

MCAM: Multimodal Causal Analysis Model for Ego-Vehicle-Level Driving Video Understanding

Tongtong Cheng^{*,1}, Rongzhen Li^{†,2}, Yixin Xiong^{*,1}, Tao Zhang¹, Jing Wang³, and Kai Liu^{†,1}

¹Department of Computer Science, Chongqing University, China

²National Elite Institute of Engineering, Chongqing University, China

³College of Computer Science and Technology, National University of Defense Technology, China

Abstract

Accurate driving behavior recognition and reasoning are critical for autonomous driving video understanding. However, existing methods often tend to dig out the shallow causal, fail to address spurious correlations across modalities, and ignore the ego-vehicle level causality modeling. To overcome these limitations, we propose a novel Multimodal Causal Analysis Model (MCAM) that constructs latent causal structures between visual and language modalities. Firstly, we design a multi-level feature extractor to capture long-range dependencies. Secondly, we design a causal analysis module that dynamically models driving scenarios using a directed acyclic graph (DAG) of driving states. Thirdly, we utilize a vision-language transformer to align critical visual features with their corresponding linguistic expressions. Extensive experiments on the BDD-X, and CoVLA datasets demonstrate that MCAM achieves SOTA performance in visual-language causal relationship learning. Furthermore, the model exhibits superior capability in capturing causal characteristics within video sequences, showcasing its effectiveness for autonomous driving applications. The code is available at <https://github.com/SixCorePeach/MCAM>

1. Introduction

With the advancement of autonomous driving technologies, understanding driving videos has become a critical research area [17, 24, 29, 46]. This is due to the cost-effective acquisition of videos, which provide richer feature information than images [9, 19]. Moreover, understanding driving videos is crucial for comprehending autonomous driving behavior, supporting the wider adoption of intelligent vehicles. Recent studies have explored vision-to-language methods [6, 12, 15, 16, 43], which translate driving be-

haviors from visual data into high-level textual narratives. However, these methods [1, 13, 39] primarily focus on simple behaviors, neglecting the causal relationships between driving actions and their environmental context. Furthermore, the learning process is often hindered by spurious correlations, which could severely degrade the accuracy of causal analysis. Thus, developing a robust model with causality that mitigates spurious correlations is essential for advancing autonomous driving technologies.

Current approaches to driving behavior understanding can be categorized into two main streams: caption generation [34, 47] and retrieval-based matching [26, 27, 30, 41, 42]. For instance, Jin et al. [17] proposed ADAPT, a Video Swin Transformer-based architecture [23], which pairs driving behavior descriptions with comprehensive predictions. Similarly, Fan et al. [13] introduced MLLM-SUL, a framework combining multi-scale feature extraction with large language models (LLMs) like LLaMA for cross-modal video understanding. As illustrated in Figure 1, such methods establish causal links between actions (e.g., stopping) and environmental changes (e.g., traffic lights turning from yellow to red). While these approaches provide insights into driving behaviors, their causal reasoning is often limited to probabilistic correlations rather than detailed causal inference. On the other hand, retrieval-based methods, such as RAG-driver [46] and RAC3 [39], leverage retrieval mechanisms to enhance scene understanding by searching for similar scenarios. However, these methods heavily depend on large-scale datasets and fail to effectively mitigate the problem of spurious correlations inherent.

To reduce the detrimental impact of spurious causal relationships—stemming from spatial-temporal correlations in driving video understanding—and to establish authentic causal relationships, researchers have focused on a sophisticated realm of causal learning that emphasizes attention mechanisms [25] and causal reasoning [37, 38, 48], with particular emphasis on exploring and understanding spatio-temporal causal relationships. For example, Cong et al. [10] introduced spatio-temporal attention to capture intra-frame

*Equal contribution

†Corresponding authors

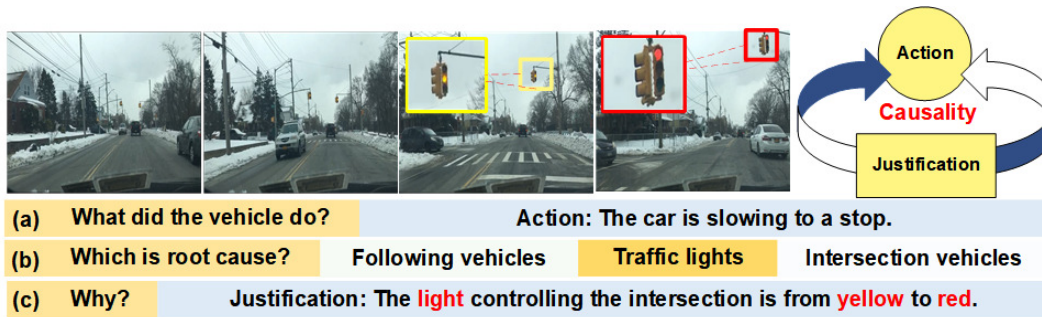


Figure 1. The relationship of justification and action exists in ego-vehicle level driving video understanding. In the initial state, the vehicle moves along the main road on the left side. There is a clear road with a visible traffic light in the distance. As the traffic light turns from yellow to red, the vehicle slows down and stops, preventing a violation of traffic regulations and reducing the risk of an accident. The yellow box and red box are added by manually to support the observation results.

and inter-frame relationships. Chen et al. [8] proposed the LLCP, which segments vehicle states into environmental context, neighboring vehicle’s behaviors, and historical states. In addition, Liu et al. [24] proposed the CMCIR framework, utilizing causal intervention techniques to facilitate event-level comprehension in driving scenarios. Despite their significant advancements, these approaches fail to adequately model the dynamic state transitions of causal entities, particularly overlooking ego-vehicle behavior understanding and the establishment of causal relationships.

