

LearnDA: Learnable Knowledge-Guided Data Augmentation for Event Causality Identification

Xinyu Zuo^{1,2}, Pengfei Cao^{1,2}, Yubo Chen^{1,2}, Kang Liu^{1,2}, Jun Zhao^{1,2},
Weihua Peng³ and Yuguang Chen³

¹National Laboratory of Pattern Recognition, Institute of Automation, CAS, Beijing, China

²School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China

³Beijing Baidu Netcom Science Technology Co., Ltd

{xinyu.zuo, pengfei.cao, yubo.chen, kliu, jzhao}@nlpr.ia.ac.cn
{pengweihua, chen yuguang}@baidu.com

Abstract

Modern models for event causality identification (ECI) are mainly based on supervised learning, which are prone to the data lacking problem. Unfortunately, the existing NLP-related augmentation methods cannot directly produce available data required for this task. To solve the data lacking problem, we introduce a new approach to augment training data for event causality identification, by iteratively generating new examples and classifying event causality in a dual learning framework. On the one hand, our approach is knowledge guided, which can leverage existing knowledge bases to generate well-formed new sentences. On the other hand, our approach employs a dual mechanism, which is a learnable augmentation framework, and can interactively adjust the generation process to generate task-related sentences. Experimental results on two benchmarks EventStoryLine and Causal-TimeBank show that 1) our method can augment suitable task-related training data for ECI; 2) our method outperforms previous methods on EventStoryLine and Causal-TimeBank (+2.5 and +2.1 points on F1 value respectively).

1 Introduction

Event causality identification (ECI) aims to identify causal relations between events in texts, which can provide crucial clues for NLP tasks, such as logical reasoning and question answering (Girju, 2003; Oh et al., 2013, 2017). This task is usually modeled as a classification problem, i.e. determining whether there is a causal relation between two events in a sentence. For example in Figure 1, an ECI system should identify two causal relations in two sentences: (1) **attack** $\xrightarrow{\text{cause}}$ **killed** in S1; (2) **statement** $\xrightarrow{\text{cause}}$ **protests** in S2.

Most existing methods for ECI heavily rely on annotated training data (Mirza and Tonelli, 2016;




Figure 1: S1 and S2 are *causal sentences* that contain *causal events*. S3 is produced by EDA based on S1. The dotted line indicates the causal relation.

Riaz and Girju, 2014b; Hashimoto et al., 2014; Hu and Walker, 2017; Gao et al., 2019). However, existing datasets are relatively small, which impede the training of the high-performance event causality reasoning model. According to our statistics, the largest widely used dataset EventStoryLine Corpus (Caselli and Vossen, 2017) only contains 258 documents, 4316 sentences, and 1770 causal event pairs. Therefore, data lacking is an essential problem that urgently needs to be addressed for ECI.

Up to now, data augmentation is one of the most effective methods to solve the data lacking problem. However, most of the NLP-related augmentation methods are a task-independent framework that produces new data at one time (Zhang et al., 2015; Guo et al., 2019; Xie et al., 2019b). In these frameworks, data augmentation and target task are modeled independently. This often leads to a lack of task-related characteristics in the generated data, such as task-related linguistic expression and knowledge. For example, easy data augmentation (EDA) (Wei and Zou, 2019) is the most representative method that relies on lexical substitution, deletion, swapping, and insertion to produce new data. However, solely relying on such word operations often generates new data that dissatisfies task-related qualities. As shown in Figure 1, S3 is produced by EDA, it lacks a linguistic expression that expresses the causal semantics between *kill* and *attack*. Therefore, how to

interactively model data augmentation and target task to generate new data with task-related characteristics is a challenging problem on ECI.

Specific to ECI, we argue that an ideal task-related generated causal sentence needs to possess two characteristics as follows. (1) The two events in the causal sentence need to have a causal relation. We call such property as **Causality**. For example, there is usually a causal relation between an *attack* event and a *kill* event, while nearly no causal relation between an *attack* event and a *born* event. (2) The linguistic expressions of the causal sentence need to be well-formed to express the causal semantic of events. We call such property as **Well-formedness**, which consists of a) canonical sentence grammar, b) event-related entities with semantic roles (e.g. the *attack* was carried out by a *police* in S1), and c) cohesive words that express complete causal semantics (e.g. *in a* and other words except for events and entities in S1).

To this end, we propose a learnable data augmentation framework for ECI, dubbed as **Learnable Knowledge-Guided Data Augmentation** (LearnDA). This framework regards sentence-to-relation mapping (*the target task*, ECI) and relation-to-sentence mapping (*the augmentation task*, sentence generation) as dual tasks and models the mutual relation between them via dual learning. Specifically, LearnDA can use the duality to generate task-related new sentences learning from identification and makes it more accurate to understand the causal semantic learning from generation. On the one hand, LearnDA is knowledge guided. It introduces diverse causal event pairs from KBs to initialize the dual generation which could ensure the **causality** of generated causal sentences. For example, the knowledge of *judgment* $\xrightarrow{\text{cause}}$ *demonstration* from KBs can be used to construct a novel causal sentence, which is also helpful to understand the causal semantic of *statement* $\xrightarrow{\text{cause}}$ *protests*. On the other hand, LearnDA is learnable. It employs a constrained generative architecture to generate **well-formed** linguistic expressions via iteratively learning in the dual interaction, which expresses the causal semantic between given events. Methodologically, it gradually fills the remaining missing cohesive words of the complete sentences under the constraint of given events and related entities.

In experiments, we evaluate our model on two benchmarks. We first concern the standard evaluation and show that our model achieves the state-of-

the-art performance on ECI. Then we estimate the main components of LearnDA. Finally, our learnable augmentation framework demonstrates definite advantages over other augmentation methods in generating task-related data for ECI.

In summary, the contributions as follows:

- We propose a new learnable data augmentation framework to solve the data lacking problem of ECI. Our framework can leverage the duality between identification and generation via dual learning which can learn to generate task-related sentences for ECI.
- Our framework is knowledge guided and learnable. Specifically, we introduce causal event pairs from KBs to initialize the dual generation, which could ensure the causality of generated causal sentences. We also employ a constrained generative architecture to gradually generate well-formed causal linguistic expressions of generated causal sentences via iteratively learning in the dual interaction.
- Experimental results on two benchmarks show that our model achieves the best performance on ECI. Moreover, it also shows definite advantages over previous data augmentation methods.

2 Related Work

To date, many researches attempt to identify the causality with linguistic patterns or statistical features. For example, some methods rely on syntactic and lexical features (Riaz and Girju, 2013, 2014b). Some focus on explicit causal textual patterns (Hashimoto et al., 2014; Riaz and Girju, 2014a, 2010; Do et al., 2011; Hidey and McKeown, 2016). And some others pay attention on statistical causal association and cues (Beamer and Girju, 2009; Hu et al., 2017; Hu and Walker, 2017).

