

# Griffin: Aerial-Ground Cooperative Detection and Tracking Dataset and Benchmark

Jiahao Wang<sup>1</sup>, Xiangyu Cao<sup>2</sup>, Jiaru Zhong<sup>3</sup>, Yuner Zhang<sup>4</sup>, Haibao Yu<sup>5</sup>, Lei He<sup>1</sup> and Shaobing Xu<sup>1\*</sup>

**Abstract**—Despite significant advancements, autonomous driving systems continue to struggle with occluded objects and long-range detection due to the inherent limitations of single-perspective sensing. Aerial-ground cooperation offers a promising solution by integrating UAVs’ aerial views with ground vehicles’ local observations. However, progress in this emerging field has been hindered by the absence of public datasets and standardized evaluation benchmarks. To address this gap, this paper presents a comprehensive solution for aerial-ground cooperative 3D perception through three key contributions: (1) *Griffin*, a large-scale multi-modal dataset featuring over 200 dynamic scenes (30k+ frames) with varied UAV altitudes (20-60m), diverse weather conditions, and occlusion-aware 3D annotations, enhanced by CARLA-AirSim co-simulation for realistic UAV dynamics; (2) A unified benchmarking framework for aerial-ground cooperative detection and tracking tasks, including protocols for evaluating communication efficiency, latency tolerance, and altitude adaptability; (3) AGILE, an instance-level intermediate fusion baseline that dynamically aligns cross-view features through query-based interaction, achieving an advantageous balance between communication overhead and perception accuracy. Extensive experiments prove the effectiveness of aerial-ground cooperative perception and demonstrate the direction of further research. The dataset and codes are available at <https://github.com/wang-jh18-SVM/Griffin>.

## I. INTRODUCTION

While significant progress has been made in autonomous driving technologies, current single-perspective systems still struggle with fundamental challenges of severe occlusions and limited field-of-view in complex environments. To address these limitations, an increasing number of cooperative perception strategies have emerged, including vehicle-to-vehicle (V2V) [1], [2] and vehicle-to-infrastructure (V2I) [3], [4] cooperation. These approaches have demonstrated feasible solutions and significant enhancements in perception capabilities. Nevertheless, their practical implementation often requires substantial infrastructure investment and widespread adoption of connected vehicles, which may

This work was supported by National Natural Science Foundation of China, Science Fund for Creative Research Groups (Grant No. 52221005) and National Natural Science Foundation of China (Grant No. 52131201).

\*Corresponding author.

<sup>1</sup>School of Vehicle and Mobility, Tsinghua University, Beijing, China. [wjh22@mails.tsinghua.edu.cn](mailto:wjh22@mails.tsinghua.edu.cn), [shaobxu@tsinghua.edu.cn](mailto:shaobxu@tsinghua.edu.cn)

<sup>2</sup>Xingjian College, Tsinghua University, Beijing, China.

<sup>3</sup>School of Mechanical Engineering, Beijing Institute of Technology, Beijing, China.

<sup>4</sup>Faculty of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, USA.

<sup>5</sup>Department of Computer Science, The University of Hong Kong, Hong Kong, China.

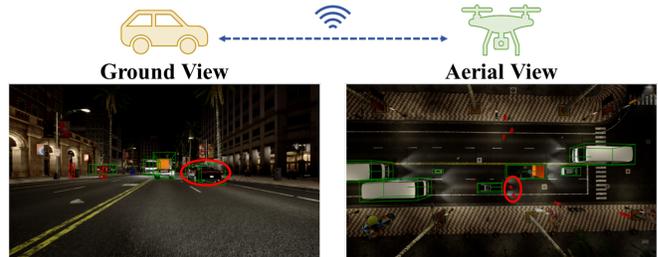


Fig. 1: Motivation for aerial-ground cooperative perception. While pedestrians (circled in red) are occluded in the ground vehicle’s view, they can be clearly observed from the drone’s perspective. This approach enables improved perception in complex environments and offers a more economically viable and rapidly deployable alternative to V2V and V2I solutions.

present economic barriers. In contrast, vehicle-to-drone or so called, aerial-ground cooperative (AGC) systems leverage the aerial panoramic views with ground-level detailed observations. The integration of unmanned aerial vehicles (UAVs), or drones, provides unique advantages for rapid deployment in critical applications, including smart cities, emergency response, and security patrols, thereby offering a new paradigm for dynamic environment perception.

Despite the promising potential, developing effective aerial-ground cooperative perception systems still faces two critical challenges. The first challenge stems from dynamic perspective mismatch: unlike V2V and V2I collaborations, where sensors primarily move on a horizontal plane, drones introduce altitude variations and dynamic pitch/roll angle changes, complicating the alignment of cross-view features. Second, existing drone-view 3D perception datasets are still compromised. As shown in Table I, datasets like CoPerception-UAV [5] and UAV3D [6] lack occlusion-aware annotations, resulting in bounding boxes that include targets in invisible regions. Meanwhile, many datasets employ oversimplified fixed-angle [8], [9] or fixed-altitude [6], [7], [8], [9] camera models, which do not reflect the real-world drone dynamics influenced by factors such as wind disturbances and target acceleration.