To address this gap, this paper introduces the Multimodal Causal Analysis Model (MCAM), a framework designed for comprehensive driving video understanding. MCAM employs a Driving State Directed Acyclic Graph (DSDAG) to decompose complex driving behavior into discrete states, enabling fine-grained causal reasoning and dynamic interaction modeling. The framework consists of three components: (1) a Multi-level Feature Extractor (MFE) captures local and global features from videos; (2) a Causal Analysis Module (CAM) constructs causal relationships using DSDAG; and (3) a Vision-Language Transformer (VLT) integrates these features and relationships to generate detailed descriptions and reasoning. The main contributions are as follows:

- **Driving State Directed Acyclic Graph (DSDAG).** We propose Driving State Directed Acyclic Graph (DSDAG) to model dynamic driving interactions and state transitions, providing a structured theoretical foundation for the Causal Analysis Module (CAM).
- **Multimodal Causal Analysis Model (MCAM).** We propose a causality-analysis ego-vehicle level driving video understanding framework, named Multimodal Causal Analysis Model (MCAM), to discover true causal relationships through driving state modeling and achieve robust ego-vehicle level driving video understanding performance. To the best of our knowledge, we are the first to introduce causal analysis structures for ego-vehicle level

driving video understanding tasks.

- **Comprehensive Evaluation.** Extensive experiments on the BDD-X and CoVLA datasets show the effectiveness of our MCAM for Mitigating the interference of spurious correlations to achieve robust and promising ego-vehicle level driving behavior understanding.

2. Related Work

2.1. Causal Representation

Building on the structural causal model proposed by Pearl et al.[32], researchers have further explored effective causal modules. For instance, Yang et al.[45] introduced a causal attention mechanism that leverages front-door adjustment to enhance the causal learning capabilities of transformer architectures. Additionally, Wei et al.[40] developed a causal scene separation module to distinguish between causal and non-causal visual scenes. They also proposed a question-guided refiner to select consecutive video frames for modeling causal relationships. Furthermore, Rahman et al.[33] presented the Modular-DCM algorithm, which integrates causal structures with pre-trained conditional generative models to effectively eliminate the influence of confounding variables and accurately estimate causal effects in high-dimensional data. Similarly, Jin et al.[18] introduced CAUSALCOT, which applies back-door adjustment to improve causal reasoning and has achieved promising results. Temp-adapter[7] proposed a contrastive learning approach to align driving videos and text at both temporal and semantic levels, significantly improving the model’s scene understanding capabilities. To be different, we employ a feature disentanglement method to establish direct, indirect, and confounding causal relationships between scene features and behavioral features.

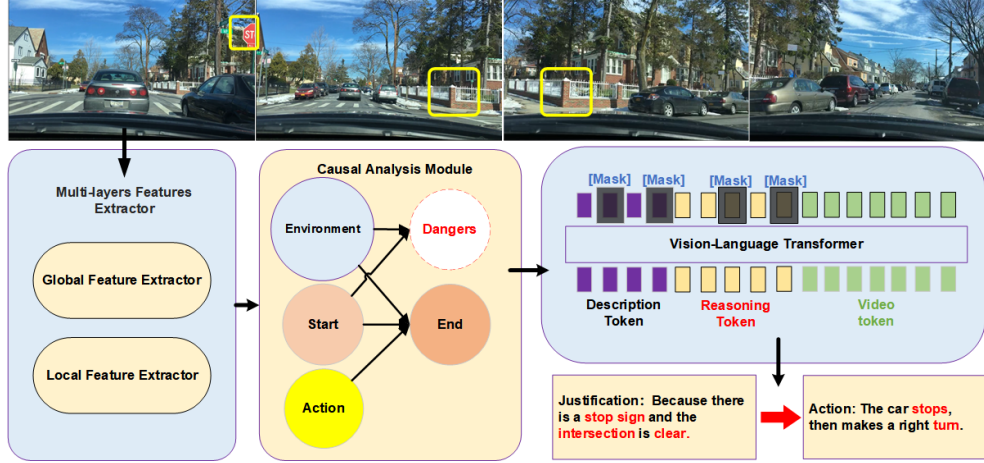


Figure 2. Overview of the MCAM. The MCAM framework consists of three key components: the MFE module, which extracts multi-level features from the video clip frames; the CAM, which builds and analyzes causal relationships within the extracted key features; the VLT module, which generates coherent and contextually relevant captions by integrating visual and textual information. The yellow boxes are added by manually to support the observation results.

2.2. Video Understanding

Video Understanding provides users with explanations of events and their logical relationships within videos[14][44], delivering cohesive and abstract information[21]. Numerous methods based on LLMs have been developed to enhance video understanding in complex environments. For instance, Ataallah et al.[4] proposed MiniGPT4-Video, a model that processes image and text tokens to align interactive content and projects them into the LLMs space, thereby improving video comprehension. Xu et al.[43] developed DriveGPT4, a multimodal LLM that leverages multimodal data to interpret scenes and enhances understanding through human interaction. Additionally, Yuan et al.[46] presented RAG-Driver, which integrates Retrieval-Augmented Generation (RAG) to identify similar scenes, thereby boosting the accuracy of video understanding. While these approaches leverage the extensive memory and computational power of large-scale LLMs[11][28] to achieve SOTA performance, MCAM adopts a more efficient and compact architecture with significantly fewer parameters.

3. Methodology

The framework of the MCAM is shown in Fig 2, which is an ego-vehicle visual driving video understanding architecture. In this section, we present the detailed implementations of MCAM.

3.1. Multi-level Feature Extractor

Conventional feature learning methods predominantly rely on Transformer-based architectures, which excel in capturing global dependencies but are relatively less efficient at extracting local features. To address this, we propose an

efficient Multi-level Feature Extractor (MFE) that simultaneously captures both global and local features. This approach enhances the model’s ability to understand complex patterns through multi-level feature representation.

