

RoadSceneVQA: Benchmarking Visual Question Answering in Roadside Perception Systems for Intelligent Transportation System

Runwei Guan^{1,2,*}, Rongsheng Hu³, Shangshu Chen⁴, Ningyuan Xiao¹, Xue Xia¹, Jiayang Liu¹,
Beibei Chen⁵, Ziren Tang¹, Ningwei Ouyang⁶, Shaofeng Liang¹, Yuxuan Fan¹,
Wanjie Sun⁷, Yutao Yue^{1,2,*}

¹Information Hub, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

²Institute of Deep Perception Technology, JITRI, Wuxi, China

³School of Artificial Intelligence and Computer Science, Jiangnan University, Wuxi, China

⁴National Research Center of Cultural Industries, Central China Normal University, Wuhan, China

⁵School of Marxism, Jilin University, Changchun, China

⁶School of Advanced Technology, Xi'an Jiaotong-Liverpool University, Suzhou, China

⁷School of Remote Sensing and Information Engineering, Wuhan University, Wuhan, China

Abstract

Current roadside perception systems mainly focus on instance-level perception, which fall short in enabling interaction via natural language and reasoning about traffic behaviors in context. To bridge this gap, we introduce RoadSceneVQA, a large-scale and richly annotated visual question answering (VQA) dataset specifically tailored for roadside scenarios. The dataset comprises 34,736 diverse QA pairs collected under varying weather, illumination, and traffic conditions, targeting not only object attributes but also the intent, legality, and interaction patterns of traffic participants. RoadSceneVQA challenges models to perform both explicit recognition and implicit commonsense reasoning, grounded in real-world traffic rules and contextual dependencies. To fully exploit the reasoning potential of Multi-modal Large Language Models (MLLMs), we further propose CogniAnchor Fusion (CAF), a vision-language fusion module inspired by human-like scene anchoring mechanisms. Moreover, we propose the Assisted Decoupled Chain-of-Thought (AD-CoT) to enhance the reasoned thinking via CoT prompting and multi-task learning. Based on the above, we propose the baseline model RoadMind. Experiments on RoadSceneVQA and CODA-LM benchmark show that the pipeline consistently improves both reasoning accuracy and computational efficiency, allowing the MLLM to achieve state-of-the-art performance in structural traffic perception and reasoning tasks.

Datasets — <https://github.com/GuanRunwei/RS-VQA>

Introduction

With the advancement of Multi-modal Large Language Models (MLLMs), the capabilities for perception and reasoning in traffic scenarios have significantly improved (Zhou et al. 2024b). As a critical component of traffic perception, roadside perception offers a distinct advantage over vehicle-based perception by providing a top-down perspective that enables clearer observation of the states and behaviors of

traffic participants, and a more comprehensive understanding of the overall scene (Bejarbaneh, Du, and Naghdy 2024).

However, current roadside perception systems primarily focus on automated tasks such as detection (Guan et al. 2023), tracking (Liang et al. 2025), trajectory prediction (Dai et al. 2025), and traffic flow forecasting (Huang et al. 2023). While these approaches have achieved a certain degree of success, they largely lack human-in-the-loop perception, which mostly emphasizes instance-level recognition, but falls short in event-level or holistic scene-level understanding. Moreover, these systems often exhibit limited scalability and interpretability in complex environments, and lack the flexibility to recognize unforeseen objects and events. Recently, Vision-Language Models (VLMs), empowered by large-scale pretraining, have demonstrated strong reasoning capabilities (Guan et al. 2025a) and generalization across a wide range of multi-modal tasks (Sun et al. 2024, 2025; Zhang et al. 2025a). These models have shown promising performance in vehicle-based perception. Moreover, the reasoning ability of MLLMs, enabled by prompting, allows for rapid adaptation to novel traffic events.

While VLM-based roadside perception benchmarks have achieved remarkable progress in visual grounding or image captioning (Zhou et al. 2025; Guan et al. 2024), their reliance on explicit localization fundamentally limits higher-level traffic understanding. Existing benchmarks measure perceptual accuracy, yet fail to evaluate whether models comprehend implicit traffic regulations, such as determining “Is any pedestrian violating traffic rules?”, which requires synthesizing signal states, spatial contexts, and behavioral dynamics. This semantic gap persists because localization tasks disregard causal relationships (e.g., a vehicle may be in a crosswalk but not violating if signals permit).

Building upon the aforementioned, we construct RoadSceneVQA, a large-scale Visual Question Answering (VQA) dataset tailored for roadside perception. Leveraging the rich object distributions and contextual diversity of the impressive Rope3D dataset (Ye et al. 2022), which provides roadside detection data across various traffic scenarios. It

*CA: {runwayrguan, yutaoyue}@hkust-gz.edu.cn
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Overview of RoadSceneVQA with samples.

| Datasets | Venues | QA Generation | QA/Caption items | Scenes | Reasoning | Domains |
|-----------------------------------|-----------------------|---------------------|------------------|-----------|-----------|-----------------|
| Talk2Car (Deruyttere et al. 2019) | EMNLP ₂₀₁₉ | Manual | 2.4k | 2.4k | ✗ | Driving |
| NuScenes-QA (Qian et al. 2024) | AAAI ₂₀₂₄ | Template | 83.3k | 0.9k | ✗ | Driving |
| DriveLM (Sima et al. 2024) | ECCV ₂₀₂₄ | Template + Manual | 15.4k | 188k | ✓ | Driving |
| Talk2Radar (Guan et al. 2025c) | ICRA ₂₀₂₅ | Manual | 8.7k | 8.7k | ✗ | Driving |
| nuPrompt (Wu et al. 2025) | AAAI ₂₀₂₅ | LLM | 40.8k | 0.9k | ✗ | Driving |
| DriveBench (Xie et al. 2025) | ICCV ₂₀₂₅ | Template + Manual | 20.5k | - | ✓ | Driving |
| TUM-VideoQA (Zhou et al. 2025) | ICML ₂₀₂₅ | LLM + Manual | 87.3k | 1k | ✗ | Roadside |
| RoadSceneVQA (ours) | 2025 | LLM + Manual | 34.7k | 26 | ✓ | Roadside |

Table 1: Summary of vision-language-based datasets oriented at traffic scenarios.

covers a wide range of traffic events, object states, lighting conditions, and weather environments. RoadSceneVQA includes both explicit and implicit question-answer pairs centered on traffic scene reasoning, posing challenges to models in terms of scene understanding and commonsense reasoning. More importantly, we introduce an agile Collaborative Human-Machine Annotation (CH-MA) system to enable high-quality and scalable annotation, which is critical for constructing this comprehensive VQA dataset.