Recently, more attention is paid to the causality between events. Mirza and Tonelli (2014) annotated Causal-TimeBank of event-causal relations based on the TempEval-3 corpus. Mirza et al. (2014), Mirza and Tonelli (2016) extracted event-causal relation with a rule-based multi-sieve approach and improved the performance incorporating with event temporal relation. Mostafazadeh et al. (2016) annotated both temporal and causal relations in 320 short stories. Caselli and Vossen (2017) annotated the EventStoryLine Corpus for

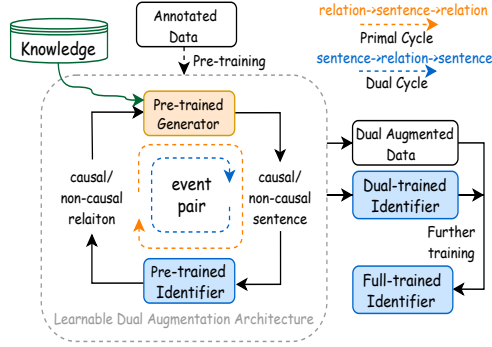


Figure 2: Overview of the learnable knowledge-guided dual data augmentation for ECI.

event causality identification. Dunietz et al. (2017) presented BECauSE 2.0, a new version of the BE-CauSE corpus (Dunietz et al., 2015) of causal relation and other seven relations. Gao et al. (2019) modeled document-level structures to identify causality. Liu et al. (2020) identified event causality with the mention masking generalization.

Unlike computer vision, the augmentation of text data in NLP is pretty rare (Chaudhary, 2020). Zuo et al. (2020) solved the data lacking problem of ECI with the distantly supervised labeled training data. However, including the distant supervision, most of the existing data augmentation methods for NLP tasks are task-independent frameworks (Related work of data augmentation and dual learning are detailed in Appendix B). Inspired by some generative methods which try to generate additional training data while preserving the class label (Anaby-Tavor et al., 2019; Yang et al., 2019; Papanikolaou and Pierleoni, 2020), we introduce a new learnable framework for augmenting task-related training data for ECI via dual learning enhanced with external knowledge.

3 Methodology

As shown in Figure 2, LearnDA jointly models a knowledge guided sentence generator (input: *event pair and its causal/non-causal relation*, output: *causal/non-causal sentence*) and an event causality identifier (input: *event pair and its sentence*, output: *causal/non-causal relation*) with dual learning. LearnDA iteratively optimizes identifier and generator to generate task-related training data, and then utilize new data to further train the identifier. Therefore, we first present the main idea of dual learning, which is the architecture of learnable dual augmentation, including the states, actions, policies, and

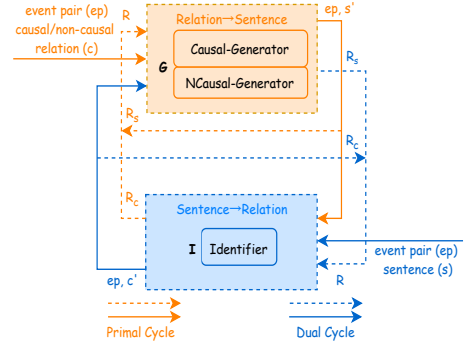


Figure 3: The architecture of learnable dual augmentation. *Causal* and *NCausal* represent the causal and non-causal sentence generator respectively. Red parts are the process of $\langle \text{event pair}, \text{relation} \rangle \rightarrow \text{sentence} \rightarrow \text{relation}$ (primal cycle), while blue parts are the process of $\langle \text{event pair}, \text{sentence} \rangle \rightarrow \text{relation} \rightarrow \text{sentence}$ (dual cycle). Solid and dashed lines denote the main process and reward feedback direction respectively.

rewards. Then, we briefly introduce the knowledge guided sentence generator, especially the processes of knowledge guiding and constrained sentence generation. Finally, we describe the event causality identifier and training processes of LearnDA.

3.1 Architecture of Learnable Dual Augmentation

The architecture of learnable dual augmentation is shown in Figure 3. Specifically, I denotes the event causality identifier, and G denotes the sentence generator which consists of two independent generators. They produce causal and non-causal sentences on the relation c of input event pair ep .

Generally, G generates a sentence s' which expresses the causal or non-causal relation c of the input event pair ep . Then it receives the reward R that consists of a semantic alignment reward R_s from itself and a causality reward R_c from I (primal cycle). Similarly, I identifies the causal or non-causal relation c' of the input event pair ep with its sentence s . Then it receives the reward R consists of a causality reward R_c from itself and a semantic alignment reward R_s from G (dual cycle).

I and G are optimized interactively with dual reinforcement learning. Specifically, for G , an action is the generation from relation to sentence, a state is denoted by the representation of input event pair and its relation, a policy is defined by the parameters of generator. For I , an action is the identification from sentence to relation, a state is denoted by the representation of input event pair and its

sentence, a policy is defined by the parameters of identifier. Inspired by Shen and Feng (2020), we utilize a probability distribution over actions given states to represent the policies, i.e., the probability distribution of the generation of G and identification of I . As aforementioned, we introduce two rewards, causality (R_c) and semantic alignment (R_s) rewards, which encourage G to generate task-related sentences with the feedback from identifier, while further optimize I with the feedback from generator. Definitions are as following:

Causality Reward (R_c) If the relation of input event pair can be clearly expressed by the generated sentence, it will be easier to be understood by identifier. Therefore, we use the causal relation classification accuracy as the causality reward to evaluate the causality of generated sentences, while tune and optimize the identifier itself:

$$R_c(ep, s) = \begin{cases} p(c'|s; \theta_I) & \text{Correct classification} \\ -p(c'|s; \theta_I) & \text{Otherwise,} \end{cases} \quad (1)$$

where θ_I is the parameter of I , $p(c'|s; \theta_I)$ denotes the probability of relation classification, s denotes the input sentence and c' is the classified relation.

Semantic Alignment Reward (R_s) We hope that the semantic of the generated sentence can be consistent with the relation of the input event pair. Additionally, if the relation of the input event pair can be more accurately classified, the semantic of the new generated sentence can be considered more consistent with it. Therefore, we measure the semantic alignment by means of the probability of constructing a sentence with similar semantic to the input relation, and the reward is:

$$R_s(ep, c) = p(s'|c; \theta_G) = \frac{1}{|T_s|} \sum_{t \in T_s} p(t|c; \theta_G), \quad (2)$$

where θ_G is the parameter of G , c is the input relation, t is one of the generated tokens T_s of the generated sentence s' , and $p(t|c; \theta_G)$ is the generated probability of t . Specifically, there are two independent G with different θ_G . In detail, θ_G^c is employed to generate causal sentence when the input c is causal relation, and non-causal sentence is generated via θ_G^{nc} when c is non-causal relation.