In the Multi-level Feature Extractor, we utilize 3DResNet for local information extraction, and VidSwin for global information extraction. Given raw video frames of size $B \times F \times H \times W \times 3$, where B is the batch size, F is the number of frames, and each frame has dimensions $H \times W \times 3$, we feed the input to both the VidSwin and 3DResNet to extract global and local features, respectively. Then the features will be refined for size $B \times \frac{F}{2} \times \frac{H}{32} \times \frac{W}{32} \times 8C$, and after the windows setting, the feature size change to $B \times \frac{FC}{2} \times \frac{H \times W}{32 \times 32}$, which C equals 64. The feature from the local features extractor will be $B \times F \times \frac{H}{16} \times \frac{W}{16} \times 16C$, to keep the consistency of features, we set the down-sampling times as 4, adjust the channels by the 1×1 convolution, and fused the features by linear layer.

3.2. Causal Analysis Module

Existing methods are inadequate in addressing spurious causal interference in driving video understanding, as they neglect the importance of modeling the ego vehicle’s state. Therefore, we propose a causal analysis module to tackle this issue. Following Pearl et al.[32], we consider the status transformation graph G , consisting of a start safe status X_s and the unsafe status W , the action Y , the end safe status X_e , and changed environment Z . The safety state generally denotes the state of motion or stillness, encompassing relatively stable conditions such as stopped and uniform motion. Actions include behaviors that alter the vehicle’s state, such as accelerating, decelerating, braking, lane




Figure 3. the (a) is Driving Status Directed Acyclic Graph (DSDAG), and (b) is the Refined and involved causal graph from DSDAG.

changing, turning, and reversing. The environment encompasses weather, surrounding vehicles, traffic signals, pedestrians, obstacles, and all other elements that could impact vehicle operation. An unsafe state denotes a foreseeable collision or fall that could result if the vehicle's current driving state remains unchanged due to environmental changes.

Furthermore, we construct the driving status causal graph, which comes from the DSDAG, as shown in Fig 3a. In this graph, X_s represents the start status, Y denotes the action of the ego vehicle, X_e represents the end status, and Z represents the changed environment respectively. The following is hypothesis and defined based on the preceding discussion:

Driving Action (Y): This include the general operation when the ego car driving in the urban road.

$$Y = \{accelerating, decelerating, turning, \dots\} \quad (1)$$

Driving Environment (Z): This include the traffic condition of road which ego car stands .

$$Z = \{weather, road\ type, congestion, \dots\} \quad (2)$$

Vehicle State (U): This encompasses the current state of the vehicle, including:

$$U = \{speed, direction, loading\ condition, \dots\} \quad (3)$$

Safe Driving State (U_v): Safe driving status is determined by the combination of the vehicle's state U_s and the environmental conditions Z_s :

$$U_v = \{U_s, Z_s\} \quad (4)$$

Initial Safe Driving State (X_s), and Ended Safe Driving State (X_e): They are the general situation in the safe driving state.

$$\{X_s, X_e\} \in U_s \quad (5)$$

Hypothesis 1: The environment Z remains invariant at the instants corresponding to the initial state X_s and the final state X_e . However, during the vehicle's transition from X_s to X_e , the environment undergoes gradual and predictable changes.

Hypothesis 2: Throughout the transition from the initial state X_s to the final state X_e , the vehicle executes a sequence of actions Y . The change in the vehicle's state is jointly driven by the environment Z and the actions Y .

Hypothesis 3: The state transition graph encodes specific causal relationships through its nodes and edges. Specifically, the action Y is determined by the current state X_s and the environment Z , while the new state X_e is influenced by the action Y and the environment Z .

We posit that the various elements with same kind in the environment are linearly related. Defined the $\Gamma = \{\gamma_1, \gamma_2, \gamma_3, \dots, \gamma_n\}$ in the overall driving environment, each γ_i represents the key coefficient of factors during vehicle driving. For $\forall \xi, \Gamma_i$, the specific driving environment Z_ξ is defined as follows:

$$Z_\xi = \{U_v, \Gamma_i, \xi\}, \xi \perp \Gamma_i \quad (6)$$

where, U_v represents the vehicle state, and Z_s can be considered as the initial environment state when $\xi = \emptyset$, and Γ_i is part of the overall environment Γ . The specific driving environment Z_ξ combines the vehicle state, the relevant environmental factors, and a random noise term ξ that is dependent of the environmental factors. Accordingly (Fig 3b.(1)), the hidden danger state should be the result of the combined action of environmental changes and unchanged vehicle states. We define the set of dangerous states $W = \{w_1, w_2, \dots, w_p\}$, including collisions, falls, loss of control of vehicles, violations of road traffic regulations, etc. Here, p is a finite positive integer. The hidden danger state W_{hidden} is determined by the function F_W as follows:

$$W_{hidden} = F_W(U_s, Z_\xi), W_{hidden} \in W \quad (7)$$

The role of function F_W in determining whether the initial state is compatible with the environment is self-evident. To avoid the occurrence of dangerous state W , drivers will choose safety action Y at the appropriate time to adjust the vehicle's state. Define the set $Y_c = \{c_1, c_2, \dots, c_p\}$ of Y , which includes simple or complex program actions such as shifting gears, changing lanes, starting, stopping, etc. Un-

der the influence of environment Z_ξ , initial state X_s will develop towards dangerous state W_m , while under the action of safe driving action Y_c , initial state X_s will turn towards safe state X_e . The following formula:

$$U_e = F_Y(X_s, Z_\xi, Y_c) \quad (8)$$

$$X_e = \{U_e, Z_e\} \quad (9)$$

at this point, the secure state X_e transitions to the next initial state X_s , forming a logical closed loop.