Last but not least, current MLLMs typically adopt a simple token-level concatenation strategy for the fusion of visual and language tokens. However, this design introduces two limitations: (1) Irrelevant visual tokens, such as background noise, are forcibly mixed with key textual tokens, disrupting the model’s ability to accurately localize objects; (2) Imbalanced information interaction, where dominant visual features overwhelm the textual signals, impairs the model’s ability to follow language instructions (Yang et al. 2024). To enhance MLLMs’ contextual awareness for commonsense reasoning, we draw inspiration from human vision-language collaborative cognition and propose a plug-and-play visual-language fusion module, termed CogniAnchor Fusion (CAF). CAF pre-anchors potential regions of interest and enables the LLM to reason precisely and efficiently. Moreover, advanced MLLM such as GPT-4o (Hurst et al. 2024) have demonstrated strong capabilities in knowledge-driven reasoning, and have been applied to autonomous driving for Chain-of-Thought (CoT) generation, enabling more stable and interpretable reasoning paradigms (Liao et al. 2025). To address the challenges of complex roadside scenarios, we propose a novel method called Assisted Decoupled Chain-of-Thought (AD-CoT), designed to support reliable scene understanding and reasoning on

lightweight edged MLLMs (smaller than 8B). Building on these, we further introduce our baseline MLLM, RoadMind, specifically tailored for multi-modal roadside reasoning.

In summary, our contributions are as follows,

1. **RoadSceneVQA**, a large-scale VQA dataset designed for roadside perception, which comprises 34,736 samples of Q-A pair, spanning perception and reasoning tasks.
2. **CH-MA**, an agile human-machine collaborative annotation framework for VQA datasets, which improves annotation efficiency while enhancing accuracy and quality.
3. **CogniAnchor Fusion**, a visual-language fusion module inspired by human scene cognition mechanisms, tailored for MLLMs. It enhances the reasoning capabilities while maintaining high computational efficiency.
4. **RoadMind**, a multi-modal large language model empowered by **Assisted Decoupled Chain-of-Thought (AD-CoT)** for roadside traffic perception and reasoning.

Related Works

Open-Ended VQA based on MLLMs. Visual Question Answering (VQA) has significantly evolved from conventional selection-based and multi-class frameworks to open-ended, human-like systems (Kim et al. 2025). Early VQA models treat the task as multi-class classification, predicting answers from predefined sets by fusing CNN-extracted visual features with textual embeddings (Yang et al. 2016; Antol et al. 2015; Malinowski, Rohrbach, and Fritz 2015). Despite promising results, their fixed answer sets hinder handling complex queries and open-ended reasoning. With the rise of large-scale datasets like Visual Genome (Krishna et al. 2017), researchers adopt encoder-decoder architectures. Incorporating attention mechanisms, the mod-

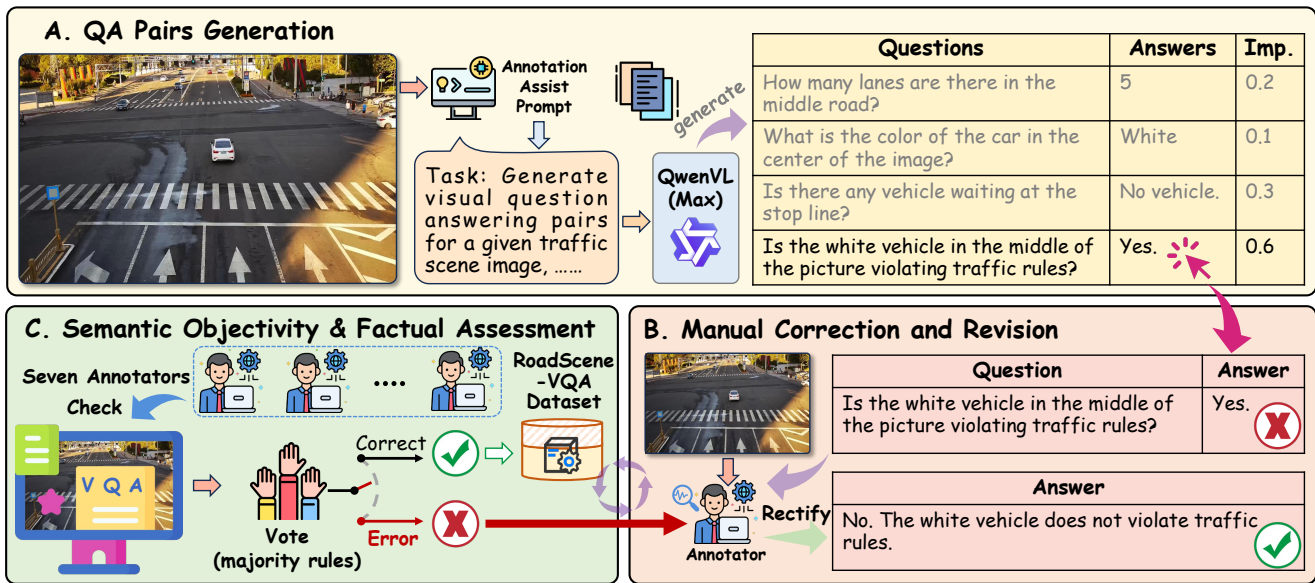


Figure 2: Annotation of RoadSceneVQA dataset by the agile Collaborative Human-Machine Annotation system.

