

MECD+: Unlocking Event-Level Causal Graph Discovery for Video Reasoning

Tieyuan Chen , Huabin Liu , Yi Wang , *Member, IEEE*, Yihang Chen , Tianyao He, Chaofan Gan, Huanyu He, and Weiyao Lin , *Senior Member, IEEE*

Abstract—Video causal reasoning aims to achieve a high-level understanding of videos from a causal perspective. However, it exhibits limitations in its scope, primarily executed in a question-answering paradigm and focusing on brief video segments containing isolated events and basic causal relations, lacking comprehensive and structured causality analysis for videos with multiple interconnected events. To fill this gap, we introduce a new task and dataset, Multi-Event Causal Discovery (MECD). It aims to uncover the causal relations between events distributed chronologically across long videos. Given visual segments and textual descriptions of events, MECD identifies the causal associations between these events to derive a comprehensive and structured event-level video causal graph explaining why and how the result event occurred. To address the challenges of MECD, we devise a novel framework inspired by the Granger Causality method, incorporating an efficient mask-based event prediction model to perform an *Event Granger Test*. It estimates causality by comparing the predicted result event when premise events are masked versus unmasked. Furthermore, we integrate causal inference techniques such as front-door adjustment and counterfactual inference to mitigate challenges in MECD like causality confounding and illusory causality. Additionally, context chain reasoning is introduced to conduct more robust and generalized reasoning. Experiments validate the effectiveness of our framework in reasoning complete causal relations, outperforming GPT-4o and VideoChat2 by 5.77% and 2.70%, respectively. Further experiments demonstrate that causal relation graphs can also contribute to downstream video understanding tasks such as video question answering and video event prediction.

Index Terms—Causal discovery, causal reasoning, video understanding, video reasoning.

Received 14 January 2025; revised 20 October 2025; accepted 25 October 2025. Date of publication 28 October 2025; date of current version 4 February 2026. This work was supported in part by the National Natural Science Foundation of China under Grant U21B2013 and Grant 62325109, in part by the Shanghai ‘The Belt and Road’ Young Scholar Exchange under Grant 24510742000, in part by the National Key R&D Program of China under Grant 2022ZD0160102, and in part by the Zhongguancun Academy under Grant 20240313. Recommended for acceptance by N. Sebe. (*Corresponding authors: Huabin Liu; Weiyao Lin.*)

Tieyuan Chen and Weiyao Lin are with Shanghai Jiao Tong University, Shanghai 200240, China, and also with Zhongguancun Academy, Beijing 100095, China (e-mail: tieyuanchen@sjtu.edu.cn; wylin@sjtu.edu.cn).

Huabin Liu, Tianyao He, Chaofan Gan, and Huanyu He are with Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: huabinliu@sjtu.edu.cn).

Yi Wang is with Shanghai AI Laboratory, Shanghai 200232, China (e-mail: wangi@pjlab.org.cn).

Yihang Chen is with Shanghai Jiao Tong University, Shanghai 200240, China, and also with Monash University, Melbourne, VIC 3800, Australia.

Digital Object Identifier 10.1109/TPAMI.2025.3626549

I. INTRODUCTION

VIDEO causal reasoning aims to understand and analyze video content from a causal perspective. It requires models to comprehend temporal relationships, anticipate actions, and adapt to dynamic visual elements across frames. This capability is essential in various real-world applications, including autonomous driving [1], activity recognition [2], video surveillance [3], and even complex decision-making scenarios like robotic navigation [4]. Among these applications, Video Question Answering (VQA) [5], [6], [7] represents one of the most prominent tasks in video causal reasoning, where models are tested on their causal ability to understand videos through causal questions such as explanations, predictions, and counterfactual assumptions. Traditional VQA tasks can be viewed as attempts to discover single causal links within videos, as they typically identify one reason from given multiple options that explain the event described in the question.

Recent advancements in VQA have expanded beyond traditional tasks by incorporating more sophisticated causal reasoning methods. Notable trends include enhancing causal chain inference from simple one-to-one relations to multi-to-one chains, as seen in multiple correct answers VQA [8], [9], [10], and improving answer causal discovery through temporal grounding techniques, as seen in Grounded Video Question Answering (grounded VQA) [11], [12], [13]. Recent research also endeavors to infer a more comprehensive causal graph focusing on a particular event within its contexts [14], [15].

Despite recent advancements, as shown in Table I, current video causal reasoning tasks remain limited in scope, primarily focusing on QA-based approaches that discover a single causal relation within a video. However, when the scene to be analyzed is more complex, these tasks often fail to meet the required understanding standards. For example, in a video where a person slips and falls while carrying a tray of drinks, the cause might be combined with a wet floor, slippery shoes, and losing balance due to walking too quickly. Besides, these tasks often lack the ability to perform fine-grained event-level reasoning, which is crucial in real-world scenarios. Such detailed reasoning typically leads to clearer insights into video content. Most importantly, they fail to provide a comprehensive, structured causal representation. For instance, in traffic surveillance videos, a detailed analysis of events occurring at different times is essential to determine which events, or combinations of events, resulted in a specific accident among the chain of traffic accidents.

TABLE I
VIDEO REASONING TASKS COMPARISON. OUR MECD TASK UNIQUELY ENABLES THE REASONING OF A COMPREHENSIVE CAUSAL GRAPH AT THE EVENT LEVEL THROUGH CAUSAL REASONING.

Task	Work	Causal Reasoning	Event-Level	Multi-Chains	Complete Graph
Common VQA	Next-QA [6]	✓			
Multi-Answer VQA	ProViQ [8]			✓	
Grounded VQA	NextGQA [11]	✓	✓		
Causal Event Prediction	VAR [15]		✓	✓	
Causal Relation Based VQA	REXTIME [16]	✓	✓	✓	
Causal Graph Reasoning	Our MECD	✓	✓	✓	✓

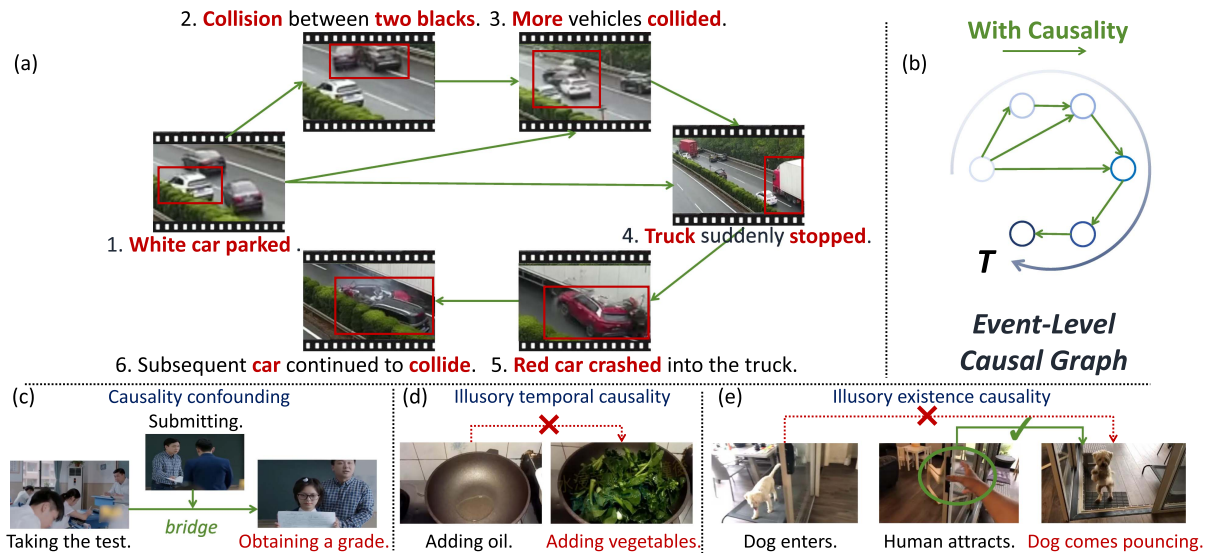


Fig. 1. (a) Illustration of Multi-Event Causal Discovery Task, where a complete causal graph of traffic surveillance videos is shown. Our task aims to determine whether a causal relation exists between events and outputs a structured causal graph in (b). (c) Example of causality confounding. (d) and (e) Illustration of illusory causality.

To address this gap, we set up a new task: **Multi-Event Causal Discovery (MECD)**, which aims to uncover causal relations among events that distribute chronologically.

As illustrated in Fig. 1, given multiple chronologically arranged event segments in a video along with their corresponding textual descriptions (Fig. 1(a)), MECD requires identifying causal relations between these events to derive a comprehensive and structured event-level causal graph (Fig. 1(b)), indicating why and how every event happens. Meanwhile, we contribute a new dataset for the training and evaluation of MECD by collecting moderately long-form videos involving multiple events and manually annotating the ground truth causal relation between any events pair.

However, to our knowledge, no available solutions can directly comprehend causal relations at the event level, necessitating the development of a new framework.

To this end, we draw inspiration from the *Granger Causality Method* [17], [18], [19] for solution, which is widely used in traditional causal discovery for low-dimensional time-series data (e.g., stock prices, weather patterns). The main idea is that temporal causality can be effectively estimated by predictive ability. Specifically, applied to videos, if Event A occurs prior to Event B, we consider A to be a cause of B only if A could facilitate the prediction of B. We term this criterion the *Event Causality Test*.

However, compared to simple low-dimensional data, the inputs of MECD involve much more complex modalities, including both visual and textual content, which may introduce bias in the estimation of causality using such a predictive paradigm. Specifically, we observe that directly applying *Event Causality Test* to video causal discovery presents two main problems:

1) *Causality confounding* indicates that the original causal relations between events are disrupted or interfered with by other relay or adjacent events. Such confounding stems from the fact that many causal relations flow through an intermediary event that acts as a bridge. For example, in Fig. 1(c), “submitting the paper” merely mediates the effect of “taking the test” on “obtaining a grade”, yet may be mistakenly treated as the sole cause, thereby obscuring the true causal dependency.

2) *Illusory Causality*, including both illusory temporal and existence causality. Illusory temporal causality exists when events exhibit a close correlation in temporal distribution. Such correlation may mislead the test of real causality. For example, in Fig. 1(d), “adding oil when cooking” often precedes “adding vegetables to stir-fry”, yet the former does not cause the latter. As for *illusory existence causality*, it occurs when some objects in early events may serve as necessary existence conditions of a later event. For instance, as shown in Fig. 1(e), while “a large brown dog enters the room” is a prerequisite for

“the dog runs towards the camera”, the former does not cause the latter.

Building upon the preceding discussion, we introduce a novel framework named VGCM (Video Granger Causality Model) to tackle MECD. This framework executes the *Event Granger Test* via an efficient mask-based event prediction model. According to the Granger Causality theorem in machine learning [17], [18], [19], if the inclusion of premise events can improve the prediction of a subsequent event, then they are causally related. Besides, the video data used in this article is theoretically and rigorously compliant with the conditions for Granger Causality, as it represents a continuous-in-time, deterministic, and complete system where potential confounders are mitigated through causal inference. Following this theory, VGCM deduces the causality of a premise event by comparing the predicted features of the result event when the premise is either masked or unmasked. Furthermore, to mitigate the challenges of causality confounding and illusory causality discussed earlier, we integrate two additional causal inference techniques—front-door adjustment [20], [21] and counterfactual inference [20], [22]—into our framework. Specifically, these techniques compensate for or remove the causal effects of previous or subsequent adjacent bridge events to eliminate confounding. Simultaneously, they address the illusory causality issue by incorporating the chain of thoughts [23], [24], [25] and existence-only descriptions.

Furthermore, we improve model robustness by applying context chain reasoning in the *Event Causality Test*, enhancing interactions among in-context causal chains. In reality, there are numerous instances where partial causes, when combined, contribute to the formation of the complete cause leading to the result. To enable the model to learn this reasoning pattern commonly found in reality, we randomly mask multiple premise events during training.

Extensive experiments validate the effectiveness and generalizability of our proposed framework VGCM in discovering structured causal relations for given long-form videos. Simultaneously, it has been proven that the complete causal graph inferred by our VGCM significantly promotes the downstream video understanding tasks such as video question answering and video event prediction.

