

CATS-V2V: A Real-World Vehicle-to-Vehicle Cooperative Perception Dataset with Complex Adverse Traffic Scenarios

Hangyu Li^{1,*}, Bofeng Cao^{1,*}, Zhaohui Liang¹, Wuzhen Li¹, Juyoung Oh¹, Yuxuan Chen¹, Shixiao Liang¹, Hang Zhou¹, Chengyuan Ma¹, Jiayi Liu¹, Zheng Li¹, Peng Zhang¹, KeKe Long¹, Maolin Liu², Jackson Jiang², Chunlei Yu², Shengxiang Liu², Hongkai Yu³, Xiaopeng Li¹
¹University of Wisconsin-Madison ²wuwei-ai ³Cleveland State University




Figure 1. One frame of CATS-V2V dataset. The middle-upper shows the combined point cloud and 3D bounding box annotations, while the middle-lower presents the HD map and BEV annotations. On sides are the seven camera views of each vehicle with projected annotations.

Abstract

Vehicle-to-Vehicle (V2V) cooperative perception has great potential to enhance autonomous driving performance by overcoming perception limitations in complex adverse traffic scenarios (CATS). Meanwhile, data serves as the fundamental infrastructure for modern autonomous driving AI. However, due to stringent data collection requirements, existing datasets focus primarily on ordinary traffic scenarios, constraining the benefits of cooperative perception. To address this challenge, we introduce CATS-V2V, the first-of-its-kind real-world dataset for V2V cooperative perception under complex adverse traffic scenarios. The dataset was collected by two hardware time-synchronized vehicles, covering 10 weather and lighting conditions across 10 diverse locations. The 100-clip dataset includes 60K frames of 10 Hz LiDAR point clouds and 1.26M multi-view 30 Hz camera images, along with 750K anonymized yet high-precision RTK-fixed GNSS and IMU records. Correspondingly, we provide time-consistent 3D bounding box annotations for objects, as well as static scenes to construct a 4D BEV representation. On this basis, we propose a

target-based temporal alignment method, ensuring that all objects are precisely aligned across all sensor modalities. We hope that CATS-V2V, the largest-scale, most supportive, and highest-quality dataset of its kind to date, will benefit the autonomous driving community in related tasks.

1. Introduction

Cooperative perception has emerged as a promising approach to overcoming the sensing limitations of autonomous vehicles in rare but safety-critical corner cases [12, 23]. Onboard sensors such as cameras and LiDARs are inherently constrained by occlusion, noise, and limited range, which can severely degrade perception performance in Complex Adverse Traffic Scenarios (CATS) [5, 19].

CATS encompass adverse weather and illumination conditions, irregular work zones, and visually challenging scenes that often cause sensor failures and decision-making instability [25, 37, 42]. These scenarios constitute the long-tail distribution of real-world driving, where the robustness and generalization of autonomous driving systems are most critically tested [6, 7].



Figure 2. 10 CATS scenarios at 10 locations: (a) Clear day at an arterial road; (b) Rainy day at a highway; (c) Snowy day at a highway; (d) Dawn at an arterial road; (e) Dusk with direct sunlight at an arterial road; (f) Clear night at a collector road; (g) Rainy night at an arterial road; (h) Snowy night at an arterial road; (i) Foggy day at a local street; (j) Overcast day with work zone at an arterial road.

To mitigate these issues, cooperative perception (CP) allows vehicles and infrastructure to share perceptual information, extending situational awareness beyond a single vehicle’s view [2, 28, 32]. Depending on the source, it can be categorized into Vehicle-to-Vehicle (V2V), Vehicle-to-Infrastructure (V2I), or hybrid V2X modes. Among them, V2V cooperative perception offers the greatest flexibility and scalability since it does not depend on fixed locations or costly infrastructure investment.

Despite its potential, developing and validating CP under real-world CATS remains difficult due to the lack of suitable datasets. Existing datasets [24, 29, 34, 39, 40, 44] have established a pioneering foundation, yet they mainly capture ordinary weather, lighting, and traffic conditions, where single-vehicle perception already works well. As a result, current CP research under CATS largely depends on digital simulations [13, 18] or indoor mockups [25].

While a recent V2I dataset [36] includes certain CATS conditions, such examples remain extremely rare and are limited to static infrastructure viewpoints. In contrast, V2V cooperation involves dynamic agents with diverse viewpoints and continuously changing spatial relations. While this flexibility is key for scalable CP, it also brings challenges in calibration, time synchronization, and temporal alignment, especially in CATS. To date, considering its stringent collecting conditions, no dataset captures real-world V2V perception under CATS.

To fill this gap, we introduce CATS-V2V, the first real-world V2V CP dataset covering CATS. It supports tasks on perception and localization with strict time synchronization, spatial-temporal alignment, and cross-view consistency. Data were collected across ten weather and lighting conditions and ten diverse locations (see Fig. 2) using two experimental vehicles equipped with synchronized sensors, including a 10 Hz LiDAR, seven multi-view 30 Hz cameras, and one 125 Hz inertial navigation system (INS).