els enable more free-form text generation (Nam, Ha, and Kim 2017; Anderson et al. 2018). However, semantic coherence and factual accuracy remain challenges due to difficulties in grounding visual concepts with language. In the advent of MLLMs, by integrating pre-trained knowledge with visual understanding, MLLMs generate natural, open-ended responses similar to human answers (Sima et al. 2024; Zhang et al. 2025b). Currently, VQA in traffic scenarios primarily focuses on explicit QA (Qian et al. 2024; Wu et al. 2025), with limited efforts exploring context-aware reasoning, which are constrained to on-vehicle settings (Xie et al. 2025). Existing roadside VQA benchmarks tend to emphasize object localization, typically outputting bounding boxes (Zhou et al. 2025), which is difficult to reason effectively about the behavior of traffic participants or to draw decision-level conclusions. To address these limitations, we propose RoadSceneVQA, which fills this gap by implicitly incorporating commonsense knowledge into VQA tasks for traffic, enabling more holistic and actionable scene understanding.

Roadside Perception for ITS. Roadside perception is vital for Intelligent Transportation Systems (ITS), offering a complementary perspective to on-board vehicle sensors for traffic monitoring and decision-making. Recent advances have addressed a range of tasks. Object detection (Ye et al. 2022; Yu et al. 2022) and semantic segmentation (Muhammad et al. 2022) from roadside cameras have been explored to identify traffic participants in varying conditions. Occupancy-based methods aim to construct bird’s-eye-view representations for scene understanding (Chang et al. 2023). Moreover, traffic flow prediction models enable proactive traffic management and congestion control (Xu et al. 2021). Despite these advances, most existing researches focus on perception only or fixed outputs such as bounding boxes, which often require post-processing or human intervention for scene understanding. In contrast, VQA offers a natu-

ral and flexible interface for reasoning about traffic scenes, yet current VQA datasets are predominantly vehicle-centric and lack the complexity of roadside contexts. To bridge this gap, we propose RoadSceneVQA, a comprehensive VQA dataset tailored for roadside perception. It introduces diverse, context-rich, and cognitively challenging question-answer pairs covering traffic understanding, behavior reasoning, and commonsense inference, paving the way for next-generation multi-modal reasoning in ITS applications.

RoadSceneVQA Dataset

RoadSceneVQA presents the first vision-language benchmark that shifts traffic intelligence evaluation from perceptual recognition to regulation-aware cognitive reasoning from roadside view. Unlike existing VQA datasets which predominantly query explicit object attributes or rely on video temporal cues (e.g., ‘Is the vehicle speeding?’), RoadSceneVQA pioneers image-based compliance judgment, demanding models to answer questions like ‘Is the pedestrian violating traffic rules given current signal states and crosswalk topology?’. This necessitates a tripartite integration of: (i) visual-semantic grounding (e.g., correlating pedestrian position with signal phase), (ii) regulatory knowledge internalization, and (iii) counterfactual causality (e.g., inferring ‘Would the cyclist have violated if the light were green?’). In this section, we first introduce **agile Collaborative Human-Machine Annotation system**, a user-friendly and efficient framework for VQA annotation. Then, we provide a statistical analysis of the dataset across multiple dimensions.

Agile Collaborative Human-Machine Annotation

To construct a natural and accurate question-answering dataset that emulates human-like reasoning, we propose an agile Collaborative Human-Machine Annotation system (CH-MA). This framework is designed to address several

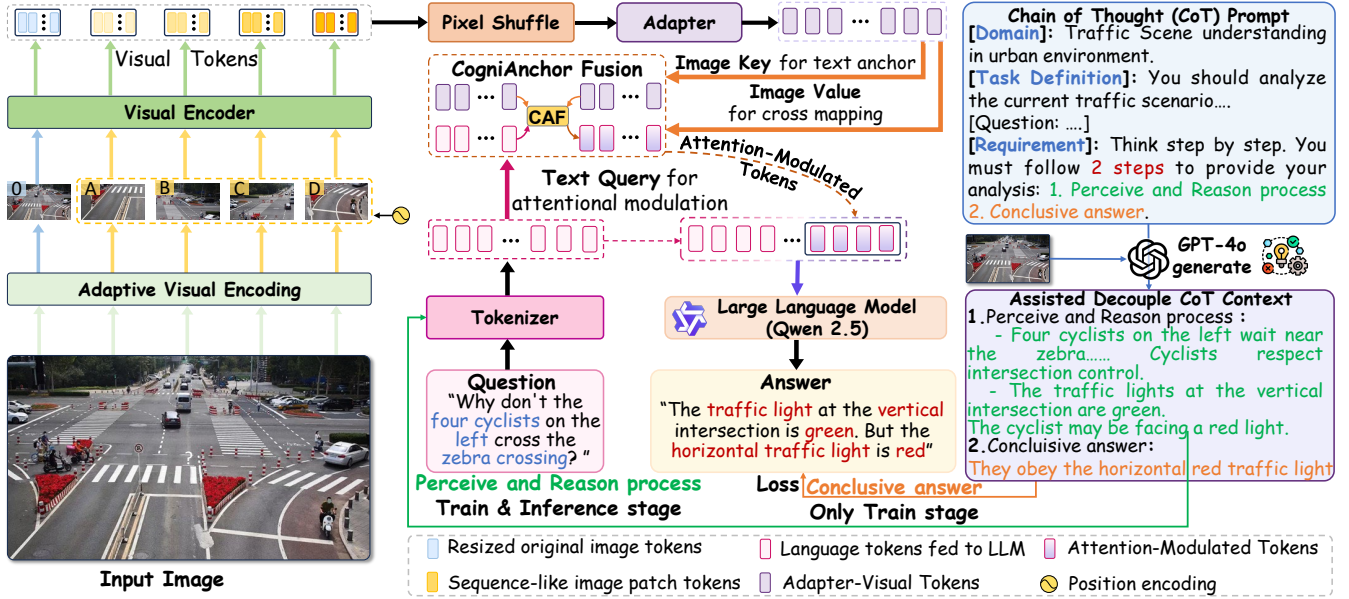


Figure 4: The architecture of RoadMind Model, where we propose CogniAnchor Fusion and Assisted Decoupled CoT.