3.2 Knowledge Guided Sentence Generator

As shown in Figure 4, knowledge guided sentence generator (KSG) first introduces diverse causal and non-causal event pairs from KBs for *causality*. Then, given an event pair and its causal or non-causal relation, it employs a constrained gen-

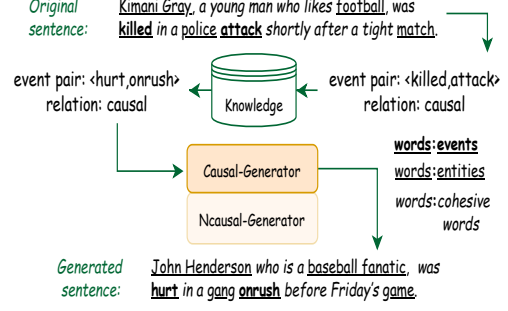


Figure 4: Flow diagram of the knowledge guided sentence generator (KSG). We take causal sentence generation via lexical knowledge expanding as an example.

erative architecture to generate new *well-formed* causal/non-causal sentences that contain them.

Knowledge Guiding KSG introduces event pairs that are probabilistic causal or non-causal from multiple knowledge bases in two ways. (1) *Lexical knowledge expanding*: expanding annotated event pairs via external dictionaries, such as WordNet (Miller, 1995) and VerbNet (Schuler, 2005). (2) *Connective knowledge introducing*: introducing event pairs from external event-annotated documents (KBP corpus) assisted with FrameNet (Baker et al., 1998) and Penn Discourse Treebank (PDTB2) (Group et al., 2008). As shown in Table 1, we illustrate how to extract event pairs from multiple knowledge bases. Then, inspired by Borde et al. (2013), we filter the extracted event pairs by converting them into triples $\langle e_i, \text{causal/non-causal}, e_j \rangle$ and calculating the causal-distance by maximizing L in a causal representation space:

$$L = \sum_{(e_i, e_j) \in T} \sum_{(e'_i, e'_j) \in T'} [\lambda + d(e'_i, e'_j) - d(e_i, e_j)]_+, \quad (3)$$

where T and T' are the causal and non-causal triples set respectively, and e is the representation of event. After that, the higher probability of causal relation, the shorter distance between two events, and we sort event pairs in ascending order by their distances. Finally, we keep the top and bottom $\alpha\%$ sorted event pairs to obtain the causal and non-causal event pairs sets for generation.

Constrained Sentence Generator Given an event pair, constrained sentence generator produces a well-formed sentence that expresses its causal or non-causal relation in three stages: (1) *assigning event-related entities* ensures the logic of the semantic roles of events, (2) *completing sentences* ensures the completeness of causal or non-causal

Knowledge	How to extract event pair	Why causal or non-causal
Lexical knowledge expanding		
WordNet	1) Extracting the synonyms and hypernyms from WordNet of each event in <i>ep</i> . 2) Assembling the items from the two groups of two events to generate causal/non-causal event pairs.	Items in each group are the synonyms and hypernyms of the annotated causal/non-causal event pairs.
VerbNet	1) Extracting the words from VerbNet under the same class as each event in <i>ep</i> . 2) Assembling the items from the two groups of two events to generate causal/non-causal event pairs.	Items in each group are in the same class of the annotated causal/non-causal event pairs.
e.g.	$\langle (killed, attack), causal \rangle \xRightarrow{Synonyms} kill \xRightarrow{Synonyms} hurt, attack \xRightarrow{Synonyms} onrush \Rightarrow \langle (hurt, onrush), causal \rangle$ <i>Original sentence:</i> Kimani Gray, a young man who likes football, was killed in a police attack shortly after a tight match.	
Connective knowledge introducing		
FrameNet PDTB2	1) Extracting causal/non-causal connectives from FrameNet ¹ and PDTB2. 2) Extracting any two events connected by causal/non-causal connectives on KBP corpus to obtain causal/non-causal event pairs and original sentences respectively.	Introduced event pairs are connected by causal/non-causal connectives.
e.g.	Looting <i>because</i> someone beat up someone, like the Travon Martin case. $\xRightarrow{because} \langle (loot, beat_up), causal \rangle$ <i>Original sentence:</i> Looting because someone beat up someone, like the Travon Martin case.	

Table 1: Extracting causal and non-causal event pairs from multiple knowledge bases.

semantic expression, (3) *filtering sentences* ensures the quality and diversity of generated sentences.

Assigning Event-related Entities. Event related entities play different semantic roles of events in sentences, which is an important part of event-semantic expression. Hence, as shown in Figure 4, given an event pair, we firstly assign logical entities for input events to guarantee the logic of semantic roles in the new sentences, such as *gang* is a logical entity as the body of the event *onrush*. Logically, entities of the same type play the same semantic roles in similar events. Moreover, as shown in Table 1, there is a corresponding original sentence for each extracted event pair. Therefore, in new sentence, we assign the most similar entity in the same type from candidate set² for each entity in the original sentence. For example, we assign *gang* for *onrush* in new sentence which is similar with the *police* related to *attack* in the original sentence. Specifically, we put the candidate entities in the same position in the original sentence to obtain their BERT embeddings. Then we select entities via the cosine similarity between their embeddings: $\mathcal{E}(ent) = \frac{1}{|ent|} \sum_{w \in ent} \mathcal{E}(w)$, where *ent* is the entity and $\mathcal{E}(w)$ is the BERT embedding of *ent*.

Completing Sentences. A well-formed sentence requires a complete linguistic expression to express the causal or non-causal semantics. Therefore, we complete sentences by filling the cohesive words between given events and assigned entities with masked BERT (Devlin et al., 2019). All words except events and entities are regarded as cohesive words. Specifically, we insert a certain number of the special token [MASK] between events and

entities, and then predict the [MASK]³ tokens as new words. As shown in Figure 4, we fill cohesive tokens via two independent generators to express causal and non-causal semantic according to the relation of given events. For example, in a guiding a causal semantic filled by the causal generator.

Filtering Sentences. Inspired by Yang et al. (2019), we design a filter to select new sentences that are balanced between high quality and high diversity with two key factors: 1) **Perplexity** (PPL): we take the average probability of the filled cohesive words in the new sentence s' as its perplexity: $PPL(s') = \frac{1}{|T(s')|} \sum_{t \in T(s')} P(t)$, where T is the set of filled cohesive words. 2) **Distance** (DIS): we calculate the cosine similarity between generated sentence s' and annotated data D_m as its distance: $DIS(s', D_m) = \frac{1}{|D_m|} \sum_{s \in D_m} \frac{\mathcal{E}(s') \cdot \mathcal{E}(s)}{\|\mathcal{E}(s')\| \times \|\mathcal{E}(s)\|}$, where D_m is m random selected annotated sentences and \mathcal{E} is the BERT sentence representation of the [CLS] token. A new sentence should have both appropriate high PPL which indicates the quality of generation, and appropriate high DIS which indicates the difference from the original sentences. Therefore, we select the top $\beta\%$ of the newly generated sentences according to *Score* for the further training of identifier as following: $Score(s') = \mu PPL(s') + (1 - \mu) DIS(s', D_m)$, where the μ is an hyper-parameter.

