

Box-QAymo: Box-Referring VQA Dataset for Autonomous Driving

Djamahl Etchegaray
The University of Queensland
Brisbane, Australia
uqdetche@uq.edu.au

Yuxia Fu
The University of Queensland
Brisbane, Australia
yuxia.fu@uq.edu.au

Zi Huang
The University of Queensland
Brisbane, Australia
huang@itee.uq.edu.au

Yadan Luo
The University of Queensland
Brisbane, Australia
y.luo@uq.edu.au

Abstract

Interpretable communication is essential for safe and trustworthy autonomous driving, yet current vision-language models (VLMs) often operate under idealized assumptions and struggle to capture user intent in real-world scenarios. Existing driving-oriented VQA datasets are limited to full-scene descriptions or waypoint prediction, preventing the assessment of whether VLMs can respond to localized user-driven queries. We introduce Box-QAymo, a box-referring dataset and benchmark designed to both evaluate and finetune VLMs on spatial and temporal reasoning over user-specified objects. Users express intent by drawing bounding boxes in complex scenes. Specifically, we propose a hierarchical evaluation protocol that begins with binary sanity-check questions to assess basic model capacities, and progresses to (1) attribute prediction for box-referred objects, (2) motion understanding of target instances, and (3) spatiotemporal motion reasoning over inter-object dynamics across frames. To support this, we crowd-sourced fine-grained object classes and visual attributes that reflect the complexity drivers encounter, and extract object trajectories to construct temporally grounded QA pairs. Rigorous quality control through negative sampling, temporal consistency checks, and difficulty-aware balancing guarantee dataset robustness and diversity. Our comprehensive evaluation reveals significant limitations in current VLMs when queried about perception questions, highlighting the gap in achieving real-world performance. This work provides a foundation for developing more robust and interpretable autonomous driving systems that can communicate effectively with users under real-world conditions. Project page and dataset are available at <https://djamahl99.github.io/qaymo-pages/>.

CCS Concepts

• **Computing methodologies** → **Scene understanding; Object recognition.**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXXX.XXXXXXX>

Keywords

Autonomous Driving, Vision-Language Models, Visual Question Answering.

ACM Reference Format:

Djamahl Etchegaray, Yuxia Fu, Zi Huang, and Yadan Luo. 2018. Box-QAymo: Box-Referring VQA Dataset for Autonomous Driving. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 Introduction

Safety-critical autonomous driving systems depend on effective human-AI communication to foster trust and support widespread adoption. However, recent high-profile incidents involving autonomous vehicles [3] have exposed critical gaps in user interaction. Studies show that 60% of consumers remain hesitant to adopt fully autonomous vehicles [1, 2], with primary concerns reflecting issues with system transparency. These findings underscore the need for frameworks to assess how effectively systems can communicate their understanding and intentions to human users.

Despite this need, progress in evaluating human-AI communication has been limited by the concentration of domain-specific data. Researchers lack access to semantically rich data required to develop and benchmark communication frameworks. Major datasets like Waymo Open [25], nuScenes [5] and KITTI [10] represent millions of dollars in labeling, however, these efforts are constrained by limited class vocabularies. The result is that academic progress in evaluating human-AI communication for autonomous driving lags behind perception capabilities, despite being necessary to enable adoption.

Existing vision-language models adapted for autonomous driving focus predominantly on motion planning while neglecting the equally critical challenge of communicating perceptual understanding to users. Recent works like Drive-LM [26] and Talk2Drive [8] demonstrate impressive open-source datasets that integrate LLMs into the planning stage with user intentions. Commercially, Wayve has released LINGO [18] and Waymo has released EMMA models that can handle many tasks, including 3D object detection, planning, and reasoning. All of these models require substantial data for training and evaluation to effectively encompass many different scenarios. This creates a significant disconnect between research