et al. 2024) to generate a Chain-of-Thought prompt \mathbf{T}^{OC} and a structured answer \mathbf{T}^{SA} , both conditioned on the image \mathbf{I} and the original question \mathbf{T} . We then concatenate the \mathbf{T} with its generated \mathbf{T}^{CoT} and obtain \mathbf{T}^{OC} :

$$\mathbf{T}^{OC} = \text{Concat}(\mathbf{T}, \mathbf{T}^{CoT}). \quad (2)$$

This enriched textual input \mathbf{T}^{OC} is tokenized and encoded into text embeddings. Meanwhile, the visual features extracted from InternViT are passed through a lightweight MLP adapter to align their representation space with that of the text. The aligned image features and the encoded textual representation are then fused via the CogniAnchor Fusion module, resulting in text-conditioned visual tokens. These tokens are concatenated with the language embeddings and jointly fed into Qwen 2.5, to produce the final answer \mathbf{A} :

$$\mathbf{A} = \text{Qwen2.5}(\text{Concat}(\text{CAF}(\mathbf{V}, \mathbf{T}^{OC}), \mathbf{T}^{OC})) \quad (3)$$

CongiAnchor Fusion

Most MLLMs adopt a naive token-level concatenation that merges visual and textual tokens. However, this paradigm requires the model to learn cross-modal spatial correspondences from scratch, often causing inefficient grounding of visual evidence performance. To address these limitations, we propose CogniAnchor Fusion (CAF), a cognitively inspired fusion mechanism that mirrors the typical human process of answering visual questions, which obtains language-anchored candidate visual regions for the following fusion with language features. Besides, to ensure efficiency, CAF is formulated as a linear-based cross attention. However, linear attention often suffers from non-injectivity, which leads to semantic ambiguity and impair the model's ability to accurately understand candidate regions within complex scenes. Therefore, CAF is inspired by InLine Attention (Han et al.

2024), enabling scalable integration with existing MLLMs while maintaining efficiency in reasoning.

Given the textual input as $\mathbf{Q}^T \in \mathbb{R}^{N \times d}$ for the source driving the query. $\{\mathbf{K}^I, \mathbf{V}^I\} \in \mathbb{R}^{N \times d}$ are the key and value upon visual features to calculate the attention matrix and form the representation. For the query vector \mathbf{Q}_i^T , the function of attention weight $\text{AW}(\cdot)$ is defined as:

$$\begin{aligned} \text{AW}(\mathbf{Q}_i^T, \mathbf{K}) &= [\phi(\mathbf{Q}_i^T)^\top \phi(\mathbf{K}_1), \dots, \phi(\mathbf{Q}_i^T)^\top \phi(\mathbf{K}_N)]^\top \\ &\quad - \frac{1}{N} \sum_{s=1}^N \phi(\mathbf{Q}_i^T)^\top \phi(\mathbf{K}_s) + \frac{1}{N}, \end{aligned} \quad (4)$$

where $\phi(\cdot)$ is the kernel function. $\frac{1}{N} \sum_{s=1}^N \phi(\mathbf{Q}_i^T)^\top \phi(\mathbf{K}_s)$ indicates the averaging operation over all similarity scores. $\frac{1}{N}$ is a constant term added to ensure that attention weights sum to 1 (meet the normalization). Finally, the language-driven image attention features can be formulated as:

$$\begin{aligned} \text{AW}(\mathbf{Q}_i^T, \mathbf{K}^I)^\top \mathbf{V}^I &= \phi(\mathbf{Q}_i^T)^\top \left[\sum_{j=1}^N \phi(\mathbf{K}_j^I) \mathbf{V}_j^\top \right] \\ &\quad - [\phi(\mathbf{Q}_i^T)^\top \sum_{j=1}^N \phi(\mathbf{K}_j^I) - 1] \frac{1}{N} \sum_{j=1}^N \mathbf{V}_j^I. \end{aligned} \quad (5)$$

Here, by adjusting the calculation order (aggregating key-value pairs first and then interacting with the query), the complexity is reduced from $O(N^2)$ of traditional Softmax attention to $O(N)$.

In conclusion, CAF enhances the causality of the generated responses, leveraging the reasoning capability of a large MLLM to empower a smaller and deployable MLLM.

| Models | LLM | Params | Exact Match | ROUGH-L | BLEU-4 | METEOR | CIDEr | SPICE | GPT-Score |
|----------------------|-------------|--------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| MiniCPM-o 2.6 (nsft) | LLaMA3 | 8B | 0.021 | 0.193 | 0.080 | 0.324 | 0.661 | 0.124 | 0.428 |
| EM-VLM4AD | T5-Large | 0.8B | 0.086 | 0.249 | 0.134 | 0.299 | 1.142 | 0.148 | 0.372 |
| InternVL3 | Qwen 2.5 | 0.9B | 0.142 | 0.386 | 0.166 | 0.385 | 1.656 | 0.170 | 0.403 |
| TinyLLaVA | OpenELM | 0.9B | 0.098 | 0.283 | 0.127 | 0.319 | 1.150 | 0.147 | 0.394 |
| RoadMind | Qwen 2.5 | 0.9B | 0.144 | 0.372 | 0.179 | 0.322 | 1.867 | 0.188 | 0.440 |
| MobileVLM v2 | MobileLLaMA | 1.7B | 0.081 | 0.296 | 0.137 | 0.338 | 1.399 | 0.162 | 0.417 |
| InternVL3 | Qwen 2.5 | 2B | 0.151 | 0.403 | 0.155 | 0.315 | 1.834 | 0.201 | 0.465 |
| MiniCPM-V 2.0 | MiniCPM | 2.8B | 0.136 | 0.389 | 0.149 | 0.352 | 1.717 | 0.194 | 0.477 |
| RoadMind | Qwen 2.5 | 2B | 0.142 | 0.405 | 0.158 | 0.347 | 1.705 | 0.219 | 0.489 |
| Qwen2.5-VL | Qwen2.5 | 7B | 0.152 | 0.403 | 0.159 | 0.385 | 1.689 | 0.213 | 0.497 |
| InternVL3 | Vicuna | 8B | 0.161 | 0.411 | 0.162 | 0.398 | 1.735 | 0.208 | 0.532 |
| MiniCPM-o 2.6 | LLaMA3 | 8B | 0.147 | 0.388 | 0.165 | 0.394 | 1.826 | 0.199 | 0.527 |
| RoadMind | Qwen 2.5 | 8B | 0.157 | 0.425 | 0.157 | 0.411 | 1.836 | 0.221 | 0.554 |