This manuscript extends our NeurIPS 2024 paper [26] and makes the following contributions: The release and deeper analysis of MECD+, a larger multi-video source dataset. An enhancement of the *Event Causality Test* through robust context chain reasoning and efficient non-regressive complete graph reasoning. Additional metrics and further analysis of the MECD task, such as input modalities, hallucination issues, and downstream applications.

II. RELATED WORK

Video Causal Reasoning: Many existing tasks have tried to carry out causal reasoning in videos. Among these, the most common task is Video Question Answering (VQA), aiming to deduce a reason according to the question, grounded-QA

methods such as SeViLA, NextGQA and Momentor [5], [11], [27] grounded a single reason happens before the result. However, the traditional VQA task does not extend to abducting multiple reasons, merely abducting a single causal chain from reason to the provided result. Some studies have enhanced video reasoning paths by introducing Chain-of-Thought (CoT) to better accomplish VideoQA [28], [29], [30] or Rumor Detection [31], however, these works remain constrained by single-chain causal inference. Recently, some methods, such as Glance [32], VideoTree [12], and TimeCraft [13] have constructed more detailed causal chains through reasoning from coarse-grained to fine-grained levels or bidirectional paths. Although these works improve the causal reasoning process, they only abduct the causal chain for a single result.

Beyond the above classical setting, multiple correct answers VQA [8], [9], [10] focuses on a more flexible setting which could exist multiple correct reasons corresponding to the result, however, most of the choices have correlations or paradoxes between them. Therefore this task still does not suggest a generalized complete causal graph.

Furthermore, many tasks are based on VQA for further causal reasoning attempts. Physical state reasoning task CLEVRER [7] and CATER [33] explored causal reasoning based on physics and other basic laws in virtual scenes. However, they haven't been committed to extending to the general video. Neural-symbolic paradigm AAR [34] and LMLN [35] symbolized data and derived inference rules using external knowledge. However, they can only reason within a defined symbol domain. The most similar task VAR [15] predicted explanation events with premise events, and the causality is introduced during its prediction process. However, it hasn't been committed to discovering the complete causal graph. Another concurrent study REX-TIME [16] determined the causal relation between two events in a video and then designed QA questions targeted at the relations. However, it only reasoned through incomplete causal chains. Moreover, unlike the causal relation reasoning tasks mentioned above, the video spatial relationship reasoning task [36] aims at inferring the semantic connection between two objects formed by a specific action or state, rather than the causal relation that is our primary focus.

Most of these tasks introduced above are coarse video-level reasoning tasks, ours is more fine-grained event-level reasoning. Besides, we devote ourselves to abducting a complete causal graph rather than causal chains related to a single event. In conclusion, to the best of our knowledge, current video reasoning tasks haven't been committed to discovering complete causality among complex multi-event videos. Consequently, there exists a necessity need for a more comprehensive task.

Causal Discovery: Current causal discovery predominantly focuses on two main application areas: time series statistical data and natural language processing (NLP).

The primary objective of causal discovery in the context of time series statistical data is to infer the causal relation between the independent variable and the dependent variable. Traditional causal discovery methods are mainly divided into three categories: Constraint-based, Score-based, and Granger

Causality method. Constraint-based methods utilize conditional independence tests to identify causal relations [37], [38], and Score-based methods search through the space of all possible causal structures to optimize a specified metric [39], [40]. The constraint-based and score-based methods require stringent assumptions about data distribution, which would severely limit their generalization capabilities. Besides the two methods above, the more widely used Granger Causality method discovers causal relations by calculating the degree to which the earlier occurred event contributes to predicting the latter occurred event [41], [42].

The task of causal reasoning in natural language processing (NLP) generally involves inferring the causal relations between words within a text paragraph. TCR [43] proposes employing constrained conditional models and formulates the extraction of causal relations between events as an integer linear programming problem. The Following methods usually incorporate external knowledge, either through knowledge graphs or prompts. For example, RFBFN [44] and CauSeRL [45] utilized knowledge related to candidate causal events and external causal statements from the knowledge graph. However, these approaches are generally limited to handling sentences with a single causal relation pair. GESI [46] further advances this by constructing heterogeneous graph-based models to capture more complex causal relations, potentially spanning across multiple paragraphs. However, GESI suffers from slower inference speeds and increased complexity, which can be prohibitive for higher-dimensional video data applications. Additionally, following the Granger Causality method, CAUSE [47] proposes capturing event dependencies by fitting neural point processes and then employing axiomatic attribution to quantify the contribution of preceding events to subsequent predictions.

Therefore, for the task of causal discovery in video, it is not feasible to simply apply methods toward time series statistical data or NLP directly. Instead, we leverage the principles of causal discovery and design a pipeline that satisfies the specific characteristics of video data. Specifically, due to the straightforward nature of Granger Causality and its potential for integration with research related to masks in video understanding, it is selected as the foundational principle for reasoning in video data.

III. MULTI-EVENT CAUSAL DISCOVERY

To quantify the ability of causal discovery of a given model in multi-event videos, we propose the task of Multi-Event Causal Discovery (MECD). In formulation, given a video \mathcal{E} that contains chronologically organized N events, $\mathbb{E} := \{e_1, \dots, e_N\}$, this task aims at determining whether any previous event e_n ($n < N$) has a causal relation with the last one (i.e., e_N). Specifically, an event $e_n = \{v_n, c_n\}$ consists of a video clip v_n and the corresponding caption c_n . Without loss of generality, relations of previous events to the last one can be expressed as $r = [r_1, \dots, r_{N-1}]$, where r_k ($k < N$) is set to “1” to indicate the existence of e_k ’s causal relation with e_N , and “0” otherwise. Notably, this setting (*defined as causal chains discovery*) is generalizable to causal relations of any two events as long as we cut off the video and treat the latter one as the last event.



Fig. 2. Constituent of the MECD Dataset. We present 5 main video categories and the verb word cloud of the dataset.

A. Task Data

Data Source: The Multi Events Causal Discovery (MECD) task contains videos with multiple events and intricate causal relations. We carefully reorganize many widely used long-term daily videos including the ActivityNet Captions [48], EgoSchema [49], and NExTVideo [6] dataset.

ActivityNet Captions dataset [48] is built on ActivityNet v1.3 which includes 20 k YouTube videos. ActivityNet Captions is widely used in dense video captioning with annotated captions and timestamps of each sub-event. The EgoSchema dataset [49] contains over 5,000 video language understanding questions spanning over 250 hours of egocentric video data. NExTVideo [6] consists of 6,000 videos that are longer and richer in objects and interactions in daily life.

We carefully select 1,438 videos (5.6 k events) encompassing complicated causal relations and a wide range of scenarios as a new dataset: the MECD dataset, where 1,139 and 299 videos are randomly split for training and testing, respectively. In our dataset, 1,106 videos are sourced from ActivityNet, 102 from EgoSchema, and 230 from NExTVideo.

Specifically, each video in the MECD dataset contains 4 to 11 events, with a minimum of 2 premise events exhibiting causal relations with the last one. Fig. 2 presents the main categories and word clouds of video types.

Data Cleaning: As illustrated in ReXTime [16], only around ten percent of event pairs in daily video datasets such as ActivityNet Caption [48] are causal. Therefore, we further clean our dataset by excluding non-causal videos to guarantee the quality of our MECD dataset. We conducted data cleaning by five annotators with strict selection criteria: If two or more annotators (out of five) label a video as containing at most one causal chain, then the video will be excluded, e.g. video that describes steps such as handcrafting. A large portion of videos lacking causal relations were excluded—43.5% from ActivityNet, 24.3% from NExTVideo, and 74.2% from EgoSchema. The results reveal that NExTVideo clips exhibit stronger causal coherence, while EgoSchema clips often lack it due to their ego-procedural nature.

Dataset Annotation: The annotations of MECD dataset include 4 attributes. For videos sourced from the ActivityNet Captions dataset, The “duration”, “sentence”, and “timestamps” attributes in annotations remain the same as the original annotations for the dense video captioning task. For videos sourced from EgoSchema and NExTVideo, we select samples that are




Fig. 3. Statistics of the MECD Dataset. As shown in Figures (a1) and (a2), our dataset mainly analyzes videos that are longer than 120 seconds and contain five or more events. As shown in Figures (b1) and (b2), the proportion of causal and non-causal relations between events in the video of MECD is relatively balanced, moreover emphasizing the existence of causal relations between adjacent events.

drawn from Event-Bench’s [50] collection, where all clips have been verified to contain discrete events within standardized 30-second intervals. We then caption each event using the SOTA [51] video captioning model Gemini-Pro [52]. Additionally, our annotators further refine the “sentence” and “timestamps” attributes to: (1) refine low-quality captions resulting from background clutter or ambiguous frames, and (2) improve event continuity that might be compromised by strict 30-second division.

Specifically, in the context of our task, a new attribute, “causality”, is introduced. This attribute represents the causal relations between all premise events $\{e_1, \dots, e_{N-1}\}$ and the final event e_N . To obtain this attribute, relations among events are first annotated using the GPT-4 API [53], and subsequently refined by five human annotators. Through a cross-annotation process [54], [55], ground truth causalities are determined by the majority of the annotators’ causal relation choices, thus mitigating potential inaccuracies and subjective biases to a certain extent. Besides, a Pearson correlation coefficient of 0.93 was achieved among the annotations, strongly indicating a high degree of objectivity in the annotation process. Our experiments also confirmed that the model’s ability to learn causality is robust to minor variations in the annotations. Annotation examples of MECD are shown in Fig. 4, our MECD dataset is carefully annotated to support the challenging task proposed with complete premise information.

Furthermore, while the accuracy of deducing causality between the final event e_N and its preceding events is a good general indicator, it doesn’t fully capture the model’s ability. To rigorously demonstrate the model’s capacity to deduce the complete causal graph, an additional attribute, “complete causality”, is introduced for the test set. This attribute represents all causal relations between any two events, and is annotated and refined in the same way as the “causality” attribute. Specifically, “complete causality” is a list of length $(N-1)$, where N is the number of events in the video, and the k -th item (length = $(k+1)$) in this complete causality list represents the causal relations between $\{e_1, \dots, e_{k+1}\}$ and e_{k+2} as shown in Fig 4. The task of discovering the “complete causality” is defined as *complete causal graph discovery*.

Data Statistics: Our MECD dataset is primarily designed for causal reasoning in moderately long-duration videos, as illustrated in Fig. 3(a1) and (a2), the videos under consideration are predominantly those exceeding two minutes in duration, and all videos contain at least four distinct sub-events, events can be considered as complete actions with scenes in terms of our MECD task.

We also present the ratio of causal relations within every video in Fig. 3(b1), and the impact of events’ positions on their causality in Fig. 3(b2). The proportion of events with or without causal relations is generally balanced, with the proportion of events without causal relations slightly higher. Moreover, by examining the last event as a case study, we investigated the correlation between causality and positional relations among events. Our findings indicate that the likelihood of adjacent events existing causal relations is higher, while the probability of other positions is comparable.

IV. METHOD

In this section, we present our proposed Video Granger Causality Model (VGCM) to address Multi-Event Causal Discovery (MECD), as shown in Fig. 5. This model establishes the global connections across all events, and deduces the causality of a premise event by comparing the output features when it is masked or not, under the concept of the *Event Causality Test*. However, masking out an event may lead to the problem of confounding and illusion. We further employ causal inference methods to mitigate confounding by compensating for or removing the effects of preceding or subsequent causal events. Meanwhile, during inference, the extra chain of thoughts and existence-only descriptions help alleviate potential illusions. Additionally, we extend the modeling of event-causal relations to the global level and utilize context chain causal reasoning to increase the interaction of global information.

A. VGCM: Video Granger Causality Model

Building upon the Granger Causality method introduced in [56], [57], [58], our core motivation for constructing VGCM is



Fig. 4. Annotation Examples of the MECD Dataset. Newly annotated attributes “causality”, “complete causality” and the formerly existing attributes “sentences”, “duration”, and “timestamps” are shown along with the video frames.