3.3 Training of LearnDA for ECI

We briefly describe the training processes of LearnDA for ECI, including the pre-training of generator and identifier, the dual reinforcement training, and the further training of identifier.

²We collect entities from annotated data and KBP corpus.

³The inserted [MASK] is 1.2 times the number of words between events and entities in the original sentence.

Algorithm 1 Dual Reinforcement Training of \mathcal{G} \mathcal{I} .

Require: A set of knowledge guided event pairs $\{(ep, s, c)\}$
A pre-trained generator \mathcal{G} and identifier \mathcal{I}
Repeat: Early stop on the development set according to \mathcal{I} .

- 1: **Loop:** PRIMAL CYCLE
- 2: **for** event pair (ep_i, s_i, c_i) in batch **do**
- 3: Generator generates the sentence s'_i of ep_i ;
- 4: Identifier re-predicts the causality c'_i of ep_i ;
- 5: Computing the reward as:
- 6: $R_{primal}^s = \lambda R_s(ep_i, c_i) + (1 - \lambda) R_c(ep_i, s'_i)$.
- 7: Computing the stochastic gradient of θ_G :
- 8: $\nabla_G = R_{primal}^s \cdot \nabla_{\theta_G} L_G(ep_i, c_i)$.
- 9: **end for**
- 10: Model batch updates: $\theta_G \leftarrow \theta_G + \eta \cdot \nabla_G$
- 11: **end Loop:**
- 12:
- 13: **Loop:** DUAL CYCLE
- 14: **for** event pair (ep_i, s_i, c_i) in batch **do**
- 15: Identifier predicts the causality c'_i of ep_i ;
- 16: Generator re-generates the sentence s'_i of ep_i ;
- 17: Computing the reward as:
- 18: $R_{dual}^s = \gamma R_c(ep_i, s_i) + (1 - \gamma) R_s(ep_i, c'_i)$.
- 19: Computing the stochastic gradient of $\theta_{\mathcal{I}}$:
- 20: $\nabla_{\mathcal{I}} = R_{dual}^s \cdot \nabla_{\theta_{\mathcal{I}}} L_I(ep_i, s_i)$.
- 21: **end for**
- 22: Model batch updates: $\theta_{\mathcal{I}} \leftarrow \theta_{\mathcal{I}} + \eta \cdot \nabla_{\mathcal{I}}$
- 23: **end Loop:**

Event Causality Identifier First of all, we formulate event causality identification as a sentence-level binary classification problem. Specifically, we design a classifier based on BERT (Devlin et al., 2019) to build our identifier. The input of the identifier is the event pair ep and its sentence s . Next, we take the stitching of manually designed features (same lexical, causal potential, and syntactic features as Gao et al. (2019)) and two event representations as the input of top MLP classifier. Finally, the output is a binary vector to predict the causal/non-causal relation of the input event pair ep .

Pre-training We pre-train the identifier and generator on labeled data before dual reinforcement training. On the one hand, we train identifier via the cross-entropy objective function of the relation classification. On the other hand, for generators, we keep the events and entities in the input sentences, replace the remaining tokens with a special token [MASK], and then train it via the cross-entropy objective function to re-predict the masked tokens. Specifically, causal generator and non-causal generator are pre-trained on causal and non-causal labeled sentences respectively.

Dual Reinforcement Training As shown in Algorithm 1, we interactively optimize the generator and identifier by dual reinforcement learning. Specifically, we maximize the following objective

functions:

$$L_G(ep, c) = \begin{cases} p(s'|c; \theta_G) = \frac{1}{|T_s|} \sum_{t \in T_s} p(t|c; \theta_G) \\ p(s'|c; \theta_{NG}) = \frac{1}{|T_s|} \sum_{t \in T_s} p(t|c; \theta_{NG}), \end{cases} \quad (4)$$

$$L_I(ep, s) = p(c'|s; \theta_I), \quad (5)$$

where θ_G and θ_{NG} is the parameters of causal and non-causal sentence generators respectively, T_s is the masked tokens. Finally, after dual data augmentation, we utilize generated sentences to further train the dual-trained identifier via the cross-entropy objective function of relation classification.

4 Experiments

4.1 Experimental Setup

Dataset and Evaluation Metrics Our experiments are conducted on two main benchmark datasets, including: **EventStoryLine** v0.9 (ESC) (Caselli and Vossen, 2017) described above; and (2) **Causal-TimeBank** (Causal-TB) (Mirza and Tonelli, 2014) which contains 184 documents, 6813 events, and 318 causal event pairs. Same as previous methods, we use the last two topics of ESC as the development set for two datasets. For evaluation, we adopt Precision (P), Recall (R), and F1-score (F1) as evaluation metrics. We conduct 5-fold and 10-fold cross-validation on ESC and Causal-TB respectively, same as previous methods to ensure comparability. All the results are the average of three independent experiments.

Parameters Settings In implementations, both the identifier and generators are implemented on BERT-Base architecture⁴, which has 12-layers, 768-hiddens, and 12-heads. We set the learning rate of generator pre-training, identifier pre-training/further training, and dual reinforcement training as 1e-5, 1e-5, and 1e-7 respectively. We set the ratio of the augmented data used for training to the labeled data, α , β , μ , λ and γ as 1:2, 30%, 50%, 0.2, 0.5 and 0.5 respectively tuned on the development set. And we apply early stop and SGD gradient strategy to optimize all models. We also adopt a negative sampling rate of 0.5 for training the identifier, owing to the sparseness of positive examples. (See Appendix D for more details.)

Compared Methods Same as previous state-of-the-art work. For ESC, we prefer 1) **LSTM** (Cheng and Miyao, 2017), a dependency path based

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

sequential model that models the context between events to identify causality; 2) **Seq** (Choubey and Huang, 2017), a sequence model explores complex human designed features for ECI; 3) **LR+** and **ILP** (Gao et al., 2019), document-level models adopt document structures for ECI. For Causal-TB, we prefer 1) **RB**, a rule-based system; 2) **DD**, a data driven machine learning based system; 3) **VR-C**, a verb rule based model with data filtering and gold causal signals enhancement. These models are designed by Mirza and Tonelli (2014); Mirza (2014) for ECI.

Owing to our methods are constructed on BERT, we build BERT-based methods: 1) **BERT**, a BERT-based baseline, our basic proposed event causality identifier. 2) **MM** (Liu et al., 2020), the BERT-based SOTA method with mention masking generalization. 3) **MM+Aug**, the further re-trained MM with our dual augmented data. 4) **KnowDis** (Zuo et al., 2020) improved the performance of ECI with the distantly labeled training data. We compare with it to illustrate the quality of our generated ECI-related training data. 5) **MM+ConceptAug**, to make a fair comparison, we introduce causal-related events from ConceptNet that employed by MM, and generate new sentences via KnowDis and LearnDA to further re-train MM (see Appendix C for details). Finally, we use **LearnDA_{Full}** indicates our full model, which is the dual-trained identifier further trained via dual augmented data.

