

In-context Contrastive Learning for Event Causality Identification

Chao Liang¹ Wei Xiang² Bang Wang^{1*}

¹ School of Electronic Information and Communications,
Huazhong University of Science and Technology, Wuhan, China
{liangchao111, wangbang}@hust.edu.cn

² Faculty of Artificial Intelligence in Education,
Central China Normal University, Wuhan, China.
xiangwei@ccnu.edu.cn

Abstract

Event Causality Identification (ECI) aims at determining the existence of a causal relation between two events. Although recent prompt learning-based approaches have shown promising improvements on the ECI task, their performance are often subject to the delicate design of multiple prompts and the positive correlations between the main task and derivate tasks. The in-context learning paradigm provides explicit guidance for label prediction in the prompt learning paradigm, alleviating its reliance on complex prompts and derivative tasks. However, it does not distinguish between positive and negative demonstrations for analogy learning. Motivated from such considerations, this paper proposes an **In-Context Contrastive Learning (ICCL)** model that utilizes contrastive learning to enhance the effectiveness of both positive and negative demonstrations. Additionally, we apply contrastive learning to event pairs to better facilitate event causality identification. Our ICCL is evaluated on the widely used corpora, including the EventStoryLine and Causal-TimeBank, and results show significant performance improvements over the state-of-the-art algorithms.¹

1 Introduction

Event Causality Identification (ECI) is to detect whether there exists a causal relation between two event mentions in a document. It is of great importance for many Natural Language Processing (NLP) applications, such as question answer (Breja and Jain, 2020), machine reading comprehension (Berrant et al., 2014), and etc. Furthermore, It also has many practical applications in real-world scenarios, such as event prediction (Preethi et al., 2015; Radinsky et al., 2012) and strategy optimization (Balgı et al., 2022). Fig. 1 illustrates an event causality

* Corresponding author: Bang Wang

¹ We release the code at: <https://github.com/ChaoLiang-HUST/ICCL>.

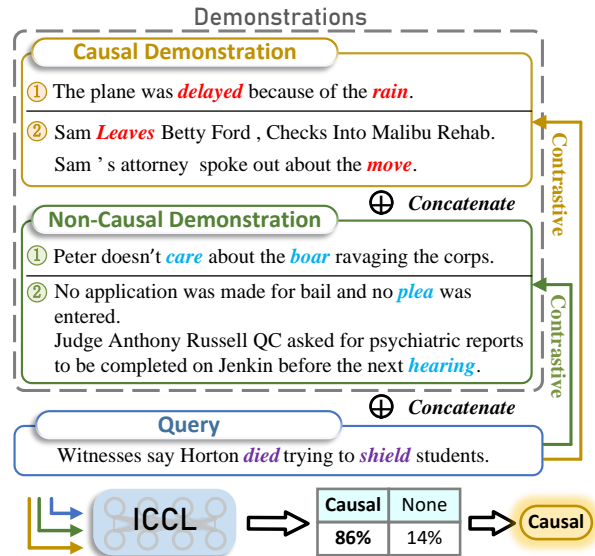


Figure 1: Illustration of our motivation. The event pairs are highlighted in different colors.

example from the Event StoryLine Corpus (ESC). We concatenated two causal demonstrations and two non-causal demonstrations before the query to be predicted, and enhanced the analogy between the query and demonstrations through contrastive. Ultimately, our ICCL model determined the causality between the two events, "*died*" and "*shield*", in the query.

Some graph-based methods have been proposed for the ECI task (Zhao et al., 2021; Phu and Nguyen, 2021; Pu et al., 2023), which apply a graph structure to represent events and their potential relations. For example, Zhao et al. (2021) initialize event nodes' embeddings using a document-level encoder and employ a graph inference mechanism to update their embeddings. Pu et al. (2023) incorporate causal label information and event pair interaction information to enhance the representation learning for event nodes in the graph. These methods follow the traditional representation learning for classification yet on a graph structure.

Recently the prompt learning paradigm (Liu

et al., 2023) has shown its great successes in many NLP tasks, as it can well leverage the potentials of a pre-trained language model (PLM). Some researchers have applied the prompt learning for the ECI task (Liu et al., 2021b; Shen et al., 2022). For example, the DPJL model (Shen et al., 2022) designs a main cloze task but also designs two derivative prompt tasks. Although the DPJL has achieved new state-of-the-art performance, it involves the delicate design of multiple prompts and relies on the positive correlations between the main task and derivative tasks.

The in-context learning paradigm (Dong et al., 2022) includes some demonstrations with their ground-truth labels into the query prompt to learn some patterns hidden in demonstrations when making its prediction. However, it does not distinguish between positive and negative demonstrations for analogy. Motivated from such considerations, we propose to use contrastive learning on the in-context demonstrations to enhance the effectiveness of analogy, as illustrated in Fig. 1. Besides, we also argue that the semantic of event mentions are the most important for the causal relation identification between them. As such we apply contrastive learning to the representation of event mentions in in-context demonstrations, so as to distinguishing the semantic between causal and non-causal event pairs and facilitating event causality predictions.

In this paper, we propose an **In-Context Contrastive Learning (ICCL)** model for the ECI task. The ICCL model contains three modules. The prompt learning module reformulates an input event pair and some retrieved demonstrations into a prompt template, as the input for PLM encoding. The in-context contrastive module optimizes the representation of event mention by simultaneously maximizing its agreement with positive demonstrations and minimizing with negative ones, via a contrastive loss. The causality prediction module predicts answer word to identify causal relations. Experiments are conducted on the widely used EventStoryLine and Causal-TimeBank corpora, and results have shown that our ICCL achieves the new state-of-the-art performance for the ECI task.

2 Related work

2.1 Event Causality Identification

Event Causality Identification (ECI) is an essential task in information extraction, attracting sig-

nificant attention due to its wide range of potential applications. Early methods mainly relied on designing task-oriented neural network models (Liu et al., 2021b; Zuo et al., 2021a). For example, Liu et al. (2021b) improve the capability of their neural model to identify previously unseen causal relations by incorporating event-agnostic and context-specific patterns derived from the ConceptNet (Speer et al., 2017). With further exploration of graph structures and the emergence of large-scale PLMs, recent studies have increasingly adopted graph-based and prompt-based learning approaches to address the ECI task.

Graph-based approaches usually model the ECI task as a node classification problem, employing graph neural networks to learn event node representations based on contextual semantics at the document level (Phu and Nguyen, 2021; Cao et al., 2021; Fan et al., 2022). For example, Fan et al. (2022) establish explicit connections between events, mentions and contexts to construct a co-occurrence graph for node representation learning and causal relation identification. In addition to node classification, some studies approach the ECI task as a graph-based edge prediction problem (Zhao et al., 2021; Chen et al., 2022). For example, Zhao et al. (2021) initialize event node embeddings using a document-level encoder based on a PLM and employ a graph inference mechanism to predict causal edges through graph updating.

