



RoadSocial: A Diverse VideoQA Dataset and Benchmark for Road Event Understanding from Social Video Narratives

Chirag Parikh*, Deepti Rawat*, Rakshitha R. T., Tathagata Ghosh, and Ravi Kiran Sarvadevabhatla
CVIT & iHub-Data, IIIT Hyderabad, India

<https://roadsocial.github.io>

Abstract

We introduce RoadSocial, a large-scale, diverse VideoQA dataset tailored for generic road event understanding from social media narratives. Unlike existing datasets limited by regional bias, viewpoint bias and expert-driven annotations, RoadSocial captures the global complexity of road events with varied geographies, camera viewpoints (CCTV, handheld, drones) and rich social discourse. Our scalable semi-automatic annotation framework leverages Text LLMs and Video LLMs to generate comprehensive question-answer pairs across 12 challenging QA tasks, pushing the boundaries of road event understanding. RoadSocial is derived from social media videos spanning 14M frames and 414K social comments, resulting in a dataset with 13.2K videos, 674 tags and 260K high-quality QA pairs. We evaluate 18 Video LLMs (open-source and proprietary, driving-specific and general-purpose) on the RoadSocial-QA benchmark. We also demonstrate RoadSocial’s utility in improving road event understanding capabilities of general-purpose Video LLMs.

1. Introduction

A road event typically refers to any incident, activity, or condition occurring on or around the roadway that affects traffic flow, safety, or road usage. The ability to recognize and interpret road events is essential for safe and reliable intelligent vehicles and transportation systems. In this regard, large-scale video datasets of road events are used to develop assistive models [2, 3, 8, 21, 24, 29]. Many recent datasets contain videos with accompanying question-answer text pairs and other text metadata [16, 22, 23, 27]. Such datasets have become a de facto choice for training Video Large Language Models (Video LLMs) [7, 13, 27, 40].

However, current video-based road event understanding approaches are limited by region-specific datasets, neglect-

ing the diversity of global road scenarios. Most datasets focus on dashcam views for autonomous driving, overlooking other camera types such as CCTV, handheld, and drone-based. They also lack annotations on generic events (e.g. defensive driving, near-misses). Due to the reliance on regionally-biased expert annotators, the broader and richer contextual insights from real-world social discourse on road events are absent. Furthermore, existing evaluation frameworks fail to test the Video LLMs’ ability to distinguish informative road event details from misleading information, essential for developing reliable, hallucination-resistant road event understanding systems.

To address these limitations and to enable foundational video language models for *generic* road event understanding, we introduce **RoadSocial**, a large-scale and diverse Video Question Answer (VideoQA) dataset. RoadSocial is obtained by processing social media videos and the narratives accompanying these videos. The inherent diversity of social media in terms of geographical locations, camera viewpoints, road event types and social commentary addresses shortcomings of video datasets mentioned previously. Specifically, we make the following contributions:

- RoadSocial: a large-scale, diverse VideoQA resource for road events, derived from social media videos spanning **14M** frames and **414K** social comments, resulting in a dataset with **13.2K** videos, **674** unique tags, and **260K** high-quality QA pairs.
- A semi-automatic annotation framework using Text LLM and Video LLM that processes social media video narratives and generates comprehensive QA pairs across 12 distinct challenging tasks.
- A robust evaluation framework incorporating *non-road event* videos and irrelevant questions to assess the robustness of Video LLMs to hallucinations.
- A demonstration of RoadSocial’s utility in improving road event understanding capabilities of general-purpose Video LLM.
- Critical insights into 18 Video LLMs’ performance on road event understanding, obtained from their evaluation

*Equal contribution.

Dataset	Viewpoint Type	Video Frames	Duration (mins)	Social Comments	Countries	Video Tags	QAs	TG QA	AV QA	IC QA	Loc. QA	Internet sourced
RoadSocial (Ours)	6	14M	7.9K	414K	100	674	260K	✓	✓	✓	✓	✓
Lingo-QA [16]	1	.1M	1.8K	-	1	7	419K	-	-	-	-	-
SUTD-TrafficQA [34]	3	10M	6.7K	-	<4	-	62.5K	-	-	-	-	✓
DRAMA [15]	1	.02M	.6K	-	1	-	102K	-	-	-	-	-
Rank2Tell [25]	1	.02M	39	-	1	-	>118	-	-	-	-	-
ROAD [29]	1	.1M	170	-	1	43	-	-	-	-	-	-
MM-AU [5]	1	2.2M	1.2K	-	>50	58	58.6K	✓	-	-	-	✓
DriveLM [27, 40]	1	.03M	5.7K	-	43	-	375K	-	-	-	-	-
BDD-OIA [35]	1	.02M	1.9K	-	1	25	-	-	-	-	-	-
BDD-X [8]	1	8.4M	4.6K	-	1	-	-	-	-	-	-	-

Table 1. **Comparison of RoadSocial with existing road event understanding datasets.** TG: Temporal Grounding, AV: Adversarial, IC: Incompatible, Loc: Geographical Location. Internet-sourced videos do not contain LiDAR or CAN bus data. **Orange**: Additional annotations added to existing datasets. **Blue**: New datasets.

on our RoadSocial-QA benchmark.

2. Related Works

Several video datasets describe road events through actions of surrounding entities [3, 29], interactions between traffic participants [15, 21, 25, 35], or explanations of normal or safety-critical driving scenarios [8, 16, 27], including dangerous driving behaviors or accidents [5, 34].

However, the diversity of these datasets is often limited by their geographical scope. Although some datasets [5, 34, 40] include crowd-sourced videos from a range of locations, their textual annotations reflect local expertise, which may lack a comprehensive understanding of global traffic norms and behaviors. In contrast, our dataset is sourced from global social media video posts which addresses this shortcoming. Existing works typically rely on a pool of manual annotators, a process that is labor-intensive and lacks scalability [5, 16, 27, 34, 35]. We propose a scalable, semi-automatic annotation framework that leverages the capabilities of powerful Video LLMs and Text LLMs to process social media content from around the world and generate high-quality QA pairs associated with road event videos.

Existing video language models built on previous road understanding benchmarks are often trained on specific camera viewpoints, usually vehicle-mounted [5, 15, 16, 25, 27, 31, 40] or CCTV [9]. Such models may not generalize well across different viewpoint types and geographical regions, limiting their effectiveness for understanding road events in a broad context. In contrast, our dataset contains videos captured in diverse and uncontrolled camera settings (drone, handheld, CCTV etc.) across the world. Coupled with social discourse, our dataset is a viable alternative.

Existing VideoQA datasets focus primarily on ego-centric tasks [5, 15, 16, 25, 27, 40], limiting perspectives to the ego-vehicle. While Xu *et al.* [34] explore complex traffic scenarios, none of the existing works assess model robustness against misleading inputs or hallucinations. We

address this gap by introducing novel QA tasks to evaluate (a) robustness to hallucinations with non-road-event videos and irrelevant questions (b) comprehension across camera viewpoint types and (c) geographical awareness. These tasks enable holistic Video LLM evaluation for *general-purpose* road event understanding. A comparison of RoadSocial with existing datasets is shown in Tab. 1.

3. RoadSocial Dataset

RoadSocial is a dataset created from social media videos in unconstrained, real-world environments. These videos are accompanied by rich social commentary that reflects facts and varied cultural perspectives on road events worldwide.

3.1. Data Collection

We crowdsourced diverse road event data from X (formerly Twitter), leveraging its global community for real-world insights. Unlike other platforms, X is characterized by an active social discourse on road events that includes the general public, road event enthusiasts, and road enforcement authorities. Our strategy focused on popular road event related social media handles worldwide, using multilingual keywords to scrape tweet data from 2012 onwards, filtering for videos with substantial commentary. The resulting dataset captures varied road events—traffic violations, accidents, safe driving, and infrastructure awareness—across different environments and locations.

3.2. Annotation Strategy: QAs and Tags

Our annotation strategy merges LLM-based automation with expert verification to produce high-quality QA pairs and video tags. We start by identifying representative road event samples, then use a hybrid approach to generate QA pairs that blend video semantics with social media context. QA pairs are refined, categorized into predefined tasks, and supplemented with video-level tags, all verified by experts. Additionally, we create incompatible QA pairs for non-road

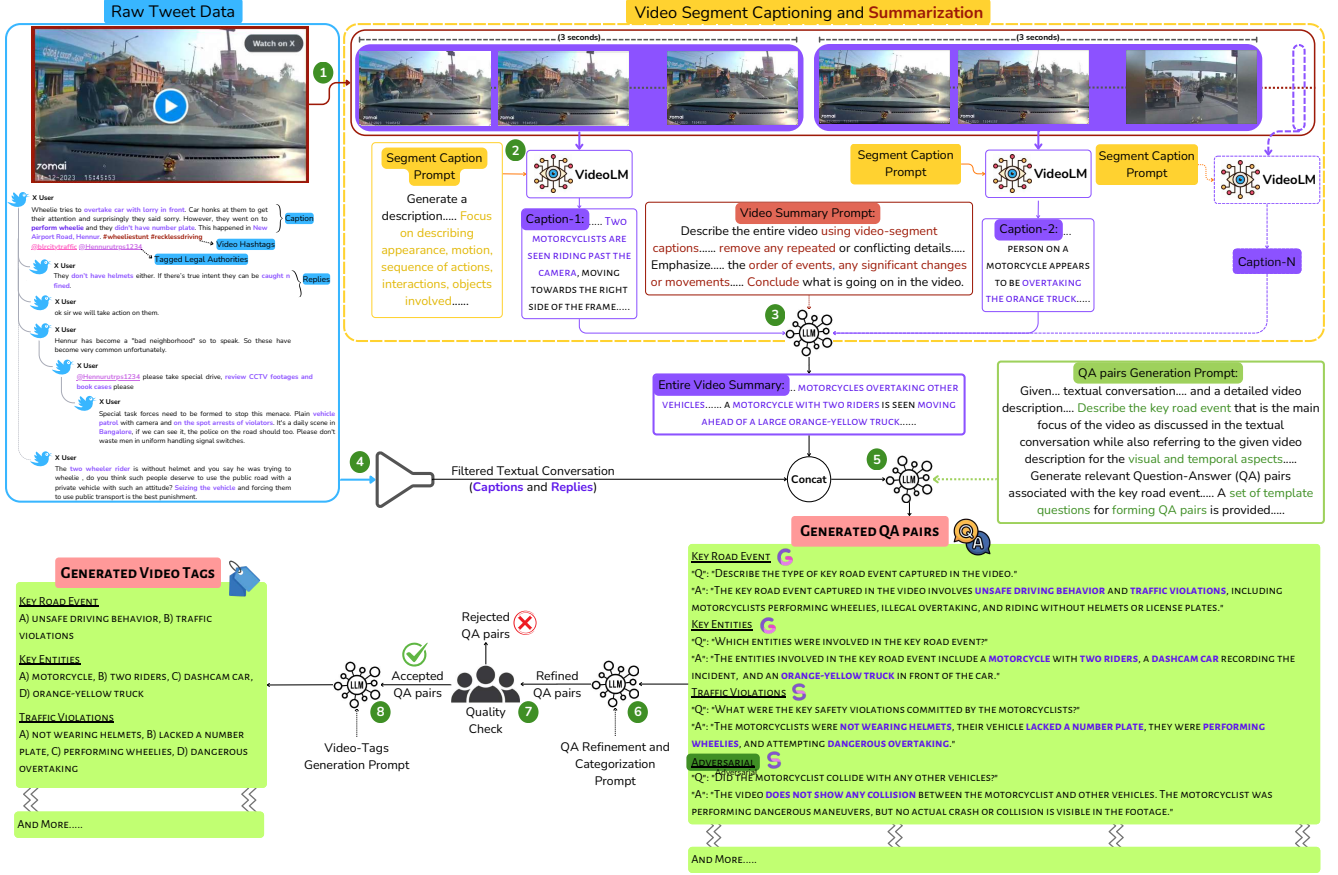


Figure 2. **RoadSocial Annotation Pipeline:** The steps involved in the annotation pipeline are depicted from 1 to 8. **Raw Tweet Data** consists of the video and the Twitter conversation. Step 1 includes splitting the video into 3-second segments (in purple shaded boxes). Step 2 involves feeding the video segments to Video LLM and prompting it to generate corresponding captions numbered from 1 to N. These captions are aggregated and summarized by an LLM to generate **entire video summary** in Step 3. Step 4 filters the raw tweet textual data and extracts the captions, replies, hashtags, and tagged legal authorities’ user handles (highlighted in blue). This filtered conversation data and the entire video visual summary are fed to LLM and prompted to generate generic (G) and specific (S) QA pairs in Step 5. All important aspects of the key road event mentioned in the raw tweet text, video segment captions, the entire video summary, and the generated QA pairs are highlighted in purple. The **generated QA pairs** are refined and categorized into pre-defined tasks in step 6. These QA pairs are verified by expert annotators to either include or exclude them from the dataset in Step 7. The human-verified QA pairs are then used as input to **generate video-level tags** in Step 8.

undergo final quality checks for relevance to the video (7). Detailed prompts and sample outputs are provided in Appendix B.5.

Video-level Tag Generation: To categorize videos by key aspects of road events, we generate diverse video-level tags (e.g. traffic violation, wheelie, unsafe overtaking) using verified answers from step 7. A Text LLM [1] scans these answers to generate top-k tags most relevant to each QA task (8). This structured tagging approach ensures that the generated QA pairs, tags are robust and reflect the diverse scenarios present in the dataset. Details about the tag generation prompt and the resulting tags distribution are provided in Appendix B.10.

Incompatible QA Generation: To assess the reliability

and resistance of Video LLMs to hallucinations, we generate incompatible QA pairs for non-road event videos. This involves sampling questions from road event QA pairs to create mismatched questions for unrelated videos. Answers are generated using the Hybrid Approach mentioned previously, with modified prompts treating these mismatched questions as templates. Further details on the prompt modifications are provided in Appendix B.6.

3.3. Dataset Statistics

Our final dataset comprises over **14M** video frames from more than **13.2K** videos (totaling **7.9K** minutes of video footage) with **260K** QA pairs and **674** unique video tags (total **100K+**).

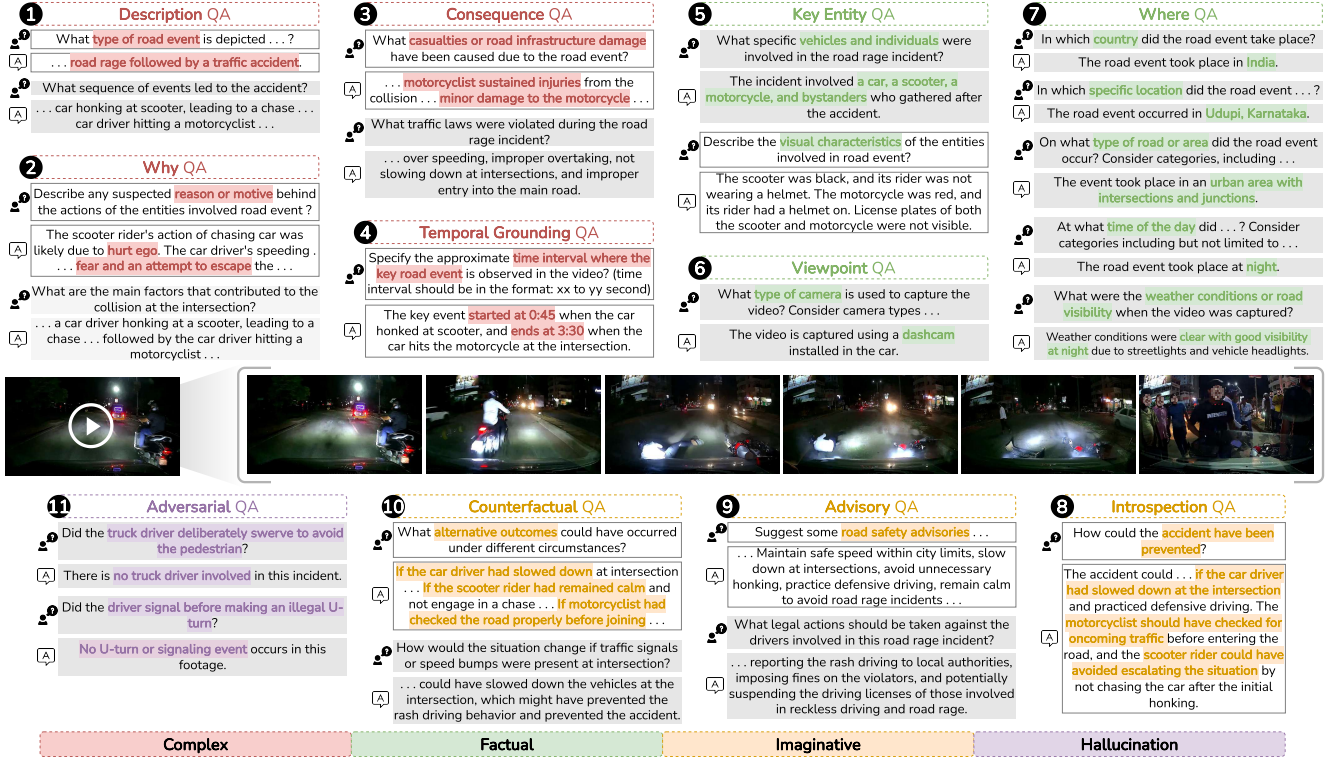


Figure 3. Examples of QA Pairs grouped by tasks and color-coded by task category. Gray outlined questions are generic while gray fill shading indicates specific questions. Highlighted text indicates key information. (Sec. 3.4).

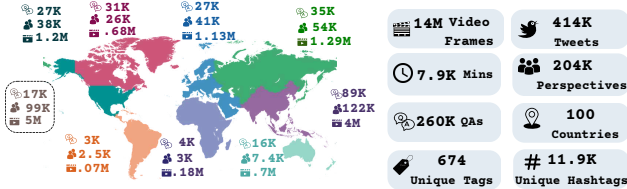


Figure 4. **The diversity of RoadSocial dataset:** The number of QA pairs, social commentary (tweets), and video frames spread across different regions is shown. Overall statistics of the raw tweet data, generated QA pairs, and tags in our dataset is also shown. Total incompatible QA pairs and related numbers for non-road event videos are specified inside a light brown box at left.

The dataset exhibits significant diversity across several dimensions, including geographical distribution (Fig. 4), QA types (Fig. 3), and video tags (Fig. 1). Fig. 4 shows the global coverage of our dataset attributes, depicting the diverse perspectives involved in the QA pair generation process.

It includes **414K** multilingual tweet captions and replies corresponding to **204K** unique user handles (from across **100** countries) sharing facts and opinions about the road or traffic events. Tab. 1 compares key attributes of our dataset with related road event understanding datasets. The distri-

bution of QA pairs corresponding to each task category is shown in Fig. 5. The distribution of video tags along different attributes is shown by word clouds in Fig. 1. The videos durations range from 0.13 seconds to 3885.44 seconds with an average of 35.6 seconds.

3.4. QA Tasks Taxonomy

We developed a question-answer (QA) taxonomy for structured evaluation of Video Large Language Models (Video LLMs). The taxonomy consists of 12 distinct tasks organized into four reasoning categories: Complex, Factual, Imaginative, and Hallucination (Fig. 5). These categories assess various aspects of road events, ranging from key entity identification (see 5 in Fig. 3) to hypothetical scenario exploration (Fig. 3 10). Our taxonomy extends beyond conventional road datasets by incorporating previously under-represented tasks, such as Viewpoint QA (analysis of camera perspectives capturing road events) and Where QA (geographic location identification of road events). As an additional novelty, our approach uniquely incorporates Adversarial QA and Incompatible QA. Adversarial QA tests a model's ability to recognize and reject misleading assumptions or false details in questions by identifying non-occurring road events *e.g.* Fig. 3 11). Incompatible QA on non-road-event videos helps evaluate models' robustness to



Figure 5. **QA Task Taxonomy:** The QA pairs in RoadSocial are broadly grouped into 4 categories (highlighted in blue) which are further subdivided into 12 tasks (shown in green). Total QA pair count for each category is shown in blue squared box. Some of these tasks are further subdivided into granular sub-tasks (highlighted in orange) to facilitate coarse to fine-grained understanding of road events along different aspects.

hallucination by identifying irrelevant video-question pairs. Fig. 3 illustrates representative QA pairs for each category including the generic and specific questions (described in Sec. 3.2). For a detailed QA task description, please refer to Appendix B.9.

4. Experiments

We evaluate a wide range of Video LLMs (both open-source and proprietary, driving-specific and general-purpose) on our road event understanding benchmark.

4.1. Data Setup

Evaluation Benchmark: RoadSocial-QA consists of **13.2K** videos encompassing **260K** QA pairs, with an average of **20** QA pairs per video. To evaluate zero-shot reasoning capabilities of Video LLMs, we split our dataset into **12K** training and **1.2K** test videos, resulting in **234K** and **26K** QA pairs respectively. The video splits maintain geographical diversity across the dataset, with the test set serving as our primary evaluation benchmark.

For model evaluation, we provide the model with video frames and a task-specific question, following the format: video frames + model’s default system prompt (if any) + our task-specific question (Fig. 6). Detailed prompting structures are described in Appendix C.1.

4.2. Model Setup

Our evaluation encompasses 18 Video LLMs, comprising 15 open-source general-purpose models, 2 proprietary general-purpose models, and 1 open-source driving-specific model. We evaluate their zero-shot performance on the test split of RoadSocial-QA using each model’s official configuration for open-ended response generation. The results, pre-

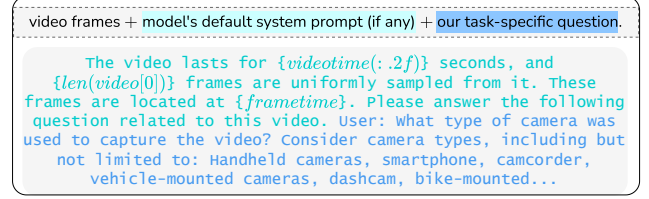


Figure 6. An example of prompting a Video LLM [39].

sented in Tab. 2, analyze model performance across different tasks. All evaluation runs were conducted on a computing cluster equipped with NVIDIA H100 GPUs. Detailed information about model configurations, prompting templates and evaluation timelines is provided in Appendix C.2.

4.3. Evaluation Metrics

To assess the similarity between model-generated and ground-truth open-ended responses, we adopt GPT-3.5 score [18] as our primary evaluation metric for all tasks (except Temporal Grounding), following established practices in recent literature [12, 14, 27, 36]. To ensure statistical robustness, we conduct multiple evaluation runs and report mean of the GPT-3.5 scores. For Temporal Grounding QAs, time interval is extracted from the model-generated response and compared with the ground-truth time interval range using the mean Average Precision (mAP) evaluation metric. The complete evaluation protocols, including prompt templates, and scoring criteria are provided in Appendix C.2.

4.4. Analysis

Overall Performance Trends: Refer to last 3 columns of Tab. 2. Tarsier-34B [32] achieves the highest overall score (63.5) across ALL QA tasks whereas IXC-2.5-7B [38] leads the benchmark on road-event related tasks (RT) (66.4) among open-source models. These models even outperform larger models such as InternVL2-76B [4] and Qwen2-VL-72B [33]. Additionally, all general-purpose models surpass driving-specific Video LLM across all tasks, revealing significant performance gaps in general road event understanding within the driving-focused model. Predictably, Video LLMs face greater difficulty with Generic QAs compared to Specific QAs because generic questions require the model to infer the context independently, unlike specific questions. Among closed-source models, GPT-4o [19] stands out as the top performer, achieving the highest scores across all models. A radar plot with representative Video LLMs can be viewed in Fig. 7.

Performance Across Task Categories: The analysis reveals distinct patterns across different reasoning categories, highlighting strengths and weaknesses among models in various types of reasoning tasks.

Model	Params	Factual			Complex				Imaginative			Hallucination		Overall	Overall	Overall	Overall
		WR	KE	VP	DS	WY	CQ	TG	AD	IN	CF	AV	IC	(ALL)	(RT)	(Generic)	(Specific)
Dolphin [13]	9B	61.3	34.5	67.8	35.8	25.2	37.2	0.01	49.8	39.1	45.5	71.8	21.3	40.8	42.5	29.8	46.5
GPT-4o [19]	-	77.0	66.6	84.3	70.2	70.8	72.1	7.8	77.7	76.4	77.0	90.0	67.6	69.8	70.0	69.5	74.4
Gemini-1.5-Pro [30]	-	77.7	56.7	85.4	61.9	61.4	60.1	18.6	72.1	70.2	75.7	72.3	48.7	63.4	64.7	60.1	68.3
InternVL2 [4]	76B	72.4	51.3	81.4	57.1	59.0	62.1	1.07	70.5	67.0	69.2	58.6	27.6	56.4	59.1	55.5	65.1
Qwen2-VL [33]	72B	76.8	56.6	85.1	60.2	64.0	67.6	0.01	71.9	72.4	71.6	37.0	40.2	58.6	60.3	58.3	68.8
LLaVA-Video [39]	72B	75.8	52.4	76.8	52.4	55.0	52.2	9.94	68.3	63.7	64.9	83.5	24.7	56.7	59.6	51.1	63.3
LLaVA-OV [10]	72B	75.1	54.1	78.7	53.0	53.3	54.1	3.99	67.8	61.9	63.1	45.1	19.9	52.5	55.5	51.8	63.0
VITA [6]	8x7B	66.6	52.1	71.6	48.1	55.6	56.3	2.27	66.7	66.0	62.4	56.3	22.0	52.2	54.9	49.8	60.4
Tarsier [32]	34B	73.7	58.1	78.2	58.2	59.0	58.8	0.32	71.6	71.1	67.4	83.2	82.3	63.5	61.8	58.4	66.1
ARIA [11]	25.3B	75.4	53.1	86.2	58.4	56.9	70.2	8.96	75.1	74.7	74.0	86.4	29.2	62.4	65.4	56.7	68.5
InternVL2 [4]	8B	67.7	51.7	78.0	55.7	59.3	60.9	0.77	66.7	66.8	70.0	68.1	26.1	56.0	58.7	53.7	64.0
Mini-CPM-V 2.6 [37]	8B	77.7	57.6	80.6	55.0	50.5	57.5	0.4	61.6	52.3	59.3	73.5	30.0	54.7	56.9	51.0	62.0
IXC-2.5 [38]	7B	78.5	58.7	85.4	61.7	65.3	68.5	0.69	73.9	75.6	75.7	85.8	29.2	63.3	66.4	60.7	70.3
Tarsier [32]	7B	69.9	54.7	72.3	52.0	53.4	55.2	0.11	69.5	69.3	63.5	79.1	67.3	58.9	58.1	54.0	61.7
LongVU [26]	7B	73.0	53.0	76.3	51.1	50.2	55.0	0.84	59.7	55.8	58.2	48.9	32.7	51.2	52.9	47.7	59.7
Qwen2-VL [33]	7B	75.5	52.8	76.1	52.7	57.7	56.4	0.59	69.2	71.6	65.9	37.5	39.6	54.6	56.0	52.6	63.9
LLaVA-Video [39]	7B	74.6	50.1	76.7	52.1	50.1	50.3	1.43	60.4	53.8	58.7	61.8	23.5	51.1	53.6	47.6	59.7
LLaVA-OV [10]	7B	73.4	51.2	77.2	50.7	51.7	51.2	0.97	62.8	55.4	58.6	45.4	21.1	50.0	52.6	48.4	59.8
LLaVA-OV ft.	7B	80.9	64.1	85.7	64.1	68.7	65.1	4.49	74.2	70.9	71.7	95.4	87.6	69.4	67.8	65.1	69.7

Table 2. **Video LLMs benchmarked on RoadSocial-QA.** Standard prompting with task-specific instructions were employed for zero-shot evaluation of Video LLMs on 12 QA tasks. Video LLMs are grouped as open-source (**driving-specific** and **general-purpose**), and closed-source models. Further, we **fine-tune** a Video LLM - LLaVA-OV-7B and report its performance at the end of the table. Abbreviations used for QA tasks include Factual (F), Complex (C), Imaginative (I), Hallucination (H), Where (WR), Key Entities (KE), Viewpoint (VP), Description (DS), Why (WY), Consequence (CQ), Temporal Grounding (TG), Advisory (AD), Introspection (IN), Counterfactual (CF), Adversarial (AV), Incompatible (IC), and Road-event related Tasks (RT). RT includes all tasks except IC which corresponds to non-road event videos. GPT-3.5 score is reported for all tasks except Temporal Grounding (TG) for which average mAP@.3:7 (%) is reported. Overall average scores are reported for ALL QA tasks (F, C, I, and H), Road-event related Tasks (RT), Generic QAs, and Specific QAs under each task. All reported scores (scale 0 to 100) are colored based on their value from **low** to **high**. VideoLLMs show per-query latencies of 1-25s (7B-76B) on H100 GPUs.

In factual reasoning, models perform well in Where (WR) and Viewpoint (VP) QA tasks, both of which yield consistently high scores. For VP tasks, this may be partly due to our prompt that offers a limited set of viewpoint options, essentially transforming the question into a multiple-choice format rather than a free-form open-ended question. Empirically, performance declines when these choices are absent from the prompt, as noted in Appendix C.1. Meanwhile, WR tasks perform well due to their inherently specific questions.

Most Video LLMs encounter difficulties with complex reasoning tasks, such as Description (DS), Why (WY), Consequence (CQ), and Temporal Grounding (TG) reasoning, as well as Key Entity (KE) tasks. These results indicate that many models struggle with identifying key road event that is the main focus of the video.

Temporal Grounding (TG) proves to be particularly challenging, with most models achieving average mAP scores below 1%, highlighting a major limitation in temporal localization for Video LLMs. The highest-performing model, Gemini-1.5-Pro [30], achieves 18.6%. In comparison,

LLaVA-Video-72B [39] leads among open-source models with 9.94%, potentially benefiting from its default prompt, which incorporates time-based instructions (Fig. 6). Empirical analysis shows two common reasons for TG underperformance: some models, such as Tarsier-34B [32], struggle with instruction following, leading to unexpected or incoherent answers, while others, such as LLaVA-OV [10], lack the capability to associate the sequence of events with time in the video (details in Appendix C.3).

In imaginative reasoning, models show promising capabilities, with several models achieving over 70% accuracy in Advisory (AD) and Introspection (IN) tasks. This indicates that models can effectively use their pre-trained knowledge to reason about hypothetical scenarios.

Robustness and Hallucination Assessment: The evaluation of model robustness through Adversarial (AV) and Incompatible (IC) QAs reveals interesting behavioral patterns. Some models, such as GPT-4o [19] and IXC-2.5-7B [38] demonstrate exceptional robustness to adversarial queries, suggesting effective mechanisms for identifying misleading information. However, most models struggle on

Incompatible QAs indicating their tendency to generate hallucinated responses for irrelevant Video and QA pairs. Notably, Tarsier-34B [32] outperforms all models by a good margin indicating inherent capability to identify misleading information and reject out-of-domain queries.

Error Analysis and Future Directions: (1) *Temporal Confusion*: Models frequently struggle with temporal localization, particularly evident in the poor Temporal Grounding (TG) scores. (2) *Complex Reasoning Gaps*: While many models perform well in factual reasoning tasks, they often struggle with QAs requiring in-depth contextual understanding. (3) *Context Integration*: The observed performance gap between Generic and Specific QAs suggests that models struggle to autonomously infer context for generic questions. Future models could benefit from improved mechanisms to integrate prior domain knowledge with visual data for more accurate general context recognition. (4) *Hallucination in Response Generation*: Although some models demonstrate resilience to adversarial queries, hallucination remains a problem for Incompatible (IC) QAs, where irrelevant or out-of-domain answers are generated. Enhancing model training with stricter grounding mechanisms may reduce hallucinations, especially when faced with ambiguous or misleading inputs.