Table 2: Overall performances on RoadSceneVQA dataset. nsft: No Supervised Fine-Tuning.

Assisted Decoupled Chain-of-Thought (AD-CoT)

Due to the limited capacity of lightweight roadside MLLMs, their reasoning capabilities are inherently constrained. To address this, we propose AD-CoT that combines knowledge distillation with cognitive transfer, where GPT-4o serves as a soft supervisory prior to enhance the model’s ability to perform context-aware reasoning in complex traffic scenarios. Exactly, given an input image and a question, we first feed the image along with a CoT prompt into GPT-4o. The model generates an assisted reasoning context, which includes both a Perceive and Reason Process and a Conclusive Answer, where the former is then concatenated with the original question and passed into our lightweight model, RoadMind, as enriched input. Meanwhile, the conclusive answer from GPT-4o is paired with the human-annotated ground truth answer to construct a multi-task learning objective, enabling RoadMind to simultaneously learn both reasoning logic and answer accuracy. The loss function L_{MTL} is presented as follows:

$$L_{MTL} = \frac{1}{\sigma_{\text{hard}}^2} \left[\sum_{l=1}^L \log p(\mathbf{y}_l^{\text{hard}} | \mathbf{y}_{<l}^{\text{hard}}, \mathbf{x}, \mathbf{q}) \right] + \frac{1}{\sigma_{\text{soft}}} \left[\sum_{l=1}^{\min(L, L')} D_{\text{KL}}(p_l^{\text{GPT}} || \hat{p}_l) \right] + \log \sigma_{\text{hard}} + \log \sigma_{\text{soft}}, \quad (6)$$

where \mathbf{x} denotes the image input while \mathbf{q} is question context. \mathbf{y}^{hard} is the ground truth answer. L is the length of ground truth answer while L' is the length of GPT-4o’s answer. D_{KL} is Kullback-Leibler divergence function. p_l^{GPT} is the probability distribution of GPT-4o at position l while \hat{p}_l is the probability distribution of the model at position l . σ_{hard} and σ_{soft} are two learnable weights for uncertainty modelling.

Experiments

Settings of Experiments

Dataset. Our proposed RoadSceneVQA is adopted for baseline experiments, including 30,058 samples for training and 4,677 samples for test. To evaluate the generalization of models, CODA-LM (Chen et al. 2025), another VQA

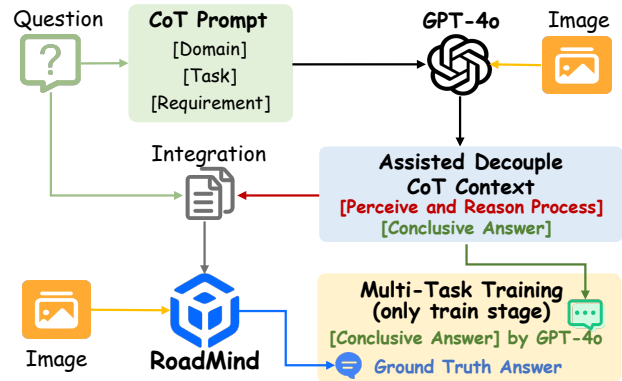


Figure 5: The pipeline of Assisted Decouple CoT.

dataset for ego-car VQA is also included, containing 20,495 samples in train set while 2,123 samples in test set.

Models. Besides RoadMind models (0.9B, 2.0 and 8.0B), QwenVL (Bai et al. 2025), InternVL (Chen et al. 2024), MiniCPM series (Yao et al. 2024), MobileVLM v2 (Chu et al. 2024), TinyLLaVA (Zhou et al. 2024a) and EM-VLM4AD (Gopalkrishnan, Greer, and Trivedi 2024) are included, containing containing MLLM for general and specific fields. Besides, for the proposed CogniAnchor Fusion, we set the head number of attention as 8.

Training and Evaluation. For the training of RoadSceneVQA, we resize the image as 448×448 (px). We finetune the pretrained models for one epoch with the initial learning rate of $1e-5$. We freeze the visual encoder while unfreeze LLM and MLP projector. We set the max sequence length as 16384. We adopt AdamW optimizer and set the weight decay as 0.05 with the cosine scheduler and warm-up ratio of 0.03. We train models on 4 A100 GPUs with the batch size per GPU as 1. For CODA-LM, we keep the settings consistent with (Chen et al. 2025). We include Exact Match, ROUGH-L, (Lin 2004), BLEU-4 (Papineni et al. 2002), METEOR (Lavie and Agarwal 2007), CIDEr (Vedantam, Lawrence Zitnick, and Parikh 2015), SPICE (Anderson et al. 2016) and GPT-Score (Sima et al. 2024) as evaluation



Figure 6: Inference results upon RoadMind-8B.

| Sample Types | ROUGH-L | METEOR | CIDEr | GPT-Score |
|-------------------|---------|--------|-------|-----------|
| Perception | | | | |
| TPP | 0.404 | 0.391 | 1.894 | 0.561 |
| TLP | 0.417 | 0.348 | 1.877 | 0.583 |
| EWP | 0.486 | 0.437 | 1.929 | 0.601 |
| FOMS | 0.432 | 0.408 | 1.752 | 0.530 |
| Reasoning | | | | |
| VR | 0.218 | 0.177 | 1.376 | 0.338 |
| TPBR | 0.323 | 0.216 | 1.438 | 0.429 |
| PS | 0.287 | 0.253 | 1.393 | 0.371 |

Table 3: Performances upon RoadMind-8B. The abbreviations in **Sample Types** are derived from Figure 3(e).

metrics, which are widely used in text generation field.

