
MECD: Unlocking Multi-Event Causal Discovery in Video Reasoning

Tieyuan Chen^{1*}, Huabin Liu^{1*}, Tianyao He¹, Yihang Chen¹, Chaofan Gan¹,
Xiao Ma², Cheng Zhong², Yang Zhang², Yingxue Wang³, Hui Lin³, Weiyao Lin^{1†}

¹ Shanghai Jiao Tong University, ² Lenovo Research, AI Lab,

³ China Academic of Electronics and Information Technology

{tieyuanchen, huabinliu, wylin}@sjtu.edu.cn

<https://github.com/tychen-SJTU/MECD-Benchmark>

Abstract

Video causal reasoning aims to achieve a high-level understanding of video content from a causal perspective. However, current video reasoning tasks are limited in scope, primarily executed in a question-answering paradigm and focusing on short videos containing only a single event and simple causal relationships, lacking comprehensive and structured causality analysis for videos with multiple events. To fill this gap, we introduce a new task and dataset, **Multi-Event Causal Discovery (MECD)**. It aims to uncover the causal relationships between events distributed chronologically across long videos. Given visual segments and textual descriptions of events, MECD requires identifying the causal associations between these events to derive a comprehensive, structured event-level video causal diagram explaining why and how the final result event occurred. To address MECD, we devise a novel framework inspired by the Granger Causality method, using an efficient mask-based event prediction model to perform an *Event Granger Test*, which 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 address challenges in MECD like causality confounding and illusory causality. Experiments validate the effectiveness of our framework in providing causal relationships in multi-event videos, outperforming GPT-4o and VideoLLaVA by 5.7% and 4.1%, respectively.

1 Introduction

Video causal reasoning aims to achieve a high-level understanding and analysis of video content from a causal perspective. Video Question Answering (VQA) [1–5] represents one of the most prominent tasks in causal reasoning, where models are tested on their causal ability to understand video content through causal questions such as explanations, predictions, and counterfactual assumptions. Recently, some studies have sought to move beyond the single QA task, attempting to construct more complex and challenging video reasoning tasks and methodologies. For example, CLEVRER [5], V-CDN [6] and CATER [7] explored more difficult causal reasoning tasks in virtual scenes by constructing object-aware features or using graph neural networks. Neural-symbolic paradigm AAR [8] and LMLN [9] extended to derive inference rules by symbolizing data. VAR [10] and BiGED [11] extended to daily video causal reasoning by introducing causality during prediction.

However, current video causal reasoning tasks are still limited in scope (primarily QA-based) and mainly focus on short videos containing only a single event or a few events. Most importantly, they

*Equal Contribution.

†Corresponding Author.

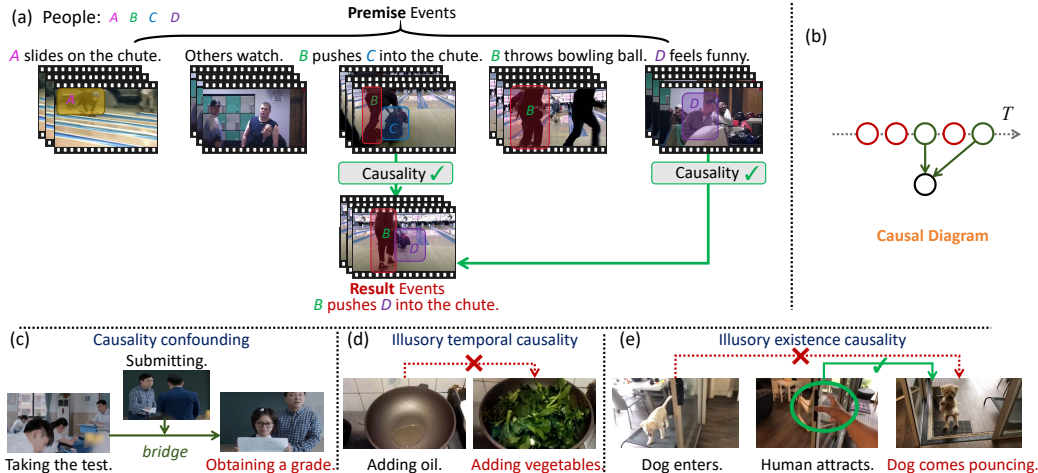


Figure 1: (a): Illustration of Multi-Event Causal Discovery Task, where the 3rd and 5th premise events account for the occurrence of the final event. The objective of our task is to determine whether a causal relation exists between events and outputs a structured causal diagram. (c): Example of causality confounding. (d)&(e): Illustration of illusory causality.

cannot provide a comprehensive and structured causal representation for multi-event video reasoning, which is typically required in real-world scenarios. For instance, in traffic surveillance videos, it is necessary to cross-analyze events happening at different times to determine which events, or combinations of events, led to the final traffic accident event.

To address this gap, we set up a new task: **Multi-Event Causal Discovery (MECD)**, which aims to uncover causal relationships among events that distribute chronologically in long videos. 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 associations between these events to derive a comprehensive and structured event-level causal diagram (Fig. 1(b)), indicating why and how the final result event happens. Meanwhile, we contribute a new dataset for the training and evaluation of MECD by collecting long-form videos involving multiple events and manually annotating real causal relationships between events for them. However, to our knowledge, no available solutions can directly comprehend causal relationships at the event level, necessitating the development of a new framework to tackle this complex task.

To this end, we draw inspiration from the *Granger Causality Method* [12–14] 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 relationships between events are disrupted or interfered with by other relay or adjacent events. Such confounding stems from the fact that many causal relationships flow through an intermediary event that acts as a bridge. As shown in Fig. 1(c), event "submitting the paper" serves as a necessary bridge between "taking the test" and "obtaining a grade." In this case, this bridge event might be mistakenly regarded as the only cause of the result event, while another cause, "taking the test," is overlooked. However, the bridge event can only occur if "taking the test" happens first. Therefore, we cannot identify the real causality between events that linked by such bridges following a simple predictive criterion, and eliminating such confounding is thus crucial for an accurate causal discovery.

(2) **Illusory Causality**, which includes 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. As shown in Fig. 1(d), the event "adding oil

when cooking" often occurs before "adding vegetables to stir-fry," but there is no real causality between them. 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 (Fig. 1(e)), consider determining the causal relation between "a large brown dog enters the room" (at the start of the video) and "the dog runs towards the camera." (at the end of the video). Although the presence of the dog in the former event is a prerequisite for the subsequent event, it does not directly cause the dog to rush towards the camera.

Building upon the preceding discussion, we introduce a novel framework to tackle MECD. This framework executes the *Event Granger Test* via an efficient mask-based event prediction model. It 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 [15–17] and counterfactual inference [15, 18, 19]—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 issue of illusory causality through the incorporation of an extra chain of thought [20–22] and existence-only descriptions during inference. Extensive experiments validate the effectiveness of our proposed framework in predicting structured causal relationships for given long-form videos.

2 Related Work

Video causal reasoning Many tasks in the past have tried to carry out causal reasoning in videos. Among these, the most common task is Video Question Answering (VQA) [1–4], aiming to give a reasonable answer to the question, methods such as SeViLA and LocAns [3, 4] made abductions based on the result, they grounded a single reason in previous time. However, VQA does not extend to abduct multiple reasons, merely creating a single causal link from reason to result.