RoadSocial improves road event understanding capability of general-purpose Video LLM: We utilize the train split of our dataset and fine-tune a general-purpose Video LLM. Specifically, we selected LLaVA-OV-7B [10] parameter model as our baseline and employed standard instruction fine-tuning strategy wherein QA pairs are structured into instruction-tuned triplets (question, video, response). We adhere to the official training guidelines and optimized the model using a global batch size of 16 distributed over 16 NVIDIA H100 GPUs. During this phase, all key components (Vision tower, MLP adapter and LLM) were fine-tuned to optimize performance. Our evaluation results (last row of Tab. 2) shows that the fine-tuned LLaVA-OV-7B [10] model attains a significant jump of **19.4%** in overall average score across all QA tasks and stands on par with the best performers. Specifically, the performance gains are significant across complex reasoning (DS, WY, CQ), introspection (IN), and Hallucination (AV, IC) tasks, showcasing our dataset’s utility for improving road event understanding capabilities of general-purpose Video LLM.

Ethical and Privacy Considerations: Our data collection adheres to ethical guidelines, using only publicly available social media content. Our QA generation process includes rigorous checks to exclude potentially harmful, biased or inappropriate content in QA pairs, ensuring the dataset supports fair and responsible research in road event understanding. Additional details can be found in Appendix B.5.

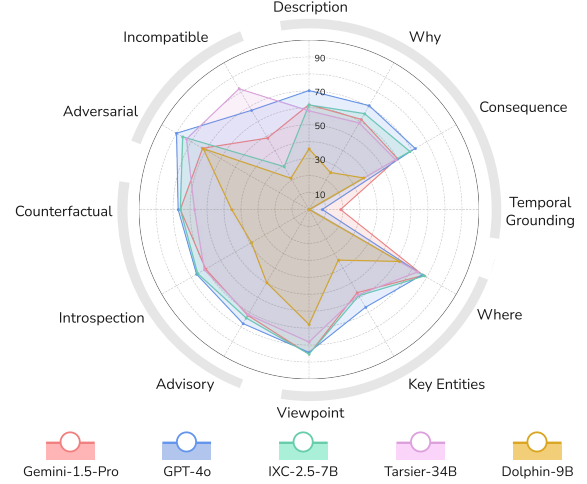


Figure 7. Comparison between representative Video LLMs on RoadSocial benchmark across different QA tasks.

5. Conclusion

RoadSocial redefines the landscape for general-purpose road event understanding. With a first-of-its-kind VideoQA dataset spanning **14M** frames and **414K** social comments, our dataset provides **13.2K** videos with **260K** high-quality QA pairs and **674** unique video tags (total **100K+**). By capturing diverse camera viewpoints, geographical contexts, and socially-informed QAs, RoadSocial delivers a comprehensive dataset that captures the complexity of real-world road scenarios across varied cultural and environmental contexts. Leveraging social media content, it addresses the limitations of traditional datasets by incorporating unique perspectives and nuanced social discourse. Our scalable semi-automatic annotation framework, powered by Text and Video LLMs, supports the creation of rich QA pairs across 12 challenging tasks. Given its scalable nature, our annotation framework can easily ingest and process social media posts generated continuously over time, enabling even larger dataset size with sustained quality. Our robust evaluation framework tests model resilience to irrelevant inputs, hallucinations, cross-viewpoint comprehension, and geographical awareness. Our evaluation across 18 Video LLMs provides critical performance insights across a spectrum of road event QA tasks.

While RoadSocial is a rich resource for road event understanding, its reliance on social media data may introduce biases, skewing coverage towards regions with higher social media use. Apart from addressing these concerns, we envision several future directions for expanding RoadSocial such as increasing language diversity and establishing additional benchmark tasks. We believe RoadSocial will be instrumental in driving progress towards safer and more inclusive intelligent transportation systems.

References

- [1] Anthropic. Claude 3.5 model. <https://www.anthropic.com/news/claude-3-5-sonnet>, 2024. Accessed: 2024-11-08. 3, 4, 13, 31
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 1
- [3] Rohan Chandra, Xijun Wang, Mridul Mahajan, Rahul Kala, Rishitha Palugulla, Chandrababu Naidu, Alok Jain, and Dinesh Manocha. Meteor: A dense, heterogeneous, and unstructured traffic dataset with rare behaviors. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9169–9175. IEEE, 2023. 1, 2
- [4] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Zhong Muyan, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Intern vl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 24185–24198, 2023. 6, 7
- [5] Jianwu Fang, Lei-lei Li, Junfei Zhou, Junbin Xiao, Hongkai Yu, Chen Lv, Jianru Xue, and Tat-Seng Chua. Abductive ego-view accident video understanding for safe driving perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22030–22040, 2024. 2
- [6] Chaoyou Fu, Haojia Lin, Zuwei Long, Yunhang Shen, Meng Zhao, Yifan Zhang, Xiong Wang, Di Yin, Long Ma, Xiawu Zheng, et al. Vita: Towards open-source interactive omni multimodal llm. *arXiv preprint arXiv:2408.05211*, 2024. 7
- [7] Akshay Gopalkrishnan, Ross Greer, and Mohan M. Trivedi. Multi-frame, lightweight & efficient vision-language models for question answering in autonomous driving. *ArXiv*, abs/2403.19838, 2024. 1
- [8] Jinkyu Kim, Anna Rohrbach, Trevor Darrell, John Canny, and Zeynep Akata. Textual explanations for self-driving vehicles. In *Proceedings of the European conference on computer vision (ECCV)*, pages 563–578, 2018. 1, 2
- [9] Quan Kong, Yuki Kawana, Rajat Saini, Ashutosh Kumar, Jingjing Pan, Ta Gu, Yohei Ozao, Bal’azs Opra, D. Anastasiu, Yoichi Sato, and Norimasa Kobori. Wts: A pedestrian-centric traffic video dataset for fine-grained spatial-temporal understanding. *ArXiv*, abs/2407.15350, 2024. 2
- [10] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024. 7, 8, 31, 32
- [11] Dongxu Li, Yudong Liu, Haoning Wu, Yue Wang, Zhiqi Shen, Bowen Qu, Xinyao Niu, Guoyin Wang, Bei Chen, and Junnan Li. Aria: An open multimodal native mixture-of-experts model. *arXiv preprint arXiv:2410.05993*, 2024. 7
- [12] Stephanie Lin, Jacob Hilton, and Owain Evans. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*, 2021. 6
- [13] Yingzi Ma, Yulong Cao, Jiachen Sun, Marco Pavone, and Chaowei Xiao. Dolphins: Multimodal language model for driving. *arXiv preprint arXiv:2312.00438*, 2023. 1, 7
- [14] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. 6
- [15] Srikanth Malla, Chiho Choi, Isht Dwivedi, Joon Hee Choi, and Jiachen Li. Drama: Joint risk localization and captioning in driving. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1043–1052, 2023. 2
- [16] 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. *arXiv preprint arXiv:2312.14115*. 1, 2
- [17] Ana-Maria Marcu, Long Chen, Jan Hünemann, Alice Karnsund, Benoit Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, et al. Lingoqa: Visual question answering for autonomous driving. In *ECCV*, 2024. 31, 32
- [18] OpenAI. Openai: Introducing chatgpt. <https://openai.com/blog/chatgpt/>, 2022. 6, 9, 31
- [19] OpenAI. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. 6, 7
- [20] OpenAI. Vector embeddings. <https://platform.openai.com/docs/guides/embeddings>, 2024. Accessed: 2024-11-08. 3
- [21] Chirag Parikh, Rohit Saluja, C. V. Jawahar, and Ravi Kiran Sarvadevabhatla. Idd-x: A multi-view dataset for ego-relative important object localization and explanation in dense and unstructured traffic. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14815–14821, 2024. 1, 2
- [22] Ehsan Qasemi, Jonathan M Francis, and Alessandro Oltramari. Traffic-domain video question answering with automatic captioning. *ArXiv*, abs/2307.09636, 2023. 1
- [23] Tianwen Qian, Jingjing Chen, Linhai Zhuo, Yang Jiao, and Yu-Gang Jiang. Nuscenescs-qa: A multi-modal visual question answering benchmark for autonomous driving scenario. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4542–4550, 2024. 1
- [24] Vasili Ramanishka, Yi-Ting Chen, Teruhisa Misu, and Kate Saenko. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7699–7707, 2018. 1
- [25] Enna Sachdeva, Nakul Agarwal, Suhas Chundi, Sean Roelofs, Jiachen Li, Mykel Kochenderfer, Chiho Choi, and Behzad Dariush. Rank2tell: A multimodal driving dataset for joint importance ranking and reasoning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7513–7522, 2024. 2
- [26] Xiaoqian Shen, Yunyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, et al. Longvu: Spatiotemporal adaptive compression for long video-language understanding. *arXiv preprint arXiv:2410.17434*, 2024. 7

- [27] Chonghao Sima, Katrin Renz, Kashyap Chitta, Li Chen, Hanxue Zhang, Chengen Xie, Ping Luo, Andreas Geiger, and Hongyang Li. Drivelm: Driving with graph visual question answering. *arXiv preprint arXiv:2312.14150*, 2023. 1, 2, 6
- [28] 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*, 2024. 31, 32
- [29] Gurkirt Singh, Stephen Akrigg, Manuele Di Maio, Valentina Fontana, Reza Javanmard Alitappeh, Salman Khan, Suman Saha, Kossar Jeddissaravi, Farzad Yousefi, Jacob Culley, et al. Road: The road event awareness dataset for autonomous driving. *IEEE transactions on pattern analysis and machine intelligence*, 45(1):1036–1054, 2022. 1, 2
- [30] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 7
- [31] Xiaoyu Tian, Junru Gu, Bailin Li, Yicheng Liu, Chenxu Hu, Yang Wang, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. Drivevlm: The convergence of autonomous driving and large vision-language models. *ArXiv*, abs/2402.12289, 2024. 2
- [32] Jiawei Wang, Liping Yuan, and Yuchen Zhang. Tarsier: Recipes for training and evaluating large video description models. *ArXiv*, abs/2407.00634, 2024. 6, 7, 8
- [33] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 3, 6, 7, 13
- [34] Li Xu, He Huang, and Jun Liu. Sutd-trafficqa: A question answering benchmark and an efficient network for video reasoning over traffic events. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9878–9888, 2021. 2
- [35] Yiran Xu, Xiaoyin Yang, Lihang Gong, Hsuan-Chu Lin, Tz-Ying Wu, Yunsheng Li, and Nuno Vasconcelos. Explainable object-induced action decision for autonomous vehicles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9523–9532, 2020. 2
- [36] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K. Wong, Zhenguo Li, and Hengshuang Zhao. Drivegpt4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 9:8186–8193, 2023. 6
- [37] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024. 7
- [38] Pan Zhang, Xiao wen Dong, Yuhang Zang, Yuhang Cao, Rui Qian, Lin Chen, Qipeng Guo, Haodong Duan, Bin Wang, Linke Ouyang, Songyang Zhang, Wenwei Zhang, Yining Li, Yang Gao, Peng Sun, Xinyue Zhang, Wei Li, Jingwen Li, Wenhai Wang, Hang Yan, Conghui He, Xingcheng Zhang, Kai Chen, Jifeng Dai, Yu Qiao, Dahua Lin, and Jiaqi Wang. Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output. *ArXiv*, abs/2407.03320, 2024. 6, 7
- [39] Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data, 2024. 6, 7, 3, 31
- [40] Yunsong Zhou, Linyan Huang, Qingwen Bu, Jia Zeng, Tianyu Li, Hang Qiu, Hongzi Zhu, Minyi Guo, Yu Qiao, and Hongyang Li. Embodied understanding of driving scenarios. *arXiv preprint arXiv:2403.04593*, 2024. 1, 2



RoadSocial: A Diverse VideoQA Dataset and Benchmark for Road Event Understanding from Social Video Narratives

Supplementary Material

Contents

1. Introduction	1
2. Related Works	2
3. RoadSocial Dataset	2
3.1. Data Collection	2
3.2. Annotation Strategy: QAs and Tags	2
3.3. Dataset Statistics	4
3.4. QA Tasks Taxonomy	5
4. Experiments	6
4.1. Data Setup	6
4.2. Model Setup	6
4.3. Evaluation Metrics	6
4.4. Analysis	6
5. Conclusion	8
List of Figures	3
A Data Collection	9
B Annotation Strategy: QAs and Tags	9
B.1. Identifying Representative Road Event Samples	9
B.2. Template Question Generation	9
B.3. QA Generation via Hybrid Approach	13
B.4. Specific QA Generation	13
B.5. QA Refinement and Categorization	14
B.6. Incompatible QA Generation	24
B.7. Adversarial QA Generation	28
B.8. Temporal Grounding QA generation	31
B.9. Final QA Task Taxonomy	31
B.10 Video-level Tag Generation	31
C Experiments	31
C.1. Data Setup	31
C.2. Model Setup	31
C.3. Qualitative Analysis	31
C.4. RoadSocial’s Utility for Planning/AV tasks	31
C.5. Video-only QAs falls short	32
D Acknowledgement	32

List of Figures

- Diverse Video Attributes in RoadSocial Dataset:** The total count of unique tags for each attribute is shown in circled boxes, alongside word clouds highlighting these values. For each attribute, we display examples with 2-3 keyframes from videos. The figure captures the diversity of road events, environmental conditions, geographical locations, viewpoints, interactions between road entities, and traffic violations. The varied scenarios under each attribute showcase the rich complexity of our dataset.

- RoadSocial Annotation Pipeline:** The steps involved in the annotation pipeline are depicted from ① to ⑧. Step ① includes splitting the video into 3-second segments (in purple shaded boxes). Step ② involves feeding the video segments to Video LLM and prompting it to generate corresponding captions numbered from 1 to N. These captions are aggregated and summarized by an LLM to generate entire video summary in Step ③. Step ④ filters the raw tweet textual data and extracts the captions, replies, hashtags, and tagged legal authorities' user handles (highlighted in yellow) from conversation data and the entire video visual summary are fed to LLM and prompted to generate generic (G) and specific (S) QA pairs in Step ⑤. All important aspects of the key road event mentioned in the raw tweet text, video segment captions, the entire video summary, and the generated QA pairs are highlighted in purple. The generated QA pairs are categorized into pre-defined tasks in step ⑥. These QA pairs are verified by expert annotators to either include or exclude them from the dataset in Step ⑦. The human-verified QA pairs are then used as input to GPT-4o mini to generate relevant tweets based on the provided context and task instructions.

- Examples of QA Pairs grouped by tasks and color-coded by task category. Gray outlined questions are generic while gray fill shading indicates specific questions. Highlighted text indicates key information. (Sec. 3.4).

- The diversity of RoadSocial dataset:** The number of QA pairs, social commentary (tweets), and video frames spread across different regions is shown. Overall statistics of the raw tweet data, generated QA pairs, and tags in our dataset is also shown. Total incompatible QA pairs and related numbers for non-road event videos are specified inside a light brown box at left.

- QA Task Taxonomy:** The QA pairs in RoadSocial are broadly grouped into 4 categories (highlighted in blue) which are further subdivided into 12 tasks (shown in green). Total QA pair count for each category is shown in blue squared box. Some of these tasks are further subdivided into granular sub-tasks (highlighted in orange) to facilitate coarse to fine-grained understanding of road events along different aspects.

- An example of prompting a Video LLM [39].

- Comparison between representative Video LLMs on RoadSocial benchmark across different QA tasks.

- Multilingual Traffic Keyword Dictionary for Tweet Mining:** A comprehensive dictionary of traffic-related keywords and hashtags, designed for identifying road event content on social media. Terms span traffic incidents, emergency services, recording devices, and location-specific templates. Effective usage involves combining terms across categories [*Traffic_General*] + [*Incidents*] and creating location-specific searches.

- Visualization of Our Dataset's Social Media Sources:** (a) Wordcloud of 3,385 unique hashtags mined iteratively from Twitter handles in our dataset, starting from initial accounts and expanding through network analysis of commonly used hashtags. (b) Wordcloud of Twitter handles from the 2,382 accounts discovered through this iterative hashtag mining process.

- Overview of Our Text Embedding and Clustering Pipeline:** Left: RegEx-based cleaning is performed to separate tweet text from hashtags and URLs. Then GPT-3 embeddings () were computed separately for both cleaned text and hashtags before combination. Right: Resulting multilingual clusters of semantically similar road events via hybrid hierarchical k-means clustering (). Refer back to Appendix B.1

11	Hybrid Approach for QA Generation Combining Visual and Social Context: Left: Input video is segmented into 3-second clips, with Qwen2-VL (🦾) generating captions (📄) for each segment. Middle: Claude 3.5 Sonnet (🦾) synthesizes these captions into a comprehensive video summary (📄). Right: Final prompt combines this summary with cleaned social media text (caption & replies) to generate relevant QA pairs using template questions. The prompt to generate caption for a video-segment (📄) is illustrated in Fig. 12. The full caption output (📄) for a video in our dataset is illustrated in Fig. 16 - 17. The prompt to generate summary of a video from its segment captions (📄) is illustrated in Fig. 13. The full summary output (📄) for the same video is illustrated Fig. 18. Also, the prompt that utilizes video summary, clean tweet text and template questions, to generate QA pairs corresponding to a video (📄), is illustrated in Fig. 14 - 15. 🗨️ represents the initially generated QA pairs which will further be modified and refined as discussed in upcoming subsections.	13
12	Prompt design for generating segment-wise captions using Qwen2-VL Video LLM. The model generates detailed descriptions for each 3-second video segment, capturing temporal visual information. Refer back to Fig. 11 or Appendix B.3.	14
13	Prompt template for generating cohesive video summaries using Claude 3.5 Sonnet. The Text LLM combines segment-wise captions to create a comprehensive temporal description of the entire video. Refer back to Fig. 11 or Appendix B.3.	14
15	Complete QA generation prompt utilizing both video summary and social media context. Template questions guide Claude 3.5 Sonnet to generate relevant question-answer pairs capturing both visual and social context. Refer back to Fig. 11.	16
17	Example of generated captions for each of the three-second segment of the video via the pipeline demonstrated in Fig. 11	23
18	Example of generated video summary from the captions via Text LLM, via the pipeline demonstrated in Fig. 11.	24
19	Demonstrating the importance of hybrid information sources for QA generation. While the video summary (Fig. 18) captures basic visual elements (<i>e.g.</i> ‘nighttime’, ‘streetlight illumination’), the tweet conversation provides crucial contextual information (shown by colored boxes) missing from the visual description alone. This example illustrates why our approach combines both video summaries and social media context to generate diverse and socially-informed QA pairs.	25
20	Prompt design for generating specific QA pairs from generic templates. The expert-driven prompt transforms generic questions into contextually specific ones while maintaining alignment with predefined categories (<i>e.g.</i> , Camera Device, Road Event Type, Actions). The prompt takes generic QA pairs and video summaries as input and generates specific questions that capture detailed aspects of road events while preserving the taxonomic structure. Refer back to Appendix B.4.	25
22	QA refinement prompt: The prompt implements comprehensive guidelines for (1) removing social media references, (2) standardizing temporal information, (3) enforcing speculative language for non-obvious causation, (4) maintaining objective observations, and (5) ensuring human-like sentence-form responses. Refer back to Appendix B.5	28
24	QA categorization prompt design. The prompt (1) matches each refined question against 18 predefined template questions used in QA generation, (2) assigns similarity scores (0-5), and (3) provides examples demonstrating proper template matching for various question types. Refer back to Appendix B.5	30
25	Prompt design for road event classification. The prompt implements binary classification (road/non-road) with confidence scoring (0-1) and reasoning requirements for video content. Refer back to Appendix B.6.	32
26	Video summarization prompt for non-road events. The prompt ensures structured description of visual content focusing on key details, temporal sequences, and object interactions. Refer back to Appendix B.6.	32
27	Incompatible QA generation prompt. The prompt generates explanations for why road event questions are incompatible with non-road video content while maintaining established response formats. Refer back to Appendix B.6.	33

28	Adversarial QA generation prompt. The prompt instructs the generation of questions about non-occurring events while maintaining road context, with examples demonstrating (1) proper introduction of irrelevant elements, (2) explicit negation in answers, and (3) preservation of video-centric response format. Refer back to Appendix B.7.	34
29	QA Task Taxonomy and Template Question ID Mapping for Video LLM Evaluation: The taxonomy consists of 12 QA tasks organized into four reasoning categories: Complex (red), Factual (green), Imaginative (orange), and Hallucination (purple). The 19 template question IDs (blue circles) map to QA tasks designed for evaluating road event understanding. For Incompatible QAs, which evaluate model robustness on non-road event videos, we employ a separate generation pipeline (Appendix B.6) without template ID mapping. Refer back to Appendix B.9.	35
30	Examples of QA Pairs grouped by tasks and color-coded by task category (for an advertisement video captured via multiple viewpoints). Gray fill shading indicates specific questions while the non-shaded QAs are generic. Highlighted text indicates key information. Refer back to Appendix B.9.	36
31	Examples of QA Pairs grouped by tasks and color-coded by task category (for an hydroplaning incident captured via handheld camera). Gray fill shading indicates specific questions while the non-shaded QAs are generic. Highlighted text indicates key information. Refer back to Appendix B.9.	36
32	Examples of QA Pairs grouped by tasks and color-coded by task category (for a near-miss incident captured in Japan). Gray fill shading indicates specific questions while the non-shaded QAs are generic. Highlighted text indicates key information. Refer back to Appendix B.9.	37
34	Tag extraction prompt design for different template question IDs. The prompt employs conditional logic based on question IDs to generate appropriate tags: camera type (<code>q_idx=0</code>), road event type (<code>q_idx=1</code>), country (<code>q_idx=2</code>), and specific location (<code>q_idx=3</code>). Each condition includes specific instructions and examples for tag generation, ensuring standardized output format <code>'tags': [",","]]</code> . Note: <code>ques_idx=1</code> is a command providing tag generation instructions for <code>q_idx=1</code> . This command can be found in Fig. 35. Refer back to Appendix B.10.	42
35	Refer to Fig. 34 for details.	43
36	The diagram shows an example of task-specific prompt utilized for the evaluation of Video LLMs. The code snippet at the top demonstrates how this is done. First, for a specific question, we find its QA type via its template ID, then for that template ID, if we have a predefined constraint, we append that to the original question. Original question + constraint examples for Description QA, Why QA and Temporal Grounding QA tasks is shown. Rest of the tasks have only original questions and no predefined constraints. Refer back to Appendix C.1.	44
37	Evaluation prompt for assessing model-generated answers. The prompt implements (1) structured comparison between predicted and ground-truth answers, (2) fine-grained scoring on a 0-100 scale, and (3) requirement for explanatory justification. The output format ensures programmatic processing while maintaining evaluation transparency. Refer back to Appendix C.2.	45
38	Model performance comparison on Temporal Grounding task: Top: Frames from a video showing a car accident sequence. Middle: Models are asked to specify the temporal interval of the key road event. Ground truth (in gray) indicates the event spans 14-292 seconds. Bottom: Model responses (colored boxes) demonstrate varying approaches: while some attempt to provide specific intervals (e.g., 15-20 seconds, 0-3 seconds), others offer vague temporal descriptions. Red circles around model icons indicate that despite different response styles, all models fail to accurately identify the correct time interval. This example illustrates the significant challenge Video LLMs face in precise temporal localization of road events. Refer back to Appendix C.3.	46

39	Model performance comparison on Temporal Grounding task: Top: Sequential frames from a CCTV video showing a nighttime road scene. Middle: Models are asked to specify the temporal interval of the key road event, with ground truth spanning 35-40 seconds (gray box). Bottom: Model responses (colored boxes) demonstrate varying approaches: most provide specific time intervals (e.g., 23:01-23:11, 05:01-05:06) while Gemini additionally describes the event type ('person hit by motorcycle'). Red circles around model icons indicate that despite different response styles, all models fail to provide the correct interval. GPT-4o's response (57 to 42 seconds) even shows incorrect temporal ordering. This example highlights Video LLMs' consistent difficulty with precise temporal localization. Refer back to Appendix C.3.	47
40	Model performance comparison on Temporal Grounding task: Top: Dashcam footage showing a night-time near-miss incident between a taxi and cyclist. Middle: Models are asked to specify the temporal interval of the key road event, with ground truth spanning 5-11 seconds (gray box). Bottom: Model responses (colored boxes) show diverse approaches: while Gemini-1.5 Pro (green circle) correctly identifies both the event type and provides a reasonable time estimate (0:05-0:10), other models either give incorrect intervals (IXC: 22-25s, GPT-4o: 52-59s), overly precise timing (Dolphin: 0.0-0.001s), or incomplete responses (Tarsier: '11 seconds'). This example demonstrates that even when models accurately describe the event (taxi passing close to cyclist), precise temporal localization remains challenging, with only one model achieving high accuracy. Refer back to Appendix C.3.	48
41	Model performance comparison on Incompatible QA task: Top: Video of a Reddit interface (non-road-event content). Middle: Models are asked about attentive driving implications, while the ground truth (gray box) correctly mentions that the content is unrelated to driving safety. Bottom: Model responses showcase varying levels of hallucination: most models (red circles) fabricate driving scenarios and safety implications despite the irrelevant content, while Tarsier (green) correctly identifies that the video is not related to road event. Although, GPT-4o (orange) correctly identifies the computer interface, it still attempts to relate it to driving. This example highlights a critical challenge in Video LLM robustness - the tendency to hallucinate road safety contexts even when presented with completely unrelated visual content. Refer back to Appendix C.3.	49
42	Model performance comparison on Incompatible QA task: Top: Frames showing indoor welding activity in a workshop. Middle: Models are asked about roadside maintenance accident prevention, while the ground truth (gray box) correctly indicates that the content shows indoor welding, not roadside maintenance. Bottom: Model responses (colored boxes) demonstrate varying degrees of hallucination: while Tarsier (green) correctly acknowledges insufficient information to discuss roadside maintenance, Dolphin and IXC (red circles) fabricate elaborate safety measures despite the obvious indoor setting. Gemini and GPT-4o's (dark orange) detailed response about welding safety, while technically accurate, still fails to address the fundamental context mismatch. This example illustrates how models can generate plausible but irrelevant safety recommendations when presented with visually similar but contextually different scenarios. Refer back to Appendix C.3.	50
43	Model performance comparison on Incompatible QA task: Top: Frames showing a person building and subsequently demolishing a brick wall. Middle: Models are asked about psychological factors behind overtaking behavior, while the ground truth (gray box) correctly mentions this as unrelated to overtaking. Bottom: Model responses show varying levels of hallucination and context confusion: Dolphin and IXC (red circles) completely ignore the brick wall context and fabricate scenarios about road safety, while Tarsier and Gemini (green) correctly acknowledges the construction setting and clearly states the content mismatch. This example demonstrates how models can struggle with maintaining contextual accuracy, with some generating elaborate but irrelevant psychological analyses despite clearly unrelated visual content. Refer back to Appendix C.3.	51
44	Model performance comparison on Why QA task: Top: CCTV footage showing an intersection incident between a truck and motorcyclist. Middle: Models are asked about potential reasons behind the road entities' actions, with ground truth (gray box) indicating rush and lack of caution as primary factors. Bottom: Model responses demonstrate varying levels of reasoning and detail: While Dolphin (red circle) provides an over-simplified response ('Because traffic moving normally'), other models offer increasingly complex analyses. Gemini generates a comprehensive analysis considering multiple factors (weather conditions, road visibility, driver attention), while GPT-4o provides a structured but possibly over-analyzed response with enumerated factors. This example illustrates the challenge of providing appropriate depth in causal reasoning without over-speculation. Refer back to Appendix C.3.	52

45	Model performance comparison on Why QA task: Top: Dashcam footage showing a traffic scenario with lane changing incidents. Middle: Models are asked about the primary reason for the incident, with ground truth (gray box) identifying improper lane changing and insufficient vehicle distance as key factors. Bottom: Model responses show varying levels of analytical accuracy and specificity: Dolphin offers an oversimplified and irrelevant response ('Because the light is green'), while Tarsier-34B provides a vague description without specific reasoning. IXC-2.5 attempts causal analysis but misidentifies the vehicles involved, and Gemini-1.5 Pro introduces unobserved elements (pedestrian crossing). GPT-4o demonstrates appropriate caution by acknowledging the difficulty in determining exact causes without clearer incident details. This example highlights the challenges in balancing between definitive causal analysis and appropriate uncertainty when visual evidence is ambiguous. Refer back to Appendix C.3.	53
46	Model performance comparison on Why QA task: Top: Split-screen dashcam footage showing both driver behavior (phone use) and road view leading to an incident. Middle: Models are asked about potential reasons behind the driver's actions, with ground truth (gray box) identifying distraction from phone use and possible panic/distress. Bottom: Model responses demonstrate varying depths of causal analysis: Dolphin provides an oversimplified response about road view, while Gemini-1.5 Pro offers a comprehensive multi-factor analysis incorporating both observed behaviors (phone distraction) and possible underlying causes. IXC-2.5 stays focused on direct observables, while GPT-4V extensively analyzes multiple scenarios but maintains grounding in the visible evidence (phone conversation). This example shows how models balance between observable evidence (phone use) and inferring potential psychological states, with varying success in maintaining relevance to the visual content. Refer back to Appendix C.3.	54
47	Model performance comparison on Description QA task: Top: Video frames showing an encounter between an elephant and a van on a rural road. Middle: Models are asked to describe the type of road event, with ground truth (gray box) identifying it as both unsafe driving behavior and an animal-related wildlife encounter. Bottom: Model responses show varying accuracy in event categorization and detail: Tarsier-3LB incorrectly describes a collision, while IXC-2.5 (green circle) provides a well-balanced response that correctly categorizes the event as 'animal-related incident' while acknowledging the safety implications for all parties. Gemini-1.5 Pro and GPT-4o offer accurate but differently focused descriptions, with Gemini emphasizing the physical interaction and GPT-4o highlighting the broader safety context. This example demonstrates models' varying abilities to balance between event classification, factual description, and safety implications in unusual road scenarios. Refer back to Appendix C.3.	55
48	Model performance comparison on Description QA task: Top: Video frames show a collision between a car and a bike on a curvy road. Middle: Models are asked to describe the actions of the entities involved in the road event, with ground truth (gray box) identifying it near-miss incident that further led to the collision. Bottom: Model responses show varying accuracy in event categorization and detail: All the models fail to answer this question due to incorrect identifications. GPT-4o fails to identify the motorcycle that was initially overtaking the auto that crashed a car. Tarsier-34B incorrectly identifies overtaking between the car and the truck. Refer back to Appendix C.3.	56
49	Model performance comparison on Description QA task: Top: Video frames show a traffic violation involving. Middle: Models are asked to describe the type of road event depicted in the video, with ground truth (gray box) identifying it as a vehicle driving the wrong way on a one-way street. Bottom: All models fail to recognize the violation. Refer back to Appendix C.3.	57
50	Model performance comparison on Description QA task: Top: Video frames showing a road safety awareness video aimed towards pedestrians or motorcyclists at night. Middle: Models are asked to describe the theme of the video, with ground truth (gray box) indicating that it is a safety awareness video. Bottom: Model responses show varying accuracy in event categorization and detail: All the models except Dolphin successfully capture the global context or theme of the video. Refer back to Appendix C.3.	58

51	The image shows a qualitative analysis of the performance of GPT-4o Video LLM for two types of questions - a generic question about the actions of the road entities, and a specific question about how the lack of proper infrastructure affects different road users. The ground truth (GT) answers are provided, and the predicted answers by the model are shown using icons - a red circle indicates the model's prediction does not align well with the ground truth, while a green icon indicates the model performs well. GPT-4o seems to be performing well in specific questions than generic one. This performance gap could be because generic questions require the model to infer the context while specific questions directly reference the event and entities, making it easier for models to answer them.	59
52	A similar phenomena between the gap between generic and specific QAs is reflected in Gemini, as seen in the previous example.	60
53	A similar phenomena indicating the gap between generic and specific QAs is reflected in Tarsier, similar to what was observed in the Gemini and GPT-4o in previous examples.	61
54	The phenomena of the model performing better in specific QAs than their generic counterparts persist in IXC as well.	62
55	Geographical location (country of origin) distribution of video tags. Tags with fewer than five videos are omitted from the radar plot for clarity and to reduce clutter.	63
56	Road Event Video Tags distribution. Tags with fewer than five videos are omitted from the radar plot for clarity and to reduce clutter.	63
57	Traffic Violation Video Tags distribution. Tags with fewer than four videos are omitted from the radar plot for clarity and to reduce clutter.	63
58	Road Entity Action Video Tags distribution. Tags with fewer than 57 videos are omitted from the radar plot for clarity and to reduce clutter.	63
59	Viewpoint Video Tags distribution.	64
60	Time of Day Tags distribution.	64
61	Road Type Video Tags distribution. Tags with fewer than 15 videos are omitted from the radar plot for clarity and to reduce clutter.	64
62	Weather Condition Video Tags distribution. Tags with fewer than five videos are omitted from the radar plot for clarity and to reduce clutter.	64

A. Data Collection

To identify relevant handles, we first created a multilingual keyword dictionary covering traffic terminology, emergency services, and regional variations (examples in Fig. 8). Using this dictionary, we manually identified key handle and analyzed their commonly used hashtags. Through hashtag mining and network analysis of these accounts, we discovered related accounts. This approach resulted in a total 2,382 accounts. We then scraped their content (videos, captions, and replies) from 2012 onwards. We programmatically filtered out tweets with fewer than four replies, retaining only those with substantial discussion. Representative hashtags and the handles are shown in Fig. 9. This systematic approach ensured the collection of road event content with significant community interaction. Full list of keywords, hashtags and handles will be released with dataset.