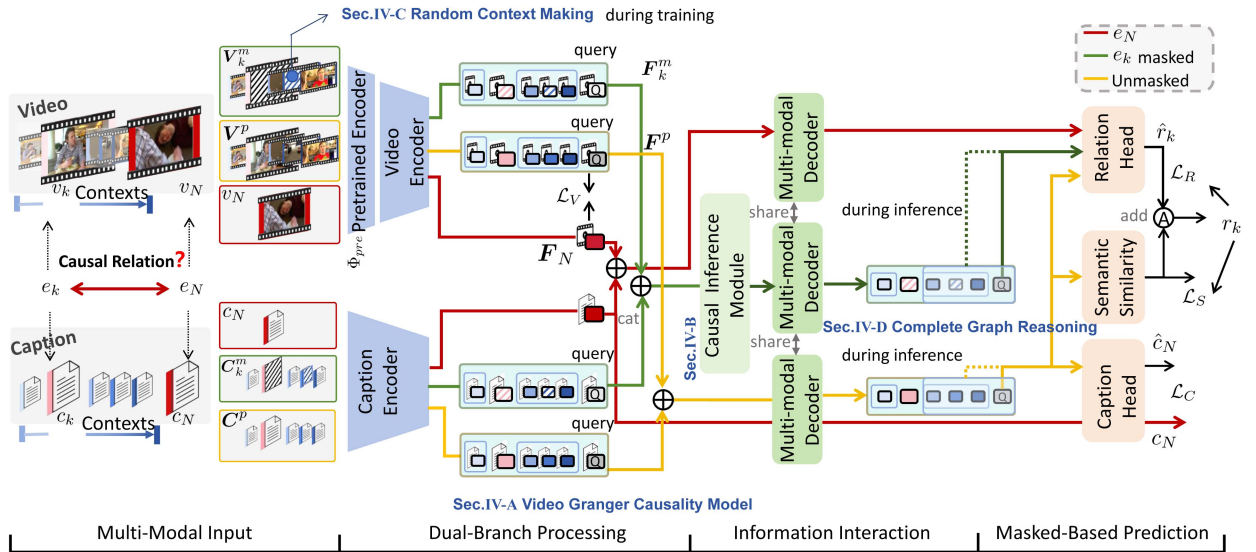


Fig. 5. Video Granger Causality Model. Two data streams $\{V^P, C^P\}$ (with certain e_k) and $\{V_k^m, C_k^m\}$ (without e_k) serve as input, video and text embeddings are concatenated after being separately embedded. During dual-branch processing, these two streams are encoded, and further multi-modal information interactions are conducted. The final mask-based prediction module conducts reasoning with two decoded prediction features and ground truth features embedded by $\{v_N, c_N\}$. Further causal inference is conducted for mitigating causality confounding and illusory as in Section IV-B, and random context masking and non-regressive complete graph reasoning are introduced for robust and efficient reasoning as in Sections IV-C and IV-D.

Event Causality Test: To compare the prediction of the last event using all the premise events with or without a certain event in it. If the results exhibit obvious divergence, it indicates that the current event is causally related to the result event.

We design VGCM to take in both the video clips and the captions to maximize information utilization. As illustrated in Fig. 5, our proposed VGCM is a multi-modal transformer-based structure with a video encoder and caption encoder, and a multi-modal decoder with causal relation head to discover causal relations through the predicting process and the comparison of predicting results.

Based on this, we denote \mathbb{E}^P as the set of all the *premise* events $\mathbb{E}^P := \mathbb{E} \setminus e_N$, and $\mathbb{E}_k^m := \mathbb{E}^P \setminus e_k$ as the event set where the premise event e_k ($k < N$) is masked. Notably, we mask the event e_k by setting all zeros to its video clip v_k and assign constant characters to the caption c_k .

Following [15], [59], [60], [61], we first pretrain a video encoder Φ_{pre} under an action recognition task to extract the

features of the video clips. We essentially create two paths, one for the unmasked event set \mathbb{E}^P (orange path in Fig. 5) while the other for the set with one event (i.e., e_k) masked \mathbb{E}_k^m (green path in Fig. 5). The video clips and captions are first separately encoded using Enc_V and Enc_C to obtain compact features, then their features are sent to a multi-modal decoder Dec that shares weights for both paths to fuse the information. Afterward, several model heads are employed for feature comparison and loss measurement. V^P and C^P are the video clip and caption matrix concatenated from all premise events set \mathbb{E}^P , similarly, V_k^m and C_k^m are from \mathbb{E}_k^m .

$$\begin{aligned}
 F^P &= \text{Enc}_V(\Phi_{pre}(V^P)), \\
 O^P &= [\text{Dec}(\text{Cat}(F^P, \text{Enc}_C(C^P)))]_{N-1}, \\
 F_k^m &= \text{Enc}_V(\Phi_{pre}(V_k^m)), \\
 O_k^m &= [\text{Dec}(\text{Cat}(F_k^m, \text{Enc}_C(C_k^m)))]_{N-1}, \quad (1)
 \end{aligned}$$

where Enc_V and Enc_C represent the encoder module of video clips and captions, respectively. Dec is a multi-modal decoder that shares weights for both paths. Cat indicates the concatenate operation, and $[-]_{N-1}$ indicates the $(N-1)$ -th slice at dimension 0. \mathbf{F}^p and \mathbf{F}_k^m are features after encoding, and \mathbf{O}^p and \mathbf{O}_k^m are the output features, which are then used for comparison of difference. Incorporating both visual and linguistic representations, the decoder conducts cross-modal reasoning and leverages the context from the unmasked premise events to posit a meaningful representation of the most likely explanatory result event.

Subsequently, the feature \mathbf{O}^p deduced from the unmasked events is sent to the caption head for supervised event prediction. Additionally, in order to compare the difference of the predictions, $\mathbf{O}^p, \mathbf{O}_k^m$ are directed to the relation head for reasoning. The result event e_N is encoded the same way as e_k to get feature $\mathbf{F}_N = \text{Enc}_V(\Phi_{pre}(v_N))$ and the output $\mathbf{O}_N = \text{Dec}(\text{Cat}(\mathbf{F}_N, \text{Enc}_C(\mathbf{C}_N)))$, \mathbf{O}_N is aggregated for reasoning (red path in Fig. 5). The relation head consists of a semantic query module and a self-enhancement module, where outputs are concatenated and then passed through the cross-reasoning layer g_r for further interaction. Last but not least, the auxiliary similarity is measured between \mathbf{O}^p and \mathbf{O}_k^m as a supplement to the output information of the relation head. After the reasoning process, the prediction output of the causal relation \hat{r}_k can be represented by:

$$\hat{r}_k = g_r(\text{Cat}(\Phi_{\text{att}}^C(\text{Cat}(\mathbf{O}_k^m, \mathbf{O}_N), \text{Cat}(\mathbf{O}^p, \mathbf{O}_N)), \Phi_{\text{att}}^I(\text{Cat}(\mathbf{O}_k^m, \mathbf{O}_N))))), \quad (2)$$

where Φ_{att}^C represents cross-attention, Φ_{att}^I represents self-attention, g_r represents linear layer. The training objective consists of two main directions as previously discussed:

To reconstruct the textual and visual representation of the result event e_N , we introduce caption loss and reconstruction loss, respectively. Caption loss \mathcal{L}_C ensures an accurate prediction of the result caption \hat{c}_N given all the premise events \mathbb{E}^p . Simultaneously, visual reconstruction loss \mathcal{L}_V forces the encoder to “imagine” a representation of the result video clip \hat{v}_N that better aligns with the original representation v_N . These losses allow the model to predict visual and textual representations that are close to the original representations, which better supports our method of inferring causal relations by comparing the results of the two-stream predictions.

For the objective of causal discovery, we introduce causal relation loss and an auxiliary semantics similarity loss. Causal relation loss \mathcal{L}_R supervised the output relations \hat{r}_k . Meanwhile, the semantics similarity loss \mathcal{L}_S is introduced to guarantee the semantics similarity of result event prediction under the presence or absence of a causal-relation-free premise event. The complete loss function is:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_C(c_N, \hat{c}_N) + \lambda_R \mathcal{L}_R(r_k, \hat{r}_k) + \lambda_V \mathcal{L}_V(\mathbf{F}_N^p, \mathbf{F}_N) \\ & + \lambda_S \text{sign}(r_k) \mathcal{L}_S(\mathbf{O}_k^m, \mathbf{O}^p), \end{aligned} \quad (3)$$

where λ_R, λ_V , and λ_S are weights for trade off. \mathcal{L}_C and \mathcal{L}_R are the cross-entropy losses, \mathcal{L}_V and \mathcal{L}_S are the mse losses, \mathbf{F}_N^p is the N th slice of \mathbf{F}^p indicating the feature of e_N .

Overall, the causal relation loss \mathcal{L}_R provides direct supervision for output causal relations using standard Cross-Entropy loss, while $\mathcal{L}_C, \mathcal{L}_V$, and \mathcal{L}_S aim to further enhance the causal discovery capability from different aspects. Specifically, \mathcal{L}_C and \mathcal{L}_V introduce auxiliary caption and reconstruction losses to facilitate event prediction, implemented by Label Smoothing loss and MSE loss respectively. They are incorporated because the Granger Causality Method determines whether an earlier event aids in predicting a subsequent one. \mathcal{L}_S , the similarity loss (implemented by InfoNCE), prompting causal relation discovery by comparing output causal feature similarity when a premise event is masked. The rationale behind this loss is that masking a non-causal event should result in a prediction of the result event similar to that of the unmasked stream.

B. Causal Inference

In Section IV-A, we employ the concept of Granger Causality to design our VGCM model under the principle of *Event Causality Test* which may, however, introduce causality confounding and illusory. Below we introduce these issues in detail, as well as how we manage to solve the problems.

Causality confounding is a phenomenon where the original causal relations across events are impacted due to modification (i.e., masking) of some intermediate events (i.e., e_k). Existing disentangled representation learning works [62], [63] disentangled different attributes of a variable under strict assumptions but failed in disentangling different variables.

Specifically, when e_k is masked for the comparison in VGCM, the causal relations between e_k 's adjacent events and the last event are impacted, leading to a confounding of causal effects. Notably, for brevity, we only employ e_k 's previous one event e_{k-1} and its subsequent one event e_{k+1} for analysis, but the same analysis also applies to all the previous or subsequent events. To be specific, there exist two distinct kinds of confounding when e_k is absent: **1)** Causal effects of e_{k-1} to e_N may be lost, as its connection to e_N is built upon e_k , (green path in Fig. 6(a1)). **2)** Causal effects of e_{k+1} to e_N may be redundant, as e_k may be a necessary prior of e_{k+1} 's connection with e_N , (red path in Fig. 6(a1)).

Illusory causality is another issue that may lead to some spatial or temporal misunderstandings, including illusory temporal and existence causality. **1)** Illusory temporal causality is the situation that events could have tight temporal ordering, but they in fact have no causal relations. **2)** Additionally, illusory existence causality occurs when an object introduced in the premise event is a necessary condition for the result, but the premise event does not semantically serve as a reason. Notably, we find that illusory in multi-event videos is much more significant than two independent events, which also tends to be exacerbated by causality confounding.

Overall, *causality confounding* and *illusory causality* both bring difficulties for relation modeling of events in videos.

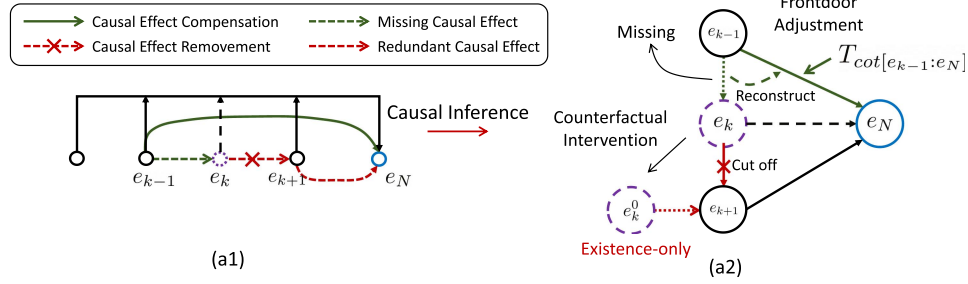


Fig. 6. Causal Effect of the Adjacent Events and Causality Graph. (a1) shows the causality of e_k analyzed, the causal effect in red needs to be mitigated because the causal effect residual remains through e_{k+1} . While the green needs to be compensated because the causal effect of e_{k-1} is affected. (a2) shows the causal inference methods (Front-door Adjustment and Counterfactual Intervention) corresponding to the two causal effects.