Many tasks were based on VQA task for further causal reasoning attempts. CLEVRER [5], CATER [7] and V-CDN [6] explored causal reasoning based on physics and other basic laws in virtual scenes. However, these tasks haven't been committed to extending to the general video causal reasoning. AAR [8] and LMLN [9] symbolized data and derived inference rules using external knowledge. However, they can only reason within a defined symbol domain. The most similar VAR [10] predicted explanation events with premise events, and the causality was introduced during its prediction process. However, firstly it hasn't been committed to discovering the complete causal diagram. Besides, there is no explicit utilization of causal methods which constrains its ability.

All tasks above are for causal reasoning in short videos, while ours aims to handle long-duration videos. Besides, most of these are coarse video-level tasks, ours is more fine-grained event-level reasoning. Additionally, we want to establish a whole causal diagram rather than a single causal link. In conclusion, all these tasks haven't been committed to discovering causality among complex multi-event videos. Consequently, there exists a necessity need for a more comprehensive task.

Causal discovery in low-dimensional temporal data Traditional causal discovery methods of simple temporal data are mainly divided into three categories. Constraint-based methods use conditional independence tests to identify causal relations [23–25]. Score-based methods search through the space of all possible causal structures to optimize a specified metric [26–28]. The Granger Causality method discovers causal relations by calculating the degree to which the earlier occurred event contributes to the prediction of the latter occurred event [29–31]. The constraint-based and score-based methods require stringent assumptions about data distribution, making them less suitable for video data. The Granger Causality methods are more suitable yet face challenges when applied directly to video data, our method reaches better performance by utilizing causal inference methods.

3 Benchmark

3.1 MECD task settings

Our Multi-Event Causal Discovery (MECD) task is designed to test the ability of causal discovery in multi-event videos. Given a video \mathcal{E} that contains chronologically organized N events, $\mathbb{E} := \{e_1, \dots, e_N\}$, the task aims at determining whether any previous event e_n ($n < N$) has a causal

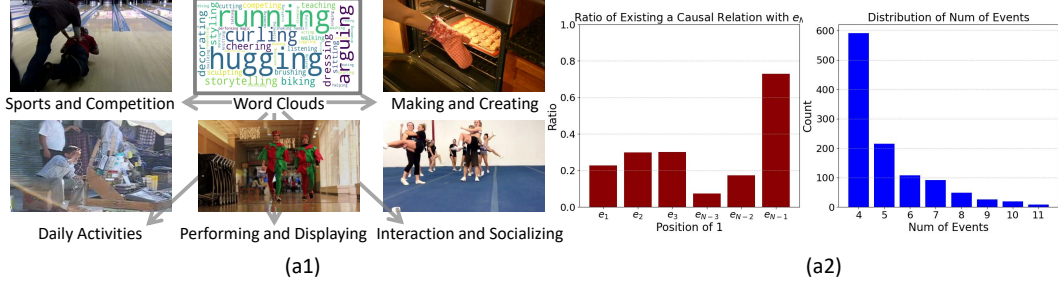


Figure 2: **Constitute of MECD dataset.** In (a1), we present 5 main video categories of the dataset. The word cloud is also summarized for video types. In (a2), the left chart indicates the impact of positions of events on their causality where we find the second last event tends to be more significant; while the right chart plots the number of events in videos.

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 $\mathbf{r} = [r_1, \dots, r_{N-1}]$, 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 is generalizable to causal relations of any of two events as long as we cut off the video and treat the latter one as the last event.

3.2 MECD task dataset

Data Source The Multi Events Causal Discovery (MECD) task contains videos with multiple events and intricate causal relationships. The ActivityNet Captions dataset [32] is built on ActivityNet v1.3 which includes 20k 120-second YouTube untrimmed videos. We carefully reorganize the ActivityNet Captions dataset and select 1,105 lifestyle videos that span diverse scenarios. We call this new dataset as MECD dataset, where 806 and 299 videos are randomly split for training and testing, respectively. 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 (a1) presents the main categories and word clouds of video types. Please refer to Appendix Sec. B.4 for more dataset examples.

Data Cleaning We further clean our dataset by excluding non-causal videos. For example, videos that describe multiple non-causal action steps such as washing hands and shaving were excluded.

Dataset Annotation The annotations of MECD dataset include 4 attributes. The “duration”, “sentence”, and “timestamps” attributes in annotations remain the same as the ActivityNet Captions dataset. Specifically, in the context of our task, a new attribute “relation” is introduced. To obtain this attribute, relations among events are firstly annotated by GPT-4 API [33], and subsequently refined by five human annotators. Through a cross-annotation process, gt labels are determined by the majority of the annotators’ causal relation choices, thus mitigating potential inaccuracies and subjective biases to a certain extent. We also present the impact of positions of events on their causality and number of events in videos in Fig. 2 (a2), annotation pipeline is in Appendix Sec. B.3.

4 Methodology

In this section, we present our Video Granger Causality Model (VGCM), as shown in Fig. 3. 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. In this context, we further utilize causal inference methods to address these by compensating or removing the effect of previous or subsequent causal events to mitigate the confounding meanwhile during inference the extra chain of thoughts and existence-only descriptions relieve the illusion.

4.1 VGCM: Video Granger Causality Model

Building upon the Granger Causality method introduced in [34–36], our core motivation for constructing VGCM is *Event Causality Test*: To compare the prediction result 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 premise event is causally related to the result event.

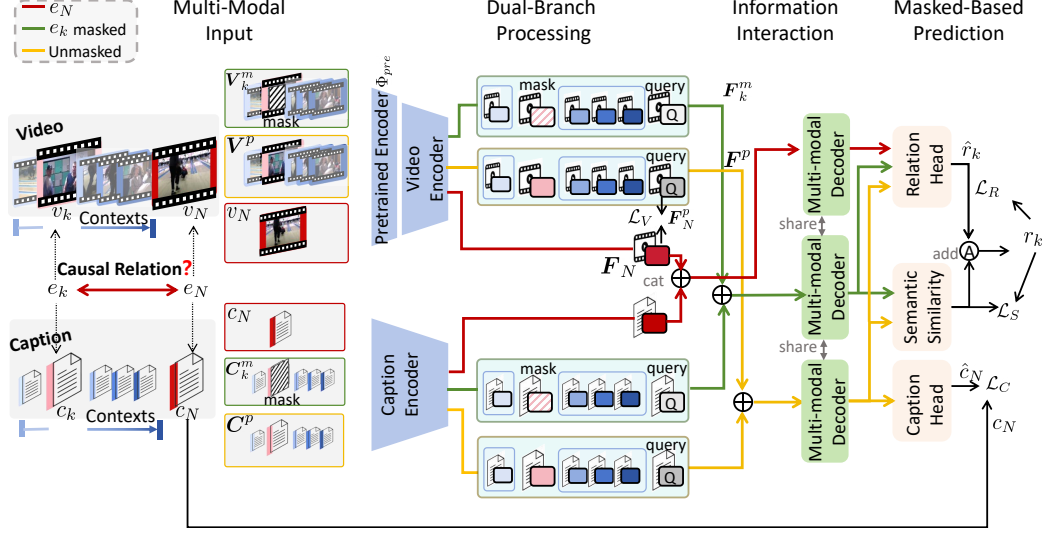


Figure 3: **Video Granger Causality Model.** Two data streams V^p and V^m serve as input, video and text embeddings are concatenated after being separately embedded. The VGCM incorporates two classifiers, the caption head takes the unmasked stream to accomplish the event-predicting task, while the relation head discovers the causal relations with two embedding streams.