Quantitative Results

Overall. Table 2 presents the performance of multi-scale MLLMs on the RoadSceneVQA. Overall, across three model size tiers, proposed RoadMind consistently achieves the best performance on most evaluation metrics. Notably, it excels in the GPT-Score, indicating stronger capabilities in contextual understanding and reasoning, as well as generating responses that are more aligned with human preferences. Furthermore, RoadMind (0.9B) outperforms both the zero-shot MiniCPM-o 2.6 (8B) and the fully fine-tuned MobileVLM v2 (1.7B) in terms of GPT-Score. This shows that our design, particularly the CAF module and the AD-CoT strategy, effectively enhancing the reasoning ability of lightweight models. These enable them to reach or surpass the performance of significantly larger counterparts.

Sample Types. Table 3 reports the performance of RoadMind-8B across different sample categories. The model performs significantly better on perception-based questions compared to reasoning-based ones, highlighting the relative difficulty of high-level contextual inference. Within the perception category, the model achieves higher accuracy on weather and traffic light recognition, whereas its performance on understanding traffic participants and road-side infrastructure remains comparatively lower.

| Models | GTS \uparrow | Regional Perception \uparrow | | | | STS \uparrow |
|--------------------|----------------|--------------------------------|--------------|--------------|--------------|----------------|
| | | All | Vehicle | VRU | Sign | |
| MiniGPT-v2-7B | 11.58 | 15.93 | 18.74 | 13.58 | 15.71 | 10.0 |
| Shikra-7B | 12.24 | 22.94 | 28.29 | 17.88 | 20.0 | 10.2 |
| QwenVL2-7B | 18.22 | 26.62 | 35.48 | 24.16 | 20.86 | 22.06 |
| MiniCPM2.5-8B | 41.12 | 57.2 | 61.91 | 54.82 | 59.43 | 48.48 |
| InternVL1.5-20B | 38.38 | 61.53 | 63.77 | 53.14 | 50.57 | 41.18 |
| RoadMind-8B | 48.50 | 70.65 | 74.25 | 59.78 | 47.43 | 54.28 |

Table 4: Comparison of models on CODA-LM dataset. **GTS**: General Text-Score; **STS**: Suggestion Text-Score.

| NSFT | LoRA | SFT | SFT + CAF | SFT + CAF + AD-CoT |
|-------|-------|-------|-----------|--------------------|
| 0.428 | 0.452 | 0.527 | 0.533 | 0.549 |

Table 5: Migration performance (GPT-Score) on MiniCPM-o 2.6 for CAF and AD-CoT.

Generalization. Table 4 shows that RoadMind-8B achieves strong performance on the CODA-LM benchmark, even surpassing InternVL 1.5-20B, which highlights the remarkable capability of RoadMind across diverse reasoning tasks.

Ablation Studies

First, Table 5 presents migration performance of MiniCPM-o 2.6 for CAF and AD-CoT, indicating the impressive transfer ability of these two approaches. **Second,** Table 6 presents the superior performances of CAF, with the trade-off be-

| VL Fusion | Params | FLOPs | ROUGH-L | METEOR | SPICE |
|-------------------|---------------|---------------|--------------|--------------|--------------|
| CAF+Concat | 0.924K | 61.08M | 0.425 | 0.411 | 0.221 |
| SCA+Concat | 0.768K | 18.79M | 0.389 | 0.401 | 0.197 |
| LCA+Concat | 1.049M | 371.86M | 0.397 | 0.386 | 0.201 |
| CA+Concat | 1.063M | 495.41M | 0.418 | 0.422 | 0.217 |
| Concat | - | - | 0.366 | 0.397 | 0.187 |
| Image as query | - | - | 0.404 | 0.374 | 0.193 |

Table 6: Ablation experiments on approaches of vision-language token fusion in RoadMind-8B. (SCA: (Guan et al. 2025b), LCA: (Wang et al. 2020), CA: (Vaswani et al. 2017))

tween precision and computational cost. **Third**, Table 7 shows each settings of AD-CoT matter while AD-CoT performs better than MCoT (Zhang et al. 2024).

| CoT | Perception | | Reasoning | |
|-----------------|--------------|--------------|--------------|--------------|
| | METEOR | GPT-Score | METEOR | GPT-Score |
| AD-CoT | 0.420 | 0.568 | 0.339 | 0.445 |
| Only OQ Input | 0.392 | 0.523 | 0.301 | 0.401 |
| Only GT (train) | 0.428 | 0.549 | 0.325 | 0.439 |
| Only CA (train) | 0.387 | 0.491 | 0.331 | 0.395 |
| MCoT | 0.422 | 0.536 | 0.319 | 0.431 |

Table 7: Ablation studies in CoT. OQ: original question; GT: ground truth; CA: conclusive answer (from GPT-4o).

Visualization Results

Figure 6 presents qualitative reasoning results from RoadMind-8B. The results show that RoadMind is capable of accurately perceiving contextual scene elements and target states, and generating rule-compliant, context-aware responses in accordance with traffic regulations.