Notably, these two issues are coupled in that *causality confounding* tends to exacerbate *illusory causality* by misallocating attention to temporal ordering and causal effect. Therefore, illusory causality can be partially relieved by solving the problem of causality confounding.

When considering taking the illusory causality, the chain of thoughts [23], [24], [25] has been shown in LLMs to lead the model to logical thinking which is similar to human thought process, the chain of thoughts $T_{cot[e_{k-1}:e_N]}$ provides a step-by-step process of reasoning the e_N from e_{k-1} . Specifically, $T_{cot[e_{k-1}:e_N]}$ is obtained using GPT-4 API [53] by feeding it with e_{k-1} , e_N along with a prompt asking it to provide the probable reasoning chain. We consider utilizing it in causal inference to eliminate the attention bias on temporal correlations introduced by non-causal temporal knowledge.

Besides, as the illusory existence causality is caused by the objects' correlation between the events, we address this influence by keeping objects in the green path in Fig. 5 the same as those in the orange path. We introduce an alternative event $e_k^0 = \{v_k^0, c_k^0\}$ of e_k to briefly recaps the objects in e_k . Specifically, c_k^0 is obtained using GPT-4 API [53] by feeding it with c_k along with a prompt asking it to extract the objects from c_k and organize them as the sentence such as "There are objects A, B and C.". Consequently, we opt to employ c_k^0 to approximate e_k^0 in our VGCM model while omitting v_k^0 , as c_k^0 is sufficient already to convey the information of objects. By providing e_k^0 , all the necessary objects are still available in this path, thus effectively mitigating the illusory existence causality, facilitating the model to focus more on essential and causality-related semantic information.

To tackle the issues above, we introduce two causal inference methods: the front-door adjustment [64] for the missing causal effect of e_{k-1} and counterfactual inference [64] for the redundant causal effect of e_{k+1} . Meanwhile, the chain of thoughts $T_{cot[e_{k-1}:e_N]}$ and the descriptions of existence c_k^0 are also provided to carefully address illusory causality, which in turn mitigates confounding.

We establish a causality graph in Fig. 6(a2) for an improved elaboration. On masking e_k , the causality confounding that requires compensation F^C or removal F^R can be expressed as:

$$F^C = P(e_N|e_k) - P(e_N|do(e_k)),$$

$$F^R = P(e_N|e_{k+1}) - P(e_N|do(e_{k+1})), \quad (4)$$

where $P(e_N|e_k)$ and $P(e_N|e_{k+1})$ represents the process by which we predict e_N from e_k and e_{k+1} in the orange path in Fig. 5, and $do(\cdot)$ represents do-operation in causal inference [20] that cuts off the causal relation between the event and its causes.

We aggregate the subsequent events e_{k+1} , the current event e_k and the chain of thoughts $T_{cot[e_{k-1}:e_N]}$ using a linear layer g_{do} for aggregation and the cross-attention and self-attention, according to the study in [21], [65], $P(e_N|do(e_k))$ can be implemented as:

$$P(e_N|do(e_k)) = g_{do}(\text{Cat}(\Phi_{att}^C(e_k, e_{k+1}, e_{k+1}), \Phi_{att}^I(e_k, e_k, e_k), \text{Enc}_c(T_{cot[e_{k-1}:e_N]}))), \quad (5)$$

Here, we re-use the cross-attention Φ_{att}^C and the self-attention Φ_{att}^I modules as in (2) to cut off the causal effect from e_{k-1} to e_k through do-operation, e_k only interacts with subsequent events in predicting e_N . Then the missing causal effect F^C can be compensated since the causal-view operation and illusory temporal causality can be suppressed at the same time with the introduction of the chain of thoughts. Similarly, the redundant causal effect F^R can be removed by applying counterfactual intervention, then $P(e_N|do(e_{k+1}))$ can be represented by:

$$P(e_N|do(e_{k+1})) = P(e_N|e_{k+1})[P(e_{k+1}|e_k) - P(e_{k+1}|e_k^0)], \quad (6)$$

$P(e_N|do(e_{k+1}))$ effectively cuts off the redundant causal effect between e_{k+1} and e_N for the reason that the causes of e_{k+1} are replaced with counterfactual description e_k^0 , then the illusory existence causality can be suppressed simultaneously.

To refine the originally decoded feature O_k^m from the path with premise events masking:

$$O_k'^m = O_k^m - \text{Dec}(F^C) + \text{Dec}(F^R), \quad (7)$$

where $O_k'^m$ is the refined feature that replaces O_k^m for further deduction of the model. With the refinement feature $O_k'^m$, our VGCM model effectively compensates the connections between e_{k-1} and e_N that were originally lost due to the removal of e_k , and effectively removes the redundant causal effect between e_{k+1} and e_N as well.

C. Context Chain Reasoning

As discussed above, when conducting the *Event Causality Test*, a certain event e_k is masked, and the prediction result of



Fig. 7. An Example of Context Chain Reasoning. Given the knowledge that paint is stuck on the clothes, it is easier to infer the causal relation between being prohibited from taking the car and taking off one's jacket.

the last event e_N utilizing all the premise events with or without e_k is compared.

Although the *Event Causality Test* complies with the Granger Causality methodology, when applied to event-level reasoning in videos, it slightly overlooks the contextual relations between different events within the same video. This overlook is disadvantageous to the model's reasoning ability according to the research in [66], [67], particularly in inferring the intricate causal patterns of the entire causal graph.

In reality, there are numerous instances where partial causes, when combined, contribute to the formation of the complete cause leading to the result as shown in Fig. 7. Being prevented from getting into the car and taking off a jacket appears to have no causal relation. However, when considering prior knowledge that the jacket is covered in paint, this context aids in making a correct judgment through causal reasoning. Thus, we employ context chain reasoning to fully enhance the model's ability to leverage contextual events better. Therefore, we conduct the *Event Causality Test* by considering several premise events simultaneously during the training phase to enhance the context reasoning ability.

Specifically, during the training phase, we adopt a strategy of concurrently masking multiple events, $e_k = \{e_{k0}, \dots, e_{ki}\}$, (where $0 < i < (n - 1)$), to enhance the model's consideration of contextual events during reasoning. The ground truth causal relation is represented by the labels of the multiple events $e_k = \{e_{k0}, \dots, e_{ki}\}$ that are masked simultaneously, which are combined using an OR operation: $r_k = \bigvee_{j=0}^i r_{kj}$.

Furthermore, given that causal data in real-world events is relatively scarce [68], the imbalance in the data resulting from the OR operation of the labels can mitigate the disadvantage of classification bias on model training [69], further enhancing the model's generalization in judging causal relations.

D. Non-Regressive Complete Graph Reasoning

Through the *Event Causality Test* introduced in the previous sections, we can infer the causal chains between all premise events and the final event e_N . To further infer the complete event-level causal graph during the inference phase, a straightforward method treats each event as a result event and conducts *Event Causality Test* using a regressive approach [70], ultimately synthesizing a complete causal graph. However, this regressive approach is highly inefficient because inferring a causal graph for a video with N events requires the model to perform forward

propagation $(N - 1)(N - 2)/2$ times. This quadratic time complexity is suboptimal, especially given that our dataset contains an average of more than five events per video. To tackle this, we employ the following efficient non-regressive complete graph reasoning method during inference.

To maintain generality, we focus on the task of inferring the causal relation $\hat{r}_{[i:j]}$ between any of the i -th and j -th events, (where $1 \leq i < j \leq (N - 1)$). As the *Event Causality Test* introduced above, the mask-based reasoning process of relation head involves two prediction features under different conditions for comparison: O^p and O_i^m , which correspond to the j -th slice of the decoder output features defined in (1). Two output features of event O^p (when e_i is not masked) and O_i^m (when e_i is masked) can be represented by:

$$\begin{aligned} O^p &= [\text{Dec}(\text{Cat}(\mathbf{F}_i^p, \text{Enc}_C(\mathbf{C}^p)))]_j, \\ O_i^m &= [\text{Dec}(\text{Cat}(\mathbf{F}_i^m, \text{Enc}_C(\mathbf{C}_k^m)))]_j, \end{aligned} \quad (8)$$

where k equals i , as k iterates from 1 to $N - 1$ during inference of causal relations with e_N , no additional forward propagation is required. Our approximation method specifically involves utilizing the features O^p and O_i^m without additional masking of e_j , following the approach outlined in [70]. The ground truth result event e_j is encoded in the same manner as e_i , yielding the encoded feature $\mathbf{F}_j = \text{Enc}_V(\Phi_{pre}(v_j))$ and the output feature $\mathbf{O}_j = \text{Dec}(\text{Cat}(\mathbf{F}_j, \text{Enc}_C(\mathbf{C}_j)))$. Following a similar mask-based reasoning as (2), after reasoning, the output causal relation $\hat{r}_{[i:j]}$ can be represented by:

$$\begin{aligned} \hat{r}_{[i:j]} &= g_r(\text{Cat}(\Phi_{att}^C(\text{Cat}(\mathbf{O}_i^m, \mathbf{O}_j), \text{Cat}(\mathbf{O}^p, \mathbf{O}_j))), \\ &\quad \Phi_{att}^I(\text{Cat}(\mathbf{O}_i^m, \mathbf{O}_j))), \end{aligned} \quad (9)$$

As a result, just like inferring causal chains only related to the result event e_N , where inference requires just $(N - 1)$ forward propagation, this non-regressive complete graph reasoning method equips the VGCM with rapid inference speed.

V. EXPERIMENTS

A. Implementation Details

Our encoder, Enc_V , Enc_C , and multi-modal video decoder, Dec , are built upon Videobert [71], a joint model for video and language representation learning. Our model contains only 144M parameters, significantly smaller than 7B VideoLLMs. λ_C , λ_R , λ_V , λ_S are set to be 1.0, 4.0, 0.25, 0.05. Maximum input lengths of the caption, the chain of thoughts, and the existence-only descriptions are set to 50.

During the pre-training process, visual features of each video are extracted using a ResNet-200 model [72] pre-trained on the ActivityNet dataset [73] under the action recognition task, following the approach in [59], [60], [61]. Given the potential benefits of prior domain knowledge for the Granger Causality Causal method [74], we pre-trained our model for the dense video captioning task on a 3.1k-video dataset from the ActivityNet Captioning dataset [48], with each video sample containing more than four events.

Comparisons: We compare VGCM with baseline [71] and multi-modal base models such as CLIP-L [75], and a representative video reasoning model VAR [15]. Besides, we also conduct experiments on powerful LLM, including Mixtral-8x22B-Instruct [76], GPT-4 [53], Gemini-Pro [52], and so on. Additionally, ImageLLMs and VideoLLMs utilized for comparison include widely accepted GPT4-o [53], Gemini-Pro [52], VideoLLaVA [77], VideoChat2 [78], Long-VideoLLM LongVA [79], Qwen2.5-VL [80], and so on.

In the setting of few-shot learning (In-Context Learning), LLMs and ImageLLMs are evaluated, following the principle utilized in causal discovery tasks [81], [82], [83] in NLP. To more comprehensively assess the capacity for causal reasoning and to leverage the benefits of open-source models, the performance of VideoLLMs and all multi-modal base models is assessed under a strong fine-tuning paradigm. For all VideoLLMs, the vision-language projector was fully fine-tuned using Lora on the training set of the MECD dataset. Similarly, for multi-modal base models, the vision and text encoders are frozen, and a causal relation classifier, composed of the MLP layers and sigmoid function, is fully tuned.

Metrics: Our model has two output formats: one that outputs causal chains related to the result event, focusing on the model’s multi-cause reasoning ability (*towards causal chains discovery task*), and another that outputs a complete causal graph of all events (*towards complete causal graph discovery task*), ultimately reflecting the generalization ability in causal reasoning. Therefore, we designed two metrics to evaluate these aspects respectively.

We evaluate the model’s causal discovery capability by utilizing the top-1 accuracy of the output causal relation chains related to the final result event. Following the approach in natural language processing [68], [91], we further introduce the “Neg” metric to represent the accuracy when the model predicts no causal relation and the “Pos” metric to represent the accuracy when the model predicts the existence of a causal relation. This allows for a more nuanced analysis of whether the model tends to misjudge causal or non-causal relations.