We design VGCM to take in both the video clips and the captions to maximize information utilization. As illustrated in Fig. 3, 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 [10, 37–39], we firstly 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. 3) while the other for the set with one event (*i.e.*, e_k) masked \mathbb{E}_k^m (green path in Fig. 3). 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 .

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

where Enc_V and Enc_T 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 prediction. Additionally, in order to compare the difference of the prediction result, $\mathbf{O}^p, \mathbf{O}_k^m$ are directed to the relation head for causal relation discovery. 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 $\mathbf{O}_N = \text{Dec}(\text{Cat}(\mathbf{F}_N, \text{Enc}_C(C_N)))$, \mathbf{O}_N is aggregated for reasoning (red path in Fig. 3). 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:

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

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:

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

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 Nth slice of \mathbf{F}^p , which represents the feature of e_N .

4.2 Causal Inference in VGCM

In Sec. 4.1, 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 [40, 41] disentangled different attributes of a variable by supervising high-order distribution 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. 4 (a1)). **2)** Causal effects of e_{k+1} to e_N may be redundant, as e_k must be a necessary prior of its causality, (**red path** in Fig. 4 (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. 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 [20–22] 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 [33] 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. 3 the same as those in the

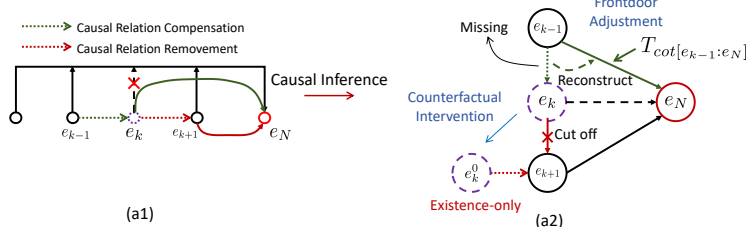


Figure 4: **Causal Effect of the Adjacent Events and Causality Diagram.** (a1) shows the causality of the third event analyzed, the red causal effect needs to be compensated while the blue needs to be mitigated. (a2) shows the causal inference methods corresponding to the two causal effects.

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 [33] 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 [42] for the missing causal effect of e_{k-1} and counterfactual inference [42] 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 diagram in Fig. 4 (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. 3, and $do(\cdot)$ represents do-operation in causal inference [15] 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 [43, 17], $P(e_N|do(e_k))$ 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 at the same time.

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.

Table 1: **Main results.** Experiments validate the effectiveness of our VGCM framework in providing causal relations in multi-event videos, outperforming GPT-4o and VideoLLaVA by 5.7% and 4.1%, respectively. * indicates the fine-tuned result, and ‡ indicates without causal inference.

	Paradigm	Method	Ave SHD ↓	Accuracy ↑
-	Random Guess	Guess all causal.	6.95	42.4
		Guess all non-causal.	5.36	57.6
Few-shot	LLM Base	Gemini-1.5-Pro [44]	4.91	59.3
		GPT-4-0613 [33]	4.92	59.6
	VLLM Base	MiniGPT4-video [45]	5.16	56.8
		MiniGPT-4 [46]	5.14	57.5
		Video-llama [47]	5.10	60.6
		VideoChat2 [48]	4.89	60.7
		VideoLLaVA [49]	4.85	62.5
GPT-4o [33]	4.69	65.5		
Fine-tuned	Multi-modal	VAR [10]	4.96	57.3
		VideoBERT [50]	4.95	60.9
		CLIP (ViT-L/14) [51]	4.77	62.9
	VLLM Base	VideoChat2* [48]	4.77	66.9
		VideoLLaVA* [49]	4.73	67.1
	Ours	VGCM‡	4.51	67.0
		VGCM	4.19	71.2
-	Humans	Deductive Reasoning	2.05	87.2

5 Experiments

5.1 Main results

Implementation details. including the pretraining process, detailed architecture of VGCM, and hyper-parameters settings can be found in Appendix Sec. A due to space constraints.

Baselines. We mainly compared our model with basic multi-modal models such as baseline model VideoBERT [50] and widely used CLIP-L [51] and the most similar reasoning model VAR [10]. Besides, we also conduct experiments on powerful LLM, including GPT-4 [33] and Gemini-Pro [44]. VLLM utilized for comparison includes widely accepted GPT-4o [33], VideoLLaVA [49], MiniGPT-4 [46], Video-llama [47], VideoChat2 [48] and MiniGPT4-video [45]. Specifically, LLMs and VLLMs are conducted under the few shot setting (In-Context Learning) following the causal discovery tasks in NLP [52–54], additionally, we reported the performance of fine-tuned VideoLLaVA and VideoChat2.

Metrics. We utilize the top-1 accuracy of the output causal relation chains with respect to the final event to evaluate the model’s capability in causal discovery. Although our VGCM is designed to discover the causal relations leading to the final event, when truncating the video during inference and redefining the final event as the new result, VGCM can generate a comprehensive causal diagram for the entire video without introducing additional training. Consequently, in addition to the primary metric accuracy, we introduce Structural Hamming Distance (SHD) [55, 56] as a supplementary metric. SHD measures the degree of matching between comprehensive causal graphs by summing the number of incorrect causal relations. In the MECDD test set, the average number of causal relations in video causal graphs is 12.31, and a lower Ave SHD value of the test set indicates better performance.

Results. We report the quantitative results in Tab. 1. Our VGCM without causal inference reaches an accuracy of 66.9%, demonstrating basic reasoning capabilities. Furtherly, the complete VGCM reaches a better performance with an accuracy of 71.2%, outperforming the GPT-4, GPT-4o, fine-tuned VideoLLaVA [49] by 11.6%, 5.7%, and 4.1%. Additionally, we explored the effect of altering the input format of the two modalities in Appendix Sec. C.1, indicating that VGCM is not dependent on the input format. The results compared with the two metrics indicate that for most models, accuracy is already adequate to represent their causal discovery capabilities. However, the additional metric Ave SHD indicates that Gemini and GPT-4 exhibit a superior overall capacity for discovering complete relations. An example of the output complete causal diagram is visualized in Figure 7.

GPT-4 [33] stands out as one of the most advanced LLM models, however, we found that even being provided with sufficient few-shot examples (detailed in Appendix Sec. C.2), its accuracy remains at

Table 2: **Ablation Study.** Adj indicates the front-door adjustment, and inter indicates the counterfactual intervention.

Base designs			Causal methods		Acc
\mathcal{L}_C	\mathcal{L}_V	\mathcal{L}_S	Adj	Inter	
✓	✓	✓			64.8
✓	✓	✓			65.1
✓	✓				65.3
✓	✓	✓			67.0
✓	✓	✓	✓		68.7
✓	✓	✓		✓	69.3
✓	✓	✓	✓	✓	71.2

Table 3: Illusory existence causality experiment. w/o C indicates without counterfactual intervention.

