

DRIVEXQA: Cross-modal Visual Question Answering for Adverse Driving Scene Understanding

Mingzhe Tao^{1,*}, Ruiping Liu^{1,*†}, Junwei Zheng¹, Yufan Chen¹, Kedi Ying¹, M. Saquib Sarfraz^{1,3}, Kailun Yang², Jiaming Zhang^{2‡}, and Rainer Stiefelhagen¹

¹Karlsruhe Institute of Technology ²Hunan University ³Mercedes-Benz Tech Innovation

Abstract

Fusing sensors with complementary modalities is crucial for maintaining a stable and comprehensive understanding of abnormal driving scenes. However, Multimodal Large Language Models (MLLMs) are underexplored for leveraging multi-sensor information to understand adverse driving scenarios in autonomous vehicles. To address this gap, we propose the DRIVEXQA, a multimodal dataset for autonomous driving VQA. In addition to four visual modalities, five sensor failure cases, and five weather conditions, it includes 102,505 QA pairs categorized into three types: global scene level, allocentric level, and ego-vehicle centric level. Since no existing MLLM framework adopts multiple complementary visual modalities as input, we design MVX-LLM, a token-efficient architecture with a Dual Cross-Attention (DCA) projector that fuses the modalities to alleviate information redundancy. Experiments demonstrate that our DCA achieves improved performance under challenging conditions such as foggy (GPTScore: 53.5 vs. 25.1 for the baseline). The dataset and source code are at <https://github.com/jtjmd/DRIVEXQA>.

1. Introduction

Autonomous vehicles require a holistic understanding of complex driving environments through multiple sensor modalities to ensure safe navigation. While existing MLLMs have achieved remarkable success in general vision-language tasks [1, 27], they remain underexplored for leveraging multi-sensor information in safety-critical autonomous driving scenarios. Unlike general visual understanding tasks, driving scenarios demand a precise understanding of spatial relationships, environmental conditions, and sensor reliability under adverse conditions such as foggy, rainy, and various sensor failures. As illustrated

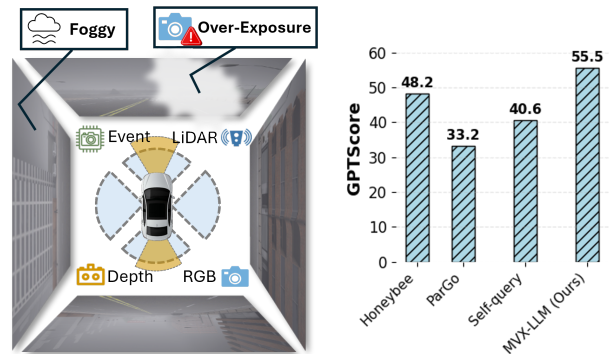


Figure 1. **Left:** Two corner cases of adverse driving scenes (*foggy* condition causing poor visibility and camera *over-exposure* resulting in degraded image quality). **Right:** Performance (GPTScore) comparison on DRIVEXQA dataset.

in Fig. 1, adverse driving conditions such as *foggy* and camera *over-exposure* represent critical challenges that existing systems fail to address comprehensively. These abnormal scenarios are commonplace in real-world deployment yet severely compromise perception quality and system reliability.

Visual Question Answering (VQA) systems serve as critical auxiliary cognition components in autonomous driving, enabling vehicles to understand and reason about complex driving scenarios through natural language interactions. However, current autonomous driving VQA systems primarily focus on normal conditions with limited sensor diversity and inadequate handling of abnormal situations. Existing datasets, including NuScenes-QA [36], DriveLM [40], and LingoQA [31], lack systematic coverage of adverse environmental conditions and multi-modal sensor fusion strategies under sensor degradation scenarios. MLLMs [5, 47] as auxiliary cognition systems in autonomous driving should ideally leverage multiple complementary sensor modalities to adapt to adverse conditions through cross-modal compensation, where degraded information from one sensor can be compensated by reliable

*Equal contribution.

†Project lead.

‡Corresponding author.

signals from others. While recent VLA models and BEV-based Transformers fuse multi-sensor data for action prediction, no existing MLLM-based VQA framework simultaneously supports multiple complementary visual modalities with token efficiency and robustness under sensor degradation. Current approaches [42, 50] typically rely on single-modality inputs (predominantly RGB) or simplistic fusion mechanisms that fail to leverage the complementary nature of different sensor modalities effectively. This limitation becomes particularly problematic under adverse conditions where individual sensors may provide degraded or unreliable information, when the interpretable reasoning capabilities of VQA systems are most needed to support critical driving decisions. For instance, in heavy fog or nighttime driving, cameras lose visibility and LiDAR returns become unreliable, exactly the scenarios where a driver most urgently needs the system to explain what it sees and why it recommends a particular action.

To address these limitations, we first propose the **DRIVEXQA** dataset for evaluating MLLMs in adverse driving scene understanding. DRIVEXQA contains 7,885 frames of driving scenes, systematically covering diverse weather conditions including *rainy*, *night*, *foggy*, *cloudy*, and *sunny* scenarios, along with five types of sensor failures: *motion blur* (MB), *overexposure* (OE), *underexposure* (UE), *LiDAR jitter* (LJ), and *event low-resolution* (EL). Second, we propose MVX-LLM, a token-efficient architecture that fuses multiple sensor modalities from **Multiple Views** effectively with a **Dual Cross-Attention (DCA)** projector. This architecture is designed to address abnormal conditions in autonomous driving scenarios through integrated abnormal condition modeling and cross-modal vision-language fusion. As demonstrated in Figure 1, our MVX-LLM achieves the highest GPT score of 55.5, significantly outperforming baseline methods without injected adverse condition knowledge. While Honeybee and ParGo [7, 43] represent conventional projector approaches that are not effective to deal with multiple modalities, and Self-query [48] incorporates abnormal handling but lacks cross-modal vision-language integration, our approach combines both capabilities of dealing with multiple modalities and comprehensive cross-modal vision-language integration. The main contributions of this work are:

- A comprehensive cross-modal driving VQA dataset DRIVEXQA considering adverse conditions including diverse weather variations and sensor failure scenarios, with a systematic hierarchical QA taxonomy covering global scene, allocentric, and ego-vehicle centric levels.
- A token-efficient multi-modal architecture MVX-LLM for robust sensor fusion under adverse conditions.
- Extensive evaluation demonstrating effectiveness under challenging weather conditions and sensor degradation scenarios.