2.2 Prompt-based Causality Identification

Recently, with the help of large-scale PLMs, such as the BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019) and etc, prompt learning has emerged as a new paradigm for various NLP tasks (Xiang et al., 2022; Ding et al., 2021). It converts downstream tasks into the similar form as pre-training task, which aligns objectives between the two stages. This alignment helps bridging the gap between PLM and task and can directly enhance the performance of a downstream task. Moreover, researchers have also devised appropriate prompts to reframe ECI task as a cloze task (Shen et al., 2022; Liu et al., 2021b). For example, Shen et al. (2022) propose a derivative prompt joint learning model that leverages potential causal knowledge within PLMs based on the causal cue words detection. Liu et al. (2021b) use an event mention masking generalization mechanism to encode some event causality patterns for causal relation reasoning. Although prompt-based methods are constrained by complex

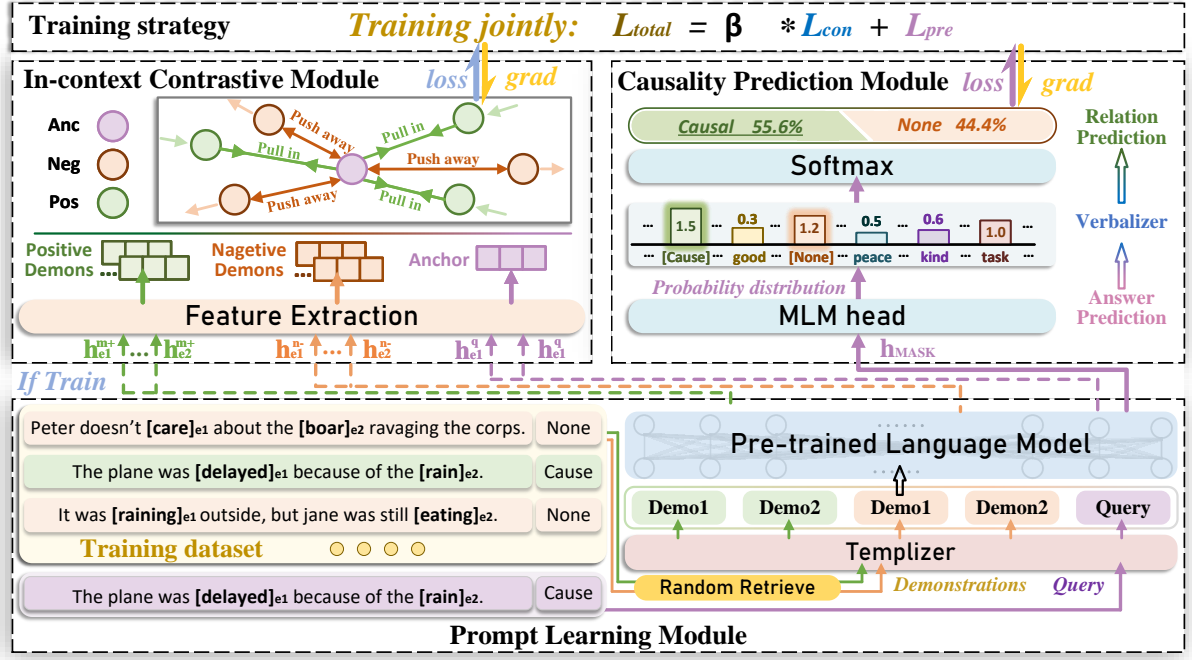


Figure 2: Illustration of our ICCL framework.

prompts and derivate tasks, these prompt-based models effectively leverage the implicit knowledge of PLMs to address the ECI task.

3 Method

Fig. 2 illustrates our ICCL model, including the prompt learning module, the in-context contrastive module, and the causality prediction module.

3.1 Task Formulation

We apply the prompt learning paradigm to transform the ECI task into a causal relation cloze task, utilizing a PLM to predict answer words for causal relation identification. As the event mentions are annotated by a few words in a sentence, we use the event mentions E_1 and E_2 of an event pair as well as their raw sentences S_1 and S_2 , as the input $x = \{E_1, E_2, S_1, S_2\}$, where $E_1 \in S_1$ and $E_2 \in S_2$. The virtual answer words $\langle \text{causal} \rangle$ and $\langle \text{none} \rangle$ indicating whether there is a causal relation between the input event pair, are used as the output $y \in \{\langle \text{causal} \rangle, \langle \text{none} \rangle\}$. We note that in cases where E_1 and E_2 originate from the same sentence, S_1 and S_2 refer to the same sentence.

3.2 Prompt Learning Module

As illustrated in the bottom of Fig. 2, we first reformulate each input instance $x = \{E_1, E_2, S_1, S_2\}$ into a kind of in-context prompt template $T(x)$,

as the input of a PLM for encoding. The in-context prompt input contains a query instance and K retrieved demonstrations. The query instance is the input event instance, denoted as $q = \{E_1^q, E_2^q, S_1^q, S_2^q\}$, with the causal relation between two events to be identified. The demonstrations are retrieved from the training dataset, consisting of an event mention pair and their raw sentences, as well as the relation label between them, denoted as $d_k = \{E_1^k, E_2^k, S_1^k, S_2^k, y^k\}$. We randomly select M demonstrations labeled with $\langle \text{causal} \rangle$ relation and N demonstrations labeled with $\langle \text{none} \rangle$ relation, denoted as d_m^+ and d_n^- , respectively.

We design a prediction prompt template $T_p(q)$ for the query instance q and an analogy prompt template $T_a(d_k)$ for its retrieved demonstrations d_k , respectively. Both of them are constructed by concatenating the raw sentences with a simple cloze template, as follows:

$$T_p(q) = S_1^q + S_2^q + [\text{start}] + E_1^q + [\text{MASK}] + E_2^q + [\text{end}].$$

$$T_a(d_k) = S_1^k + S_2^k + [\text{start}] + E_1^k + y^k + E_2^k + [\text{end}].$$

where $E_1^q, E_2^q, S_1^q, S_2^q$ are the two event mentions and their raw sentences, and the PLM-specific token $[\text{start}]$ and $[\text{end}]$ are used to indicate the beginning and ending of the cloze template. For prediction prompt template $T_p(q)$, a PLM-specific to-

ken [MASK] is inserted between two event mentions for relation prediction; For analogy prompt template $T_a(d_k)$, it is replaced by the virtual word of the relation label y^k for each demonstrations, i.e. <causal> or <none>.