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

Table 4: Illusory temporal causality experiment. w/o F indicates without front-door adjustment and Ave indicates average.

Method	r_0 Acc	r_{N-1} Acc	Ave r Acc
VAR [10]	53.8 (-3.5)	54.6 (-3.7)	57.3
VGCM (w/o F)	63.6 (-3.3)	63.7 (-3.2)	66.9
VGCM	68.0 (-0.7)	68.4 (-0.3)	68.7

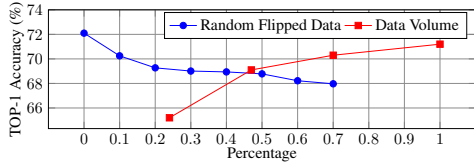


Figure 5: **Dataset robustness.** Accuracy decreases slightly when increasing noise, and increases slowly when increasing the training data.

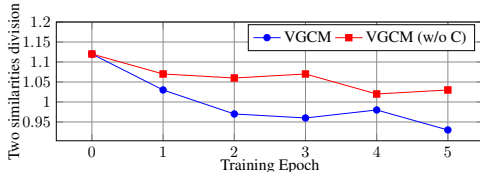


Figure 6: **Causality discovered analysis.** The similarity of masking causal premise events is obviously lower through counterfactual intervention.



Figure 7: Complete causal diagram.

Table 5: Open-set ability of VGCM.

Method	TOP-1 Accuracy
VAR [10]	54.8
VGCM [‡]	59.2
VGCM	64.4

only 59.6%. Possible explanations may be due to task contamination [57], GPT-4 mainly performs well on datasets released before the training date, while our task is novel. Moreover, other reasons may include the causal hallucination problem of establishing a threshold for differentiating between scenarios with and without causality [58]. For further insights into GPT-4’s failure cases, refer to Appendix Sec. B.2.

As illustrated in Tab. 6, we have assessed the inference speed of various models, with our VGCM achieving a swift 0.76 seconds per sample. The proposed method incurs an overhead of only 8.57% over the Videobert baseline. It is noteworthy that our inference speed is 3 to 6 times faster than that of all Video LLMs. The inference speed experiments were conducted on 1 NVIDIA A6000 GPU.

5.2 Ablation Study

Video Granger Causality Model design We designed our causal discovery model based on the Granger Causality, three auxiliary losses are applied. The performances in Tab. 2 indicate that our VGCM benefits from the design of \mathcal{L}_V and \mathcal{L}_C , for they support our method of inferring causal relations by facilitating the model with event prediction ability. \mathcal{L}_S also benefits our model by supervising the causal feature similarity of e_N with and without non-causal event e_k masked.

Front-door adjustment with chain of thoughts candidate The method does improve reasoning ability in Tab. 2. We conduct an experiment in Tab. 4 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 with chain of thought candidate 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, visualization can be found in Fig. 8 Example 1.

Counterfactual intervention with existence-only descriptions The performance in Tab. 2 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 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 Tab. 3 and Fig. 6. 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 existence illusory is solved with counterfactual intervention for the reason that the division is below 1 of VGCM, example visualization can be found in Fig. 8 Example 2.

5.3 Robustness Analysis

Model Robustness To prove our model’s robust reasoning ability, we split the MECD dataset into five categories, and conduct an experiment similar to the open-set setting with cross-validation. VGCM

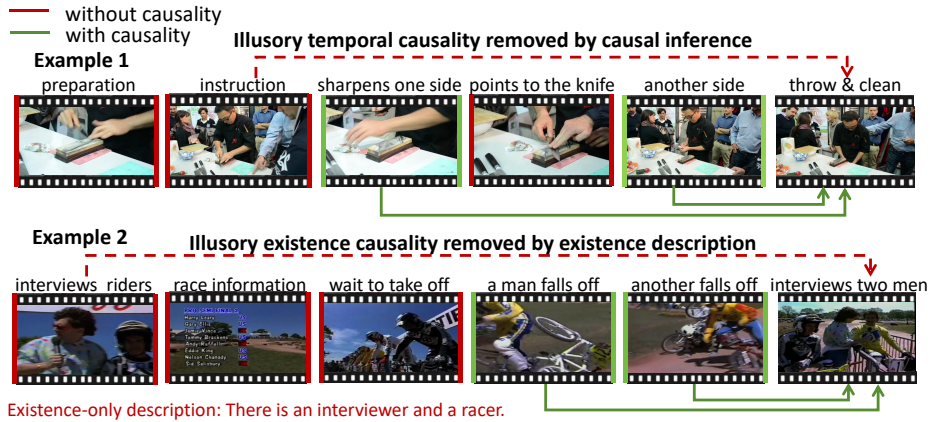


Figure 8: **Successful abduction examples of our VGCM.** Results indicate that after utilizing causal inference methods, illusory causality is suppressed and robust abduction ability is facilitated.

Table 6: **Inference speed.** Our VGCM is 3-6 times faster than all Video LLMs while slightly slower than the baseline.

Model	Inference Speed
Videobert [50]	0.70
Our VGCM	0.76
VideoLLaVA [49]	2.12
VideoChat2 [48]	2.96
MiniGPT4-video [45]	3.98
MiniGPT-4 [46]	4.72

Table 7: **Model’s generalizability test.** Higher VQA Acc and VQA Score are reached when prompted with causal relations from our VGCM.

Output Causal Relations	VQA Acc	VQA Score
w/o (Standard QA setting for VLLMs)	43.17	2.82
w Gemini-Pro [44]	49.10	2.90
w GPT-4 [33]	49.36	2.89
w VideoChat2 [48]	51.01	2.95
w VideoLLaVA [49]	51.88	2.93
w Our VGCM	62.21	3.12

reaches an average accuracy of 64.4%, outperforms VGCM without causal inference and VAR by 5.2% and 9.6%, details can be found in Appendix Sec. D.1.

Moreover, to further validate the generalization capabilities of our model, we evaluate the quality of output causal relations on a related and representative video reasoning task: Video Question Answering (VQA) as shown in Tab. 7. Specifically, during inference on the multi-event subset of ActivityNet-QA [59] (The part that overlaps with the MECD test set), we prompted MiniGPT4-video [45] with additional causal relations outputs alongside the standard question inputs. This paradigm facilitates the VLLMs in considering the task from a causal perspective. As shown in the table below, when prompted with these additional causal relations, the answering accuracy of MiniGPT4-video [45] improved by our VGCM surpasses other strong VLLMs like VideoChat2 [48]. These findings confirm that our model can provide accurate causal perception for videos, significantly improving performance on related video reasoning tasks.

Dataset Robustness We study the subjectivity and data volume of our proposed MECD dataset, which is shown in Tab. 5. 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 data, the increment from 600 examples to 806 examples yields a very modest improvement, indicating the adequacy of our dataset.

6 Conclusion

We proposed a novel task, multi-event video causal discovery (MECD), which focuses on event-level causal discovery in long-term videos. Besides, we built the MECD dataset with long-term daily life video datasets with causal relations to support this task and proposed the first video events causal discovery framework VGCM in the principles of Granger Causality. Additionally, our proposed VGCM was facilitated with deeper reasoning ability through causal inference with the chain of thoughts and existence-only descriptions. Our VGCM significantly outperforms GPT-4o and VideoLLaVA by 5.7% and 4.1%, respectively, demonstrating its robust reasoning ability.