Inspired by the Griffin, a mythical creature that unites the lion’s terrestrial strength and the eagle’s aerial dominance, we aim to harness the combined power of aerial and ground perspectives to overcome these challenges and enhance perception for autonomous driving. To this end, we present the following contributions for aerial-ground cooperative 3D perception:

TABLE I: A detailed comparison between representative cooperative perception datasets

Mode	Dataset Name	Year	Source	Cams (/Agent)	BBox	Tracking ID	Occ-Aware	Frames (k)	Altitude (m)
Veh-Veh	OPV2V [1]	2022	Joint Sim	Multiple	3D	✓	✓	11	–
	V2V4Real [2]	2023	Real	Multiple	3D	✓	–	310	–
Veh-Inf	DAIR-V2X [3]	2022	Real	Single	3D	×	–	22	20-25
	V2X-Seq [4]	2023	Real	Single	3D	✓	–	15	20-25
Air-Air	CoPerception-UAVs [5]	2022	Joint Sim	Multiple	3D	✓	×	5.2	40,60,80
	UAV3D [6]	2024	Joint Sim	Multiple	3D	✓	×	20	60
	AeroCollab3D [7]	2024	Joint Sim	Single	3D	✓	×	3.2	50
	Air-Co-Pred [8]	2024	Sim	Single	3D	✓	✓	8.0	50
Veh-Air	V2U-COO <sup>†</sup> [9]	2024	Sim	Single	3D	×	–	9.3	80(R)&70(L)
	CoPeD [10]	2024	Real	Single	2D	×	–	203	2-10
	<b>Griffin (Ours)</b>	<b>2025</b>	<b>Joint Sim</b>	<b>Multiple</b>	<b>3D</b>	<b>✓</b>	<b>✓</b>	<b>30</b>	<b>20-60</b>

Note:

- ‘Cams (/Agent)’ refers to the camera numbers on each agent. ‘Occ-Aware’ (Simulation-only) flags whether annotation visibility is considered. Altitude represents the set height of infrastructure sensors or the cruising altitude of UAVs in corresponding modes.
- In the Source column, ‘Joint Sim’ refers to the co-simulation of CARLA and AirSim / SUMO, while ‘Sim’ refers to using only CARLA.
- <sup>†</sup> Attributes are derived from the original paper, as the V2U-COO dataset is not publicly available. ‘80(R)&70(L)’ in the Altitude column represent the fixed altitudes of the right and left drones, respectively.

- **The Griffin Dataset:** We release *Griffin*, the first publicly available aerial-ground cooperative 3D perception dataset. It encompasses over 200 dynamic scenes (over 30K frames, 270K images) from CARLA-AirSim co-simulation, with instance-aware occlusion quantification, varying cruising altitudes, and realistic simulation of drone dynamics under various conditions.
- **Benchmark:** We present a benchmarking framework for evaluating aerial-ground cooperative 3D object detection and tracking. It includes implementations of classic baselines and provides a suite of metrics to evaluate accuracy, communication cost, and robustness under varying latency conditions.
- **AGILE Framework:** We propose AGILE, the first open-source Aerial-Ground Instance-LEvel intermediate fusion framework for joint detection and tracking. Our approach demonstrates the effectiveness of aerial-ground collaboration, maintaining perception accuracy while adhering to practical communication constraints.

The remainder of this paper is organized as follows: Section II reviews related work on cooperative perception methods and datasets. Section III details our dataset construction methods, followed by benchmark framework including evaluation tasks and metrics in Section IV and baseline methods in Section V. Experimental results and analysis are presented in Section VI, with conclusions in Section VII.

## II. RELATED WORK

### A. Cooperative Perception

Cooperative perception systems are generally categorized into three fusion strategies [11], [12], [13]: early, intermediate, and late fusion. Early fusion [14], [15] directly integrates raw sensor data, which retains rich semantic information but incurs significant bandwidth overhead. In contrast, late fusion [3], [15] minimizes communication costs by exchanging detection results; however, its performance degrades under occlusion or ambiguous observations. Intermediate fusion

offers a compromise by transmitting certain network features. Recent works [16], [17], [18] have highlighted the advantages of transmitting Bird’s Eye View (BEV) features derived from point clouds or images to enhance 3D object detection, while studies such as [5], [19], [20] explore the transmission of instance-level features or queries to reduce communication overhead and alleviate bandwidth constraints.

Despite recent advancements, cooperative perception methods designed explicitly for aerial-ground collaboration remain underexplored, primarily due to the lack of publicly available datasets tailored for such scenarios. Minaeian et al. [21] were among the first to investigate this area, proposing a vision-based solution for target detection and localization involving multiple unmanned aerial/ground agents. More recently, UVCPNet [9] introduced a comprehensive framework for aerial-ground collaboration. This framework demonstrated the effectiveness of aerial-ground cooperative perception with substantial statistical evidence. However, the dataset used in the study is not publicly available, which limits the broader research community’s ability to build upon existing work.

### B. Multi-agent Cooperative Perception Datasets

The development of datasets exploring diverse communication modes has significantly advanced cooperative perception algorithms. Table I demonstrates a comparison between recent datasets across different cooperation modes. For V2V scenarios, OPV2V [1] provides comprehensive simulated data with precise annotations, while V2V4Real [2] bridges the gap to real-world scenarios using synchronized LiDAR and camera streams. For V2I collaboration, datasets like V2X-Sim [22], DAIR-V2X [3] and V2X-Seq [4] offer both simulated and real-world testbed. Roadside-focused datasets like RCooper [23] further address occlusion challenges through infrastructure sensors.