The dataset correspondingly provides a time-consistent 3D bounding box annotations with global object IDs and static HD maps for constructing 4D BEV scenes (see Fig. 1). Furthermore, we propose a target-based temporal alignment method alongside with the dataset, which ensures all objects are precisely aligned across modalities.

Our key contributions are summarized as follows:

- The first real-world V2V cooperative perception dataset covering diverse weather (sunny, cloudy, rainy, snowy) and lighting (clear, direct sunlight, low light, night) conditions, including dynamic work zones.
- Broad task support with multiple sensor data (LiDAR, camera, INS) and rich annotations (3D boxes, HD maps, traffic information) including object detection, object tracking, trajectory prediction, localization and mapping.
- Target-based method enabling alignment across hardware time synchronized high-frame-rate multi-view cameras, LiDAR, and INS, achieving highest data quality.

We will release the dataset and code for target-based temporal alignment method to foster further research.

2. Related work

A variety of V2X datasets have been publicly released, yet they vary greatly in realism and coverage of CATS. While simulation-based datasets can easily model adverse conditions, real-world datasets remain dominated by normal weather and lighting. Moreover, most existing real-world V2X datasets still fall short of the quality in multi-modal richness, synchronization, and alignment achieved by well-known single-vehicle perception datasets such as KITTI [9], Waymo [26], and nuScenes [1]. Therefore, in this section, we provide a comprehensive review and comparison of existing simulation-based V2X datasets with CATS and real-world V2X datasets, including both V2V and V2I settings, as summarized in Tab. 1.

Table 1. Comparison of CATS-V2V to some existing V2X cooperative perception datasets.

Dataset	S/R	V2X	CATS	Sync	Align	BEV	Clip×Time	LiDAR	Cam	MV×FPS
V2XSet-w[18]	Sim	V2I	✓	Soft	-	-	55×25s	2.8K	-	-
OPV2V-w[18]	Sim	V2V	✓	Soft	-	-	73×13s	2.7K	-	-
DeepAccident[27]	Sim	Both	✓	Soft	-	-	691×8s	285K	1.71M	6×10
SCOPE[8]	Sim	Both	✓	Soft	-	o ^a	44×40s	500K	800K	5×10
Adver-City[14]	Sim	Both	✓	Soft	-	-	110×22s	120K	480K	4×10
DAIR-V2X[39]	Real	V2I	o ^b	30 ms	Stamp	-	100×20s	39K	39K	1×20 ^c
V2X-Seq[40]	Real	V2I	-	30 ms	-	✓	95×16s	30K	30K	1×10
TUMTraf-V2X[44]	Real	V2I	-	25 ms	Stamp	✓	8×10s	2K	5K	1×15 ^c
HoloVIC[24]	Real	V2I	-	25 ms ^d	Stamp	✓	? ^e	? ^e	? ^e	2×25 ^c
V2X-Radar[36]	Real	V2I	✓	20 ms	-	-	40×20s	16K	32K	1×10
V2V4Real[34]	Real	V2V	-	50 ms	Stamp	✓	67×15s	20K	40K	2×10
V2X-Real[29]	Real	Both	-	50 ms ^f	Stamp	-	68×12s	33K	171K	4×10
CATS-V2V(ours)	Real	V2V	✓	1 ms	Target	✓	100×30s	60K	1.26M	7×30

S/R = Sim/Real; V2X = V2I(nfrastructure), V2V, or both; BEV = BEV annotations; MV×FPS = Multi-View cameras on each vehicle×Frame-Per-Second.

^a It provides a BEV camera’s images but not structural representations.

^b It provides limited rainy scenarios as CATS.

^c The vehicle camera runs on 20 Hz / 15 Hz / 25 Hz, but only 10 Hz data is sampled for publication.

^d Only roadside units are claimed to be synchronized under 5 ms. Considering the characteristics of NTP that it employs, delays >25 ms can be expected.

^e The scale is unknown, as the figures in their paper are inconsistent, and the dataset is not publicly available.

^f The LiDARs triggered synchronously with GPS, but not the cameras. Considering their 10 Hz capture frequency, a 50 ms error can be expected.

2.1. Simulated V2X datasets with CATS

Simulation-based V2X datasets are typically built on high-fidelity simulators such as CARLA [4] and OpenCDA [31] to capture virtual sensor data.

OPV2V-w and V2XSet-w [18] extend the original OPV2V [33] and V2XSet [32] datasets to simulate CATS, incorporating virtual weather and lighting variations. Several simulation frameworks [10, 11, 15] were leveraged to emulate fog, rain, and snow effects on LiDAR sensors.

SCOPE [8] is a recently introduced large-scale simulation dataset involving 24 connected and autonomous vehicles (CAVs). It supports both realistic LiDAR and multi-view camera models and further reduces the sim-to-real gap by using two Digital Twin maps reconstructed from real-world environments.

DeepAccident [27] and Adver-City [14] focus on safety-critical cooperative perception, reconstructing accident-prone scenarios from real-world crash data with CATS.