7 Acknowledgement

The paper is supported in part by the National Natural Science Foundation of China (No. 62325109, U21B2013) and the Lenovo Academic Collaboration Project.

References

- [1] Difei Gao, Luwei Zhou, Lei Ji, Linchao Zhu, Yi Yang, and Mike Zheng Shou. Mist: Multi-modal iterative spatial-temporal transformer for long-form video question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14773–14783, 2023.
- [2] Ruoyue Shen, Nakamasa Inoue, and Koichi Shinoda. Text-guided object detector for multi-modal video question answering. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1032–1042, 2023.
- [3] Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. *arXiv preprint arXiv:2305.06988*, 2023.
- [4] Tianwen Qian, Ran Cui, Jingjing Chen, Pai Peng, Xiaowei Guo, and Yu-Gang Jiang. Locate before answering: Answer guided question localization for video question answering. *IEEE Transactions on Multimedia*, 2023.
- [5] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Clevrer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019.
- [6] Yunzhu Li, Antonio Torralba, Anima Anandkumar, Dieter Fox, and Animesh Garg. Causal discovery in physical systems from videos. *Advances in Neural Information Processing Systems*, 33:9180–9192, 2020.
- [7] Rohit Girdhar and Deva Ramanan. Cater: A diagnostic dataset for compositional actions and temporal reasoning. *arXiv preprint arXiv:1910.04744*, 2019.
- [8] Tao Zhuo, Zhiyong Cheng, Peng Zhang, Yongkang Wong, and Mohan Kankanhalli. Explainable video action reasoning via prior knowledge and state transitions. In *Proceedings of the 27th acm international conference on multimedia*, pages 521–529, 2019.
- [9] Yang Jin, Linchao Zhu, and Yadong Mu. Complex video action reasoning via learnable markov logic network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3242–3251, 2022.
- [10] Chen Liang, Wenguan Wang, Tianfei Zhou, and Yi Yang. Visual abductive reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15565–15575, 2022.
- [11] Clement Tan, Chai Kiat Yeo, Cheston Tan, and Basura Fernando. Abductive action inference. *arXiv preprint arXiv:2210.13984*, 2022.
- [12] Anil Seth. Granger causality. *Scholarpedia*, 2(7):1667, 2007.
- [13] Mariusz Maziarz. A review of the granger-causality fallacy. *The journal of philosophical economics: Reflections on economic and social issues*, 8(2):86–105, 2015.
- [14] Ali Shojaie and Emily B Fox. Granger causality: A review and recent advances. *Annual Review of Statistics and Its Application*, 9:289–319, 2022.
- [15] Judea Pearl. Causal inference. *Causality: objectives and assessment*, pages 39–58, 2010.
- [16] Xu Yang, Hanwang Zhang, and Jianfei Cai. Deconfounded image captioning: A causal retrospect. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):12996–13010, 2021.

- [17] Xu Yang, Hanwang Zhang, Guojun Qi, and Jianfei Cai. Causal attention for vision-language tasks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9847–9857, 2021.
- [18] Long Chen, Yuhang Zheng, Yulei Niu, Hanwang Zhang, and Jun Xiao. Counterfactual samples synthesizing and training for robust visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [19] Wenjie Wang, Fuli Feng, Xiangnan He, Hanwang Zhang, and Tat-Seng Chua. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1288–1297, 2021.
- [20] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- [21] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022.
- [22] Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. A survey of chain of thought reasoning: Advances, frontiers and future. *arXiv preprint arXiv:2309.15402*, 2023.
- [23] Andreas Gerhardus and Jakob Runge. High-recall causal discovery for autocorrelated time series with latent confounders. *Advances in Neural Information Processing Systems*, 33:12615–12625, 2020.
- [24] Charles K Assaad, Emilie Devijver, and Eric Gaussier. Discovery of extended summary graphs in time series. In *Uncertainty in Artificial Intelligence*, pages 96–106. PMLR, 2022.
- [25] Joris M Mooij, Sara Magliacane, and Tom Claassen. Joint causal inference from multiple contexts. *The Journal of Machine Learning Research*, 21(1):3919–4026, 2020.
- [26] Xiangyu Sun, Oliver Schulte, Guiliang Liu, and Pascal Poupart. Nts-notears: Learning nonparametric dbns with prior knowledge. *arXiv preprint arXiv:2109.04286*, 2021.
- [27] Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. Dynotears: Structure learning from time-series data. In *International Conference on Artificial Intelligence and Statistics*, pages 1595–1605. PMLR, 2020.
- [28] Tian Gao, Debarun Bhattacharjya, Elliot Nelson, Miao Liu, and Yue Yu. Idyno: Learning nonparametric dags from interventional dynamic data. In *International Conference on Machine Learning*, pages 6988–7001. PMLR, 2022.
- [29] Ruichu Cai, Siyu Wu, Jie Qiao, Zhifeng Hao, Keli Zhang, and Xi Zhang. Thp: Topological hawkes processes for learning granger causality on event sequences. *arXiv preprint arXiv:2105.10884*, 2021.
- [30] Wei Chen, Jibin Chen, Ruichu Cai, Yuequn Liu, and Zhifeng Hao. Learning granger causality for non-stationary hawkes processes. *Neurocomputing*, 468:22–32, 2022.
- [31] Tsuyoshi Idé, Georgios Kollias, Dzung Phan, and Naoki Abe. Cardinality-regularized hawkes-granger model. *Advances in Neural Information Processing Systems*, 34:2682–2694, 2021.
- [32] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Nieves. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715, 2017.
- [33] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

- [34] David Lopez-Paz, Krikamol Muandet, and Benjamin Recht. The randomized causation coefficient. *J. Mach. Learn. Res.*, 16:2901–2907, 2015.
- [35] Jean-François Ton, Dino Sejdinovic, and Kenji Fukumizu. Meta learning for causal direction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9897–9905, 2021.
- [36] Hebi Li, Qi Xiao, and Jin Tian. Supervised whole dag causal discovery. *arXiv preprint arXiv:2006.04697*, 2020.
- [37] Jie Lei, Liwei Wang, Yelong Shen, Dong Yu, Tamara L Berg, and Mohit Bansal. Mart: Memory-augmented recurrent transformer for coherent video paragraph captioning. *arXiv preprint arXiv:2005.05402*, 2020.
- [38] Teng Wang, Ruimao Zhang, Zhichao Lu, Feng Zheng, Ran Cheng, and Ping Luo. End-to-end dense video captioning with parallel decoding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6847–6857, 2021.
- [39] Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, and Zheng-Jun Zha. Object relational graph with teacher-recommended learning for video captioning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13278–13288, 2020.
- [40] Aapo Hyvarinen and Hiroshi Morioka. Nonlinear ica of temporally dependent stationary sources. In *Artificial Intelligence and Statistics*, pages 460–469. PMLR, 2017.
- [41] Aapo Hyvarinen and Hiroshi Morioka. Unsupervised feature extraction by time-contrastive learning and nonlinear ica. *Advances in neural information processing systems*, 29, 2016.
- [42] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [43] Yang Liu, Guanbin Li, and Liang Lin. Cross-modal causal relational reasoning for event-level visual question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [44] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [45] Kirolos Ataallah, Xiaoqian Shen, Eslam Abdelrahman, Essam Sleiman, Deyao Zhu, Jian Ding, and Mohamed Elhoseiny. Minigt4-video: Advancing multimodal llms for video understanding with interleaved visual-textual tokens. *arXiv preprint arXiv:2404.03413*, 2024.
- [46] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.
- [47] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023.
- [48] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22195–22206, 2024.
- [49] Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. *arXiv preprint arXiv:2311.10122*, 2023.
- [50] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473, 2019.
- [51] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.