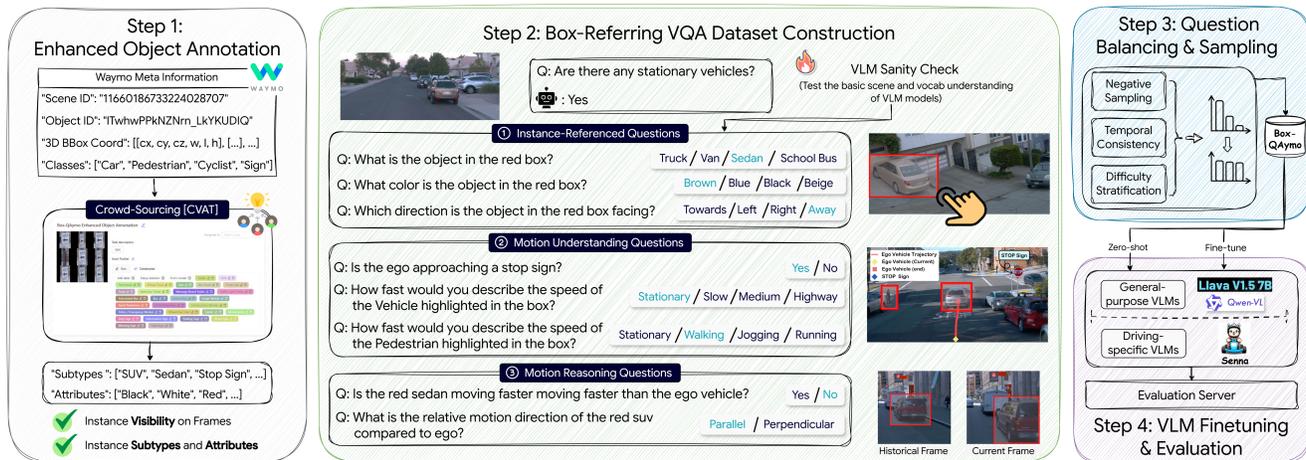


Figure 1: Overview of the Box-QAymo dataset pipeline for evaluating vision-language models (VLMs). Step 1 extracts 3D metadata from Waymo which is enhanced with human-annotated semantics. Step 2 introduces box-referenced visual question answering (VQA) tasks spanning instance recognition, motion interpretation, and temporal trajectory reasoning. Step 3 implements rigorous quality control through negative sampling, temporal consistency filtering, and difficulty-aware balancing to ensure a robust and challenging dataset. Step 4 benchmarks general and domain-specific VLMs in zero-shot and fine-tuned settings.

evaluation and real-world deployment scenarios. In practice, autonomous systems must operate with uncertain, incomplete perceptions while simultaneously communicating their confidence and reasoning to users. Current evaluation frameworks fail to capture this dual challenge, leading to models that can plan effectively but cannot explain their limitations or reasoning processes. This limitation becomes particularly problematic when these models encounter edge cases or unusual scenarios in which clear communication of uncertainty becomes critical to maintaining user trust and safety.

The limited semantic diversity in existing autonomous driving datasets constrains the development of models capable of communicating about complex driving scenarios. The vocabularies of popular datasets are restricted to basic objects, such as, vehicles, pedestrians, and cyclists. These vocabulary constraints force models to communicate about complex, nuanced scenarios using overly simplified categorical descriptions. A “vehicle” classification cannot convey the difference between a parked truck, a merging sedan, or an emergency vehicle, distinctions that are crucial for appropriate user communication and trust. This limited object vocabulary directly undermines the development of communication systems that can provide users with the detailed, contextual understanding necessary for informed decision-making in collaborative driving scenarios.

Existing benchmarks either assume unrealistic access to perfect scene understanding or limit evaluation to oversimplified object categories that fail to capture the semantic richness required for effective human communication. To our knowledge, existing datasets do not provide hierarchical evaluation that can assess model understanding across the spectrum from basic perception to complex temporal reasoning about dynamic scenes. This evaluation gap means that researchers cannot adequately measure the progress toward communication systems that would actually be deployable

in real-world scenarios. Without benchmarks that reflect true perceptual constraints and semantic complexity, the field lacks the tools necessary to develop reliable, communicative autonomous systems. Addressing this evaluation gap requires datasets that combine realistic perceptual constraints with rich semantic annotations and structured assessment frameworks.

In this paper, we introduce a novel dataset and evaluation framework (Box-QAymo) that addresses these limitations through crowd-sourced semantic diversity, temporal reasoning assessment, and hierarchical evaluation structures. Our approach expands semantic diversity through crowd-sourced fine-grained class labeling, moving beyond the 3-10 class limitations of existing datasets. We incorporate temporal reasoning evaluation to assess whether multimodal language models can understand object trajectories from visual information alone. Our hierarchical question format provides structured evaluation from binary classification through attribute prediction to complex motion understanding. Our dataset comprises 1,662 binary, 5,403 attribute, and 13,714 motion Q&A pairs, respectively. By combining semantic richness with temporal understanding and structured evaluation, we can identify specific limitations in current models while providing a pathway for systematic improvement. Our evaluation of general-purpose VLMs and Senna demonstrates the effectiveness of this approach, revealing systematic performance degradation from binary to more complex reasoning tasks, with substantial improvements achievable through targeted finetuning. Our framework establishes a foundation for developing autonomous driving systems that can communicate effectively with users, ultimately supporting the trust and adoption necessary for widespread deployment of autonomous vehicles.

2 Related Work

Recent efforts in building Visual Question Answering (VQA) datasets for autonomous driving have primarily focused on three core components: perception, prediction, and planning. Perception-focused questions commonly cover environmental elements and conditions like traffic signals [7, 12, 20, 24, 26, 27]. The questions range from closed-form answers [26] to free-form scene descriptions [12, 27]. Prediction-related questions often emphasise the behavioural intentions of surrounding objects to ensure safe path planning for the ego vehicle [12, 17, 19, 20, 24]. Moreover, some datasets incorporate object recognition tasks, involving both static (e.g., type, colour) [19, 20, 24, 30] and dynamic (e.g., speed) [12, 17, 24] attributes, as well as counting and identifying vulnerable or critical objects [12, 17, 24, 26]. Some datasets require reasoning about spatial relationships between objects [20, 21, 24, 30, 32] and ranking their importance [17, 22]. Planning-related questions can be grouped into meta-actions, justifications, waypoints, attention cues, and user instructions. Meta-actions [7, 12, 19, 24, 26, 27, 32] and waypoints [26] correspond to future ego-vehicle trajectories, while justification [14, 28, 31, 32] and attention cues [13, 19, 24] are designed to enhance the model’s reasoning and chain-of-thought [26] capabilities. To facilitate human-AI collaboration and promote safe decision-making, some datasets [9, 13, 16] introduce user instructions, while others [19, 27] offer counterfactual reasoning annotations to expose unsafe behaviours.