Based on the above, the current task is to infer the reasons for the vehicle's driving behavior. This approach differs from traditional causal reasoning, which typically infers future events based on observed phenomena. Instead, it analyzes the underlying causes based on the outcomes. This can be expressed as the following equation:

$$U_s = F_{X_s}(U_v, Z_s) \quad (10)$$

$$Y_c = F_Y(Z_\xi|U_s), \xi \perp U_s \quad (11)$$

$$W = F_W(Z_\xi|U_s, do(Y_c = \emptyset)) \quad (12)$$

$$X_e = F_{X_e}(Z_\xi|U_s, do(Y_c = c)) \quad (13)$$

where ξ is random noise at the safe status X_s , F is computing function. As the Fig 3b.(2), in the driving video understanding task, the output result is the one or several factors in the environment that have the most significant impact on the original state. So the formula is as follows:

$$N = \{1, 2, 3, \dots, n, \dots, +\infty\} \quad (14)$$

$$Z_\xi = \{\gamma_1, \gamma_2, \dots, \gamma_n | n \in N\} \quad (15)$$

$$V = \{\gamma_x, \gamma_y, \gamma_z, \dots\} \quad (16)$$

when the environment Z_s interacts with the initial state X_s in a certain scenario and transitions to the final state X_e through driving action Y , certain variables V in the environment Z_ξ will serve as key influencing factors for determining the appearance of Y_c , while the influence of other variables can be ignored. These key influencing factors V_γ provide the primary basis for understanding the occurrence of action Y_c occurs.

$$\begin{aligned} U_e &= F_{X_e}(Z_\xi|U_s, do(Y_c = c)) \\ &= F_{X_e}(V|U_s, do(Y_c = c)) \\ F_V(V) &= \sum_{i=1} \sum_{j=1} P(U_e|Y_c = c_j, V = \gamma_i) \end{aligned} \quad (17)$$

$$\text{argmax}\{F_V(V)\} \quad (18)$$

where the probability of V could be computed by the distribution of V_γ and specific Y_c . During the driving status transformation process, these variables are closely related

to traffic regulations, road conditions, pedestrians, and vehicles. As shown in Fig 4, we implement the CAM to fusion the multi-level features as follow.

Given inputs come from MFE:

$$\begin{aligned} &F_{init_{global}}, F_{init_{local}}, F_{end_{global}}, \\ &F_{end_{local}}, F_{whole_{global}}, F_{whole_{local}} \end{aligned} \quad (19)$$

where F_{init} is the is the initial frame of the video, used to record the initial state of the ego-vehicle, the F_{end} is the ended frame of the video, used to record the ended state of the ego-vehicle.

Global-Local Fusion is performed as follows:

$$\begin{aligned} F_{init} &= W_{X_s} (F_{init_{global}} + F_{init_{local}}) \\ F_{end} &= W_{X_e} (F_{end_{global}} + F_{end_{local}}) \\ F_{pot} &= W_w (F_{whole_{global}}, F_{whole_{local}}) \\ F_{act} &= W_Y (F_{whole_{global}}, F_{whole_{local}}) \\ F_{ori} &= W_o (F_{whole_{global}}, F_{whole_{local}}) \end{aligned} \quad (20)$$

Causal fusion combines these features:

$$H = \text{Concat}(F_{init}, F_{end}, F_{pot}, F_{act}) \quad (21)$$

Attention mechanism computes weights α :

$$\alpha = \text{Softmax}(W_H H + b_H) \quad (22)$$

The final output feature F is computed by

$$F = \alpha \odot F_{ori} \quad (23)$$

3.3. Vision-Language Transformer

Lastly, we propose the Vision-Language Transformer (VLT) to translate causal feature relationships into language expressions. Due to the causal feature size is $B \times 4FC \times \frac{H}{32} \times \frac{W}{32}$, we tokenize the sentence, which includes both the description and the reasoning, and pad it to a fixed length. The tokenized sentence is then embedded, and a multi-layer perceptron (MLP) is used to align the causal features. To maintain the correct relationship between the video and the corresponding sentences and to prevent model hallucination, we employ the sparse attention mask[23]. The sparsity constraint is defined as follows:

$$L_{sparse} = \lambda \times \sum_{i=1}^M \left(\sum_{j=1}^M (|V_{(i,j)}|) \right) \quad (24)$$

where V is the relationship matrix between the causal features and word embeddings, and M is the length of the causal feature vector, λ is regularization parameter. This setting is crucial for ensuring consistency between the objects and the corresponding sentences, thus mitigating the risk of model hallucination. Additionally, the overall loss

in MCAM includes signal losses (if there are any signal inputs) and sentence prediction losses. The sentence prediction losses are composed of Cross Entropy Loss for optimization, KL divergence, the $Loss_{total}$ is as follow:

$$\begin{aligned}
 L_{signal} &= \frac{1}{2N} \sum_{i=1}^N |y_i - \hat{y}_i| + (y_i - \hat{y}_i)^2 \\
 L_{caption} &= -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^C y_{i,c} \log(\hat{y}_{i,c}) \\
 &\quad + \beta \cdot \sum_{i=1}^N \sum_{c=1}^C P(y_{i,c}) \log\left(\frac{P(y_{i,c})}{Q(y_{i,c})}\right) \\
 L_{total} &= L_{signal} + L_{caption} \tag{25}
 \end{aligned}$$

where N is the length of sequence, the C is the class of words, the y_i, \hat{y}_i is the sign and signal label, the $P(y_{i,c}), Q(y_{i,c})$ is the distribution of prediction, and ground truth, respectively.