4.2 Our Method vs. State-of-the-art Methods

Table 2 shows the results of ECI on EventStoryLine and Causal-TimeBank. From the results:

1) Our **LearnDA_{Full}** outperforms all baselines and achieves the best performance (52.6%/51.9% on F1 value), outperforming the no-bert (ILP/VR-C) and bert (MM/KnowDis) state-of-the-art methods by a margin of 7.9%/8.7% and 2.5%/2.1% respectively, which justifies its effectiveness. Moreover, BERT-based methods demonstrate high recall value, which is benefited from more training data and their event-related guided knowledge.

2) Comparing KnowDis with **LearnDA_{Full}**, we note that training data generated by LearnDA is more helpful to ECI than distant supervision with external knowledge (+2.9%/+2.1%). This shows that LearnDA can generate more ECI-related data.

3) Comparing **MM+ConceptNet** with MM, with the same knowledge base, our dual augmented data can further improve the performance

Methods	P	R	F1
ESC			
LSTM (Cheng and Miyao, 2017)	34.0	41.5	37.4
Seq (Choubey and Huang, 2017)	32.7	44.9	37.8
LR+ (Gao et al., 2019)	37.0	45.2	40.7
ILP (Gao et al., 2019)	37.4	55.8	44.7
BERT	36.1	56.0	43.9
KnowDis (Zuo et al., 2020)	39.7	66.5	49.7
MM (Liu et al., 2020)	41.9	62.5	50.1
MM+ <i>ConceptAug</i> (Ours)	41.2	66.5	50.9*
MM+ <i>Aug</i> (Ours)	41.0	69.3	51.5*
LearnDA_{Full} (Ours)	42.2	69.8	52.6*
Causal-TB			
RB (Mirza and Tonelli, 2014)	36.8	12.3	18.4
DD (Mirza and Tonelli, 2014)	67.3	22.6	33.9
VR-C (Mirza, 2014)	69.0	31.5	43.2
BERT	38.5	43.9	41.0
MM (Liu et al., 2020)	36.6	55.6	44.1
KnowDis (Zuo et al., 2020)	42.3	60.5	49.8
MM+ <i>ConceptAug</i> (Ours)	38.8	59.2	46.9*
MM+ <i>Aug</i> (Ours)	39.2	61.9	48.0*
LearnDA_{Full} (Ours)	41.9	68.0	51.9*

Table 2: Results on event causality identification. * denotes a significant test at the level of 0.05.

(+0.8%/+2.8%), which illustrates that LearnDA can make more effective use of external knowledge by generating task-related training data.

4) Comparing **MM+Aug** with MM, we note that training with our dual augmented data can improve the performance by 1.4%/3.9%, even though MM is designed on BERT-Large (LearnDA is constructed on BERT-Base) and also introduces external knowledge. This indicates that the augmented data generated by our LearnDA can effectively alleviate the problem of data lacking on the ECI.

4.3 Effect of Learnable Dual Augmentation

We analyze the effect of the learnable dual augmentation for event causality identification. 1) *For identifier*. Comparing **LearnDA_{Dual}** with BERT in Table 3, we note that the performance of the proposed identifier is improved (+2.6%) after the dual training only with the same labeled data. This indicates that the identifier can learn more informative expressions of causal semantic from generation with dual learning. 2) *For generator*. Comparing **BERT_{DualAug}** with **BERT_{Aug}** in Table 3, we note that the dual augmented data is high quality and more helpful to ECI (+2.6%). This indicates generator can generate more ECI task-related data learned from identifier with dual learning.

Figure 5 illustrates the learnability of our LearnDA. Specifically, as the number of training rounds of dual learning increases, the generated data gradually *learns* task-related information, fur-

Method	P	R	F
BERT (Our basic identifier)	36.1	56.0	43.9
BERT _{OrgAug}	36.6	59.7	45.4*
BERT _{DualAug}	37.8	65.6	48.0*
LearnDA _{Dual}	36.8	63.0	46.5*
LearnDA _{DualAug-w/o.KB}	37.5	67.0	48.1*
-LearnDA _{DualAug-w/.intro}	39.0	66.0	49.0*
-LearnDA _{DualAug-w/.verbnet}	39.4	66.7	49.5*
-LearnDA _{DualAug-w/.wordnet}	39.6	67.6	49.9*
LearnDA _{Full}	42.2	69.8	52.6*

Table 3: Ablation results on event causality identification on ESC. * denotes a significant test at the level of 0.05. BERT_{OrgAug} and BERT_{DualAug} denote the BERT is further trained on no-dual and dual augmented data respectively; LearnDA_{Dual} denotes our identifier is only trained by dual learning without further training; LearnDA_{DualAug-w/o.KB} denotes the LearnDA_{Dual} is further trained by dual augmented data without knowledge guiding; LearnDA_{DualAug-w/.<kb>} denotes LearnDA_{Dual} is further trained by dual augmented data guided with knowledge base *kb*.




Figure 5: The impact of the training rounds of dual learning on event causality identification on ESC. In each round, we generate new training data by the generator at the current round. The performance is achieved by further training the identifier at the current round with the aforementioned newly generated data.

ther improving the performance accordingly.

4.4 Effect of Knowledge Guiding

Table 3 also illustrates the effect of knowledge guiding on ECI depending on different knowledge bases. 1) Comparing LearnDA_{Full} with LearnDA_{DualAug-w/o.KB}, we note that the augmented data guided by external knowledge can further improve the performance of ECI. 2) Specifically, *lexical expanding* and *connective introducing* (Sec 3.2) can both make the representation of causal relation more generalized, further making it easier for the identifier to understand the causality. 3) Moreover, the expanding is more effective than the introducing, because the former brings a wider range of effective knowledge, thus the guidance of

Method	P	R	F
BERT (Our identifier)	36.1	56.0	43.9
TextSurface _{BERT}	37.0	57.5	45.0*
BackTranslation _{BERT}	36.8	61.0	45.9*
EDA _{BERT}	36.6	62.4	46.1*
LearnDA _{BERT}	37.8	65.6	48.0*

Table 4: Results of different data augmentation methods on event causality identification on ESC dataset. * denotes a significant test at the level of 0.05.

	Gold	EDA	BackTrans	LearnDA
Causality	3.80	3.20	3.70	3.60
Well-formedness	3.95	2.75	3.83	3.64
Diversity (Man/Auto)	0.0/1.0	3.08/0.70	2.80/0.85	3.51/0.66

Table 5: Manual (4-score rating (0, 1, 2, 3)) and automatic (BLEU score) evaluation of the generated sentences via different methods from causality, well-formedness and diversity. Causality and well-formedness are assessed manually, while diversity is assessed manually and automatically.

causal-related knowledge is better.