Although several datasets address either object-centric questions or user-instructed interactions, none have unified both in a way that supports explicit referential grounding. To fill this gap, we propose Box-QAymo, a dataset where users highlight specific objects via bounding boxes and pose instance-grounded, spatiotemporal questions. This design enables a more intuitive form of human-AI interaction and allows for direct evaluation of a model’s referential understanding through box-referenced queries.

3 Dataset Overview

Our dataset construction methodology follows a three-stage approach as seen in Figure 1: (1) enhanced object annotation using crowd-sourced labeling, (2) box-referring VQA dataset generation based on object characteristics and temporal dynamics, and (3) sampling and balancing strategies for VQA distribution and difficulty control.

3.1 Enhanced Object Annotation

Base Dataset Selection. We build upon the Waymo Open Dataset (WOD) [25], which provides high-quality 3D annotations for autonomous driving scenarios. We split the WOD validation set evenly into finetuning/validation splits with 101 scenes each. While previous works have opted for nuScenes [6] for its existing class diversity, we focus on WOD for its superior scene diversity and LiDAR density. While WOD offers robust geometric annotations, its semantic labeling is limited to coarse categories (*i.e.*, pedestrians, cyclists, vehicles, and signs). To enable fine-grained visual question answering, we augment these annotations with detailed semantic labels through crowdsourcing.

Crowdsourced Semantic Labeling. Labeling is necessary to enrich objects with enough semantic description for a VLM to single

them out. To accomplish this, we employ a crowdsourced annotation strategy to obtain fine-grained semantic labels for approximately 50% of objects in our dataset, skipping objects with poor visibility. The taxonomy established in Argoverse 2.0 [29] was followed, providing significantly richer semantic categories. For vehicle instances, annotators also provide color labels to enable color-based reasoning questions.

The annotation pipeline is as follows:

- (1) Objects are exported as crops with 3D bounding boxes overlaid for each camera and timestamp they are visible.
- (2) We collect up to 9 best crops based on the visibility of the object in each crop.
- (3) We fuse these crops as a 3x3 gallery, giving the annotators multiple views to help them with occluded and distant objects.
- (4) The object galleries are collected and uploaded to CVAT.
- (5) Annotators label instances iteratively and ignore those that are not visible.

3.2 Box-Referring VQA Dataset Construction

Table 1: Overview of question subcategories and examples in the Box-QAymo dataset. The taxonomy follows a hierarchical structure, starting from basic binary VLM sanity checks and extending to fine-grained object classification and complex spatiotemporal reasoning tasks. This design enables systematic evaluation of VLM capabilities across varying levels of visual understanding.

Subcategory	Example Questions
Movement Status	"Are there any stationary vehicles?"
Orientation	"Are there any vehicles moving towards the camera?"
Fine-grained Classification	"What type of object is in the red box?"
Color Recognition	"What color is the object highlighted in red?"
Facing Direction	"What direction is the object in the red box facing?"
Speed Assessment	"How fast is the blue sedan moving?"
Movement Direction	"What direction is the object in the red box moving?"
Relative Motion Analysis	"Is the green pickup truck traveling faster than the ego vehicle?"
Traffic Element Recognition	"Is the ego vehicle approaching a stop sign?"
Trajectory Analysis	"Are the ego vehicle and the truck on a collision course?"
Relative Motion Direction	"What is the relative motion direction of the hatchback compared to ego?"
Path Conflict Detection	"Is there a vehicle in the ego vehicle’s future path?"

Our question generation follows a principled hierarchical structure, progressing from simple binary queries to complex spatiotemporal reasoning tasks. This design enables systematic evaluation of model capabilities spanning different levels of visual understanding.

3.2.1 VLM Sanity Check We generate binary (yes/no) questions about object presence and basic characteristics, without referring to bounding boxes. These questions are designed to test the common sense and fundamental scene understanding capabilities of VLMs. We use these simple questions as a sanity check to filter out models that lack basic visual-language alignment. Our binary questions include **Movement Status** and **Orientation** (see Table 1).

3.2.2 Instance-Grounded Questions To enable VLMs to answer questions about specific objects, some datasets [24, 27] incorporate object coordinates into the questions. However, VLMs struggle with interpreting precise numerical coordinates, while excelling at visual markers [26]. Building on evidence that even simple visual prompts,

such as red circles, can effectively guide VLM attention[23], we design instance-grounded questions that use red bounding boxes to highlight target objects. This approach tests whether VLMs can (1) correctly identify the referenced object among multiple instances, (2) classify the fine-grained object types (e.g., sedan vs SUV), and (3) recognize detailed attributes (color, orientation, state). In autonomous driving, this capability is crucial: passengers may need to reference specific vehicles or pedestrians when querying the system’s perception. Our instance-grounded questions include **Fine-grained Classification** and **Color Recognition** (see Table 1).