B. Annotation Strategy: QAs and Tags

B.1. Identifying Representative Road Event Samples

Our annotation strategy begins with identifying representative samples that capture the diversity of road events in our multilingual dataset. As shown in Fig. 10, we first implement a text preprocessing pipeline where tweets undergo cleaning to remove URLs while preserving essential content. For instance, a tweet `Cyclist nearly hit by car #OxfordStreet @MetPolice https://t.co/xyz` is reduced to `Cyclist nearly hit by car @MetPolice`. Concurrently, we extract and process hashtags separately, maintaining their semantic value by removing only the # symbol (e.g. `#RoadSafety #CyclingUK #NearMiss` becomes `RoadSafety CyclingUK NearMiss`). For tweets lacking hashtags, we introduce a placeholder `#NoHashtag`. Using OpenAI’s GPT-3 text embeddings API [18], we generate separate embeddings for cleaned text and processed hashtags. Our empirical analysis suggested that separately computing embeddings for cleaned text and hashtags, followed by their combination through averaging, yielded more representative sample clusters compared to alternatives such as embedding raw text or cleaned text alone.

These combined embeddings then undergo a hierarchical k-means clustering with a divisive approach (Fig. 10). The process begins with a single cluster and iteratively creates sub-clusters based on silhouette scores. Specifically, after each k-means step, if the score improves or remained stable, we proceed to divide sub-clusters further; if it decreases significantly (indicating poor separation), we halt further splits for that branch of the hierarchy. This recursive process continues until reaching either a minimum cluster size or a predefined depth, with empirical analysis suggest-

ing optimal results at 95 clusters. This approach effectively groups similar road events across languages. For example, one cluster combines near-miss incidents like `Bike’s near-miss with bus (Thailand)` and `Close call with cyclist on Main Street (Australia)`, while another groups illegal overtaking events such as `Car illegal overtaking from China` and `Dangerous overtaking by bus on a bike lane from Australia`. Weather-related incidents form distinct clusters including `Car hydroplaning in heavy rain on I-95 (USA)` and `Vehicle sliding on icy road conditions (Canada)`. From each cluster, we select five representative samples using a center-based approach. By computing the Euclidean distance between each sample and its cluster center, we identify the samples that best represent the cluster’s core characteristics while maintaining linguistic and regional diversity. This systematic approach, validated through manual review, ensures our QA generation is grounded in well-represented events across our dataset.

B.2. Template Question Generation

To develop comprehensive template questions for our dataset, we implemented an iterative approach based on analysis of representative video samples and their associated social media discourse. Following our hierarchical clustering process (Fig. 10), we selected 5 representative videos from each of the 95 distinct clusters, creating a diverse corpus of 475 videos for detailed examination.

Formulating fundamental questions: In the initial phase, we conducted manual analysis of the selected videos and their associated tweet conversations, focusing on fundamental aspects of road events. We began by formulating basic questions such as `What road event took place in the video?`

Formulating analysis questions: We expanded our template set based on patterns observed in social media discussions. For example, in videos involving accidents and near-misses, conversations were frequently centered on causal analysis. This observation led us to develop questions specifically probing the potential causes and motivations behind road events, such as `What was the primary reason behind the occurrence of this incident?` Similarly, discussion around post-crash measures in relevant scenarios, led the inclusion of template questions addressing response actions such as `What measures should be taken after witnessing an event like this?`

Template refinement: The template refinement process was inherently iterative, with each round of video analysis contributing to the evolution of our question set. A key consideration was maintaining question generalizability while preserving specificity where necessary. For instance,

```

"Traffic_General":
{
    "English": ["traffic", "road", "highway", "street", "accident", "incident"],
    "Spanish": ["tráfico", "carretera", "autopista", "calle", "accidente", "incidente"],
    "French": ["circulation", "route", "autoroute", "rue", "accident", "incident"],
    "German": ["verkehr", "straße", "autobahn", "unfall", "vorfall"],
    "Japanese": ["交通", "道路", "高速道路", "事故", "通行"],
    "Chinese": ["交通", "公路", "高速", "事故", "道路"],
    "Hindi": ["यातायात", "सड़क", "राजमार्ग", "दुर्घटना", "हादसा"],
    "Korean": ["교통", "도로", "고속도로", "사고", "통행"],
    "Russian": ["движение", "дорога", "автострада", "авария", "происшествие"],
    "Arabic": ["طريق سريعة", "شارع", "حادث", "طريق", "مرور"]
    ...
},

"Emergency_Services":
{
    "English": ["police", "highway patrol", "traffic police", "emergency"],
    "Spanish": ["policía", "guardia civil", "policía de tráfico", "emergencia"],
    "French": ["police", "gendarmerie", "police routière", "urgence"],
    "German": ["polizei", "verkehrspolizei", "notfall", "autobahnpolizei"],
    "Japanese": ["警察", "道路警察", "緊急", "パトロール"],
    "Chinese": ["警察", "交警", "紧急", "巡逻"],
    "Hindi": ["पुलिस", "यातायात पुलिस", "आपातकालीन", "गश्ती"],
    "Korean": ["경찰", "도로경찰", "긴급", "순찰"],
    "Russian": ["полиция", "дорожная полиция", "патруль", "чрезвычайный"],
    "Arabic": ["شرطة المرور", "طوارئ", "دورية", "شرطة"]
    ...
},

"Incidents":
{
    "English": ["crash", "collision", "roadblock", "traffic jam", "construction"],
    "Spanish": ["choque", "colisión", "bloqueo", "atasco", "construcción"],
    "French": ["collision", "embouteillage", "blocage", "construction"],
    "German": ["zusammenstoß", "kollision", "stau", "baustelle"],
    "Japanese": ["衝突", "渋滞", "封鎖", "工事"],
    "Chinese": ["碰撞", "堵塞", "封锁", "施工"],
    "Hindi": ["टक्कर", "भीड़", "जाम", "निर्माण"],
    "Korean": ["충돌", "교통체증", "봉쇄", "공사"],
    "Russian": ["столкновение", "пробка", "блокировка", "строительство"],
    "Arabic": ["بناء", "حظر", "ازدحام", "تصادم"]
    ...
},

"Recording_Devices":
{
    "English": ["dashcam", "CCTV", "traffic camera", "surveillance"],
    "Spanish": ["cámara de coche", "CCTV", "cámara de tráfico", "vigilancia"],
    "French": ["caméra embarquée", "vidéosurveillance", "caméra routière"],
    "German": ["dashcam", "überwachungskamera", "verkehrskamera"],
    "Japanese": ["ドライブレコーダー", "監視カメラ", "交通カメラ"],
    "Chinese": ["行车记录仪", "监控", "交通摄像头"],
    "Hindi": ["डैशकैम", "सीसीटीवी", "यातायात कैमरा"],
    "Korean": ["블랙박스", "CCTV", "교통카메라"],
    "Russian": ["видеорегистратор", "камера наблюдения", "дорожная камера"],
    "Arabic": ["كاميرا المرور", "كاميرا مراقبة", "كاميرا السيارة"]
    ...
},

```



```

"Hashtag_Templates":
{
  "English": ["#TrafficAlert", "#RoadIncident", "#TrafficUpdate"],
  "Spanish": ["#AlertaTráfico", "#IncidenteVial", "#ActualizaciónTráfico"],
  "French": ["#AlerteCirculation", "#IncidentRoute", "#InfoTrafic"],
  "German": ["#Verkehrsmeldung", "#verkehrsInfo", "#StauAlert"],
  "Japanese": ["#交通情報", "#事故情報", "#渋滞情報"],
  "Chinese": ["#交通提醒", "#事故通知", "#路况"],
  "Hindi": ["#यातायातसूचना", "#सड़कदुर्घटना", "#ट्रैफिकअपडेट"],
  "Korean": ["#교통알림", "#사고정보", "#교통정보"],
  "Russian": ["#дорожнаяситуация", "#дтп", "#пробкиСейчас"],
  "Arabic": ["#حالة_الممر", "#حادث_طريق", "#تنبيه_مروري"]
  ...
},

"Location_Specific": {
  "template": {
    "English": "[City]Traffic, [City]Roads, [City]Alert",
    "Spanish": "[Ciudad]Tráfico, [Ciudad]Vial",
    "French": "[Ville]Circulation, [Ville]Route",
    "German": "[Stadt]Verkehr, [Stadt]Straßen",
    "Japanese": "[都市]交通, [都市]道路",
    "Chinese": "[城市]交通, [城市]道路",
    "Hindi": "[शहर]यातायात, [शहर]सड़क",
    "Korean": "[도시]교통, [도시]도로",
    "Russian": "[Город]движение, [Город]дороги",
    "Arabic": "[مدينة]طرق [مدينة]مرور"
    ...
  }
}

Search_combinations =
{
  "basic": "[Language_Hashtag] + [City_Name] + [Incident_Type]",
  "advanced": "[Emergency_Service] + [Recording_Device] + [Location]",
  "monitoring": "[City_Name] + [Traffic_General] + [Emergency_Service]"
  ...
}

```

Figure 8. **Multilingual Traffic Keyword Dictionary for Tweet Mining:** A comprehensive dictionary of traffic-related keywords and hashtags, designed for identifying road event content on social media. Terms span traffic incidents, emergency services, recording devices, and location-specific templates. Effective usage involves combining terms across categories *[Traffic_General]* + *[Incidents]* and creating location-specific searches.

certain questions (*e.g.* those about accident causation) were not universally applicable across our diverse video corpus. This recognition prompted us to reformulate the questions to ensure broader applicability. For instance, accident-related questions were reframed conditionally: If the road event involves an accident or a near-miss incident, explain how it could have been prevented. We also incorporated universally applicable questions about recording devices (*e.g.* What type of camera was used to capture the video?) and geographical context (*e.g.* In which country did this road event

take place?).

Spatial and temporal aspects in questions: Furthermore, we carefully structured the questions to address both spatial and temporal aspects of road events. Spatial questions could be answered through single-frame analysis (*e.g.* In which country did this road event take place? or What were the weather conditions or road visibility when the video was captured?). While temporal questions inquire about event sequences and interactions (*e.g.* Describe the actions performed by the road entities involved in the road



Figure 9. **Visualization of Our Dataset’s Social Media Sources:** (a) Wordcloud of 3,385 unique hashtags mined iteratively from Twitter handles in our dataset, starting from initial accounts and expanding through network analysis of commonly used hashtags. (b) Wordcloud of Twitter handles from the 2,382 accounts discovered through this iterative hashtag mining process.

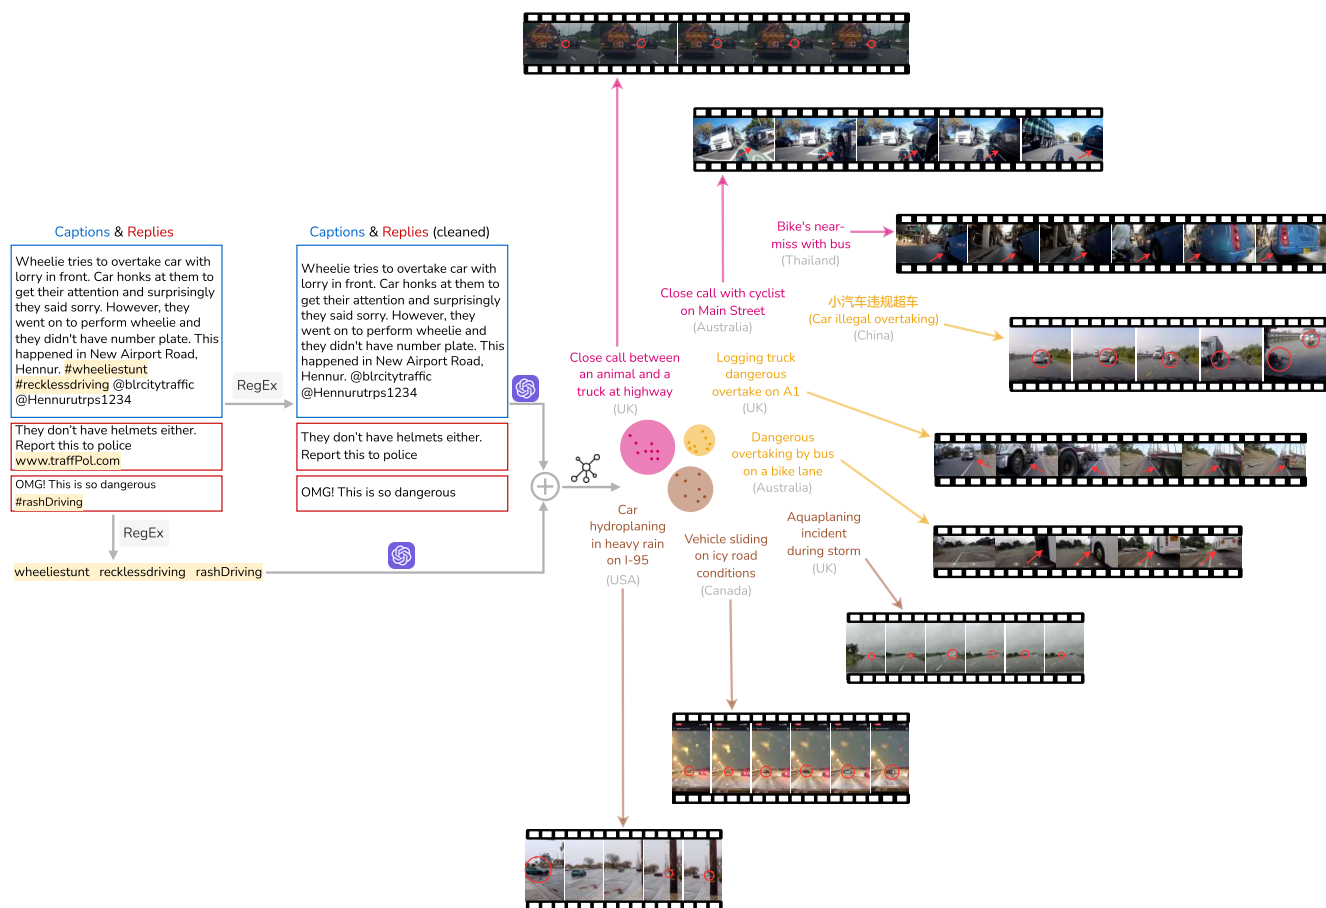


Figure 10. **Overview of Our Text Embedding and Clustering Pipeline:** Left: RegEx-based cleaning is performed to separate tweet text from hashtags and URLs. Then GPT-3 embeddings (🧠) were computed separately for both cleaned text and hashtags before combination. Right: Resulting multilingual clusters of semantically similar road events via hybrid hierarchical k-means clustering (🌀). Refer back to Appendix B.1

event or Specify the approximate time interval where the key road event is observed in the video?). This dual approach en-

ures comprehensive coverage of both spatial and temporal dimensions of road events.

The final set of 18 carefully curated questions, shown

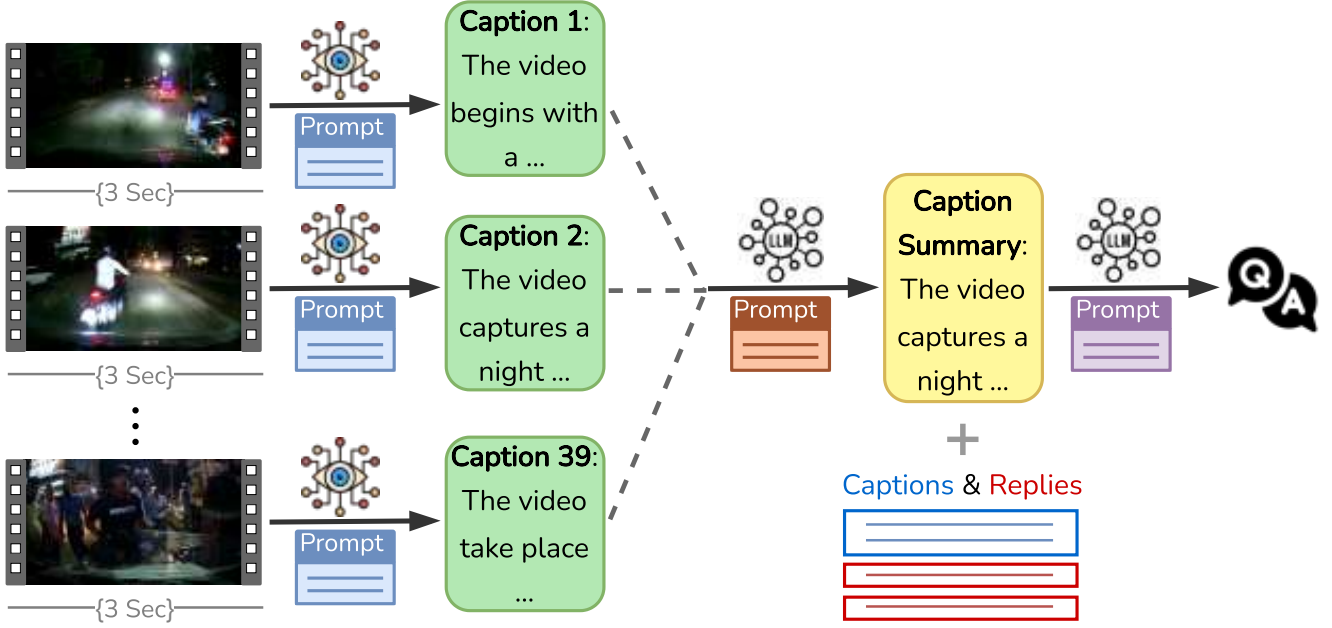


Figure 11. **Hybrid Approach for QA Generation Combining Visual and Social Context:** Left: Input video is segmented into 3-second clips, with Qwen2-VL (🔍) generating captions (🟢) for each segment. Middle: Claude 3.5 Sonnet (🔍) synthesizes these captions into a comprehensive video summary (🟡). Right: Final prompt combines this summary with cleaned social media text (caption & replies) to generate relevant QA pairs using template questions. The prompt to generate caption for a video-segment (🔍) is illustrated in Fig. 12. The full caption output (🟢) for a video in our dataset is illustrated in Fig. 16 - 17. The prompt to generate summary of a video from its segment captions (🔍) is illustrated in Fig. 13. The full summary output (🟡) for the same video is illustrated in Fig. 18. Also, the prompt that utilizes video summary, clean tweet text and template questions, to generate QA pairs corresponding to a video (🔍), is illustrated in Fig. 14 - 15. 🗣️ represents the initially generated QA pairs which will further be modified and refined as discussed in upcoming subsections.

in Fig. 14 - 15, were integrated into our LLM prompting strategy which is described in the next subsection.

B.3. QA Generation via Hybrid Approach

To generate question-answer pairs for each video, we developed a hybrid approach that leverages both visual content and social media context. Our pipeline, illustrated in Fig. 11, consists of three main stages that systematically combine video understanding with social context.

First, we extract visual semantics by splitting each video into 3-second segments and employ Qwen2-VL Video LLM [33] to generate detailed captions for each segment. The prompt (🔍) to generate caption (🟢) for a video-segment is illustrated in Fig. 12. This temporal segmentation ensures capture of fine-grained details and event progression. Next, these segment-wise captions are processed by Claude 3.5 Sonnet Text LLM [1] to generate a cohesive, visually-rich summary of the entire video. The prompt (🔍) to generate summary of a video from its segment captions is illustrated in Fig. 13. Finally, we combine this generated summary (🟡) with cleaned tweet text (captions & replies) to create contextually rich QA pairs using template questions through Claude 3.5 Sonnet. The prompt (🔍) that utilizes

video summary, clean tweet text and template questions, to generate QA pairs corresponding to a video, is illustrated in Fig. 14 - 15. References for inputs and outputs at each stage of this pipeline are provided in the Fig. 11. Fig. 19 demonstrates the utility of social conversation in QA formation.

B.4. Specific QA Generation

To create a comprehensive question set with varying difficulty levels, we developed an approach for generating specific questions from generic template set (Fig. 14 - 15). While generic questions like What actions were performed by the road entities involved in the key road event? require complex temporal reasoning and synthesis of multiple observations, specific questions such as How was the truck involved in the accident? focus on particular entities and events, offering more straightforward path for answer formulation.

We developed a specialized prompt that instructs the LLM (Claude 3.5 Sonnet) to act as an expert with comprehensive knowledge of driving norms across different geographical regions. The prompt takes two inputs: the generic QA pairs generated from our initial template-based

Video Segment Caption Generation Prompt

Generate a detailed and accurate description of a video.
Use the following details to create a clear and complete narrative:

Instructions for writing the detailed description:

1. Focus on describing key visual details such as appearance, motion, sequence of actions, objects involved, and interactions between elements in the video.
2. Emphasize important points like the order of events, appearance and actions of people or objects, and any significant changes or movements.
3. Give a thorough description, highlighting the key visual and temporal details while keeping it clear and easy to understand.

Figure 12. Prompt design for generating segment-wise captions using Qwen2-VL Video LLM. The model generates detailed descriptions for each 3-second video segment, capturing temporal visual information. Refer back to Fig. 11 or Appendix B.3.

Segment Caption Summarization Prompt

We split a video into segments and extracted detailed captions for each segment. The captions for all segments can be found as follows, in the order of their occurrence. For example, 'Caption 1' corresponds to the caption generated for the first video segment.

Generate a detailed and accurate description of the entire video as a paragraph, based on all the given video captions. Make sure not to lose any important information. `{input_captions}`

Use the following details to create a clear and complete narrative:

Instructions for writing the detailed description:

1. Focus on describing key visual details such as appearance, motion, sequence of actions, objects involved, and interactions between elements in the video.
2. Check for consistency between captions, and prioritize details that match the captions. Ignore any conflicting or irrelevant details from the captions.
3. Combine and organize information from all captions into one clear and detailed description, removing any repeated or conflicting details.
4. Emphasize important points like the order of events, appearance and actions of people or objects, and any significant changes or movements.
5. Do not mention that the information comes from captions.
6. Give a thorough description, highlighting the key visual and temporal details while keeping it clear and easy to understand. Use your intelligence to combine and refine the captions into an informative description of the entire video.
7. Also, use your common sense to conclude what is going on in the video.

Figure 13. Prompt template for generating cohesive video summaries using Claude 3.5 Sonnet. The Text LLM combines segment-wise captions to create a comprehensive temporal description of the entire video. Refer back to Fig. 11 or Appendix B.3.

approach (Fig. 14 - 15) and the corresponding video summary (e.g. Fig. 18) to generate contextually appropriate specific questions. The prompt (📄) to generate specific QA pairs is illustrated in Fig. 20.

B.5. QA Refinement and Categorization

QA Refinement: To ensure our QA pairs are strictly video-centric and maintain high quality, we developed a comprehensive refinement process that addresses the challenges inherent in social media discourse. Social media discussions often contain non-visual information such as personal iden-

tifiers, historical references, and specific temporal details that cannot be directly verified through video content alone. To address this challenge, we designed a refinement prompt (Fig. 21 - 22) for Claude 3.5 Sonnet Text LLM.

The refinement prompt takes the generic and specific QA pairs generated in previous step as input and applies multiple filtering criteria to ensure video-centricity. The process eliminates references to social media context (e.g. 'based on replies', 'as mentioned in comments') while preserving essential information about entities and events. It standardizes temporal references, converting specific dates and times-

QA Generation Prompt

You are an expert in understanding Road and Traffic Events with extensive knowledge of safe driving norms across various geographical regions.

You are provided with a textual conversation related to a video posted on Twitter as well as a detailed summary of that video. The textual conversation includes a caption and multiple replies in some cases. However, you do not have access to the actual video.

Task:

- Describe the key road and traffic events discussed in the textual conversation while also referring to the detailed video summary. (A key road event is the main focus of the video that is being discussed in the textual conversation)
- Generate relevant Question-Answer (QA) pairs by analyzing key aspects discussed in the textual conversation while also referring to the detailed video summary.
- In addition to the provided template questions, feel free to generate additional QA pairs that are contextually appropriate.

Below is a set of **template questions** for forming QA pairs:

<Question-1> What type of camera was used to capture the video? </Question-1> (Type-of-Camera e.g., dashcam, vehicle-mounted camera, hand-held camera, cell-phone camera, CCTV camera, surveillance camera, drone camera, multiple-cameras i.e., not a fixed view point, etc. Do not specify the name of the camera model, just specify its type.)

<Question-2> Describe the type of key road event captured in the video. </Question-2> (Type-of-Road-Event e.g., safe/unsafe road infrastructure or driving behavior, dangerous, rash, or aggressive driving, road rage, traffic violation, accident/crash, post-crash, near-miss, awareness of road safety, defensive driving, etc.)

<Question-3> In which country did this road event take place? </Question-3> (Country-of-Origin e.g., India, UK, US, Japan, China, etc. Do not justify how you got the answer)

<Question-4> In which state, district, city/town/village, or locality did the road event occur? </Question-4> (Location could be the name of a state, district, landmark, type of locality, like city/town/village, etc. Specific-Location e.g., Hyderabad city, Big Ben London, etc.)

<Question-5> On which type of road or area, this event have taken place? </Question-5> (Type-of-Road e.g., urban area, rural area, highway, flyover, turn, intersection, tunnel, bridge, T-junction, roundabout, hilly or mountain area, etc. Do not justify how you got the answer. Do not specify the name or address of the region where the event took place, just specify its type.)

<Question-6> When did this road event happen? </Question-6> (Time-of-Day e.g., morning, afternoon, evening, night, etc. Do not specify the exact date or time in the generated answer.)

<Question-7> What were the weather conditions or road visibility when the video was captured? </Question-7> (e.g., sunny, rainy, windy, foggy, low visibility, etc.)

<Question-8> List down all the road entities involved in the key road event. </Question-8> (A road entity can include road infrastructure objects like traffic signs, lane markings, barricades, etc. Road entities can also include road users like cars, bikes, pedestrians, drivers, etc.)

<Question-9> Describe the visual characteristics of the listed road entities above </Question-9> (e.g., what was the vehicle's color?, was the headlight, brake light, or turn signal on?, what was the license plate number?, was the rider wearing helmet or seat belt?, etc.)

<Question-10> Describe the actions performed by the listed road entities above. </Question-10> (e.g., illegal overtaking, speeding, swerving, yielding, cutting, etc.)

<Question-11> Describe any suspected reason or motive behind the actions of the involved road entities. </Question-11> (e.g., thrill, road rage, impressing others, in a rush, aggressive, impatient, etc.)

<Question-12> If the road event involves an accident or a near-miss incident, what was the primary reason behind its occurrence? </Question-12> (e.g., road rage, etc.)

<Question-13> If the road event involves an accident or a near-miss incident, Explain how it could have been prevented. </Question-13> (e.g., by slowing down at the intersection, checking the rearview mirror, etc.)

<Question-14> If the road event involves an accident, list down any casualties or road infrastructure damage during the event. </Question-14> (e.g., people in the car died, bikers got injured, pedestrians got hit by car, divider was damaged, etc. Do not specify the exact number of casualties (e.g., 5 pedestrians or 3 people) in the generated answer.)

<Question-15> List down all traffic rule violations associated with this road event </Question-15> (e.g., illegal overtaking, illegal overtaking by crossing solid lane markings, hiding license plates, license plate not visible, helmet rule violation, no helmet, wrong-side driving, triple riding violation, red light violation, drunk driving, etc.)

<Question-16> What measures should be taken upon witnessing an unsafe driving situation during this road event? </Question-16> (e.g., reporting any traffic violation, or unsafe road infrastructure to local government authorities or police, fines, jail time, license ban, vehicle confiscation, etc. List only the most relevant measures.)

<Question-17> List down all the road safety advisories corresponding to the listed road entities. </Question-17>

<Question-18> List down all the Counterfactuals related to different road events or driving situations that could have happened under different circumstances. </Question-18> (e.g., the biker would have met an accident if the truck steered a little towards the right, the incident could have been worse if there were pedestrians by the roadside, If the car had not been speeding, it would have safely stopped before the intersection and avoided being hit by the truck, etc.)

Guidelines for Response:

- DO NOT give any reference of the video summary and the textual conversation when answering the questions. Also, avoid using phrases like 'based on the replies', 'based on the comments', 'based on the conversation', 'based on the text', 'mention', 'conversation', 'caption', 'replies', 'comment', 'post', 'twitter', 'user', <name of the person>, etc., that may indicate that the generated answer is based on the textual conversation. Instead, refer to the information as being inferred from the video.

- Do not justify how you got the answer.
- You can choose to ignore if some information seems irrelevant in the summary or the textual conversation.
- All answers should be concise, and important keywords should be highlighted using ***bold*** formatting.
- Do not generate a QA pair if the question is not applicable to the road event.