4.5 Our Augmentation vs. Other NLP Augmentations

In this section, we conduct a comparison between our augmentation framework and other NLP-related augmentation methods to further illustrate the effectiveness of LearnDA.

Effectiveness of Our Augmentation We train our identifier with augmented data produced by different NLP-related augmentation methods. As shown in Table 4, the augmented data generated by our LearnDA is more efficient for ECI, which is consistent with the previous analysis. The LearnDA can generate well-formed task-related new sentences that contain more event causal knowledge. Specifically, 1) *text surface transformation* brings a slight change to the labeled data, thus it has relatively little impact on ECI; 2) *Back translation* introduces limited new causal expressions by translation, thus it slightly increases the recall value on ECI; 3) *EDA* can introduce new expressions via substitution, but the augmented data is not canonical and cannot accurately express the causality, therefore, its impact on ECI is also limited.

Quantitative Evaluation of Task-relevance

We select five Ph.D. students majoring in NLP to manual score the 100 randomly selected augmented sentences given their corresponding original sentences as reference (Cohen’s kappa = 0.85). Furthermore, we calculate the BLEU (Papineni et al., 2002) value to further evaluate the

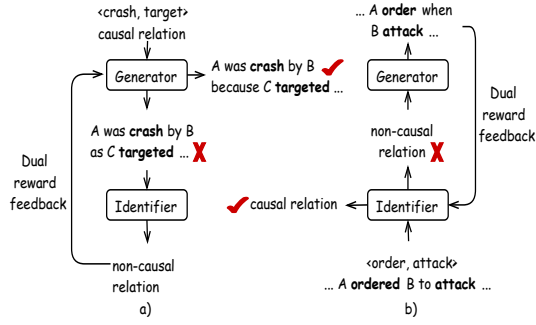


Figure 6: The modification of dual learning.

diversity. As aforementioned, the task-relevance of new sentences on ECI is manifested in causality and well-formedness, while the diversity indicates the degree of generalization. As shown in Table 5, we note the sentences generated by LearnDA are equipped with the above three properties that are close to the labeled sentences. Specifically, the sentences produced by EDA has a certain degree of causality and diversity due to the lexical substitution assisted by external knowledge. However, they cannot well express the causality due to the grammatical irregularities. Correspondingly, new sentences generated via back translation are very similar to the original sentences, while the diversity is poor.

4.6 Case Study

We conduct a case study to further investigate the effectiveness of our LearnDA. Figure 6 illustrates the modification process of dual learning. For example as a), given two causal events, the generator is expected to generate a causal sentence. However, the generator without dual learning produces a non-causal sentence. Fortunately, with dual learning, the identifier judges the generated sentence as a non-causal one and guides the generator to produce a causal sentence with the feedback. Similarly, as shown in b), given a causal sentence, the identifier is expected to output a causal relation, but no dual-trained one cannot do. Correspondingly, the generator constructs feedback of low confidence to guide the identifier to output a causal relation.

5 Conclusion

This paper proposes a new learnable knowledge-guided data augmentation framework (LearnDA) to solve the data lacking problem on ECI. Our framework can leverage the duality between generation and identification via dual learning to gener-

ate task-related sentences for ECI. Moreover, our framework is knowledge guided and learnable. Our method achieves state-of-the-art performance on EventStoryLine and Causal-TimeBank datasets.

Acknowledgments

We thank anonymous reviewers for their insightful comments and suggestions. This work is supported by the National Key Research and Development Program of China (No.2018YFB1005100), the National Natural Science Foundation of China (No.U1936207, 61806201). This work is also supported by Beijing Academy of Artificial Intelligence (BAAI2019QN0301) and the joint project with Beijing Baidu Netcom Science Technology Co., Ltd.

References

- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2019. [Not enough data? deep learning to the rescue!](#) *ArXiv*, abs/1911.03118.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. [The Berkeley FrameNet project](#). In *36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics, Volume 1*, pages 86–90, Montreal, Quebec, Canada. Association for Computational Linguistics.
- Brandon Beamer and Roxana Girju. 2009. [Using a bi-gram event model to predict causal potential](#). In *International Conference on Intelligent Text Processing and Computational Linguistics*, pages 430–441. Springer.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Advances in neural information processing systems*, pages 2787–2795.
- Ruisheng Cao, Su Zhu, Chen Liu, Jieyu Li, and Kai Yu. 2019. [Semantic parsing with dual learning](#). pages 51–64.
- Ruisheng Cao, Su Zhu, Chenyu Yang, Chen Liu, Rao Ma, Yanbin Zhao, Lu Chen, and Kai Yu. 2020. [Un-supervised dual paraphrasing for two-stage semantic parsing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6806–6817, Online. Association for Computational Linguistics.
- 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, Vancouver, Canada. Association for Computational Linguistics.
- Amit Chaudhary. 2020. [A visual survey of data augmentation in nlp](#).
- Yubo Chen, Shulin Liu, Xiang Zhang, Kang Liu, and Jun Zhao. 2017. [Automatically labeled data generation for large scale event extraction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 409–419, Vancouver, Canada. Association for Computational Linguistics.
- Fei Cheng and Yusuke Miyao. 2017. [Classifying temporal relations by bidirectional LSTM over dependency paths](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–6, Vancouver, Canada. Association for Computational Linguistics.
- Prafulla Kumar Choubey and Ruihong Huang. 2017. [A sequential model for classifying temporal relations between intra-sentence events](#). pages 1796–1802.
- Claude Coulombe. 2018. [Text data augmentation made simple by leveraging nlp cloud apis](#). *ArXiv*, abs/1812.04718.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Quang Do, Yee Seng Chan, and Dan Roth. 2011. [Minimally supervised event causality identification](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2015. [Annotating causal language using corpus lexicography of constructions](#). In *Proceedings of The 9th Linguistic Annotation Workshop*, pages 188–196, Denver, Colorado, USA. Association for Computational Linguistics.
- Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2017. [The BECauSE corpus 2.0: Annotating causality and overlapping relations](#). In *Proceedings of the 11th Linguistic Annotation Workshop*, pages 95–104, Valencia, Spain. Association for Computational Linguistics.
- 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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Roxana Girju. 2003. [Automatic detection of causal relations for question answering](#). In *Proceedings of the ACL 2003 Workshop on Multilingual Summarization and Question Answering*, pages 76–83, Sapporo, Japan. Association for Computational Linguistics.
- PDTB Research Group et al. 2008. The pdtb 2.0. *Annotation Manual. Technical Report IRCS-08-01, Institute for Research in Cognitive Science, University of Pennsylvania*.
- Hongyu Guo, Yongyi Mao, and Richong Zhang. 2019. [Augmenting data with mixup for sentence classification: An empirical study](#). *ArXiv*, abs/1905.08941.
- Chikara Hashimoto, Kentaro Torisawa, Julien Kloetzer, Motoki Sano, István Varga, Jong-Hoon Oh, and Yutaka Kidawara. 2014. [Toward future scenario generation: Extracting event causality exploiting semantic relation, context, and association features](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 987–997, Baltimore, Maryland. Association for Computational Linguistics.
- Christopher Hidey and Kathy McKeown. 2016. [Identifying causal relations using parallel Wikipedia articles](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1424–1433, Berlin, Germany. Association for Computational Linguistics.
- Zhichao Hu, Elahe Rahimtoroghi, and Marilyn Walker. 2017. [Inference of fine-grained event causality from blogs and films](#). pages 52–58.
- Zhichao Hu and Marilyn Walker. 2017. [Inferring narrative causality between event pairs in films](#). pages 342–351.
- Jian Liu, Yubo Chen, and Jun Zhao. 2020. [Knowledge enhanced event causality identification with mention masking generalizations](#). In *IJCAI-20*, pages 3608–3614. International Joint Conferences on Artificial Intelligence Organization. Main track.
- George A Miller. 1995. [Wordnet: a lexical database for english](#). *Communications of the ACM*, 38(11):39–41.
- Paramita Mirza. 2014. [Extracting temporal and causal relations between events](#). pages 10–17.
- Paramita Mirza, Rachele Sprugnoli, Sara Tonelli, and Manuela Speranza. 2014. [Annotating causality in the TempEval-3 corpus](#). In *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL)*, pages 10–19, Gothenburg, Sweden. Association for Computational Linguistics.