Rather than requiring models to infer user intent from ambiguous text alone, bounding-box inputs provide explicit spatial references similar to how humans naturally point to objects when asking questions. This design tests whether models can reason about referred objects or rely on scene-level statistics or textual correlations.

3.2.3 Motion Reasoning Questions. Beyond box-referring comprehension, we investigate whether VLMs can understand motion relationships of referred objects. To address this, we introduce more challenging motion reasoning questions that require understanding object movement over time and spatial relationships between the ego vehicle and surrounding objects. We categorize these as motion understanding questions, divided into implicit and explicit based on their input format.

Motion reasoning questions test VLMs’ spatiotemporal understanding through two approaches: implicit questions use single frames to infer motion from visual cues (e.g. motion blur, wheel orientation), while explicit questions provide frame sequences for direct temporal analysis. This design reveals whether models truly process temporal information or rely on static scene understanding. Our motion reasoning questions include **Speed Assessment**, **Relative Motion Analysis**, **Traffic Element Recognition**, **Trajectory Analysis** and **Path Conflict Detection** (see Table 1).

These questions have answers varying from binary to multiple choice. The multiple choice questions prompt the model to understand speed, heading and movement direction.

3.3 Question Balancing and Sampling

To ensure data set balance and prevent trivial solutions, we implement several sampling strategies.

Negative sampling: We find samples that do not fit the criteria for the ground truth answer and create additional samples from these criteria - i.e. negatives with varying object types, locations or temporal contexts. In fact, we find for many trajectory questions negative samples are much easier to come by than positive ones, as criteria for determining if a trajectory matches a question are usually quite strict, so finding positives is actually much harder.

Temporal consistency. We ensure that trajectory-based questions maintain logical consistency across frame sequences through: (1) linearly interpolating object positions to handle missing object information in Waymo, (2) requiring movement patterns to satisfy criteria over multiple consecutive frames to become positive, (3) smooth velocity calculations using temporal derivatives, and (4) selecting timestamps for questions based on temporal convergence patterns rather than instantaneous criteria.

Difficulty stratification. Questions are categorized by complexity, from basic object detection to complex multi-object trajectory

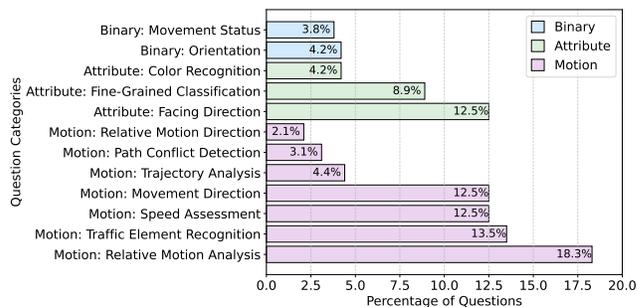


Figure 2: Question distribution across semantic categories in the Box-QAymo dataset. The taxonomy spans binary status checks, object attribute recognition, and motion-centric reasoning. Notably, motion-related questions form the majority, emphasizing the dataset’s focus on dynamic scene understanding in autonomous driving scenarios.

reasoning. The hierarchical structure allows for granular analysis of model capabilities, from basic object recognition to complex spatio-temporal reasoning in autonomous driving scenarios.

Answer format. Each multiple-choice question (MCQ) includes multiple-choice answers with 2-4 options, preventing binary guessing while maintaining evaluation precision.

4 Experiments

4.1 Experimental Setup

Models Evaluated. We evaluate three VLMs with varying specialization: **LLaVA-1.5 7B** [15], a general-purpose multimodal model trained on diverse internet-scale image-text pairs; **Qwen-7B** [4], a recently released large-scale vision-language model demonstrating strong zero-shot performance; and **Senna** [12], a domain-specific model tailored for autonomous driving scenarios, designed to integrate structured driving cues into multimodal reasoning.

Evaluation Protocol. We evaluated the models in our three-question hierarchies using F1 as the primary metric, supplemented by per-class precision and recall analysis. We report both overall performance and category-specific breakdowns to identify model strengths and limitations.

Steering Prompts. When running inference, we parse a format prompt to help steer the model to output the desired output format. This could be something like "Respond with only the full name of your selection (e.g., brown). Choose the best option out of ‘red’, ‘yellow’, ‘green’ or ‘brown’.". This supplies the model with the choices for multiple choice questions, allowing it to give a valid response.

Response Parsing. We employ a multistage parser to extract valid answers from free-form model outputs despite the use of steering prompts. This includes: (1) text normalization, (2) exact choice matching, and (3) ranked list handling. Outputs that do not contain a match are marked as invalid.

Finetuning. We finetuned models with LoRA [11] for one epoch on our finetuning split using LoRA with rank 128 and alpha 256. We used a learning rate of $2e-4$ with cosine scheduling, batch size 8, and bf16 precision. Finetuned models are denoted in tables and figures with † for single-frame and ‡ for two-frame models. The

training data was converted to LLaVA chat format and we adapted LLaVA’s LoRA finetuning script to suit.