Simulation datasets inherently provide accurate localization, perfect annotations, and precise time synchronization via software. Moreover, because sensor data are generated instantaneously, there is no need for high-frame-rate cameras and temporal alignment between modalities. However, they differ significantly from real-world conditions, not only in time, sensor information, and annotation errors, but particularly regarding the physical interaction of CATS with sensor hardware, which simulation cannot yet perfectly reproduce.

2.2. Real-world V2I datasets

DAIR-V2X [39] is one of the first real-world V2X datasets, establishing cooperation between a roadside unit and a vehicle. Subsequently, V2X-Seq [40] extended this setup to consecutive intersections with additional HD map and object tracking annotations. Although the data-collection vehicle in both datasets was equipped with a 20 Hz camera, only compressed front-view images downsampled to 10 Hz based on the timestamps were released. Moreover, both datasets suffer from imperfect time synchronization, with errors exceeding 30 ms.

TUMTraf-V2X [44] introduces a multi-view intersection dataset consisting of four infrastructure cameras and a single vehicle. The dataset is relatively small, with only eight 10-second clips. Similarly, only the downsampled images from the vehicle’s front-view camera were released. Although annotated for weather and time of day, the dataset only includes two conditions (clear daytime and nighttime) and thus lacks CATS coverage. Their earlier work, TUMTraf-A9 [3], captured snow conditions but solely from the infrastructure side, without vehicle cooperation.

HoloVIC [24] is a large-scale dataset covering four intersections equipped with multiple LiDARs and cameras. However, due to limited accessibility and statistical inconsistencies in the paper, information such as the number of clips, segment durations, and data scale remains unclear. It employs Network Time Protocol (NTP) for time synchronization, achieving reported timestamp accuracy of about 5




Figure 3. Sensor configurations of our two data-collecting vehicles.

ms, but only for roadside units. Given the characteristics of NTP, errors exceeding 25 ms are expected between roadside units and the vehicle. Two cameras are also downsampled and aligned by selecting the closest timestamps to the LiDAR frames.

V2X-Radar [36] is the most recent V2I dataset, introducing a radar modality and certain coverage of CATS. However, as a trade-off, it only provides 10 Hz images from one front-view camera, and no alignment could be performed among modalities. It also lacks BEV and corresponding tracking annotations. Although the vehicle platform is described as being hardware-synchronized, timestamp errors of approximately 20 ms were still reported.

2.3. Real-world V2V datasets

V2V4Real [34] is the first and one of the very few real-world V2V cooperative perception datasets that leverage the perception of two vehicles for collaboration. Each vehicle is equipped with a 32-beam LiDAR and two cameras facing forward and backward, but without hardware-level time synchronization. Images are then aligned to LiDAR frames using the nearest timestamps, which causes some delays. The dataset annotates only five classes of vehicles and does not include vulnerable road users such as pedestrians or cyclists, which may be due to both scene composition and the limited LiDAR resolution.

V2X-Real [29] is currently the only real-world dataset supporting both V2I and V2V cooperation. It employs 128-beam LiDARs to produce dense point clouds, with

LiDARs synchronized to GPS. However, four on-vehicle ZEDi cameras and two Axis infrastructure cameras lack hardware synchronization. As in other V2X datasets, images are aligned to LiDAR frames by closest timestamps. The dataset does not include CATS, or BEV and tracking annotations, and its camera configuration provides limited field-of-view overlap, despite a nominal 360° coverage.

3. CATS-V2V dataset

Despite these pioneering efforts, existing real-world V2V datasets still lack coverage of CATS. More broadly, current real-world V2X datasets remain limited in quality, with imperfect time synchronization, multi-modal alignment, and annotation richness. In response, we introduce our CATS-V2V dataset to address these gaps.

The section describes our vehicle configuration, data acquisition, preprocessing, and annotation procedures. More details could be found in the Supplementary Material.

3.1. Vehicle configurations

We use two Lincoln MKZ sedans to collect data, as shown in Fig. 3, with sensors integrated through the roof rack. During the data collection process, they are driven by humans.

3.1.1. Sensors setup

Each of our two vehicles is equipped with a 128-beam mechanical spinning LiDAR, seven automotive-grade cameras, and one deeply-coupled Inertial Navigation System (INS). The detailed specifications are shown in Tab. 2.

Table 2. Detailed sensor specifications on vehicles

Sensor	Details
LiDAR	Black: 10 Hz RoboSense Ruby 128-beam LiDAR, dual return mode, 250m range, vertical FOV $-25^\circ \sim 15^\circ$, angle resolution 0.2° ; Red: 10 Hz RoboSense Ruby Plus 128-beam LiDAR, dual return mode, 250m range, vertical FOV $-25^\circ \sim 15^\circ$, angle resolution 0.2° ;
Camera	Black: $3 \times$ (front and rear) 30 Hz OMNIVISION OX08B40, YUV422 8bit, 3840×2160 , 140dB HDR, LFM; $4 \times$ (side) Sony ISX031, YUV422 8bit, 1920×1536 , 120dB HDR, LFM; Red: $3 \times$ (front and rear) 30 Hz OMNIVISION OX08B40, YUV422 8bit, 3840×2160 , 140dB HDR, LFM; $4 \times$ (side) OMNIVISION OX03C10, YUV422 8bit, 1920×1080 , 140dB HDR, LFM;
INS	Deeply-coupled INS integration with Epson G320 IMU, 1 cm + 1 ppm (RMS) accuracy with RTK, $0.5^\circ/\text{h}$ bias instability.