- 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, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.
- Paramita Mirza and Sara Tonelli. 2016. [CATENA: CAusal and TEmporal relation extraction from NATural language texts](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 64–75, Osaka, Japan. The COLING 2016 Organizing Committee.
- Nasrin Mostafazadeh, Alyson Grealish, Nathanael Chambers, James Allen, and Lucy Vanderwende. 2016. [CaTeRS: Causal and temporal relation scheme for semantic annotation of event structures](#). In *Proceedings of the Fourth Workshop on Events*, pages 51–61, San Diego, California. Association for Computational Linguistics.
- Jong-Hoon Oh, Kentaro Torisawa, Chikara Hashimoto, Motoki Sano, Stijn De Saeger, and Kiyonori Ohtake. 2013. [Why-question answering using intra- and inter-sentential causal relations](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1733–1743, Sofia, Bulgaria. Association for Computational Linguistics.
- Jong-Hoon Oh, Kentaro Torisawa, Canasai Kruengkrai, Ryu Iida, and Julien Kloetzer. 2017. [Multi-column convolutional neural networks with causality-attention for why-question answering](#). In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 415–424. ACM.
- Yannis Papanikolaou and A. Pierleoni. 2020. [Dare: Data augmented relation extraction with gpt-2](#). *ArXiv*, abs/2004.13845.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Mehwish Riaz and Roxana Girju. 2010. [Another look at causality: Discovering scenario-specific contingency relationships with no supervision](#). In *2010 IEEE Fourth International Conference on Semantic Computing*, pages 361–368. IEEE.
- Mehwish Riaz and Roxana Girju. 2013. [Toward a better understanding of causality between verbal events: Extraction and analysis of the causal power of verb-verb associations](#). In *Proceedings of the SIGDIAL 2013 Conference*, pages 21–30, Metz, France. Association for Computational Linguistics.
- Mehwish Riaz and Roxana Girju. 2014a. [In-depth exploitation of noun and verb semantics to identify causation in verb-noun pairs](#). In *Proceedings of the 15th Annual Meeting of the Special Interest Group on Discourse and Dialogue (SIGDIAL)*, pages 161–170, Philadelphia, PA, U.S.A. Association for Computational Linguistics.
- Mehwish Riaz and Roxana Girju. 2014b. [Recognizing causality in verb-noun pairs via noun and verb semantics](#). In *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL)*, pages 48–57, Gothenburg, Sweden. Association for Computational Linguistics.
- Dana Ruiter, Cristina España-Bonet, and Josef van Genabith. 2019. [Self-supervised neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1828–1834, Florence, Italy. Association for Computational Linguistics.
- Karin Kipper Schuler. 2005. [Verbnet: A broad-coverage, comprehensive verb lexicon](#).
- Lei Shen and Yang Feng. 2020. [CDL: Curriculum dual learning for emotion-controllable response generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 556–566, Online. Association for Computational Linguistics.
- Shang-Yu Su, Chao-Wei Huang, and Yun-Nung Chen. 2019. [Dual supervised learning for natural language understanding and generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5472–5477, Florence, Italy. Association for Computational Linguistics.
- Shang-Yu Su, Chao-Wei Huang, and Yun-Nung Chen. 2020. [Towards unsupervised language understanding and generation by joint dual learning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 671–680, Online. Association for Computational Linguistics.
- Mingming Sun, Xu Li, and Ping Li. 2018. [Logician and orator: Learning from the duality between language and knowledge in open domain](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Brussels, Belgium. Association for Computational Linguistics.
- William Yang Wang and Diyi Yang. 2015. [That’s so annoying!!!: A lexical and frame-semantic embedding based data augmentation approach to automatic categorization of annoying behaviors using #petpeeve tweets](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2557–2563, Lisbon, Portugal. Association for Computational Linguistics.
- Jason Wei and Kai Zou. 2019. [EDA: Easy data augmentation techniques for boosting performance on text classification tasks](#). In *Proceedings of the*

2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 6382–6388, Hong Kong, China. Association for Computational Linguistics.

Yingce Xia, Tao Qin, Wei Chen, Jiang Bian, Nenghai Yu, and Tie-Yan Liu. 2017. [Dual supervised learning](#). In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3789–3798. JMLR. org.

Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L. Yuille, and Quoc V. Le. 2019a. [Adversarial examples improve image recognition](#). *ArXiv*, abs/1911.09665.

Qizhe Xie, Zihang Dai, Eduard H. Hovy, Minh-Thang Luong, and Quoc V. Le. 2019b. [Unsupervised data augmentation for consistency training](#). *arXiv: Learning*.

Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. [Exploring pre-trained language models for event extraction and generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5284–5294, Florence, Italy. Association for Computational Linguistics.

Hai Ye, Wenjie Li, and Lu Wang. 2019. [Jointly learning semantic parser and natural language generator via dual information maximization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2090–2101, Florence, Italy. Association for Computational Linguistics.

Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. [Character-level convolutional networks for text classification](#). In *NIPS*.

Xinyu Zuo, Yubo Chen, Kang Liu, and Jun Zhao. 2020. [KnowDis: Knowledge enhanced data augmentation for event causality detection via distant supervision](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1544–1550, Barcelona, Spain (Online). International Committee on Computational Linguistics.

A Supplementary Experiment Results

A.1 Statistics of Dual Augmented Data

	Annotated data	Augmented data
#causal ep.	1170	3588
#causal sent.	1770	10442
#Ave sent.	1.5	2.9

Table 6: Statistics of causal event pairs and causal sentences in labeled data (ESC) and dual augmented data. (**#causal ep.** denotes the number of causal event pairs after removing duplicates, **#causal sent.** denotes the number of causal sentences, **#Ave sent.** denotes the average number of causal sentences containing the same causal event pair.)