The in-context prompt template $T(x)$ is constructed by concatenating the prediction prompt template $T_p(q)$ and some analogy prompt templates $T_a(d_k)$ of its retrieved demonstrations, as follows:

$$T(x) = [\text{CLS}] + T_a(d_1^+) [\text{SEP}] \dots T_a(d_M^+) [\text{SEP}] + T_a(d_1^-) [\text{SEP}] \dots T_a(d_N^-) [\text{SEP}] + T_p(q) [\text{SEP}].$$

where the PLM-specific token [CLS] and [SEP] are used to indicate the beginning and ending of an input, and some [SEP] tokens are used as separators between the query and those demonstrations. Note that, the causal demonstrations d_m^+ are positioned before the none causal demonstrations d_n^- . We provide a specific example of in-context prompt template input in Appendix C.

After the PLM encoding, we obtain a hidden state $\mathbf{h} \in \mathbb{R}^d$ for each input tokens, where d is the dimension of hidden states. We denote the hidden state of input [MASK] token as \mathbf{h}_{mask} for causality prediction. The hidden states of input event pair in query instance, retrieved causal and non-causal demonstrations are denoted as $[\mathbf{h}_{e_1}^q, \mathbf{h}_{e_2}^q]$, $[\mathbf{h}_{e_1}^{m+}, \mathbf{h}_{e_2}^{m+}]$ and $[\mathbf{h}_{e_1}^{n-}, \mathbf{h}_{e_2}^{n-}]$, respectively, which are next used for in-context contrastive learning.

3.3 In-context Contrastive Module

The in-context contrastive module optimizes the representation of event mention by simultaneously maximizing its agreement with positive demonstration samples and minimizing with negative ones, via a contrastive loss. In the training phase, we use the input query instance as an anchor. The retrieved demonstrations with the same relation label as the query are positive samples, while those with different relation label are negative samples. We assume that the query’s label is <causal>, so the causal demonstrations d_m^+ being treated as positives, and non-causal ones d_n^- as negatives.

Motivated by the fact that the offsets of pre-trained word embeddings can model the relationship between them (Mikolov et al., 2013; Pennington et al., 2014; Chen et al., 2016), such as $\mathbf{h}_{king} - \mathbf{h}_{man} \approx \mathbf{h}_{queen} - \mathbf{h}_{woman}$. We use the offsets between event mentions’ hidden states to represent their relation for contrastive learning, as

follows:

$$\mathbf{z}^q = \mathbf{h}_{e_1}^q - \mathbf{h}_{e_2}^q, \quad (1)$$

$$\mathbf{z}_m^+ = \mathbf{h}_{e_1}^{m+} - \mathbf{h}_{e_2}^{m+}, \quad (2)$$

$$\mathbf{z}_n^- = \mathbf{h}_{e_1}^{n-} - \mathbf{h}_{e_2}^{n-}, \quad (3)$$

where $\mathbf{z}^q, \mathbf{z}_m^+, \mathbf{z}_n^-$ are the relation vector of event pair in query instance, positive and negative demonstrations, respectively.

We adopt supervised contrastive learning on the relation vector of event pair for its representation optimization (Khosla et al., 2020). Specifically, it pulls together the anchor towards positive samples in embedding space, while simultaneously pushing it apart from negative samples. The supervised contrastive loss is computed as follows:

$$L_{con} = -\log \frac{\sum_{m=1}^M \exp(\text{sim}(\mathbf{z}^q, \mathbf{z}_m^+)/\tau)}{\sum_{d \in \mathcal{D}} \exp(\text{sim}(\mathbf{z}^q, d)/\tau)}, \quad (4)$$

where $\mathcal{D} = \{\mathbf{z}_m^+\}_{m=1}^M \cup \{\mathbf{z}_n^-\}_{n=1}^N$, M and N represent the number of positive and negative demonstrations, respectively.

3.4 Causality Prediction Module

The causality prediction module uses the [MASK] token of input query instance to predict an answer word for causal relation identification. Specifically, we input the hidden state \mathbf{h}_{mask} into the masked language model classifier, and estimate the probability of each word in its vocabulary dictionary \mathcal{V} for the [MASK] token, as follows:

$$P([\text{MASK}] = v \in \mathcal{V} | T(x)), \quad (5)$$

We add two virtual words into PLM’s vocabulary dictionary as the answer space, viz. <causal> and <none>, to indicate whether a causal relation exists or not. Then a softmax layer is applied on the prediction scores of the two virtual answer words to normalize them into probabilities:

$$P_i(v_i \in \mathcal{V}_a | T(x)) = \frac{\exp(p_{v_i})}{\sum_{j=1}^n \exp(p_{v_j})}, \quad (6)$$

where $\mathcal{V}_a = \{\text{<causal>, <none>}\}$.

In the training phase, we tune parameters of PLM and MLM classifier based on in-context prompt and newly added virtual words. We adopt the cross entropy loss as the loss function:

$$L_{pre} = -\frac{1}{L} \sum_{l=1}^L \mathbf{y}^{(l)} \log(\hat{\mathbf{y}}^{(l)}) + \lambda \|\theta\|^2, \quad (7)$$

where $y^{(l)}$ and $\hat{y}^{(l)}$ are answer label and predicted label of the l -th training instance respectively. λ and θ are the regularization hyper-parameters. We use the AdamW optimizer (Loshchilov and Hutter, 2017) with $L2$ regularization for model training.

3.5 Training strategy

We jointly train the in-context contrastive module and the causality prediction module. The loss function of our ICCL model is optimized as follows:

$$L_{total} = L_{pre} + \beta * L_{con}, \quad (8)$$

where β is the weight coefficient to balance the importance of contrastive loss and prediction loss. We conduct some experiments to explore the impact of different β values on model performance. The experimental results and analysis are presented in Appendix D.

4 Experiment Setting

4.1 Datasets

Our experiments are conducted on two widely used datasets for the ECI task: EventStory-Line 0.9 Corpus (ESC) (Caselli and Vossen, 2017) and Causal-TimeBank Corpus (CTB) (Mirza and Tonelli, 2014).

EventStoryLine contains 22 topics and 258 documents collected from various news websites. In total, there are 5,334 event mentions in ECS dataset. Among them, 5,625 event pairs are annotated with causal relations. Specifically, 1,770 causal relations are intra-sentence causalities, while 3,855 ones are cross-sentence causalities. Following the standard data splitting Gao et al. (2019), we use the last two topics as the development set, and conduct 5-fold cross-validation on the remaining 20 topics. The average results of precision (P), recall (R), and F1 score are adopted as performance metrics.