Each vehicle is equipped with a 360° mechanical spinning LiDAR mounted at the top-center position. To mitigate interference from water droplets and snowflakes under CATS, the LiDARs are operated in dual-return mode, providing both the strongest and last points.

Seven cameras are installed on each vehicle, including two front-view (wide-angle and telephoto), one rear-view, and four side-view ones. The specific fields of view (FOVs) are illustrated in Fig. 3. All cameras feature High Dynamic Range (HDR) with LED Flicker Mitigation (LFM) capability, which enhances robustness under CATS.

Each vehicle is also equipped with a high-precision INS connected to a nearby Real-Time Kinematic (RTK) base station within 10 km, leading to an accuracy better than 2 cm (1 cm + 1 ppm). During data collection, both systems maintain a fixed RTK status. In occasionally obstructed environments (e.g., beneath trees or bridges) the tactical-grade IMU ensures short-term localization without drift.

3.1.2. Time synchronization

Time synchronization is crucial for multi-modal cooperative perception, as it ensures consistency across all sensors and agents. In our system, all sensors on both vehicles are hardware-synchronized to GPS time and triggered at integer seconds. The overall synchronization topology is illustrated in Fig. 4.

Each vehicle is equipped with an FPGA card responsible for camera triggering and image acquisition. The card re-




Figure 4. Time synchronization topology for all sensors and module controllers of both vehicles.

ceives the GPRMC (contains absolute time) and PPS (Pulse Per Second, integer trigger) signals from the onboard INS, achieving an offset within 20 ns to the GPS time. It then serves as the master clock, synchronizing all other sensors and module controllers within the local network via the Precision Time Protocol (PTP) through an Ethernet switch, maintaining a time deviation below 1 ms relative to GPS.

Overall, our system achieves a 1 ms synchronization accuracy across all sensors and vehicles, which is an order of magnitude improvement over existing datasets, which typically exhibit over 20 ms offsets, indicating a 40 cm misalignment of objects relatively moving at 72 km/h.

3.1.3. Calibration

We use the factory intrinsic calibration for cameras to achieve the highest accuracy, while for extrinsic parameters, we follow two open-source tools [17, 43] and conduct both LiDAR-Camera and LiDAR-INS calibrations.

The LiDAR-Camera extrinsics are estimated by directly aligning the dense LiDAR intensity mapping with the corresponding grayscale camera image [17]. On the other hand, the LiDAR-INS extrinsics are derived from LiDAR odometry with short-term IMU inertial integration (as the origin of the INS coordinate locates at the IMU) [43].

3.2. Data acquisition

We conduct data collection across ten diverse locations, including highways, arterial roads, collector roads, and local streets (see Fig. 2). For highway environments, we capture both mainline and ramp scenarios. On lower-level roads, our dataset includes signalized, all-way-stop, and unprotected intersections. Residential and campus areas are touched, where vulnerable road users (VRUs) such as pedestrians and cyclists are more frequent.

Each location contains ten repeated runs recorded under varying weather and lighting conditions (see Fig. 2), all referenced to the same HD maps. All sensor data are recorded using ROS2 bag files. The cameras, LiDARs, and INS each publish hardware timestamps through modified drivers.

3.3. Preprocessing

We then perform motion compensation for LiDAR frames to deskew distortions caused by ego-motion during scanning (0.1 s at 10 Hz, up to 2 m error at 72 km/h). The process is essential for accurate multi-frame registration and object localization [38], yet existing datasets [22, 30] acknowledge but do not address the issue, partially due to the absence of per-point timestamps or precise poses.

To ensure consistent annotations between the two vehicles, we register their point clouds into a unified coordinate system. The transformation between the two LiDAR coordinate systems is derived from the vehicle poses and the LiDAR-INS extrinsics, as expressed in (1a):

$$T_{L_1 L_2}^{\text{init}} = (T_{I_1 L_1})^{-1} (T_{W I_1})^{-1} T_{W I_2} T_{I_2 L_2}, \quad (1a)$$

$$T_{L_1 L_2}^{\text{init}} \xrightarrow{\text{GICP}} T_{L_1 L_2}^{\text{refine}}, \quad (1b)$$

where T_{WI} denotes the pose of the INS in the world coordinate frame, and T_{IL} represents the extrinsic calibration between the INS and the LiDAR.

However, small accumulated errors from four transformations can propagate (see Fig. 5). Therefore, we refine it using Generalized Iterative Closest Point (GICP) [16] to get $T_{L_1 L_2}^{\text{refine}}$, setting $T_{L_1 L_2}^{\text{init}}$ as the initial estimation.