2. Related Work

2.1. Autonomous Driving VQA Datasets and Benchmarks

Visual Question Answering (VQA) is a task that enables machines to comprehend visual content and provide natural language responses to structured queries. In the context of autonomous driving, VQA systems primarily serve as auxiliary reasoning components that support onboard decision-making by translating complex perceptual inputs into interpretable natural language representations, facilitating both system diagnostics and human oversight. Previous VQA datasets [2, 16, 19, 22] established foundational evaluation principles for visual reasoning, introducing compositional question structures and systematic evaluation of spatial and relational understanding. Building on these foundations, driving-specific VQA datasets have progressively expanded in scope and modality coverage. NuScenes-QA [36] pioneered driving-specific VQA with multi-modal inputs, including RGB and LiDAR for dynamic outdoor environments. NuScenes-MQA [21] introduced MarkupQA for unified evaluation, DriveLM [40] combined graph-based VQA for spatial reasoning, and MapLM [6] addressed traffic scene evaluation, with further specialized work on uncertainty [32], object-level reasoning [9], and 3D comprehension [11, 51]. NuPlanQA [34] presented large-scale multi-view driving scene understanding across multiple subtasks (*e.g.*, traffic light detection, spatial relations recognition, ego-vehicle maneuver prediction), but lacks depth cameras and systematic sensor failure coverage. LingoQA [31] introduced video-based driving VQA with temporal reasoning, but focuses exclusively on RGB inputs without complementary sensor modalities. Recent specialized datasets, including DRAMA [29], VLAAD [33], and additional work [10], have addressed hazard prediction, interaction reasoning, and cooperative driving.

Despite this progress, existing datasets suffer from significant limitations that hinder robust evaluation under real-world deployment conditions. VQA model performance evaluation studies [38, 39] identified limitations including insufficient spatial reasoning and environmental robustness, and multi-modal datasets [4, 41, 48] only partially address the need for rigorous sensor fusion evaluation. As summarized in Table 1, these works share four critical gaps: (1) lack of comprehensive multi-modal sensor coverage, particularly event cameras; (2) absence of sensor failure scenarios; (3) insufficient hierarchical question structures; and (4) limited scale for training and evaluation. Our DRIVEXQA is specifically designed to address all of these limitations.

Table 1. Comparison of autonomous driving VQA datasets. **MV**: Multi-View RGB cameras. **SF**: Sensor Failure scenarios. **#VQA**: Number of VQA pairs. **FT**: Fine-tuned Evaluation. **Hierarchical**: Hierarchical question structure. Δ means partially involved.

Dataset	multi-view	modalities	SF	#VQA	FT	Hierarchical
DRAMA [29]	×	1	×	17K	✓	×
VLAAD [33]	×	1	×	64K	✓	×
LingoQA [31]	×	1	×	419K	✓	×
NuPlanQA [34]	✓	1	×	1M	✓	Δ
NuScenes-QA [36]	✓	2	×	460K	✓	×
NuScenes-MQA [21]	✓	2	×	-	✓	×
DriveLM [40]	✓	2	×	-	×	Δ
DRIVEXQA (Ours)	✓	4	✓	102K	✓	✓

2.2. Multi-Modal Large Language Models for Sensor Fusion

The emergence of large language models has fundamentally transformed multi-modal understanding, with foundational works [1, 8, 25, 27] establishing vision-language integration paradigms. LLaVA [27] pioneered two-stage training with feature alignment pretraining and instruction tuning. Within vision-language models, efficient visual token management became critical. Honeybee [7] introduces locality-enhanced projectors for computational efficiency while preserving spatial context. ParGo [43] advances this with Partial-Global projectors integrating both global and partial visual representations to preserve fine-grained details while mitigating overemphasis on prominent regions. Parameter-efficient methods such as LoRA [17], QLoRA [13], and related strategies [30] enable practical fine-tuning of large-scale models under constrained budgets.

However, real-world autonomous driving requires multiple complementary sensors beyond vision-language pairs. Cross-modal alignment research emerged to handle semantic correspondence across sensor modalities through contrastive learning [15, 26, 49] and multi-scale integration [12], enabling consistent alignment across diverse spatial and temporal characteristics. Dynamic environmental conditions further complicate multi-sensor fusion. Condition-aware fusion strategies [3, 44–46] introduced adaptive mechanisms that adjust fusion weights based on environmental context and sensor reliability. Despite these advances, existing multi-modal LLM frameworks fundamentally cannot accept multiple complementary sensor modalities as input, making them inadequate for autonomous driving scenarios requiring robust performance under sensor failures and adverse conditions. Our work addresses this limitation by proposing MVX-LLM, a specialized architecture that fuses multiple sensor inputs for robust multi-sensor autonomous driving applications.

3. DRIVEXQA Dataset

3.1. Dataset Generation

Our DRIVEXQA dataset comprises 7,885 frames of driving scenes with 102,505 question-answer pairs, following a sys-




Figure 2. Statistics of DRIVEXQA dataset showing distribution of weather conditions (left) and distribution of sensor failure types (right): MB (Motion Blur), OE (Overexposure), UE (Underexposure), LJ (LiDAR Jitter), EL (Event Low-resolution).