Causal-TimeBank comprises 184 documents sourced from English news articles, with a total of 7,608 annotated event pairs. Among them, 318 are annotated with causal relations. Specifically, 300 causal relations are intra-sentence causalities, while only 18 ones are cross-sentence causalities. Following the standard data splitting (Liu et al., 2021a), we employ a 10-fold cross-validation and the average results of precision (P), recall (R), and F1 score are adopted as performance metrics. Following Phu and Nguyen (2021), we only conduct intra-sentence event causality identification exper-

iments on CTB, as the number of cross-sentence event causal pairs is quite small.

4.2 Parameter Setting

We use the pre-trained RoBERTa (Liu et al., 2019) model with 768-dimension base version provided by the HuggingFace transformers² (Wolf et al., 2020). Our implementation is based on PyTorch framework³, running on NVIDIA GTX 3090 GPUs. The training process costs approximately 5 GPU hours on average. We set the learning rate to 1e-5, batch size to 16. The contrastive loss ratio β is set to 0.5, the temperature parameter τ is set to 1.0, and the number of demonstrations is set to 4, viz. $(M, N) = (2, 2)$. All trainable parameters are randomly initialized from normal distributions.

4.3 Competitors

We compare our ICCL with the following competitors: ILP (Gao et al., 2019), KnowMMR (Liu et al., 2021b), RichGCN (Phu and Nguyen, 2021), CauSeRL (Zuo et al., 2021a), LSIN (Cao et al., 2021), LearnDA (Zuo et al., 2021b), GESI (Fan et al., 2022), ERGO (Chen et al., 2022), DPJL (Shen et al., 2022), SemSIn (Hu et al., 2023). The detailed introduction of competitors can be found in Appendix B.

5 Result and Analysis

5.1 Overall Result

Table 1 compares the overall performance between our ICCL and the competitors on the ESC and CTB corpus. We can observe that the ILP cannot outperform other competitors, including the RichGCN, GESI, ERGO, and SemSIn. This can be attributed to their utilization of some graph neural networks for document structure encoding, enabling them to learn global contextual semantic for causality prediction. We can also observe that the DPJL adopting a kind of derivative prompt learning can significantly outperform the other competitors in intra-sentence causality identification. The outstanding performance can be attributed to its applying the prompt learning paradigm that transforms the ECI task to directly predict a PLM vocabulary word, other than fine-tuning a task-specific neural model upon a PLM. Although some other competitors have used external knowledge bases for rela-

²<https://github.com/huggingface/transformers>

³pytorch.org

Model	PLM	EventStoryLine									Causal-TimeBank		
		Intra			Cross			Intra and Cross			Intra		
		P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)	P(%)	R(%)	F1(%)
ILP (Gao et al., 2019)	-	38.8	52.4	44.6	35.1	48.2	40.6	36.2	49.5	41.9	-	-	-
LearnDA (Zuo et al., 2021b)	BERT	42.2	69.8	52.6	-	-	-	-	-	-	41.9	68.0	51.9
RichGCN (Phu and Nguyen, 2021)	BERT	49.2	63.0	55.2	39.2	45.7	42.2	42.6	51.3	46.6	39.7	56.5	46.7
DPJL (Shen et al., 2022)	RoBERTa	65.3	70.8	67.9	-	-	-	-	-	-	63.6	66.7	64.6
GESI (Fan et al., 2022)	BERT	-	-	50.3	-	-	49.3	-	-	49.4	-	-	-
ERGO (Chen et al., 2022)	Longformer	57.5	72.0	63.9	51.6	43.3	47.1	48.6	53.4	50.9	62.1	61.3	61.7
SemSIn (Hu et al., 2023)	BERT	64.2	65.7	64.9	-	-	-	-	-	-	52.3	65.8	58.3
ICCL	BERT	64.9	69.6	67.1	56.3	58.4	57.2	59.0	61.9	60.4	60.5	58.4	59.1
	ERNIE	66.8	68.5	67.5	63.7	56.2	59.5	64.8	60.0	62.1	64.8	66.0	64.7
	DeBERTa	67.6	73.7	70.4	61.8	58.4	59.9	61.7	63.2	63.3	66.7	64.4	64.9
	RoBERTa	67.5	73.7	70.4	60.3	62.7	61.3	62.6	66.1	64.2	63.7	68.8	65.4

Table 1: Comparison of overall results on the ESC and CTB corpus.

tion identification, the prompt learning paradigm can better leverages potential causal knowledge in PLMs.

Finally, our ICCL with different PLMs has achieved significant performance improvements overall competitors in terms of much higher F1 score with all intra-sentence, inter-sentence, and overall event causality identification on both ESC and CTB corpus. We attribute its outstanding performance to applying contrastive learning on in-context demonstrations, by which our ICCL can better distinguish the semantic of causal and non-causal event pairs for causality prediction. Furthermore, we can also observe that using different PLMs do result in some performance variations, which are further discussed in Appendix A. Finally the ICCL based on RoBERTa has achieved the best performance, as such we implement the remaining ablation experiments with RoBERTa.

5.2 Ablation Study

To examine the effectiveness of contrastive learning and in-context learning, we design the following ablation study. Table 2 compares their performance.

- Prompt is prompt learning model, without demonstrations or contrastive module.
- In-context is in-context learning model, including retrieved demonstrations but without contrastive module.
- ProCon w/o Demos is prompt based contrastive model, but without demonstrations. We select positive and negative samples within batch insted of demonstrations, and use hidden state of [MASK] as input to contrastive module.
- ProCon w/ Demos is in-context based contrastive model with retrieved demonstrations, but still use the hidden state of [MASK] as input to contrastive module.

- EvtCon is event based prompt contrastive model, the only difference with ProCon w/o Demos is using hidden states of event pairs as contrastive module inputs.

In-context learning: The first observation is that models incorporating in-context learning perform better. For example, the three models, In-context, ProCon w/ Demos and ICCL outperform Prompt, ProCon w/o Demos and EvtCon, respectively. This indicates the inclusion of demonstrations to explicitly guide the label prediction is highly effective in improving model performance. Furthermore, models with in-context learning show notable performance gains in challenging cross-sentence causality identification. That’s because randomly selected demonstrations are predominantly composed of cross-sentence samples, which are more abundant in datasets. Therefore, PLMs develop a more comprehensive understanding of cross-sentence causality.