Output Format:

Your entire response must be formatted in JSON as shown below:

```
{
  "Summary": "",
  "QA_pairs": [
    {"Q": "", "A": ""},
    {"Q": "", "A": ""},
    {"Q": "", "A": ""},
    {"Q": "", "A": ""}
  ]
}
```

Figure 15. Complete QA generation prompt utilizing both video summary and social media context. Template questions guide Claude 3.5 Sonnet to generate relevant question-answer pairs capturing both visual and social context. Refer back to Fig. 11.

tamps to general indicators (e.g. 'morning', 'night'). For non-obvious causation, it enforces the use of speculative language (e.g. 'potential', 'likely') while maintaining factual observations for directly visible events. Additionally,

the process mandates human-like sentence-form responses and removes precise measurements such as exact speeds or weather metrics that cannot be reliably inferred from video content.



Caption 1: The video begins with a view from inside a vehicle at night, showing a road illuminated by streetlights. Various vehicles, including trucks and motorcycles, are visible on the road ahead, moving in the same direction as the viewer's vehicle. The scene is dark, with the headlights of other vehicles creating bright spots on the road surface. As the video progresses, the vehicles continue to move forward, maintaining their positions relative to each other. The road appears to be a multi-lane highway, with vehicles occupying different lanes. The background shows some roadside structures and trees, adding context to the location. The video maintains a consistent perspective from the driver's seat, providing a first-person view of the journey.

The video continues to show the same nighttime scene from inside a vehicle. The road remains illuminated by streetlights, and various vehicles, including trucks and motorcycles, are visible on the road ahead, moving in the same direction as the viewer's vehicle. The scene remains dark, with the headlights of other vehicles creating bright spots on the road surface.

Caption 2: The video captures a nighttime driving scene from the perspective of a vehicle's dashboard camera. The view is through the windshield, showing the road ahead illuminated by various lights from other vehicles and streetlights. The road is busy with multiple lanes of traffic, including cars, motorcycles, and trucks. The vehicles are moving at different speeds, with some appearing to be stationary or moving slowly due to traffic conditions. The scene is dynamic, with the motion of the vehicles creating a sense of movement and activity. The background features buildings and streetlights, adding to the urban nighttime atmosphere. The video maintains this consistent view of the busy road, providing a continuous snapshot of the driving experience during the night.

Caption 3: The video captures a nighttime scene from the perspective of a vehicle's dashboard camera. The view is directed forward, showing a busy street illuminated by streetlights and vehicle headlights. A prominent three-wheeled vehicle with bright blue and red lights on top is seen ahead, driving in the same direction as the viewer's vehicle. To the right of this three-wheeled vehicle, a motorcyclist wearing a helmet and dark clothing rides alongside. The background features various buildings and signs, adding to the urban atmosphere. The scene remains consistent, with minimal changes in the positions of the vehicles and the surrounding environment, emphasizing the steady movement and typical night-time traffic scenario.

Caption 4: The video begins with a view from inside a vehicle at night, focusing on the road ahead. A colorful auto-rickshaw with blue and purple lights is seen driving ahead on the right side of the road. The background features streetlights and buildings, creating a typical urban night scene. The vehicle follows the auto-rickshaw as it moves forward. The scene continues with the same view from inside the vehicle, maintaining the focus on the road ahead. The colorful auto-rickshaw remains visible, now slightly ahead and to the right of the vehicle's position. The streetlights and buildings continue to line the road. The vehicle follows the auto-rickshaw as it moves forward.

The video progresses with the same nighttime setting, showing the road ahead illuminated by streetlights. The colorful auto-rickshaw is no longer visible, but another vehicle is seen driving ahead on the right side of the road. The vehicle follows this new vehicle as it moves forward, maintaining the consistent urban night scene with streetlights and buildings lining the road.

Caption 5: The video begins with a view from inside a car at night, showing a road illuminated by streetlights. A white car is visible in the distance, moving away from the viewer's perspective. As the video progresses, the white car continues to move further down the road, eventually turning right onto another street.

The narrative develops as the white car moves further down the road, now approaching an intersection where it turns left. The surrounding environment remains dark, with streetlights casting a dim glow on the road. As the white car continues to move, it passes through the intersection and continues straight ahead, eventually moving out of the frame to the right. The video concludes with the road empty, maintaining the same nighttime setting.

Caption 6: The video begins with a view from inside a car at night, focusing on the illuminated dashboard. The car is driving on a dark road with streetlights casting a bright glow on the asphalt. The surroundings are dimly lit, with occasional red lights visible in the distance. As the car moves forward, the road ahead appears to be clear with no other vehicles in sight. The scene continues with the same view, maintaining the focus on the illuminated dashboard and the dark road. The road remains clear, and the surrounding environment is still dimly lit with streetlights providing the main source of light. The video progresses with the car continuing its journey down the dark road.

Caption 7: The video begins with a view from inside a car at night, driving on a road illuminated by streetlights. The road is mostly empty, with only a few distant vehicles visible. As the car moves forward, it passes through an intersection where a red traffic light is visible on the right side. The car continues to drive straight ahead, maintaining its speed and direction. The scene remains consistent throughout, with the car moving steadily down the road. The video continues to show the same nighttime scene from inside the car. The car continues to move forward on the road, passing through another intersection where a red traffic light is visible on the right side. The road remains mostly empty, with only a few distant vehicles visible. The car maintains its steady speed and direction throughout the sequence.

Caption 8: The video begins with a view from inside a car at night, driving on a well-lit road. The road is illuminated by streetlights and the headlights of other vehicles, creating a bright path ahead. The car moves forward, passing through an intersection where traffic lights are visible, although their colors are not discernible due to the nighttime setting. The surroundings include buildings and trees lining the sides of the road, adding to the urban atmosphere. As the car continues its journey, it passes through another intersection with traffic lights, again with no clear indication of their color. The road remains well-lit, and the surroundings remain consistent with buildings and trees on either side.

Caption 9: The video begins with a view from inside a vehicle at night, driving on a road illuminated by streetlights. The road is mostly empty, with only a few distant vehicles visible. As the vehicle moves forward, it passes through an intersection where a red traffic light is visible on the left side of the frame. The scene remains consistent with minimal changes in the surroundings, maintaining the same road conditions and lighting throughout. The narrative continues with the vehicle still driving on the same road at night. The road remains mostly empty, with occasional distant vehicles passing by. The vehicle approaches another intersection where a red traffic light is visible on the right side of the frame. The scene remains consistent with minimal changes in the surroundings, maintaining the same road conditions and lighting throughout. The video progresses with the vehicle continuing to drive on the same road at night. The road remains mostly empty, with occasional distant vehicles passing by. The vehicle approaches yet another intersection where a red traffic light is visible on the right side of the frame. The scene remains consistent with minimal changes in the surroundings, maintaining the same road conditions and lighting throughout.

Caption 10: The video begins with a view from inside a vehicle at night, driving on a well-lit road. The road is illuminated by streetlights, and other vehicles are visible in the distance, some with their headlights on. The scene remains consistent as the vehicle continues to move forward, maintaining its position on the road. The narrative develops through a continuation of the same nighttime setting, with the vehicle still moving forward on the well-lit road. The road is lined with streetlights, and other vehicles can be seen in the distance, some with their headlights on. The scene remains consistent, with no significant changes in the vehicle's position or the surrounding environment.

Caption 11: The video begins with a view from inside a car at night, focusing on the road ahead. The car's headlights illuminate the road, which is lined with streetlights and signs. The scene remains consistent as the car moves forward, with other vehicles occasionally passing by or parked on the side of the road. The background features buildings and trees, adding to the urban nighttime setting. The narrative continues with the same view from inside the car, maintaining the focus on the road ahead. The car's headlights continue to illuminate the road, and the surrounding environment stays consistent with streetlights, signs, buildings, and trees. Other vehicles are seen passing by or parked on the side of the road, and the overall scene remains unchanged.

Caption 12: The video begins with a view from inside a vehicle at night, driving on a well-lit road.

The road is illuminated by streetlights and the headlights of other vehicles, creating a bright path ahead. The vehicle moves forward, passing various street signs and billboards on the side of the road. The scene remains consistent as the vehicle continues to drive down the road, maintaining its speed and direction.

The video develops through the continuation of the nighttime drive on the same well-lit road. The vehicle moves steadily forward, passing more street signs and billboards. The scene remains consistent with the previous clips, showing no significant changes in the environment or the vehicle's movement.

Caption 13: The video begins with a view from inside a vehicle at night, driving on a two-lane road. The road is illuminated by streetlights and the headlights of other vehicles, including cars and motorcycles. The surroundings are dark, with some buildings and trees visible on the sides of the road. As the vehicle moves forward, the background changes slightly, but the overall scene remains consistent with the same lighting and road conditions. The video continues to show the same view from inside the vehicle, maintaining the focus on the road ahead. The dashboard still displays the headlights of other vehicles, including a motorcycle with red lights, are visible. The road is well-lit by streetlights, and the surroundings remain dark with some buildings and trees visible on the sides. The vehicle continues to move forward, and the background changes slightly, but the overall scene remains consistent with the same lighting and road conditions.

Caption 14: The video begins with a view from inside a vehicle at night, showing the road ahead illuminated by the car's headlights. The road is wet, likely due to rain, and there are other vehicles visible in the distance, including motorcycles and cars. As the vehicle moves forward, it passes through an intersection where traffic lights are visible, and other vehicles can be seen waiting or moving around. The scene continues with the vehicle driving along the same wet road, maintaining its speed and direction.

The video then shows the vehicle continuing to drive along the wet road at night. A motorcycle with red tail lights appears in front of the vehicle, and the rider is wearing a dark jacket. The motorcycle moves slightly to the left, and the vehicle follows closely behind. The motorcycle eventually turns off the road, and the vehicle continues straight ahead. The scene transitions to another part of the road where the vehicle drives past a sign on the side of the road and continues along the wet road.

Caption 15: The video begins with a view from inside a vehicle at night, driving on a two-lane road. The road is illuminated by the vehicle's headlights and the lights of other vehicles ahead. The road has white lane markings and yellow barriers on the sides. As the vehicle moves forward, it passes through various intersections with traffic lights and streetlights. The surroundings include buildings and trees lining the road. The vehicle continues to drive straight, passing more intersections and streetlights, maintaining a steady pace.

The scene transitions to another view from inside a vehicle at night, again on a two-lane road. The road is illuminated by the vehicle's headlights and the lights of other vehicles ahead. The road has white lane markings and yellow barriers on the sides. The vehicle moves forward, passing through intersections with traffic lights and streetlights. The surroundings include buildings and trees lining the road. The vehicle continues to drive straight, passing more intersections and streetlights, maintaining a steady pace.

Caption 16: The video begins with a view from inside a vehicle at night, driving on a two-lane road. The road is illuminated by the vehicle's headlights and streetlights, casting a bright glow on the asphalt. The road has white lane markings and a yellow divider on the right side. The background shows other vehicles' lights, including red taillights and white headlights, indicating traffic in both directions. On the left side of the road, there are buildings and streetlights, adding to the urban nighttime setting.

The video progresses with the vehicle still moving forward on the two-lane road. The road's features, such as the white lane markings and yellow divider, remain unchanged. The background continues to show other vehicles' lights, including red taillights and white headlights, indicating ongoing traffic. The left side of the road still features buildings and streetlights.

Towards the end of the video, the vehicle approaches an intersection where other vehicles are present, including a bus and a truck. The road's features, such as the white lane markings and yellow divider, remain consistent. The background continues to show buildings and streetlights, maintaining the urban nighttime setting.

Caption 17: The video begins with a view from inside a vehicle at night, driving on a road illuminated by streetlights and the headlights of other vehicles. The road is marked with white lines and has a yellow divider on the right side. Buildings with lit signs line the sides of the road, and various vehicles, including cars and motorcycles, are visible in the distance. The scene captures the typical nighttime urban environment with ongoing traffic.

As the vehicle continues to move forward, the surroundings remain consistent with the previous scene. The road's white lines and yellow divider are still visible, and the buildings with lit signs continue to

line the sides. The vehicle passes by a gas station on the left side, and more vehicles, including cars and motorcycles, are seen in the distance. The scene maintains the same nighttime urban environment with ongoing traffic.

Caption 18: The video begins with a view from inside a vehicle at night, driving on a multi-lane road. The road is illuminated by streetlights and the headlights of other vehicles, creating a bright path ahead. The vehicle moves forward, passing various street signs and buildings on both sides of the road. The scene transitions to another view of the same road, still at night, with the same dashboard text and logo visible. The vehicle continues to move forward, passing more street signs and buildings, and eventually overtakes a white car. The video then shows a close-up view of the white car's rear, with its brake lights illuminated, indicating it is slowing down or stopping. The vehicle continues to pass the white car, which remains stationary. The video concludes with a view of the white car from behind, with its brake lights still illuminated, as the vehicle passes by.

Caption 19: The video captures a nighttime scene on a multi-lane road, illuminated by streetlights and the headlights of vehicles. A white car with the license plate 'KL 12 J 1638' is prominently featured in the center lane, moving forward. To the left of the white car, a motorcyclist is seen riding alongside it. The background reveals other vehicles and street signs, contributing to the urban setting. The sequence of images shows minimal movement, suggesting a steady flow of traffic. The overall atmosphere is calm and orderly, with no significant changes in the scene, emphasizing the routine nature of the journey.

Caption 20: The video begins with a nighttime scene on a road, illuminated by streetlights and the headlights of vehicles. A white car is prominently featured in the foreground, with its brake lights on, indicating it is stationary or moving slowly. To the left of the car, a motorcyclist wearing a white helmet and jacket is seen riding a green motorcycle. The rider appears to be maintaining a safe distance from the car. The background shows other vehicles and streetlights lining the road, creating a typical urban night-time setting. The video progresses with the same nighttime setting on the road. The white car is still stationary or moving slowly, with its brake lights on. The motorcyclist continues to ride alongside the car, maintaining a consistent distance. The background remains unchanged, with other vehicles and streetlights visible. As the video concludes, the motorcyclist starts to move away from the car, heading towards the left side of the frame, while the car remains stationary or moving slowly.

Caption 21: The video begins with a view from inside a car driving on a two-lane road at night. The road is illuminated by streetlights and the car's headlights, casting a bright light on the asphalt. The road is flanked by yellow barriers on both sides, and there are buildings and trees visible on the left side of the road. As the car moves forward, the surroundings remain consistent, with the road stretching into the distance and occasional streetlights and traffic signs appearing on the right side. The car continues to move forward, maintaining its position on the road. The scene continues with the same view from inside the car, now moving further down the two-lane road at night. The road remains well-lit by streetlights and the car's headlights, with yellow barriers on both sides. Buildings and trees are still visible on the left side of the road. The car continues to move forward, passing by more streetlights and traffic signs on the right side of the road. The surroundings remain consistent, with the road stretching into the distance.

Caption 22: The video begins with a view from inside a car at night, driving on a well-lit road. The road is illuminated by streetlights and the car's headlights, casting a bright light on the asphalt. On either side of the road, there are buildings with lit windows and signs, indicating commercial establishments. The car moves forward, passing various street signs and traffic lights, which are visible in the distance. The car continues to move along the road, maintaining a steady pace. The scene continues with the same nighttime setting, showing the car moving forward on the well-lit road. The surroundings remain consistent with buildings on both sides, illuminated by streetlights and the car's headlights. The car passes more street signs and traffic lights, and other vehicles can be seen in the distance, including a truck on the left side of the road. The car maintains a steady pace, and the video captures the motion of the vehicle as it moves along the road.

Caption 23: The video begins with a view from inside a car at night, driving on a well-lit road. The road is illuminated by streetlights and the car's headlights, creating a clear path ahead. On the left side of the road, there are buildings and signs, while on the right side, there are yellow barriers and some greenery. The dashboard of the car is visible at the bottom of the frame, showing the speedometer and other indicators. The car continues to move forward, maintaining a steady pace as it travels down the road. The scene remains consistent with minimal changes in the surroundings, emphasizing the focus on the road and the car's movement.

The video maintains this steady pace and consistent setting throughout, highlighting the importance of safe driving practices.

Caption 24: The video begins with a view from inside a car at night, driving on a multi-lane road. The road is illuminated by streetlights and the headlights of other vehicles. On the left side of the road, there are buildings and signs, while the right side has a yellow divider. As the car moves forward, it passes through an intersection where traffic lights are visible. The scene transitions to a wider view of the road, showing more cars and streetlights. The car continues to move forward, passing through another intersection with traffic lights and various signs. The video then shows a closer view of the road ahead, with multiple lanes and vehicles moving in both directions. The car continues to drive forward, passing through another intersection with traffic lights and various signs.

Caption 25: The video captures a nighttime view from inside a vehicle, focusing on the road ahead. The scene is illuminated by streetlights and the headlights of other vehicles. The dashboard of the car is visible at the bottom of the frame, showing the car's speedometer and other indicators. The road is busy with multiple cars moving in both directions. Some cars are closer to the camera, while others are further away, creating a sense of depth. The vehicles vary in size and color, with some appearing as bright red taillights due to their distance. The background features buildings with lit windows and signs, contributing to the urban atmosphere. The sequence of images shows the car moving forward, passing through intersections, and navigating around other vehicles. The overall scene conveys a typical city night drive, with the focus remaining on the road and the surrounding traffic.

Caption 26: The video begins with a view from inside a car at night, showing a street illuminated by streetlights and the headlights of the vehicle. In the distance, a motorcycle with its red taillight visible is seen moving away from the camera's perspective. The background features buildings with lit windows and various signs, including a large billboard with an advertisement. The scene captures the typical urban nightlife with minimal traffic. The narrative continues with the same view from inside the car, maintaining the focus on the street ahead. The motorcycle with the red taillight remains in the distance, and the surrounding environment stays consistent with the previous scene. The buildings and signs continue to be visible, reinforcing the urban setting.

Caption 27: The video begins with a nighttime view from inside a vehicle, showing a busy street illuminated by various lights. A person on a motorcycle is seen riding ahead of the vehicle, with other cars and motorcycles visible on the road. The rider is wearing a white shirt and a helmet. As the vehicle moves forward, the rider maintains a steady pace, and the surrounding environment includes buildings with lit billboards and signs. The scene continues with the same nighttime setting, where the rider on the motorcycle is now seen from behind, still wearing the white shirt and helmet. The vehicle moves closer to the rider, and the rider appears to be slowing down or stopping. The rider then loses control of the motorcycle, causing it to fall onto its side. The rider falls off the motorcycle and lands on the ground, appearing disoriented. The video concludes with the rider still on the ground, trying to recover from the fall, while the motorcycle remains on its side on the road.

Caption 28: The video begins with a nighttime view from inside a vehicle, showing a busy street illuminated by the headlights of various vehicles. A person wearing a white helmet and a dark jacket is seen lying on the ground in front of the vehicle's windshield, appearing to be injured or unconscious. The scene remains static for a few moments before the person starts to move, slowly sitting up and eventually standing up. The person then walks away from the vehicle, leaving the scene. The narrative continues with the same nighttime setting, where the person who was previously lying on the ground is now standing near the vehicle. The person appears to be assessing the situation or looking for help. The person stands still for a moment before walking away from the vehicle, moving towards the right side of the frame. The scene remains static for a few moments, with the person still walking away.

Caption 29: The video begins with a nighttime scene on a busy street, illuminated by the headlights of a vehicle. A person wearing a white shirt and dark pants is seen walking towards the vehicle from the left side of the frame. As the person approaches, more individuals gather around the vehicle, including a man in a light-colored shirt who appears to be inspecting or interacting with the vehicle. The background shows buildings with lit windows and various street signs, indicating an urban setting. The scene develops as the group of people continues to gather around the vehicle, with some standing closer and others moving around. A man in a light-colored shirt leans on the hood of the vehicle, while other individuals stand nearby, engaged in conversation or observation. The background continues to show buildings with lit windows and various street signs, maintaining the urban setting.

Caption 30: The video begins with a nighttime scene on a busy street where a group of people has gathered around a damaged car. The car, which appears to have been involved in an accident, is surrounded by several individuals who seem to be assessing the situation or discussing the incident. A person wearing a helmet and another in a white shirt are among those standing near the car. In the background, buildings and streetlights illuminate the area, adding to the urban setting. The background remains consistent with buildings and streetlights, maintaining the urban atmosphere.

The video then transitions to a split-screen view, showing two different perspectives of the same scene. On the left side, the same group of people is seen standing around the damaged car, with the person in the white shirt still present. On the right side, a close-up view of the damaged car is shown, highlighting its front bumper and hood, which appear to be significantly damaged.

Caption 31: The video captures a nighttime scene on a busy street where a group of people has gathered around a silver car that appears to have been involved in an accident. The car is positioned on the side of the road, with its front end damaged and the hood lifted. The surrounding area is illuminated by streetlights and the headlights of other vehicles, creating a contrast between the darkness of the night and the artificial light sources. Several individuals are standing near the car, some of whom are wearing white shirts with red badges or pins on their chests. One person is holding a helmet, suggesting they might be a motorcyclist. The background features buildings and other vehicles, adding to the urban setting. Throughout the video, the group of people remains mostly stationary, observing the situation. The overall atmosphere is one of concern and curiosity as the individuals gather around the damaged vehicle, possibly discussing the circumstances of the accident.

Caption 32: The video begins with a nighttime scene on a busy street where several people are gathered around a damaged car. The car has visible damage to its front bumper and hood, indicating it may have been involved in an accident. A person wearing a white shirt and helmet is seen making hand gestures, possibly explaining or discussing the situation. Another individual in a black shirt with 'INTENSITY' written on it stands near the car, gesturing towards the damaged vehicle. The background shows a bustling urban environment with illuminated buildings and streetlights, adding to the sense of urgency and activity.

The scene continues to focus on the same busy street at night. The person in the black shirt with 'INTENSITY' written on it remains near the damaged car, pointing towards it and gesturing as if explaining something. The individual in the white shirt and helmet is also present, standing close to the car. The background remains consistent with the previous scene, showing the same urban environment with illuminated buildings and streetlights.

Caption 33: The video captures a nighttime scene on a busy street, illuminated by the glow of streetlights and vehicle headlights. A group of people is gathered around a motorcycle that has been involved in an accident. The motorcycle lies on its side on the road, with its helmet placed nearby. Several individuals are standing around the motorcycle, engaged in conversation and assessing the situation. One person, dressed in a black shirt with 'INTEGRITY' written on it, is holding a mobile phone, possibly capturing the scene or communicating with someone. Another individual, wearing a blue shirt, stands near the motorcycle, seemingly observing the situation. The background reveals a bustling urban environment with buildings, street signs, and other vehicles passing by. The scene remains largely unchanged, with the group continuing to discuss the incident and the motorcycle remaining stationary on the road.

Caption 34: The video captures a nighttime scene on a busy street where an accident has occurred. The view is from inside a vehicle, with the windshield showing the aftermath of the incident. A motorcycle lies damaged in front of the car, and several people are gathered around it. One person, wearing a black shirt with 'INTENSITY' written on it, is seen moving around the area, possibly assessing the situation or looking for something. The background is illuminated by streetlights and the headlights of other vehicles, creating a contrast between the darkness of the night and the bright lights. The scene remains largely unchanged throughout the video, with minimal movement from the individuals present.

Caption 35: The video begins with a nighttime scene on a busy street, illuminated by the glow of streetlights and vehicle headlights. A group of people is gathered around a motorcycle accident, with some standing near the damaged motorcycle and others observing from a distance. The background features buildings with lit windows and billboards, adding to the urban atmosphere. A white car is parked nearby, and other vehicles are visible in the distance. The scene remains largely static, with minimal movement from the individuals present. The narrative continues with the same nighttime setting on the busy street. The group of people remains focused on the motorcycle accident, with one person walking towards the camera and another walking away. The white car and other vehicles continue to be visible in the background. The scene maintains its static nature, with minimal changes in the positions of the individuals and vehicles.

Caption 36: The video captures a nighttime scene on a busy street, illuminated by the glow of streetlights and vehicle headlights. The view is from inside a car, with the windshield reflecting the surrounding environment. In the foreground, two men stand near a damaged motorcycle, one wearing a white shirt and the other in a blue shirt. They appear to be inspecting the damage or discussing the situation. The background shows several people standing around, some observing the scene while others are engaged in conversation. Cars and motorcycles are visible on the road, indicating ongoing traffic. As the video progresses, the focus remains on the two men near the damaged motorcycle. The man in the white shirt continues to inspect the motorcycle, while the man in the blue shirt stands nearby, occasionally looking at the motorcycle. The background remains consistent with several people standing around, some engaged in conversation, and vehicles moving on the road.

Caption 37: The video captures a nighttime scene on a busy street where a group of people has gathered around an overturned motorcycle. The setting is illuminated by the glow of streetlights and vehicle headlights, creating a dramatic atmosphere. The motorcycle lies on its side, with its wheels detached and scattered nearby. Several individuals are standing around the motorcycle, engaged in conversation or observing the situation. One person, dressed in a white shirt and dark pants, stands near the motorcycle, while others are scattered around the area. In the background, buildings with lit windows line the street, and a few vehicles can be seen moving through the scene. The overall mood of the video is somber and reflective, as the people appear to be discussing the incident that led to the motorcycle's overturning.

Caption 38: The video captures a nighttime scene on a busy street where a group of people has gathered around an accident involving a motorcycle. The scene is illuminated by the headlights and taillights of various vehicles, creating a dynamic and somewhat chaotic atmosphere. The individuals present are engaged in conversation and appear to be discussing the incident. A person wearing a helmet stands near the motorcycle, which lies on the ground, indicating that it was involved in the accident. The background features buildings with lit windows and signs, adding to the urban setting. The overall mood of the video is somber and reflective, emphasizing the importance of safe driving practices.

Caption 39: The video takes place at night on a busy street, illuminated by various lights from vehicles and buildings. A person wearing a black shirt with the word 'INTENSITY' in white letters stands near a motorcycle that has been knocked over. The individual appears to be engaged in an activity involving the motorcycle, possibly adjusting or inspecting it. In the background, several people are walking around, some of whom are looking towards the motorcycle. The scene is bustling with activity, and the atmosphere suggests a sense of urgency or concern. As the video progresses, the person in the black shirt continues to interact with the motorcycle, making adjustments or repairs. The background remains busy with pedestrians moving about, and the overall atmosphere stays consistent with the previous scene. The person's actions are focused and deliberate, indicating a serious engagement with the task at hand.

Figure 17. Example of generated captions for each of the three-second segment of the video via the pipeline demonstrated in Fig. 11

The refinement prompt (Fig. 21 - 22) implements specific guidelines for different types of questions and answers. For instance, when describing road events, location information is stripped to focus solely on event characteristics. Entity descriptions maintain specificity when clearly visible (e.g. 'school bus' vs. generic 'vehicle') while avoiding unverifiable details. This structured approach ensures that final QA pairs remain answerable solely through video content while retaining the ability to describe complex road events through observable facts and reasonable inferences. For example, while we remove specific speed measurements like '80 km/h', we retain qualitative assessments like 'high speed'. Similarly, instead of stating driver was angry, we describe observable behaviors like vehicle swerved across lanes potentially intimidating a cyclist.

Eliminating potentially harmful and biased tweet con-

tent via multi-stage filtering: *Prompt-level:* Strict road-event focused prompting and curated template questions within the prompt (Supp, Fig. 8). *LLM-level:* Built-in guardrails of LLMs used in our data generation pipeline eliminate harmful and inappropriate content to some extent. *Post QA generation level:* QA pairs with overly subjective or speculative content are removed. (Supp, Fig. 14-15, Sec: 2.5). *Human verification level:* Independent annotators review each QA pair against predefined criteria. They reject pairs that violate our guidelines, e.g. including unverifiable details (Supp, Fig. 14-15). Final acceptance is determined through majority voting, minimizing individual bias. *VLM Training:* The measures described above have ensured VLMs trained on our data do not inherit any biases. In future, it might be possible to additionally mitigate social media bias by leveraging world knowledge from external sources.

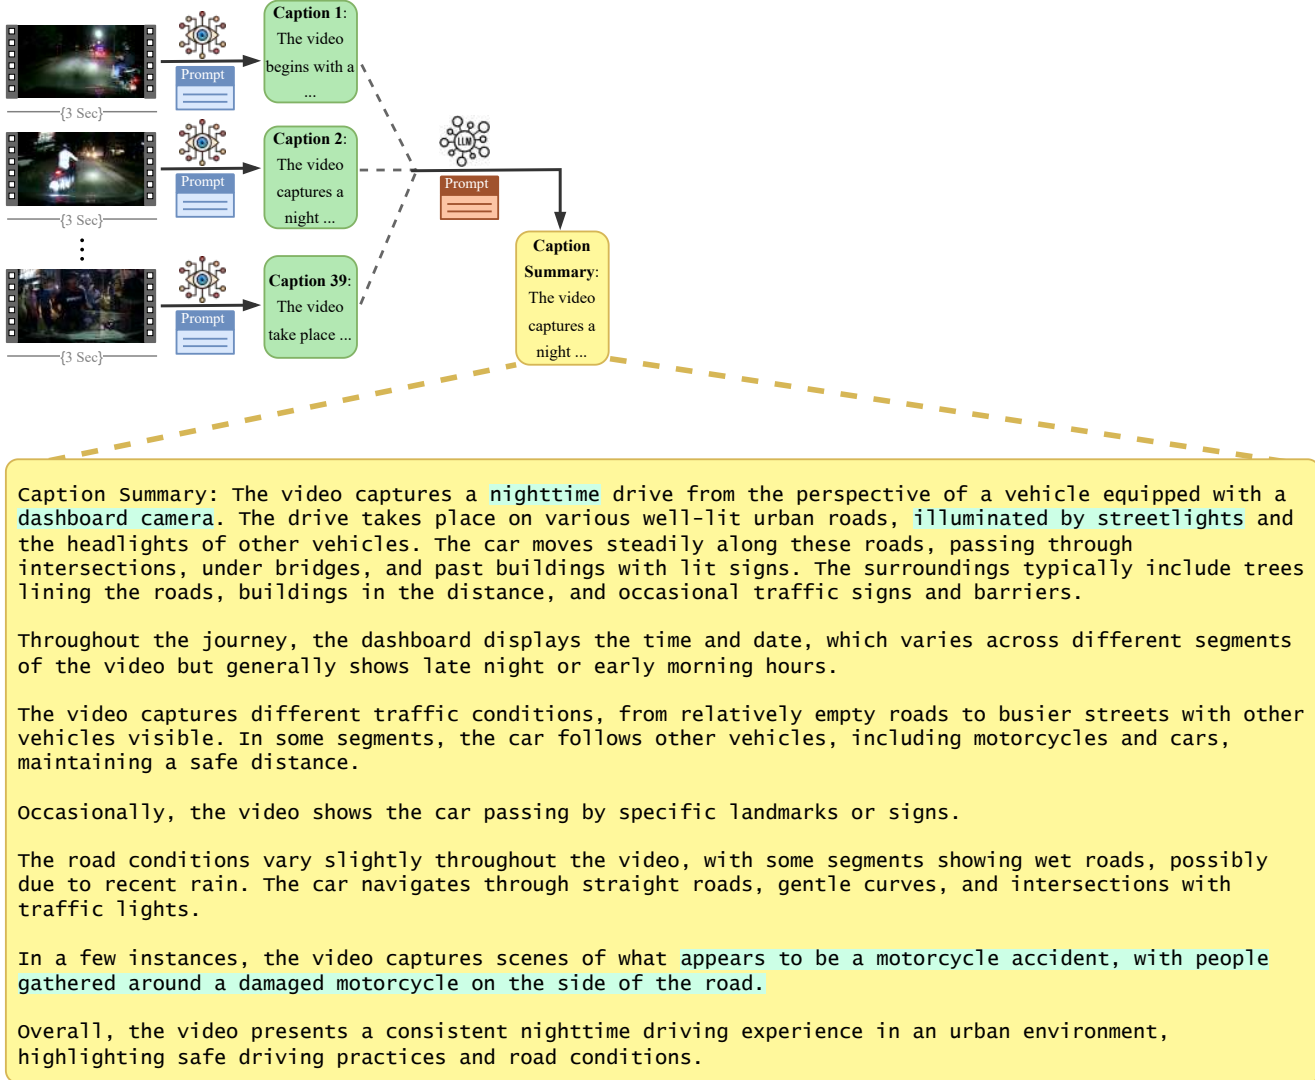


Figure 18. Example of generated video summary from the captions via Text LLM, via the pipeline demonstrated in Fig. 11.