3.4. Data annotation

We provide precise 3D bounding boxes, with time-consistent object dimensions for rigid targets. Each dynamic object is also assigned a globally unique ID, enabling cross-frame and cross-agent re-identification.

Dynamic objects are categorized into Vehicles and VRUs. The Vehicle category includes *Car*, *Van*, *Truck*, *Trailer*, *Bus*, and *Others*; The VRU category includes *Pedestrian*, *Scooter*, *Bicycle*, and *Motorcycle* rider.

Additionally, we introduce a virtual link attribute to indicate physical coupling between two annotated objects. For instance, a link between a *Car* and a *Trailer* object may represent a pickup truck towing a trailer.




Figure 5. Comparison between two deskewed point clouds register (a) with initial estimation; (b) after GICP refinement.

4. Supporting tasks

Thanks to our hardware-synchronized multi-sensor setup, high-quality temporal alignment, and scale of diverse multi-pass data (detailed in Sec. 3), the CATS-V2V dataset is well-suited for a wide range of computer vision and multi-modal perception tasks, especially for rare cases.

In addition to providing high-quality 2D and 3D annotations, we offer a set of task-specific data conversion and segment tools to facilitate research and practical usage across different domains.

In the following Tab. 3, we summarize the five major categories of vision tasks[35][21][41][20] supported by CATS-V2V. Detailed data processing pipelines and annotation formats for each task are provided in the Supplementary Material.

Due to the extensive range of supported tasks, we do not provide benchmarking results for every task within this dataset release. Instead, we will introduce comprehensive evaluations in a separate future study based on CATS-V2V.

5. Target-based temporal alignment

Despite fine calibrations and millisecond time synchronization, precise temporal alignment across modalities is not

Table 3. List of supporting tasks of CATS-V2V dataset.

Category	Tasks	Modalities	CATs-V2V Support
Perception	2D/3D Detection	LiDAR; Multi-view Camera	2D and 3D bounding box annotations
	2D/3D Tracking	LiDAR; Multi-view Camera	2D and 3D bounding box annotations; Global object IDs across frames
Spatial Understanding	Map Generation	LiDAR; IMU/INS; Multi-view Camera	High-frequency IMU and RTK-fixed INS; HD maps and BEV annotations
	SLAM/Odometry	All	High-frequency IMU and RTK-fixed INS
	3D Reconstruction	All	HD Maps and geometry transformations based-on RTK-fixed INS
Multi-Modal Learning	Joint Compression	LiDAR; Multi-view Camera	Hardware-synchronized sensor frames with deskewing and temporal alignment
Cross-Modal Learning	Depth Estimation	LiDAR; Multi-view Camera	Geometry transformations based-on RTK-fixed INS; Multi-frame fused point clouds for depth supervision
	View Synthesis	IMU/INS; Multi-view Camera	Overlapping camera frames; Multi-frame fused point clouds
Domain Adaptation	Scene Transfer	Multi-view Camera	Filtered sets of frames from identical scenes captured under different conditions

guaranteed. In this section, we analyze this limitation and progressively introduce our frame-based and target-based temporal alignment method.

5.1. Stamp-based alignment

While only a few real-world V2X datasets are equipped with multi-view or high-frame-rate cameras [24, 29, 34, 39, 44], they mostly report that LiDAR and camera frames are aligned with the closest timestamps (see Tab. 1).

While this approach provides a convenient approximation, it neglects the inherent characteristics of mechanical spinning LiDARs, which acquire points continuously over a full revolution rather than instantaneously. As a result, each azimuth within a single scan corresponds to a slightly different timestamp. Consequently, stamp-based alignment introduces temporal misalignment, as shown in Fig. 6.

5.2. Frame-based alignment

In contrast, our dataset adopts a frame-based alignment strategy that combines 30 Hz multi-view cameras with LiDAR scans containing per-point timestamps. For each camera, we select the subset of LiDAR points falling within its FOV and associate them with the camera frame whose timestamp is closest to the acquisition times of those points.

We take one of our vehicle’s sensor configurations as an example (see Fig. 6). Since cameras operate at $3 \times$ the LiDAR frame rate, the two forward cameras are aligned with the starting timestamp of each LiDAR scan, the front-

right camera is aligned with the subsequent image frame, the three rear-facing cameras use the next frame, and the front-left camera corresponds to the ending timestamp.

At preprocessing (Sec. 3.3), we perform motion compensation to correct distortions by compensating all points to the timestamp of the first scan point. This procedure effectively deskews the static background of each LiDAR frame.

However, for frame-based alignment with cameras, the LiDAR points must further be compensated to the corresponding camera frame timestamp, so that both the static environment and dynamic objects are geometrically consistent across modalities.

5.3. Target-based alignment

While frame-based alignment greatly improves temporal consistency across modalities, it still suffers from rare misalignment in multi-camera overlap or wide-FOV scenarios.