Contrastive learning: We can observe that models with a contrastive module exhibit better performance. For example, both ProCon w/ Demos and EvtCon preform better than Prompt. Additionally, both ProCon w/o Demos and ICCL preform better than In-context. This can be attributed to the utilization of the contrastive learning paradigm, which enables the PLM to concentrate on event pairs or [MASK] and enhances PLM’s ability to model them. Furthermore, it also helps discriminatively model positive and negative demonstrations, strengthening analogy between the query and all demonstrations. Additionally, we also observe that EvtCon usually outperforms ProCon w/o Demos. That’s because hidden state of [MASK] serves as input for both contrastive and prediction module in the case of ProCon w/o Demos, yet the optimization directions of two modules do not

Model	EventStoryLine									Cause-TimeBank		
	Intra			Cross			Intra and Cross			Intra		
	p (%)	r (%)	f1 (%)	p (%)	r (%)	f1 (%)	p (%)	r (%)	f1 (%)	p (%)	r (%)	f1 (%)
Prompt	67.2	69.7	68.2	58.6	59.8	59.0	61.3	62.9	61.7	58.9	55.3	56.6
In-context	66.0	72.4	68.9	57.7	60.9	59.1	60.4	64.5	62.2	60.3	58.0	58.7
ProCon w/o Demos	60.8	77.9	68.2	54.2	65.6	59.3	56.4	69.4	62.1	51.5	71.8	58.9
ProCon w/ Demos	67.1	73.5	70.1	58.0	61.9	59.8	60.9	64.5	63.1	66.9	60.2	62.5
EvtCon	62.1	78.2	69.0	52.3	68.9	59.1	55.3	71.8	62.1	55.8	65.6	59.8
ICCL	67.5	73.7	70.4	60.3	62.7	61.3	62.6	66.1	64.2	63.7	68.8	65.4

Table 2: Results of ablation study on the ESC and CTB corpus.

completely align.

5.3 Numbers of demonstrations

To further investigate the impact of demonstrations, we conducted an experiment that compared the performance of In-context and ICCL with varying numbers of causal and non-causal demonstrations. The results are showcased in Fig. 3.

With more demonstrations, F1-score of both models initially exhibited improved performance, further validating the effectiveness of using demonstrations as explicit guidance. However, as the input length becomes too long, performance of In-context declines, while the performance of ICCL continues to improve. This can be attributed to the effectiveness of contrastive module used in ICCL, which aids the PLM in better focusing on event pairs, even with longer input. Additionally, the causal/non-causal ratio of 2/1 performs better compared to that of 1/2. That’s because the dataset contains a limited number of causal samples. Increasing the number of causal demonstrations helps the model better learn the features of causal examples, mitigating the data imbalance issue.

We can also observe that performance metrics of In-context model, particularly precision, exhibit minimal changes when the number of demonstrations varies. While as for our ICCL model, the precision and recall vary based on the ratio of causal and non-causal demonstrations. More non-causal demonstrations results in higher recall, while the opposite scenario leads to higher precision. These findings emphasize that the critical role of the contrastive module in enhancing analogy and enabling the PLM to effectively utilize positive and negative demonstrations.

5.4 Few shot

Some researchers have reported the robustness of prompt paradigm in using fewer training data (Wang et al., 2021; Ding et al., 2021). Since




Figure 3: Comparison of ICCL and In-context model when using different numbers of causal and non-causal demonstrations on ESC corpus.

our ICCL also employs a prompt-based method to predict the label, we examine its performance in low-resource scenarios and replicate the performance of ERGO as a benchmark for comparison. Fig. 4 shows the performance comparison between ERGO and our ICCL on ESC corpus.

As expected, the performance of ICCL gradually decreases as the amount of training data decreases. However, the decrease in performance is relatively slow, with an F1 score decrease of about 10% when training data is reduced by 80%, whereas the performance of ERGO declined by nearly 25%. Notably, even with only 20% of the training data, ICCL (F1: 51.9%) outperforms ERGO (F1: 50.9%) and many other competitors with full training data. These results confirm the effectiveness of ICCL even with fewer training data.

We also showcase the intra-sentence causality identification performance among different PLMs




Figure 4: Results of few shot on ESC corpus. We replicated ERGO and get its few-shot results in the figure.

Model	EventStoryLine			Cause-TimeBank		
	P (%)	R (%)	F1 (%)	P (%)	R (%)	F1 (%)
BERT (Gao et al., 2023)	38.1	56.8	45.6	41.1	45.8	43.5
RoBERTa (Gao et al., 2023)	42.1	64.0	50.8	39.9	60.9	48.2
T5 (Our implementation)	36.2	49.2	40.7	7.7	52.1	12.1
gpt-3.5-turbo (Gao et al., 2023)	27.6	80.2	41.0	6.9	82.6	12.8
gpt-4 (Gao et al., 2023)	27.2	94.7	42.2	6.1	97.4	11.5

Table 3: Intra-sentence causality identification results of different PLMs and LLMs on the ESC and CTB corpus.

and several zero-shot models in the Table 3. We can not only find that our fine-tuned generative model, T5 (Our implementation), perform significantly worse than autoencoder models like BERT-base (Gao et al., 2023) and RoBERTa-base (Gao et al., 2023), which confirms the conclusion drawn by Gao et al. (2023) that generative models may not be well-suited for causal reasoning tasks like ECI. We can also observe that although the ChatGPT models, such as gpt-3.5-turbo and gpt-4, have more comprehensive pre-training and larger model scales, these zero-shot models exhibit a significant performance gap compared to fine-tuned models like T5-base and et al. This demonstrates the importance of fine-tune, indicating that it is challenging to address causal reasoning tasks like ECI in a zero-shot scenario. For more detailed analysis, please refer to Appendix A.

5.5 Embedding Visualization

In order to verify the impact of contrastive module with event pairs as input, we compare the learned event pairs’ embeddings ($h_{e_1} - h_{e_2}$) of different models on ESC test dataset by t-distributed stochastic neighbor embedding (t-SNE) (Hinton and Roweis, 2002). In Fig. 5, we color-coded the points to represent *True Negative (TN)*, *False Pos-*

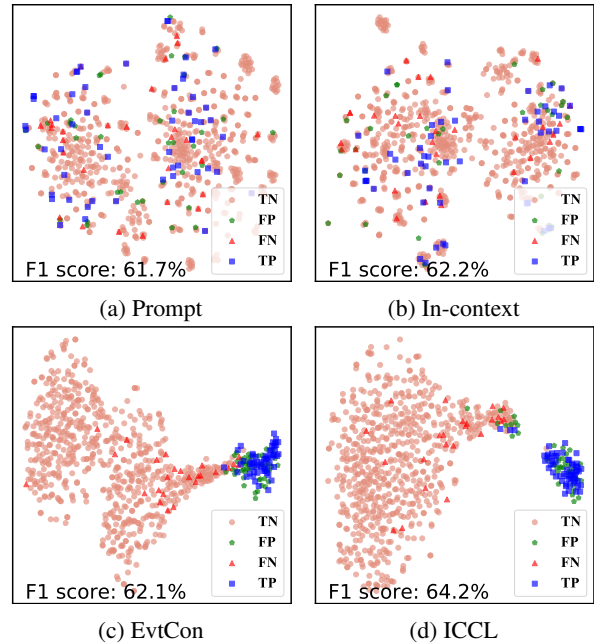


Figure 5: Visualization of the event pairs’ embedding encoded by different models on ESC corpus

itive (FP), *False Negative (FN)* and *True Positive (TP)* samples.