QA Categorization: We developed a categorization framework for our refined QA pairs to match each question from the refined QA pairs against the same set of 18 template questions that were used in our hybrid QA generation approach (see Appendix B.3). These template questions span a range of complexity levels, from basic observational queries (e.g., camera type, weather conditions) to complex analytical questions addressing causation, prevention, and counterfactual scenarios.

In this matching process, we prompt Claude 3.5 Sonnet Text LLM to assign a similarity score from 0 to 5 to each QA pair, where 5 indicates perfect alignment with a template question and 0 indicates no meaningful similarity. For example, a question like `What recording device was used?` would receive a high similarity score with

the template `What type of camera was used to capture the video?`, while a question about video purpose would receive a score of 0 as it doesn't align with any template. To ensure categorization quality, QA pairs receiving low similarity scores undergo expert review for potential refinement or removal. This human-in-the-loop validation helps maintain the integrity of our categorization while ensuring comprehensive coverage across all aspects of road event analysis. The complete categorization process and scoring mechanism is implemented through carefully designed prompt () shown in Fig. 23 - 24.

B.6. Incompatible QA Generation

To evaluate Video LLMs' resilience to hallucination and their ability to discriminate between road and non-road events, we developed an approach for generating incompat-

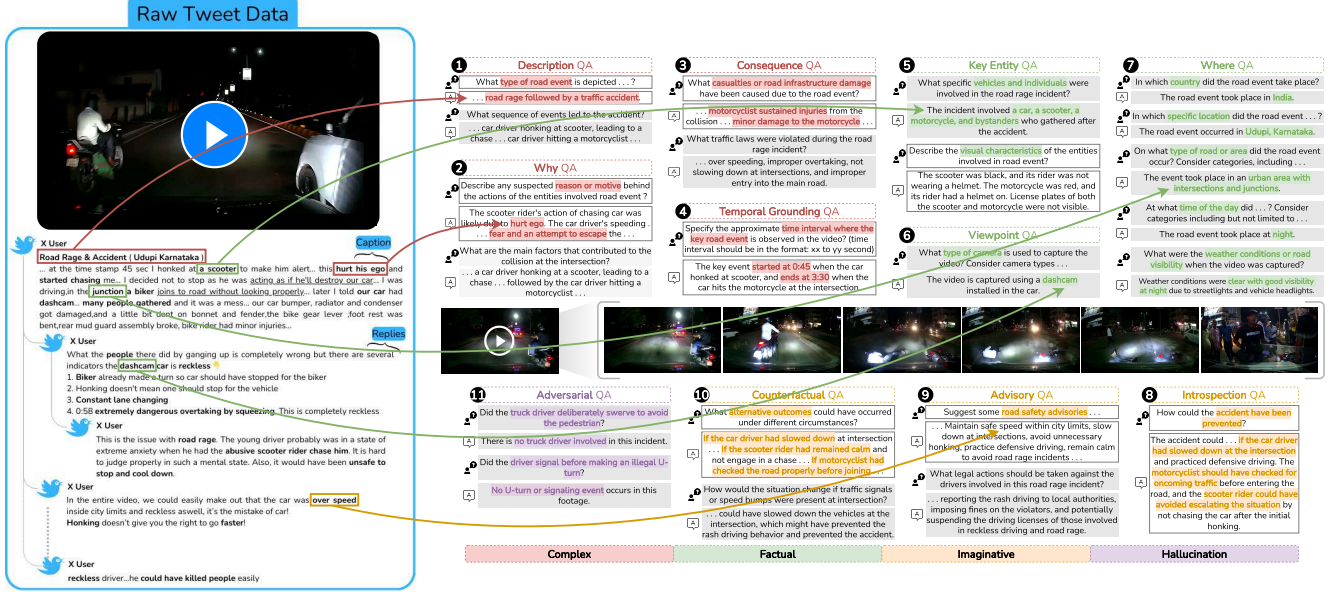


Figure 19. **Demonstrating the importance of hybrid information sources for QA generation.** While the video summary (Fig. 18) captures basic visual elements (e.g. ‘nighttime’, ‘streetlight illumination’), the tweet conversation provides crucial contextual information (shown by colored boxes) missing from the visual description alone. This example illustrates why our approach combines both video summaries and social media context to generate diverse and socially-informed QA pairs.

Specific QA Generation Prompt

You are an expert on understanding Road and Traffic Events and have extensive knowledge of safe driving norms across different geographical regions.

You are provided with a set of question-answer (QA) pairs and a video summary in JSON format. These QA pairs are designed for road event understanding through videos.

Your task is to regenerate the QA pairs that are more specific and contextually appropriate.

Output Format:

Your entire response must be in JSON format as follows:

```
{
  "Summary": "",
  "QAS": [
    {"Q": "", "A": "", "Question-Type": ""},
    ...
  ]
}
```

Figure 20. Prompt design for generating specific QA pairs from generic templates. The expert-driven prompt transforms generic questions into contextually specific ones while maintaining alignment with predefined categories (e.g., Camera Device, Road Event Type, Actions). The prompt takes generic QA pairs and video summaries as input and generates specific questions that capture detailed aspects of road events while preserving the taxonomic structure. Refer back to Appendix B.4.

ible QA pairs. The process involves a three steps: video classification, summarization, and QA generation.

Identifying non-road event videos: First, we employ a specialized prompt (Fig. 25) to classify videos as road or non-road events. The prompt (Fig. 25) defines road events as ‘any

incident, activity, or condition occurring on or around the roadway that affects traffic flow, safety, or road usage’ and assigns a confidence score between 0 and 1, providing detailed reasoning for the classification.

Non-road event video summarization: Videos confirmed as

QA Refinement Prompt

The objective of the input QA pairs is to help understand different aspects of the road and traffic events or provide insights into different safe driving practices based on the posted video. However, access to the video and the textual conversation is not available.

YOUR TASK IS TO:

1. Rewrite the answers given in JSON format such that they DO NOT give any reference of the video summary and the textual conversation when answering the questions. Also, avoid using phrases like 'based on the replies', 'based on the comments', 'based on the conversation', 'based on the text', 'mention', 'conversation', 'caption', 'replies', 'comment', 'post', 'twitter', 'youtube', 'hashtag', 'user', <name of the person>, 'QA', etc., that may indicate that the generated answer is based on the textual conversation. Instead, refer to the information as being inferred from the video.
2. If any part of the answer lacks sufficient information, either remove that specific part or exclude the entire question and answer pair.
3. Make corrections in the existing QA pairs as per the following additional instructions:

SPECIFIC INSTRUCTIONS:

Describe the Type of Road Event: When the question asks for the "type of road event", do not include the location in the answer. Only mention the event type and describe it. Following are the examples of 'valid' and 'invalid' answers. Note that the 'Invalid' answer contains the name of the location.

Valid Example

Question: Describe the type of road event captured in the video.

Answer: The video captures multiple traffic violations and unsafe driving behavior.

Invalid Example

Question: Describe the type of road event captured in the video.

Answer: The video captures multiple traffic violations and unsafe driving behavior during a typical morning in Hoboken.

A few examples of the "type of road events or traffic events:

- Road Safety awareness and safe driving education through demonstration, advertisement, social initiative, or traffic rule or policy change: It can be about front and back seatbelt usage, safe driving advisory around vulnerable road users, driving tests, impact of over speeding, wearing helmet, car not rushing lane change, etc.
- Road Rage: Attacking and causing damage to a vehicle, riders using derogatory body language, absurd language by riders, cutting in and stops another vehicle's path and getting into argument, bus drivers fighting on the road, etc.
- Dangerous or Rash Driving: Using phone while riding, abrupt lane changes, zig zagging behavior, wheeling stunt, standing on bike while riding, overtaking during turn, person jaywalking in front of a car, etc.
- Accident: Car overturned, collision between vehicle and road infrastructure, bus skids due to hydroplaning, collision between vehicles, etc.
- Near-miss: Near-miss incident between car and pedestrian, Near-miss incident between motorcycle and truck, etc.
- Faulty Road Infrastructure: Traffic cone fallen in middle of road, no amber light only green or red on the intersection, etc.
- Animal Related Accident: Deer jumped onto a biker, dog ran over by truck, etc.
- Traffic Violations: Overtaking by crossing the solid yellow line, wrong side driving, riding on footpath, not wearing helmet, triple riding on motorcycle, red light violation, parking vehicle on footpath, etc.
- Post Crash: Report of crashed vehicles, overturned vehicle on slippery road, multiple vehicle collisions on slippery road, etc.
- Defensive Driving: maintaining safe following distance, avoiding distractions, speed control, lane discipline, yielding to pedestrians, etc.

Listing Road Entities: When the question asks for "List down all the road entities observed/present/visible in the video.", A few examples of how you should rephrase this question are: "What entities are involved in the road event?" or "What entities are observed during the road event?" or "What entities are involved in the accident?" or "What entities are involved in the traffic violation?"

or "What entities are involved in the near-miss incident?", etc.

Speculative Answers: If the answer involves speculating the cause of an event that is not directly obvious just by watching the video, use terms like "potential", "likely" or "probably" to indicate speculation. Following are the examples of 'valid' and 'invalid' answers. Note that the 'Invalid' answer contains the cause or reason that is highly unlikely to be known by watching only a video.

Valid Examples

Question: What type of road event is depicted in the video?

Answer: The video depicts a dangerous and rash driving incident involving a potentially drunk driver ramming into a person standing near a toll booth, resulting in serious injuries to the person.

Question: Describe the actions of the road entities involved in the event.

Answer: The driver of the car was probably asleep and that's why diverted from its lane, leading to a head-on collision with the white van.

Invalid Examples

Question: What type of road event is depicted in the video?

Answer: The video depicts a dangerous and rash driving incident involving a drunk driver ramming into a person standing near a toll booth, resulting in serious injuries to the person.

Question: Describe the actions of the road entities involved in the event.

Answer: The driver of the car was asleep and that's why diverted from its lane, leading to a head-on collision with the white van.

Objective Information: Express only factual or objective information about the road event, avoiding subjective feelings. Following are the examples of 'valid' and 'invalid' answers. Note that in the first example, "intentionally" is a subjective term because just by watching the video, one wouldn't know if the action performed by the vehicle which is "swerving" is intentional or not. Also note that "intimidation" is again a subjective feeling that might not directly be visible just by watching a video. Therefore either remove it or use a speculative term like "potential", "likely", etc.

Valid Examples

Question: What type of road event is depicted in the video?"

Answer: The video captured an instance of aggressive driving, where a vehicle swerved across lanes potentially intimidating a cyclist.

Question: Can you describe the actions of the drivers in the road event?

Answer: The actions performed include collision and verbal confrontation.

Invalid Examples

Question: What type of road event is depicted in the video?"

Answer: The video captured an instance of aggressive driving, where a vehicle intentionally swerved across lanes to intimidate a cyclist.

Question: Can you describe the actions of the drivers in the road event?

Answer: The actions performed include intentional collision and verbal confrontation.

Weather Conditions: Do not include specifics about the weather condition in the answers. Following are the examples of 'valid' and 'invalid' answers. Note that the 'Invalid' answer contains the specific number like '161' AQI that is unlikely to be known from a video.

Valid Example

Question: What were the weather conditions or road visibility when the video was captured?

Answer: The sky was hazy.

non-road events then undergo detailed visual summarization using a structured prompt (Fig. 26) that captures key

visual details, temporal sequences, and object interactions.

Generating mismatched QA pairs: For generating incom-

```

##### Invalid Example #####
Question: What were the weather conditions or road visibility when the video was captured?
Answer: The sky was hazy and the Air Quality Index (AQI) was 161 when the video was captured.

Specific Entity Involvement: If more information is provided in the text regarding the road entities involved in the accident, mention them in answer. Do not mention the name of the road entities if unsure. The answer in "Could be better example" is correct but as shown in "Good example".

##### Could be better Example #####
Question: What type of road event is depicted in the video?
Answer: The video depicts a head-on collision.

##### Good Example #####
Question: What type of road event is depicted in the video?
Answer: The video depicts a post crash event showing the damage caused by a head-on collision between a school van and a school bus. It highlights the issue of unsafe school transportation and the need for ensuring child safety on roads.

Exact Speed: Do not include exact speed in the answers.
Following are the examples of 'valid' and 'invalid' answers. Note that the 'Invalid' answer contains the specific number '80 km/h' which is unlikely to be known by watching a video.

##### Valid Example #####
Question: Describe the actions performed by the listed road entities above.
Answer: The car in the video was likely driving at a high speed.

##### Invalid Example #####
Question: Describe the actions performed by the listed road entities above.
Answer: The car in the video was likely driving at a high speed, possibly around 80 km/h.

Sentence Form Answers: Answers should be in sentence form, not lists or dictionaries.
Following are the examples of 'valid' and 'invalid' answers. Note that the 'Invalid' answer contains a list. The response should sound like a human has written it as shown in the 'valid' answer.

##### Valid Example #####
Question: List down all the road entities observed during the road event.
Answer: vehicles and a disabled traffic signal.

##### Invalid Example #####
Question: List down all the road entities observed during the road event.
Answer: ['vehicles', 'disabled traffic signal']

```

Figure 22. **QA refinement prompt**: The prompt implements comprehensive guidelines for (1) removing social media references, (2) standardizing temporal information, (3) enforcing speculative language for non-obvious causation, (4) maintaining objective observations, and (5) ensuring human-like sentence-form responses. Refer back to Appendix B.5

patible QA pairs, we sample questions generated for our road event videos and apply them to these confirmed non-road videos. We modify our hybrid QA generation approach by incorporating additional prompting constraints. The modified prompt explicitly acknowledges the video's non-road nature and requires the model to articulate the incompatibility between road-event questions and the video content based on the non-road even summary generated in the previous step. This approach generates responses that highlight the fundamental mismatch between the question's assumptions and the video's actual content.

This methodology serves multiple purposes: testing QA generation pipeline robustness, evaluating Video LLMs'

ability to recognize and reject inappropriate questions, and generating training data for improving model discrimination. Fig. 25, Fig. 26 and Fig. 27 illustrates the prompts utilized for generating Incompatible QA pairs for non-road event videos.

B.7. Adversarial QA Generation

To evaluate Video LLMs' ability to recognize and reject misleading assumptions, we developed an approach for generating adversarial QA pairs. These QA pairs specifically test models' capabilities in identifying non-occurring road events and avoiding hallucination by introducing questions about events, objects, or actions that are not present in

QA Categorization Prompt

You are an AI assistant tasked with matching each of the 'Q' from the user-provided QA pairs with a set of template questions and finding which template question, a 'Q' is the most closest too.

User provides a set of QA pairs as a list of dictionaries format, in the form:

```
[{'Q': '', 'A': ''},  
 {'Q': '', 'A': ''},  
 {'Q': '', 'A': ''}]
```

These are the final (Generic & Specific) QA pairs obtained after the Refinement process

Each of the QA pairs might belong to one of the following template questions:

<Question-1> What type of camera was used to capture the video? </Question-1> (Type-of-Camera e.g., dashcam, vehicle-mounted camera, hand-held camera, cell-phone camera, cctv camera, surveillance camera, drone camera, multiple-cameras i.e., not a fixed view point, etc. Do not specify the name of the camera model, just specify its type.)

<Question-2> Describe the type of key road event captured in the video. </Question-2> (Type-of-Road-Event e.g., safe/unsafe road infrastructure or driving behavior, dangerous, rash, or aggressive driving, road rage, traffic violation, accident/crash, post-crash, near-miss, awareness of road safety, defensive driving, etc.)

<Question-3> In which country did this road event take place? </Question-3> (Country-of-Origin e.g., India, UK, US, Japan, China, etc. Do not justify how you got the answer)

<Question-4> In which state, district, city/town/village, or locality did the road event occur? </Question-4> (Location could be the name of a state, district, landmark, type of locality, like city/town/village, etc. Specific-Location e.g., Hyderabad city, Big Ben London, etc.)

<Question-5> On which type of road or area, this event have taken place? </Question-5> (Type-of-Road e.g., urban area, rural area, highway, flyover, turn, intersection, tunnel, bridge, T-junction, roundabout, hilly or mountain area, etc. Do not justify how you got the answer. Do not specify the name or address of the region where the event took place, just specify its type.)

<Question-6> When did this road event happen? </Question-6> (Time-of-Day e.g., morning, afternoon, evening, night, etc. Do not specify the exact date or time in the generated answer.)

<Question-7> What were the weather conditions or road visibility when the video was captured? </Question-7> (e.g., sunny, rainy, windy, foggy, low visibility, etc.)

<Question-8> List down all the road entities involved in the key road event. </Question-8> (A road entity can include road infrastructure objects like traffic signs, lane markings, barricades, etc. Road entities can also include road users like cars, bikes, pedestrians, drivers, etc.)

<Question-9> Describe the visual characteristics of the listed road entities above </Question-9> (e.g., what was the vehicle's color?, was the headlight, brake light, or turn signal on?, what was the license plate number?, was the rider wearing helmet or seat belt?, etc.)

<Question-10> Describe the actions performed by the listed road entities above. </Question-10> (e.g., illegal overtaking, overspeeding, swerving, yielding, cutting, etc.)

<Question-11> Describe any suspected reason or motive behind the actions of the involved road entities. </Question-11> (e.g., thrill, road rage, impressing others, in a rush, aggressive, impatient, etc.)

<Question-12> If the road event involves an accident or a near-miss incident, what was the primary reason behind its occurrence? </Question-12> (e.g., road rage, etc.)

<Question-13> If the road event involves an accident or a near-miss incident, Explain how it could have been prevented. </Question-13> (e.g., by slowing down at the intersection, checking the rearview mirror, etc.)

<Question-14> If the road event involves an accident, list down any casualties or road infrastructure damage during the event. </Question-14> (e.g., people in the car died, bikers got injured, pedestrians got hit by car, divider was damaged, etc. Do not specify the exact number of casualties (e.g., 5 pedestrians or 3 people) in the generated answer.)

```

<Question-15> List down all traffic rule violations associated with this road event </Question-15> (e.g.,
illegal overtaking, illegal overtaking by crossing solid lane markings, hiding license plates, license
plate not visible, helmet rule violation, no helmet, wrong-side driving, triple riding violation, red
light violation, drunk driving, etc.)

<Question-16> What measures should be taken upon witnessing an unsafe driving situation during this road
event? </Question-16> (e.g., reporting any traffic violation, or unsafe road infrastructure to local
government authorities or police, fines, jail time, license ban, vehicle confiscation, etc. List only the
most relevant measures.)

<Question-17> List down all the road safety advisories corresponding to the listed road entities.
</Question-17>

<Question-18> List down all the Counterfactuals related to different road events or driving situations
that could have happened under different circumstances. </Question-18> (e.g., the biker would have met an
accident if the truck steered a little towards the right, the incident could have been worse if there
were pedestrians by the roadside, If the car had not been speeding, it would have safely stopped before
the intersection and avoided being hit by the truck, etc.)

### Task:
1. Read each of the questions 'Q' in the JSON, and compare it with all the questions in the 'Sample
Questions', one by one.

2. Find out the most similar sample question corresponding to the 'Q', and output the index of that
sample question. Also, assign similarity score between 0 to 5, where 0 indicates no similarity and 5
indicates the highest similarity.

For example, if <Question-5> is most similar to the 'Q', then output: "sample_question_index": "5",
"sample_question_similarity_score": "4"

3. If a 'Q' is not similar to any of the sample questions, then assign the "sample_question_index": "-1"
and "sample_question_similarity_score": "0"

### Example Response:
[
  {"Q": "what type of camera recorded the video?", "A": "The video was recorded using a bicycle-mounted
camera.", "template_question_index": "1", "template_question_similarity_score": "4.9"},
  {"Q": "what road entities are involved in the road event.", "A": "The road entities involved are
construction workers, paving stones, traffic cones, construction materials, and a partially completed
paved area.", "template_question_index": "8", "template_question_similarity_score": "4.7"},
  {"Q": "What was the purpose of capturing this video?", "A": "The video was captured to provide proof
for reporting the incident to legal authorities.", "template_question_index": "-1",
"template_question_similarity_score": "0"},
]

```

Figure 24. QA categorization prompt design. The prompt (1) matches each refined question against 18 predefined template questions used in QA generation, (2) assigns similarity scores (0-5), and (3) provides examples demonstrating proper template matching for various question types. Refer back to Appendix B.5

the video.

The generation process employs Claude 3.5 Sonnet Text LLM to analyze the generated QA pairs (from Appendix B.5) associated with a video and create new questions that maintain the road safety context while introducing irrelevant elements. For example, given a video showing a simple traffic violation, an adversarial question might ask about non-existent casualties or emergency responses. The answers are carefully crafted to explicitly state the absence of these elements while

maintaining a video-centric perspective (e.g. The video shows no emergency vehicles or medical response as there were no casualties in this traffic violation incident).

This approach differs from Incompatible QA generation as it maintains the road event context while testing for fine-grained discrimination. While Incompatible QAs evaluate model robustness on completely unrelated videos, Adversarial QAs test the model's ability to reject false premises within relevant road scenarios. Representative examples

of adversarial QA pairs are shown in Fig. 29, 30 and 31. Fig. 28 demonstrates prompt (Fig. 28) to generate adversarial QA pairs from existing QA pairs.

B.8. Temporal Grounding QA generation

The answers in “Description QA”, auto generated by our annotation pipeline, provides details of key video events (e.g. ‘car hitting a biker’). Annotators manually mark the temporal extent (start, end) of these key events.

B.9. Final QA Task Taxonomy

We developed a structured taxonomy to evaluate Video Large Language Models’ (Video LLMs) capabilities across different reasoning categories. Our taxonomy organizes question-answer (QA) pairs into four primary reasoning categories: Complex, Factual, Imaginative and Hallucination reasoning. These categories are further divided into 12 distinct QA tasks designed to assess specific aspects of video understanding. Through our QA generation pipeline (discussed till now), each QA pair is assigned one of 19 template question IDs. We then map these template IDs to specific QA tasks within our taxonomy, as illustrated in Fig. 29. This systematic mapping enables structured evaluation of Video LLMs across different reasoning capabilities while ensuring comprehensive coverage of our taxonomy’s tasks.

B.10. Video-level Tag Generation

To enable efficient retrieval and analysis of videos based on content characteristics, we developed a tag generation system that generates diverse video-level tags from refined QA pairs (Appendix B.5). Our method employs Claude 3.5 Sonnet Text LLM [1] to analyze QA pairs and generate tags across multiple categories (as shown in Fig 1, Main Paper).

For each QA pair, based on its template question category (Appendix B.9), the LLM generates specific tags following structured guidelines provided in prompt Fig. 33 - 34).

The resulting video tags or video attributes provide fine-grained details about road scenarios, enabling efficient video retrieval and analysis. Their distribution is shown in Fig. 55 - 60.

C. Experiments

C.1. Data Setup

For model evaluation, we provide the model with video frames and a task-specific question, following the format: video frames + model’s default system prompt (if any) + our task-specific question. An example prompt for LLaVA-Video [39] in the specified format is given in Fig. 6. Also, the prompting structure for QA tasks in our dataset is described in Fig. 36.

Model	DriveLM Planning [28]	Lingo-QA Eval [17]
LLaVA-OV [10]	31.7	37.0
LLaVA-OV ft.	40.1 (+8.3%)	41.6 (+4.6%)

Table 3. Performance comparison of LLaVA-OV [10] with and without fine-tuning on RoadSocial dataset.

C.2. Model Setup

For evaluating model-generated responses against ground-truth answers in our open-ended QA tasks, we employ GPT-3.5 score [18]. Our evaluation method prompts GPT-3.5 to act as an expert assessor, analyzing the semantic alignment between predicted and ground-truth answers. For each QA pair, the system generates a similarity score on a scale of 0 to 100, where higher scores indicate closer alignment with the ground truth. To ensure interpretability, each score is accompanied by a detailed explanation justifying the rating. This approach provides transparent insights into the evaluation process while maintaining reproducibility. Fig. 37 demonstrates GPT evaluation prompt (Fig. 37). For the Temporal Grounding task, we employ separate metrics (average mAP) better suited to temporal alignment assessment.

Evaluation Details: For zero-shot evaluation of 18 models across 12 tasks on the test set, we utilized a cluster of 16 NVIDIA H100 GPUs. We maintained a batch size of 1 for all model evaluations and used the sampling frame rate and input video resolution parameters as recommended in their respective official repositories. The evaluation process was parallelized across multiple GPUs, resulting in an average evaluation time of 4 hours per model. For the closed-source models (Gemini 1.5 Pro [gemini-1.5-pro-latest], GPT-4o), the evaluation on 12 tasks required approximately two days per model due to API rate limitations.

Our evaluation period spanned October-November 2024, with specific access windows for different models: Gemini 1.5 Pro and GPT-4V (November 1-6, 2024). Claude 3.5 Sonnet was used from October 1-29 and GPT-3.5-turbo was used from October 18 to November 10, 2024.

C.3. Qualitative Analysis

The qualitative results are shown Fig. 38 onwards.

C.4. RoadSocial’s Utility for Planning/AV tasks

RoadSocial contains a significant number of egocentric road videos. Planning-related QAs in RoadSocial are distributed across multiple tasks, such as “Advisory”, “Counterfactual”, “Description” and “Why”, covering critical and planning-related road events like dangerous driving, near-misses, and defensive driving. Similarly, perception-related QAs for such videos refer to ego-relative important objects (called “Key Entities” in our dataset). To demonstrate the utility of the mentioned QA types for planning and perception tasks, we fine-tuned a Video-LLM (LLaVA-OV)

Prompt to Determine if a Video is related to Road Event or Not

You are an expert on understanding Road and Traffic Events and have extensive knowledge of safe driving norms across different geographical regions. A road event typically refers to any incident, activity, or condition occurring on or around the road way that affects traffic flow, safety, or road usage.

You are provided with a video as input. Determine if the video is related to a road event. Respond with "True" if it is related to such an event, or "False" if it is not. Provide reasoning for the decision. Also provide a score between 0 - 1 where 1 stands for highest confidence for the video to be road event and 0 means not a road event.

Your response must be in JSON format as specified below:

```
{
  "Reason": "",
  "is_road_or_traffic_event": <"True", "False">,
  "Score": "";
}
```

Figure 25. Prompt design for road event classification. The prompt implements binary classification (road/non-road) with confidence scoring (0-1) and reasoning requirements for video content. Refer back to Appendix B.6.

Prompt to Generate Summary of Non-Road Event Video

Generate a detailed and accurate description of a video. Use the following details to create a clear and complete narrative:

Instructions for writing the detailed description:

1. Focus on describing key visual details such as appearance, motion, sequence of actions, objects involved, and interactions between elements in the video.
2. Emphasize important points like the order of events, appearance and actions of people or objects, and any significant changes or movements.
3. Give a thorough description, highlighting the key visual and temporal details while keeping it clear and easy to understand.

Figure 26. Video summarization prompt for non-road events. The prompt ensures structured description of visual content focusing on key details, temporal sequences, and object interactions. Refer back to Appendix B.6.

Fine-tuning dataset	Factual			Complex			Imaginative			Hallucination		Overall	Overall	Overall	Overall	
	WR	KE	VP	DS	WY	CQ	TG	AD	IN	CF	AV	IC	(ALL)	(RT)	(Generic)	(Specific)
RoadSocial	80.6	63.9	85.7	64.0	68.7	65.0	4.49	74.1	70.8	71.7	95.4	84.7	69.1	67.7	64.2	70.6
Video-only	67	53.4	68.3	32.7	34.7	43.5	0.08	55.8	42.1	51.5	83.2	71.4	50.3	44.9	45.3	43.6

Table 4. LLaVA-OV fine-tuned on RoadSocial with Video and Commentary v/s fine-tuned on only Video-based QA dataset.

[10] on our dataset. We evaluated the model on representative autonomous driving benchmarks: PlanningQA task in DriveLM[28] and Action/Scenery QA task in LingoQA[17]. The substantial improvement in performance after fine-tuning on RoadSocial (Tab. 3) demonstrates our dataset’s utility for evaluating video-based planning/AV tasks.

C.5. Video-only QAs falls short

We regenerated entire QA dataset using only video-based summaries and used it to fine-tune LLaVA-OV model[10]. We compared this model with the LLaVA-OV model described in paper (fine-tuned using QA obtained from video and social conversations). The performance gap (Tab. 4) highlights the crucial role of social commentary in enhancing QA data quality.

D. Acknowledgement

Our sincere gratitude goes to late B.V. Khadiravana for his invaluable help in optimizing the execution of our experiments. The project was supported by iHub-Data and Mobility at IIIT Hyderabad.

QA Generation Prompt for Non-Road Event Video

A road event refers to any incident, activity, or condition on or around the roadway that affects traffic flow, safety, or road usage. This includes events such as accidents, traffic violations, unsafe driving behavior, road rage, road safety awareness, and other related occurrences.

You are an expert in understanding road events with knowledge of traffic rules, road infrastructure, and driving behaviors across various geographical regions.

You are provided with a textual conversation related to a video posted on Twitter, including a caption and, in some cases, replies. Additionally, you have a detailed summary of the video. However, you do not have access to the actual video footage.

It is confirmed that the video is not related to any road event. Your task is to explain why it is unrelated.

Please respond in the following JSON format:

```
{
  "Reason": "The video does not involve any road event because....",
}
```

Guideline for Response:

DO NOT give any reference of the video summary and the textual conversation when generating the response. Also, avoid using phrases like 'based on the replies', 'based on the comments', 'based on the conversation', 'based on the text', 'mention', 'conversation', 'summary', 'caption', 'replies', 'comment', 'post', 'twitter', 'user', <name of the person>, etc., that may indicate that the generated response is based on the textual conversation. Instead, refer to the information as being inferred from the video.

Figure 27. Incompatible QA generation prompt. The prompt generates explanations for why road event questions are incompatible with non-road video content while maintaining established response formats. Refer back to Appendix B.6.

Adversarial QA Generation Prompt

You are an expert in understanding Road and Traffic Events, with extensive knowledge of safe driving norms across various geographical regions. You are provided with a set of question-answer (QA) pairs and a summary in JSON format, corresponding to a video. These QA pairs are designed for road event understanding through videos, where the goal is to evaluate models' ability to answer questions based on visual input.

Your task is to:

1. Regenerate QA pairs by introducing questions that ask about events, objects, or actions that are not present or are irrelevant to the actual QA pairs. These new questions are designed to test the model's ability to avoid hallucinating answers to irrelevant questions.
2. Ensure the answers to these irrelevant questions explicitly state that the events, objects, or actions mentioned in the question are not present or irrelevant in the video.
3. Ensure the answers to these irrelevant questions sound like they have been answered purely by looking at the video.

