, potentially revealing more redundancy. Integrating these modalities could further refine how to model and measure redundancy, as well as other important data quality dimensions.
- 4) Different AV tasks prioritize data quality differently. Other tasks, such as prediction and planning, may emphasize different metrics over redundancy, which is critical for safe decision-making in end-to-end autonomous driving systems.

### ACKNOWLEDGMENT

The authors would like to thank Dr. Junhua Ding from the Department of Data Science, Drs. Heng Fan and Yunhe Feng from the Department of Computer Science and Engineering at the University of North Texas for their precious feedback and suggestions. This work was supported in part by the National Science Foundation (NSF) under Award 2505686.

### REFERENCES

- [1] A. Guo, J. Huang, C. Lv, L. Chen, F.-Y. Wang *et al.*, “Advances in autonomous vehicle testing: The state of the art and future outlook on driving datasets simulators and proving grounds,” *Authorea Preprints*, 2024.
- [2] H. Liu, Z. Cao, X. Yan, S. Feng, and Q. Lu, “Autonomous vehicles: A critical review (2004-2024) and a vision for the future,” *TechRxiv* [174857767.78237989/v1](https://arxiv.org/abs/174857767.78237989/v1), 2025.
- [3] Y. Han, H. Zhang, H. Li, Y. Jin, C. Lang, and Y. Li, “Collaborative perception in autonomous driving: Methods, datasets, and challenges,” *IEEE Intelligent Transportation Systems Magazine*, vol. 15, no. 6, pp. 131–151, 2023.

- [4] L. Chen, P. Wu, K. Chitta, B. Jaeger, A. Geiger, and H. Li, “End-to-end autonomous driving: Challenges and frontiers,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 12, pp. 10164–10183, 2024.
- [5] H. Li, Y. Li, H. Wang, J. Zeng, H. Xu, P. Cai, L. Chen, J. Yan, F. Xu, L. Xiong *et al.*, “Open-sourced data ecosystem in autonomous driving: the present and future,” *arXiv preprint arXiv:2312.03408*, 2023.
- [6] M. Liang, J.-C. Su, S. Schuler, S. Garg, S. Zhao, Y. Wu, and M. Chandraker, “Aide: An automatic data engine for object detection in autonomous driving,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 14695–14706.
- [7] K. Kim, W. Kakani, and H. Kim, “Automatic pruning and quality assurance of object detection datasets for autonomous driving,” *Preprints*, 2025.
- [8] L. Li, W. Shao, W. Dong, Y. Tian, Q. Zhang, K. Yang, and W. Zhang, “Data-centric evolution in autonomous driving: A comprehensive survey of big data system, data mining, and closed-loop technologies,” *arXiv preprint arXiv:2401.12888*, 2024.
- [9] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. 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*, 2020, pp. 2446–2454.
- [10] N. Li, F. Song, Y. Zhang, P. Liang, and E. Cheng, “Traffic context aware data augmentation for rare object detection in autonomous driving,” in *2022 international conference on robotics and automation (ICRA)*. IEEE, 2022, pp. 4548–4554.
- [11] A. Goknil, P. Nguyen, S. Sen, D. Politaki, H. Niavis, K. J. Pedersen, A. Suyuthi, A. Anand, and A. Ziegenbein, “A systematic review of data quality in cps and iot for industry 4.0,” *ACM Computing Surveys*, vol. 55, no. 14s, pp. 1–38, 2023.
- [12] M. Liu, E. Yurtsever, J. Fossaert, X. Zhou, W. Zimmer, Y. Cui, B. L. Zagar, and A. C. Knoll, “A survey on autonomous driving datasets: Statistics, annotation quality, and a future outlook,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [13] 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, pp. 11621–11631.
- [14] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *The international journal of robotics research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [15] M. Alibeigi, W. Ljungbergh, A. Tonderski, G. Hess, A. Lilja, C. Lindström, D. Motorniuk, J. Fu, J. Widahl, and C. Petersson, “Zenseact open dataset: A large-scale and diverse multimodal dataset for autonomous driving,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 20178–20188.