tematic 80%/10%/10% train/validation/test split. Scenes are collected in CARLA (v0.9.14) [14] and sensor failures are injected via post-processing. QA pairs are generated by prompting GPT-4o with structured scene-metadata templates. Manual quality control covering a stratified 5% sample was then conducted to ensure answer quality, correcting weather misclassification and inaccurate counts, and validating ego-centric questions by cross-referencing LiDAR and multi-view RGB data. 13 QA pairs are generated for each driving scene according to our hierarchical framework (2 global scene level, 8 allocentric level, 3 ego-vehicle centric level), ensuring proportional representation across question categories. The question-answer pairs demonstrate balanced linguistic complexity with an average question length of 11.4 words and answer length of 12.9 words, reflecting concise yet information-dense annotations that support reliable automatic evaluation across diverse question types and difficulty levels. Fig. 2 presents the broad coverage of our dataset across environmental conditions and sensor degradation scenarios. The balanced weather distribution covers all realistic automotive scenarios, while the systematic inclusion of sensor failures enables evaluation of system resilience under realistic implementation cases with 35.5% corner cases. The dataset covers balanced weather conditions across Rainy (20.0%), Night (20.1%), Sunny (20.0%), Cloudy (19.9%), and Foggy (20.0%). Sensor failures in-

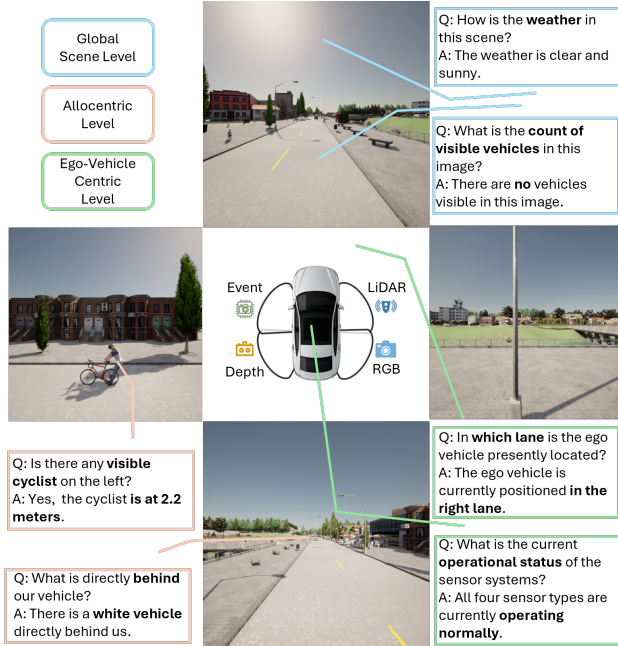


Figure 3. Hierarchical XQA examples on DRIVEXQA dataset. The framework demonstrates three semantic levels: *Global Scene Level*, *Allocentric Level*, and *Ego-Vehicle Centric Level*.

clude Motion Blur (42.7%), Underexposure (14.2%), Overexposure (14.3%), LiDAR Jitter (14.2%), and Event Low-resolution (14.6%), representing 35.5% of total samples as corner cases.

3.2. Hierarchical XQA

We develop a comprehensive XQA dataset structured across three receptive fields, *i.e.* global scene level, allocentric level, and ego-vehicle centric level. Exact 13 questions are generated from each driving scene following our hierarchical framework (2 global scene level, 8 allocentric level, 3 ego-vehicle centric level).

3.2.1. Global Scene Level

Global scene questions assess overall environmental conditions and traffic patterns that cannot be reliably inferred from any single sensor modality alone. These questions target weather-level perception (*e.g.*, visibility estimation, precipitation recognition) and macro-level traffic state assessment, requiring holistic fusion of multi-modal inputs to reason about scene-wide factors that directly influence driving strategy and system reliability. These questions focus on weather impact and traffic density, providing essential context for the decision-making of autonomous driving. This level considers broad environmental factors that influence overall driving strategy and system reliability.

3.2.2. Allocentric Level

Allocentric questions cover spatial relationships in the scene through structured queries, including quantitative spatial analysis, distance measurements, traffic sign interpretation, road condition assessment, and object categorization. This level leverages multi-view visual information from four camera perspectives (front, back, left, right) to enable comprehensive spatial reasoning and cross-view consistency validation (*e.g.*, “*What is the distance to the cyclist on the left?*”). Multi-view perspectives are crucial at this level because they offer complete scene coverage, enabling both inter-object relationship reasoning and spatial awareness for navigation planning.

3.2.3. Ego-Vehicle Centric Level

Egocentric questions focus on direct environment perception and vehicle state, including lane positioning analysis, sensor health monitoring, and surrounding vehicle behavior prediction. These queries address safety-critical aspects requiring a precise understanding of the ego-vehicle’s operational environment and immediate decision-making requirements.

Fig. 3 demonstrates our hierarchical QA structure through concrete examples. The global level question assesses weather impact on visibility, the allocentric level analyzes traffic density and spatial relationships, while the egocentric level focuses on lane positioning and operational status. This systematic approach ensures a comprehensive evaluation of autonomous driving capabilities across different analysis perspectives and environmental conditions.

4. MVX-LLM Framework

4.1. Overall Architecture

To solve the XQA task, we propose the MVX-LLM framework (Fig. 4) that processes multi-modal sensor inputs through specialized encoders: a shared CLIP-based [37] Vision Transformer for RGB, depth, and event camera data from four perspectives, and a PointNet++ [35] encoder for LiDAR point clouds. Due to the heterogeneous modalities, LiDAR data is processed separately through hierarchical set abstraction operations and projected to unified 512-dimensional representations without involving in the cross-modal fusion. Our framework employs a Dual Cross-Attention (DCA) mechanism for 2D modalities, specifically designed for robust multi-modal fusion when individual sensors experience degradation due to adverse weather or hardware failures.

4.2. Dual Cross-Attention Mechanism

Autonomous driving scenarios demand adaptive cross-modal fusion where individual sensors may experience degradation due to adverse weather or sensor failures.