In practice, a single object may appear simultaneously in multiple camera views whose frame-based alignments correspond to different LiDAR scanning times. Similarly, wide-angle camera images may contain a few objects whose alignment should correspond to a previous or a next frame.

To address this issue, we introduce a target-based temporal alignment strategy. After annotation (or LiDAR-based object detection), we compute the average timestamp of all points belonging to each object and associate the object with the nearest camera frame. The corresponding LiDAR points are then motion-compensated to that timestamp.




Figure 6. Illustration of alignment between camera and LiDAR frames due to the inherent characteristics of mechanical spinning LiDARs.

Table 4. Quantitative evaluation of our proposed temporal alignment methods.

Method base	average IoU	Recall@IoU=0.3	Recall@IoU=0.5	Recall@IoU=0.7	center-offset (px)
Stamp	0.3736	0.6206	0.3906	0.1172	61.54
Frame (ours)	0.4493 (↑ 20.3%)	0.6795 (↑ 9.5%)	0.5766 (↑ 47.6%)	0.2494 (↑ 113%)	50.26 (↓ 18.3%)
Target (ours)	0.4623 (↑ 23.7%)	0.6932 (↑ 11.7%)	0.5947 (↑ 52.3%)	0.2768 (↑ 136%)	49.76 (↓ 19.1%)

5.4. Quantitative Evaluation

To quantitatively evaluate our proposed alignment strategies, we conduct thorough experiments on one representative clip selected from the 100 recorded scenes.

This clip is captured under favorable conditions without strong sunlight, or inclement weather, so as to minimize perception noise and annotation uncertainty. We manually annotate 2D bounding boxes for all visible dynamic objects across five camera views (not on the two right-view cameras, as we drive on the right lane and they contain few objects), providing ground truth for alignment evaluation.

For benchmarking, we project the annotated 3D bounding boxes onto the images using three different temporal alignment methods: (1) Stamp-based alignment, (2) Frame-based alignment (ours), and (3) Target-based alignment (ours). The projected boxes are then compared against the manually annotated 2D boxes on the corresponding images.

We evaluate the alignment accuracy using three standard metrics: average IoU, Recall at different IoU levels, and mean center-point deviation between the projected and annotated 2D bounding boxes. As shown in Tab. 4, higher average IoU and Recall among different IoU levels, along with lower center offset, indicate better temporal alignment across modalities progressively with our proposed frame-based and target-based temporal alignment methods.

6. Conclusion

In this work, we present **CATS-V2V**, the first real-world cooperative perception dataset with **Complex Adverse Traffic Scenarios** under the **Vehicle-to-Vehicle (V2V)** collaboration domain. It is a large-scale dataset collected with two vehicles across ten diverse weather, lighting, and traffic conditions (rain, snow, direct sunlight, nighttime, etc) at ten locations, including highways, arterial roads, and urban intersections in campus and residential areas. Specifically, the dataset provides anonymized point clouds, images, and pose records from one LiDAR, seven cameras, and an INS for each vehicle, as well as time-consistent global-identified 3D bounding box annotations and HD Maps. We additionally propose a target-based temporal alignment method to improve the quality of our dataset along with hardware time synchronization. Results on one selected clip from the dataset have demonstrated that our method significantly improves alignment performance across modalities. Furthermore, we offer task-specific data conversion and segment tools to facilitate our supporting tasks. We hope this largest-scale, most supportive, and highest-quality V2V dataset with CATS to date could promote research in related communities. Future plans include combining roadside infrastructure and diverse emerging automotive sensors to provide a richer dataset covering various corner cases and developing tools for converting it into motion and trajectory datasets.