In addition to the primary metrics, we introduce Structural Hamming Distance (SHD) [92], [93] as a supplementary metric to evaluate the model’s generalization ability in causal reasoning. SHD measures the degree of matching between comprehensive causal graphs by summing the number of incorrect causal relations. In the MECD test set, the average number of causal relations in video causal graphs is 12.31, and a lower Ave SHD value indicates better performance. An example of an evaluation of SHD is illustrated in Fig. 8.

B. Main Results

As shown in Table II, our proposed VGCM achieves state-of-the-art results in causal chains and causal graph reasoning tasks, with an accuracy of 71.28%. and an average SHD of 3.94. This surpasses existing MLLMs, including PLLaVA [88] and VideoChat2 [78], as well as proprietary models such as Gemini-Pro [52] and GPT-4o [53]. Specifically, VGCM demonstrates



Fig. 8. An Example of the Evaluation of SHD. The figure illustrates that the prediction of the complete causal graph on the right contains two missing edges and one extra edge. Consequently, the total number of incorrectly predicted causal relation edges (SHD value) is 3.

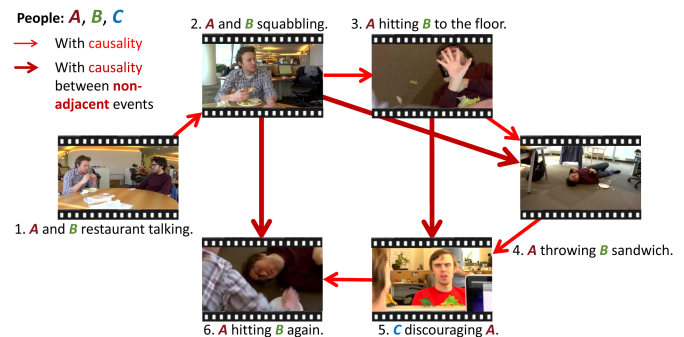


Fig. 9. Complete Causal Graph. The red arrows between events indicate causality, pointing from the cause to the result. Moreover, Events begin at the upper left corner and are sorted in a clockwise direction, with the causal relations between non-adjacent events being represented by thicker arrows.

improvements of over 0.74 in SHD and 2.70% in accuracy compared to the closest model, which exhibits reasoning capabilities most similar to human-level performance.

In the following **a) to d)**, we analyze the performances of VGCM, LLMs, VideoLLMs, and human reasoning. Furthermore, we perform an illusory test to assess the hallucination tendencies.

a) Our Performances: As shown in Table II, our VGCM outperforming the SOTA Proprietary LLM GPT-4 [53], ImageLLM GPT4-o [53], open-source basic multi-modal model SIGLIP [89], VideoLLM VideoChat2 [78] by 6.40%, 5.77%, 6.43%, and 2.70%, respectively.

Given that the Neg metric exceeds the Pos metric, our model initially shows a propensity to misclassify non-causal event pairs as causal. However, this issue is significantly alleviated following context chain reasoning, indicating that our model reduces the incidence of misjudged causes by enhancing the representation of contextual event relations.

Additionally, our VGCM further highlights its advantages in terms of complete causal graph reasoning, particularly without additional data supervision annotations. VGCM achieves performance with only an average of 3.94 false causal relations for a complete causal graph. An example is shown in Fig. 9.

As illustrated in Table III, we have assessed the inference speed of various models, with our VGCM achieving a swift

TABLE II
MAIN RESULTS. VGCM REACHES SOTA PERFORMANCE ON CAUSAL CHAINS AND COMPLETE CAUSAL GRAPH DISCOVERY. “w/o CR” INDICATES WITHOUT CONTEXT CHAIN REASONING.

Paradigm	Method	SHD ↓	Neg ↑	Pos ↑	Acc ↑	
Random	Guess all causal.	6.95	0.00	100.00	42.39	
	Guess all non-causal.	5.36	100.00	0.00	57.61	
In-context	<i>Open-Source LLM</i>					
	DeepSeek-Coder [84]	5.11	67.27	55.01	61.97	
	Mixtral-8x22B [76]	4.88	64.54	65.60	64.78	
	Qwen-Max-0428 [85]	4.95	69.28	58.57	64.80	
	<i>Proprietary LLM</i>					
	Gemini-1.5-Pro [52]	4.75	71.96	57.32	64.80	
	GPT-4-0613 [53]	4.72	74.07	56.80	64.88	
	<i>Proprietary ImageLLM</i>					
	Gemini-1.5-Pro [52]	4.70	68.97	62.52	65.15	
	GPT-4o [53]	4.69	70.33	61.53	65.51	
	<i>Open-Source VideoLLM</i>					
	MiniGPT-4 [86]	5.12	62.71	51.92	58.29	
	MiniGPT4-Video [87]	5.00	65.72	53.85	59.88	
	VideoLLaVA [77]	4.95	64.65	57.46	62.78	
	LongVA [79]	4.82	66.93	56.10	62.84	
	PLLaVA [88]	4.86	69.14	55.88	63.00	
	Qwen2.5-VL [80]	4.82	65.37	60.75	63.59	
	VideoChat2 [78]	4.75	70.63	55.57	63.85	
	<i>Open-Source Basic Multi-modal model</i>					
	Videobert [71]	4.95	62.33	57.82	60.92	
	VAR [15]	4.86	59.79	63.44	61.46	
	CLIP (ViT-L) [75]	4.77	63.88	61.54	62.87	
	SIGLIP [89]	4.80	64.13	65.39	64.85	
	BLIP-2 (ViT-L) [90]	4.75	63.75	66.90	65.54	
	Fine-tuning	<i>Open-Source VideoLLM</i>				
		PLLaVA [88]	4.74	69.78	63.35	65.71
		VideoLLaVA [77]	4.73	66.69	67.31	67.12
		Qwen2.5-VL [80]	4.69	69.37	67.15	68.06
VideoChat2 [78]		4.68	70.03	65.86	68.58	
<i>Ours</i>						
VGCM (w/o CR)		4.19	76.58	63.89	71.20	
VGCM		3.94	73.01	68.63	71.28	
Human		Only Textual Input	2.23	87.47	79.66	83.50
		Only Visual Input	2.09	88.09	85.72	87.02
	Complete Input	2.05	88.32	85.66	87.19	

TABLE III
INFERENCE SPEED. VGCM IS 3-6 TIMES FASTER THAN ALL VIDEOLLMS WHILE SLIGHTLY SLOWER THAN THE BASELINE. W/O NR INDICATES WITHOUT NON-REGRESSIVE COMPLETE GRAPH REASONING.

Model	Inference Speed (seconds/sample)
Videobert [71]	0.70
Our VGCM	0.76
CLIP (ViT-B/32) [75]	0.79
PLLaVA [88]	1.89
VideoLLaVA [77]	2.12
VideoChat2 [78]	2.96
VGCM (w/o NR)	3.39
MiniGPT4-Video [87]	3.98

0.76 seconds per sample. VGCM with non-regressive complete graph reasoning incurs an overhead of only 8.57% over the Videobert [71] baseline. VGCM’s inference speed is 3 to 6 times faster than that of all VideoLLMs. This shows that VGCM with the non-regressive approach achieves both precise and efficient reasoning capabilities.




Fig. 10. Failure Abduction Examples of GPT-4. Many failures of GPT’s causal reasoning are due to confusion with illusions and the conflation of subjective emotions with objective laws.

b) *LLM Performances*: As shown in Table II, Proprietary LLMs GPT-4 [53] and Gemini-Pro [52] have demonstrated the best performances among all LLMs. However, they are still limited by the influence of hallucinations and the conflation of subjective emotions with objective laws. As presented in Fig. 10, GPT-4 incorrectly infers that all premise events have a causal relation with the result event of the team winning.




Fig. 11. Comparison between Output Causal Chains of Mistral-22B-Instruct and Mixtral-8x22B-Instruct. LLM is facilitated with knowledge in a specific field of electrical appliances after utilizing a large-scale Mixture of Experts.



Fig. 12. Comparison between Output Causal Chains of EVA-CLIP and VideoChat2. VideoChat2 could be attributed to utilizing spatiotemporal cues to infer causal relations.

Moreover, the Neg metric generally surpasses the Pos, which may be attributed to the training corpus's frequent inclusion of data instances where two events are causally linked, alongside a scarcity of examples where events are not causally related [68]. Meanwhile, the balanced performance of Mixtral-8x22B-Instruct [76] may benefit from the use of a large-scale mixed expert model. As illustrated in Fig. 11, after utilizing the MoE, Mixtral-22B [76] successfully infers the causal knowledge involved in repairing electrical appliances.

c) VideoLLM Performances: As shown in Table II, under the paradigm of In-Context Learning, models such as VideoChat2 [78] and PLLaVA [88] have demonstrated superior performance, likely due to the inclusion of Next-QA [6] and CLEVRER [7] datasets in their pre-training data. Furthermore, under the paradigm of Fine-tuning, the performances are enhanced, though they still fall short of the performance of our VGCM. Moreover, under more frames of input, Long-VideoLLMs (LongVA, Qwen2.5-VL) do not outperform general VideoLLMs. However, their relatively better SHD performance indicates that long-sequence perception enhances the ability to reason about overall causal relations.

Besides, the performance of VideoLLMs suggests a reduction in the disparity between the Pos and Neg metrics, which may be attributed to the mitigation of hallucinations and the decrease in the influence of caption ambiguity due to adequate visual information. Furthermore, the enhancement in the Pos metric for VideoChat2 [78] could be attributed to the utilization of spatiotemporal cues to infer causal relations [78]. As shown in Fig. 12, after utilizing spatiotemporal cues, compared to the same encoders with MLP structure model EVA-CLIP [94], VideoChat2 [78] successfully discovered the long-term causal relations between the initial reason event “paying in cash” and the result event “having a quarrel”, reasoning leads to the ultimate source of results.

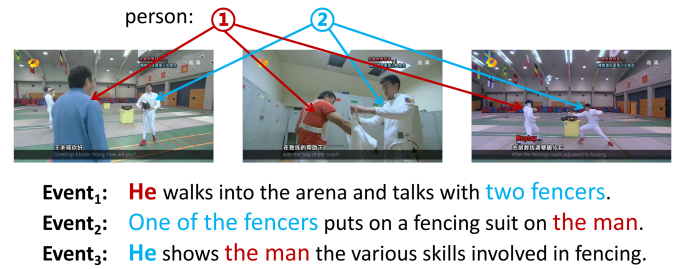


Fig. 13. Examples of Confusion due to Unclear Caption. The red arrows and text descriptions correspond to Person 1, while the blue ones correspond to Person 2. Evidently, the pronoun “He” successively refers to different individuals.



Fig. 14. Illusory Test Visualization Examples. LLMs and ImageLLMs suffer from serious illusory problems, leading to erroneous reasoning of non-existent causal relations between similar objects and scenes.

The reasons for the performance gap between VideoLLMs and VGCM are summarized here: (1) VGCM is tailored for causal reasoning, while general VideoLLMs lack causal understanding due to insufficient exposure during pretraining; (2) VGCM models bidirectional dependencies using the Granger Causality framework, enabling robust causal inference, while the autoregressive nature of VideoLLMs enforces an autoregressive bias that limits their ability to capture global event relations.

d) Human Performances: As shown in Table II, the average performance of ten volunteers reaches 87.19%. When only inputting one modality of information, higher accuracy is achieved when visual information is the sole input. We posit that this may be attributed to ambiguous pronouns can affect the judgment of their referents. An example of caption confusion is shown in Fig. 13, however, this problem can be released with additional visual input.

e) Illusory Test: To further assess the hallucination tendencies, we collect an additional validation set of 100 examples where causality does not exist between any events. Some examples include pairs of events with frequent statistical regularity, while others include cases where the prior event only serves as a necessary condition but does not directly lead.

As shown in Table IV, we can observe that the performances of LLMs drop significantly when judging the causal relations between examples with conditional correlation or statistical correlation in the temporal domain. Visualization examples are shown in Fig. 14, in example one, GPT-4o mistakenly believes that several unrelated cooking steps are causal. Although OpenAI [53] mentioned that GPT's upgrading process reduces

TABLE IV

ILLUSORY CAUSALITY TEST. † INDICATES WITH FINE-TUNING PARADIGM, WHILE * INDICATES WITH IN-CONTEXT LEARNING PARADIGM, THE CHANGE REFLECTS THE ACCURACY VARIANCE COMPARED WITH THE MECD TEST SET. LLMs AND IMAGELLMs SUFFER SERIOUS ILLUSIONS, HOWEVER, VGCM AND VIDEOLLMs DO NOT SUFFER.