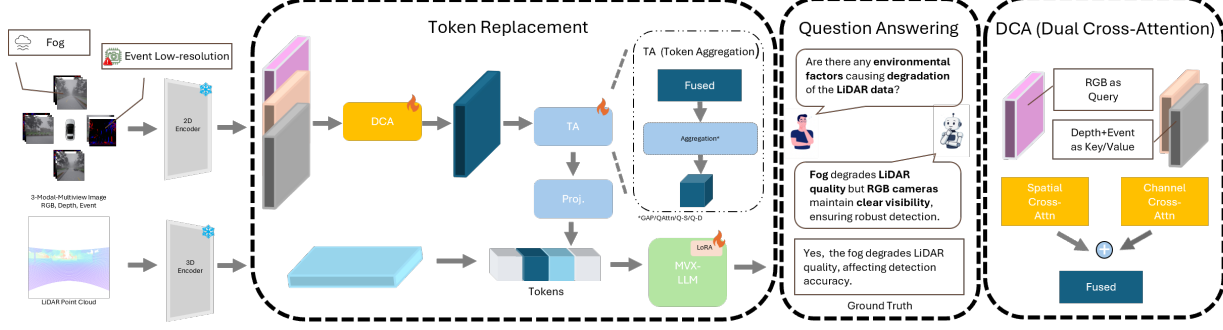


Figure 4. Overview of MVX-LLM. The framework processes multi-modal sensor inputs (RGB, Depth, Event cameras from four viewpoints, and LiDAR point clouds) through specialized encoders. The DCA mechanism integrates RGB, Depth, and Event features before token replacement. The Question Answering component utilizes the fused representations to generate hierarchical responses across Global Scene, Allocentric, and Ego-Vehicle Centric levels under adverse driving conditions.

Traditional concatenation-based approaches [18] treat all modalities equally, lacking the ability to adapt to sensor failures. DCA mechanism addresses this limitation by operating through complementary spatial and channel cross-attention pathways that enable adaptive feature weighting based on environmental conditions and sensor availability. All modalities undergo spatial alignment from a native 49-token representation (derived from the 77 spatial grid output of the CLIP-based ViT encoder) to a 48-token configuration for computational efficiency and attention head compatibility:

$$\mathbf{F}'_{rgb} = \mathbf{F}_{rgb} \mathbf{W}_{align}^{rgb} \in \mathbb{R}^{B \times 48 \times 512} \quad (1)$$

$$\mathbf{F}'_{depth} = \mathbf{F}_{depth} \mathbf{W}_{align}^{depth} \in \mathbb{R}^{B \times 48 \times 512} \quad (2)$$

$$\mathbf{F}'_{event} = \mathbf{F}_{event} \mathbf{W}_{align}^{event} \in \mathbb{R}^{B \times 48 \times 512} \quad (3)$$

The spatial cross-attention branch establishes spatial correspondences between modalities, enabling cross-modal information exchange at specific spatial locations:

$$\mathbf{F}_s = \text{MultiHeadAttn}_s(\mathbf{Q}_{rgb}, \mathbf{K}_{multi}, \mathbf{V}_{multi}) \quad (4)$$

where $\mathbf{Q}_{rgb} = \mathbf{F}'_{rgb} \in \mathbb{R}^{B \times 48 \times 512}$ and:

$$\mathbf{K}_{multi} = \mathbf{V}_{multi} = [\mathbf{F}'_{depth}; \mathbf{F}'_{event}] \in \mathbb{R}^{B \times 96 \times 512} \quad (5)$$

The channel cross-attention branch operates on transposed features to enable cross-channel feature enhancement. Concatenated multi-modal features require alignment to resolve dimensional incompatibility with RGB queries:

$$\mathbf{F}_c = \text{MultiHeadAttn}_c(\mathbf{Q}_{rgb}^T, \mathbf{K}'_{multi}, \mathbf{V}'_{multi}) \quad (6)$$

where:

$$\mathbf{Q}_{rgb}^T = (\mathbf{F}'_{rgb})^T \in \mathbb{R}^{B \times 512 \times 48} \quad (7)$$

$$\mathbf{K}'_{multi} = \mathbf{V}'_{multi} = [\mathbf{F}'_{depth}; \mathbf{F}'_{event}]^T \mathbf{W}_{align} \quad (8)$$

with learned transformation $\mathbf{W}_{align} \in \mathbb{R}^{96 \times 48}$ ensuring $\mathbf{K}'_{multi} \in \mathbb{R}^{B \times 512 \times 48}$. The final fused representation combines both pathways:

$$\mathbf{F}_{fused} = \frac{\mathbf{F}_s + \mathbf{F}_c^T}{2} \quad (9)$$

This dual-pathway design captures both spatial-level and channel-level cross-modal interactions, providing a comprehensive multi-modal fusion capability.

4.3. Token Aggregation

The token aggregation module processes the fused multi-modal features through specialized fusion heads that compress the visual tokens from multi-modalities into compact token embeddings suitable for language model integration. Following TokenFusion [44] principles, we compare four distinct aggregation strategies to address different aspects of spatial feature consolidation while maintaining computational efficiency.

4.3.1. Global Average Pooling (GAP)

The Global Average Pooling head provides the baseline aggregation mechanism through simple mean pooling across all spatial tokens:

$$\mathbf{F}_{gap} = \frac{1}{N} \sum_{i=1}^N \mathbf{F}_{fused}[i, :] \in \mathbb{R}^{B \times 512} \quad (10)$$

where $N = 48$ represents the number of spatial tokens. GAP serves as the computational efficiency baseline, requiring no additional parameters.

4.3.2. Query Attention (QAttn)

The Query Attention mechanism implements learnable query-based aggregation through multi-head cross-attention:

$$\mathbf{F}_{qattn} = \text{MultiHeadAttn}(\mathbf{Q}_{learn}, \mathbf{F}_{fused}, \mathbf{F}_{fused}) \quad (11)$$

where $\mathbf{Q}_{learn} \in \mathbb{R}^{1 \times 512}$ represents learnable query. The learnable queries enable the model to focus on task-relevant spatial regions during training, significantly improving performance over uniform pooling strategies.

4.3.3. Spectral QAttn

The Spectral QAttn variant incorporates frequency domain processing through depthwise convolution before query attention:

$$\mathbf{F}_{enhanced} = [\mathbf{F}_{fused}; \text{DWConv1D}(\mathbf{F}_{fused})] \quad (12)$$

$$\mathbf{F}_{spectral} = \text{MultiHeadAttn}(\mathbf{Q}_{learn}, \mathbf{F}_{enhanced}, \mathbf{F}_{enhanced}) \quad (13)$$

4.3.4. DepthGate QAttn

The DepthGate QAttn variant incorporates depth-guided confidence estimation inspired by DFormerv2 [46]:

$$\mathbf{G}_{depth}^{(i)} = \sigma(\mathbf{F}'_{depth}[i, :] \cdot \mathbf{w}_{gate} + b_{gate}) \quad (14)$$

$$\mathbf{F}_{depthgate} = \text{MultiHeadAttn}(\mathbf{Q}_{learn}, \mathbf{F}_{gated}, \mathbf{F}_{gated}) \quad (15)$$

Based on experimental validation, QAttn achieves the highest performance (GPTScore: 55.5) and serves as our final implementation choice.