Conclusion

We propose RoadSceneVQA, a large-scale VQA dataset tailored for traffic scenario understanding, which encompasses a rich collection of scene-level, event-level, and instance-level perception and reasoning samples. Moreover, we propose RoadMind, a MLLM specifically designed for road-side interaction perception, which integrates two key components: CogniAnchor Fusion (CAF), inspired by human visual cognition for anchoring relevant regions, and Assisted Decoupled Chain-of-Thought (AD-CoT), a reasoning strategy that enhances contextual understanding and inference. Overall, these components significantly improve VQA performance in various complex traffic environments. Our results demonstrate that RoadSceneVQA serves as a challenging and effective benchmark for advancing interactive scene understanding in real-world traffic perception tasks.

Acknowledgment

This research was partially funded by the Postdoctoral Research Project of Guangzhou City with Hong Kong University of Science and Technology (Guangzhou) (Grant No. P00603).

References

Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. In *European conference on computer vision*, 382–398. Springer.

Anderson, P.; He, X.; Buehler, C.; Teney, D.; Johnson, M.; Gould, S.; and Zhang, L. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6077–6086.

Antol, S.; Agrawal, A.; Lu, J.; Mitchell, M.; Batra, D.; Zitnick, C. L.; and Parikh, D. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, 2425–2433.

Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Bejarbaneh, E. Y.; Du, H.; and Naghdy, F. 2024. Exploring shared perception and control in cooperative vehicle-intersection systems: A review. *IEEE transactions on intelligent transportation systems*.

Chang, C.; Zhang, J.; Zhang, K.; Zhong, W.; Peng, X.; Li, S.; and Li, L. 2023. BEV-V2X: Cooperative birds-eye-view fusion and grid occupancy prediction via V2X-based data sharing. *IEEE Transactions on Intelligent Vehicles*, 8(11): 4498–4514.

Chen, K.; Li, Y.; Zhang, W.; Liu, Y.; Li, P.; Gao, R.; Hong, L.; Tian, M.; Zhao, X.; Li, Z.; et al. 2025. Automated evaluation of large vision-language models on self-driving corner cases. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 7817–7826. IEEE.

Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024. Internvl: Scaling up vision foundation models and alignment for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 24185–24198.

Chu, X.; Qiao, L.; Zhang, X.; Xu, S.; Wei, F.; Yang, Y.; Sun, X.; Hu, Y.; Lin, X.; Zhang, B.; et al. 2024. Mobilevlm v2: Faster and stronger baseline for vision language model. *arXiv preprint arXiv:2402.03766*.

Dai, W.; Wu, S.; Wu, W.; Wang, Z.; Lyu, S.; Liao, H.; Yu, L.; Ding, W.; Guan, R.; and Yue, Y. 2025. Large Foundation Models for Trajectory Prediction in Autonomous Driving: A Comprehensive Survey. *arXiv preprint arXiv:2509.10570*.

Deruyttere, T.; Vandenhende, S.; Grujicic, D.; Van Gool, L.; and Moens, M. F. 2019. Talk2Car: Taking Control of Your Self-Driving Car. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2088–2098.

Gopalkrishnan, A.; Greer, R.; and Trivedi, M. 2024. Multi-Frame, Lightweight & Efficient Vision-Language Models for Question Answering in Autonomous Driving. In *First Vision and Language for Autonomous Driving and Robotics Workshop*.

Guan, R.; Hu, R.; Zhou, Z.; Xue, T.; Man, K. L.; Smith, J.; Lim, E. G.; Ding, W.; and Yue, Y. 2025a. Referring flexible image restoration. *Expert Systems with Applications*, 274: 126857.

Guan, R.; Jia, L.; Yao, S.; Yang, F.; Xu, S.; Purwanto, E.; Zhu, X.; Man, K. L.; Lim, E. G.; Smith, J.; et al. 2025b. Watervg: Waterway visual grounding based on text-guided vision and mmwave radar. *IEEE Transactions on Intelligent Transportation Systems*.

Guan, R.; Man, K. L.; Chen, F.; Yao, S.; Hu, R.; Zhu, X.; Smith, J.; Lim, E. G.; and Yue, Y. 2024. FindVehicle and