Two-frame Motion. We investigate whether models can interpret object motion better if we provide the previous camera frame alongside the current, giving the model a chance to spot the differences and interpret the motion of the referred object. The prompt includes both frames to offer temporal context for the referred object.

4.2 Overall Performance

Figure 3 presents the overall performance of all models on the validation set. Finetuned models (\dagger) outperform their zero-shot counterparts in most cases, with the most significant gains observed in motion-related reasoning, highlighting the importance of task-specific adaptation. Notably, Senna’s performance is very low for many tasks and zero for attribute tasks. Qwen-VL is the best of the class, flanked only fine-tuned LLaVA.

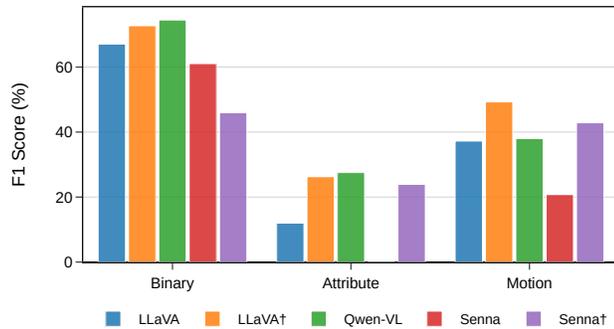


Figure 3: Overall F1 performance across question types for general-purpose and domain-specific VLMs.

Table 2: Performance on binary characteristic questions by subcategory. \dagger indicates single-frame finetuning. **Bold** indicates best performance, underline indicates second-best.

Question Type	Model	Precision (%)	Recall (%)	F1 (%)
Movement Status	LLaVA	<u>76.43</u>	62.55	68.77
	LLaVA \dagger	71.99	70.37	69.81
	Qwen-VL	76.59	76.13	76.03
	Senna	65.23	61.11	62.20
	Senna \dagger	70.43	56.17	46.91
Orientation	LLaVA	72.03	59.85	65.37
	LLaVA \dagger	77.86	76.02	75.62
	Qwen-VL	<u>73.49</u>	<u>73.05</u>	<u>72.92</u>
	Senna	65.65	61.90	59.94
	Senna \dagger	68.66	55.02	44.96

4.3 Discussion

Model Response Quality. Our evaluation reveals stark differences in models’ ability to generate parsable responses, with Senna achieving only 34% valid responses compared to LLaVA (90%) and Qwen-VL (99%). This parsing failure rate represents a fundamental limitation - models that cannot reliably follow structured output formats are unsuitable for safety-critical autonomous driving applications where consistent, interpretable outputs are essential.

Hierarchy Complexity Validation. Model evaluation across binary, attribute and motion questions can be found in Table 2, Table 3

Table 3: Performance on attribute questions by subcategory. \dagger indicates single-frame finetuning. **Bold** indicates best performance, underline indicates second-best.

Question Type	Model	Precision (%)	Recall (%)	F1 (%)
Color Recognition	LLaVA	22.37	17.14	15.10
	LLaVA \dagger	46.52	<u>37.60</u>	<u>38.46</u>
	Qwen-VL	<u>42.69</u>	50.49	42.49
	Senna	0.00	0.00	0.00
	Senna \dagger	39.16	29.30	28.54
Facing Direction	LLaVA	<u>22.98</u>	25.05	14.49
	LLaVA \dagger	9.40	24.95	13.66
	Qwen-VL	30.04	29.64	21.34
	Senna	0.00	0.00	0.00
	Senna \dagger	22.26	<u>27.08</u>	<u>20.72</u>
Fine-grained Classification	LLaVA	11.41	5.51	6.36
	LLaVA \dagger	<u>26.87</u>	29.88	26.73
	Qwen-VL	19.60	<u>24.93</u>	18.81
	Senna	0.00	0.00	0.00
	Senna \dagger	27.92	21.02	<u>22.50</u>

Table 4: Performance on motion questions by subcategory. \dagger indicates single-frame finetuning. **Bold** indicates best performance, underline indicates second-best.