Output Format:

Your response must be in the following JSON format:

```
{
  "QAs": [
    {"Q": "<irrelevant question>", "A": "<Response explaining the irrelevance or absence of the asked event, object, or action>"},
    {"Q": "<Another irrelevant question>", "A": "<Response to the question making it clear that the premise of the question is false or does not exist>"},
    ...
  ]
}
```

The output QA pairs should align with the original context of the video and should serve as ground truth to evaluate whether models correctly identify irrelevant or non-existent events in road and traffic event videos.

Figure 28. Adversarial QA generation prompt. The prompt instructs the generation of questions about non-occurring events while maintaining road context, with examples demonstrating (1) proper introduction of irrelevant elements, (2) explicit negation in answers, and (3) preservation of video-centric response format. Refer back to [Appendix B.7](#).

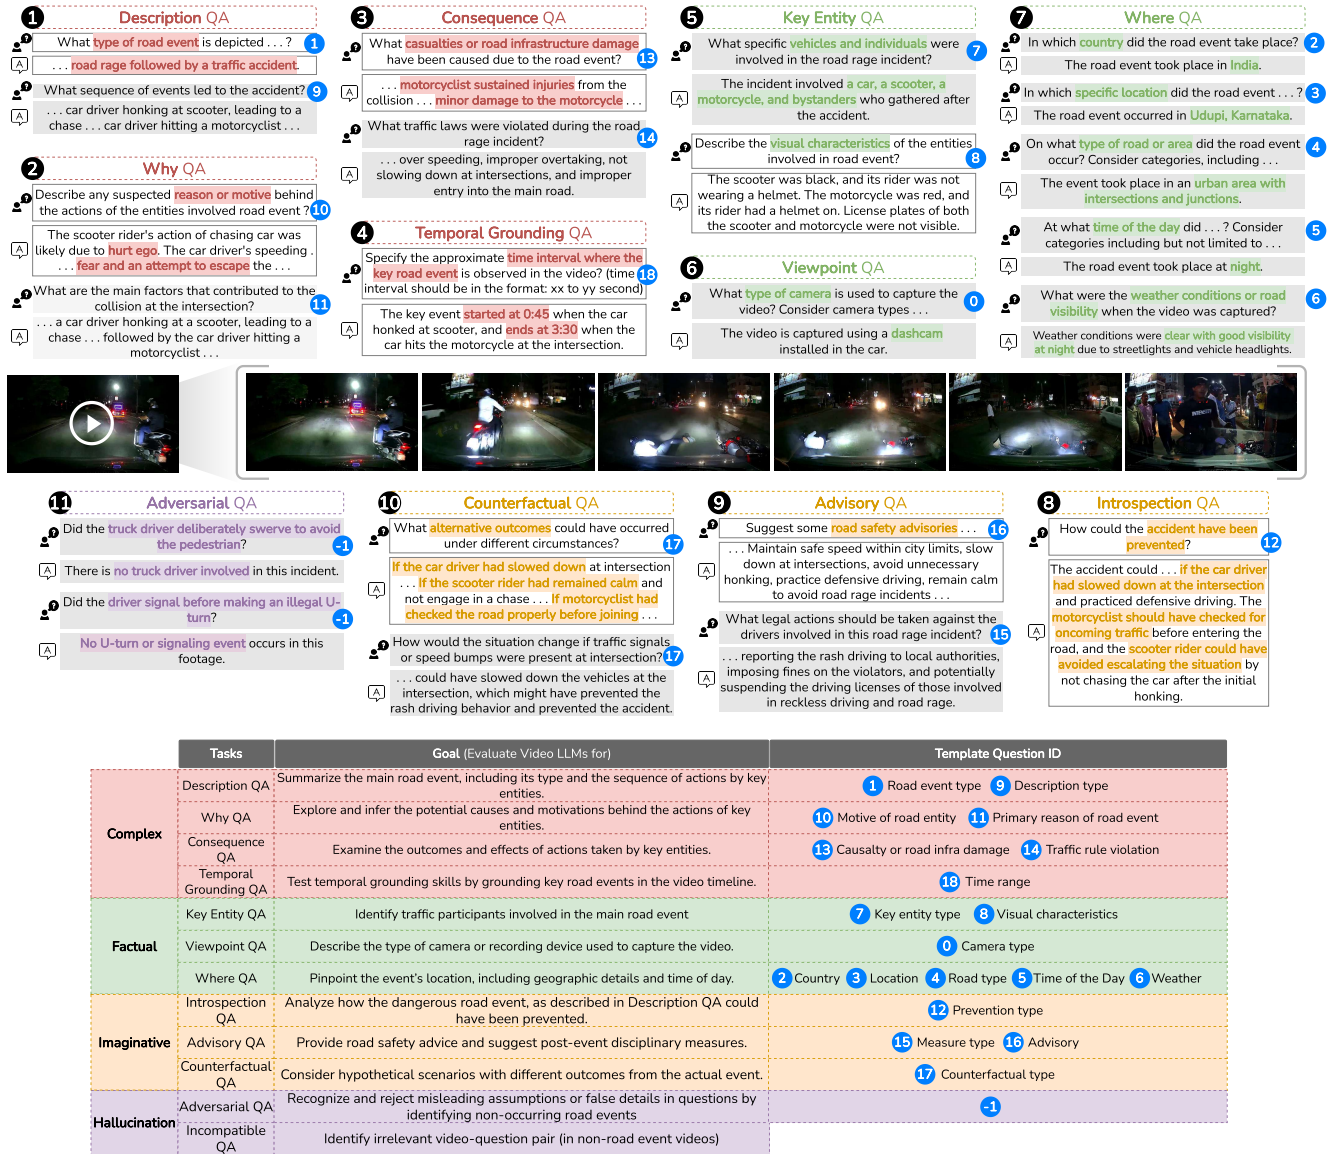


Figure 29. **QA Task Taxonomy and Template Question ID Mapping for Video LLM Evaluation:** The taxonomy consists of 12 QA tasks organized into four reasoning categories: Complex (red), Factual (green), Imaginative (orange), and Hallucination (purple). The 19 template question IDs (blue circles) map to QA tasks designed for evaluating road event understanding. For Incompatible QAs, which evaluate model robustness on non-road event videos, we employ a separate generation pipeline (Appendix B.6) without template ID mapping. Refer back to Appendix B.9.

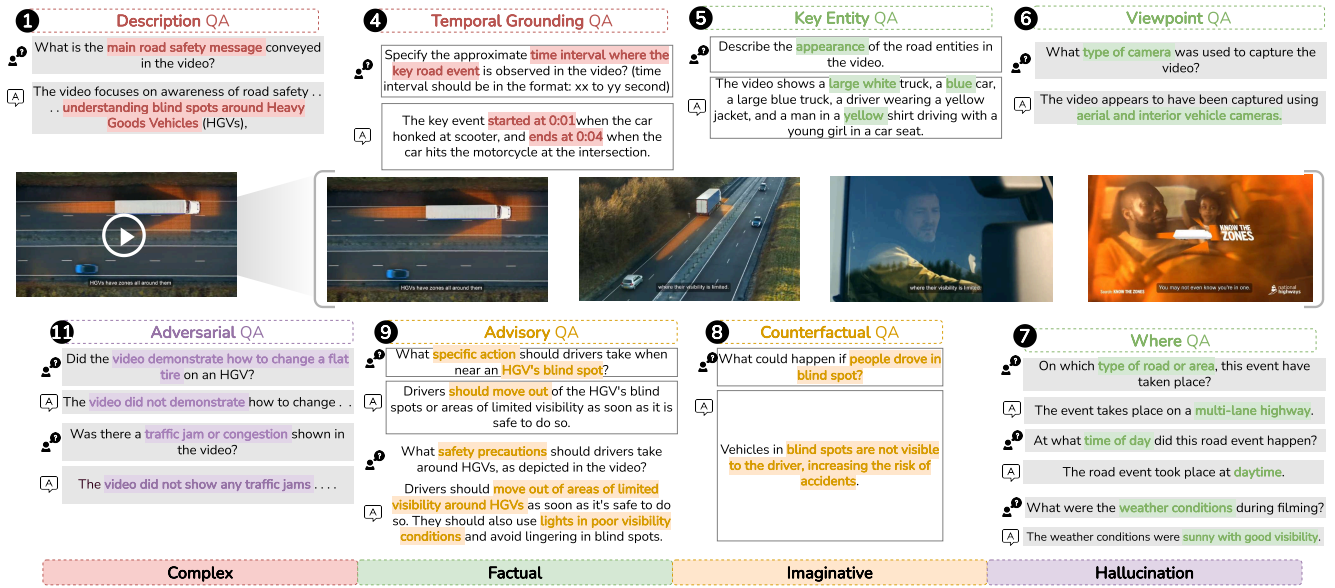


Figure 30. Examples of QA Pairs grouped by tasks and color-coded by task category (for an advertisement video captured via multiple viewpoints). Gray fill shading indicates specific questions while the non-shaded QAs are generic. Highlighted text indicates key information. Refer back to Appendix B.9.

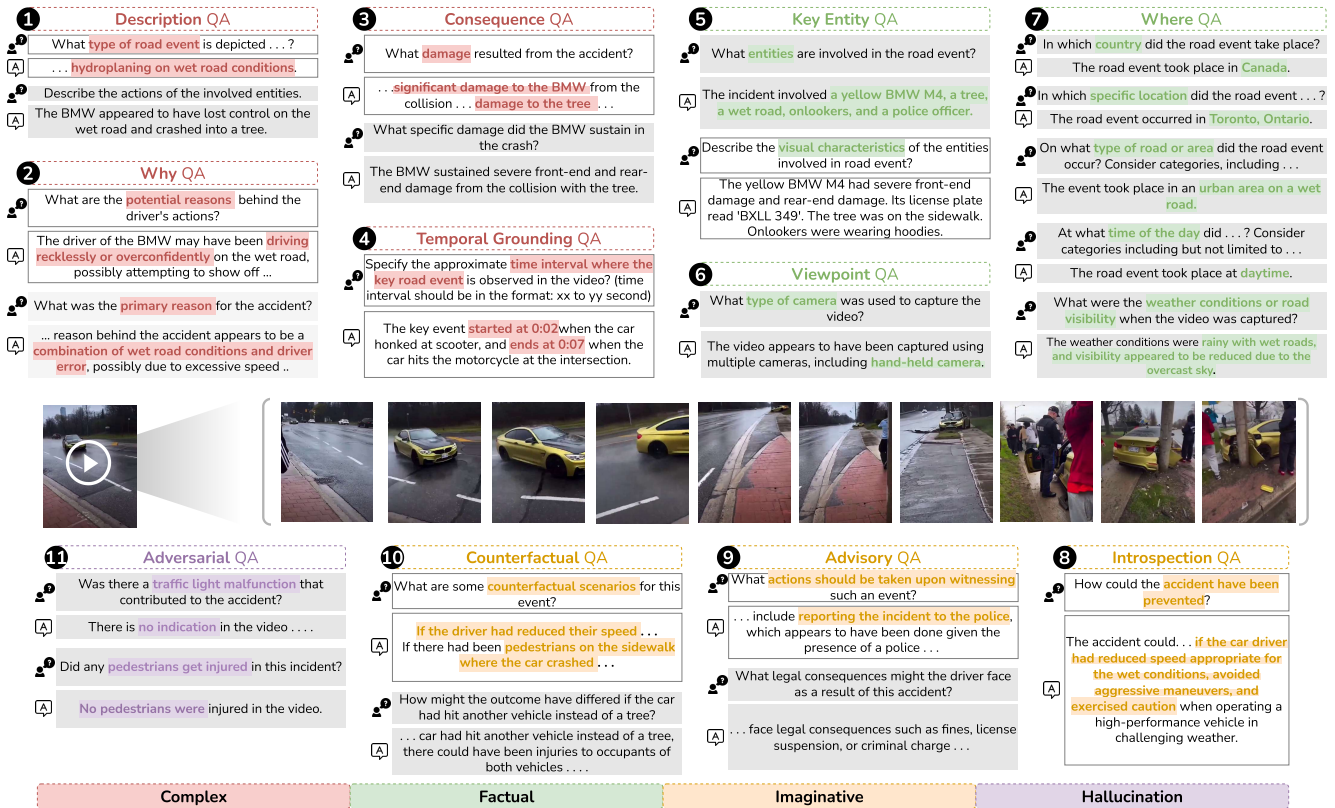


Figure 31. Examples of QA Pairs grouped by tasks and color-coded by task category (for an hydroplaning incident captured via handheld camera). Gray fill shading indicates specific questions while the non-shaded QAs are generic. Highlighted text indicates key information. Refer back to Appendix B.9.



Figure 32. Examples of QA Pairs grouped by tasks and color-coded by task category (for a near-miss incident captured in Japan). Gray fill shading indicates specific questions while the non-shaded QAs are generic. Highlighted text indicates key information. Refer back to Appendix B.9.

Tag Generation Script: For Each Question Template ID

```
if q_idx == 0:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the type of camera that is used to record the video. Do not specify
    the name of the camera model. Few examples are 'dashcam', 'vehicle-mounted camera', 'hand-held
    camera', 'cell-phone camera', 'cctv camera', 'surveillance camera', 'drone camera', 'satellite
    view'. If the information on type of camera is not available in the QA pair, respond with
    'unknown'. Ensure the response is strictly generated in the given format with no
    explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 1:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, describe the type of road event captured in the video. {ques_idx_1}. If the
    information about the road event is not available in the QA pair, respond with 'unknown'. Ensure
    the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 2:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the name of the country where this event is taking place. For
    example, answers could be 'USA', 'UK', 'India', 'Canada', 'Australia', etc. If the
    information about the country is not available in the QA pair, respond with 'unknown'. Ensure the
    response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 3:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the location where this event is taking place. Location could be
    the name of a state, district, landmark, type of locality, like city/town/village, etc. Some
    examples are, 'Big Ben London', 'Shanghai city', etc. If the information on location is not
    available in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in
    the given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}
```

elif `q_idx == 4:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify on which type of road or area, this event have taken place. For example, answers could be 'urban area', 'rural village area', 'highway', 'flyover', 'turn', 'intersection', 't-junction', 'roundabout', 'hilly or mountain area', etc. If the information is not available in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

elif `q_idx == 5:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify when did this road event happen. For example, 'morning', 'afternoon', 'evening', 'night', etc. Do not specify the exact date or time in the generated answer. If the information is not present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

elif `q_idx == 6:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify what was the weather conditions or road visibility. For example, 'sunny', 'rainy', 'windy', 'foggy', 'low visibility', etc. If the information is not present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

elif `q_idx == 7:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify all the road entities observed during the road event. A road entity can include road infrastructure objects like 'traffic signs', 'lane markings', 'barricades', etc. Road entities can also include road users like 'cars', 'two-wheelers', 'pedestrians', 'drivers', etc. If the information about road entity is not present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation:

{'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

```

elif q_idx == 8:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the visual characteristics of the road entities in the tag. Few
    examples of visual characteristics are information about 'what was the vehicle's color', 'was
    the headlight, brake light or turn signal on', 'what was the license plate number', 'was the
    rider wearing helmet or seat belt', etc. If the information about visual characteristics are not
    present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the
    given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 9:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the actions performed by the road entities in the tag. Few examples
    of actions are 'illegal overtaking', 'over speeding', 'swerving', 'yielding', 'cutting', etc. Few
    examples of what are not actions: 'frustrating', 'intimidating' (i.e., emotions or motive behind
    the actions should not be considered in generation of tags). If the information about the actions
    performed by road entities is not present in the QA pair, respond with 'unknown'. Ensure the
    response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 10:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the likely motives behind the actions of road entity, in the tag.
    Few examples of motives could be, 'thrill', 'road rage', 'impressing others', 'in a rush',
    'aggressive', 'impatient', 'drink and drive', etc. If the information about motives behind the
    action of the road entities is not present in the QA pair, respond with 'unknown'. Ensure the
    response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 11:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify the primary reason behind the accident or a near-miss collision, in
    the tag. Few examples of reasons could be, 'thrill', 'road rage', 'impressing others', 'in a
    rush', 'aggressive', 'impatient', 'drink and drive', etc. If the information about reasons is not
    present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the
    given format with no explanation: {'tags': ['', '', '']}
    ### QA pair ###
    Q: {question}
    A: {answer}

```


elif `q_idx == 12:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify how the accident or a near-miss collision could have been prevented, in the tag. Few examples of preventive measure could be, 'by slowing down at the intersection', 'checking the rearview mirror', etc. The information about preventive measure might not be present in the QA pair, in that case, you have to come up with a creative response on how could that incident be prevented. If you cannot determine the response, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

elif `q_idx == 13:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify all the casualties or road infrastructure damage during the event in the tag. Few examples are, 'people in the car died', 'bikers got injured', 'pedestrians got hit by car', 'divider was damaged', etc. If the information about casualties or damage is not present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

elif `q_idx == 14:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify all the traffic rule violations associated with the road event in the tag. Few examples are, 'illegal overtaking', 'illegal overtaking by crossing solid lane markings', 'hiding license plates', 'license plate not visible', 'helmet rule violation', 'wrong-side driving', 'triple riding violation', 'red light violation', 'drunk driving', etc. If the information about traffic rule violation is not present in the QA pair, respond with 'unknown'. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

elif `q_idx == 15:`

Go through the following 'QA pair', focus on extracting or generating tags as per the instructions.

Instructions

Based on the QA pair, specify what measures should be taken upon witnessing an unsafe driving road like this, in the tag. Few examples could be, 'reporting traffic violation to local government authorities or police', 'unsafe road infrastructure to local government authorities or police', 'fines', 'jail time', 'license ban', 'vehicle confiscation', etc. Think about the response and only include up to 5 relevant responses in the tag. Ensure the response is strictly generated in the given format with no explanation: {'tags': ['', '', '']}

QA pair

Q: {question}

A: {answer}

```

elif q_idx == 16:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify all the road safety advisory tags corresponding to the listed road
    entities. Think about the response and only include up to 5 relevant responses in the tag. Ensure
    the response is strictly generated in the given format with no explanation: {'tags': ['','','']}
    ### QA pair ###
    Q: {question}
    A: {answer}

elif q_idx == 17:
    Go through the following 'QA pair', focus on extracting or generating tags as per the
    instructions.
    ### Instructions ###
    Based on the QA pair, specify all the counterfactual reasoning tags related to different road
    events or driving situations that could have happened under different circumstances. Few examples
    of what counterfactual reasoning tags could look like 'the biker would have met an accident if
    the truck steered a little towards the right', 'the incident could have been worse if there were
    pedestrians by the roadside', 'If the car had not been speeding, it would have safely stopped
    before the intersection and avoided being hit by the truck'. Think about the response and only
    include up to 5 relevant responses in the tag. Ensure the response is strictly generated in the
    given format with no explanation: {'tags': ['','','']}
    ### QA pair ###
    Q: {question}
    A: {answer}

```

Figure 34. Tag extraction prompt design for different template question IDs. The prompt employs conditional logic based on question IDs to generate appropriate tags: camera type (q_idx=0), road event type (q_idx=1), country (q_idx=2), and specific location (q_idx=3). Each condition includes specific instructions and examples for tag generation, ensuring standardized output format 'tags': [",",",","]. Note: [ques_idx_1](#) is a command providing tag generation instructions for q_idx=1. This command can be found in [Fig. 35](#). Refer back to [Appendix B.10](#).

Some examples of traffic events and their corresponding Tags:

Tag: Safe driving education

For traffic events like: - drive around pothole guideline - front and back seatbelt usage - safe driving advisory around kids - safe driving around cyclists - truck moving safely on narrow road in dense traffic - biker and pillion wearing helmet - car not rushing lane change - car slowing down to ensure safety of biker in rear - advice about driving around big vehicles - safe driving on highways - safe driving around bus pick up and drop point - safe driving advice around Toll plaza - advice around tail light importance - traffic police regulating road safety - guideline about stopping at amber light before crossing the intersection - helmet wearing advisory - advise to pedestrian to walk on footpath - road infra advisory

Tag: Road Rage

For traffic events like: - riders using derogatory body language - absurd language by riders - cutting in and stops another vehicle's path and getting into argument - bus drivers fighting on the road - police and biker arguing - group of bikers hitting pedestrians - car driver and auto - Two wheeler and car rage - car drivers rage - group of bikers hitting pedestrians

Tag: Dangerous or Rash driving

For traffic events like: - using phone while riding - abrupt lane changes - zig zagging behavior - wheeling stunt - standing on bike while riding - rider switching seats with pillion passenger while riding - horizontal lane cutting - vehicle squeezing in between large vehicles - overtaking during turn - child holding bike steering or handle while adult is riding along - abrupt overtaking - pedestrian dangerous road crossing - person jaywalking in front of a car - reckless turning

Tag: Accident

For traffic events like: - car overturn - collision between biker and road infrastructure - collision between car and road infrastructure - bus skids due to hydroplaning - collision between car and auto - biker skids - truck and bike collision - biker colliding with road object - biker crashes into other biker - kid crashes into the car - bike and car collision - car and bus collision - car and car collision - Three wheeler and car collision - collision between bus and road infrastructure - car and pedestrians collision - collision between truck and road infrastructure - truck and car collision - farm truck and car collision

Tag: Near-miss

For traffic events like: - Near-miss between car and pedestrian - car and biker near-miss - Near-miss between car and car - biker and biker near-miss - Near-miss between biker and truck - car and cart near-miss

Tag: Faulty road infrastructure

For traffic events like: - traffic congestion on both lanes - bad design of traffic flow - construction site obstruction - traffic cone fallen in middle of road - advisory about faulty road infrastructure like no amber light only green or red on the intersection

Tag: Animal related accident

For traffic events like: - deer jumped onto a biker - bull on highway hit by car at night - monkey and three wheeler incident - puppies ran over by truck - dog and bike crash

Tag: Traffic Violations

For traffic events like: - illegal left or right turn - overtaking by crossing the solid yellow line - hiding the license plate numbers - wrong side driving - riding on footpath - helmet rule violation - triple riding violation - red light violation - illegal unlocking of load vehicle - parking vehicle on footpath

Tag: Post crash

For traffic events like: - crashed car - crashed bus - vehicle crashed into pole - overturned vehicle on slippery road - multiple vehicle collisions on slippery road

Tag: Defensive driving

For traffic events like: - safe following distance - avoiding distractions - speed control - lane discipline - yielding to pedestrians - proper signaling - following traffic laws - avoiding aggressive driving

Figure 35. Refer to Fig. 34 for details.

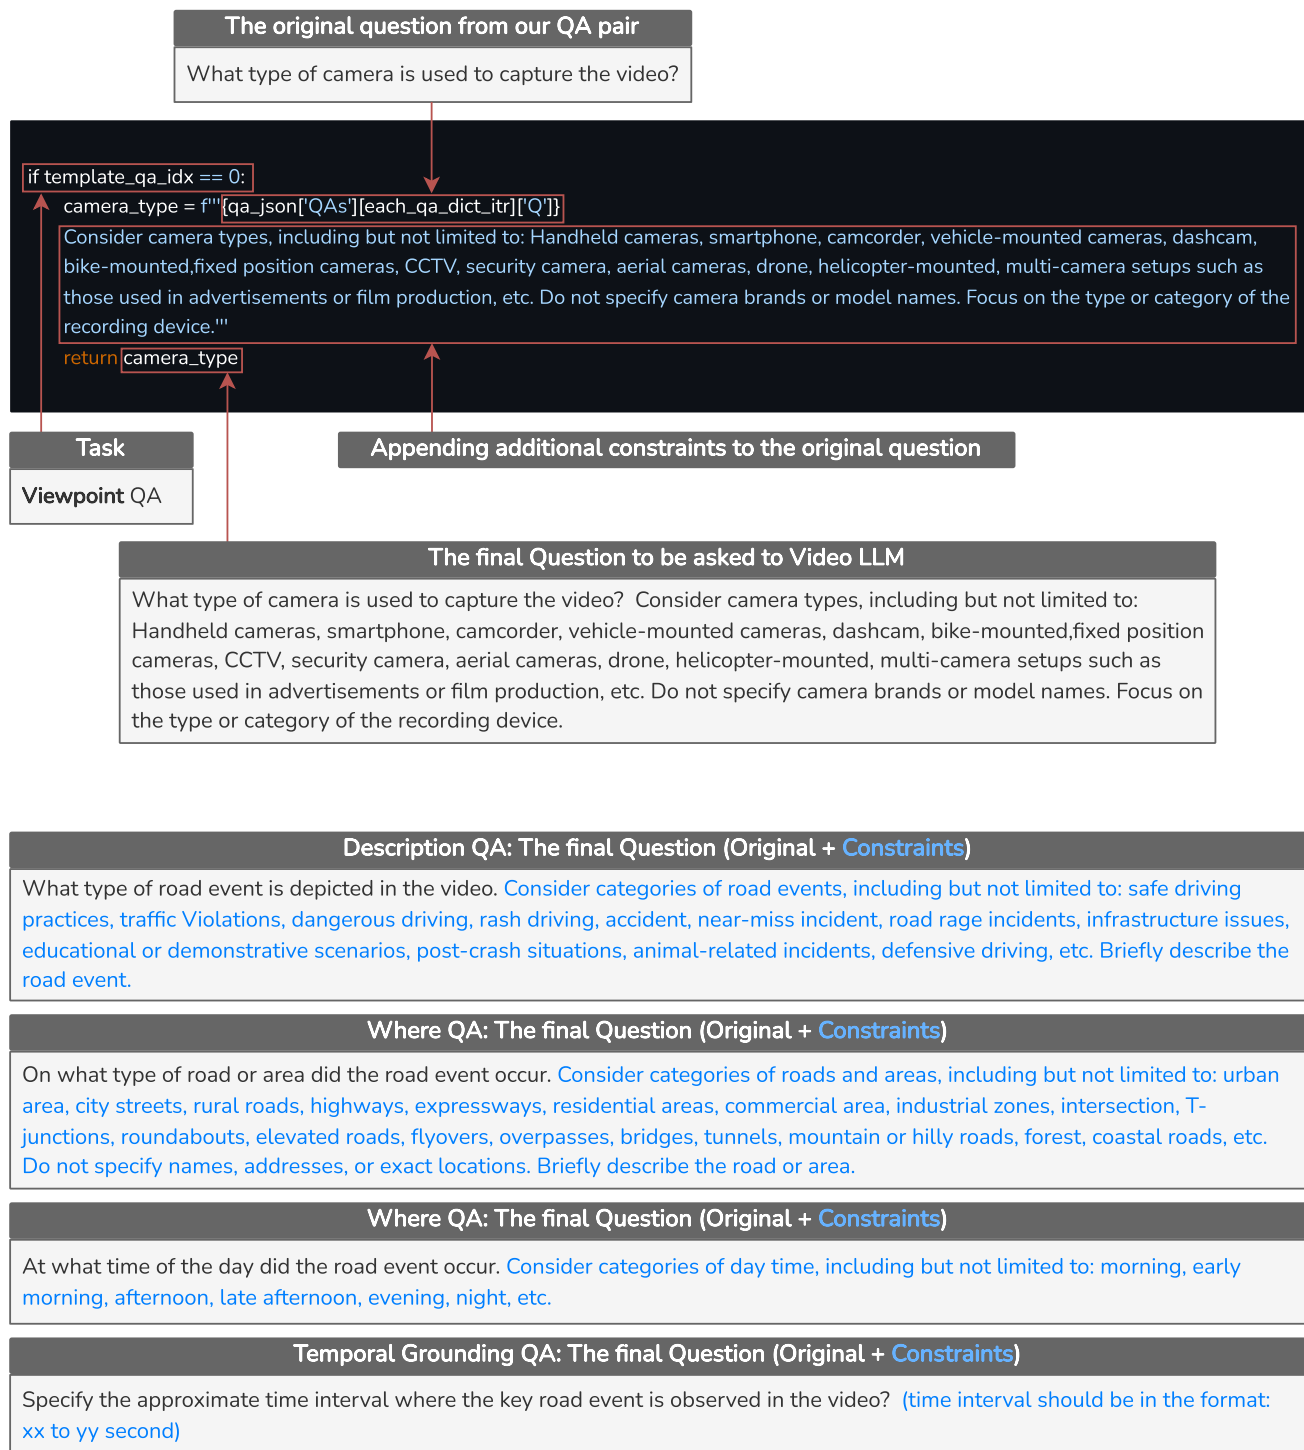


Figure 36. The diagram shows an example of task-specific prompt utilized for the evaluation of Video LLMs. The code snippet at the top demonstrates how this is done. First, for a specific question, we find its QA type via its template ID, then for that template ID, if we have a predefined constraint, we append that to the original question. Original question + constraint examples for Description QA, Why QA and Temporal Grounding QA tasks is shown. Rest of the tasks have only original questions and no predefined constraints. Refer back to Appendix C.1.

GPT-Eval Prompt

System_message:

You are an expert evaluator who rates the predicted answer based on the correct answer for a given question.

Prompt:

Evaluate the following question-answer pair:

Question: {question}

Correct Answer: {GT}

Predicted Answer: {answer}

Rate the Predicted Answer based on the Correct Answer on a scale from 0 to 100, with higher scores indicating that the Predicted Answer is closer to the Correct Answer. Your rating should be accurate to single digits like 62, 78, 41, etc.

Please generate the response in the form of a Python dictionary string with keys "score", where its value is in INTEGER, not STRING, and "explanation" giving reasoning behind the score.

For example, your response should look like this: {"score": 45, "explanation": "..."}

Figure 37. Evaluation prompt for assessing model-generated answers. The prompt implements (1) structured comparison between predicted and ground-truth answers, (2) fine-grained scoring on a 0-100 scale, and (3) requirement for explanatory justification. The output format ensures programmatic processing while maintaining evaluation transparency. Refer back to Appendix C.2.



Figure 38. **Model performance comparison on Temporal Grounding task:** Top: Frames from a video showing a car accident sequence. Middle: Models are asked to specify the temporal interval of the key road event. Ground truth (in gray) indicates the event spans 14-292 seconds. Bottom: Model responses (colored boxes) demonstrate varying approaches: while some attempt to provide specific intervals (e.g., 15-20 seconds, 0-3 seconds), others offer vague temporal descriptions. Red circles around model icons indicate that despite different response styles, all models fail to accurately identify the correct time interval. This example illustrates the significant challenge Video LLMs face in precise temporal localization of road events. Refer back to Appendix C.3.



Figure 39. **Model performance comparison on Temporal Grounding task:** Top: Sequential frames from a CCTV video showing a nighttime road scene. Middle: Models are asked to specify the temporal interval of the key road event, with ground truth spanning 35-40 seconds (gray box). Bottom: Model responses (colored boxes) demonstrate varying approaches: most provide specific time intervals (e.g., 23:01-23:11, 05:01-05:06) while Gemini additionally describes the event type ('person hit by motorcycle'). Red circles around model icons indicate that despite different response styles, all models fail to provide the correct interval. GPT-4o's response (57 to 42 seconds) even shows incorrect temporal ordering. This example highlights Video LLMs' consistent difficulty with precise temporal localization. Refer back to Appendix C.3.