As shown in Table 6, our dual augmented data is significantly more quantitative than the labeled data. Specifically, the causal event pairs are increased by 3.1 times, the causal sentences are increased by 5.9 times and the average number of causal sentences corresponding to each causal event pair is also increased.

A.2 Effectiveness of Different Quantities of Augmented Training Data

Ratio	P	R	F1
1:1	37.3	64.7	47.3*
1:2	37.8	65.6	48.0*
1:3	37.0	64.8	47.1*
1:4	36.2	64.2	46.3*

Table 7: Performance of identifier (BERT) trained with different ratios of labeled data and dual augmented data.

* denotes a significant test at the level of 0.05.

We change the quantity of dual augmented data for training to explore the influence of augmentation ratio on ECI. As shown in Table 7, when the ratio is 1:2, the effective knowledge brought by dual augmented data is maximized. And as the ratio increasing, the dual augmented data will bring noises, which obstructs the model to identify event causality and may change the data distribution from original data (Xie et al., 2019a). This suggests that too much augmented data is not better and that there is a trade-off between introducing knowledge and reducing noise.

A.3 Effectiveness of Extracting Event Pairs with Different Filtering Ratios

Table 8 tries to show the effectiveness of extracting event pairs with different filtering ratios on ECI. With the ratio of retained event pairs increasing,

α	P	R	F1	∇
30%	37.8	65.6	48.0*	-
40%	37.0	65.7	47.3*	-0.7
50%	36.2	65.0	46.5*	-1.5

Table 8: Performance of identifier (BERT) trained with different extracting event pairs filtered in different α . * denotes a significant test at the level of 0.05.

the augmented data hurts ECI’s performance. This proves the effectiveness of filtering, which further improves the causality of the generated sentences.

A.4 Effectiveness of Generated sentences with Different Filtering Ratios

β	P	R	F1	∇
50%	37.8	65.6	48.0*	-
60%	37.3	65.3	47.5*	-0.5
70%	36.9	64.9	47.0*	-1.0
80%	36.6	64.5	46.7*	-1.3

Table 9: Performance of identifier (BERT) trained with new generated sentences filtered in different β . * denotes a significant test at the level of 0.05.

Table 9 tries to show the effectiveness of generated sentences with different filtering ratios. With the ratio of retained generated sentences increasing, the contribution of filtered generated sentences for ECI decreases gradually. This proves the effectiveness of filtering, which can balance the overall quality of the sentences against diversity.

B Supplementary Related Work

B.1 Dual Learning

For many Natural Language Processing (NLP) tasks, there exist many primal and dual tasks, such as open information narration (OIN) and open information extraction (OIE) (Sun et al., 2018), natural language understanding (NLU) and natural language generation (NLG) (Su et al., 2019, 2020), semantic parsing and natural language generation (Ye et al., 2019; Cao et al., 2019, 2020), link prediction and entailment graph induction (Cao et al., 2019), query-to-response and response-to-query generation (Shen and Feng, 2020) and so on. The duality between the primal task and the dual task is considered as a constraint that both problems must share the same joint probability mutually. Recently, inspired by Xia et al. (2017) who implemented the duality in a neural-based dual learning system, the above primal-dual tasks are implemented in two different ways: 1) providing additional labeled samples via bootstrapping, and 2) adding rewards at

the training stage for each agent. We observe that the event causality identification and the sentence generation are dual to each other. Therefore, we apply a dual learning framework in the second way to optimize identification and generation interactively for generating ECI-related data.

B.2 Data Augmentation for NLP

The scarcity of annotated data is a thorny problem in machine learning. Unlike computer vision, the augmentation of text data in NLP is pretty rare. Existing text data augmentation methods for NLP tasks are almost task-independent frameworks and can be roughly summarized into the following categories (Chaudhary, 2020): (1) Lexical substitution tries to substitute words without changing the meaning (Zhang et al., 2015; Wei and Zou, 2019; Wang and Yang, 2015; Xie et al., 2019b); (2) Back translation tries to paraphrase a text while retraining the meaning (Xie et al., 2019b); (3) Text surface transformation tries to match transformations using regex (Coulombe, 2018); (4) Random noise injection tries to inject noise in the text to make the model more robust (Wei and Zou, 2019); (5) Generative method tries to generate additional training data while preserving the class label (Anaby-Tavor et al., 2019; Yang et al., 2019); (6) Distantly supervision and self-supervision try to introduce new training data from unlabeled text (Chen et al., 2017; Ruiter et al., 2019). As aforementioned, these frameworks cannot directly produce new suitable task-related examples for ECI. However, (1), (3), and (4) cannot guarantee the causality and well-formedness of new examples for ECI. Additionally, (2) and (5) are not easy to directly use external knowledge bases to generalize the event-related causal commonsense. Furthermore, (6) needs to design proprietary processing methods to generate ECI task-related training data. Zuo et al. (2020) solved the data lacking problem of ECI with the distantly supervised labeled training data. However, including the distant supervision, most of the existing text data augmentation methods for NLP tasks are task-independent frameworks. Therefore, we introduce a new learnable framework for augmenting task-related training data for ECI via dual learning enhanced with external knowledge.

C Generation with ConceptNet

To make a fair comparison, we introduce causal-related events from ConceptNet based on causal-

related concepts, and obtain the causal sentence via the method in KonwDis (Zuo et al., 2020) to further re-train MM (Liu et al., 2020). Specifically, firstly, we obtain triples based on cause-related semantic relations from ConceptNet, such as *Causes*, *HasSubevent*, *HasFirstSubevent*, *HasLastSubevent*, *MotivatedByGoal*, and *CausesDesire* relations. Secondly, we assemble any two events from obtained causal triples to generate causal event pairs set and filter them via the filter of KonwDis. Next, we employ filtered causal event pairs to collect preliminary noisy labeled sentences from external documents via the DistantAnnotator of KonwDis. Then, we use the CommonFilter of KnowDis assisted with causal commonsense knowledge to pick out labeled sentences that express causal semantics between events. Finally, the refined causal sentences are input into LearnDA to generated ECI-related dual augmented training data and further train the MM to obtain MM+*ConceptAug*.

D Main Experimental Environments and Other Parameters Settings

D.1 Experimental Environments

We deploy all models on a server with 250GB of memory and 4 TITAN Xp GPUs. Specifically, the configuration environment of the server is ubuntu 16.04, and our framework mainly depends on python 3.6.0 and PyTorch 1.0.

D.2 Other Parameters Settings

All the final hyper-parameters for evaluation are averaged after 3 independent tunings on the development set. Moreover, the whole dual learning framework which includes event causality identifier and knowledge guided sentence generator takes approximately 5 minutes per epoch when training. According to the early stop strategy, the training rounds for different folds are different, and it takes about 20-30 rounds.