- [52] Xinyi Wang, Wanrong Zhu, Michael Saxon, Mark Steyvers, and William Yang Wang. Large language models are latent variable models: Explaining and finding good demonstrations for in-context learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- [53] Kun Luo, Tong Zhou, Yubo Chen, Jun Zhao, and Kang Liu. Open event causality extraction by the assistance of llm in task annotation, dataset, and method. In *Proceedings of the Workshop: Bridging Neurons and Symbols for Natural Language Processing and Knowledge Graphs Reasoning (NeusymBridge)@ LREC-COLING-2024*, pages 33–44, 2024.
- [54] Aniket Vashishtha, Abbavaram Gowtham Reddy, Abhinav Kumar, Saketh Bachu, Vineeth N Balasubramanian, and Amit Sharma. Causal inference using llm-guided discovery. *arXiv preprint arXiv:2310.15117*, 2023.
- [55] Verónica Rodríguez-López and Luis Enrique Sucar. Knowledge transfer for causal discovery. *International Journal of Approximate Reasoning*, 143:1–25, 2022.
- [56] Konstantina Biza, Ioannis Tsamardinos, and Sofia Triantafillou. Tuning causal discovery algorithms. In *International Conference on Probabilistic Graphical Models*, pages 17–28. PMLR, 2020.
- [57] Changmao Li and Jeffrey Flanigan. Task contamination: Language models may not be few-shot anymore. *arXiv preprint arXiv:2312.16337*, 2023.
- [58] SM Tonmoy, SM Zaman, Vinija Jain, Anku Rani, Vipula Rawte, Aman Chadha, and Amitava Das. A comprehensive survey of hallucination mitigation techniques in large language models. *arXiv preprint arXiv:2401.01313*, 2024.
- [59] Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. Activitynet-qa: A dataset for understanding complex web videos via question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9127–9134, 2019.
- [60] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [61] Chang Gong, Di Yao, Chuzhe Zhang, Wenbin Li, and Jingping Bi. Causal discovery from temporal data: An overview and new perspectives. *arXiv preprint arXiv:2303.10112*, 2023.
- [62] Blai Meléndez Catalán, Emilio Molina, and Emilia Gómez Gutiérrez. Bat: An open-source, web-based audio events annotation tool. 2017.
- [63] Weiyao Lin, Huabin Liu, Shizhan Liu, Yuxi Li, Hongkai Xiong, Guojun Qi, and Nicu Sebe. Hieve: A large-scale benchmark for human-centric video analysis in complex events. *International Journal of Computer Vision*, 131(11):2994–3018, 2023.
- [64] Renáta Németh, Domonkos Sik, and Fanni Máté. Machine learning of concepts hard even for humans: The case of online depression forums. *International Journal of Qualitative Methods*, 19:1609406920949338, 2020.
- [65] Sungmin Kang, Juyeon Yoon, and Shin Yoo. Large language models are few-shot testers: Exploring llm-based general bug reproduction. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pages 2312–2323. IEEE, 2023.
- [66] Toufique Ahmed and Premkumar Devanbu. Few-shot training llms for project-specific code-summarization. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–5, 2022.

Appendix

A Implementation details

Pretraining process For each video event, visual features are extracted using ActivityNet pretrained ResNet200 [60], following [10, 37–39]. Prior domain knowledge could benefit the Granger Causality Causal discovery method [61], so we fully pre-trained our model for the dense video captioning task on a 3.1k ActivityNet Captioning video dataset, each video sample contains more than 4 events.

Training set All the experiments are conducted on 1 NVIDIA A40 GPU. We train our model for 20 epochs with a learning rate of $16e-5$ about 6 hours. Our optimizer is consistent with BertAdam [50] optimizer, with 3 epochs of warm-up. The open-set experiment set can be found in Appendix Sec. D.1. We report the average results during all experiments under three random seeds (2023, 2024, 2025). The ablation of two modalities can be found in Appendix Sec. C.1.

Model details Our encoder Enc_V , Enc_C , and multi-modal video decoder Dec are built upon Videobert [50], a joint model for video and language representation learning. The details of the GPT-4 API prompt can be found in Sec. D.2 in the Appendix.

Hyperparameters 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.

Implementation of VAR** We migrate the VAR to our task through an effective method: We mask any event e_k , ($k < N$), and then utilize the fully trained VAR to perform event prediction of e_k . If the prediction results \hat{e}_k is obviously various from e_k , it is considered that the event e_k is non-causal. Then r_k is labeled as 0; in the opposite case, r_k is labeled as 1. We also report the average results of VAR under three random seeds (2023, 2024, 2025).

Implementation of LLMs As for GPT-4 and Gemini-Pro, We report the average results of three calls.

Implementation of VLLMs We report the average results of VLLMs under three random seeds (2023, 2024, 2025). When VLLMs do not output r in the required format, we order them to re-answer until the outputs match the format to measure their best performance.

B Additional Visualization

B.1 Successful abduction examples of our VGCM

In Fig. 9, additional examples are presented to showcase the performance of our VGCM, particularly excelling in complex abduction scenarios. The first example successfully discovers that there is no causal relation between “*We see the targets in front of a backdrop.*” and “*The instructor walks over to the targets.*”, despite the backdrop being a necessary object of the result event. This abduction avoids the illusory existence causality.

The second example successfully discovers that there is no causal relation between “*The video shows different cricket matches taking place where Sri Lanka is playing against teams from different countries.*” and “*The stadium is filled with spectators cheering for the cricketers.*”, despite the spectators’ cheering often happening after the game playing. This abduction avoids the illusory temporal causality. Both instances align with the foundational principles motivating our method design.

The third example shows the 83.3% accuracy of video causal relations abduction. Notably, it correctly discerns the most complex causal relations, however, it fails to realize that person B doesn’t hit the tennis ball can contribute to the anticipation of the result event of continuous hitting. This indicates that VGCM might still require refinement in understanding causality within higher-level semantics, especially in the mining of some obscure mental or emotional influences. We will strive to explore further solutions in the follow-up work.