	Method	Accuracy	Change
LLM*	GLM-4 (0520) [95]	54.96	-12.54
	Llama-3-70b [96]	58.39	-9.42
	GPT-4-0613 [53]	60.10	-13.97
	Gemini-1.5-Pro [52]	62.85	-9.11
	Mixtral-8x22B [76]	65.11	+0.57
	ImageLLM*	Gemini-1.5-Pro [52]	58.29
GPT-4o [53]		60.41	-9.92
VideoLLM*	Minigt-4 [86]	61.69	-1.02
	VideoChat2 [78]	65.28	-5.35
	PLLaVA [88]	67.01	-2.13
	VideoLLaVA [77]	67.15	+2.50
Base†	Videobert [71]	60.74	-1.59
	CLIP (ViT-L/14) [75]	63.00	-0.88
Ours†	VGCM	76.23	-0.35

TABLE V

ABLATION STUDY. ADJ INDICATES FRONT-DOOR ADJUSTMENT, INTER INDICATES COUNTERFACTUAL INTERVENTION, AND CONTEXT INDICATES CONTEXT CHAIN REASONING.

Base designs			Causal methods			Metrics	
\mathcal{L}_C	\mathcal{L}_V	\mathcal{L}_S	Adj	Inter	Context	Acc ↑	SHD ↓
	✓	✓				64.8	4.66
✓		✓				65.1	4.71
✓	✓					65.3	4.64
✓	✓	✓				67.0	4.59
✓	✓	✓	✓			68.7	4.27
✓	✓	✓		✓		69.3	4.32
✓	✓	✓	✓	✓		71.2	4.19
✓	✓	✓	✓	✓	✓	71.3	3.94

the hallucination issue in various other tasks, both GPT-4 and GPT-4o still show serious hallucination problems. In comparison, VideoLLMs and VGCM exhibit reduced susceptibility to illusions with the introduction of extra visual information.

C. Ablation Study

We conduct ablation studies to evaluate various components of the VGCM model, including the loss function, encoder design, causal methods, and input modalities, as described in Sections **a) to d)**. Besides we perform a detailed analysis of the architecture and input modalities of VideoLLMs. Additionally, we examine the annotation subjectivity and data volume of the MECD dataset.

a) Loss Function: We design our causal discovery model based on Granger Causality and applied three auxiliary losses \mathcal{L}_V , \mathcal{L}_C , and \mathcal{L}_S to enhance its reasoning capabilities. As shown in Table **V**, the benefits of \mathcal{L}_V and \mathcal{L}_C on our VGCM model

TABLE VI

ENCODER ABLATION STUDY. VIDEOBERT IS UTILIZED AS THE ENCODER BY DEFAULT, WE ALSO COMPARED THE PERFORMANCE OF UTILIZING OTHER MULTI-MODAL PRE-TRAINED ENCODERS SUCH AS CLIP.

	Encoder	SHD ↓	Acc ↑
VGCM	Univl [97]	4.24	68.33
	CLIP-B [75]	4.15	69.10
	CLIP-L [75]	4.17	69.99
	All-in-one-B+ [98]	4.24	70.01
	Videobert [71]	3.94	71.28

TABLE VII

ILLUSORY TEMPORAL CAUSALITY EXPERIMENT. W/O F INDICATES WITHOUT FRONT-DOOR ADJUSTMENT, THE MINOR ACCURACY DROPS OF PARTICULAR POSITION INDICATE THE ILLUSORY IS RELEASED.

Method	Δ Acc of r_0	Δ Acc of r_{N-1}	Overall Acc
VAR [15]	-3.5	-3.7	57.3
VGCM (w/o F)	-3.3	-3.2	66.9
VGCM	-0.7	-0.3	68.7

are evident, as they facilitate event prediction, which in turn supports the inference of causal relations. A stronger event prediction ability enables the model to better determine whether the existence of a certain event is beneficial for predicting the result event, thereby inferring causal relations. Additionally, \mathcal{L}_S contributes by supervising the prediction of e_N with and without certain non-causal event e_k masked, ensuring that the prediction of the result event is independent of the occurrence or non-occurrence of causally unrelated premise events.

b) Encoder: For our encoders, Enc_V and Enc_C , we explored the possibility of substituting them with other vision-language pre-trained models, including Univl [97], CLIP [75], and All-in-one-B+ [98] as indicated in Table **VI**.

The highest performance was achieved when Videobert [71] was employed as the encoder, the core reason is as follows. The mask-prediction training approach in Videobert [71] is particularly effective for learning contextual relations, making it well-suited for subsequent causal reasoning guided by Granger Causality. The principle of Granger Causality is to assess whether preceding event information improves the prediction of future outcomes—a concept that aligns with Videobert's objective of modeling contextual dependencies.

c) Causal Methods:

i) Front-door Adjustment: The method does improve reasoning ability in Table **V**. We conduct an experiment in Table **VII** for further proof. Since events closer to the result event are higher as the cause, the model likely learns these biased time-domain tendencies. So we compare the accuracy of VGCM without front-door adjustment and VGCM in determining the first relation r_1 and the last relation r_{N-1} . The results demonstrate that temporal illusory causality is greatly mitigated.

ii) Counterfactual Intervention: The performance in Table **V** shows that counterfactual intervention with existence-only descriptions does facilitate the model with powerful reasoning ability. We dive into further analysis on the basis that when

TABLE VIII

ILLUSORY EXISTENCE CAUSALITY EXPERIMENT. W/O C INDICATES WITHOUT COUNTERFACTUAL INTERVENTION, THE SIMILARITY DIVISION RESULTS INDICATE THAT ILLUSORY IS SUPPRESSED AS MODELS PAY MORE ATTENTION TO CAUSALITY RATHER THAN SIMPLE SEMANTICS.

Method	starting division	Ending division
VGCM (w/o C)	1.12	1.04
VGCM	1.12	0.93

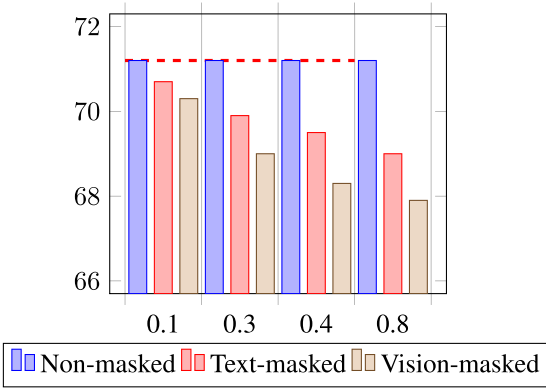


Fig. 15. Experiment of Input Modalities of the VGCM. The x-axis represents the proportion of masked textual input, and the y-axis represents the accuracy of causal reasoning. When 80% of textual or visual information is masked, VGCM still performs well in inferring causal relations, and visual information plays a more significant role.

a non-causal event is masked, the causal feature F_k^m fed into the causal relation head should be similar to the unmasked feature F^p , instead, a bigger gap appears when masking a causal event. For stronger proof, we measure the difference in feature similarity in Table VIII. We define the similarities division as the quotient of the similarity (F_k^m , F^p) with a non-causal e_k masked over with a causal e_k masked. In the experiment, we find that the similarity division is always above 1 without the counterfactual intervention, however, the division is below 1 with the help of counterfactual intervention.

iii) *Context Chain Reasoning*: The results presented in Table V indicate that, despite not enhancing the accuracy of causal chain discovery, context chain reasoning enhances the model's overall causal reasoning ability obviously.

d) *Input Modalities*: Typically, the input of VGCM consists of a textual input with an average of 13.5 words and a visual input of 50 frames. To investigate the influence of the text modality, we employ a masking strategy for the input caption of the premise event, gradually increasing the masking ratio from 10% to 80%. The results presented in Fig. 15 indicate that VGCM does not rely on the textual modality input.

In contrast, the experimental results suggest a more obvious performance decrease towards less visual modality input in the causality discovery task, as shown in Fig. 15.

e) *Architecture and Input Analysis of VideoLLMs*: We analyze the impact of the LLM module and input modalities in VideoLLMs, focusing on two aspects: how LLMs enhance

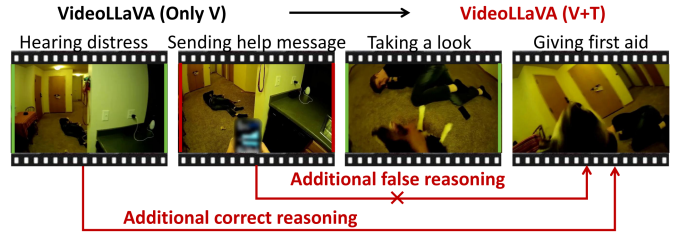


Fig. 16. Input Modality Test of VLLMs. When prompting VideoLLaVA with both textual and visual modalities compared with only visual input, VideoLLaVA successfully discovered more causal relations, however illusory problem arises.




Fig. 17. Dataset Robustness. Accuracy decreases slightly when increasing noise to the causal relation annotations of the training set, and increases slowly when training data increases.



Fig. 18. Visualizations for Downstream Video Question Answering. VideoLLMs offer more reasonable answers with the additional help of causal graphs reasoned from VGCM.

reasoning abilities compared to simple MLP layers, and the visual-to-textual information ratio in the MECD dataset.

i) *Reasoning Abilities with LLMs Versus MLPs*: We compare the performances between VideoLLMs and corresponding basic multi-modal models in Table IX. We found the understanding

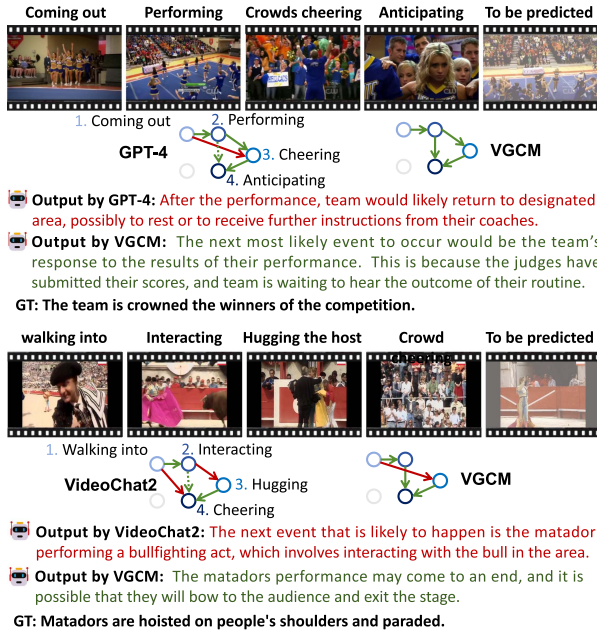


Fig. 19. Visualizations for the Downstream Event Prediction Task. VideoLLMs offer more reasonable event prediction results with the help of causal graphs output from our VGCM.

TABLE IX
PERFORMANCE COMPARISON BETWEEN THE BASE MULTI-MODAL MODEL AND CORRESPONDING VIDEO LLM. AN OBVIOUS GAIN IS ACHIEVED AFTER UTILIZING AN LLM REASONER.

Method	SHD ↓	Accuracy ↑
EVA-CLIP [94]	4.75	65.76
VideoChat2 [78]	4.68	68.58 (+2.82)
CLIP [75]	4.77	62.87
PLLaVA [88]	4.74	65.71 (+2.84)
LanguageBind [99]	4.92	59.83
VideoLLaVA [77]	4.73	67.12 (+7.29)

capabilities of LLMs significantly enhance the causal discovery of multi-modal foundation models versus simple MLP layers. Although visual and textual information are encoded by the same encoder, the models are equipped with enhanced reasoning capabilities through a more sophisticated 7B language model.

ii) *Input Modalities of VideoLLMs*: We compare the performances between VideoLLMs with all modalities input and only vision input as shown in Table X. When additional text information is input into VideoLLMs, it serves to supplement the limitations of visual information and emphasizes key details. However, ambiguous text cues can exacerbate illusions, which in turn leads to a decline in the Pos metric that contrasts with the Acc trend. The modest increase in the overall indicator of causal reasoning ability (+1.43) underscores a high Visual-Textual Information Ratio of the MECD dataset [78].