In contrast, the cooperative perception datasets featuring aerial perspectives from UAVs remain limited. Recent multi-UAV datasets, such as UAV3D [6], AeroCollab3D [7], and

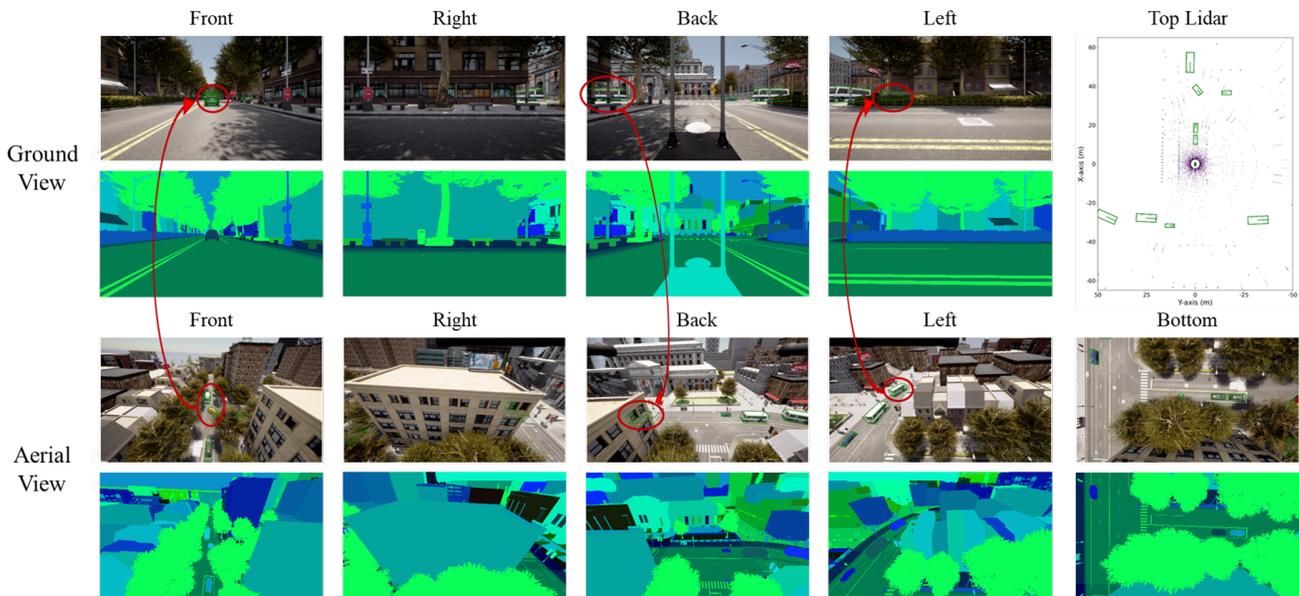


Fig. 2: An example from *Griffin* with visualized annotation. The ground vehicle platform is equipped with four cameras and one LiDAR, while the aerial drone platform has five cameras. Both platforms include instance segmentation ground truth sensors, shown in the lower row. Bounding boxes represent annotations from cooperative perspectives, indicating that one agent may be able to ‘see’ occluded objects from the other’s view, same as the case in Fig. 1. We use red circles and arrows to highlight those cases.

Air-Co-Pred [8] primarily focus on Air-Air cooperation but suffer from fixed altitude constraints. For datasets designed for AGC scenarios, while some, such as GRACO [24], are tailored for SLAM tasks, very few focus on detection or tracking objectives. Datasets like V2U-COO [9] pioneer vehicle-UAV cooperative perception but rely on predefined UAV poses without realistic motion dynamics. CoPeD [10], though large-scale, targets low-altitude robot scenarios and directly uses unrefined automatic annotations generated by foundation models.

Notably, most simulation datasets neglect frame-level object visibility by including occluded objects, leading to perception inaccuracies. This issue arises because the native interface of CARLA delivers data of all objects in the scene regardless of visual accessibility. To address all these gaps mentioned, our AGC dataset *Griffin* introduces co-simulation for realistic multi-agent dynamics, occlusion-aware 3D annotations with tracking IDs, and diverse altitude settings, offering a comprehensive platform for advancing aerial-ground collaborative perception.

### III. DATA SETUP

#### A. Data Collection

Our aerial-ground cooperative perception data collection framework adopts a modular architecture, as illustrated in Fig. 3. The framework comprises two primary components: (1) a server based on Unreal Engine 4 (UE4), which integrates the CARLA and AirSim platforms, and (2) a Python client featuring four specialized managers. Specifically, the Traffic Manager handles vehicle planning and control, the Aerial Manager generates drone trajectories, the Scene Manager configures random environments, and the Sensor Man-

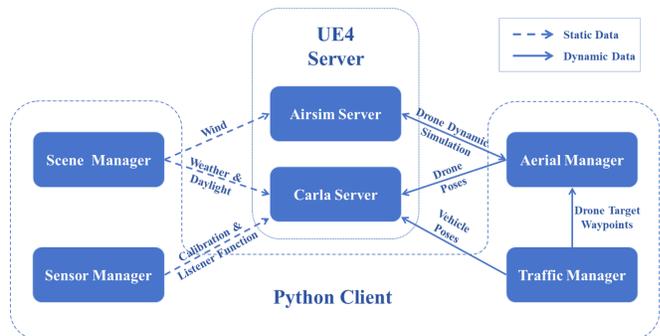


Fig. 3: Data collection framework. Dashed lines represent static data flows (used for scene initialization), while solid lines denote dynamic data flows (transferred per frame).

ager manages the acquisition and processing of multi-modal sensor data.

**Sensor Setting.** Our cross-platform perception system is carefully designed to balance perception performance and platform constraints, as depicted in Fig. 4. The ground platform is equipped with a comprehensive multi-modal suite consisting of four wide field-of-view (FoV) RGB cameras (108.8°, 1920×1080 resolution) positioned in cardinal directions to provide overlapping coverage. This setup is enhanced by an 80-beam LiDAR operating at 10Hz with a vertical FoV ranging from -25° to 15°, allowing dense 3D sensing.

For aerial operations, which must comply with size, weight, and power (SWaP) constraints, we exclude LiDAR payloads and instead employ a vision-centric configuration. This setup includes five downward-oriented cameras that match the sensor specifications of the ground platform.