B.2 Failure abduction examples of GPT-4

While examining the causal discovery results of GPT-4, we encountered some intriguing observations. In the example presented in Fig. 10, the GPT-4 API incorrectly infers that all premise events have a causal relation with the result event of the team winning. However, the initial appearance of the team

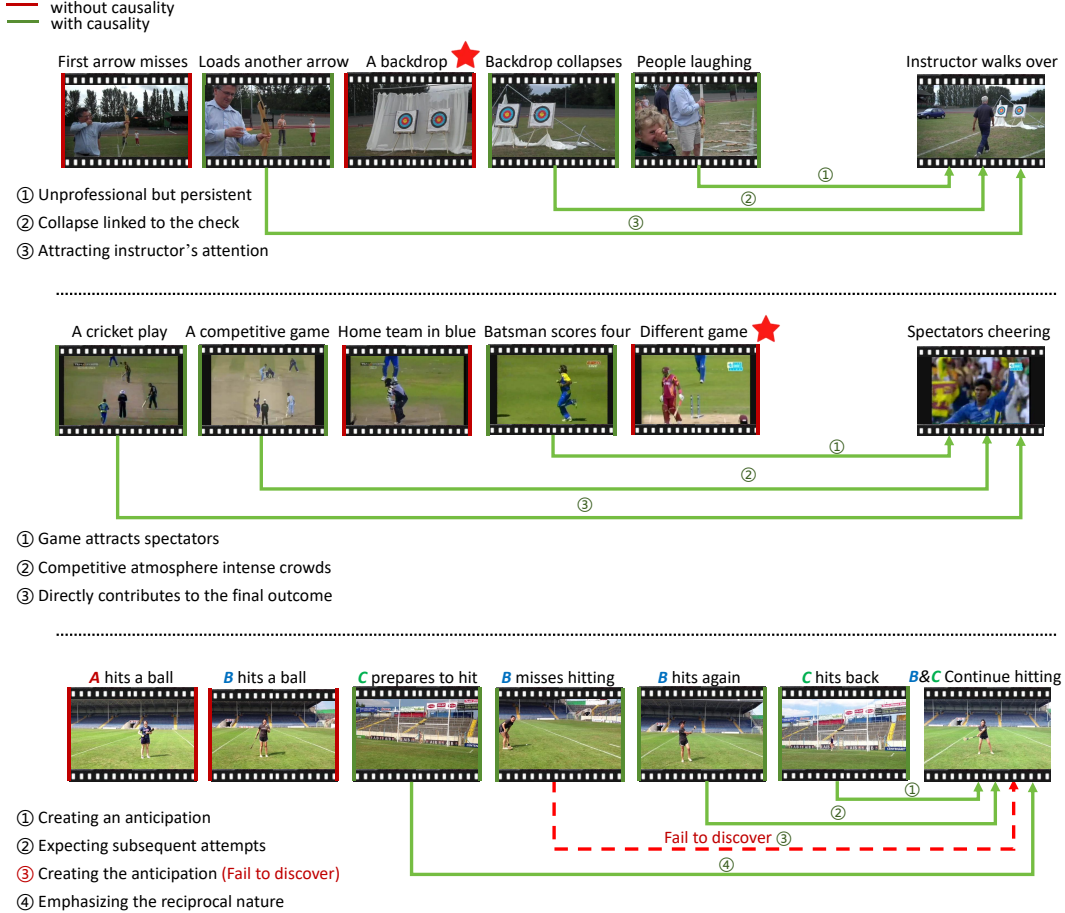


Figure 9: **More successful abduction examples of our proposed VGCM.** The relation which reveals our method of eliminating illusory causality is marked by a red five-pointed star ★. The failure case is annotated in a red dotted line - - - -.

does not directly lead to their victory, and the subsequent celebrations also lack any causal links with the outcome. Indeed, the false discovery by the GPT-4 API could stem from the illusion of causality, where the team’s mere presence is perceived as a necessary condition for the outcome. Additionally, the illusion of temporal causality may also play a role, as statistics indicate that celebrations often occur before the announcement of the competition winner. These cognitive biases could contribute to the erroneous causal inference made by the GPT-4 API in this scenario.

When we request a detailed explanation from the GPT-4 API regarding the discovered causal relation between the result event and the initial appearance of the team, the response is “*Setting up the motive for the last event.*” Obviously, the GPT-4 confuses causality with the illusion of existence causality. In contrast, our VGCM makes a correct inference in this scenario. Furthermore, when we seek detailed reasons from the GPT-4 API for the discovered causal relations between the result event and the celebrations, the answer is “*Indicating their satisfaction and confidence in their performance, implying they believe they have a good chance to win.*” Here, the GPT-4 API misinterprets causality by associating it with the expression of subjective emotions unrelated to the events in question. It may mistake the display of subjective emotions for the presence of objectively implied causality.

B.3 Annotation pipeline of MECD dataset

To improve the accuracy and mitigate subjective biases in annotating causal relations, we employ a cross-annotation strategy [62–64]. The interactive interface used by annotators during the labeling process of our MECD dataset is illustrated in Fig. 11. Each video example is endowed with a

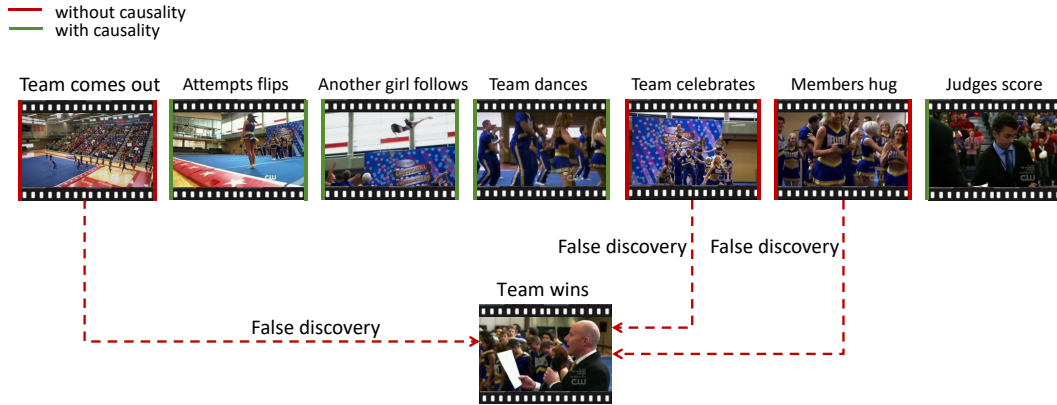




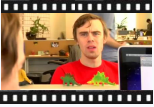



Figure 10: **Failure abduction examples of GPT-4.** Many failure cases of GPT’s causal reasoning are due to confusion with illusions and the conflation of subjective emotions with objective laws.

Annotator 4’s view:

Videos:

Captions:

1. "man is sitting on a table eating a sandwich in a restaurant talking to other man.",
2. " man is watching the man while eats and talking to him.",
3. " both men stands and play rock paper scissors and man hits the man sand this falls to the floor.",
4. " man is sitting on the chair and talks to the man on the floor and throw him the sandwich.",
5. " anoher man siting in front talk to the man in the desk.",
6. " man stands and talk to the man on the floor and hits him again and the man spit blood."

Annotated reference:

GPT-4 Relation: 00111


Annotator 1 Relation: 01111

Annotator 2 Relation: 01111

Annotator 3 Relation: 01111

Annotator 5 Relation: 11111

Tips:

1. Illusory existence causality and illusory temporal causality should be avoided.
2. Click  to play the video.

Your Annotation: _____ (0 for without causality, 1 for with causality)

Figure 11: **Annotation pipeline of MECD dataset.** Illustration of the interactive interface used by annotators during the labeling process of our MECD dataset. Key information is provided during annotation.

“relation” attribute. First, GPT-4 [33] provides an initial annotation of attribution, which is then further refined by five human annotators. Ground truth labels are determined based on the majority choices of the annotators regarding causal relations. This methodology ensures the creation of a more reliable and objective dataset.