- VehicleFinder: a NER dataset for natural language-based vehicle retrieval and a keyword-based cross-modal vehicle retrieval system. *Multimedia Tools and Applications*, 83(8): 24841–24874.
- Guan, R.; Man, K. L.; Zhao, H.; Zhang, R.; Yao, S.; Smith, J.; Lim, E. G.; and Yue, Y. 2023. MAN and CAT: mix attention to nn and concatenate attention to YOLO. *The Journal of Supercomputing*, 79(2): 2108–2136.
- Guan, R.; Zhang, R.; Ouyang, N.; Liu, J.; Man, K. L.; Cai, X.; Xu, M.; Smith, J.; Lim, E. G.; Yue, Y.; et al. 2025c. Talk2radar: Bridging natural language with 4d mmwave radar for 3d referring expression comprehension. *2025 IEEE International Conference on Robotics and Automation (ICRA)*.
- Han, D.; Pu, Y.; Xia, Z.; Han, Y.; Pan, X.; Li, X.; Lu, J.; Song, S.; and Huang, G. 2024. Bridging the divide: Reconsidering softmax and linear attention. *Advances in Neural Information Processing Systems*, 37: 79221–79245.
- Huang, T.; Liu, J.; Zhou, X.; Nguyen, D. C.; Azghadi, M. R.; Xia, Y.; Han, Q.-L.; and Sun, S. 2023. V2X cooperative perception for autonomous driving: Recent advances and challenges. *arXiv preprint arXiv:2310.03525*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Kim, B. S.; Kim, J.; Lee, D.; and Jang, B. 2025. Visual question answering: A survey of methods, datasets, evaluation, and challenges. *ACM Computing Surveys*, 57(10): 1–35.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123: 32–73.
- Lavie, A.; and Agarwal, A. 2007. Meteor. 228–231.
- Liang, S.; Guan, R.; Lian, W.; Liu, D.; Sun, X.; Wu, D.; Yue, Y.; Ding, W.; and Xiong, H. 2025. Cognitive Disentanglement for Referring Multi-Object Tracking. *Information Fusion*, 103349.
- Liao, H.; Kong, H.; Wang, B.; Wang, C.; Ye, W.; He, Z.; Xu, C.; and Li, Z. 2025. Cot-drive: Efficient motion forecasting for autonomous driving with llms and chain-of-thought prompting. *IEEE Transactions on Artificial Intelligence*.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Malinowski, M.; Rohrbach, M.; and Fritz, M. 2015. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE international conference on computer vision*, 1–9.
- Muhammad, K.; Hussain, T.; Ullah, H.; Del Ser, J.; Rezaei, M.; Kumar, N.; Hijji, M.; Bellavista, P.; and de Albuquerque, V. H. C. 2022. Vision-based semantic segmentation in scene understanding for autonomous driving: Recent achievements, challenges, and outlooks. *IEEE Transactions on Intelligent Transportation Systems*, 23(12): 22694–22715.
- Nam, H.; Ha, J.-W.; and Kim, J. 2017. Dual attention networks for multimodal reasoning and matching. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 299–307.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Qian, T.; Chen, J.; Zhuo, L.; Jiao, Y.; and Jiang, Y.-G. 2024. Nuscenes-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4542–4550.
- Sima, C.; Renz, K.; Chitta, K.; Chen, L.; Zhang, H.; Xie, C.; Beißwenger, J.; Luo, P.; Geiger, A.; and Li, H. 2024. Drivelm: Driving with graph visual question answering. In *European Conference on Computer Vision*, 256–274. Springer.
- Sun, P.; Song, Y.; Liu, X.; Yang, X.; Wang, Q.; Li, T.; Yang, Y.; and Chu, X. 2024. 3d question answering for city scene understanding. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 2156–2165.
- Sun, P.; Song, Y.; Zhu, X.; Liu, X.; Wang, Q.; Liu, Y.; Xia, C.; Li, T.; Yang, Y.; and Chu, X. 2025. City-VLM: Towards Multidomain Perception Scene Understanding via Multimodal Incomplete Learning. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 3448–3457.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4566–4575.
- Wang, S.; Li, B. Z.; Khabsa, M.; Fang, H.; and Ma, H. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- Wu, D.; Han, W.; Liu, Y.; Wang, T.; Xu, C.-z.; Zhang, X.; and Shen, J. 2025. Language prompt for autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 8359–8367.
- Xie, S.; Kong, L.; Dong, Y.; Sima, C.; Zhang, W.; Chen, Q. A.; Liu, Z.; and Pan, L. 2025. Are VLMs Ready for Autonomous Driving? An Empirical Study from the Reliability, Data, and Metric Perspectives. *arXiv preprint arXiv:2501.04003*.
- Xu, H.; Berres, A.; Tennille, S. A.; Ravulaparthi, S. K.; Wang, C.; and Sanyal, J. 2021. Continuous emulation and multiscale visualization of traffic flow using stationary roadside sensor data. *IEEE Transactions on Intelligent Transportation Systems*, 23(8): 10530–10541.
- Yang, L.; Zhang, R.-Y.; Wang, Y.; and Xie, X. 2024. Mma: Multi-modal adapter for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23826–23837.

Yang, Z.; He, X.; Gao, J.; Deng, L.; and Smola, A. 2016. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 21–29.

Yao, Y.; Yu, T.; Zhang, A.; Wang, C.; Cui, J.; Zhu, H.; Cai, T.; Li, H.; Zhao, W.; He, Z.; et al. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*.

Ye, X.; Shu, M.; Li, H.; Shi, Y.; Li, Y.; Wang, G.; Tan, X.; and Ding, E. 2022. Rope3d: The roadside perception dataset for autonomous driving and monocular 3d object detection task. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21341–21350.

Yu, H.; Luo, Y.; Shu, M.; Huo, Y.; Yang, Z.; Shi, Y.; Guo, Z.; Li, H.; Hu, X.; Yuan, J.; et al. 2022. 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*, 21361–21370.

Zhang, W.; Zhou, Z.; Zheng, Z.; Gao, C.; Cui, J.; Li, Y.; Chen, X.; and Zhang, X.-P. 2025a. Open3dvqa: A benchmark for comprehensive spatial reasoning with multimodal large language model in open space. *arXiv preprint arXiv:2503.11094*.

Zhang, Z.; Li, X.; Xu, Z.; Peng, W.; Zhou, Z.; Shi, M.; and Huang, S. 2025b. Mpdribe: Improving spatial understanding with marker-based prompt learning for autonomous driving. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 12089–12099.

Zhang, Z.; Zhang, A.; Li, M.; Zhao, H.; Karypis, G.; and Smola, A. 2024. Multimodal Chain-of-Thought Reasoning in Language Models. *Transactions on Machine Learning Research*, 2024.

Zhou, B.; Hu, Y.; Weng, X.; Jia, J.; Luo, J.; Liu, X.; Wu, J.; and Huang, L. 2024a. Tinyllava: A framework of small-scale large multimodal models. *arXiv preprint arXiv:2402.14289*.

Zhou, X.; Larintzakis, K.; Guo, H.; Zimmer, W.; Liu, M.; Cao, H.; Zhang, J.; Lakshminarasimhan, V.; Strand, L.; and Knoll, A. C. 2025. TUMTraffic-VideoQA: A Benchmark for Unified Spatio-Temporal Video Understanding in Traffic Scenes. *arXiv preprint arXiv:2502.02449*.

Zhou, X.; Liu, M.; Yurtsever, E.; Zagar, B. L.; Zimmer, W.; Cao, H.; and Knoll, A. C. 2024b. Vision language models in autonomous driving: A survey and outlook. *IEEE Transactions on Intelligent Vehicles*.