Question Type	Model	Precision (%)	Recall (%)	F1 (%)
Movement Direction	LLaVA	24.62	25.38	19.13
	LLaVA \dagger	20.03	<u>25.50</u>	<u>20.23</u>
	Qwen-VL	26.29	26.80	19.68
	Senna	0.00	0.00	0.00
	Senna \dagger	<u>25.46</u>	24.57	22.22
Path Conflict Detection	LLaVA	47.53	43.21	44.88
	LLaVA \dagger	50.06	50.02	30.28
	Qwen-VL	<u>57.93</u>	56.63	56.84
	Senna	63.37	46.56	<u>50.08</u>
	Senna \dagger	53.08	<u>51.57</u>	33.89
Relative Motion Analysis	LLaVA	52.17	54.71	<u>48.93</u>
	LLaVA \dagger	64.57	<u>51.80</u>	50.94
	Qwen-VL	44.33	49.41	46.73
	Senna	49.71	28.38	35.61
	Senna \dagger	44.39	50.00	47.03
Relative Motion Direction	LLaVA	51.48	52.51	45.12
	LLaVA \dagger	57.90	61.84	57.40
	Qwen-VL	<u>54.56</u>	<u>57.29</u>	44.16
	Senna	0.00	0.00	0.00
	Senna \dagger	41.06	50.00	45.09
Speed Assessment	LLaVA	15.53	18.69	12.92
	LLaVA \dagger	38.76	28.94	27.41
	Qwen-VL	16.29	<u>22.39</u>	<u>15.86</u>
	Senna	0.00	0.00	0.00
	Senna \dagger	<u>24.83</u>	13.98	13.28
Traffic Element Recognition	LLaVA	53.81	54.56	46.07
	LLaVA \dagger	73.11	<u>58.33</u>	<u>59.73</u>
	Qwen-VL	<u>65.39</u>	65.77	65.57
	Senna	<u>51.42</u>	36.42	39.07
	Senna \dagger	40.34	50.00	44.66
Trajectory Analysis	LLaVA	44.56	43.67	43.63
	LLaVA \dagger	98.84	98.97	98.90
	Qwen-VL	18.65	18.90	17.03
	Senna	35.76	15.27	20.64
	Senna \dagger	<u>92.92</u>	<u>95.26</u>	<u>93.81</u>

and Table 4, respectively. Average F1 scores decrease substantially from binary questions (66.1%) to attribute questions (18.3%) and motion questions (37.6%). This strongly validates our complexity assumptions. This pattern reinforces our assumption that binary decisions are the most accessible visual reasoning task, attribute

Table 5: Single-frame vs two-frame input comparison on motion questions. Senna'' compares single-frame finetuned Senna to two-frame finetuned Senna. Despite finetuning on temporal data, single-frame input consistently outperforms two-frame input. **Bold** indicates best performance, underline indicates second-best.

Model	Single Frame F1 (%)	Two Frames F1 (%)	Improvement (%)
LLaVA	49.23	43.81	-5.42
LLaVA No Box	46.94	<u>46.33</u>	-0.61
Qwen-VL	52.88	50.14	-2.74
Senna	20.52	21.79	1.27
Senna''	56.56	44.40	-12.16

Table 6: Effectiveness of red bounding boxes at instance-grounding. ✓ and ✗ refer to models run with and without bounding boxes, respectively.

Question Type	Model	Box: ✗ F1 (%)	Box: ✓ F1 (%)
Color Recognition	LLaVA	14.75	15.10
	Qwen-VL	40.07	42.49
Fine-grained Classification	LLaVA	7.64	6.36
	Qwen-VL	18.65	18.81
Orientation	LLaVA	13.17	14.49
	Qwen-VL	19.74	21.34

reasoning demands fine-grained visual understanding and semantic knowledge, and motion understanding requires temporal reasoning capabilities that current VLMs lack. Importantly, the intermediate performance on motion tasks (37.6%) compared to attributes (18.3%) suggests that the models may rely on static visual cues rather than true temporal understanding.

Senna's Performance. Senna's poor performance despite being driving-specific reveals the brittleness of narrow task training. The model was optimized for planning (e.g., given some instructions, respond "go straight") and scene description, making it unable to adapt to structured Q&A. This highlights the importance of diverse training tasks for robust VLMs. General models (LLaVA, Qwen-VL) show better instruction following and adaptability. There is a significant challenge in task transfer for VLMs in autonomous driving. Development in autonomous driving implies designing domain-specific models like Senna, however, due to a narrowed task focus, these models can fail to transfer to other driving-domain tasks. Models should be trained on diverse tasks within their domain rather than just 1-2 specific task types.

Single vs Two Frames Analysis. We investigated whether consecutive frames improve motion understanding by comparing single frame vs. two frame input (current + previous frame at 100-ms intervals). Counterintuitively, consecutive frames consistently degrade performance across all models, even when specifically finetuned on temporal data (Senna'' : 56.56% vs 44.40% F1 in Table 5). This suggests that current VLMs struggle to effectively integrate short-term temporal information for structured reasoning tasks, likely due to training paradigms optimized for single images or longer video sequences rather than frame pairs.

Box Grounding Effectiveness. Table 6 shows that box grounding consistently improves Qwen-VL performance in all tasks (average +1.39% F1), with the largest gains in color recognition (+2.42% F1).

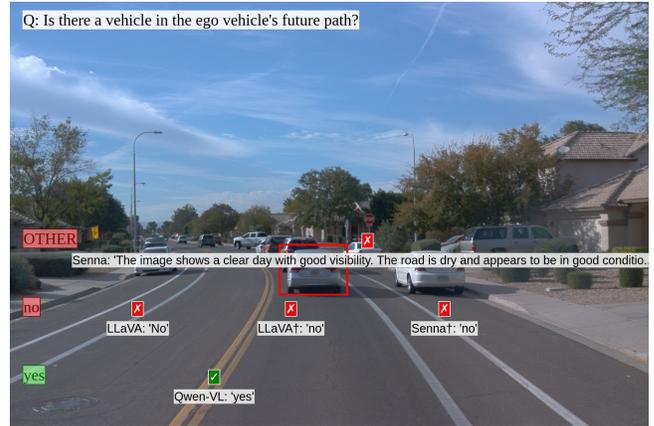


Figure 4: Example of a path conflict detection question and the outputs of pretrained and fine-tuned VLMs.