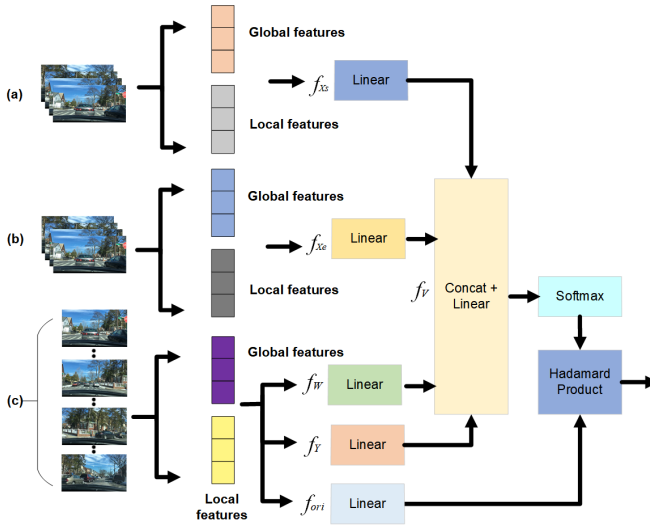


Figure 4. Overview of Causal Analysis Module, (a) is the first and second frame of video, and (b) is the last frame of video, (c) is the whole video, all of them are fed into the MFE to get the local and global feature, then, All features will be fused in special linear layer, and absorb input to f_V . the f_{ori} aims to protect the features as the residual layer.

4. Experiments

4.1. Datasets & Device

In this section, we conduct experiments on two ego-vehicle level driving understanding datasets: BDD-X [20] and CoVLA [3], to verify the effectiveness of MCAM and its components through comparing with existing methods.

Table 1. Experiment setting on All Datasets.

Datasets	Batch size	Training Epochs	Initial Learning Rate	Decays over time
BDD-X	16	40	0.0003	16
CoVLA	16	40	0.0001	16

BDD-X Dataset. The BDD-X dataset is an open-source driving-domain caption dataset, consisting of 7,000 videos paired with descriptions, reasoning, and control signals. Each video is segmented into frames of 1280×720 pixels at a uniform frame rate of 30 frames per second. The dataset is meticulously annotated, with each video segment containing no more than five actions, all described in textual form. We perform uniform sampling on video clips of varying lengths (from 3 to 28 seconds), extracting 32 frames while preserving the original frame rate information to ensure temporal accuracy. We follow the existing work[17], and split the BDD-X dataset into 21,143, 2,519, and 2,859 for training, validation and testing.

CoVLA Dataset. The CoVLA dataset consists of 10,000 videos, each paired with detailed descriptions and control signals. The descriptions encompass critical information such as vehicle status, weather conditions, traffic signal states, pedestrian activities, and essential safety guidelines for driving. Each video spans 30 seconds, with 20 captions generated per second. In this study, six fixed timestamps were uniformly sampled from each video, along with their corresponding labels. The dataset is divided into 42,000 video clips for training, 6,000 clips for validation, and 12,000 clips for testing.

Training Details. The MCAM is implemented using PyTorch 1.8 with CUDA 11.3. All video frames are pre-processed to a resolution of $224 \times 224 \times 3$, and the clips are sampled before training with 32 frames. Other settings are as follows Table 1. We train all models on Ubuntu 24.04 using single A100 GPU with 80 GB of memory for 40 epochs.

Evaluation Metrics. In this part, we evaluate MCAM through the metrics index, including BLEU-4(B4)[31], METEOR(M)[5], ROUGE(R)[22], and CIDEr(C)[35].

4.2. Main Results

Compared with SOTA methods on BDD-X. As shown in Table 2, we compare our proposed MCAM with the state-of-art method. The ADAPT was re-implement using the same training settings and hyperparameter configurations. Additionally, we compare MCAM with the following methods: S2VT[36], WAA, SAA [20], RAG-Driver[46], DriveGPT4[43]. By leveraging the construction of the DS-DAG, MCAM achieves superior accuracy on the BDD-X dataset compared to the baseline, while effectively mitigating spurious correlations. Specifically, for the narration task, MCAM achieves the B4 of 35.7%, because the causal analysis module could construct the dependence key factors with action. In the reasoning task, our MCAM achieves a



Figure 5. Visualization for comparing the result with ADAPT in BDD-X dataset, and the red boxes are added by manually to support the observation results. The first red box highlights the taillights of the vehicle ahead of the car in front, while the second red box indicates a 'STOP' sign that requires vehicles to stop and only allows right turns.

Table 2. Experiment on BDD-X dataset, the narration is driving action description and reasoning is explanation.

Method	Narration				Reasoning				Param	FPS
	B4	C	M	R	B4	C	M	R		
S2VT	30.2	179.8	27.5	-	6.3	53.4	11.2	-	86.3M	-
SAA	31.8	214.8	29.1	-	7.1	66.1	12.2	-	34.1M	-
WAA	32.3	215.8	29.2	-	7.3	69.5	12.2	-	35.7M	-
DriveGPT4 [43]	30.0	214.0	29.8	-	9.4	102.7	14.6	-	7.85B	-
RAG-Driver [46]	34.3	260.8	30.7	-	11.1	109.1	14.8	-	7.08B	-
Baseline	33.4	241.6	30.3	62.6	8.2	75.5	14.6	30.7	620.2 M	365
MCAM(ours)	35.7	252.0	30.5	63.4	9.1	94.1	14.0	32.0	885.3 M	336

B4 metric of 9.1%. In Fig 5.a, the MCAM accurately identifies the red brake lights on the vehicle and attributes the stop to the car in front. This indicates that the stop is more likely caused by the red traffic light rather than the preceding vehicle. The brake lights may interfere with the model’s judgment. In Fig 5.b, the stop sign is hard to recognize. The baseline erroneously attributes the turning action to the sign, creating a spurious relationship. Our method, however, focuses on the driving process, particularly the transition from the parking lot to another road. In Fig 5.c, both the baseline and MCAM fail to recognize the ego vehicle’s lane change requirement, instead misinterpreting the stopping behavior as parking intention. This occurs due to two key factors: (1) a yellow van in the front-right region is stationary while interacting with pedestrians, and (2) the vacated space from a departing leading vehicle presents a viable parking opportunity. While this interpretation diverges from ground truth, it represents a reasonable contextual inference. Notably, MCAM demonstrates superior performance in accurately localizing the parking target.