5. Experiments

5.1. Implementation Details

DCA mechanism operates with 8 attention heads for spatial attention (64 dimensions each) and 4 heads for channel attention (12 dimensions each), ensuring balanced computational load while maintaining full utilization of the 512-dimensional feature space. Training is conducted using standard cross-entropy loss with Adam optimizer, learning rate of $1e-4$, and batch size of 16. We employ LoRA fine-tuning with rank 16 for parameter-efficient adaptation of the language model components. All experiments use 4 NVIDIA A100-40GB GPUs with mixed-precision training. The DCA projector adds $\sim 2.6M$ parameters ($< 0.3\%$ of total model size), and per-modality token compression to a single query token keeps visual token count constant regardless of modality number, preserving inference efficiency. For evaluation, we employ multiple metrics to assess different aspects of visual question answering performance. Following standard VQA evaluation protocols, we use BLEU-4, ROUGE-L, METEOR, and CIDEr for lexical and semantic similarity assessment. Additionally, we employ sentence embedding-based similarity scores. Following the LLM-as-a-judge paradigm [23, 24], we use GPT-4o-mini [20] for semantic correctness evaluation, providing human-like judgment on factual accuracy and contextual

appropriateness. GPT-based evaluation uses a 1~5 rubric normalized to 0~100 via:

$$S = \frac{1}{N_{samples}} \sum_{i=1}^{N_{samples}} \frac{s_i - 1}{4} \times 100\%, \quad (16)$$

where s_i represents the individual GPT score for sample i , and $N_{samples}$ is the total number of evaluated samples. To assess the reliability of GPT-based evaluation for DriveXQA, we recruited four human evaluators who independently assessed a subset of 360 samples across different weather conditions and sensor failures.

5.2. Results on DRIVEXQA

Table 2 presents comprehensive comparisons across different architectural approaches and fusion strategies on the DRIVEXQA dataset. We evaluate both projector architectures and attention-based fusion mechanisms under various environmental conditions and sensor failure scenarios. Human evaluation on a subset of samples shows strong correlation with GPT-4o-mini scores across all methods and conditions, with Spearman correlation coefficients [28] consistently above 0.83, confirming the reliability of our automated evaluation approach. MVX-LLM framework with DCA mechanism significantly outperforms both token fusion methods and traditional projector architectures. Specifically, our full system (Ours) achieves the highest average GPTScore of 55.5, surpassing the best token fusion method QAttn(Spectral) (52.8) by 3% and the strongest projector architecture GAP + Honeybee (48.2) by 7.3%. Notably, our approach shows remarkable improvements in challenging conditions such as foggy (53.5 vs. 25.1 for prepend, 53.7 for QAttn(Spectral) token fusion, and 50.8 for GAP + Honeybee projector) and various sensor failure scenarios. The consistent performance gains across all environmental conditions and sensor degradation types highlight the effectiveness of our DCA mechanism and QAttn for robust multi-modal fusion, demonstrating clear advantages over both conventional token-level fusion strategies and projector-based integration approaches.

5.3. Ablation Study

5.3.1. Attention Component Ablation Analysis

Table 3 examines the contributions of spatial cross-attention (sAttn), channel cross-attention (cAttn), and Query Attention (QAttn) for multi-modal fusion. Individual mechanisms show limited performance, with channel attention alone achieving only (GPTScore: 20.4) and spatial attention alone reaching (GPTScore: 24.0), indicating that neither mechanism is sufficient for effective multi-modal fusion independently. Dual attention combinations demonstrate significant improvements. The cAttn+QAttn combination achieves (GPTScore: 34.8), while sAttn+QAttn reaches

Table 2. Results on DRIVEXQA dataset across weather conditions and sensor failures. Each cell shows GPT-4o-mini score/Human evaluation score. **Blue** indicates state-of-the-art results. Weather conditions: CL (Cloudy), FG (Foggy), NT (Night), RN (Rainy), SN (Sunny). Sensor failures: MB (Motion Blur), OE (Overexposure), UE (Underexposure), LJ (LiDAR Jitter), EL (Event Low-resolution).

Method	CL	FG	NT	RN	SN	MB	OE	UE	LJ	EL	Average
prepend	61.2/58.2	25.1/26.3	58.8/60.1	52.3/50.5	55.1/58.3	53.1/57.4	44.1/44.7	51.5/49.2	60.5/62.5	55.2/58.1	55.1/57.9
<i>Token Fusion</i>											
Self-query	40.9/43.4	40.3/42.7	34.9/37.6	41.1/43.5	48.1/46.3	39.9/36.7	28.9/28.4	36.1/39.8	39.5/42.5	44.1/47.6	40.6/43.3
GAP	51.2/54.5	53.2/56.4	50.5/53.6	50.2/47.4	50.4/53.8	54.2/57.5	49.4/52.7	51.1/48.4	54.9/57.2	54.9/57.3	50.5/53.1
QAttn(Spectral)	56.5/58.7	53.7/55.4	46.8/47.9	53.0/56.3	53.2/54.5	54.3/57.7	46.0/44.9	51.6/54.4	55.2/57.3	55.3/57.2	52.8/55.9
QAttn(DepthGate)	38.3/40.2	40.1/43.6	33.0/34.2	41.3/44.6	50.0/53.4	41.1/43.4	27.5/24.8	48.3/51.9	43.9/45.4	40.3/44.6	40.0/44.1
<i>Projector Architectures</i>											
GAP + Honeybee	48.7/50.6	50.8/53.1	42.7/44.8	48.7/50.3	49.2/47.1	49.7/53.4	40.8/43.7	42.7/44.5	51.3/54.2	48.5/46.1	48.2/51.9
GAP + Pargo	36.5/33.1	28.7/25.4	34.6/37.5	28.3/31.5	44.6/47.4	31.2/28.4	23.3/19.9	22.8/20.6	28.7/25.0	29.9/29.4	33.2/35.1
MVX-LLM (Ours)	59.6/ 61.0	53.5/55.7	51.7/53.6	55.2/58.1	58.3/61.2	55.2/58.2	51.3/53.9	57.3/59.4	57.4/60.3	52.8/55.6	55.5/58.1