We can observe that models incorporating the contrastive module with event pairs as input exhibit a clear phenomenon of event pairs representations clustering together based on labels in the embedding space, which demonstrates the effective of the contrastive module. Additionally, representations of samples predicted to have the same label tended to cluster together, highlighting the crucial role of event pairs in identifying causality.

6 Concluding Remarks

In this paper, we propose an ICCL model and apply it on the ECI task. We leverage the causality knowledge of PLM by introducing explicit guidance through the inclusion of demonstrations, rather than relying on the design of complex prompts. Meanwhile, we employ contrastive learning with event pairs as input to enhance the PLM’s attention to event pairs and strengthen the analogy between query and demonstrations. Experiments on the ESC and CTB corpus have validated that our ICCL can significantly outperform the state-of-the-art algorithms.

In future, we will try to undertake experiments to apply our proposed framework to other NLP tasks in order to explore whether it can exhibit favorable adaptability when applied to different tasks.

Limitation

Due to the input length limitations of the PLM, the number of demonstrations needs to be kept within a manageable range. However, our ICCL uses demonstrations as positive and negative samples in contrastive learning. This implies that there are limited positive and negative samples, which weakens the effectiveness of contrastive learning.

Acknowledgements

This work is supported in part by National Natural Science Foundation of China (Grant No: 62172167). The computation is completed in the HPC Platform of Huazhong University of Science and Technology.

Ethics Statement

This paper has no particular ethic consideration.

References

- Sourabh Balgi, Jose M Pena, and Adel Daoud. 2022. Personalized public policy analysis in social sciences using causal-graphical normalizing flows. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11810–11818.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Abby Vander Linden, Brittany Harding, Brad Huang, Peter Clark, and Christopher D Manning. 2014. Modeling biological processes for reading comprehension. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1499–1510.
- Manvi Breja and Sanjay Kumar Jain. 2020. Causality for question answering. In *COLINS*, pages 884–893.
- Pengfei Cao, Xinyu Zuo, Yubo Chen, Kang Liu, Jun Zhao, Yuguang Chen, and Weihua Peng. 2021. Knowledge-enriched event causality identification via latent structure induction networks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4862–4872.
- Tommaso Caselli and Piek Vossen. 2017. The event storyline corpus: A new benchmark for causal and temporal relation extraction. In *Proceedings of the Events and Stories in the News Workshop*, pages 77–86.
- Meiqi Chen, Yixin Cao, Kunquan Deng, Mukai Li, Kun Wang, Jing Shao, and Yan Zhang. 2022. Ergo: Event relational graph transformer for document-level event causality identification. *arXiv preprint arXiv:2204.07434*.
- Qian Chen, Xiaodan Zhu, Zhenhua Ling, Si Wei, Hui Jiang, and Diana Inkpen. 2016. Enhanced lstm for natural language inference. *arXiv preprint arXiv:1609.06038*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ning Ding, Yulin Chen, Xu Han, Guangwei Xu, Pengjun Xie, Hai-Tao Zheng, Zhiyuan Liu, Juanzi Li, and Hong-Gee Kim. 2021. Prompt-learning for fine-grained entity typing. *arXiv preprint arXiv:2108.10604*.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Chuang Fan, Daoxing Liu, Libo Qin, Yue Zhang, and Ruifeng Xu. 2022. Towards event-level causal relation identification. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1828–1833.
- Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. 2023. Is chatgpt a good causal reasoner? a comprehensive evaluation. *arXiv preprint arXiv:2305.07375*.
- Lei Gao, Prafulla Kumar Choubey, and Ruihong Huang. 2019. Modeling document-level causal structures for event causal relation identification. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1808–1817.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*.
- Geoffrey E Hinton and Sam Roweis. 2002. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15.
- Zhilei Hu, Zixuan Li, Xiaolong Jin, Long Bai, Saiping Guan, Jiafeng Guo, and Xueqi Cheng. 2023. Semantic structure enhanced event causality identification. *arXiv preprint arXiv:2305.12792*.
- Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673.

- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021a. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*.
- Jian Liu, Yubo Chen, and Jun Zhao. 2021b. Knowledge enhanced event causality identification with mention masking generalizations. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3608–3614.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Paramita Mirza and Sara Tonelli. 2014. An analysis of causality between events and its relation to temporal information. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2097–2106.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Minh Tran Phu and Thien Huu Nguyen. 2021. Graph convolutional networks for event causality identification with rich document-level structures. In *Proceedings of the 2021 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pages 3480–3490.
- Peter G Preethi, Vilma Uma, et al. 2015. Temporal sentiment analysis and causal rules extraction from tweets for event prediction. *Procedia computer science*, 48:84–89.
- Ruili Pu, Yang Li, Suge Wang, Deyu Li, Jianxing Zheng, and Jian Liao. 2023. Enhancing event causality identification with event causal label and event pair interaction graph. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10314–10322.
- Kira Radinsky, Sagie Davidovich, and Shaul Markovitch. 2012. Learning causality for news events prediction. In *Proceedings of the 21st international conference on World Wide Web*, pages 909–918.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Shirong Shen, Heng Zhou, Tongtong Wu, and Guilin Qi. 2022. Event causality identification via derivative prompt joint learning. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2288–2299.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Xuyi Chen, Han Zhang, Xin Tian, Danxiang Zhu, Hao Tian, and Hua Wu. 2019. Ernie: Enhanced representation through knowledge integration. *arXiv preprint arXiv:1904.09223*.
- Chengyu Wang, Jianing Wang, Minghui Qiu, Jun Huang, and Ming Gao. 2021. Transprompt: Towards an automatic transferable prompting framework for few-shot text classification. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pages 2792–2802.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Wei Xiang, Zhenglin Wang, Lu Dai, and Bang Wang. 2022. Connprompt: Connective-cloze prompt learning for implicit discourse relation recognition. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 902–911.
- Kun Zhao, Donghong Ji, Fazhi He, Yijiang Liu, and Yafeng Ren. 2021. Document-level event causality identification via graph inference mechanism. *Information Sciences*, 561:115–129.
- Xinyu Zuo, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Weihua Peng, and Yuguang Chen. 2021a. Improving event causality identification via self-supervised representation learning on external causal statement. *arXiv preprint arXiv:2106.01654*.
- Xinyu Zuo, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Weihua Peng, and Yuguang Chen. 2021b. Learnda: Learnable knowledge-guided data augmentation for event causality identification. *arXiv preprint arXiv:2106.01649*.