Generation Experience on CoVLA. On a driving video

understanding dataset of the same category, we reproduced the baseline model ADAPT and evaluated its performance using eight widely recognized metrics: BLEU-1 (B1), BLEU-2 (B2), BLEU-3 (B3), BLEU-4 (B4), METEOR (M), ROUGE (R), SIPCE(S)[2] and CIDEr (C). As depicted in Fig 6, the MCAM demonstrates effective identification of ego vehicle behaviors within videos. Specifically, as shown in Fig 6.a.(1), from the 1st to the 5th frames, the ego-vehicle maintains a straight trajectory, consistent with the label information. The MCAM not only accurately describes the vehicle’s movement state but also exhibits robust recognition capabilities for weather conditions and traffic signals. More importantly, the model almost comprehensively covers all critical attention factors in complex traffic scenarios. In Fig 6.a.(2), the model accurately identifies the number of cameras and interprets their content, offering pertinent advisories for safe vehicular operation. This capability effectively demonstrates the model’s inferential prowess. Regarding the identification of lane-changing behaviors, our method also performs excellently. Particularly in Fig 6b.(1), where labels generated by the Llama-7b model based on im-

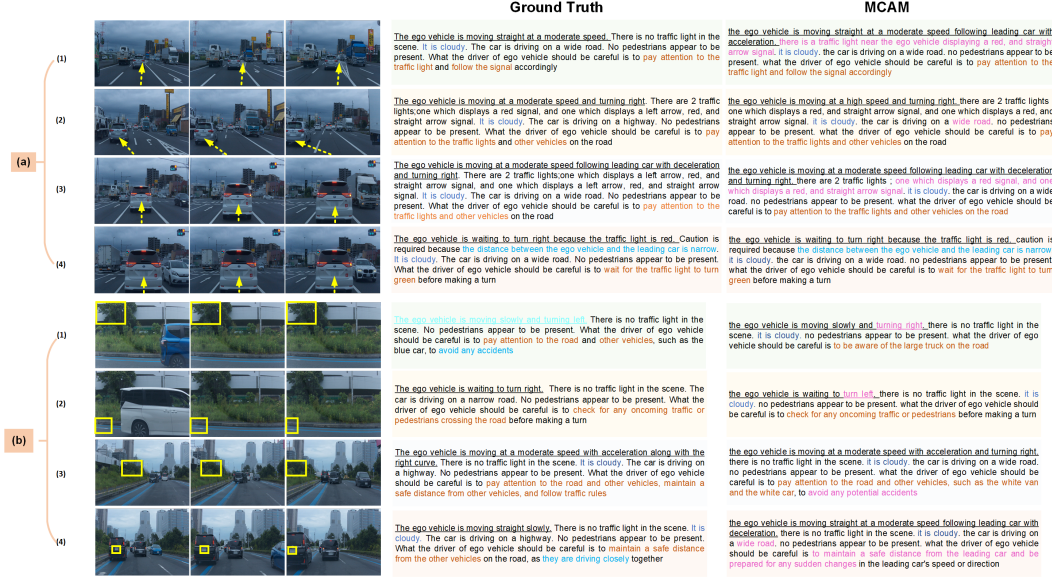


Figure 6. The visualization of MCAM on the CoVLA dataset, includes the ground truth (GT) caption, and the result from MCAM prediction. The yellow arrows and boxes is added by manually to support the observation results, and the words with underline indicate the description of ego-car, the blue words are the weather, brown words emphasis the caution of driving, the sky blue words are the reason of caution, and the words highlighted in pink indicate the discrepancy between the ground truth and our method.

Table 3. The generated experiment on CoVLA dataset.

Method	CoVLA									
	B1	B2	B3	B4	C	M	R	S	Para	FPS
Baseline	81.9	78.5	76.1	74.2	236.9	48.8	80.7	78.5	620.2M	94.2
MCAM(ours)	82.6	79.4	77.2	75.3	275.4	50.2	81.9	79.9	885.3M	86.22

Table 4. Ablation experiment on BDD-X dataset.

Method	Narration				Reasoning				Param	FPS
	B4	C	M	R	B4	C	M	R		
VidSwin + VLT	32.9	235.8	29.8	62.2	8.0	74.3	14.6	31.1	582.2M	374
3DResNet + VLT	32.8	221.8	28.97	61.4	6.4	55.3	12.5	28.1	494.0M	616
MFE + VLT	33.4	218.1	29.5	61.6	7.1	79.8	12.8	30.5	845.8M	364
VidSwin +CAM+ VLT	34.2	242.7	30.0	62.5	8.2	80.3	14.0	31.1	588.3M	336
MFE + CAM + VLT (MCAM)	35.3	251.6	30.7	63.1	9.0	92.9	13.8	31.7	885.3M	320

age understanding of specific frames incorrectly predicted the vehicle’s direction, our model, after fine-tuning, significantly corrects this misjudgment, demonstrating strong generalization ability.