- [16] Y. Zhou, H. Chen, and K. Sha, “A novel multi-layer task-centric and data quality framework for autonomous driving,” *IEEE Internet Computing*, 2026.
- [17] Q. Zhang, Y. Wei, Z. Han, H. Fu, X. Peng, C. Deng, Q. Hu, C. Xu, J. Wen, D. Hu, and C. Zhang, “Multimodal fusion on low-quality data: A comprehensive survey,” *arXiv preprint arXiv:2404.18947*, 2024.
- [18] Z. Wang, S. Fan, X. Huo, T. Xu, Y. Wang, J. Liu, Y. Chen, and Y.-Q. Zhang, “Vimi: Vehicle-infrastructure multi-view intermediate fusion for camera-based 3d object detection,” *arXiv preprint arXiv:2303.10975*, 2023.
- [19] X. Sun, S. Wang, M. Wang, S. S. Cheng, and M. Liu, “An advanced lidar point cloud sequence coding scheme for autonomous driving,” in *Proceedings of the 28th ACM International Conference on Multimedia*, 2020, pp. 2793–2801.
- [20] T. Zhao, X. Ning, K. Hong, Z. Qiu, P. Lu, Y. Zhao, L. Zhang, L. Zhou, G. Dai, H. Yang, and Y. Wang, “Ada3d: Exploiting the spatial redundancy with adaptive inference for efficient 3d object detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2023, pp. 17728–17738.
- [21] F. Li, C. Fu, D. Sun, H. Marzbani, and M. Hu, “Reducing redundancy in maps without lowering accuracy: A geometric feature fusion approach for simultaneous localization and mapping,” *ISPRS International Journal of Geo-Information*, vol. 12, no. 6, p. 235, 2023.
- [22] X. Sun, M. Wang, J. Du, Y. Sun, S. S. Cheng, and W. Xie, “A task-driven scene-aware lidar point cloud coding framework for autonomous vehicles,” *IEEE Transactions on Industrial Informatics*, vol. 19, no. 8, pp. 8731–8742, 2023.
- [23] K. Xu, Q. Feng, H. Chen, Z. Wang, X. Geng, X. Yang, M. Wu, X. Li, and W. Lin, “Dm3d: Distortion-minimized weight pruning for lossless 3d object detection,” *arXiv preprint arXiv:2407.02098*, 2024.
- [24] Q. Liu, W. Zhou, Y. Zhang, and X. Fei, “Multi-target detection based on multi-sensor redundancy and dynamic weight distribution for driverless cars,” in *2021 International Conference on Communications, Information System and Computer Engineering (CISCE)*, 2021, pp. 229–234.
- [25] R. Qian, X. Lai, and X. Li, “3d object detection for autonomous driving: A survey,” *Pattern Recognition*, vol. 130, p. 108796, 2022.
- [26] Y. Zhou, P. Sun, Y. Zhang, D. Anguelov, J. Gao, T. Ouyang, J. Guo, J. Ngiam, and V. Vasudevan, “End-to-end multi-view fusion for 3d object detection in lidar point clouds,” in *Conference on Robot Learning*. PMLR, 2020, pp. 923–932.
- [27] M. Munsif, W. Ahmad, A. Ali, M. Ullah, A. Hussain, and S. W. Baik, “Multi camera connected vision system with multi view analytics: A comprehensive survey,” *arXiv preprint arXiv:2510.09731*, 2025.
- [28] J. Zhang, “Sensor data validation and driving safety in autonomous driving systems,” *arXiv preprint arXiv:2203.16130*, 2022.
- [29] X. Song, P. Wang, D. Zhou, R. Zhu, C. Guan, Y. Dai, H. Su, H. Li, and R. Yang, “Apollocar3d: A large 3d car instance understanding benchmark for autonomous driving,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5452–5462.