**Scene Setting.** We select four representative maps from the



(a) Ground Vehicle Platform (b) Aerial Drone Platform

Fig. 4: Integrated Platform Designs

CARLA simulator, including two urban settings (Town03 and Town10HD) and two suburban settings (Town06 and Town07). Data is collected under diverse weather conditions, actor densities, and vehicle speeds. Furthermore, multiple collaboration modes are designed by selecting desired horizontal and vertical distances.

Based on variations in UAV cruising altitude, we divided the dataset into three distinct categories. The *Griffin-Random* set covers the widest altitude range, spanning from 20 to 60 meters above the vehicle. In contrast, *Griffin-25m* and *Griffin-40m* focus on specific altitudes, targeting  $25 \pm 2$  meters and  $40 \pm 2$  meters respectively.

### B. Data Post-processing

**Spatio-temporal Alignment.** There are four types of coordinate systems in *Griffin*, including the world coordinate, ego coordinate, sensor coordinate, and simulator coordinate, as detailed in Table II.

TABLE II: Different Coordinate Systems.

Name	Category	Type	Origin
World	Geodetic	ENU (R)	Fixed reference point
Ego	Drone Vehicle	FLU (R) FLU(R)	Drone center Vehicle center
Sensor	Camera Image LiDAR	RDF (R) RD (2D) FLU (R)	Camera optical center Top-left of the image LiDAR center
Simulator	CARLA AirSim	ESU (L) NED (R)	Fixed reference point Fixed reference point

Note:

- Axis direction: ENU (East-North-Up), FLU (Forward-Left-Up), RDF (Right-Down-Forward), RD (Right-Down), ESU (East-South-Up), NED (North-East-Down).
- Handedness: R (Right-handed), L (Left-handed)

Our spatial alignment converts simulator-native 3D annotations into unified right-handed coordinates, with dual output formats supporting both KITTI [25] (ego centric) and NuScenes [26] (global reference) benchmarks.

To address the time synchronization issue, we utilize CARLA’s synchronous mode during data recording to ensure consistency across data captured under the same timestamp. Additionally, during network training, we provide a code interface to simulate specified time delays and thus evaluate the algorithms’ robustness to communication latency.

**Annotation.** *Griffin* provides high-quality 3D annotations for each frame, covering six object categories: pedestrian, car, truck, bus, motorcycle, and bicycle. The annotation for each object includes a category label, tracking ID, visibility rate, and a 9-dimensional bounding box model defined by x, y, z, length, width, height, roll, pitch, and yaw.

To address the common issue of occlusion judgment in existing datasets, we develop a visibility quantification method utilizing CARLA’s instance segmentation ground truth interface. During data collection, RGB and segmentation images are recorded with identical sensor configurations to maintain spatial-temporal alignment. In the post-processing phase, we sample points within each target’s bounding box and project them onto the segmentation ground truth. The visibility rates are then calculated by comparing semantic categories and instance IDs between sampled pixels and corresponding targets. Targets with low visibility are filtered out to ensure annotation precision.

### C. Statistic and Scene Analysis

The dataset comprises 205 scene clips (104 for *Griffin-Random*, 47 for *Griffin-25m*, and 54 for *Griffin-40m*), each lasting approximately 15 seconds and corresponding to around 150 frames of image data, totaling over 30,000 samples and nearly 275,000 images.

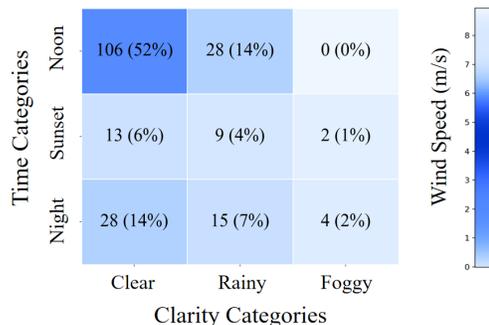


Fig. 5: **Weather Distribution.** Most weather conditions are clear at noon. Certain phenomena are either non-existent or extremely rare, such as fog at noon.

To ensure the dataset’s diversity and generalization capability, the scene design incorporates multi-dimensional environmental variables. In terms of weather conditions, as shown in Fig. 5, it includes combinations of different times of day (noon, sunset, and night), clarity levels (clear, rainy and foggy), and wind speeds ranging from negligible to substantial (0 ~ 9m/s).

Additionally, our dataset features multiple collaboration modes between the ground vehicle and UAV, resulting in diverse relative positioning patterns. Taking *Griffin-Random* as an example, the pose distribution of the drone is illustrated in Fig. 6. The horizontal distance is randomly set within 0 ~ 20 meters ahead. However, during turns or at intersections, the drone may drift to the left or right of the vehicle’s front, resulting in a symmetric lateral (y-axis) distribution and a reduced longitudinal (x-axis) distance, which aligns with expectations.

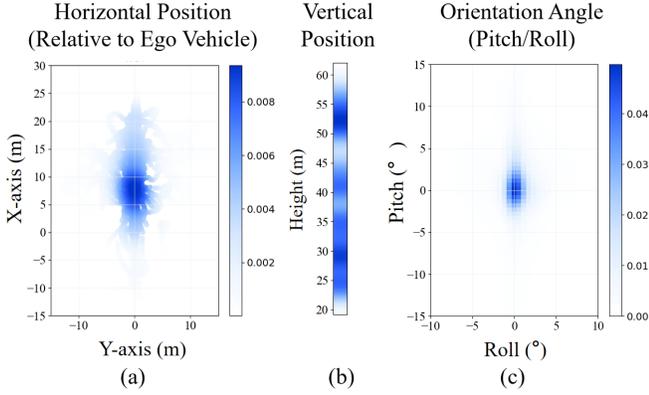


Fig. 6: UAV Pose Distribution Analysis of *Griffin-Random*. (a) Horizontal positions relative to the ego vehicle. (b) Vertical position distribution. (c) Orientation angles relative to the ground, in line with the dynamics of UAVs.