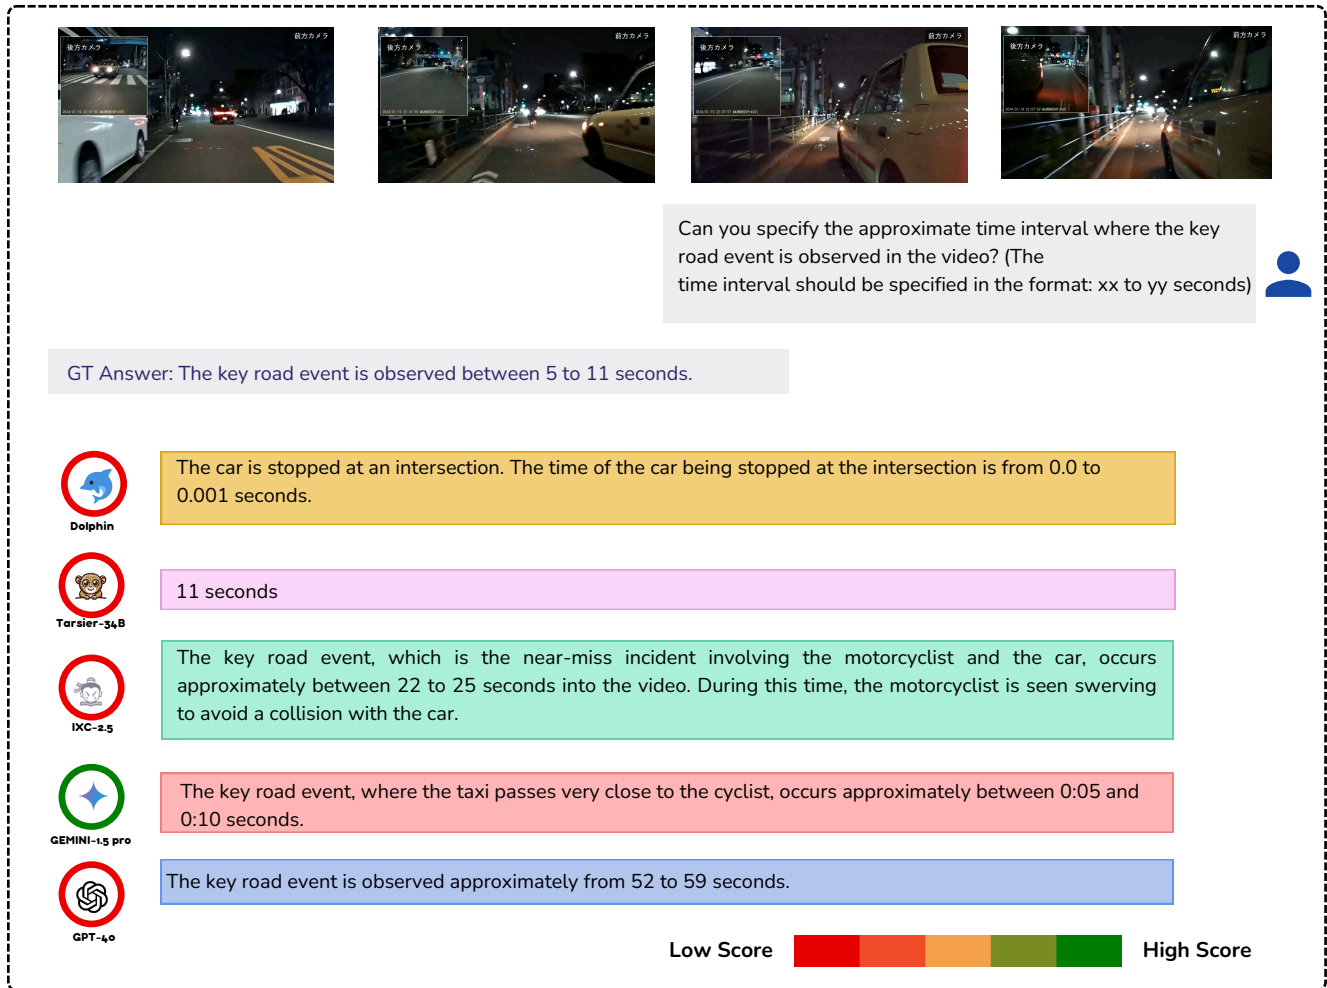


Figure 40. **Model performance comparison on Temporal Grounding task:** Top: Dashcam footage showing a nighttime near-miss incident between a taxi and cyclist. Middle: Models are asked to specify the temporal interval of the key road event, with ground truth spanning 5-11 seconds (gray box). Bottom: Model responses (colored boxes) show diverse approaches: while Gemini-1.5 Pro (green circle) correctly identifies both the event type and provides a reasonable time estimate (0:05-0:10), other models either give incorrect intervals (IXC: 22-25s, GPT-4o: 52-59s), overly precise timing (Dolphin: 0.0-0.001s), or incomplete responses (Tarsier: '11 seconds'). This example demonstrates that even when models accurately describe the event (taxi passing close to cyclist), precise temporal localization remains challenging, with only one model achieving high accuracy. Refer back to Appendix C.3.

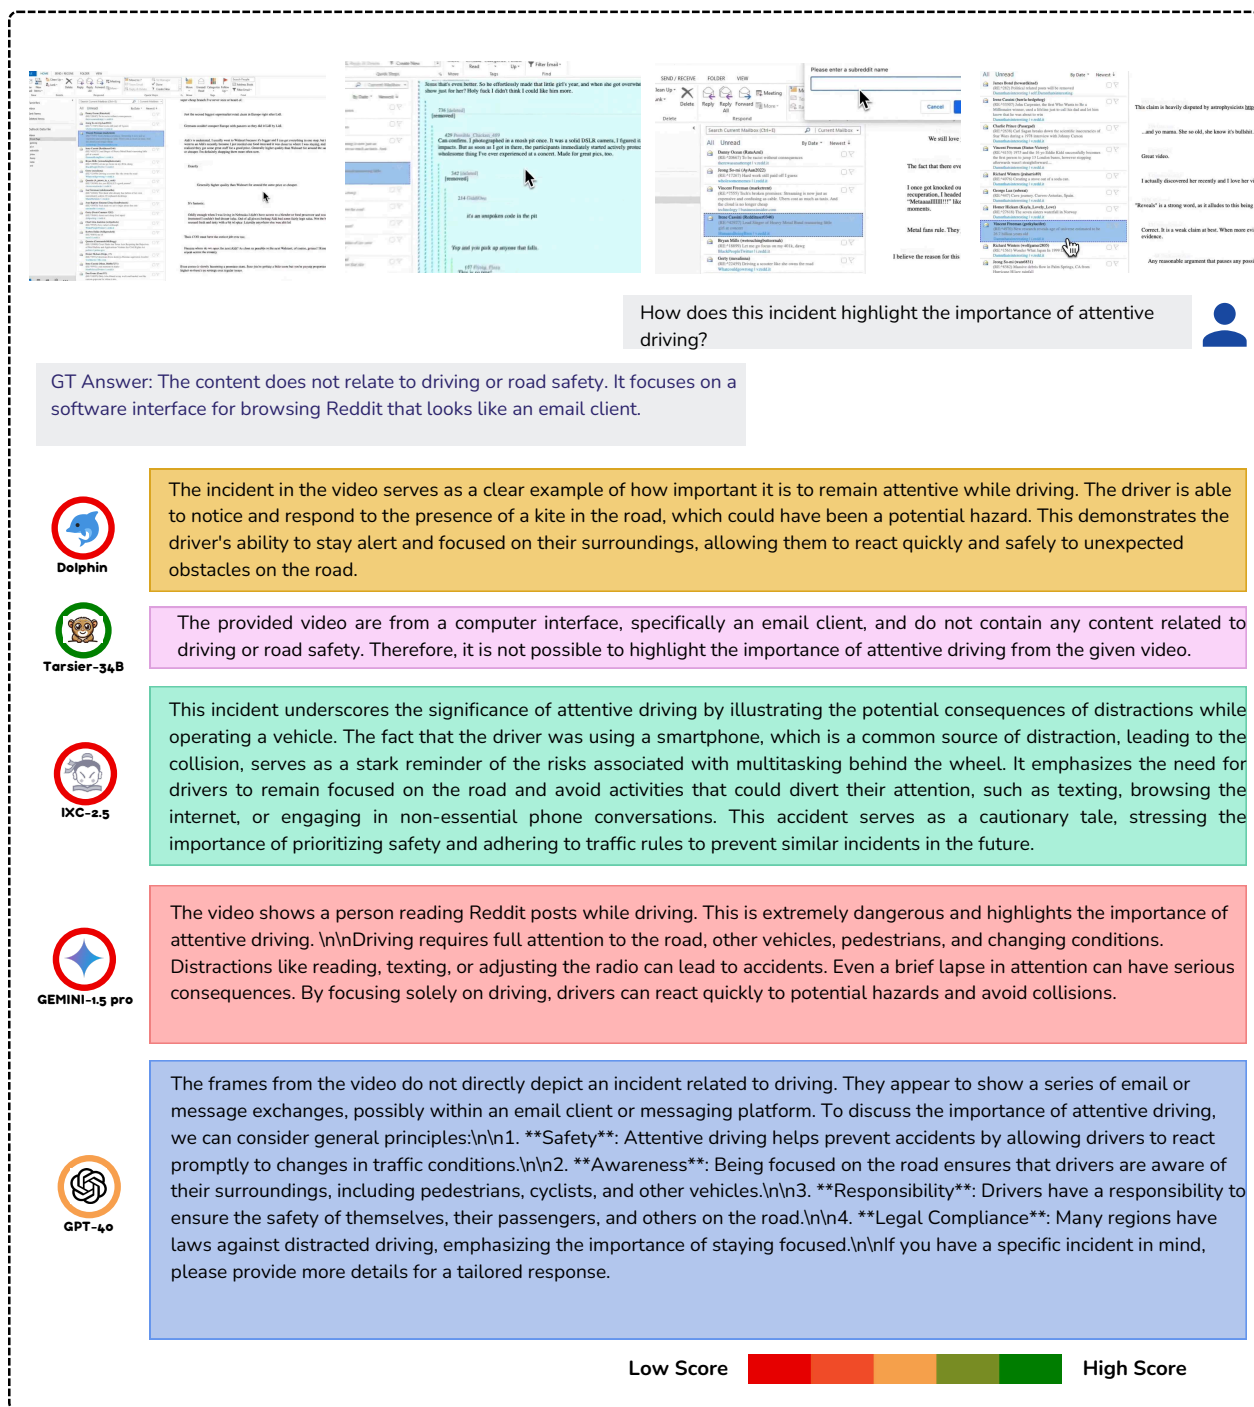


Figure 41. **Model performance comparison on Incompatible QA task**: Top: Video of a Reddit interface (non-road-event content). Middle: Models are asked about attentive driving implications, while the ground truth (gray box) correctly mentions that the content is unrelated to driving safety. Bottom: Model responses showcase varying levels of hallucination: most models (red circles) fabricate driving scenarios and safety implications despite the irrelevant content, while Tarsier (green) correctly identifies that the video is not related to road event. Although, GPT-4o (orange) correctly identifies the computer interface, it still attempts to relate it to driving. This example highlights a critical challenge in Video LLM robustness - the tendency to hallucinate road safety contexts even when presented with completely unrelated visual content. Refer back to Appendix C.3.



Figure 42. **Model performance comparison on Incompatible QA task:** Top: Frames showing indoor welding activity in a workshop. Middle: Models are asked about roadside maintenance accident prevention, while the ground truth (gray box) correctly indicates that the content shows indoor welding, not roadside maintenance. Bottom: Model responses (colored boxes) demonstrate varying degrees of hallucination: while Tarsier (green) correctly acknowledges insufficient information to discuss roadside maintenance, Dolphin and IXC (red circles) fabricate elaborate safety measures despite the obvious indoor setting. Gemini and GPT-4o's (dark orange) detailed response about welding safety, while technically accurate, still fails to address the fundamental context mismatch. This example illustrates how models can generate plausible but irrelevant safety recommendations when presented with visually similar but contextually different scenarios. Refer back to Appendix C.3.

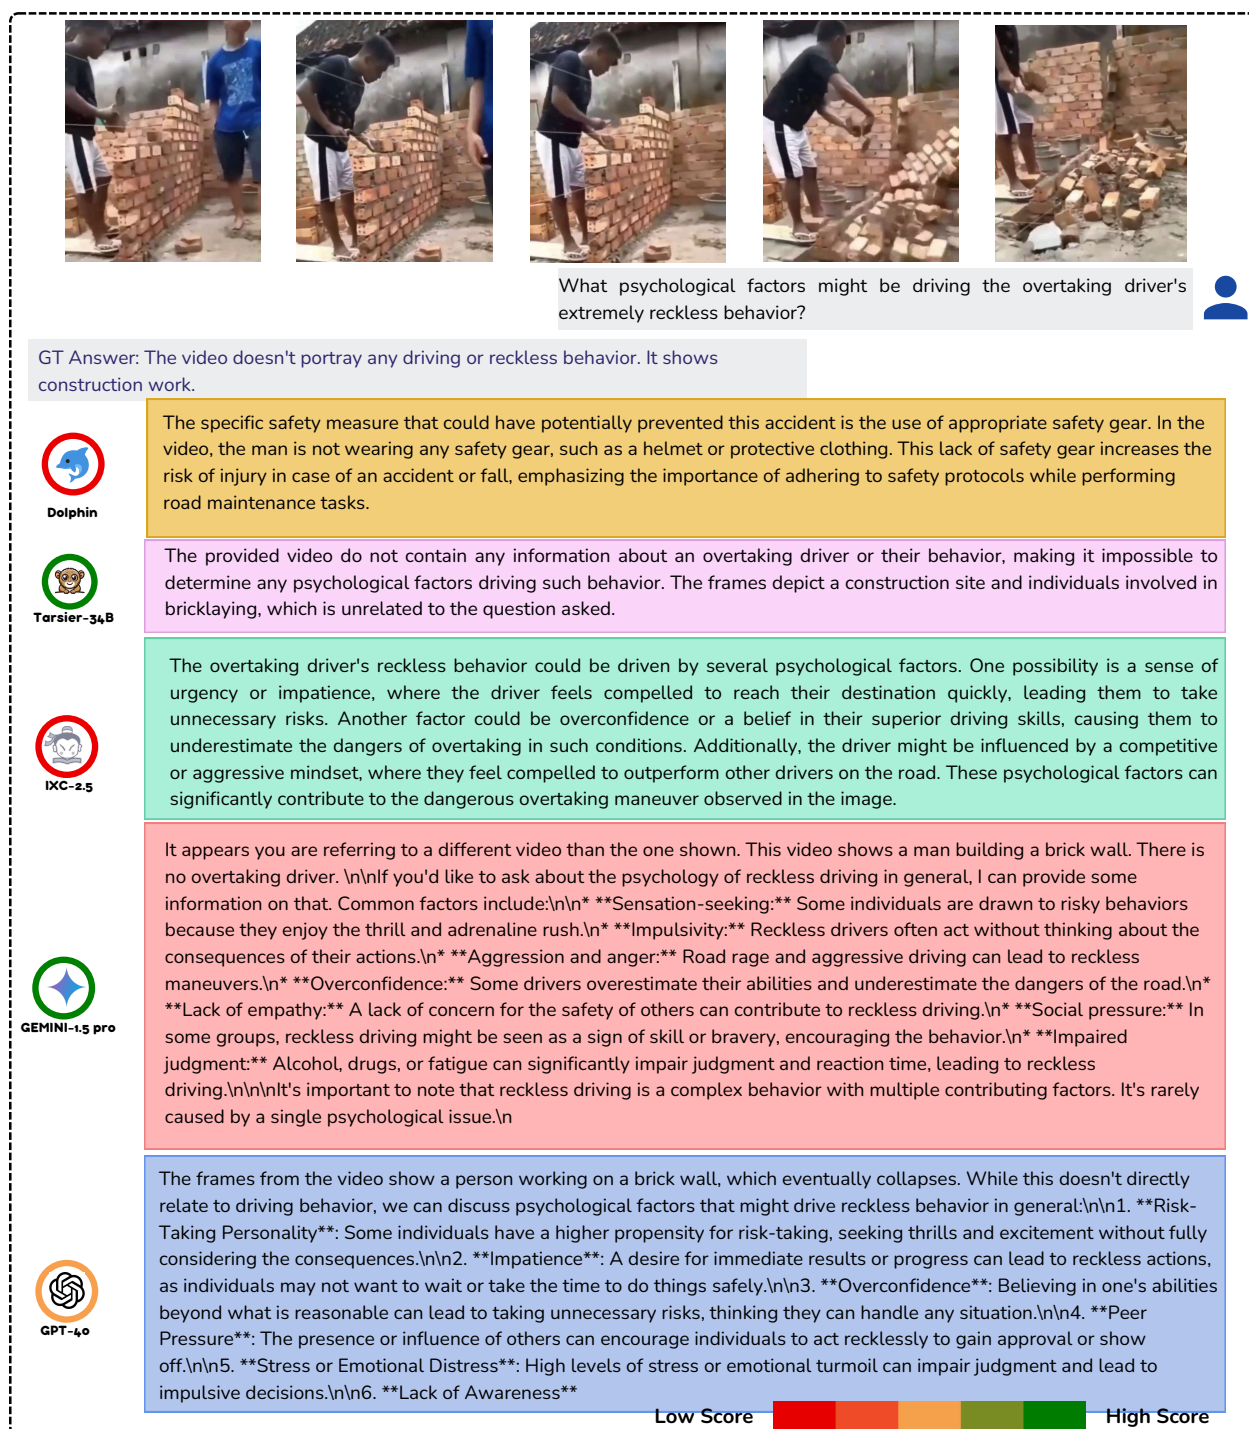


Figure 43. **Model performance comparison on Incompatible QA task:** Top: Frames showing a person building and subsequently demolishing a brick wall. Middle: Models are asked about psychological factors behind overtaking behavior, while the ground truth (gray box) correctly mentions this as unrelated to overtaking. Bottom: Model responses show varying levels of hallucination and context confusion: Dolphin and IXC (red circles) completely ignore the brick wall context and fabricate scenarios about road safety, while Tarsier and Gemini (green) correctly acknowledges the construction setting and clearly states the content mismatch. This example demonstrates how models can struggle with maintaining contextual accuracy, with some generating elaborate but irrelevant psychological analyses despite clearly unrelated visual content. Refer back to Appendix C.3.

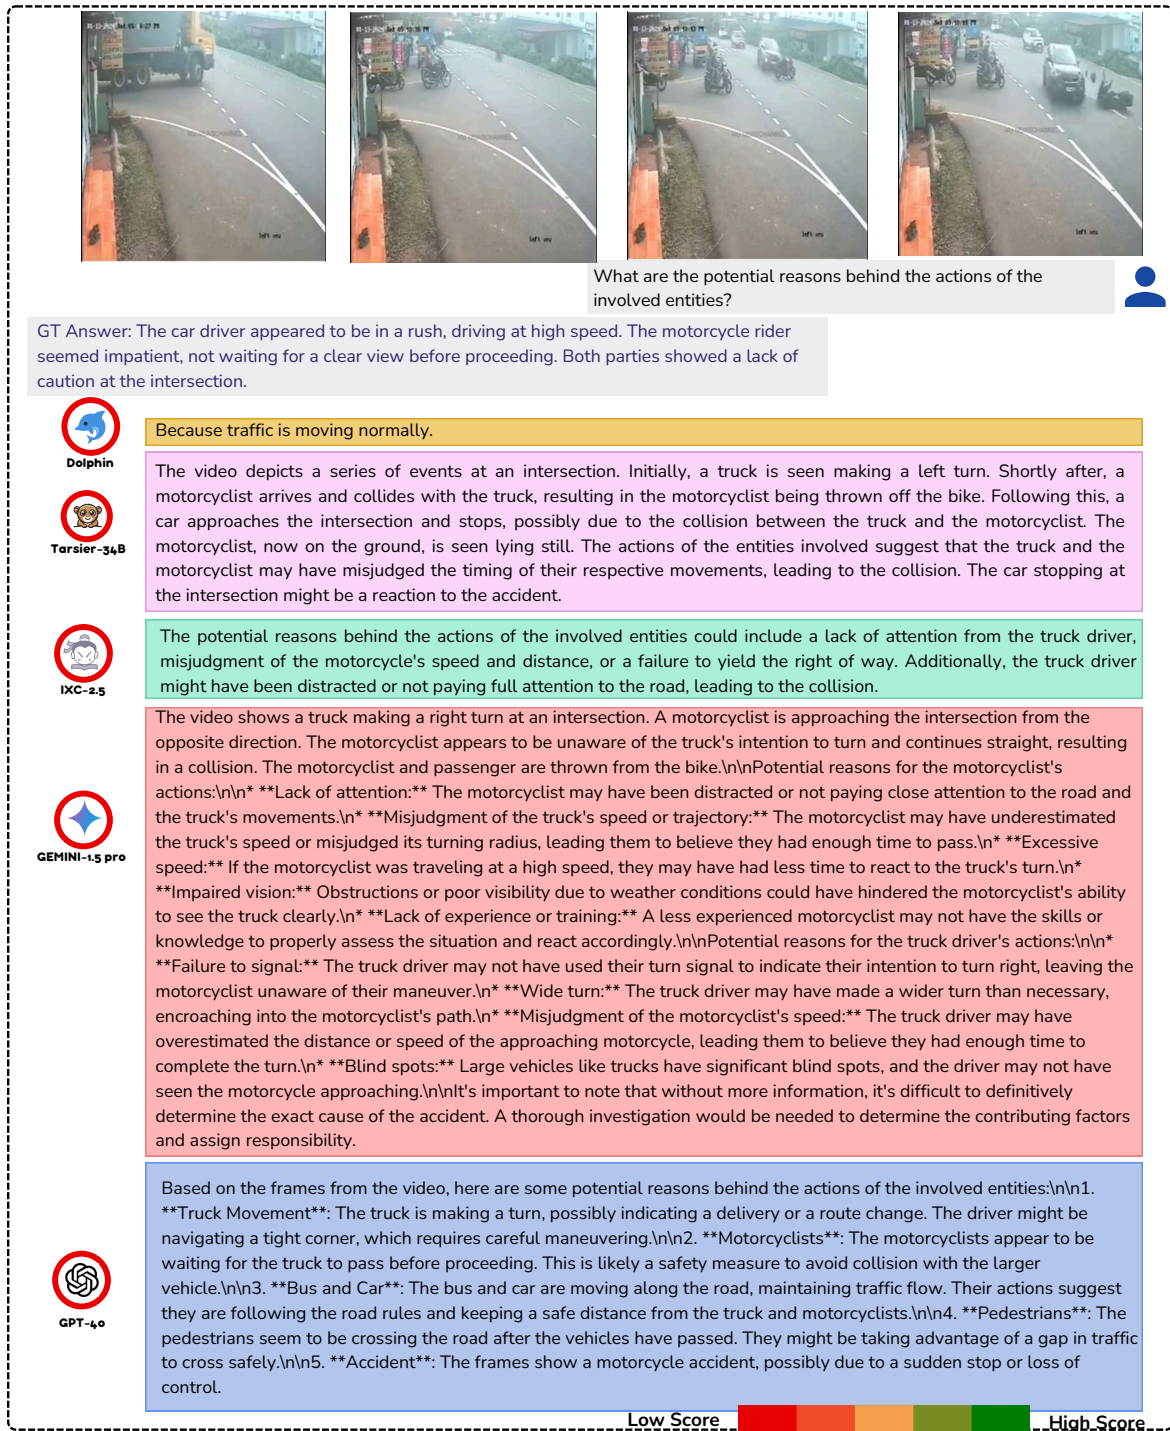


Figure 44. **Model performance comparison on Why QA task**: Top: CCTV footage showing an intersection incident between a truck and motorcyclist. Middle: Models are asked about potential reasons behind the road entities' actions, with ground truth (gray box) indicating rush and lack of caution as primary factors. Bottom: Model responses demonstrate varying levels of reasoning and detail: While Dolphin (red circle) provides an oversimplified response ('Because traffic moving normally'), other models offer increasingly complex analyses. Gemini generates a comprehensive analysis considering multiple factors (weather conditions, road visibility, driver attention), while GPT-4o provides a structured but possibly over-analyzed response with enumerated factors. This example illustrates the challenge of providing appropriate depth in causal reasoning without over-speculation. Refer back to Appendix C.3.

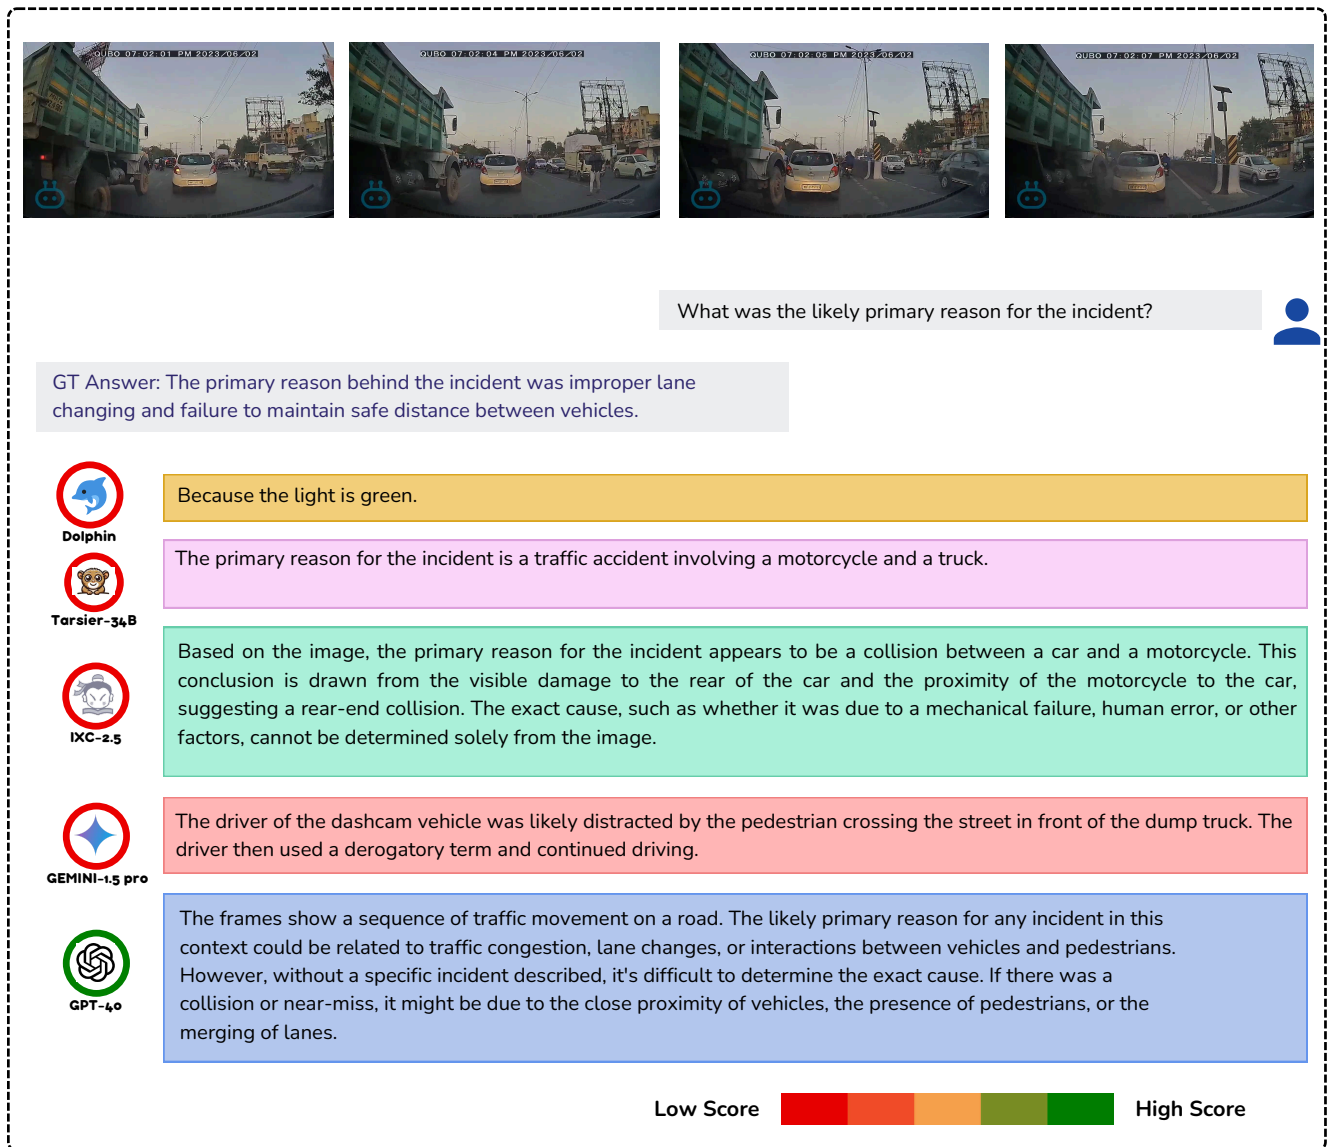


Figure 45. **Model performance comparison on Why QA task:** Top: Dashcam footage showing a traffic scenario with lane changing incidents. Middle: Models are asked about the primary reason for the incident, with ground truth (gray box) identifying improper lane changing and insufficient vehicle distance as key factors. Bottom: Model responses show varying levels of analytical accuracy and specificity: Dolphin offers an oversimplified and irrelevant response ('Because the light is green'), while Tarsier-34B provides a vague description without specific reasoning. IXC-2.5 attempts causal analysis but misidentifies the vehicles involved, and Gemini-1.5 Pro introduces unobserved elements (pedestrian crossing). GPT-4o demonstrates appropriate caution by acknowledging the difficulty in determining exact causes without clearer incident details. This example highlights the challenges in balancing between definitive causal analysis and appropriate uncertainty when visual evidence is ambiguous. Refer back to Appendix C.3.