B.4 Annotation examples of MECD

Annotation examples of MECD are shown in Fig. 12, our MECD dataset is carefully annotated to support the challenging task proposed with complete premise information.



Figure 12: **Annotation examples of MECD.** Annotation examples of MECD are shown. Newly annotated attributes “Relations”, “COT”, “Existence” and the existing caption attribute “Sentences” are shown along with the video frames.

C Additional Experiments

C.1 Modalities analysis of causality discovering

The MECD task employs both video input and corresponding captions to uncover causality. Our objective in this experiment is to evaluate the degree of reliance on these two modalities in causal discovery. Typically, each event in our MECD task consists of a textual input with an average of 13.5 words caption 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

Table 8: **VGCM performance with masked premise event caption input.** * indicates 30 frames masked at the same time.

Num of words masked	Accuracy
non-masked	71.2
2 per event	70.2
5 per event	69.7
8 per event	69.2
8 per event*	67.4
11 per event	68.9

Table 9: **VGCM performance with masked premise event visual input.** * indicates 10 words masked at the same time.

Num of frames masked	Accuracy
non-masked	71.2
5 per event	70.3
15 per event	69.0
20 per event	68.3
20 per event*	67.1
40 per event	67.9

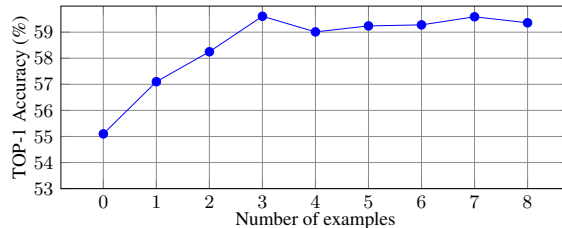


Figure 13: **The trend chart of inference accuracy as the number of examples changes under the In Context Learning paradigm.** Accuracy increases slightly when increasing the number of few-shot examples, when the number of examples > 3, the accuracy tends to remain constant.

in Tab. 8 indicate that our VGCM does not rely on the textual modality input; VGCM can also conduct causal discovery for videos without any captions.

In contrast, the experimental results suggest a more obvious performance decrease towards less visual modality input in the causality discovery task, as shown in Tab. 9. However, even with 80% masking of either modality, the results consistently outperform our strong baseline model, VideoLLaVA, underscoring the robust causal discovery capability of VGCM.

Furthermore, we conducted experiments involving simultaneous masking of both modalities of information. Interestingly, we observe a noticeable decrease in accuracy compared to when only one modality is masked. This observation highlights the importance of jointly considering both modalities in the causality discovery task.

C.2 Adequacy of the prompts provided to GPT-4

To delve deeper into the limitations of the straightforward baseline approach of prompting GPT-4, we examined the correlation between its accuracy and the number of video examples provided in the few-shot prompts. The findings, illustrated in Fig. 13, suggest that increasing the number of examples shown to GPT-4 does not effectively enhance its accuracy. This suggests that the limitation of the GPT-4 baseline is not strongly correlated with the number of presented examples but rather is more attributable to its intrinsic limitation in understanding complex causal relationships solely through text modality.

D Experiments details

D.1 Details of causality discovery experiment

In the open-set experiment of exploring reasoning ability, the five categories mainly consist of the activities below, demonstrating the colorful daily activities included in our dataset.

Sports: Arm wrestling, BMX, Beach soccer, Blow-drying hair, Capoeira, Croquet, Futsal, Ice fishing, Kite flying, Playing beach volleyball

Creating & Making: Assembling bicycle, Baking cookies, Building sandcastles, Carving jack-o-lanterns, Decorating the Christmas tree, Hanging wallpaper, Making a cake, Making an omelet, Painting fence, Putting in contact lenses

Daily Activities: Changing car wheel, Cleaning sink, Drinking coffee, Eating ice cream, Gargling mouthwash, Hanging wallpaper, Kneeling, Peeling potatoes, Putting on shoes, Washing face

Performing: Baton twirling, Bullfighting, Drum corps, Fun sliding down, Hula hoop, Playing congas, Playing drums, Playing rubik cube, Playing saxophone, Tumbling

Socializing: Beer pong, Playing blackjack, Playing field hockey, Playing harmonica, Playing piano, Playing squash, Playing water polo, Rock climbing, Smoking hookah, Belly dance

D.2 Prompts to generate auxiliary premise information

In this section, we introduce the detailed method of prompting GPT-4 [33] to generate more premise information. Firstly, we prompt the GPT-4 with the following prompts to generate the description-only sentences.

```
# Task: Each input consists of n sentences, and the text description
of each sentence has been given correspondingly (separated by " ",).
You need to offer the existence description of each sentence.
```

Besides the task description, we further append the few-shot paradigm (In-Context Learning) introduced in [52–54, 65, 66]. Similarly, we prompt the GPT-4 [33] with the following instructions to generate the chain-of-thoughts candidate sentences in the same few-shot paradigm.

```
# Task: Each video consists of n events and the text description
of each event has been given correspondingly separated by " ",).
First n-1 events might be the cause of the last event. You need to
offer the chain of thoughts you derive that causes the last event.
```

D.3 Chain of thoughts examples

In this section, we present an example of the chain of thoughts prompted, the corresponding premise event and result event descriptions are also shown below:

```
{"premise event sentence": "He continues sharpening the knife, turn it again
to further sharpen the other side and wipe it with paper towel."}
```

```
"result event sentence": "Throws the old and dirty paper towel and reach the
roll of paper towel and clean the knife."
```

```
"COT": "The repeated action of sharpening and wiping the knife underscores
the importance of both the knife's sharpness and cleaners, leading directly
to the final action of disposing of the used paper towel and getting a new
one to ensure the knife is thoroughly clean"}
```

The chain of thoughts shown above provides a logical causal chain between the event of the cleaning of the knife and the subsequent throwing of the dirty paper towel. The reasoning initiates by considering the heightened need for sharp and pristine knives achieved through sharpening. This causal chain is then expanded by suggesting that this demand could have led to the replacement of the paper towel. The chain of thoughts generated from GPT-4 serves as a candidate in the process of correct reasoning, contributing to the exploration of potential causal relations.

E Limitations

1. The video we input for causal discovery needs to provide timestamps, we encourage future work to realize causal discovery with weakly annotated inputs.

2. VGCM might still require refinement in understanding causality within higher-level semantics, especially in the mining of some obscure mental or emotional influences according to the failure cases analysis in Appendix Sec. B.1.
3. VGCM is based on the supervised paradigm of causal discovery, subsequent works may be able to extend to the unsupervised paradigm.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The main contributions and scope are summarized in Sec. 1.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Limitations are discussed in Sec. E in the Appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: We have already provided the proof of theoretical results in Sec. 5.1, Sec. 5.2 and Sec. 5.3.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Implementation details are provided in Sec. A in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Codes and datasets are released on the GitHub Page.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Implementation details are provided in Sec. A in the Appendix.

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We report the average results under three random seeds (2023, 2024, 2025).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Implementation details are provided in Sec. A in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We conducted the research in the paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There is no societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: ActivityNet Captions Dataset is with no license needed.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: Implementation details and limitations are provided in Sec. **A** and Sec. **E** in the Appendix.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.