(GPTScore: 41.1), showing that spatial attention provides more effective cross-modal alignment than channel attention when combined with learnable query mechanisms. The sAttn+cAttn combination performs best among dual mechanisms (GPTScore: 50.5), verifying the superiority of our dual cross-attention design. The complete triple attention mechanism (sAttn+cAttn+QAttn) achieves optimal performance (GPTScore: 55.5), demonstrating the complementary effects of all components: spatial cross-attention enables position-level cross-modal alignment, channel cross-attention facilitates semantic information exchange, and QAttn provides adaptive spatial attention for task-relevant focus in autonomous driving scenarios.

Table 3. Ablation study on attention mechanism components. Results demonstrate the individual and combined contributions of spatial cross-attention (sAttn), channel cross-attention (cAttn), and Query Attention (QAttn) for robust multi-modal fusion.

sAttn	cAttn	QAttn	CIDEr	BLEU-4	ROUGE-L	METEOR	Sim	GPTScore
✓			48.3	9.4	23.5	29.8	49.1	24.0
✓		✓	50.6	15.7	33.2	41.4	61.1	41.1
	✓		31.4	6.5	19.7	26.6	46.8	20.4
	✓	✓	42.9	11.3	27.0	35.6	57.0	34.8
✓	✓		105.7	24.6	46.0	56.6	70.5	50.5
✓	✓	✓	144.6	28.1	50.6	61.2	74.6	55.5

5.3.2. Multi-Modal Sensor Combination Analysis

To systematically evaluate the contribution of different sensor modality combinations in MVX-LLM framework, we conducted a comprehensive ablation study using the QAttn token aggregation mechanism. We examined various multi-modal combinations to understand their complementary effects and potential interference patterns. Table 4 presents the detailed results across multiple evaluation metrics. Two-modality combinations reveal distinct patterns in sensor complementarity. RGB+Depth achieves moderate performance (GPTScore: 23.4), while

RGB+Event demonstrates stronger complementary effects (GPTScore: 42.4). Notably, the Depth+Event combination without RGB guidance performs poorly (GPTScore: 18.9), highlighting the fundamental importance of RGB visual information as the foundation for multi-modal fusion in autonomous driving scenarios. Three-modality combinations show more complex interaction dynamics. RGB+Depth+Event achieves substantial improvement (GPTScore: 46.0), indicating that when properly grounded with RGB, multiple visual modalities can work synergistically. RGB+Depth+LiDAR (GPTScore: 50.7) demonstrates the strongest three-modality performance through effective integration of visual and geometric information, while RGB+Event+LiDAR (GPTScore: 34.2) shows more modest gains. Interestingly, Depth+Event+LiDAR without RGB guidance achieves moderate performance (GPTScore: 39.7), suggesting that geometric information from LiDAR can partially compensate for the absence of RGB visual context, though not as effectively as RGB-grounded combinations. Our complete four-modality MVX-LLM system achieves optimal performance (GPTScore: 55.5), verifying that the QAttn mechanism effectively leverages all complementary sensor information while maintaining robust fusion across diverse modalities.

5.4. Qualitative Analysis

To further illustrate the effectiveness of our multi-modal fusion approach, we analyze specific prediction cases that demonstrate the complementary nature of different sensor modalities under adverse conditions. Fig. 5 presents a representative example from the DRIVEXQA dataset showing the model’s performance under combined night conditions and camera overexposure. In the upper case, when asked about objects directly in front of the ego vehicle, our model correctly identifies the red car despite the challenging lighting and overexposure conditions. This success demonstrates the strength of our multi-modal fusion:

Table 4. Ablation study on sensor modality combinations. Results show the importance of each modality and its complementary effects for robust VQA performance under adverse conditions.

RGB	Depth	Event	LiDAR	CIDEr	BLEU-4	ROUGE-L	METEOR	Sim	GPTScore
✓	✓			163.4	19.5	31.6	38.1	53.4	23.4
✓		✓		152.6	26.1	43.4	48.2	63.1	42.4
	✓	✓		8.1	3.9	16.9	23.0	43.5	18.9
✓	✓	✓		157.6	27.3	46.2	51.7	68.7	46.0
✓	✓		✓	117.5	26.0	46.7	54.1	71.5	50.7
✓		✓	✓	50.1	13.4	30.0	37.6	57.3	34.2
	✓	✓	✓	77.4	17.7	35.3	43.2	62.9	39.7
✓	✓	✓	✓	144.6	28.1	50.6	61.2	74.6	55.5

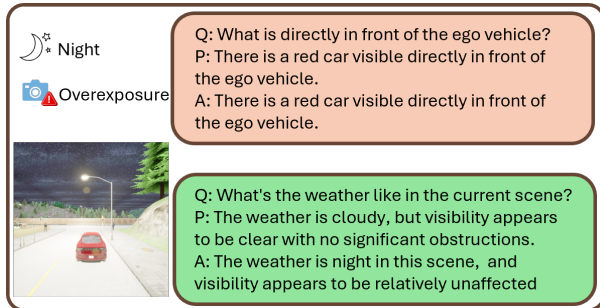


Figure 5. Qualitative analysis of multi-modal fusion performance under night conditions with camera overexposure.

while the RGB camera suffers from overexposure artifacts, complementary modalities such as LiDAR provide robust geometric information about vehicle positions and shapes. The DCA mechanism appears to leverage depth and point cloud data for object detection even when visual sensors are compromised, consistent with strong RGB+Depth+LiDAR performance (GPTScore: 50.7) in our quantitative results. Conversely, the weather recognition error suggests potential limitations when geometric modalities must compensate for compromised visual cues. This aligns with our ablation findings showing performance variations across different modality combinations. This analysis reveals that geometric sensors effectively compensate for object detection tasks, but semantic environmental understanding remains critically dependent on functional visual sensors.