As shown in Fig. 16, when prompted by additional caption information, although a missing causal relation is discovered, the illusory causality problem of VideoLLaVA [77] becomes more serious. A more serious hallucination problem led to the

TABLE X
PERFORMANCE COMPARISON BETWEEN VIDEO LLM WITH AND WITHOUT TEXTUAL INPUT. MINOR INCREMENTS INDICATE A HIGH VISUAL-TEXTUAL INFORMATION RATIO OF THE MECD DATASET.

	Method	SHD ↓	Pos ↑	Accuracy ↑
V	MiniGPT4-Video [87]	5.15	54.90	57.45
	MiniGPT4 [86]	5.14	52.89	57.75
	VideoLLaVA [77]	5.10	60.44	60.58
	PLLaVA [88]	5.03	57.10	60.97
	VideoChat2 [78]	4.88	57.42	63.03
V + T	MiniGPT4 [86]	5.12	51.92	58.29 (+0.54)
	MiniGPT4-Video [87]	5.00	53.85	59.88 (+2.43)
	VideoLLaVA [77]	4.95	57.46	62.78 (+2.20)
	PLLaVA [88]	4.86	55.88	63.00 (+2.03)
	VideoChat2 [78]	4.75	55.57	63.85 (+0.82)

TABLE XI
DOWNSTREAM VIDEO QUESTION ANSWERING TASK. HIGHER VQA ACC AND VQA SCORE ARE REACHED WHEN PROMPTED WITH CAUSAL RELATIONS REASONED FROM OUR VGCM.

Output Causal Relations	VQA Acc	VQA Score
w/o (Standard QA setting for VLLMs)	43.17	2.82
w Gemini-Pro [52]	49.10	2.90
w GPT-4 [53]	49.36	2.89
w VideoChat2 [78] (Vicuna-7B)	51.01	2.97
w VideoLLaVA [77] (Vicuna-7B)	51.88	2.93
w PLLaVA [88] (Vicuna-7B)	52.47	2.96
w VideoChat2 [78] (Mistral-7B)	55.42	2.93
w Our VGCM	62.21	3.12

mistaken perception that “Sending help messages” would result in “Giving first aid”.

f) *Dataset Annotation and Volume*: We study the subjectivity and data volume of our proposed MECD dataset, which is shown in Fig. 17. In the experiments of increasing the ratio of randomly flipped annotated causal relations (flipping only one relation of the whole causal relations of video), the accuracy decreases slightly, demonstrating the small amount of subjectivity in labeling does not have a serious impact. Besides, we analyze the scale of training data, the increment from 1,000 examples to 1,139 examples yields a very modest improvement, indicating the adequacy of MECD.

D. Downstream Tasks

Moreover, to further validate the generalized reasoning capabilities of our model, we evaluate the quality of output causal relations on two related and representative video reasoning tasks: Video Question Answering (VQA) and Event Prediction (EP) as shown in Tables XI and XII.

a) *Downstream Task-VQA*: We identified 527 examples from the ActivityNet VQA dataset [100] that overlap with the MECD test set, which we used to evaluate the VQA task for this subset that has been separated. Specifically, we prompted MiniGPT4-video [87] with additional causal relations outputs alongside the standard question inputs. This paradigm facilitates the VideoLLMs in considering the task from a causal

TABLE XII

DOWNSTREAM EVENT PREDICTION TASK. HIGHER EVENT PREDICTION ACC AND EVENT PREDICTION SCORE ARE REACHED WHEN PROMPTED WITH CAUSAL RELATIONS FROM OUR VGCM.

Output Causal Relations	EP Acc	EP Score
w/o (Standard EP setting for VLLMs)	20.83	2.08
w Gemini-Pro [52]	28.44	2.54
w GPT-4 [53]	29.32	2.60
w VideoChat2 [78] (Vicuna-7B)	28.66	2.63
w VideoLLaVA [77] (Vicuna-7B)	28.71	2.61
w PLLaVA [88] (Vicuna-7B)	36.98	2.67
w VideoChat2 [78] (Mistral-7B)	39.39	2.71
w Our VGCM	43.39	2.73

perspective. As shown in the Table XI, when prompted with these additional causal relations, the answering accuracy of MiniGPT4-video [87] improved by VGCM surpasses VLLMs. These findings confirm that VGCM can provide accurate causal perception for videos, significantly improving performance on related video reasoning tasks concerning explanation, description, and prediction [7].

The visualization of the QA task’s comparison results, as depicted in Fig. 18, illustrates the first example as a temporal order question. Prompted with the causal graph generated by our model, the model correctly infers that food preparation should precede cooking. Conversely, when prompted with the causal graph produced by VideoLLaVA [77], the answers and temporal relations are chaotic.

b) Downstream Task-Event Prediction: Similar to the VQA task, we introduce a new task: Event Prediction. We utilize the first $n-1$ events from the MECD test set as known events, requiring the model to predict the content of the subsequent n -th event based on the acquired information. The evaluation method is the same as for the VQA task: we use the GPT-4 API to score the prediction results against the ground truth event and determine whether the prediction is correct. We prompted models with the following instruction “Please answer that what could happen next with the highest probability?”.

Specifically, we adopted a strategy similar to the VQA task, prompting MiniGPT4-video [87] with additional causal relation outputs in addition to the standard question inputs. As shown in Table XII, when prompted with these additional causal relations, the event prediction accuracy of MiniGPT4-video [87], enhanced by our VGCM, outperformed other strong VLLMs. These results confirm that our model can also enhance performance on related event prediction tasks.

As shown in Fig. 19, when prompting with the causal graph output from our VGCM, the model accurately infers from a series of prior events that the players have delivered a strong performance. It likely interprets the players’ nervous expressions as anticipation of the game results and their reactions to those results. In contrast, with the causal graph generated by GPT-4 [53], the model makes incorrect predictions based on common assumptions, such as the idea that the team would receive further instructions from their coaches.

VI. CONCLUSION

We proposed a novel task, multi-event video causal discovery (MECD), which focuses on event-level complete causal graph reasoning. To support this task, we developed the MECD dataset, consisting of 1,438 long-term daily-life videos featuring intricate causal relations. Additionally, we proposed the first model for video causal discovery, the Video Granger Causality Model (VGCM), which is based on the principles of the *Event Granger Test* and incorporates advanced causal inference techniques to address issues such as illusions and confounding factors. VGCM employs robust and efficient reasoning mechanisms, including context chain reasoning and non-regressive complete graph reasoning. Our VGCM outperformed GPT-4o and VideoChat2 by 5.77% and 2.70%, respectively, highlighting its superior reasoning capabilities.

REFERENCES

- [1] S. Park et al., “VLAAD: Vision and language assistant for autonomous driving,” in *Proc. Winter Conf. Appl. Comput. Vis.*, 2024, pp. 980–987.
- [2] J. Zhang, F. Shen, X. Xu, and H. T. Shen, “Temporal reasoning graph for activity recognition,” *IEEE Trans. Image Process.*, vol. 29, pp. 5491–5506, 2020.
- [3] N. M. Robertson and I. D. Reid, “Automatic reasoning about causal events in surveillance video,” *EURASIP J. Image Video Process.*, vol. 2011, pp. 1–19, 2011.
- [4] Q. Liu, G. Wang, Z. Liu, and H. Wang, “Visuomotor navigation for embodied robots with spatial memory and semantic reasoning cognition,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 36, no. 5, pp. 9512–9523, May 2025.
- [5] S. Yu et al., “Self-chained image-language model for video localization and question answering,” 2023, *arXiv:2305.06988*.
- [6] J. Xiao, X. Shang, A. Yao, and T.-S. Chua, “NEXt-QA: Next phase of question-answering to explaining temporal actions,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 9777–9786.
- [7] K. Yi et al., “CLEVRER: Collision events for video representation and reasoning,” 2019, *arXiv:1910.01442*.
- [8] R. Choudhury et al., “Video question answering with procedural programs,” in *Proc. Eur. Conf. Comput. Vis.*, Springer, 2025, pp. 315–332.
- [9] A. Yang, A. Miech, J. Sivic, I. Laptev, and C. Schmid, “Just ask: Learning to answer questions from millions of narrated videos,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 1686–1697.
- [10] R. Choudhury, K. Niinuma, K. M. Kitani, and L. A. Jeni, “Zero-shot video question answering with procedural programs,” 2023, *arXiv:2312.00937*.
- [11] J. Xiao, A. Yao, Y. Li, and T.-S. Chua, “Can i trust your answer? visually grounded video question answering,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 13204–13214.
- [12] Z. Wang et al., “VideoTree: Adaptive tree-based video representation for LLM reasoning on long videos,” 2024, *arXiv:2405.19209*.
- [13] L. Huabin, M. Xiao, Z. Cheng, Z. Yang, and L. Weiyao, “Timecraft: Navigate weakly-supervised temporal grounded video question answering via bi-directional reasoning,” in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 92–107.
- [14] C. Tan, C. K. Yeo, C. Tan, and B. Fernando, “Abductive action inference,” 2022, *arXiv:2210.13984*.
- [15] C. Liang, W. Wang, T. Zhou, and Y. Yang, “Visual abductive reasoning,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 15 565–15 575.
- [16] J.-J. Chen et al., “ReXTime: A benchmark suite for reasoning-across-time in videos,” 2024, *arXiv:2406.19392*.
- [17] A. Seth, “Granger causality,” *Scholarpedia*, vol. 2, no. 7, 2007, Art. no. 1667.
- [18] M. Maziarz, “A review of the granger-causality fallacy,” *J. Philos. Econ.: Reflections Econ. Social Issues*, vol. 8, no. 2, pp. 86–105, 2015.
- [19] A. Shojaie and E. B. Fox, “Granger causality: A review and recent advances,” *Annu. Rev. Statist. Appl.*, vol. 9, pp. 289–319, 2022.
- [20] J. Pearl, “Causal inference,” in *Proc. Workshop Causality: Objectives Assessment*, 2010, pp. 39–58.