A Study of PLMs

The ICCL model we proposed is a PLM-sensitive model. In order to investigate the performance of our model using different PLMs and select the most suitable one, we conducted PLM ablation experiment to test performance of our model with different PLMs. Furthermore, we also cited performance of some baseline methods based on PLMs finetuned on full training datasets from the work of Gao et al. (2023) to evaluate various PLMs and summarized the results in Table 4. The introductions of main PLMs we considered are as follows:

- **BERT** (Devlin et al., 2018): The most representative PLM proposed by Google⁴, which is pre-trained using a cloze task and a next sentence prediction task.

- **RoBERTa** (Liu et al., 2019): A BERT enhanced PLM proposed by Facebook⁵, which removes the next sentence prediction objective and is pre-trained on a much larger dataset with some modified key hyper-parameters.

- **ERNIE** (Sun et al., 2019): A knowledge enhanced PLM proposed by Baidu⁶, which uses some knowledgeable masking strategies in pre-training.

- **DeBERTa** (He et al., 2020): The latest masked PLM proposed by Microsoft⁷, which improves BERT and RoBERTa models using a disentangled attention mechanism and an enhanced mask decoder.

- **T5** (Raffel et al., 2020): A generative language model proposed by Google⁸ in 2020, which is pre-trained on large-scale unsupervised datasets using an autoregressive approach and fine-tuned on task-specific annotated data. It has achieved state-of-the-art performance on multiple NLP tasks such as text generation, summarization, and translation.

As shown in Table 4, according to the research by Gao et al. (2023), it can be observed that, our fine-tuned generative model, *T5-base*, performs significantly worse than autoencoder models like *BERT-base* (Gao et al., 2023) and *RoBERTa-base* (Gao et al., 2023). Moreover, the performance of *In-context-T5* is also far inferior to the model *In-context-RoBERTa*. This confirms the conclusion

⁴<https://github.com/google-research/bert>

⁵<https://github.com/pytorch/fairseq>

⁶<https://github.com/PaddlePaddle/ERNIE>

⁷<https://github.com/microsoft/DeBERTa>

⁸<https://github.com/google-research/multilingual-t5>

drawn by Gao et al. (2023) that generative models may not be well-suited for causal reasoning tasks like ECI. Additionally, although the ChatGPT models, such as *gpt-3.5-turbo* and *gpt-4*, have more comprehensive pre-training and larger model scales, these zero-shot models exhibit a significant performance gap compared to fine-tuned models like *T5-base* and et al. This demonstrates the importance of fine-tune, indicating that it is challenging to address causal reasoning tasks like ECI in a zero-shot scenario.

Besides, we can observe that our ICCL with all four PLMs has achieved better performance than most of competitors on both ESC and CTB corpus. Even our *ICCL-BERT* outperformed many competitors with advanced PLMs, such as ERGO based on Longformer (Beltagy et al., 2020). This further demonstrates the effectiveness of our proposed method. Compared to approaches involving complex prompts or joint training across multiple tasks, our approach of utilizing simple explicit guidance and leveraging it for contextual contrastive learning better harnesses the semantic knowledge embedded in PLMs and guides their understanding of causal relationships.

We can also observe that using different PLMs do result in some performance variations. This is not unexpected. It can be attributed to that while all the four PLMs employ a kind of Transformer-based model in pre-training on large-scale corpus, their training strategies or training corpus are not entirely identical. Compared to *ICCL-BERT*, our ICCL model using ERNIE, DeBERTa, or RoBERTa achieved better performance. This is attributed to the fact that these three PLMs have made some optimizations based on BERT. For example, ERNIE utilizes a strategy of continuous learning in the pre-training stage. Finally, *ICCL-RoBERTa* achieved the best performance, which removes the next sentence prediction objective and is pre-trained on a much larger dataset with some modified key hyper-parameters. Therefore, we implement the remaining ablation experiments with RoBERTa.

B Competitors

Table 4 also presents results of more competitors. The introductions of these competitors are as follows:

- **ILP** (Gao et al., 2019) employs integer linear programming to detect causal relationships by in-

Model	EventStoryLine									Cause-TimeBank		
	Intra			Cross			Intra and Cross			Intra		
	p (%)	r (%)	f1 (%)	p (%)	r (%)	f1 (%)	p (%)	r (%)	f1 (%)	p (%)	r (%)	f1 (%)
T5	36.2	49.2	40.7	-	-	-	-	-	-	7.7	52.1	12.1
BERT †	38.1	56.8	45.6	-	-	-	-	-	-	41.4	45.8	43.5
RoBERTa †	42.1	64.0	50.8	-	-	-	-	-	-	39.9	60.9	48.2
text-davinci-002 †	23.2	80.0	36.0	-	-	-	-	-	-	5.0	75.2	9.3
text-davinci-003 †	33.2	74.4	45.9	-	-	-	-	-	-	8.5	64.4	15.0
gpt-3.5-turbo †	27.6	80.2	41.0	-	-	-	-	-	-	6.9	82.6	12.8
gpt-4 †	27.2	94.7	42.2	-	-	-	-	-	-	6.1	97.4	11.5
In-context-T5	63.3	62.6	62.7	53.7	46.6	49.3	57.0	51.5	53.7	9.2	50.4	14.8
In-context-RoBERTa	66.0	72.4	68.9	57.7	60.9	59.1	60.4	64.5	62.2	60.3	58.0	58.7
ILP (Gao et al., 2019)	38.8	52.4	44.6	35.1	48.2	40.6	36.2	49.5	41.9	-	-	-
KnowMMR (Liu et al., 2021b)	41.9	62.5	50.1	-	-	-	-	-	-	36.6	55.6	44.1
RichGCN (Phu and Nguyen, 2021)	49.2	63.0	55.2	39.2	45.7	42.2	42.6	51.3	46.6	39.7	56.5	46.7
CauSeRL (Zuo et al., 2021a)	41.9	69.0	52.1	-	-	-	-	-	-	43.6	68.1	53.2
LSIN (Cao et al., 2021)	47.9	58.1	52.5	-	-	-	-	-	-	51.5	56.2	53.7
LearnDA (Zuo et al., 2021b)	42.2	69.8	52.6	-	-	-	-	-	-	41.9	68.0	51.9
GESI (Fan et al., 2022)	-	-	50.3	-	-	49.3	-	-	49.4	-	-	-
ERGO (Chen et al., 2022)	57.5	72.0	63.9	51.6	43.3	47.1	48.6	53.4	50.9	62.1	61.3	61.7
DPJL (Shen et al., 2022)	65.3	70.8	67.9	-	-	-	-	-	-	63.6	66.7	64.6
SemSIn (Hu et al., 2023)	64.2	65.7	64.9	-	-	-	-	-	-	52.3	65.8	58.3
ICCL-BERT	64.9	69.6	67.1	56.3	58.4	57.2	59.0	61.9	60.4	60.5	58.4	59.1
ICCL-ERNIE	66.8	68.5	67.5	63.7	56.2	59.5	64.8	60.0	62.1	64.8	66.0	64.7
ICCL-DeBERTa	67.6	73.7	70.4	61.8	58.4	59.9	61.7	63.2	63.3	66.7	64.4	64.9
ICCL-RoBERTa	67.5	73.7	70.4	60.3	62.7	61.3	62.6	66.1	64.2	63.7	68.8	65.4