5.5. Implications for Sim-to-Real Transfer

Despite the sim-to-real gap, our benchmark provides a controlled stress-test environment by systematically modeling adverse weather and sensor degradations (*e.g.*, motion blur, over/under-exposure, LiDAR jitter, and low-resolution event streams). This setup mirrors common deployment-time failures and exposes multimodal fusion bottlenecks under safety-critical conditions. As a result, it supports more reliable method comparison and helps validate robustness

improvements before transitioning from simulation to real-world autonomous driving.

6. Conclusion

In this work, we introduce the DRIVEXQA dataset and MVX-LLM framework for multi-modal autonomous driving scene understanding under adverse conditions. Our dataset evaluates three semantic levels: global environmental assessment, allocentric spatial reasoning, and ego-vehicle operational analysis. MVX-LLM addresses the limitation that existing MLLMs cannot process multiple complementary modalities by introducing a Dual Cross-Attention mechanism that fuses RGB, depth, event camera, and LiDAR data. Our method achieves GPTScore 55.5, demonstrating effective multi-modal fusion for challenging driving scenarios.

6.1. Limitations and Future Work

While our approach demonstrates promising results, several limitations remain. First, DRIVEXQA is built on CARLA simulation data; while this enables systematic control over adverse conditions, physical phenomena such as LiDAR scattering in heavy rain or DVS noise in low-light are only approximated in simulation, and validation on real-world corrupted sensor data remains important future work. Second, the benchmark is limited to single-frame understanding, treating scenes independently without temporal context; we adopt this static formulation as a controlled baseline to isolate cross-modal spatial reasoning, with video-based temporal extension planned as direct future work. Third, our modality-combination ablations provide evidence for cross-modal complementarity under adverse conditions; explicit sensor-dropout evaluation at inference time remains a valuable direction to further substantiate compensation dynamics. Future work will incorporate real-world data, temporal reasoning, and lightweight onboard deployment.

Acknowledgement

This work was supported in part by National Natural Science Foundation of China under Grant No. 62503166 and No. 62473139, in part by the Hunan Provincial Research and Development Project (Grant No. 2025QK3019), in part by the State Key Laboratory of Autonomous Intelligent Unmanned Systems (the opening project number ZZKF2025-2-10), in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - SFB 1574 - 471687386, and in part by Helmholtz Association of German Research Centers, in part by the Ministry of Science, Research and the Arts of Baden-Württemberg (MWK) through the Cooperative Graduate School Accessibility through AI-based Assistive Technology (KATE) under Grant BW6-03, and in part by the Helmholtz Association Initiative and Networking Fund on the HAICORE@KIT and HOREKA@KIT partition. This research was partially funded by the Ministry of Education and Science of Bulgaria (support for INSAIT, part of the Bulgarian National Roadmap for Research Infrastructure).

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 35: 23716–23736, 2022. 1, 3
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, pages 2425–2433, 2015. 2
- [3] Tim Brödermann, Christos Sakaridis, Yuqian Fu, and Luc Van Gool. Cafuser: Condition-aware multimodal fusion for robust semantic perception of driving scenes. *IEEE Robotics and Automation Letters*, 10(4):3134–3141, 2025. 3
- [4] 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 multimodal dataset for autonomous driving. In *CVPR*, pages 11621–11631, 2020. 2
- [5] Davide Caffagni, Federico Cocchi, Luca Barsellotti, Nicholas Moratelli, Sara Sarto, Lorenzo Baraldi, Marcella Cornia, and Rita Cucchiara. The revolution of multimodal large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13590–13618, 2024. 1
- [6] Xu Cao, Tong Zhou, Yunsheng Ma, Wenqian Ye, Can Cui, Kun Tang, Zhipeng Cao, Kaizhao Liang, Ziran Wang, James M Rehg, et al. Maplm: A real-world large-scale vision-language benchmark for map and traffic scene understanding. In *CVPR*, pages 21819–21830, 2024. 2
- [7] Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. In *CVPR*, pages 13817–13827, 2024. 2, 3
- [8] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. In *ECCV*, pages 370–387, 2024. 3
- [9] Long Chen, Oleg Sinavski, Jan Hünermann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. Driving with llms: Fusing object-level vector modality for explainable autonomous driving. In *ICRA*, pages 14093–14100, 2024. 2
- [10] Hsu-kuang Chiu, Ryo Hachiuma, Chien-Yi Wang, Stephen F Smith, Yu-Chiang Frank Wang, and Min-Hung Chen. V2v-llm: Vehicle-to-vehicle cooperative autonomous driving with multi-modal large language models. *arXiv preprint arXiv:2502.09980*, 2025. 2
- [11] Tushar Choudhary, Vikrant Dewangan, Shivam Chandhok, Shubham Priyadarshan, Anushka Jain, Arun K Singh, Siddharth Srivastava, Krishna Murthy Jatavallabhula, and K Madhava Krishna. Talk2bev: Language-enhanced bird’s-eye view maps for autonomous driving. In *ICRA*, pages 16345–16352, 2024. 2
- [12] Giordano Cichetti, Eleonora Grassucci, Luigi Sigillo, and Danilo Comminiello. Gramian multimodal representation learning and alignment. In *ICLR*, 2025. 3
- [13] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *NeurIPS*, 36:10088–10115, 2023. 3
- [14] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator, 2017. 3
- [15] Benoit Dufumier, Javiera Castillo-Navarro, Devis Tuia, and Jean-Philippe Thiran. What to align in multimodal contrastive learning? In *ICLR*, 2025. 3
- [16] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, pages 6904–6913, 2017. 2
- [17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. 3
- [18] Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. An embodied generalist agent in 3d world. In *ICML*, 2024. 5
- [19] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, pages 6700–6709, 2019. 2
- [20] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 6
- [21] Yuichi Inoue, Yuki Yada, Kotaro Tanahashi, and Yu Yamaguchi. Nuscenesc-mqa: Integrated evaluation of captions and qa for autonomous driving datasets using markup annotations. In *WACV*, pages 930–938, 2024. 2, 3
- [22] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross

- Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, pages 2901–2910, 2017. 2
- [23] Dawei Li, Renliang Sun, Yue Huang, Ming Zhong, Bohan Jiang, Jiawei Han, Xiangliang Zhang, Wei Wang, and Huan Liu. Preference leakage: A contamination problem in llm-as-a-judge. In *ICML*, 2025. 6
- [24] Haitao Li, Qian Dong, Junjie Chen, Huixue Su, Yujia Zhou, Qingyao Ai, Ziyi Ye, and Yiqun Liu. Llm-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579*, 2024. 6
- [25] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, pages 19730–19742, 2023. 3
- [26] Zhenxi Lin, Ziheng Zhang, Meng Wang, Yinghui Shi, Xian Wu, and Yefeng Zheng. Multi-modal contrastive representation learning for entity alignment. In *COLING*, 2022. 3
- [27] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 1, 3
- [28] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, et al. Openeqa: Embodied question answering in the era of foundation models. In *CVPR*, pages 16488–16498, 2024. 6
- [29] Srikanth Malla, Chiho Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. Drama: Joint risk localization and captioning in driving. In *WACV*, pages 1043–1052, 2023. 2, 3
- [30] Sourab Mangrulkar, Sylvain Gugger, Lysandre Debut, Younes Belkada, Sayak Paul, and Benjamin Bossan. Peft: State-of-the-art parameter-efficient fine-tuning methods, 2022. 3
- [31] Ana-Maria Marcu, Long Chen, Jan Hünemann, Alice Karnsund, Benoit Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, Alex Kendall, Jamie Shotton, et al. Lingoqa: Visual question answering for autonomous driving. In *ECCV*, pages 252–269, 2024. 1, 2, 3
- [32] Tinghui Ouyang, Yoshinao Isobe, Saima Sultana, Yoshiki Seo, and Yutaka Oiwa. Autonomous driving quality assurance with data uncertainty analysis. In *IJCNN*, pages 1–7. IEEE, 2022. 2
- [33] SungYeon Park, MinJae Lee, JiHyuk Kang, Hahyeon Choi, Yoonah Park, Juhwan Cho, Adam Lee, and DongKyu Kim. Vlaad: Vision and language assistant for autonomous driving. In *WACV*, pages 980–987, 2024. 2, 3
- [34] Sung-Yeon Park, Can Cui, Yunsheng Ma, Ahmadreza Moradipari, Rohit Gupta, Kyungtae Han, and Ziran Wang. Nuplanqa: A large-scale dataset and benchmark for multi-view driving scene understanding in multi-modal large language models. In *ICCV*, 2025. 2, 3
- [35] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017. 4
- [36] Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. In *AAAI*, pages 4542–4550, 2024. 1, 2, 3
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 4
- [38] Kaavya Rekanar, Ciarán Eising, Ganesh Sistu, and Martin Hayes. Towards a performance analysis on pre-trained visual question answering models for autonomous driving. *arXiv preprint arXiv:2307.09329*, 2023. 2
- [39] Kaavya Rekanar, Abbirah Ahmed, Reenu Mohandas, Ganesh Sistu, Ciarán Eising, and Martin Hayes. Subjective scoring framework for vqa models in autonomous driving. *IEEE Access*, 12:141306–141323, 2024. 2
- [40] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Jens Beißwenger, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. In *ECCV*, pages 256–274, 2024. 1, 2, 3
- [41] 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 *CVPR*, pages 2446–2454, 2020. 2
- [42] Georgios Tzifas and Hamidreza Kasaei. Early or late fusion matters: Efficient rgb-d fusion in vision transformers for 3d object recognition. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9558–9565. IEEE, 2023. 2
- [43] An-Lan Wang, Bin Shan, Wei Shi, Kun-Yu Lin, Xiang Fei, Guozhi Tang, Lei Liao, Jingqun Tang, Can Huang, and Wei-Shi Zheng. Pargo: Bridging vision-language with partial and global views. In *AAAI*, 2025. 2, 3
- [44] Yikai Wang, Xinghao Chen, Lele Cao, Wenbing Huang, Fuchun Sun, and Yunhe Wang. Multimodal token fusion for vision transformers. In *CVPR*, pages 12186–12195, 2022. 3, 5
- [45] Bowen Yin, Xuying Zhang, Zhongyu Li, Li Liu, Ming-Ming Cheng, and Qibin Hou. Dformer: Rethinking rgb-d representation learning for semantic segmentation. In *ICLR*, 2024.
- [46] Bo-Wen Yin, Jiao-Long Cao, Ming-Ming Cheng, and Qibin Hou. Dformerv2: Geometry self-attention for rgb-d semantic segmentation. In *CVPR*, pages 19345–19355, 2025. 3, 6
- [47] Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *National Science Review*, 11(12): nwa403, 2024. 1
- [48] Jiaming Zhang, Ruiping Liu, Hao Shi, Kailun Yang, Simon Reiß, Kunyu Peng, Haodong Fu, Kaiwei Wang, and Rainer Stiefelhagen. Delivering arbitrary-modal semantic segmentation. In *CVPR*, pages 1136–1147, 2023. 2
- [49] Ming Zhang, Ke Chang, and Yunfang Wu. Multi-modal semantic understanding with contrastive cross-modal feature alignment. In *LREC-COLING*, 2024. 3
- [50] Xue Zhang, Si-Yuan Cao, Fang Wang, Runmin Zhang, Zhe Wu, Xiaohan Zhang, Xiaokai Bai, and Hui-Liang Shen. Rethinking early-fusion strategies for improved multispectral object detection. *IEEE Transactions on Intelligent Vehicles*, 2024. 2

[51] Chenming Zhu, Tai Wang, Wenwei Zhang, Jiangmiao Pang, and Xihui Liu. Llava-3d: A simple yet effective pathway to empowering lmms with 3d-awareness. In *ICCV*, 2025. [2](#)