In contrast, LLaVA shows mixed results: improvements in color recognition and orientation (+0.35% and +1.32% F1 respectively) but degradation in fine-grained classification (-1.28% F1). We attribute these differences to architectural design: Qwen-VL's cross-attention mechanism and larger ViT-G/14 encoder enable better separation of visual grounding cues from box artifacts, while LLaVA's simpler MLP projection with ViT-L/14 struggles to filter box-induced noise. These findings highlight that the effectiveness of visual grounding critically depends on model architecture.

Case Study. Path conflict detection is a critical capability for autonomous vehicles. Figure 4 illustrates pretrained and fine-tuned VLM model responses to the given motion query. Despite the clear visual evidence of a leading vehicle in the lane, only Qwen-VL correctly identifies the conflict. LLaVA and its fine-tuned variant fail to recognize the obstruction, responding incorrectly with "no." The domain-specific VLM Senna also fails to detect the vehicle and instead generates an irrelevant scene-level caption focused on weather and road conditions. This example highlights a key limitation of current VLMs: even with fine-tuning, many models struggle with localized, intent-driven queries requiring precise spatial grounding.

5 Conclusion

We present the Box-QAymo dataset, to benchmark VLMs under realistic constraints. Our evaluation revealed three critical failures: format adherence, hierarchical performance degradation, and temporal reasoning. These findings challenge current approaches, especially that domain-specific models perform significantly worse on our benchmark. This highlights a critical gap in the demonstration of safety and trust in autonomous driving. When passengers cannot reliably query what an autonomous vehicle perceives, the fundamental promise of AI-enabled driving is not fulfilled. This gap between current capabilities and deployment requirements suggests that achieving trustworthy human-vehicle communication will require rethinking our approach to VLM development, moving beyond task optimization toward architectures that maintain robust reasoning across diverse query types.