Table 4: Comparison of overall results on the ESC and CTB corpus. Performance of models marked with "†" after the name are cited from the research of Gao et al. (2023). We name our models in the format of Model-PLM, for example, ICCL-BERT is the version of ICCL model based on BERT.

corporating causal constraints at document level.

- KnowMMR (Liu et al., 2021b) utilizes external knowledge to extract event causality patterns.
- RichGCN (Phu and Nguyen, 2021) uses a graph convolutional network to learn context-enriched representations for event pairs based on document-level information.
- CauSeRL (Zuo et al., 2021a) employs a contrastive approach to transfer externally learned causal statements.
- LSIN (Cao et al., 2021) employs graph induction to acquire external structural and relational knowledge.
- LearnDA (Zuo et al., 2021b) utilizes knowledge bases to interactively generate training data.
- GESI (Fan et al., 2022) designs a graph convolutional network on an event co-reference graph to model causality.
- ERGO (Chen et al., 2022) constructs a relational graph where event pairs serve as nodes, capturing causal transitivity through a transformer-like network.
- DPJL (Shen et al., 2022) leverages two derivative prompt tasks to identify causality.
- SemSIn (Hu et al., 2023) uses a Graph Neural Network (GNN) to learn from event-centric struc-

tures for encoding events.

C In-context input

To help readers gain a better understanding of the in-context input generated by our Prompt module, we provide a specific example in Fig. 6.

As depicted in Fig. 6, we randomly chose two causal demonstrations and two non-causal demonstrations from the training dataset for the query. Each segment in Fig. 6 represents either a prompted demonstration or a prompted query. The initial two segments, highlighted in green font, represents demonstrations labeled as *<causal>*. The following two segments, highlighted in orange font, represents demonstrations labeled as *<none>*. Lastly, the final segment, highlighted in purple font, represents the query to predict.

Besides, we have annotated some specific tokens we used with special colors. We utilized three PLM-special tokens: *[CLS]* to indicate the beginning of the input, *[SEP]* as a sentence separator, and *[MASK]* as a placeholder for the label to predict. Furthermore, we have also devised some additional special tokens: *[start]* and *[end]* are used to indicate the beginning and end of the cloze tem-

In-context Input

Causal Demonstrations

[CLS] A Provisional trial date has been set in the case of a son **accused** of killing his mother , sister and pet dog in Millom. A preliminary **hearing** for John Jenkin , 23 , charged with the murders of his mother Alice McMeekin , 58 , and sister Katie Jenkin , 20 , was heard in Preston Crown Court this morning . [start] **accused** <causal> **hearing** [end] [SEP]

A powerful **earthquake** hit southern Iran on Sunday , causing major destruction in seven village and killing 10 people and injuring 80 . The island's airport was also **damaged** . [start] **earthquak** <causal> **damaged** [end] [SEP]

Non-causal Demonstrations

" He was shot in the head and **left** dying on the ground while his killer ran away and tried to hide The defendant ' s custody status gave Sheriff ' s detectives and our prosecutors in the Crimes Against Police Officers Section (CAPOS) additional time to fully investigate this murder and the **case** on which Deputy Ortiz was working when he was killed , " Cooley said . [start] **left** <none> **case** [end] [SEP]

"My client is ensconced in the bosom of that facility right now , " Heller **argued** after a prosecuto objected to Lohan's choice of rehab facilities . " Nothing bad is **going to happen** . " [start] **arguer** <none> **going** [end] [SEP]

Query

Breaking : man <event1> **charged** </event1> with arson after fire at Waitrose in Wellington. A man has been charged on suspicion of <event2> **arson** </event2> following a fire that devastated a Somerset supermarket. [start] **charged** [MASK] **arson** [end] [SEP]

Figure 6: Example of in-context input. The line breaks and the title of each part (ex. Causal Demonstrations) are only to make the input readable, and they are not included in the actual input.




Figure 7: Comparison of ICCL model with different value of β on the ESC corpus.

plate respectively, $[event1]$, $[event1/]$, $[event2]$, $[event2/]$ are used to highlight the events in the query, while $\langle causal \rangle$ and $\langle none \rangle$ respectively represent the causal and uncausal labels for the demonstrations.

Additionally, although the contrastive module only works during the training phase, we select appropriate demonstrations for the query in both training and testing phases. Specifically, we randomly select M samples labeled as $\langle causal \rangle$ and N samples labeled as $\langle none \rangle$ from training dataset to be demonstrations. And on the contrastive learning process, positive demonstrations are those with the same label as the query, while negative demonstrations have different labels. Furthermore, during training phase, different demonstrations are retrieved for the same query in different epochs to introduce variability and enhance the model's ability to handle diverse instances of the same query. However, during validation and testing state, demonstrations retrieved for the same query, as well as the permutation order, remain consistent across epochs which ensures fair evaluation.

D Study of β

To further explore how to balance the importance of contrastive loss and prediction loss, we investigated the performance of the ICCL model with different values of the hyperparameter β on the ESC corpus.

As shown in Fig. 7, we can observe that as β increases from 0, the performance of the model initially improves and then starts to decline. The optimal performance on both intra-sentence causality and cross-sentence causality is achieved when $\beta = 0.5$. This indicates that the introduction of contrastive learning loss does indeed help the model better focus on event pairs of the query and demonstrations, understand causalities, and achieve better performance. However, it is important to strike a balance between the contrastive learning loss and the prediction loss. Excessive emphasis on the former should be avoided as it may cause the model to overly prioritize modeling event pairs and overlook the semantic relevance of the context, which can ultimately lead to a decrease in the model's performance.