- [30] H. Su, H. Gao, X. Wang, X. Fang, Q. Liu, G. Huang, X. Li, and Q. Cao, “Object detection in adverse weather for autonomous vehicles based on sensor fusion and incremental learning,” *IEEE Transactions on Instrumentation and Measurement*, 2024.
- [31] W. Hou, W. Li, and P. Li, “Fault diagnosis of the autonomous driving perception system based on information fusion,” *Sensors*, vol. 23, no. 11, p. 5110, 2023.
- [32] X. Zhang, B. Story, and D. Rajan, “Night time vehicle detection and tracking by fusing vehicle parts from multiple cameras,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 7, pp. 8136–8156, 2022.
- [33] J. Sanchez, J.-E. Deschaud, and F. Goulette, “Cola: Coarse-label multi-source lidar semantic segmentation for autonomous driving,” *IEEE Transactions on Robotics*, 2025.
- [34] J. Fu, C. Gao, Z. Wang, L. Yang, X. Wang, B. Mu, and S. Liu, “Eliminating cross-modal conflicts in bev space for lidar-camera 3d object detection,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16381–16387.
- [35] K. Wang, T. Zhou, X. Li, and F. Ren, “Performance and challenges of 3d object detection methods in complex scenes for autonomous driving,” *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 2, pp. 1699–1716, 2023.
- [36] Y. Wang, Q. Mao, H. Zhu, J. Deng, Y. Zhang, J. Ji, H. Li, and Y. Zhang, “Multi-modal 3d object detection in autonomous driving: a survey,” *International Journal of Computer Vision*, vol. 131, no. 8, pp. 2122–2152, 2023.
- [37] Z. Yang, N. Song, W. Li, X. Zhu, L. Zhang, and P. H. Torr, “Deepinteraction++: Multi-modality interaction for autonomous driving,” *arXiv preprint arXiv:2408.05075*, 2024.
- [38] H. Chen, J. Chen, and J. Ding, “Data evaluation and enhancement for quality improvement of machine learning,” *IEEE Transactions on Reliability*, vol. 70, no. 2, pp. 831–847, 2021.
- [39] N. Sambasivan, S. Kapania, H. Highfill, D. Akrong, P. Paritosh, and L. M. Aroyo, ““everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai,” in *proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–15.
- [40] R. S. Geiger, K. Yu, Y. Yang, M. Dai, J. Qiu, R. Tang, and J. Huang, “Garbage in, garbage out? do machine learning application papers in social computing report where human-labeled training data comes from?” in *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery, 2020, p. 325–336.
- [41] S. Rangineni, “An analysis of data quality requirements for machine learning development pipelines frameworks,” *International Journal of Computer Trends and Technology*, vol. 71, no. 9, pp. 16–27, 2023.
- [42] M. Priestley, F. O’donnell, and E. Simperl, “A survey of data quality requirements that matter in ml development pipelines,” *J. Data and Information Quality*, vol. 15, no. 2, Jun. 2023.

- [43] D. Schwabe, K. Becker, M. Seyferth, A. Klaub, and T. Schaeffter, "The metric-framework for assessing data quality for trustworthy ai in medicine: a systematic review," *NPJ Digital Medicine*, vol. 7, no. 1, p. 203, 2024.
- [44] M. Rahal, B. S. Ahmed, G. Szabados, T. Fornstedt, and J. Samuelson, "Enhancing machine learning performance through intelligent data quality assessment: An unsupervised data-centric framework," *Heliyon*, 2025.
- [45] H. Chen, R. Li, J. Ding, and A. Cleveland, "Enhancing data quality in medical concept normalization through large language models," Available at SSRN 4979696, 2024.
- [46] T. Emmanuel, T. Maupong, D. Mpoeleng, T. Semong, B. Mphago, and O. Tabona, "A survey on missing data in machine learning," *Journal of Big data*, vol. 8, pp. 1–37, 2021.
- [47] S. Li, Y. Fan, Y. Ma, and Y. Pan, "Evaluation of dataset distribution and label quality for autonomous driving system," in *2021 IEEE 21st International Conference on Software Quality, Reliability and Security Companion (QRS-C)*, 2021, pp. 196–200.