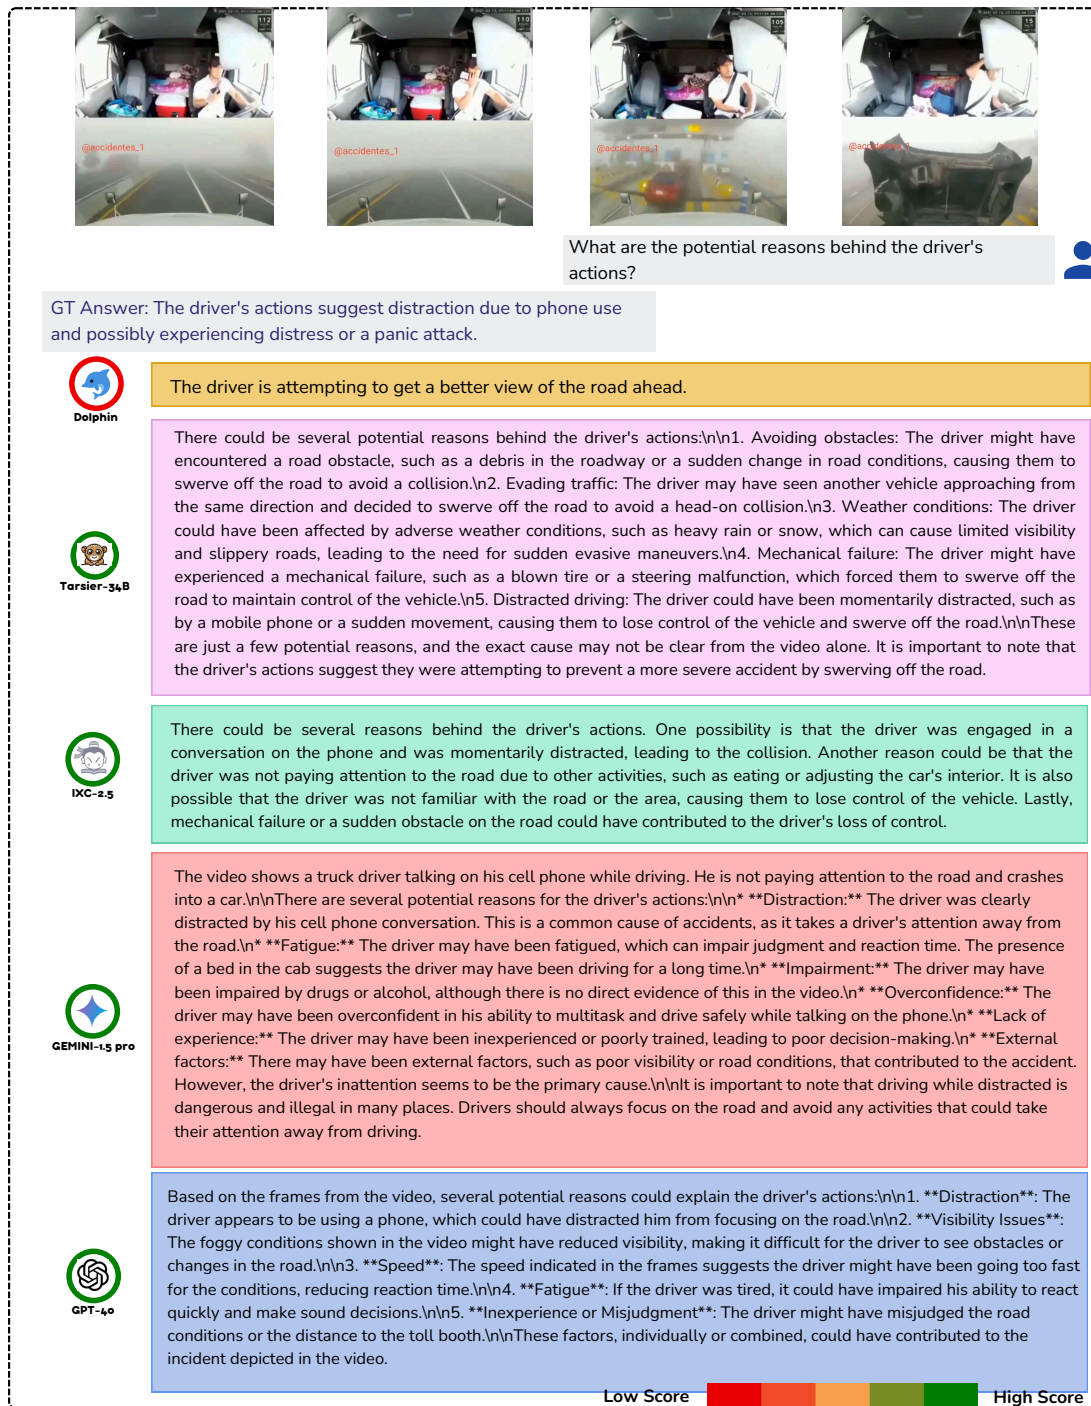


Figure 46. **Model performance comparison on Why QA task:** Top: Split-screen dashcam footage showing both driver behavior (phone use) and road view leading to an incident. Middle: Models are asked about potential reasons behind the driver's actions, with ground truth (gray box) identifying distraction from phone use and possible panic/distress. Bottom: Model responses demonstrate varying depths of causal analysis: Dolphin provides an oversimplified response about road view, while Gemini-1.5 Pro offers a comprehensive multi-factor analysis incorporating both observed behaviors (phone distraction) and possible underlying causes. IXC-2.5 stays focused on direct observables, while GPT-4V extensively analyzes multiple scenarios but maintains grounding in the visible evidence (phone conversation). This example shows how models balance between observable evidence (phone use) and inferring potential psychological states, with varying success in maintaining relevance to the visual content. Refer back to Appendix C.3.

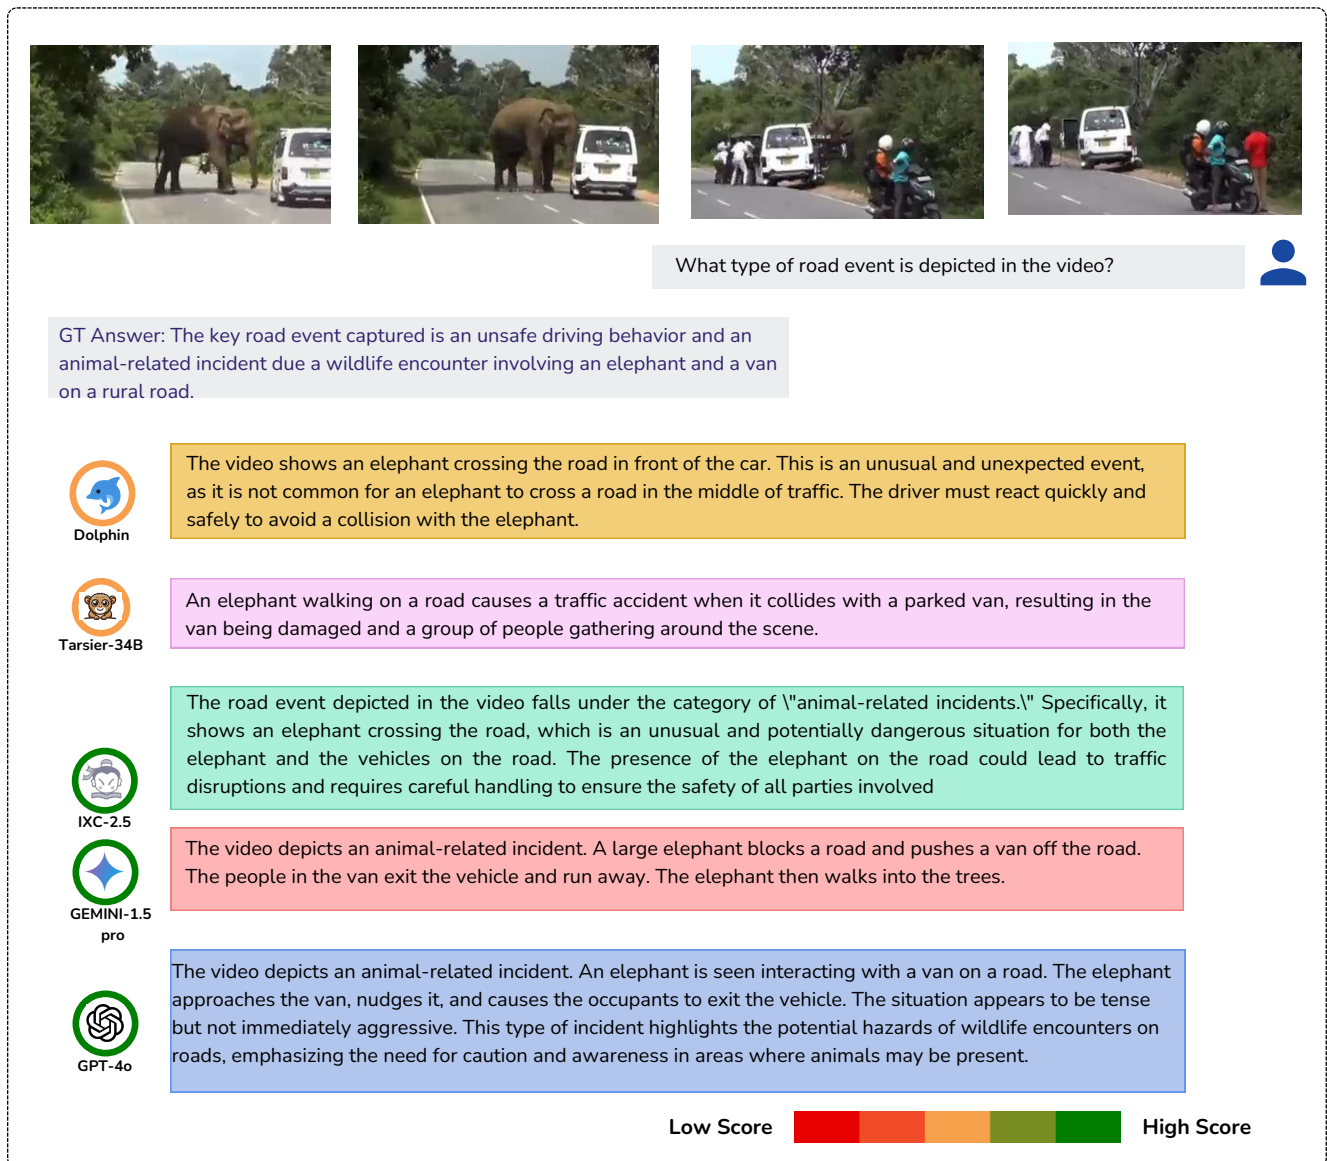


Figure 47. **Model performance comparison on Description QA task:** Top: Video frames showing an encounter between an elephant and a van on a rural road. Middle: Models are asked to describe the type of road event, with ground truth (gray box) identifying it as both unsafe driving behavior and an animal-related wildlife encounter. Bottom: Model responses show varying accuracy in event categorization and detail: Tarsier-3LB incorrectly describes a collision, while IXC-2.5 (green circle) provides a well-balanced response that correctly categorizes the event as 'animal-related incident' while acknowledging the safety implications for all parties. Gemini-1.5 Pro and GPT-4o offer accurate but differently focused descriptions, with Gemini emphasizing the physical interaction and GPT-4o highlighting the broader safety context. This example demonstrates models' varying abilities to balance between event classification, factual description, and safety implications in unusual road scenarios. Refer back to Appendix C.3.



Figure 48. **Model performance comparison on Description QA task:** Top: Video frames show a collision between a car and a bike on a curvy road. Middle: Models are asked to describe the actions of the entities involved in the road event, with ground truth (gray box) identifying it near-miss incident that further led to the collision. Bottom: Model responses show varying accuracy in event categorization and detail: All the models fail to answer this question due to incorrect identifications. GPT-4o fails to identify the motorcycle that was initially overtaking the auto that crashed a car. Tarsier-34B incorrectly identifies overtaking between the car and the truck. Refer back to Appendix C.3.



Figure 49. **Model performance comparison on Description QA task:** Top: Video frames show a traffic violation involving. Middle: Models are asked to describe the type of road event depicted in the video, with ground truth (gray box) identifying it as a vehicle driving the wrong way on a one-way street. Bottom: All models fail to recognize the violation. Refer back to Appendix C.3.



Figure 50. **Model performance comparison on Description QA task:** Top: Video frames showing a road safety awareness video aimed towards pedestrians or motorcyclists at night. Middle: Models are asked to describe the theme of the video, with ground truth (gray box) indicating that it is a safety awareness video. Bottom: Model responses show varying accuracy in event categorization and detail: All the models except Dolphin successfully capture the global context or theme of the video. Refer back to Appendix C.3.



Generic Q

Can you describe the actions of the road entities involved in the event?

GT Answer: A cyclist is waiting in the middle of the road, motorists are moving from all directions, and pedestrians are navigating the unsafe infrastructure.



GPT-4o

The frames from the video show a road scene with several entities:\n\n1. **Blue Bus**: The bus is moving straight along the road under the elevated structure. It maintains a consistent position in the lane, indicating it is traveling in a straight path.\n\n2. **Motorcyclists**: There are several motorcyclists visible in the frames. They appear to be moving in the same direction as the bus, with some overtaking it on the right side.\n\n3. **Other Vehicles**: There are a few other vehicles, including a yellow truck, visible in the background. They seem to be moving in the same direction as the bus and motorcycles.\n\n4. **Pedestrians and Bicyclists**: There are no clear indications of pedestrians or bicyclists crossing the road in these frames.\n\nOverall, the road entities are moving in an orderly fashion, with the bus and motorcycles traveling in the same direction under the elevated structure.

Specific Q

How does the lack of proper infrastructure affect different road users at this junction?

GT Answer: Cyclists are forced to wait in the middle of the road, motorists move from all directions without guidance, and pedestrians struggle to navigate safely. The absence of proper infrastructure increases the risk of accidents for all road users.



GPT-4o

The frames from the video show a busy junction with various road users, including buses, motorcycles, and pedestrians. The lack of proper infrastructure can affect them in several ways:\n\n1. **Safety Risks**: Without clear lane markings or traffic signals, there is a higher risk of accidents. Motorcyclists and pedestrians are particularly vulnerable.\n\n2. **Traffic Congestion**: The absence of designated lanes or traffic management can lead to congestion, as vehicles may not follow a structured flow.\n\n3. **Pedestrian Challenges**: Pedestrians may find it difficult to navigate safely without crosswalks or pedestrian signals, increasing the risk of accidents.\n\n4. **Inefficient Traffic Flow**: Without proper infrastructure, vehicles may not move efficiently, leading to delays and increased travel time.\n\n5. **Increased Stress**: Drivers and riders may experience higher stress levels due to the chaotic and unpredictable nature of the traffic.\n\nImproving infrastructure with clear signage, lane markings, and pedestrian pathways can enhance safety and efficiency.

Figure 51. The image shows a qualitative analysis of the performance of GPT-4o Video LLM for two types of questions - a generic question about the actions of the road entities, and a specific question about how the lack of proper infrastructure affects different road users. The ground truth (GT) answers are provided, and the predicted answers by the model are shown using icons - a red circle indicates the model's prediction does not align well with the ground truth, while a green icon indicates the model performs well. GPT-4o seems to be performing well in specific questions than generic one. This performance gap could be because generic questions require the model to infer the context while specific questions directly reference the event and entities, making it easier for models to answer them.



Generic Q

Can you describe the actions of the road entities involved in the event?



GT Answer: A cyclist is waiting in the middle of the road, motorists are moving from all directions, and pedestrians are navigating the unsafe infrastructure.



Sure! Here's a description of the actions of the road entities in the video:\n\n**Camera Vehicle:** The vehicle with the dashcam remains stationary throughout the video, likely at an intersection or stoplight.\n\n**Bus:** A blue bus is seen driving straight ahead in the center of the road.\n\n**Trucks:** Two trucks, one yellow and one white, are seen driving in the same direction as the bus, behind it.\n\n**Motorcycles/Scooters:** Several motorcycles and scooters are seen driving in various directions, some passing the camera vehicle, some going straight, and some turning.\n\n**Bicycles:** Two bicyclists are seen riding on the right side of the road, going in the same direction as the bus and trucks.\n\n**Auto-rickshaw:** An auto-rickshaw is briefly seen on the right side of the road, near the end of the video.\n\n**Pedestrian:** One pedestrian is seen walking on the left sidewalk near the beginning of the video.\n\nIt appears to be a typical traffic scene at an intersection, with various vehicles and a few pedestrians navigating the road.

Specific Q

How does the lack of proper infrastructure affect different road users at this junction?



GT Answer: Cyclists are forced to wait in the middle of the road, motorists move from all directions without guidance, and pedestrians struggle to navigate safely. The absence of proper infrastructure increases the risk of accidents for all road users.



The lack of proper infrastructure at this junction affects different road users in various ways:\n\n**Pedestrians:** There are no designated crosswalks or pedestrian signals, forcing pedestrians to navigate the intersection amidst moving traffic, increasing the risk of accidents. The uneven and broken pavement further adds to the difficulty and danger for pedestrians.\n\n**Cyclists:** Cyclists also lack dedicated lanes or paths, forcing them to share the road with larger vehicles, including buses and trucks. This increases their vulnerability to collisions, especially given the lack of clear traffic signals or right-of-way rules.\n\n**Motorcyclists:** Similar to cyclists, motorcyclists face the same dangers of sharing the road with larger vehicles without dedicated lanes. The uneven road surface and lack of clear traffic management also pose risks to their safety.\n\n**Drivers of larger vehicles (buses, trucks):** The lack of proper lane markings and traffic signals can lead to confusion and potential conflicts between larger vehicles, especially during turns or lane changes. The presence of other road users, such as pedestrians, cyclists, and motorcyclists, further complicates navigation and increases the risk of accidents.\n\n**Drivers of smaller vehicles (cars, auto-rickshaws):** Drivers of smaller vehicles also face challenges due to the lack of clear traffic management and the presence of various road users sharing the same space. The uneven road surface and potential for congestion can lead to delays and increase the risk of minor collisions.\n\nOverall, the lack of proper infrastructure at this junction creates a chaotic and unsafe environment for all road users, increasing the risk of accidents and hindering smooth traffic flow. The absence of designated lanes, crosswalks, traffic signals, and proper pavement contributes to the overall hazardous conditions.

Low Score  High Score

Figure 52. A similar phenomena between the gap between generic and specific QAs is reflected in Gemini, as seen in the previous example.

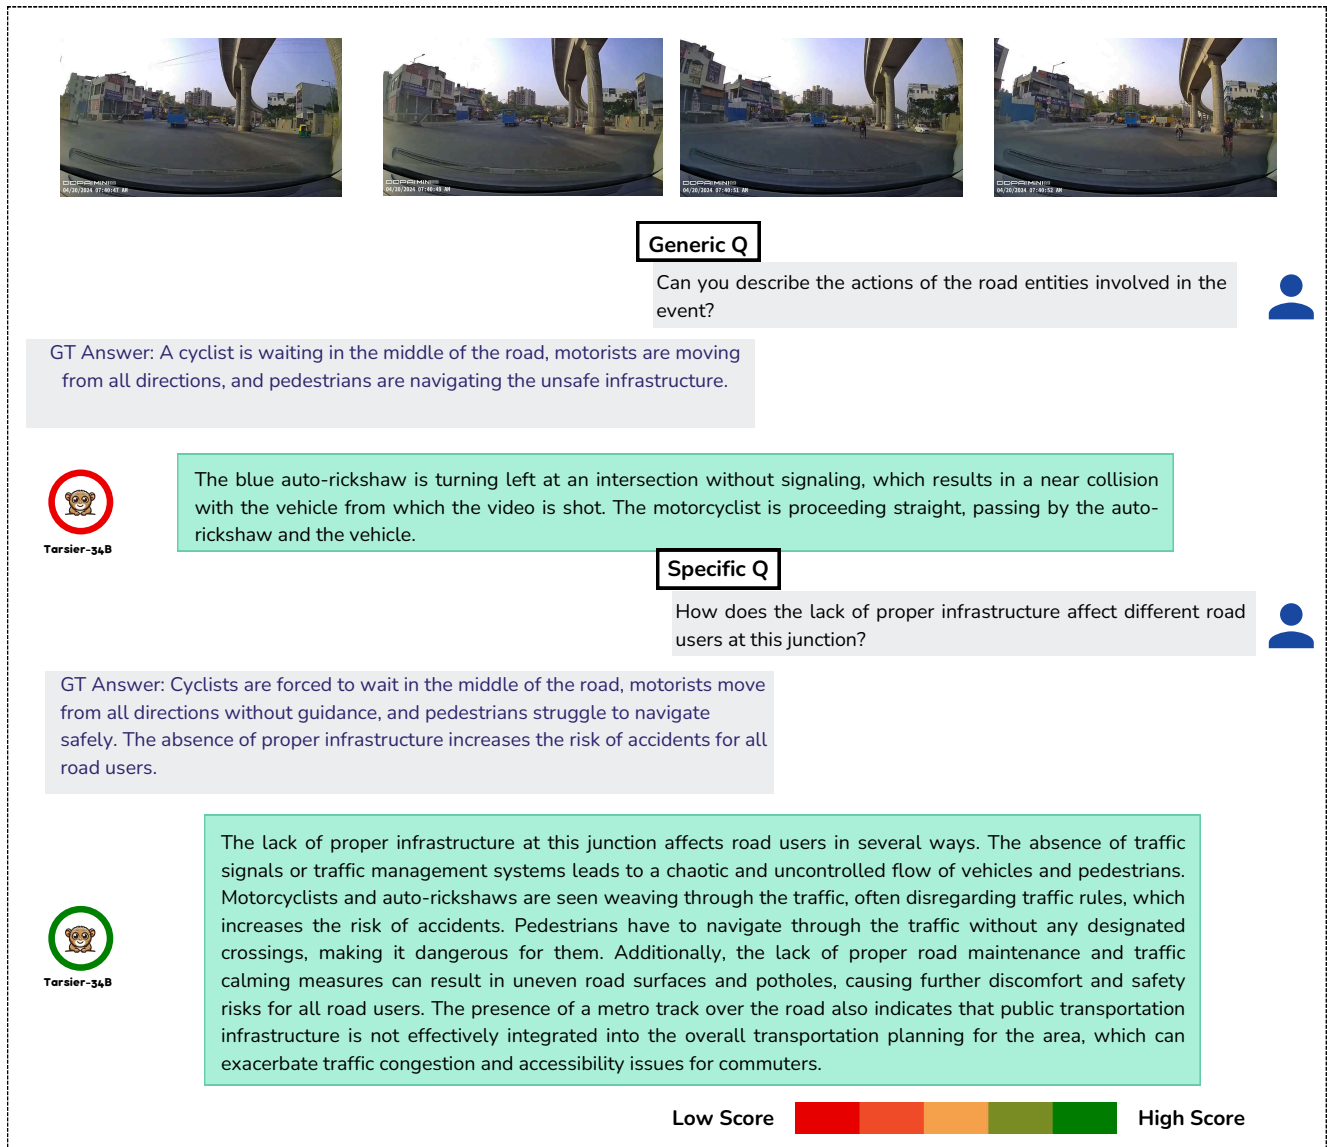


Figure 53. A similar phenomena indicating the gap between generic and specific QAs is reflected in Tarsier, similar to what was observed in the Gemini and GPT-4o in previous examples.



Generic Q

Can you describe the actions of the road entities involved in the event?



GT Answer: A cyclist is waiting in the middle of the road, motorists are moving from all directions, and pedestrians are navigating the unsafe infrastructure.



Based on the image, it appears that the car from which the photo was taken is stationary or moving slowly, as there is no significant motion blur. The motorcycle is in motion, as indicated by the position of the rider and the orientation of the bike. The bus is also in motion, traveling in the opposite direction of the motorcycle. The pedestrians seem to be going about their business, with some walking on the sidewalk and others possibly waiting to cross the street.

Specific Q

How does the lack of proper infrastructure affect different road users at this junction?



GT Answer: Cyclists are forced to wait in the middle of the road, motorists move from all directions without guidance, and pedestrians struggle to navigate safely. The absence of proper infrastructure increases the risk of accidents for all road users.



The lack of proper infrastructure at this junction affects different road users in several ways. For pedestrians, the absence of sidewalks or pedestrian crossings makes it dangerous to navigate the area on foot, as they are forced to walk on the road, increasing the risk of accidents with moving vehicles. For cyclists, the lack of dedicated bike lanes can lead to conflicts with motor vehicles, as seen in the video where the motorcyclist had to swerve to avoid a collision with a car. This can result in dangerous maneuvers and a higher likelihood of accidents. For motorists, the lack of clear traffic markings and signals can lead to confusion and difficulty in navigating the junction, increasing the risk of collisions. Overall, the lack of proper infrastructure at this junction creates a hazardous environment for all road users, leading to increased risks of accidents and injuries.

Low Score High Score

Figure 54. The phenomena of the model performing better in specific QAs than their generic counterparts persist in IXC as well.

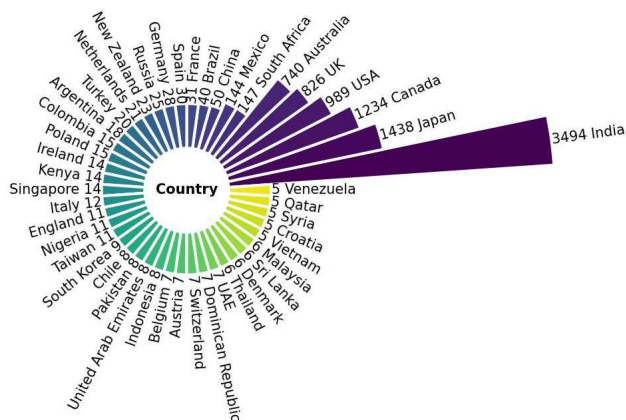


Figure 55. Geographical location (country of origin) distribution of video tags. Tags with fewer than five videos are omitted from the radar plot for clarity and to reduce clutter.

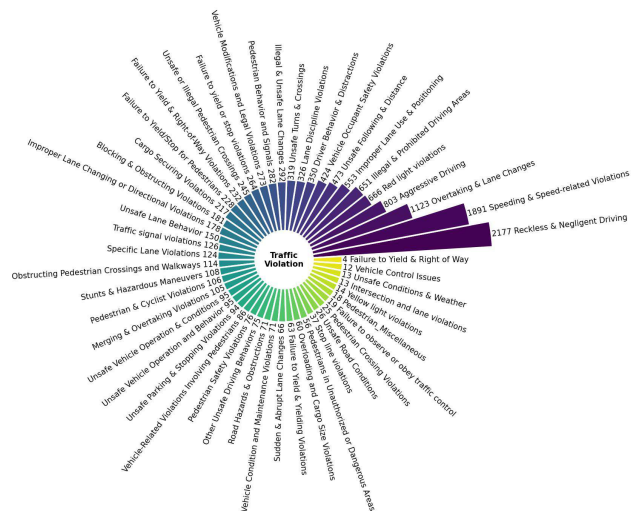


Figure 57. Traffic Violation Video Tags distribution. Tags with fewer than four videos are omitted from the radar plot for clarity and to reduce clutter.

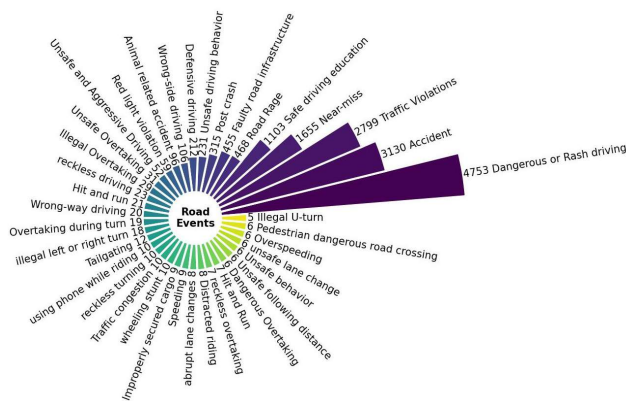


Figure 56. Road Event Video Tags distribution. Tags with fewer than five videos are omitted from the radar plot for clarity and to reduce clutter.

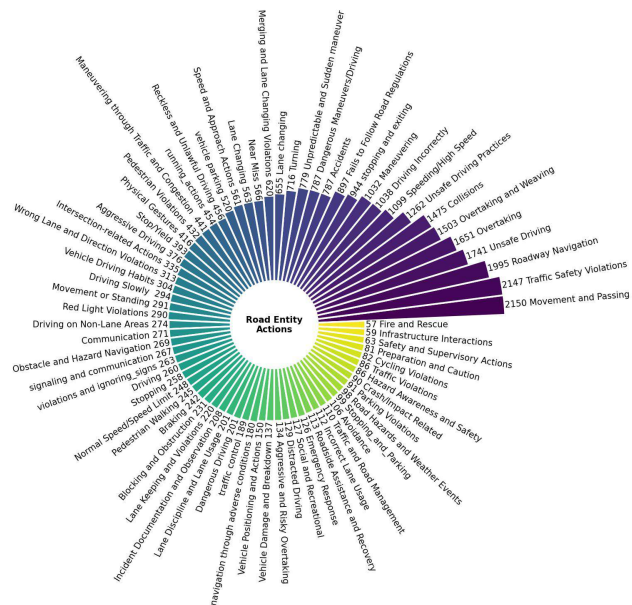


Figure 58. Road Entity Action Video Tags distribution. Tags with fewer than 57 videos are omitted from the radar plot for clarity and to reduce clutter.

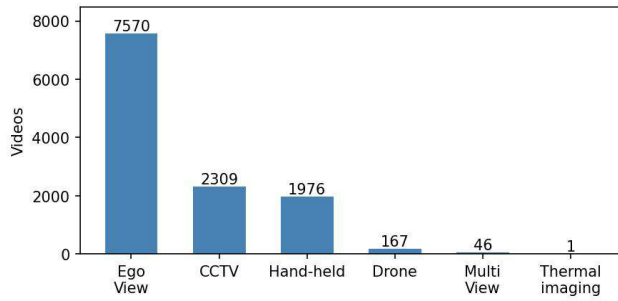


Figure 59. Viewpoint Video Tags distribution.

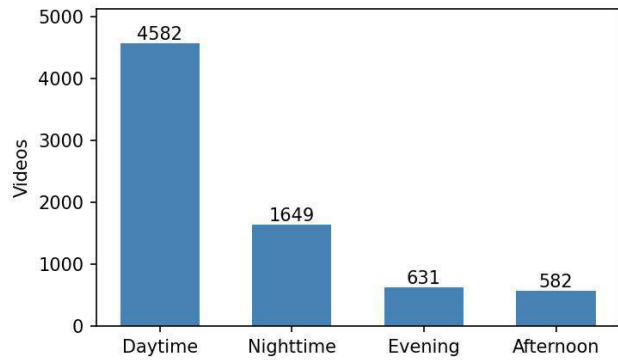


Figure 60. Time of Day Tags distribution.

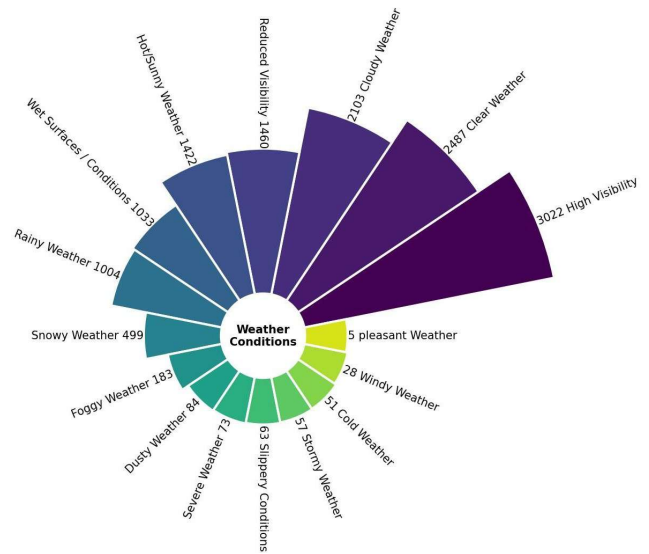


Figure 62. Weather Condition Video Tags distribution. Tags with fewer than five videos are omitted from the radar plot for clarity and to reduce clutter.

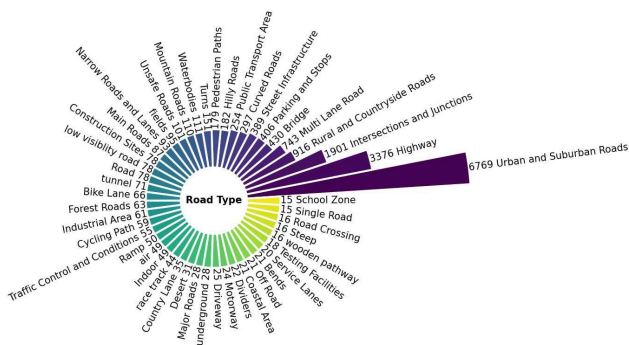


Figure 61. Road Type Video Tags distribution. Tags with fewer than 15 videos are omitted from the radar plot for clarity and to reduce clutter.