4.3. Ablation Study

As is shown in Table 4, we implement the ablation experiment from the caption including narration and reasoning, to prove the reasonable combination of MFE, and CAM along VLT. For the video understanding task, an accurate description of driving actions forms the foundation. By analyzing the observed actions and key influential objects, the cause of the actions can be determined more precisely and convincingly. The 3DResNet paired with VLT could get lower performance than VidSwin with VLT due to lack of global information. When the MFE encounters the VLT,

narration, and reasoning performance has increased. The testing process of MFE is more steady than the single encoder, the accuracy in testing will be consistency with more steps training. Differently, including CAM, the combination of VidSwin and VLT could create a more accurate relation between the observed object and cause. Interestingly, the CAM could simultaneously indicate the narration and reasoning efficiently.

5. Conclusion

In this paper, we introduce the Multimodal Causal Analysis Model (MCAM) for modeling causality in ego-vehicle scenarios. MCAM incorporates a causal analysis module to enable effective identification of diverse driving behaviors under dynamic road conditions. This integration enhances the precision of predictions and interpretations of ego-vehicle behaviors. We validate the effectiveness of our approach on two ego-vehicle level benchmarks datasets BDD-X and CoVLA, demonstrating MCAM’s superior performance and broad applicability in real-world scenarios.

Limitation and Future Work. Our work also has certain limitations. For example, the datasets existing incorrect labels will impact model’s the recognized capability. Future work will construct higher-quality datasets and explore internal causal relationships among various objects in videos to further improve prediction accuracy and efficiency.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 62472055 and No. 62172064).

References

- [1] Nayyer Aafaq, Naveed Akhtar, Wei Liu, Syed Zulqarnain Gilani, and Ajmal Mian. Spatio-temporal dynamics and semantic attribute enriched visual encoding for video captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12487–12496, 2019.
- [2] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. Spice: Semantic propositional image caption evaluation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part V 14*, pages 382–398. Springer, 2016.
- [3] Hidehisa Arai, Keita Miwa, Kento Sasaki, Yu Yamaguchi, Kohei Watanabe, Shunsuke Aoki, and Issei Yamamoto. Covla: Comprehensive vision-language-action dataset for autonomous driving. *arXiv preprint arXiv:2408.10845*, 2024.
- [4] Kirolos Ataallah, Xiaoqian Shen, Eslam Abdelrahman, Essam Sleiman, Deyao Zhu, Jian Ding, and Mohamed Elhoseiny. Minigt4-video: Advancing multimodal llms for video understanding with interleaved visual-textual tokens. *arXiv preprint arXiv:2404.03413*, 2024.
- [5] Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005.
- [6] Feiyu Chen, Cong Xu, Qi Jia, Yihua Wang, Yuhan Liu, Hao-tian Zhang, and Endong Wang. Egocentric vehicle dense video captioning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 137–146, 2024.
- [7] Guangyi Chen, Xiao Liu, Guangrun Wang, Kun Zhang, Philip HS Torr, Xiao-Ping Zhang, and Yansong Tang. Tem-adapt: Adapting image-text pretraining for video question answer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13945–13955, 2023.
- [8] Guangyi Chen, Yuke Li, Xiao Liu, Zijian Li, Eman Al Suradi, Donglai Wei, and Kun Zhang. Llcp: Learning latent causal processes for reasoning-based video question answer. In *ICLR*, 2024.
- [9] Siran Chen, Yuxiao Luo, Yue Ma, Yu Qiao, and Yali Wang. H-mba: Hierarchical mamba adaptation for multi-modal video understanding in autonomous driving. *arXiv preprint arXiv:2501.04302*, 2025.
- [10] Yuren Cong, Wentong Liao, Hanno Ackermann, Bodo Rosenhahn, and Michael Ying Yang. Spatial-temporal transformer for dynamic scene graph generation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16372–16382, 2021.
- [11] Quang Minh Dinh, Minh Khoi Ho, Anh Quan Dang, and Hung Phong Tran. Trafficvlm: A controllable visual language model for traffic video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7134–7143, 2024.
- [12] Yang Du, Yuqi Liu, and Qin Jin. Reversed in time: A novel temporal-emphasized benchmark for cross-modal video-text retrieval. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 5260–5269, 2024.
- [13] Jiaqi Fan, Jianhua Wu, Jincheng Gao, Jianhao Yu, Yafei Wang, Hongqing Chu, and Bingzhao Gao. Mllm-sul: Multimodal large language model for semantic scene understanding and localization in traffic scenarios. *arXiv preprint arXiv:2412.19406*, 2024.
- [14] Lianli Gao, Yu Lei, Pengpeng Zeng, Jingkuan Song, Meng Wang, and Heng Tao Shen. Hierarchical representation network with auxiliary tasks for video captioning and video question answering. *IEEE Transactions on Image Processing*, 31:202–215, 2021.
- [15] Xin Gu, Guang Chen, Yufei Wang, Libo Zhang, Tiejian Luo, and Longyin Wen. Text with knowledge graph augmented transformer for video captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18941–18951, 2023.
- [16] Md Mohaiminul Islam, Ngan Ho, Xitong Yang, Tushar Nagarajan, Lorenzo Torresani, and Gedas Bertasius. Video recap: Recursive captioning of hour-long videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18198–18208, 2024.
- [17] Bu Jin, Xinyu Liu, Yupeng Zheng, Pengfei Li, Hao Zhao, Tong Zhang, Yuhang Zheng, Guyue Zhou, and Jingjing Liu. Adapt: Action-aware driving caption transformer. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7554–7561, 2023.
- [18] Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, LYU Zhiheng, Kevin Blin, Fernando Gonzalez Aduato, Max Kleiman-Weiner, Mrinmaya Sachan, et al. Cladder: Assessing causal reasoning in language models. In *Thirty-seventh conference on neural information processing systems*, 2023.
- [19] Salman Khan, Izzeddin Teeti, Reza Javanmard Alitappeh, Mihaela C Stoian, Eleonora Giunchiglia, Gurkirt Singh, Andrew Bradley, and Fabio Cuzzolin. Road-waymo: Action awareness at scale for autonomous driving. *arXiv preprint arXiv:2411.01683*, 2024.
- [20] 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.
- [21] Guorong Li, Hanhua Ye, Yuankai Qi, Shuhui Wang, Laiyun Qing, Qingming Huang, and Ming-Hsuan Yang. Learning hierarchical modular networks for video captioning. *IEEE transactions on pattern analysis and machine intelligence*, 2023.
- [22] Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the*