#### IV. TASK AND METRICS

Our dataset could further support multiple cooperative perception tasks, including detection, tracking, prediction, localization, etc. In this paper, we focus on two sequential visual tasks: Aerial-Ground Cooperative (AGC) 3D Object Detection and AGC 3D Object Tracking.

##### A. AGC 3D Object Detection

**Input and Ground Truth.** The input data for the AGC 3D Object Detection task includes sequential images from multiple agents and their relative pose information. Taking the ground vehicle node  $g$  and the aerial node  $a$  as examples:

- The sequential frames of  $g$  are denoted as  $\{C_g(t_g) \mid t_g \leq T_g\}$ , and the sequential frames of  $a$  are denoted as  $\{C_a(t_a) \mid t_a \leq T_a\}$ , where  $T_g$  represents the perception timestamp of the vehicle,  $T_a \leq T_g$  represents the capture timestamp of the aerial images, and  $C(\cdot)$  denotes the capture function.
- The relative pose of the drone with respect to the vehicle is denoted as  $M_{a \rightarrow g}$ .

The detection output comprises precisely localized objects within the surrounding traffic area, represented as 3D bounding boxes (with position, dimensions, and orientation) along with their semantic labels and confidence scores. The ground truth set  $GT$  represents all actual objects in the scene that should be detected, and is formally defined as:

$$GT = (GT_g \cup GT_a) \cap R \quad (1)$$

where  $GT_g$  and  $GT_a$  are the ground truth from the vehicle and the drone respectively, and  $R$  denotes the region of interest centered around the ego vehicle.

**Evaluation Metrics.** For comprehensive evaluation, we employ standard metrics in 3D object detection [26]: Average Precision (AP) to assess detection quality at various thresholds, alongside Average Translation Error (ATE), Average Scale Error (ASE), Average Orientation Error (AOE), and Average Velocity Error (AVE) to measure prediction accuracy of object position, size, orientation, and velocity, respectively. To evaluate communication costs for cooperative

methods, we additionally utilize Bytes per second (BPS) as a key metric. This suite of metrics provides a thorough assessment of detection performance.

##### B. AGC 3D Object Tracking

**Input and Ground Truth.** The AGC 3D Object Tracking task supports two prevalent paradigms. For the joint detection and tracking method, the input consists of raw sensor data frames from both agents, following the same format previously defined for the detection task. For the tracking-by-detection approach, the input is prediction results from the AGC detection stage, including the detected 3D bounding boxes and confidence scores. The formulation of the input data for the latter method remains similar to the detection task, except that the capture function  $C(\cdot)$  is replaced by the bounding box prediction function  $\hat{B}(\cdot)$ .

The outputs and corresponding ground truth sets  $GT$  of both paradigms include the target's category, 3D bounding box, and, critically, a unique tracking ID that persists throughout the tracking sequence.

**Evaluation Metrics.** To evaluate tracking performance, we utilize established metrics [26] in the 3D multi-object tracking domain, including: Average Multi-Object Tracking Accuracy (AMOTA) and Average Multi-Object Tracking Precision (AMOTP) to assess overall tracking quality across different detection confidence thresholds, and Mostly Tracked trajectories (MT), Mostly Lost trajectories (ML) and ID Switches (IDS) to quantify identity preservation. These metrics collectively characterize the key aspects of tracking baselines.

#### V. BASELINE FRAMEWORK

We implement a series of baseline methods for AGC 3D object detection and tracking, so as to establish performance references for subsequent research. As illustrated in Fig. 7, existing cooperative perception methods can generally be categorized into early fusion, intermediate fusion, and late fusion paradigms. This benchmark implements all three categories, including our proposed AGILE (Aerial-Ground cooperative Instance-LEvel fusion) method. Together, these methods form a comprehensive framework for aerial-ground cooperative perception.

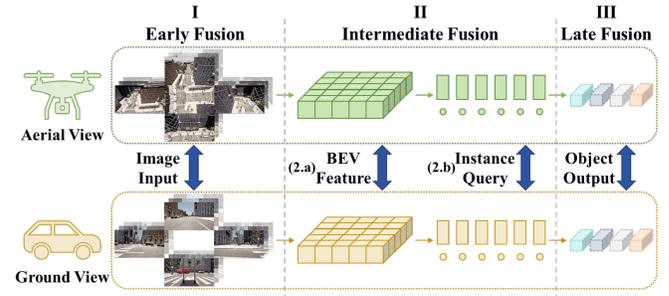


Fig. 7: Different Fusion Stages for cooperative perception

##### A. AGILE Framework

We first present AGILE, an instance-level intermediate fusion framework for joint detection and tracking tasks.

TABLE III: Detection and tracking performance under various latency conditions on the *Griffin-25m* dataset.