- [48] L. Kong, X. Xu, J. Ren, W. Zhang, L. Pan, K. Chen, W. T. Ooi, and Z. Liu, "Multi-modal data-efficient 3d scene understanding for autonomous driving," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [49] J.-A. Chen, W. Niu, B. Ren, Y. Wang, and X. Shen, "Survey: Exploiting data redundancy for optimization of deep learning," *ACM Computing Surveys*, vol. 55, no. 10, pp. 1–38, 2023.
- [50] S. Guha, F. A. Khan, J. Stoyanovich, and S. Scheluter, "Automated data cleaning can hurt fairness in machine learning-based decision making," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 12, pp. 7368–7379, 2024.
- [51] T. He, L. Zhang, F. Kong, and A. Salekin, "Exploring inherent sensor redundancy for automotive anomaly detection," in *2020 57th ACM/IEEE Design Automation Conference (DAC)*, 2020, pp. 1–6.
- [52] N. P. A. Duong, A. Almin, L. Lemarié, and B. R. Kiran, "Active learning with data augmentation under small vs large dataset regimes for semantic-kitti dataset," in *International Joint Conference on Computer Vision, Imaging and Computer Graphics*. Springer, 2022, pp. 268–280.
- [53] J. Ju, H. Jung, and J. Kim, "Extending contrastive learning to unsupervised redundancy identification," *Applied Sciences*, vol. 12, no. 4, p. 2201, 2022.
- [54] Y. Lu, B. Jiang, N. Liu, Y. Li, J. Chen, Y. Zhang, and Z. Wan, "Crossprune: Cooperative pruning for camera–lidar fused perception models of autonomous driving," *Knowledge-Based Systems*, vol. 289, p. 111522, 2024.
- [55] X. Li, Y. Zhang, Y. Shi, H. Zhu, J. Hu, and L. Peng, "Quality assessment of image dataset for autonomous driving," in *2023 IEEE International Conference on Imaging Systems and Techniques (IST)*. IEEE, 2023, pp. 1–6.
- [56] R. Rao, S. Rao, E. Nouri, D. Dey, A. Celikyilmaz, and B. Dolan, "Quality and relevance metrics for selection of multimodal pretraining data," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 956–957.
- [57] H. Hadizadeh, S. F. Yeganli, B. Rashidi, and I. V. Bajić, "Mutual information analysis in multimodal learning systems," in *2024 IEEE 7th International Conference on Multimedia Information Processing and Retrieval (MIPR)*. IEEE, 2024, pp. 390–395.
- [58] J. Terven, D.-M. Córdova-Esparza, and J.-A. Romero-González, "A comprehensive review of yolo architectures in computer vision: From yolov1 to yolov8 and yolo-nas," *Machine learning and knowledge extraction*, vol. 5, no. 4, pp. 1680–1716, 2023.
- [59] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4490–4499.
- [60] 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.
- [61] Y. Li, H. Mao, R. Girshick, and K. He, "Exploring plain vision transformer backbones for object detection," in *European conference on computer vision*. Springer, 2022, pp. 280–296.
- [62] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. L. Rus, and S. Han, "Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation," in *2023 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2023, pp. 2774–2781.
- [63] A. Song, "Nuscenes dataset in kitti format," Jan 2025. [Online]. Available: <https://www.kaggle.com/datasets/alohasong/nuscenes-dataset-in-kitti-format>
- [64] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes *et al.*, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," *arXiv preprint arXiv:2301.00493*, 2023.
- [65] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [66] T. Kieffer, "Timkie/yolo-lidar-fusion: Lidar-camera fusion for 3d object detection in autonomous driving systems," Aug 2024. [Online]. Available: <https://github.com/TimKie/YOLO-LiDAR-Fusion>
- [67] C. Zhang and A. Eskandarian, "Image-guided outdoor lidar perception quality assessment for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, 2024.