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 2
- [2] Qi Chen, Xu Ma, Sihai Tang, Jingda Guo, Qing Yang, and Song Fu. F-cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3d point clouds. In *Proceedings of the 4th ACM/IEEE Symposium on Edge Computing*, pages 88–100, 2019. 2
- [3] Christian Creß, Walter Zimmer, Leah Strand, Maximilian Fortkord, Siyi Dai, Venkatnarayanan Lakshminarasimhan, and Alois Knoll. A9-dataset: Multi-sensor infrastructure-based dataset for mobility research. In *2022 IEEE Intelligent Vehicles Symposium (IV)*, pages 965–970. IEEE, 2022. 3
- [4] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017. 3
- [5] Rui Fan, Sicen Guo, and Mohammad Junaid Bocus. Autonomous driving perception. *Cham, Switzerland: Springer*, 2023. 1
- [6] Shuo Feng, Xintao Yan, Haowei Sun, Yiheng Feng, and Henry X Liu. Intelligent driving intelligence test for autonomous vehicles with naturalistic and adversarial environment. *Nature communications*, 12(1):748, 2021. 1
- [7] Shuo Feng, Haowei Sun, Xintao Yan, Haojie Zhu, Zhengxia Zou, Shengyin Shen, and Henry X Liu. Dense reinforcement learning for safety validation of autonomous vehicles. *Nature*, 615(7953):620–627, 2023. 1
- [8] Jörg Gamberdinger, Sven Teufel, Patrick Schulz, Stephan Amann, Jan-Patrick Kirchner, and Oliver Bringmann. Scope: A synthetic multi-modal dataset for collective perception including physical-correct weather conditions. In *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*, pages 2622–2628. IEEE, 2024. 3
- [9] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The international journal of robotics research*, 32(11):1231–1237, 2013. 2
- [10] Martin Hahner, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Fog simulation on real lidar point clouds for 3d object detection in adverse weather. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15283–15292, 2021. 3
- [11] Martin Hahner, Christos Sakaridis, Mario Bijelic, Felix Heide, Fisher Yu, Dengxin Dai, and Luc Van Gool. Lidar snowfall simulation for robust 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16364–16374, 2022. 3
- [12] Yushan Han, Hui Zhang, Huifang Li, Yi Jin, Congyan Lang, and Yidong Li. Collaborative perception in autonomous driving: Methods, datasets, and challenges. *IEEE Intelligent Transportation Systems Magazine*, 15(6):131–151, 2023. 1
- [13] Ping Jiang, Xiaoheng Deng, Weishang Wu, Lixin Lin, Xuechen Chen, Chen Chen, and Shaohua Wan. Weather-aware collaborative perception with uncertainty reduction. *IEEE Transactions on Intelligent Transportation Systems*, 2024. 2
- [14] Mateus Karvat and Sidney Givigi. Adver-city: Open-source multi-modal dataset for collaborative perception under adverse weather conditions. *arXiv preprint arXiv:2410.06380*, 2024. 3
- [15] Velat Kilic, Deepti Hegde, A Brinton Cooper, Vishal M Patel, and Mark Foster. Lidar light scattering augmentation (lisa): Physics-based simulation of adverse weather conditions for 3d object detection. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2025. 3
- [16] Kenji Koide, Masashi Yokozuka, Shuji Oishi, and Atsuhiko Banno. Voxelized gicp for fast and accurate 3d point cloud registration. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11054–11059. IEEE, 2021. 6
- [17] Kenji Koide, Shuji Oishi, Masashi Yokozuka, and Atsuhiko Banno. General, single-shot, target-less, and automatic lidar-camera extrinsic calibration toolbox. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11301–11307. IEEE, 2023. 5
- [18] Baolu Li, Jinlong Li, Xinyu Liu, Runsheng Xu, Zhengzhong Tu, Jiacheng Guo, Qin Zou, Xiaopeng Li, and Hongkai Yu. V2x-dgw: Domain generalization for multi-agent perception under adverse weather conditions. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 974–980. IEEE, 2025. 2, 3
- [19] Hangyu Li, Xiaotong Sun, Chenglin Zhuang, and Xiaopeng Li. On the robotic uncertainty of fully autonomous traffic: From stochastic car-following to mobility–safety trade-offs. *Transportation Research Part C: Emerging Technologies*, 178:105254, 2025. 1
- [20] Jinlong Li, Baolu Li, Zhengzhong Tu, Xinyu Liu, Qing Guo, Felix Juefei-Xu, Runsheng Xu, and Hongkai Yu. Light the night: A multi-condition diffusion framework for unpaired low-light enhancement in autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15205–15215, 2024. 6
- [21] Yuhuan Lin, Tongda Xu, Ziyu Zhu, Yanghao Li, Zhe Wang, and Yan Wang. Your camera improves your point cloud compression. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023. 6
- [22] Katie Z Luo, Minh-Quan Dao, Zhenzhen Liu, Mark Campbell, Wei-Lun Chao, Kilian Q Weinberger, Ezio Malis, Vincent Fremont, Bharath Hariharan, Mao Shan, et al. Mixed signals: A diverse point cloud dataset for heterogeneous lidar v2x collaboration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 28763–28773, 2025. 6
- [23] Chengyuan Ma, Hangyu Li, Keke Long, Hang Zhou, Zhao-hui Liang, Pei Li, Hongkai Yu, and Xiaopeng Li. Real-time identification of cooperative perception necessity in road traffic scenarios. *Available at SSRN 4973353*, 2024. 1