Fusion Stage	Method	Latency (ms)	Detection Metrics			Tracking Metrics					Comm. Cost
			AP $\uparrow$	ATE $\downarrow$	AOE $\downarrow$	AMOTA $\uparrow$	AMOTP $\downarrow$	MT $\uparrow$	ML $\downarrow$	IDS $\downarrow$	BPS $\downarrow$
No Fusion	-	0	0.366	0.398	0.492	0.363	1.295	16	49	5	0
Early Fusion	Concat	0	0.626	0.385	0.795	0.653	0.869	34	16	38	$\sim 3 \times 10^8$
		200	0.510	0.502	0.776	0.561	1.044	25	21	34	
		400	0.416	0.522	0.714	0.418	1.189	19	32	26	
Instance Fusion	AGILE	0	0.412	0.484	0.708	0.435	1.130	21	35	23	$\sim 6 \times 10^5$
		200	0.396	0.484	0.705	0.410	1.165	16	36	20	
		400	0.373	0.505	0.709	0.366	1.212	15	40	12	
Late Fusion	Hungarian	0	0.377	0.357	0.511	0.369	1.078	14	55	16	$\sim 1 \times 10^4$
		200	0.363	0.362	0.498	0.357	1.060	15	55	11	
		400	0.339	0.350	0.458	0.310	1.129	15	58	11	

As depicted in Fig. 7 (2.a), most existing intermediate fusion approaches [5], [27], [13], [28], [29] utilize Bird’s Eye View (BEV) features as the bridge for fusion. These methods spatially align BEV representations from different perspectives and merge features from overlapping regions. While effective for V2V or V2I cooperation, this paradigm faces significant challenges in aerial-ground scenarios. The substantial dynamic pitch and roll angles of UAVs cause the generated BEV feature plane to deviate from being parallel to the actual ground, compromising fusion accuracy.

Alternatively, drawing on prior research [30], [31], AGILE adopts the instance-level intermediate fusion paradigm, as shown in Fig. 7 (2.b). Instance-level object queries are generated from BEV features, with each query comprising a feature vector encoding the target’s semantic and geometric attributes, along with an explicit 3D reference point (typically the target’s center position). These queries dynamically focus on potential targets in the scene via spatial attention mechanisms, offering fine-grained representations for subsequent cross-view cooperation. Compared to dense BEV features, these sparse instance queries are more straightforward to align across perspectives and thus better suited for aerial-ground cooperative scenarios.

Specifically, AGILE employs BEVFormer [32] as the backbone network to generate independent BEV features from both the aerial and ground views. Following TrackFormer’s design [33], instance queries are then derived from the BEV features for each perspective. These queries undergo temporal propagation and updating across frames, corresponding to persistently tracked instances. New queries are initialized per frame for emerging targets, while queries for targets exiting the field of view are discarded. This paradigm enables feature-level tracking, simultaneously producing detection results and tracking IDs.

For the aerial view, we leverage flight altitude information to estimate the ground position within the drone’s coordinate system, guiding the initialization positions of instance reference points. These queries first interact with the UAV-maintained BEV features to update their features, reference points, and confidence scores. High-confidence queries are then passed to the ground platform.

In the ground vehicle view, transmitted queries are explicitly aligned by projecting their reference points to the

vehicle coordinate system. Additionally, we implement the implicit alignment approach from UniV2X [31], where query features concatenated with rotation matrices are processed through a three-layer multilayer perceptron network (MLP) for feature refinement. Cross-perspective query matching is then performed based on Euclidean distances between reference points and feature similarities. Matched query pairs undergo feature fusion via another three-layer MLP, while unmatched high-confidence queries are preserved. Together, these queries are used to generate detection boxes and tracking IDs, optimized end-to-end by computing the loss against ground truth trajectories.

### B. Early Fusion

The early fusion method integrates raw images from drones and vehicles at the data level. After transforming drone camera extrinsics to the vehicle coordinate system, all images are concatenated and fed into BEVFormer [32] to generate unified BEV features. The downstream architecture mirrors AGILE, offering it joint detection and tracking capability. Although communication costs limit its practical applicability, this method preserves maximal input information integrity. Its performance can be considered the upper bound of cooperative perception capabilities under identical testing conditions.

### C. Late Fusion

The late fusion method first applies BEVFormer [32] independently to generate detection boxes from both aerial and ground perspectives. These results are then fused via Hungarian matching based on Euclidean distances between bounding boxes. We further implement a tracking-by-detection baseline using AB3DMOT [34], employing Kalman filtering for motion prediction and the Hungarian algorithm for cross-frame association.

## VI. EXPERIMENTS

### A. Implementation Details

We evaluated our approach with *Griffin-25m*, *Griffin-40m*, and *Griffin-Random*, each divided into training and validation sets at an 8:2 ratio. The cooperative perception evaluation is conducted within a 102.4m  $\times$  102.4m area centered on the vehicle. For a fair comparison, all baseline methods adopt

the same backbone network and training/testing strategies. The target categories are merged into three classes: vehicles, pedestrians, and two-wheelers. The BEVFormer components in AGILE, early fusion, and late fusion frameworks are all implemented based on ResNet-50. All networks are trained using the AdamW optimizer with a learning rate of  $2e-4$  and a batch size of 8, distributed across 4 NVIDIA 3090 GPUs. Due to space limitations, only representative metrics are discussed in the paper, while the complete experiment results and corresponding model weights are available in our GitHub repository.

### B. Results and Analysis

Table III presents the detection and tracking performance for car class under various latency conditions on the *Griffin-25m* dataset. The comparison includes no fusion (ground view only), early fusion, late fusion baselines, and our proposed AGILE framework.

**Performance Improvement by Cooperation.** Cooperative methods significantly outperform the no-fusion baseline, validating the effectiveness of aerial-ground collaboration. Under zero latency conditions, all fusion methods exhibit distinct improvements compared to the no-fusion approach. The early fusion method achieves a 71% increase in AP for detection and a 79.9% boost in AMOTA for tracking, but its communication cost far exceeds other methods, limiting its practical applicability. AGILE provides a more tempered improvement, with a 12.6% gain in AP, a 19.8% increase in AMOTA, and much reduced communication overhead. Late Fusion, though offering only a 3.0% AP gain, lowers communication costs by three orders of magnitude. Overall, AGILE demonstrates a favorable trade-off between performance and communication cost in both detection and tracking, highlighting the advantages of instance-level fusion for temporal association.