- [21] X. Yang, H. Zhang, G. Qi, and J. Cai, "Causal attention for vision-language tasks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 9847–9857.
- [22] W. Wang et al., "Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue," in *Proc. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2021, pp. 1288–1297.
- [23] J. Wei et al., "Chain-of-thought prompting elicits reasoning in large language models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 24824–24837.
- [24] T. Kojima et al., "Large language models are zero-shot reasoners," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 22199–22213.
- [25] Z. Chu et al., "A survey of chain of thought reasoning: Advances, frontiers and future," 2023, *arXiv:2309.15402*.
- [26] T. Chen et al., "MECD: Unlocking multi-event causal discovery in video reasoning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 92554–92580.
- [27] L. Qian et al., "Momentor: Advancing video large language model with fine-grained temporal reasoning," 2024, *arXiv:2402.11435*.
- [28] Y. Wang et al., "VideoCoT: A video chain-of-thought dataset with active annotation tool," in *Proc. 3rd Workshop Adv. Lang. Vis. Res.*, 2024, pp. 92–101.
- [29] E. A. Hao Fei and S. Wu, "Video-of-thought: Step-by-step video reasoning from perception to cognition," in *Proc. Int. Conf. Mach. Learn.*, 2024, pp. 13109–13125.
- [30] S. Han et al., "VideoEspresso: A large-scale chain-of-thought dataset for fine-grained video reasoning via core frame selection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2025, pp. 26181–26191.
- [31] R. Hong et al., "Following clues, approaching the truth: Explainable micro-video rumor detection via chain-of-thought reasoning," in *Proc. Int. Conf. World Wide Web*, 2025, pp. 4684–4698.
- [32] Z. Bai, R. Wang, and X. Chen, "Glance and focus: Memory prompting for multi-event video question answering," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 34247–34259.
- [33] R. Girdhar and D. Ramanan, "CATER: A diagnostic dataset for compositional actions and temporal reasoning," 2019, *arXiv:1910.04744*.
- [34] T. Zhuo et al., "Explainable video action reasoning via prior knowledge and state transitions," in *Proc. ACM Int. Conf. Multimedia*, 2019, pp. 521–529.
- [35] Y. Jin, L. Zhu, and Y. Mu, "Complex video action reasoning via learnable markov logic network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 3242–3251.
- [36] Y.-H. H. Tsai et al., "Video relationship reasoning using gated spatio-temporal energy graph," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 10424–10 433.
- [37] A. Gerhardus and J. Runge, "High-recall causal discovery for autocorrelated time series with latent confounders," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 12615–12625.
- [38] C. K. Assaad, E. Devijver, and E. Gaussier, "Discovery of extended summary graphs in time series," in *Proc. Uncertainty Artif. Intell.*, 2022, pp. 96–106.
- [39] X. Sun et al., "NTS-NOTEARS: Learning nonparametric DBNs with prior knowledge," 2021, *arXiv:2109.04286*.
- [40] R. Pamfil et al., "DYNOTEARS: Structure learning from time-series data," in *Proc. Int. Conf. Artif. Intell. Statist.*, PMLR, 2020, pp. 1595–1605.
- [41] R. Cai et al., "THP: Topological hawkes processes for learning granger causality on event sequences," 2021, *arXiv:2105.10884*.
- [42] W. Chen et al., "Learning granger causality for non-stationary hawkes processes," *Neurocomputing*, vol. 468, pp. 22–32, 2022.
- [43] Q. Ning, Z. Feng, H. Wu, and D. Roth, "Joint reasoning for temporal and causal relations," 2019, *arXiv:1906.04941*.
- [44] Z. Li et al., "RFBFN: A relation-first blank filling network for joint relational triple extraction," in *Proc. Conf. Assoc. Comput. Linguistics*, 2022, pp. 10–20.
- [45] X. Zuo et al., "Improving event causality identification via self-supervised representation learning on external causal statement," in *Proc. Conf. Assoc. Comput. Linguistics*, 2021, pp. 2162–2172.
- [46] C. Fan, D. Liu, L. Qin, Y. Zhang, and R. Xu, "Towards event-level causal relation identification," in *Proc. 45th Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2022, pp. 1828–1833.
- [47] W. Zhang, T. Panum, S. Jha, P. Chhalasani, and D. Page, "CAUSE: Learning granger causality from event sequences using attribution methods," in *Proc. 37th Int. Conf. Mach. Learn.*, 2020, pp. 11235–11245.
- [48] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles, "Dense-captioning events in videos," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 706–715.
- [49] K. Mangalam, R. Akshulakov, and J. Malik, "EgoSchema: A diagnostic benchmark for very long-form video language understanding," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 46212–46244.
- [50] Y. Du et al., "Towards event-oriented long video understanding," 2024, *arXiv:2406.14129*.
- [51] Z. Cheng et al., "VideoLLaMA 2: Advancing spatial-temporal modeling and audio understanding in video-LLMs," 2024, *arXiv:2406.07476*.
- [52] R. Anil et al., "Gemini: A family of highly capable multimodal models," 2023, *arXiv:2312.11805*.
- [53] J. Achiam et al., "GPT-4 technical report," 2023, *arXiv:2303.08774*.
- [54] W. Lin et al., "HiEve: A large-scale benchmark for human-centric video analysis in complex events," *Int. J. Comput. Vis.*, vol. 131, no. 11, pp. 2994–3018, 2023.
- [55] R. Németh, D. Sik, and F. Máté, "Machine learning of concepts hard even for humans: The case of online depression forums," *Int. J. Qualitative Methods*, vol. 19, 2020, Art. no. 1609406920949338.
- [56] D. Lopez-Paz, K. Muandet, and B. Recht, "The randomized causation coefficient," *J. Mach. Learn. Res.*, vol. 16, pp. 2901–2907, 2015.
- [57] J.-F. Ton, D. Sejdinovic, and K. Fukumizu, "Meta learning for causal direction," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 9897–9905.
- [58] H. Li, Q. Xiao, and J. Tian, "Supervised whole DAG causal discovery," 2020, *arXiv:2006.04697*.
- [59] J. Lei et al., "MART: Memory-augmented recurrent transformer for coherent video paragraph captioning," 2020, *arXiv:2005.05402*.
- [60] T. Wang, R. Zhang, Z. Lu, F. Zheng, R. Cheng, and P. Luo, "End-to-end dense video captioning with parallel decoding," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 6847–6857.
- [61] Z. Zhang et al., "Object relational graph with teacher-recommended learning for video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 13278–13 288.
- [62] A. Hyvarinen and H. Morioka, "Nonlinear ICA of temporally dependent stationary sources," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2017, pp. 460–469.
- [63] A. Hyvarinen and H. Morioka, "Unsupervised feature extraction by time-contrastive learning and nonlinear ICA," in *Proc. Proc. Int. Conf. Neural Inf. Process. Syst.*, 2016, pp. 3772–3780.
- [64] J. Pearl, *Causality*. Cambridge, U.K.: Cambridge Univ. Press, 2009.
- [65] Y. Liu, G. Li, and L. Lin, "Cross-modal causal relational reasoning for event-level visual question answering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 10, pp. 11624–11641, Oct. 2023.
- [66] W. Zhang et al., "Cause: Learning granger causality from event sequences using attribution methods," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 11235–11245.
- [67] D. W. Gow Jr and B. B. Olson, "Sentential influences on acoustic-phonetic processing: A granger causality analysis of multimodal imaging data," *Lang. Cogn. Neurosci.*, vol. 31, no. 7, pp. 841–855, 2016.
- [68] J. Gao, X. Ding, B. Qin, and T. Liu, "Is ChatGPT a good causal reasoner? a comprehensive evaluation," 2023, *arXiv:2305.07375*.
- [69] K. Tang, J. Huang, and H. Zhang, "Long-tailed classification by keeping the good and removing the bad momentum causal effect," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 1513–1524.
- [70] A. Tank et al., "Neural granger causality," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 8, pp. 4267–4279, Aug. 2022.
- [71] C. Sun, A. Myers, C. Vondrick, K. Murphy, and C. Schmid, "VideoBERT: A joint model for video and language representation learning," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 7464–7473.
- [72] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [73] F. C. Heilbron, V. Escorcia, B. Ghanem, and J. C. Niebles, "ActivityNet: A large-scale video benchmark for human activity understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 961–970.
- [74] C. Gong et al., "Causal discovery from temporal data: An overview and new perspectives," 2023, *arXiv:2303.10112*.
- [75] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [76] A. Q. Jiang et al., "Mixtral of experts," 2024, *arXiv:2401.04088*.
- [77] B. Lin et al., "Video-LLaVA: Learning united visual representation by alignment before projection," 2023, *arXiv:2311.10122*.
- [78] K. Li et al., "MVBench: A comprehensive multi-modal video understanding benchmark," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 22195–22 206.
- [79] P. Zhang et al., "Long context transfer from language to vision," 2024, *arXiv:2406.16852*.
- [80] S. Bai et al., "Qwen2. 5-VL technical report," 2025, *arXiv:2502.13923*.


- [81] X. Wang et al., “Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 15614–15638.
- [82] K. Luo et al., “Open event causality extraction by the assistance of LLM in task annotation, dataset, and method,” in *Proc. Workshop: Bridging Neurons Symbols Natural Lang. Process. Knowl. Graphs Reasoning*, 2024, pp. 33–44.
- [83] A. Vashishtha et al., “Causal inference using LLM-guided discovery,” 2023, *arXiv:2310.15117*.
- [84] Q. Zhu et al., “DeepSeek-Coder-V2: Breaking the barrier of closed-source models in code intelligence,” 2024, *arXiv:2406.11931*.
- [85] J. Bai et al., “Qwen technical report,” 2023, *arXiv:2309.16609*.
- [86] D. Zhu et al., “MiniGPT-4: Enhancing vision-language understanding with advanced large language models,” 2023, *arXiv:2304.10592*.
- [87] K. Ataallah et al., “MiniGPT4-video: Advancing multimodal LLMs for video understanding with interleaved visual-textual tokens,” 2024, *arXiv:2404.03413*.
- [88] L. Xu et al., “PLLaVA : Parameter-free LLaVA extension from images to videos for video dense captioning,” 2024, *arXiv:2404.16994*.
- [89] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, “Sigmoid loss for language image pre-training,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 11975–11986.
- [90] J. Li, D. Li, S. Savarese, and S. Hoi, “BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 19730–19 742.
- [91] J. Gao et al., “Enhancing complex causality extraction via improved subtask interaction and knowledge fusion,” 2024, *arXiv:2408.03079*.
- [92] V. Rodríguez-López and L. E. Sucar, “Knowledge transfer for causal discovery,” *Int. J. Approx. Reasoning*, vol. 143, pp. 1–25, 2022.
- [93] K. Biza, I. Tsamardinos, and S. Triantafillou, “Tuning causal discovery algorithms,” in *Proc. Int. Conf. Probabilistic Graphical Models*, 2020, pp. 17–28.
- [94] Q. Sun et al., “EVA-CLIP: Improved training techniques for CLIP at scale,” 2023, *arXiv:2303.15389*.
- [95] A. Zeng et al., “ChatGLM: A family of large language models from GLM-130B to GLM-4 all tools,” 2024, *arXiv:2406.12793*.
- [96] Meta, “Introducing meta llama 3: The most capable openly available LLM to date,” 2024. [Online]. Available: <https://ai.meta.com/blog/meta-llama-3/>
- [97] H. Luo et al., “UniVL: A unified video and language pre-training model for multimodal understanding and generation,” 2020, *arXiv:2002.06353*.
- [98] J. Wang et al., “All in one: Exploring unified video-language pre-training,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 6598–6608.
- [99] B. Zhu et al., “Languagebind: Extending video-language pre-training to N-modality by language-based semantic alignment,” 2023, *arXiv:2310.01852*.
- [100] Z. Yu et al., “ActivityNet-QA: A dataset for understanding complex web videos via question answering,” in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 9127–9134.



Tieyuan Chen received the bachelor’s degree from the School of Electronic Information and Electrical Engineering, Sichuan University, in 2023. He is currently working toward the PhD degree with Shanghai Jiao Tong University and Zhongguancun Academy. His research interests include causal discovery, causal reasoning, and video reasoning.



Huabin Liu received the bachelor’s degree from the School of Electronic Engineering, Xidian University, in 2019. He is currently working toward the PhD degree with Shanghai Jiao Tong University. His research interests include action understanding, video reasoning, and captioning.



Yi Wang (Member, IEEE) received the PhD degree with the Chinese University of Hong Kong, in 2021. He is a research scientist with Shanghai AI Lab. His current research interests include video understanding and multimodal large language models.



Yihang Chen received the bachelor’s degree from Electronic Information Engineering, Dalian University of Technology, in 2021. He is currently working toward the PhD degree with the joint program of Shanghai Jiao Tong University and Monash University. His current research interests include 3D vision and data compression techniques.




Tianyao He received the bachelor’s degree from the School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University. He is currently working toward the master’s degree with Shanghai Jiao Tong University. His current research interests include video analysis and multimodal comprehension.



Chaofan Gan received the bachelor’s degree in software engineering from the Huazhong University of Science and Technology, Wuhan, China, in 2022. He is currently working toward the PhD degree in electronic information and electrical engineering with Shanghai Jiao Tong University, Shanghai, China. His research interests include video understanding, diffusion models, and noise learning.



Huanyu He received the BE degree from Nanjing University, in 2016 and the ME degree from Shanghai Jiao Tong University, in 2020. He is currently working toward the PhD degree with the School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University. His current research interests include machine learning and computer vision, especially on object detection and pedestrian detection.



Weiyao Lin (Senior Member, IEEE) received the BE and ME degrees from Shanghai Jiao Tong University, Shanghai, China, in 2003 and 2005, respectively, and the PhD degree from the University of Washington, Seattle, WA, USA, in 2010, all in electrical engineering. He is currently a professor with the Department of Electronic Engineering, Shanghai Jiao Tong University. He has authored or coauthored more than 100 technical articles on top journals/conferences including the *IEEE Transactions on Pattern Analysis and Machine Intelligence*, the *International Journal of Computer Vision*, the *IEEE Transactions on Image Processing*, CVPR, NeurIPS, ICLR, and ICCV. He holds more than 20 patents. His research interests include video/image analysis, computer vision, and video/image processing applications.