References

- [1] AAA. 2023. Fear of Self-Driving Cars on the Rise. <https://newsroom.aaa.com/2023/03/aaa-fear-of-self-driving-cars-on-the-rise/>. Accessed: 2025-05-26.
- [2] AAA. 2025. Drivers Still Wary of Self-Driving Cars. <https://info.oregon.aaa.com/aaa-drivers-still-wary-of-self-driving-cars/>. Accessed: 2025-05-26.
- [3] National Highway Traffic Safety Administration et al. 2022. Summary report: standing general order on crash reporting for level 2 advanced driver assistance systems. *US Department of Transport: Washington, DC, USA* (2022).
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Ming-Hsuan Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. Qwen2.5-VL Technical Report. *CoRR* abs/2502.13923 (2025). doi:10.48550/ARXIV.2502.13923
- [5] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. 2020. nuScenes: A Multimodal Dataset for Autonomous Driving. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 11618–11628. doi:10.1109/CVPR42600.2020.01164
- [6] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. 2020. nuScenes: A Multimodal Dataset for Autonomous Driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [7] Long Chen, Oleg Sinavski, Jan Hünermann, Alice Karnsund, Andrew James Willmott, Danny Birch, Daniel Maund, and Jamie Shotton. 2024. Driving with llms: Fusing object-level vector modality for explainable autonomous driving. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 14093–14100.
- [8] Can Cui, Zichong Yang, Yupeng Zhou, Yunsheng Ma, Juanwu Lu, Lingxi Li, Yaobin Chen, Jitesh H. Panchal, and Ziran Wang. 2024. Personalized Autonomous Driving with Large Language Models: Field Experiments. In *27th IEEE International Conference on Intelligent Transportation Systems, ITSC 2024, Edmonton, AB, Canada, September 24-27, 2024*. IEEE, 20–27. doi:10.1109/ITSC58415.2024.10919978
- [9] Thierry Deruyttere, Simon Vandenhende, Dusan Grujicic, Luc Van Gool, and Marie-Francine Moens. 2019. Talk2car: Taking control of your self-driving car. *arXiv preprint arXiv:1909.10838* (2019).
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. 2012. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, June 16-21, 2012*. IEEE Computer Society, 3354–3361. doi:10.1109/CVPR.2012.6248074
- [11] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models. *CoRR* abs/2106.09685 (2021). arXiv:2106.09685 <https://arxiv.org/abs/2106.09685>
- [12] Bo Jiang, Shaoyu Chen, Bencheng Liao, Xingyu Zhang, Wei Yin, Qian Zhang, Chang Huang, Wenyu Liu, and Xinggang Wang. 2024. Senna: Bridging Large Vision-Language Models and End-to-End Autonomous Driving. *CoRR* abs/2410.22313 (2024). doi:10.48550/ARXIV.2410.22313 arXiv:2410.22313
- [13] Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. 2019. Grounding human-to-vehicle advice for self-driving vehicles. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10591–10599.
- [14] Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John Canny, and Zeynep Akata. 2018. Textual explanations for self-driving vehicles. In *Proceedings of the European conference on computer vision (ECCV)*. 563–578.
- [15] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual Instruction Tuning. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.). http://papers.nips.cc/paper_files/paper/2023/hash/6dcf27ea32ce3288914fa369fe6de0-Abstract-Conference.html
- [16] Yunsheng Ma, Can Cui, Xu Cao, Wenqian Ye, Peiran Liu, Juanwu Lu, Amr Abdelaouf, Rohit Gupta, Kyungtae Han, Aniket Bera, et al. 2024. Lampilot: An open benchmark dataset for autonomous driving with language model programs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 15141–15151.
- [17] Srikanth Malla, Chihoh Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. 2023. Drama: Joint risk localization and captioning in driving. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 1043–1052.
- [18] Ana-Maria Marcu, Long Chen, Jan Hünermann, Alice Karnsund, Benoît Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, Alex Kendall, Jamie Shotton, Elahe Arani, and Oleg Sinavski. 2024. LingoQA: Visual Question Answering for Autonomous Driving. In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024. Proceedings, Part LXXVII (Lecture Notes in Computer Science, Vol. 15135)*, Ales Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (Eds.). Springer, 252–269. doi:10.1007/978-3-031-72980-5_15
- [19] Ana-Maria Marcu, Long Chen, Jan Hünermann, Alice Karnsund, Benoît Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, Alex Kendall, Jamie Shotton, and Oleg Sinavski. 2023. LingoQA: Video Question Answering for Autonomous Driving. *CoRR* abs/2312.14115 (2023). doi:10.48550/ARXIV.2312.14115 arXiv:2312.14115
- [20] Ming Nie, Renyuan Peng, Chunwei Wang, Xinyue Cai, Jianhua Han, Hang Xu, and Li Zhang. 2024. Reason2drive: Towards interpretable and chain-based reasoning for autonomous driving. In *European Conference on Computer Vision*. Springer, 292–308.
- [21] Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. 2024. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 4542–4550.
- [22] Enna Sachdeva, Nakul Agarwal, Suhas Chundi, Sean Roelofs, Jiachen Li, Mykel Kochenderfer, Chihoh Choi, and Behzad Dariush. 2024. Rank2tell: A multimodal driving dataset for joint importance ranking and reasoning. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 7513–7522.
- [23] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. 2023. What does CLIP know about a red circle? Visual prompt engineering for VLMs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 11987–11997.
- [24] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Ping Luo, Andreas Geiger, and Hongyang Li. 2023. DriveLM: Driving with Graph Visual Question Answering. *CoRR* abs/2312.14150 (2023). doi:10.48550/ARXIV.2312.14150 arXiv:2312.14150
- [25] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Ettinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. 2020. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. Computer Vision Foundation / IEEE, 2443–2451. doi:10.1109/CVPR42600.2020.00252
- [26] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Yang Wang, Zhiyong Zhao, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. 2024. DriveVLM: The Convergence of Autonomous Driving and Large Vision-Language Models. In *Conference on Robot Learning, 6-9 November 2024, Munich, Germany (Proceedings of Machine Learning Research, Vol. 270)*, Pulkit Agrawal, Oliver Kroemer, and Wolfram Burgard (Eds.). PMLR, 4698–4726. <https://proceedings.mlr.press/v270/tian25c.html>
- [27] Shihao Wang, Zhiding Yu, Xiaohui Jiang, Shiyi Lan, Min Shi, Nadine Chang, Jan Kautz, Ying Li, and José M. Álvarez. 2025. OmniDrive: A Holistic Vision-Language Dataset for Autonomous Driving with Counterfactual Reasoning. *CoRR* abs/2504.04348 (2025). doi:10.48550/ARXIV.2504.04348 arXiv:2504.04348
- [28] Wenhui Wang, Jiangwei Xie, ChuanYang Hu, Haoming Zou, Jianan Fan, Wenwen Tong, Yang Wen, Silei Wu, Hanming Deng, Zhiqi Li, et al. 2023. Drivemlm: Aligning multi-modal large language models with behavioral planning states for autonomous driving. *arXiv preprint arXiv:2312.09245* (2023).
- [29] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. 2023. Argoverse 2: Next Generation Datasets for Self-Driving Perception and Forecasting. *CoRR* abs/2301.00493 (2023). doi:10.48550/ARXIV.2301.00493 arXiv:2301.00493
- [30] Dongming Wu, Wencheng Han, Yingfei Liu, Tiancai Wang, Cheng-zhong Xu, Xiangyu Zhang, and Jianbing Shen. 2025. Language prompt for autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 8359–8367.
- [31] Yiran Xu, Xiaoyin Yang, Lihang Gong, Hsuan-Chu Lin, Tz-Ying Wu, Yunsheng Li, and Nuno Vasconcelos. 2020. Explainable object-induced action decision for autonomous vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9523–9532.
- [32] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K. Wong, Zhenguo Li, and Hengshuang Zhao. 2024. DriveGPT4: Interpretable End-to-End Autonomous Driving Via Large Language Model. *IEEE Robotics Autom. Lett.* 9, 10 (2024), 8186–8193. doi:10.1109/LRA.2024.3440097