**Robustness to Communication Latency.** With increasing communication latency, the performance of all fusion methods declines, but AGILE demonstrates the highest robustness to latency. Although early fusion still achieves the highest precision at each latency level, it suffers the most significant percentage drop, as -33.6% in AP and -36.0% in AMOTA at 400ms latency. AGILE and late fusion exhibit similar percentage decreases in performance; however, AGILE consistently outperforms late fusion across all latency conditions. These results indicate that AGILE is more robust to latency interference, making it well-suited for dynamic communication environments in real-world scenarios.

**Generalization to Flight Altitude Changes.** Table IV presents the detection and tracking performance of different methods across datasets, highlighting the impact of UAV altitude on cooperative perception. The results demonstrate that cooperative perception methods are sensitive to changes in UAV altitude, with performance varying significantly across different flight heights.

On the *Griffin-25m* dataset, thanks to the suitable target scales observed from the UAV perspective, cooperative

TABLE IV: **Detection and tracking performance across datasets with varying UAV altitudes.** The values in parentheses indicate the relative improvement compared to the No Fusion method. Comparing the performance improvement of the same method across different datasets reflects its robustness to UAV altitude variations.

Dataset	Fusion Stage	AP $\uparrow$	AMOTA $\uparrow$
Griffin-25m	No Fusion	0.366	0.363
	Early Fusion	0.626 (+71.0%)	0.653 (+79.9%)
	Instance Fusion	0.412 (+12.6%)	0.435 (+19.8%)
	Late Fusion	0.377 (+3.0%)	0.369 (+1.7%)
Griffin-40m	No Fusion	0.351	0.371
	Early Fusion	0.495 (+41.0%)	0.543 (+46.4%)
	Instance Fusion	0.359 (+2.3%)	0.397 (+7.0%)
	Late Fusion	0.354 (+0.9%)	0.387 (+4.3%)
Griffin-Random	No Fusion	0.465	0.491
	Early Fusion	0.580 (+24.7%)	0.646 (+31.6%)
	Instance Fusion	0.408 (-12.3%)	0.426 (-13.2%)
	Late Fusion	0.375 (-19.4%)	0.394 (-19.8%)

methods achieve their maximum improvements over the no-fusion baseline. On *Griffin-40m*, the increased altitude of 40m reduces the target scale and increases depth estimation difficulty, leading to a reduction in fusion gains. AGILE and the late fusion method show marginal gains of only 2.3% and 0.9% in AP, and 7.0% and 4.3% in AMOTA, respectively. This reduction suggests that fusion algorithms, particularly those operating at later stages, struggle to detect and track small ground targets from higher altitudes. Notably, on *Griffin-Random* with random UAV altitudes between 20m and 60m, both AGILE and late fusion underperform the no-fusion baseline, likely due to inconsistent target scales and disrupted cross-view alignment caused by the varying altitudes. These results highlight the need for more adaptive fusion mechanisms to handle dynamic changes in UAV perspectives.

## VII. CONCLUSION

This paper introduces a novel framework for aerial-ground cooperative 3D detection and tracking in autonomous driving. It introduces the Griffin dataset, a multi-modal collection with occlusion-aware annotations and drone dynamics simulation, and AGILE, an efficient instance-level intermediate fusion method for joint detection and tracking tasks.

Experiments confirm the efficacy of cooperative perception while revealing several limitations of current methods, such as the generalization capability across varying UAV altitudes and the trade-off between communication bandwidth and perception performance.

Future research can focus on three key directions: (1) developing altitude-adaptive fusion mechanisms capable of handling dynamic aerial perspectives; (2) incorporating scale-awareness into cooperative perception frameworks to address altitude-related challenges; and (3) bridging the simulation-to-reality gap for reliable real-world deployments. Addressing these challenges will be crucial for achieving

robust aerial-ground cooperative perception in dynamic and unpredictable real-world conditions.