- 42nd annual meeting of the association for computational linguistics (ACL-04), pages 605–612, 2004.
- [23] Kevin Lin, Linjie Li, Chung-Ching Lin, Faisal Ahmed, Zhe Gan, Zicheng Liu, Yumao Lu, and Lijuan Wang. Swinbert: End-to-end transformers with sparse attention for video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17949–17958, 2022.
- [24] Yang Liu, Guanbin Li, and Liang Lin. Cross-modal causal relational reasoning for event-level visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10):11624–11641, 2023.
- [25] Shao-Yuan Lo, Poojan Oza, Sumanth Chennupati, Alejandro Galindo, and Vishal M Patel. Spatio-temporal pixel-level contrastive learning-based source-free domain adaptation for video semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10534–10543, 2023.
- [26] Kaijing Ma, Xianghao Zang, Zerun Feng, Han Fang, Chao Ban, Yuhan Wei, Zhongjiang He, Yongxiang Li, and Hao Sun. Llavilo: Boosting video moment retrieval via adapter-based multimodal modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2798–2803, 2023.
- [27] Avinash Madasu and Vasudev Lal. Icsvr: Investigating compositional and syntactic understanding in video retrieval models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1733–1743, 2024.
- [28] 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.
- [29] Ana-Maria Marcu, Long Chen, Jan Hünemann, Alice Karnsund, Benoit Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, Alex Kendall, Jamie Shotton, et al. Lingoqa: Visual question answering for autonomous driving. In *European Conference on Computer Vision*, pages 252–269. Springer, 2024.
- [30] WonJun Moon, Sangeek Hyun, SangUk Park, Dongchan Park, and Jae-Pil Heo. Query-dependent video representation for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 23023–23033, 2023.
- [31] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, 2002.
- [32] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [33] Md Musfiqur Rahman and Murat Kocaoglu. Modular learning of deep causal generative models for high-dimensional causal inference. *arXiv preprint arXiv:2401.01426*, 2024.
- [34] Mingkang Tang, Zhanyu Wang, Zhenhua Liu, Fengyun Rao, Dian Li, and Xiu Li. Clip4caption: Clip for video caption. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4858–4862, 2021.
- [35] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575, 2015.
- [36] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. Sequence to sequence-video to text. In *Proceedings of the IEEE international conference on computer vision*, pages 4534–4542, 2015.
- [37] Junyan Wang, Yiqi Jiang, Yang Long, Xiuyu Sun, Maurice Pagnucco, and Yang Song. Deconfounding causal inference for zero-shot action recognition. *IEEE Transactions on Multimedia*, 26:3976–3986, 2023.
- [38] Wei Wang, Junyu Gao, and Changsheng Xu. Weakly-supervised video object grounding via causal intervention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3933–3948, 2022.
- [39] Yujin Wang, Quanfeng Liu, Jiaqi Fan, Jinlong Hong, Hongqing Chu, Mengjian Tian, Bingzhao Gao, and Hong Chen. Rac3: Retrieval-augmented corner case comprehension for autonomous driving with vision-language models. *arXiv preprint arXiv:2412.11050*, 2024.
- [40] Yushen Wei, Yang Liu, Hong Yan, Guanbin Li, and Liang Lin. Visual causal scene refinement for video question answering. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 377–386, 2023.
- [41] Yicheng Xiao, Zhuoyan Luo, Yong Liu, Yue Ma, Hengwei Bian, Yatai Ji, Yujiu Yang, and Xiu Li. Bridging the gap: A unified video comprehension framework for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18709–18719, 2024.
- [42] Jilan Xu, Yifei Huang, Junlin Hou, Guo Chen, Yuejie Zhang, Rui Feng, and Weidi Xie. Retrieval-augmented egocentric video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13525–13536, 2024.
- [43] Zhenhua Xu, Yujia Zhang, Enze Xie, Zhen Zhao, Yong Guo, Kwan-Yee K Wong, Zhenguo Li, and Hengshuang Zhao. Drivept4: Interpretable end-to-end autonomous driving via large language model. *IEEE Robotics and Automation Letters*, 2024.
- [44] Liqi Yan, Cheng Han, Zenglin Xu, Dongfang Liu, and Qifan Wang. Prompt learns prompt: Exploring knowledge-aware generative prompt collaboration for video captioning. In *IJCAI*, pages 1622–1630, 2023.
- [45] Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. Causal attention for vision-language tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9847–9857, 2021.
- [46] Jianhao Yuan, Shuyang Sun, Daniel Omeiza, Bo Zhao, Paul Newman, Lars Kunze, and Matthew Gadd. Rag-driver: Generalisable driving explanations with retrieval-augmented in-context learning in multi-modal large language model. *arXiv preprint arXiv:2402.10828*, 2024.
- [47] Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. Vision-language models for vision tasks: A survey. *IEEE*

Transactions on Pattern Analysis and Machine Intelligence, 2024.

- [48] Yaochen Zhu, Yinhan He, Jing Ma, Mengxuan Hu, Sheng Li, and Jundong Li. Causal inference with latent variables: Recent advances and future perspectives. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 6677–6687, 2024.