- [24] Cong Ma, Lei Qiao, Chengkai Zhu, Kai Liu, Zelong Kong, Qing Li, Xueqi Zhou, Yuheng Kan, and Wei Wu. Holovic: Large-scale dataset and benchmark for multi-sensor holographic intersection and vehicle-infrastructure cooperative. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22129–22138, 2024. 2, 3, 7
- [25] Fatih Sezgin, Daniel Vriesman, Dagmar Steinhauser, Robert Lugner, and Thomas Brandmeier. Safe autonomous driving in adverse weather: Sensor evaluation and performance monitoring. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, pages 1–6. IEEE, 2023. 1, 2
- [26] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 2
- [27] Tianqi Wang, Sukmin Kim, Ji Wenxuan, Enze Xie, Chongjian Ge, Junsong Chen, Zhenguo Li, and Ping Luo. Deepaccident: A motion and accident prediction benchmark for v2x autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5599–5606, 2024. 3
- [28] Tsun-Hsuan Wang, Sivabalan Manivasagam, Ming Liang, Bin Yang, Wenyuan Zeng, and Raquel Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *Computer vision–ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part II 16*, pages 605–621. Springer, 2020. 2
- [29] Hao Xiang, Zhaoliang Zheng, Xin Xia, Runsheng Xu, Letian Gao, Zewei Zhou, Xu Han, Xinkai Ji, Mingxi Li, Zonglin Meng, et al. V2x-real: a large-scale dataset for vehicle-to-everything cooperative perception. In *European Conference on Computer Vision*, pages 455–470. Springer, 2024. 2, 3, 4, 7
- [30] Hao Xiang, Zhaoliang Zheng, Xin Xia, Seth Z Zhao, Letian Gao, Zewei Zhou, Tianhui Cai, Yun Zhang, and Jiaqi Ma. V2x-realo: An open online framework and dataset for cooperative perception in reality. *arXiv preprint arXiv:2503.10034*, 2025. 6
- [31] Runsheng Xu, Yi Guo, Xu Han, Xin Xia, Hao Xiang, and Jiaqi Ma. Opencda: an open cooperative driving automation framework integrated with co-simulation. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 1155–1162. IEEE, 2021. 3
- [32] Runsheng Xu, Hao Xiang, Zhengzhong Tu, Xin Xia, Ming-Hsuan Yang, and Jiaqi Ma. V2x-vit: Vehicle-to-everything cooperative perception with vision transformer. In *European conference on computer vision*, pages 107–124. Springer, 2022. 2, 3
- [33] Runsheng Xu, Hao Xiang, Xin Xia, Xu Han, Jinlong Li, and Jiaqi 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)*, pages 2583–2589. IEEE, 2022. 3
- [34] Runsheng Xu, Xin Xia, Jinlong Li, Hanzhao Li, Shuo Zhang, Zhengzhong Tu, Zonglin Meng, Hao Xiang, Xiaoyu Dong, Rui Song, et al. V2v4real: A real-world large-scale dataset for vehicle-to-vehicle cooperative perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13712–13722, 2023. 2, 3, 4, 7
- [35] Lei Yang, Tao Tang, Jun Li, Peng Chen, Kun Yuan, Li Wang, Yi Huang, Xinyu Zhang, and Kaicheng Yu. Bevheight++: Toward robust visual centric 3d object detection. *arXiv preprint arXiv:2309.16179*, 2023. 6
- [36] Lei Yang, Xinyu Zhang, Chen Wang, Jun Li, Jiaqi Ma, Zhiying Song, Tong Zhao, Ziyang Song, Li Wang, Mo Zhou, et al. V2x-radar: A multi-modal dataset with 4d radar for cooperative perception. *arXiv preprint arXiv:2411.10962*, 2024. 2, 3, 4
- [37] Keisuke Yoneda, Naoki Suganuma, Ryo Yanase, and Mohammad Aldibaja. Automated driving recognition technologies for adverse weather conditions. *IATSS research*, 43(4): 253–262, 2019. 1
- [38] David Yoon, Tim Tang, and Timothy Barfoot. Mapless online detection of dynamic objects in 3d lidar. In *2019 16th Conference on Computer and Robot Vision (CRV)*, pages 113–120. IEEE, 2019. 6
- [39] Haibao Yu, Yizhen Luo, Mao Shu, Yiyi Huo, Zebang Yang, Yifeng Shi, Zhenglong Guo, Hanyu Li, Xing Hu, Jirui Yuan, et al. Dair-v2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21361–21370, 2022. 2, 3, 7
- [40] Haibao Yu, Wenxian Yang, Hongzhi Ruan, Zhenwei Yang, Yingjuan Tang, Xu Gao, Xin Hao, Yifeng Shi, Yifeng Pan, Ning Sun, et al. V2x-seq: A large-scale sequential dataset for vehicle-infrastructure cooperative perception and forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5486–5495, 2023. 2, 3
- [41] Chunran Zheng, Wei Xu, Zuhao Zou, Tong Hua, Chongjian Yuan, Dongjiao He, Bingyang Zhou, Zheng Liu, Jiarong Lin, Fangcheng Zhu, Yunfan Ren, Rong Wang, Fanle Meng, and Fu Zhang. Fast-livo2: Fast, direct lidar-inertial-visual odometry. *IEEE Transactions on Robotics*, 41:326–346, 2025. 6
- [42] Ziqiang Zheng, Yujie Cheng, Zhichao Xin, Zhibin Yu, and Bing Zheng. Robust perception under adverse conditions for autonomous driving based on data augmentation. *IEEE Transactions on Intelligent Transportation Systems*, 24(12): 13916–13929, 2023. 1
- [43] Fangcheng Zhu, Yunfan Ren, and Fu Zhang. Robust real-time lidar-inertial initialization. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3948–3955. IEEE, 2022. 5
- [44] Walter Zimmer, Gerhard Arya Wardana, Suren Sritharan, Xingcheng Zhou, Rui Song, and Alois C Knoll. Tumtraf v2x cooperative perception dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22668–22677, 2024. 2, 3, 7