## REFERENCES

- [1] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma, "Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication," in *2022 International Conference on Robotics and Automation (ICRA)*, May 2022, pp. 2583–2589.
- [2] R. Xu, X. Xia, J. Li, H. Li, S. Zhang, Z. Tu, Z. Meng, H. Xiang, X. Dong, R. Song, H. Yu, B. Zhou, and J. Ma, "V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 13 712–13 722.
- [3] H. Yu, Y. Luo, M. Shu, Y. Huo, Z. Yang, Y. Shi, Z. Guo, H. Li, X. Hu, J. Yuan, and Z. Nie, "Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 21 329–21 338.
- [4] H. Yu, W. Yang, H. Ruan, Z. Yang, Y. Tang, X. Gao, X. Hao, Y. Shi, Y. Pan, N. Sun, J. Song, J. Yuan, P. Luo, and Z. Nie, "V2x-seq: A large-scale sequential dataset for vehicle-infrastructure cooperative perception and forecasting," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 5486–5495.
- [5] Y. Hu, S. Fang, Z. Lei, Y. Zhong, and S. Chen, "Where2comm: Communication-efficient collaborative perception via spatial confidence maps," in *Advances in Neural Information Processing Systems*, vol. 35, Dec. 2022, pp. 4874–4886.
- [6] H. Ye, R. Sunderraman, and J. S. Ji, "Uav3d: A large-scale 3d perception benchmark for unmanned aerial vehicles," *Advances in Neural Information Processing Systems*, vol. 37, pp. 55 425–55 442, Dec. 2024.
- [7] P. Tian, P. Cheng, Y. Wang, Z. Wang, Z. Wang, M. Yan, X. Yang, and X. Sun, "Ucdnet: Multi-uav collaborative 3d object detection network by reliable feature mapping," *IEEE Transactions on Geoscience and Remote Sensing*, June 2024.
- [8] Z. Wang, P. Cheng, M. Chen, P. Tian, Z. Wang, X. Li, X. Yang, and X. Sun, "Drones help drones: A collaborative framework for multi-drone object trajectory prediction and beyond," in *Advances in Neural Information Processing Systems*, vol. 37. Curran Associates, Inc., 2024, pp. 64 604–64 628.
- [9] Y. Wang, P. Cheng, P. Tian, Z. Yuan, L. Zhao, J. Tian, W. Wang, Z. Wang, and X. Sun, "Uvcpnet: A uav-vehicle collaborative perception network for 3d object detection," June 2024.
- [10] Y. Zhou, L. Quang, C. Nieto-Granda, and G. Loianno, "Coped-advancing multi-robot collaborative perception: A comprehensive dataset in real-world environments," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6416–6423, July 2024.
- [11] A. Caillot, S. Ouerghi, P. Vasseur, R. Boutteau, and Y. Dupuis, "Survey on cooperative perception in an automotive context," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 9, pp. 14 204–14 223, Sept. 2022.
- [12] 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, Nov. 2023.
- [13] B. Gao, J. Liu, H. Zou, J. Chen, L. He, and K. Li, "Vehicle-road-cloud collaborative perception framework and key technologies: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 12, pp. 19 295–19 318, Dec. 2024.
- [14] Q. Chen, S. Tang, Q. Yang, and S. Fu, "Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. Dallas, TX, USA: IEEE, July 2019, pp. 514–524.
- [15] E. Arnold, M. Dianati, R. De Temple, and S. Fallah, "Cooperative perception for 3d object detection in driving scenarios using infrastructure sensors," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 3, pp. 1852–1864, Mar. 2022.
- [16] Y. Hu, Y. Lu, R. Xu, W. Xie, S. Chen, and Y. Wang, "Collaboration helps camera overtake lidar in 3d detection," in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 9243–9252.
- [17] Y. Li, S. Ren, P. Wu, S. Chen, C. Feng, and W. Zhang, "Learning distilled collaboration graph for multi-agent perception," in *Advances in Neural Information Processing Systems*, vol. 34. Curran Associates, Inc., 2021, pp. 29 541–29 552.
- [18] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun, "V2vnet: Vehicle-to-vehicle communication for joint perception and prediction," in *Computer Vision – ECCV 2020*, ser. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2020, pp. 605–621.
- [19] Z. Chen, Y. Shi, and J. Jia, "Transiff: An instance-level feature fusion framework for vehicle-infrastructure cooperative 3d detection with transformers," in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. Paris, France: IEEE, Oct. 2023, pp. 18 159–18 168.
- [20] S. Fan, H. Yu, W. Yang, J. Yuan, and Z. Nie, "Quest: Query stream for practical cooperative perception," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, May 2024, pp. 18 436–18 442.
- [21] S. Minaeian, J. Liu, and Y.-J. Son, "Vision-based target detection and localization via a team of cooperative uav and ugvs," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 46, no. 7, pp. 1005–1016, July 2016.
- [22] Y. Li, D. Ma, Z. An, Z. Wang, Y. Zhong, S. Chen, and C. Feng, "V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 914–10 921, Oct. 2022.
- [23] R. Hao, S. Fan, Y. Dai, Z. Zhang, C. Li, Y. Wang, H. Yu, W. Yang, J. Yuan, and Z. Nie, "Rcooper: A real-world large-scale dataset for roadside cooperative perception," in *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 22 347–22 357.
- [24] Y. Zhu, Y. Kong, Y. Jie, S. Xu, and H. Cheng, "Graco: A multimodal dataset for ground and aerial cooperative localization and mapping," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 966–973, Feb. 2023.
- [25] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, Sept. 2013.
- [26] 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 *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020, pp. 11 618–11 628.
- [27] H. Yu, Y. Tang, E. Xie, J. Mao, P. Luo, and Z. Nie, "Flow-based feature fusion for vehicle-infrastructure cooperative 3d object detection," *Advances in Neural Information Processing Systems*, vol. 36, pp. 34 493–34 503, Dec. 2023.
- [28] D. Yang, K. Yang, Y. Wang, J. Liu, Z. Xu, R. Yin, P. Zhai, and L. Zhang, "How2comm: Communication-efficient and collaboration-pragmatic multi-agent perception," *Advances in Neural Information Processing Systems*, vol. 36, pp. 25 151–25 164, Dec. 2023.
- [29] C. Lin, D. Tian, X. Duan, J. Zhou, D. Zhao, and D. Cao, "V2vformer: Vehicle-to-vehicle cooperative perception with spatial-channel transformer," *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 2, pp. 3384–3395, Feb. 2024.
- [30] J. Zhong, H. Yu, T. Zhu, J. Xu, W. Yang, Z. Nie, and C. Sun, "Leveraging temporal contexts to enhance vehicle-infrastructure cooperative perception," in *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2024.
- [31] H. Yu, W. Yang, J. Zhong, Z. Yang, S. Fan, P. Luo, and Z. Nie, "End-to-end autonomous driving through v2x cooperation," in *The 39th Annual AAAI Conference on Artificial Intelligence*, 2025.
- [32] 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 *Lecture Notes in Computer Science*, ser. Lecture Notes in Computer Science. Cham: Springer Nature Switzerland, 2022, pp. 1–18.
- [33] T. Meinhardt, A. Kirillov, L. Leal-Taixé, and C. Feichtenhofer, "Trackformer: Multi-object tracking with transformers," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 8834–8844.
- [34] X. Weng, J. Wang, D. Held, and K. Kitani, "3d multi-object tracking: A baseline and new evaluation metrics," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 10 359–10 366.